SPECIAL COLLECTOR'S EDITION

SCIENTIFIC AMERICAN

The Brain

How this amazing organ defines our reality and helps us sense, think and act

INSIDE

How teen brains work

Why we sleep

Our individual consciousnesses

Long COVID: neurological disorder

Babies can do math

Dreams can foreshadow disease

SPRING 2023

SCIENTIFIC AMERICAN®
ESTABLISHED 1845

The Brain
is published by the staff of
Scientific American, with
project management by:

Editor in Chief: **Laura Helmuth**

Managing Editor: **Jeanna Bryner**

Chief Newsletter Editor:
Andrea Gawrylewski

Creative Director: **Michael Mrak**

Issue Designer: **Lawrence R. Gendron**

Senior Graphics Editor: **Jen Christiansen**

Associate Graphics Editor: **Amanda Montañez**

Photography Editor: **Monica Bradley**

Issue Photo Researcher: **Beatrix Mahd Soltani**

Copy Director: **Maria-Christina Keller**

Senior Copy Editors:
Angelique Rondeau, Aaron Shattuck

Associate Copy Editor: **Emily Makowski**

Managing Production Editor: **Richard Hunt**

Prepress and Quality Manager:
Silvia De Santis

Executive Assistant Supervisor: **Maya Harty**

Senior Editorial Coordinator: **Brianne Kane**

President: **Kimberly Lau**

Publisher and Vice President:
Jeremy A. Abbate

Vice President, Product and Technology:
Dan Benjamin

Vice President, Commercial: **Andrew Douglas**

Vice President, Content Services:
Stephen Pincock

Associate Vice President,
Business Development: **Diane McGarvey**

Programmatic Product Manager: **Zoya Lysak**

Product Managers:
Ian Kelly, Miguel Olivares

Head, Marketing: **Christopher Monello**

Senior Commercial Operations Coordinator:
Christine Kaelin

Custom Publishing Editor: **Lisa Pallatroni**

Senior Communications Manager:
Sarah Hausman

Production Controller: **Madelyn Keyes-Milch**

Advertising Production Manager:
Michael Broomes

Advertising Production Controller:
Michael Revis-Williams

Mark Ross Studios (*cover illustration*);
Bruce Rolff/Stockrek Images/Getty Images (*this page*)

FROM THE EDITOR

The Final Frontier Is Right Here

The human brain is a powerhouse. The Milky Way galaxy has hundreds of billions of stars—just a *fraction* of the 100 trillion connections in our brains that enable us to sense, think and act. Our minds are the result of countless neuronal firings that happen every millisecond. Across animals, brain size is usually tightly correlated with body size—the encephalization quotient, as it's called. Elephants score a 1 or 2 on this measure, long-finned pilot whales 2 to 3. Humans? We score between 7 and 8. Among the gains from this outsize brain is the work of neuroscientists, philosophers and even poets. Read this collection for a look at the most astounding recent brain discoveries.

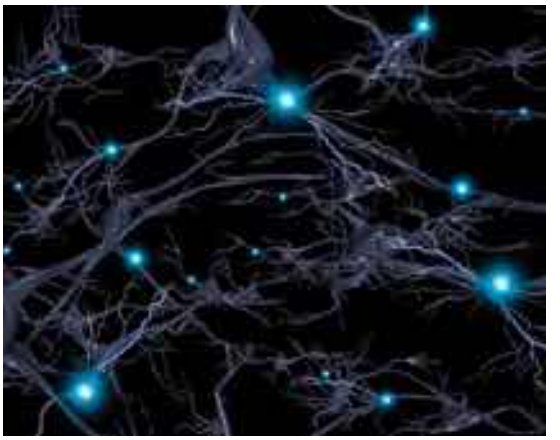

Branching and tangled neuronal cells fire countless times every millisecond.

Consciousness is perhaps humans' most complex aspect. Our perception of the world builds our individual realities—and no two are alike (*page 12*). The ventromedial prefrontal cortex may produce a fundamental model of the self and knit together versions of us from past and present—even projecting into the future (*page 30*). New discoveries for detecting consciousness in comatose people illuminate the remarkable persistence of the mind (*page 24*).

We are self-aware, yes. But how does that awareness integrate with a complicated and ever-changing world? A nonstop scanning and filtering system in the brain selects the most essential elements from our surroundings to help us thrive (*page 32*). Special brain regions are activated to help us physically navigate our surroundings (*page 62*) and recognize faces (*page 54*). From birth we have an innate ability with numbers (*page 40*), and we grow into a period of adolescence marked by sensitivity to social cues and community acceptance (*page 48*).

As with any complex machine, errors can arise. Leaks in the protective blood-brain barrier can trigger events leading to Alzheimer's disease and other dementias—a process that scientists are urgently trying to prevent (*page 68*). Acting out dreams can be a telltale early sign of brain disorders (*page 88*). Even viruses such as the one that causes COVID can produce lasting neurological problems, although precisely how is a fresh area of study (*page 74*).

The brain is wonderfully self-sustaining, even after many decades of firing away. Important maintenance includes getting ample sleep, which neuroscientists now believe flushes toxins and helps to consolidate memories (*page 98*). A positive attitude may affect how well our immune systems function (*page 104*). And here's both bad and good news on aging and the brain: after we turn 65, our risk of dementia increases (*page 96*), but stress seems to have less of an effect on our happiness (*page 110*).

Around the beginning of the 20th century Santiago Ramón y Cajal—the father of modern neuroscience—described neurons, the branching and tangled cells that make up brain matter, as "mysterious butterflies of the soul." More than 100 years later the enigma and marvel of the human brain endures. We know much more than when he first looked through his microscope and eyed that mystery. For some, the final frontier will be not in the reaches of space but as close as can be.

Andrea Gawrylewski
Chief Newsletter Editor,
editors@sciam.com

SPECIAL EDITION

SCIENTIFIC AMERICAN

Volume 32, Number 2, Spring 2023

HOW **CONSCIOUSNESS ARISES**

MAKING SENSE **OF THE WORLD**

BRAIN **DISORDERS**

HEALTHY BODY, **HEALTHY BRAIN**

DEPARTMENTS

 Scientific American Special (ISSN 1936-1513), Volume 32, Number 2, Spring 2023, published by Scientific American, a division of Springer Nature America, Inc., 1 New York Plaza, Suite 4600, New York, N.Y. 10004-1562. Canadian BN No. 127387652RT; TVQ1218059275 TQ0001. To purchase additional quantities: U.S., $13.95 each; elsewhere, $17.95 each. Send payment to Scientific American Back Issues, P.O. Box 3187, Harlan, Iowa 51537. **Inquiries: fax 212-355-0408 or telephone 212-451-8415. Printed in U.S.A.**

Zara Picken

Richard Borge

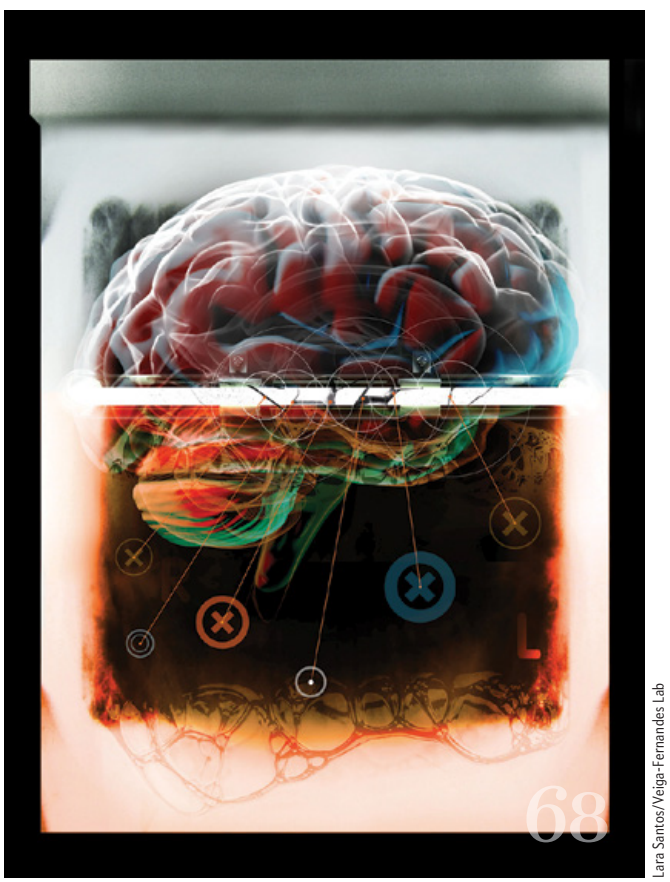

Viktor Koen

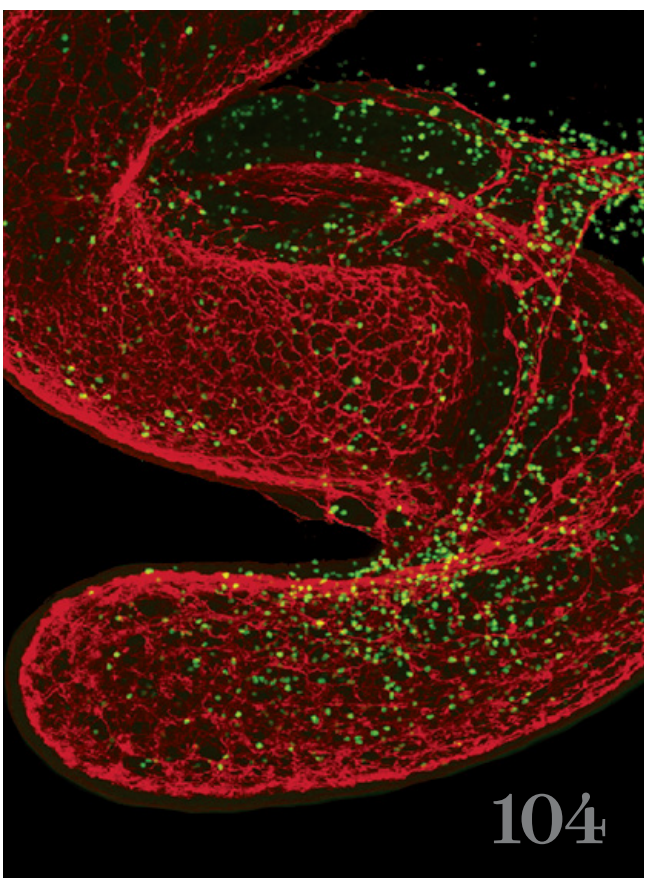

Lara Santos/Veiga-Fernandes Lab

how matter becomes mind

The new discipline of network neuroscience yields a picture of how mental activity arises from carefully orchestrated interactions among different brain areas

By Max Bertolero and Dani S. Bassett

Illustration by Mark Ross Studios

NETWORKS PERVADE OUR LIVES. EVERY DAY WE USE INTRICATE NETWORKS of roads, railways, maritime routes and skyways traversed by commercial flights. They exist even beyond our immediate experience. Think of the World Wide Web, the power grid and the universe, of which the Milky Way is an infinitesimal node in a seemingly boundless network of galaxies. Few such systems of interacting connections, however, match the complexity of the one underneath our skull.

Neuroscience has gained a higher profile in recent years, as many people have grown familiar with splashily colored images that show brain regions "lighting up" during a mental task. There is, for instance, the temporal lobe, the area by your ear, which is involved with memory, and the occipital lobe at the back of your head, which dedicates itself to vision.

What has been missing from this account of human brain function is how all these distinct regions interact to give rise to who we are. We and other researchers have drawn on a branch of mathematics called graph theory to better parse, probe and predict complex interactions in the brain that bridge the seemingly vast gap between frenzied neural electrical activity and an array of cognitive tasks—sensing, remembering, making decisions, learning a new skill and initiating movement. This new field of network neuroscience builds on and reinforces the idea that certain regions of the brain carry out defined activities. In the most fundamental sense, what the brain is—and thus who we are as conscious beings—is, in fact, defined by a sprawling network of 100 billion neurons with at least 100 trillion connecting points, or synapses.

Network neuroscience seeks to capture this complexity. We can now model the data supplied by brain imaging as a graph composed of nodes and edges. In a graph, nodes represent the units of the network, such as neurons or, in another context, airports. Edges serve as the connections between nodes—think of one neuron intertwined with the next, or contemplate airline flight routes. In our work, we reduce the human brain to a graph of roughly 300 nodes. Diverse areas can be linked together by edges representing the brain's structural connections: thick bundles of tubular wires called white matter tracts that tie together brain regions. This depiction of the brain as a unified network has already furnished a clearer picture of cognitive functioning, along with the practical benefit of enabling better diagnoses and treatment of psychiatric disorders. In the future, an understanding of brain networks may lead to a blueprint for improved artificial intelligence, new medicines, electrical-stimulation technology to alter malfunctioning neural circuitry in people with depression—and perhaps even the development of genetic therapies to treat mental illness.

THE MUSIC OF THE MIND

TO UNDERSTAND HOW networks underlie our cognitive capabilities, first consider the analogy of an orchestra playing a symphony. Until recently, neuroscientists have largely studied the functioning of individual brain regions in isolation, the neural equivalent of separate brass, percussion, string and woodwind sections. In the brain, this categorization represents an approach that dates back to Plato—quite simply, it entails carving nature at the joints and then studying the individual components that remain.

Just as it is useful to understand how the amygdala helps to process emotions, it is similarly vital to grasp how a violin produces high-pitched sounds. Still, even a complete list of brain regions and their functions—vision, motor, emotion, and so on—does not tell us how the brain really works. Nor does an inventory of instruments provide a recipe for Beethoven's *Eroica* symphony.

Network neuroscientists have begun to probe these mysteries by examining the way each brain region is embedded in a larger network of such regions and by mapping the connections between regions to study how each fits into the large, integrated network that is the brain. There are two major approaches. First, examining structural connectivity captures the instrumentation of the brain's orchestra. It is the physical means of creating the music, and the unique instrumentation of a given musical work constrains what can be played. Instrumentation matters, but it is not the music itself. Put another way, just as a collection of instruments is not music, an assemblage of wires does not represent brain function.

Second, living brains are massive orchestras of neurons that fire together in quite specific patterns. We hear a brain's music by measuring the correlation between the activity of each pair of regions, which indicates that they are working in concert. This measure of joint activity is known as functional connectivity, and we colloquially think of it as reflecting the music of the brain. If two regions fire with the same time-varying fluctuations, they are considered as functionally connected. This music is just as important as the decibels produced by a French horn or a viola. The volume of the brain's music can be thought of as the level of activity of electrical signals buzzing about one brain area or another.

At any moment, though, some areas within the three-pound organ are more active than others. We have all heard that people use a small fraction of their brain capacity. In fact, the entire brain is active at any point in time, but a given task modulates the activity of only a portion of the brain from its baseline level of activity.

That arrangement does not mean that you fulfill only half of your cognitive potential. In fact, if your entire brain were strongly active at the same time, it would be as if all the orchestra members were playing as loudly as possible—and that scenario would create chaos, not enable communication. The deafening sound would not convey the emotional overtones present in a great musical piece. It is the pitch, rhythms, tempo and strategic pauses that communicate information, both during a symphony and inside your head.

MODULARITY

JUST AS AN ORCHESTRA can be divided into groups of instruments from different families, the brain can be separated into collections of nodes called modules—a description of localized networks. All brains are modular. Even the 302-neuron network of the nematode *Caenorhabditis elegans* has a modular structure. Nodes within a module share stronger connections to one another than to nodes in other modules.

Each module in the brain has a certain function, just as every family of instruments plays a role in the symphony. We recently carried out an evaluation of a large number of independent studies—a meta-analysis—that included more than 10,000 functional magnetic resonance imaging (fMRI) experiments of subjects performing 83 different cognitive tasks and discovered that separate tasks map to different brain-network modules. There are modules occupied with attention, memory and introspective thought. Other modules, we found, are dedicated to hearing, motor movement and vision.

These sensory and motor cognitive processes involve single, contiguous modules, most of which are confined to one lobe of the brain. We also found that computations in modules do not spur more activity in other modules—a critical aspect of modular processing. Imagine a scenario in which every musician in an orchestra had to change the notes played every time another musician changed their notes. The orchestra would spiral out of control and would certainly not produce aesthetically pleasing sounds. Processing in the brain is similar—each module must be able to function mostly independently. Philosophers as early as Plato and as recent as Jerry Fodor have noted this necessity, and our research confirms it.

Even though brain modules are largely independent, a symphony requires that families of instruments be played in unison. Information generated by one module must eventually be integrated with other modules. Watching a movie with only a brain module for vision—without access to the one for emotions—would detract greatly from the experience.

For that reason, to complete many cognitive tasks, modules must often work together. A short-term memory task—holding a new phone number in your head—requires the cooperation of auditory, attention and memory-processing modules. To integrate and control the activity of multiple modules, the brain uses hubs—nodes where connections from the brain's different modules meet.

Some key modules that control and integrate brain activity are less circumspect than others in their doings. Their connections extend globally to multiple brain lobes. The frontoparietal control module spans the frontal, parietal and temporal lobes. It developed relatively recently on the timescale of evolution. The module is especially large in humans, relative to those of our closest primate ancestors. It is analogous to an orchestra conductor and becomes active across a large number of cognitive tasks.

The frontoparietal module ensures that the brain's multiple modules function in unison. It is heavily involved in what is called executive function, which encompasses the separate processes of decision-making, short-term memory and cognitive control. The last is the ability to develop complex strategies and inhibit inappropriate behavior.

Another highly interconnected module is the salience module, which hooks up to the frontoparietal control module and contributes to a range of behaviors related to attention and responses to novel stimuli. For example, take a look at two words: blue and red. If you are asked to tell someone the color of the words, you will react much faster to the one set in red. For the one set in green, your frontoparietal and salience modules activate when responding to its color because you have to suppress a natural inclination to answer "blue."

Finally, the default mode module spans the same lobes as the frontoparietal control network. It contains many hubs and is linked to a variety of cognitive tasks, including introspective thought, learning, memory retrieval, emotional processing, inference of the mental state of others, and even gambling. Critically, damage to these hub-rich modules disturbs functional connections throughout the brain and causes widespread cognitive difficulties, just as bad weather at a hub airport delays air traffic all over the country.

PERSONAL CONNECTIONS

ALTHOUGH OUR BRAINS have certain basic network components—modules interconnected by hubs—each of us shows slight variations in the way our neural circuits are wired. Researchers have devoted intense scrutiny to this diversity. In an initial phase of what is called the Human Connectome Project, 1,200 young people volunteered for a study of brain-network architecture, funded by the National Institutes of Health. (The final goal of the project is to map connectomes across the entire life span.) Each individual's structural and functional connectivity networks were probed using fMRI. These data were supplemented by a battery of testing and questionnaires to analyze 280 behavioral and cognitive traits. Participants provided information about how well they slept, how often they drank alcohol, their language and memory abilities, and their emotional states. Neuroscientists from all over the world have pored over this incredibly rich data set to learn how our brain networks encode who we are.

Using data from hundreds of participants in the Human Connectome Project, we and others have demonstrated that brain-connectivity patterns establish a "fingerprint" that distinguishes each individual. People with strong functional connections among certain regions have an extensive vocabulary, exhibit higher fluid intelligence—helpful for solving novel problems—and are able to delay gratification. They tend to have more education and life satisfaction, as well as better memory and attention. People with weaker functional connections among those same brain areas have lower fluid intelligence, histories of substance use, poor sleep and a decreased capacity for concentration.

Inspired by this research, we showed that the findings could

Decoding 100 Trillion Messages

The Milky Way has hundreds of billions of stars—just a fraction of the 100 trillion connections in our brains that enable us to sense, think and act. To unravel this complexity, network neuroscientists create a map, or "graph," consisting of nodes linked by edges that fit into modules, which are tethered to one another with highly connected nodes called hubs.

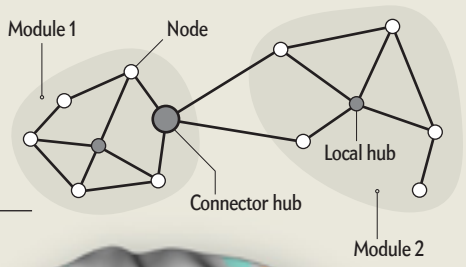

From Modules to Hubs to Thoughts

Collections of nodes form modules that devote themselves to processing vision, attention and motor behaviors, among other tasks **A**. Some of the nodes act as local hubs that link to other nodes in their own module. A node that has many linkages to a lot of modules is known as a connector hub (the type most commonly referenced in this article) **B**. Its diverse connections across the brain's modules are critical for many tasks, particularly complex behaviors **C**.

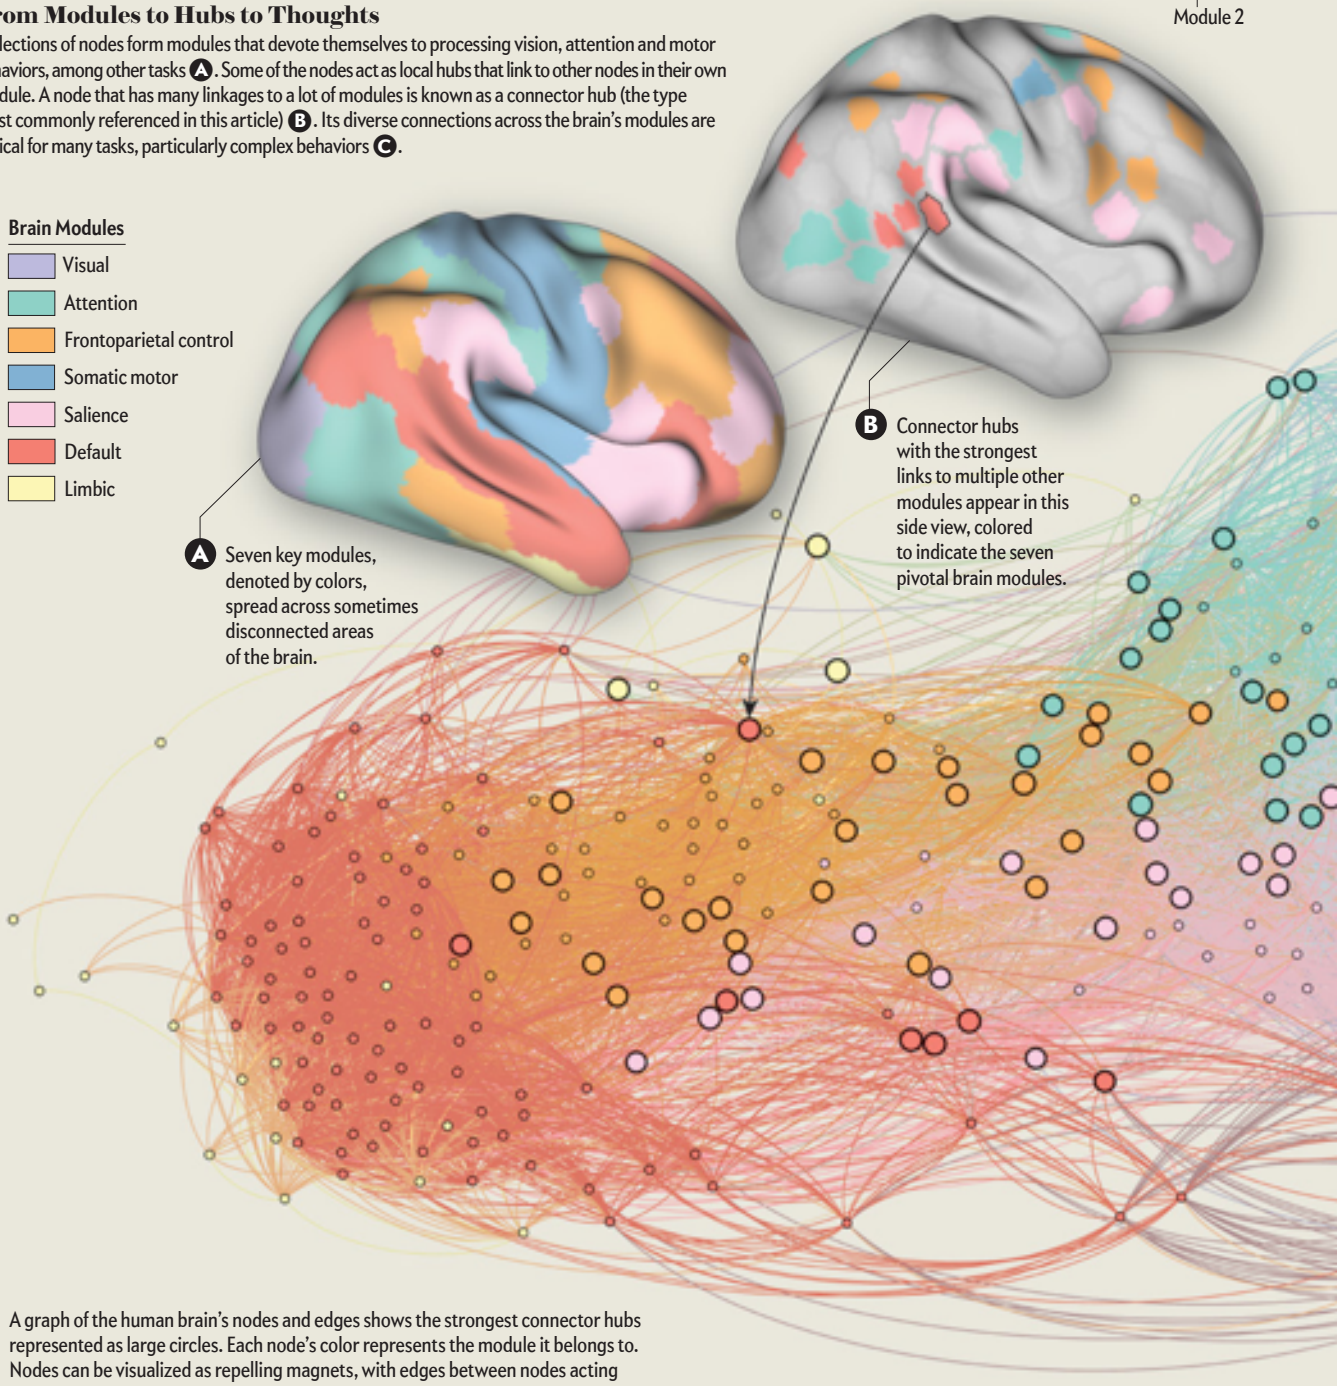

C A graph of the human brain's nodes and edges shows the strongest connector hubs represented as large circles. Each node's color represents the module it belongs to. Nodes can be visualized as repelling magnets, with edges between nodes acting as springs that hold them together. Tightly connected nodes cluster together. Connector hubs occupy the center because they are well connected to all modules.

Putting It Together

Modules for vision, attention and other cognitive functions are dedicated to specific tasks, often represented here by psychological tests. The most active tasks rise to the top. The visual module, for instance, is involved with naming, reading and observing. Many tasks require multiple modules. For example, a mental rotation task recruits both the visual and the attention modules. Some modules are entrusted with more abstract tasks. The frontoparietal module engages in switching tasks or recalling lists. The default mode module attends to subjective emotional states or passive listening when a person is at rest.

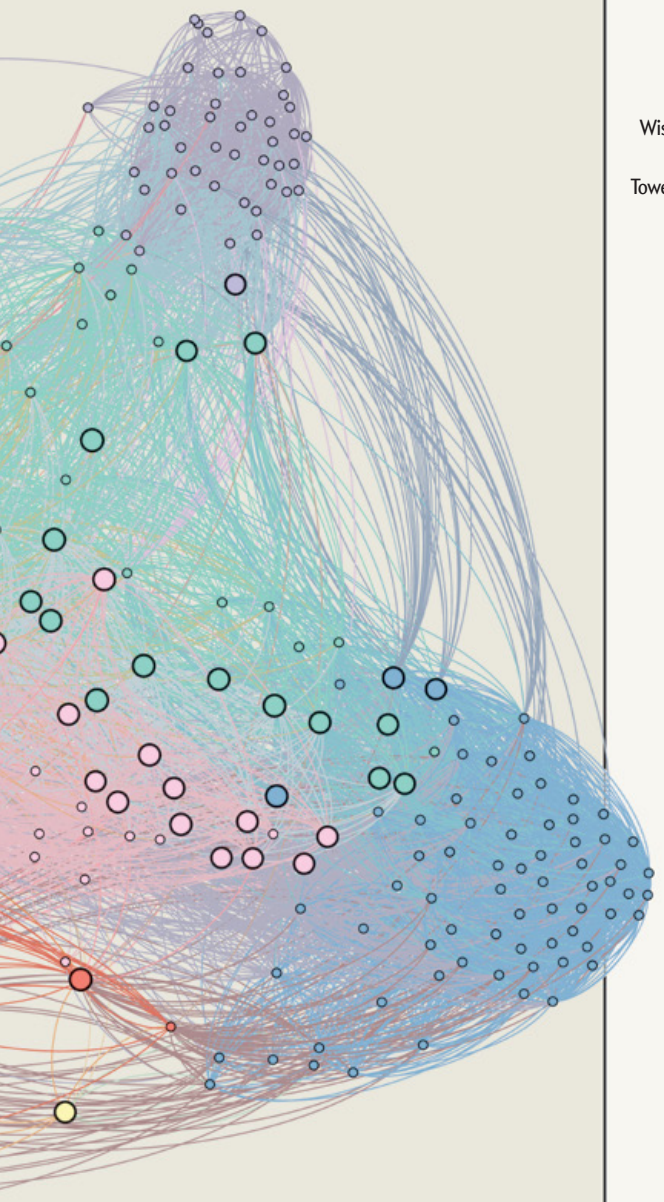

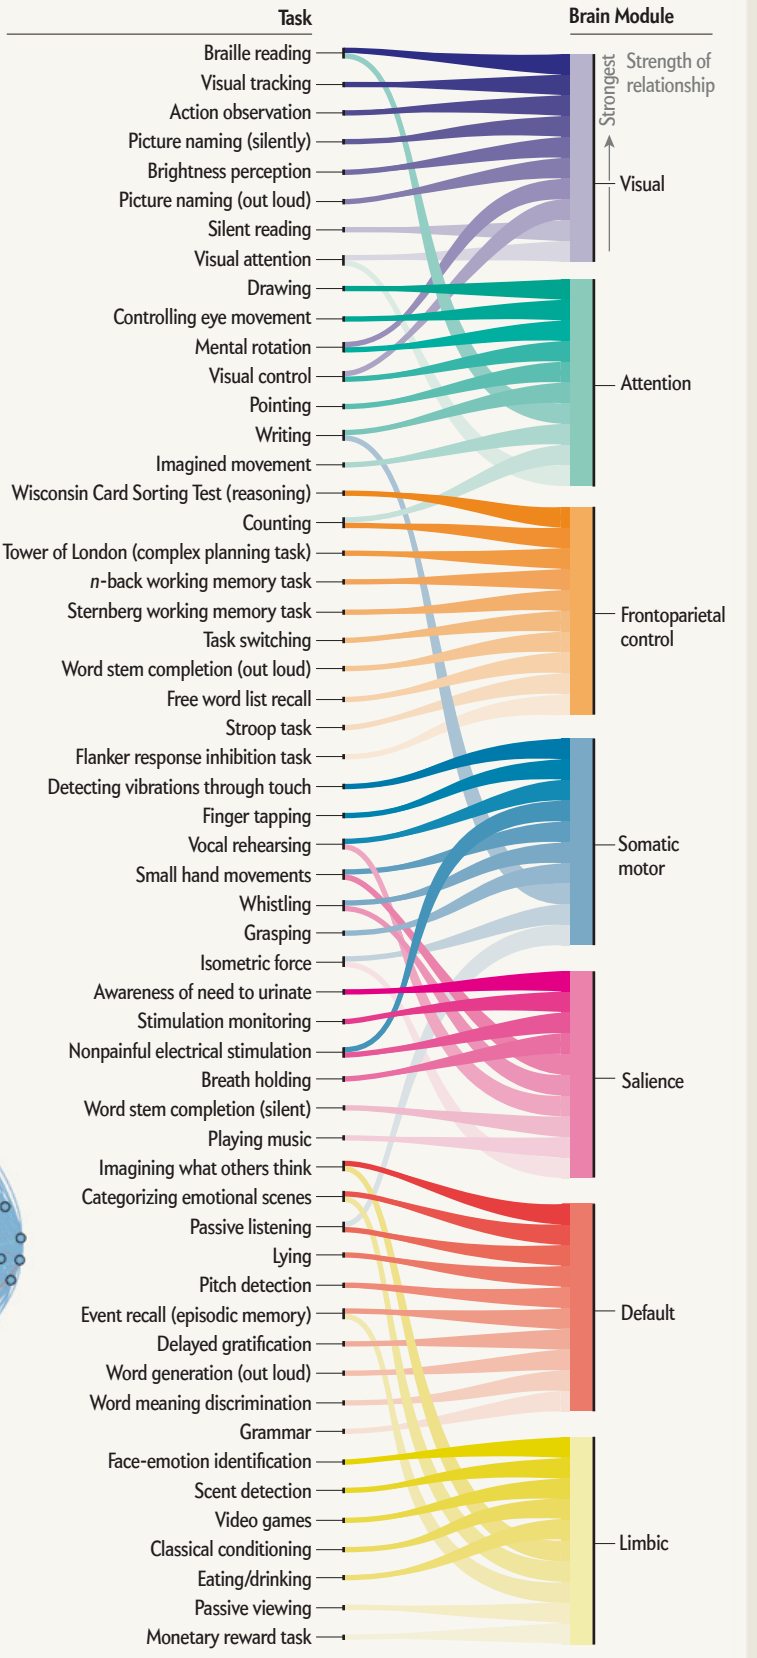

Graphics by Max Bertolero (brains and network diagram) and Jen Christiansen (task chart)

be described by particular patterns among the hub connections. If your brain network has strong hubs with many connections across modules, it tends to have modules that are clearly segregated from one another, and you will perform better on a range of tasks, from short-term memory to mathematics, language or social cognition. Put simply, your thoughts, feelings, quirks, flaws and mental strengths are all encoded by the specific organization of the brain as a unified, integrated network. In sum, it is the music your brain plays that makes you *you*.

The brain's synchronized modules both establish your identity and help to retain it over time. The musical compositions they play appear to always be similar. The likeness could be witnessed when participants in two other studies in the Human Connectome Project engaged in various tasks that involved short-term memory, recognition of the emotions of others, gambling, finger tapping, language, mathematics, social reasoning and a self-induced "resting state" in which they let their mind wander.

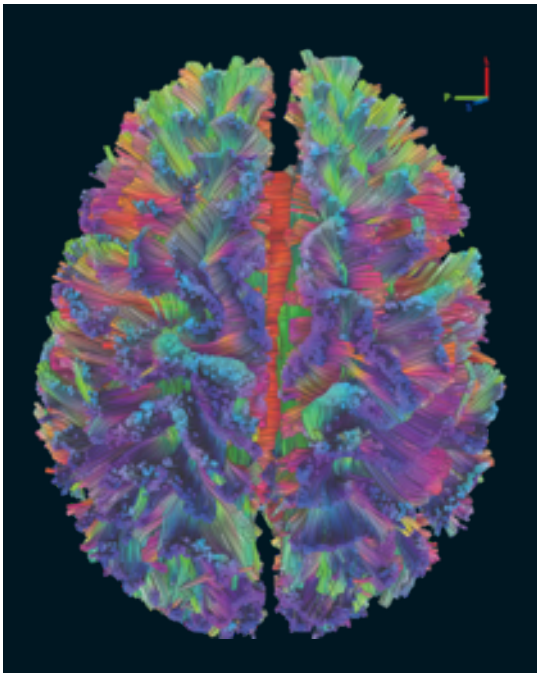

MULTITUDES of white matter connections in this scan are used to model the brain's physical pathways. Functional networks use these structural linkages to carry out an array of cognitive tasks.

Fascinatingly, the networks' functional wiring has more similarities than expected across all these activities. Returning to our analogy, it is not as if the brain plays Beethoven when doing math and Tupac when resting. The symphony in our head comes from the same musician playing the same musical genre. This consistency derives from the fact that the brain's physical pathways, or structural connections, place constraints on the routes over the brain's integrated network that a neural signal can travel. And those pathways delineate how functional connections—the ones, say, for math or language—can be configured. In the music metaphor, a bass drum cannot play the melodic line of a piano.

Changes in the brain's music inevitably occur, just as new arrangements do for orchestral music. Physical connections undergo alterations over the course of months or years, whereas functional connectivity shifts on the order of seconds when a person switches between one mental task and the next.

Transformations in both structural and functional connectivity are important during adolescent brain development, when the finishing touches of the brain's wiring diagram are being refined. This period is of critical importance because the first signs of mental disorders often appear in adolescence or early adulthood.

One area our research relates to is understanding how brain networks develop through childhood and adolescence and into adulthood. These processes are driven by underlying physiological changes, but they are also influenced by learning, exposure to new ideas and skills, an individual's socioeconomic status, and other experiences.

Brain-network modules emerge very early in life, even in the womb, but their connectivity is refined as we grow up. Consistent strengthening of the structural connections to hubs throughout the course of childhood is associated with an increase in the segregation between modules and augmentation of the efficiency with which young people perform executive tasks such as complex reasoning and self-regulation. We have also found that the segregation of modules from one another is more rapid in children who have a higher socioeconomic status, highlighting the key impact of their environment.

Although changes in structural connectivity are slow, the reconfiguration of functional connections can occur quickly, within a few seconds or minutes. These rapid shifts are instrumental for moving between tasks and for the massive amount of learning demanded by even a single task. In a set of studies that we published from 2011 to 2019, we found that networks with modules that can change readily turn up in individuals who have greater executive function and learning capacity.

To better understand what was happening, we used publicly available data from a landmark study known as MyConnectome, in which Stanford University psychology professor Russell Poldrack personally underwent imaging and cognitive appraisals three times a week for more than a year. Although modules are mostly autonomous and segregated, at times the brain will spontaneously reorganize its connections. This property, called functional network flexibility, lets a node with strong functional connections within a module suddenly establish many connections to a different module, changing the flow of information through the network. Using data from this study, we found that the rerouting of a network's connections changes from day to day in a manner that matches positive mood, arousal and fatigue. In healthy individuals, such network flexibility correlates with better cognitive function.

DISSONANT NOTES

THE CONFIGURATION of brain connections also reflects one's mental health. Aberrant connectivity patterns accompany depression, schizophrenia, Alzheimer's, Parkinson's, autism spectrum disorder, attention deficit disorder, dementia and epilepsy.

Most mental illnesses are not confined to one area of the brain. The circuitry affected in schizophrenia extends quite widely across the entire organ. The so-called disconnectivity hypothesis for schizophrenia holds that there is nothing abnormal about the individual modules. Instead the disarray relates to an overabundance of connections between modules.

In a healthy brain, modules are mostly autonomous and segregated, and the ability to bring about flexible changes in network connections is beneficial for cognitive functioning—with-

Dani S. Bassett and Matthew Cieslak/University of Pennsylvania

in certain limits. In our research, we found that in the brains of people with schizophrenia and their first-degree relatives, there is an overabundance of flexibility in how networks reconfigure themselves. Auditory hallucinations might result when nodes unexpectedly switch links between speech and auditory modules. The uninvited mix can result in what seem to be the utterings of voices in one's head.

Like schizophrenia, major depressive disorder is not caused by a single abnormal brain region. Three specific modules appear to be affected in depression: the frontoparietal control, salience and default mode modules. In fact, the symptoms of depression—emotional disinhibition, altered sensitivity to emotional events and rumination—map to these modules.

As a result, normal communication among the three modules becomes destabilized. Activities from module to module typically tug back and forth to balance the cognitive processing of sensory inputs with more introspective thoughts. In depression, though, the default mode dominates, and the afflicted person lapses into ruminative thought. The music of the brain thus becomes increasingly unbalanced, with one family of instruments governing the symphony. These observations have broadened our understanding of the network properties of depression to the extent that a connectivity pattern in a brain can allow us to diagnose certain subtypes of the disorder and determine which areas should be treated with electrical-stimulation technology.

NETWORKS EVOLVE

BESIDES STUDYING DEVELOPMENT, network neuroscientists have begun to ask why brain networks have taken their present form over tens of thousands of years. The areas identified as hubs are also the locations in the human brain that have expanded the most during evolution, making them up to 30 times the size they are in macaques. Larger brain hubs most likely permit greater integration of processing across modules and so support more complex computations. It is as if evolution increased the number of musicians in a section of the orchestra, fostering more intricate melodies.

Another way neuroscientists have explored these questions is by creating computer-generated networks and subjecting them to evolutionary pressures. We have probed the evolutionary origins of hubs. This exercise started with a network in which all edges were placed uniformly at random. Next, the network was rewired, mimicking natural selection to form segregated modules and display a property known in network science as small-worldness, in which paths form to let distant network nodes communicate with surprising ease. Thousands of such networks then evolved, each of which ultimately contained hubs strongly connected to multiple modules but also tightly interconnected to one another, forming what is called a club. Nothing in the selection process explicitly selected for a club of hubs—they simply emerged from this iterative process.

This simulation demonstrates that one potential path to evolving a brain capable of exchanging information among modules requires hubs with strong connections. Notably, real networks—brains, airports, power grids—also have durable, tightly interconnected hubs, exactly as predicted by evolutionary experiments. That observation does not mean evolution necessarily occurred in the same way as the simulation, but it shows a possible means by which one of nature's tricks might operate.

STATES OF MIND

WHEN NOBEL PRIZE–WINNING PHYSICIST Richard Feynman died in 1988, his blackboard read, "What I cannot create, I do not understand." He created a beautiful aphorism, yet it misses a pivotal idea: it should be revised to "What I cannot create *and control,* I do not understand." Absent such control, we still know enough to enjoy a symphony, even if we do not qualify to be the conductor.

When it comes to the brain, we have a basic understanding of its form and the importance of its network architecture. We know that our brain determines who we are, but we are just beginning to understand how it all happens. To rephrase mathematician Pierre-Simon Laplace's explanation of determinism and mechanics and apply it to the brain, one's present brain, and so one's mental state, can be thought of as a compilation of past states that can be used to predict the future. A neuroscientist who knew all the principles of brain function and everything about someone's brain could predict that person's mental conditions—the future, as well as the past, would be present inside the person's mind.

This knowledge could be used to prevent pain and suffering, given that many mental illnesses are associated with network abnormalities. With enough engineering ingenuity, we may develop implanted devices that alter brain networks or even generate new ones to prevent the disorganization associated with mental disorders from occurring in the first place. Such an achievement would enable us to treat diseases and help to restore brain function after stroke or injury or, potentially, enhance function in healthy individuals.

Before those futuristic scenarios materialize, two major gaps must be filled: we need to know more about how personal genetics, early-life development and environment determine one's brain's structure and how that structure leads to functional capacities. Neuroscientists have some knowledge from the human genome about the structure that gives rise to functional networks but still need to learn precisely how this process occurs. We are starting to grasp the way brain networks develop and are shaped by the environment, but we are not close to explaining the entire complexity of this process. The brain's wiring, its structural connectivity, constrains how various modules interact with one another, but our knowledge remains limited. As we fill in these gaps, chances improve for interventions to guide brain functioning into healthy trajectories.

What holds us back, for the moment, is our still blurry vision of the brain—it is as if we are outside the concert hall and have seen only sketches of the instruments. Inside each brain region that neuroscientists study are millions of neurons firing every millisecond, and we are able just to indirectly measure their average activity levels every second or so. Thus far we can roughly identify the human brain's structural connections. Luckily, scientists and engineers have taken steps to deliver ever clearer data that will enable a deeper look into perhaps the most complex network in the known universe: your brain. SA

Max Bertolero is director of research and development at Turing Medical in St. Louis, Mo. He received a doctorate in systems neuroscience from the University of California, Berkeley, and undergraduate degrees in philosophy and psychology from Columbia University.

Dani S. Bassett is J. Peter Skirkanich Professor of Bioengineering at the University of Pennsylvania. Their lab studies networks in physical, biological and social systems. In 2014 Bassett became a MacArthur Fellow.

OUR INNER UNIVERSES

Reality is constructed by the brain, and no two brains are exactly alike

By Anil K. Seth

Illustration by Brook VanDevelder

"We do not see things as they are, we see them as we are."
—from *Seduction of the Minotaur,* by Anaïs Nin (1961)

ON THE 10TH OF APRIL 2019 POPE FRANCIS, PRESIDENT SALVA KIIR OF South Sudan and former rebel leader Riek Machar sat down together for dinner at the Vatican. They ate in silence, the start of a two-day retreat aimed at reconciliation from a civil war that had killed some 400,000 people since 2013. At about the same time in my laboratory at the University of Sussex in England, Ph.D. student Alberto Mariola was starting to work on an experiment in which volunteers experience being in a room they believe is there but is not. In psychiatry clinics across the globe, people arrive complaining that things no longer seem "real" to them, whether it is the world around them or their own selves. In the fractured societies in which we live, what is real—and what is not—seems to be increasingly up for grabs. Warring sides may experience and believe in different realities. Perhaps eating together in silence can help because it offers a small slice of reality that can be agreed on, a stable platform on which to build further understanding.

We need not look to war and psychosis to find radically different inner universes. In 2015 a badly exposed photograph of a dress tore across the Internet, dividing the world into those who saw it as blue and black (me included) and those who saw it as white and gold (half my lab). Those who saw it one way were so convinced they were right—that the dress truly was blue and black or white and gold—that they found it almost impossible to believe that others might perceive it differently.

We all know that our perceptual systems are easy to fool. The popularity of visual illusions is testament to this phenomenon. Things seem to be one way, and they are revealed to be another: two lines appear to be different lengths, but when measured they are exactly the same; we see movement in an image we know to be still. The story usually told about illusions is that they exploit quirks in the circuitry of perception, so that what we perceive deviates from what is there. Implicit in this story, however, is the assumption that a properly functioning perceptual system will render to our consciousness things precisely as they are.

The deeper truth is that perception is never a direct window onto an objective reality. All our perceptions are active constructions, brain-based best guesses at the nature of a world that is forever obscured behind a sensory veil. Visual illusions are fractures in the Matrix, fleeting glimpses into this deeper truth.

Take, for example, the experience of color—say, the bright red of the coffee mug on my desk. The mug really does seem to be red: its redness seems as real as its roundness and its solidity. These features of my experience seem to be truly existent properties of the world, detected by our senses and revealed to our mind through the complex mechanisms of perception.

Yet we have known since Isaac Newton that colors do not exist out there in the world. Instead they are cooked up by the brain from mixtures of different wavelengths of colorless electromagnetic radiation. Colors are a clever trick that evolution has hit on to help the brain keep track of surfaces under changing lighting conditions. And we humans can sense only a tiny slice of the full electromagnetic spectrum, nestled between the lows of infrared and the highs of ultraviolet. Every color we perceive, every part of the totality of each of our visual worlds, comes from this thin slice of reality.

Just knowing this is enough to tell us that perceptual experience cannot be a comprehensive representation of an external objective world. It is both less than that and more than that. The reality we experience—the way things *seem*—is not a direct reflection of what is actually out there. It is a clever construction by the brain, for the brain. And if my brain is different from your brain, my reality may be different from yours, too.

THE PREDICTIVE BRAIN

IN PLATO'S *Allegory of the Cave,* prisoners are chained to a blank wall all their lives, so that they see only the play of shadows cast by objects passing by a fire behind them, and they give the shadows names because for them the shadows are what is real. A thousand years later, but still a thousand years ago, Arabian scholar Ibn al-Haytham wrote that perception, in the here and now, depends on processes of "judgment and inference" rather than involving direct access to an objective reality. Hundreds of years later again Immanuel Kant realized that the chaos of unrestricted sensory data would always remain meaningless without being given structure by preexisting conceptions or "beliefs," which for him included a priori frameworks such as space and time. Kant's term "noumenon" refers to a "thing in itself"—*Ding an sich*—an objective reality that will always be inaccessible to human perception.

Today these ideas have gained a new momentum through an influential collection of theories that turn on the idea that the brain is a kind of prediction machine and that perception of the world—and of the self within it—is a process of brain-based predictions about the causes of sensory signals.

These new theories are usually traced to German physicist and physiologist Hermann von Helmholtz, who in the late 19th century proposed that perception is a process of unconscious inference. Toward the end of the 20th century Helmholtz's notion was taken up by cognitive scientists and artificial-intelligence researchers, who reformulated it in terms of what is now generally known as predictive coding or predictive processing.

The central idea of predictive perception is that the brain is attempting to figure out what is out there in the world (or in here, in the body) by continually making and updating best guesses about the causes of its sensory inputs. It forms these best guesses by combining prior expectations or "beliefs" about the world, together with incoming sensory data, in a way that takes into account how reliable the sensory signals are. Scientists usually conceive of this process as a form of Bayesian inference, a framework that specifies how to update beliefs or best guesses with new data when

both are laden with uncertainty.

In theories of predictive perception, the brain approximates this kind of Bayesian inference by continually generating predictions about sensory signals and comparing these predictions with the sensory signals that arrive at the eyes and the ears (and the nose and the fingertips and all the other sensory surfaces on the outside and inside of the body). The differences between predicted and actual sensory signals give rise to so-called prediction errors, which are used by the brain to update its predictions, readying it for the next round of sensory inputs. By striving to minimize sensory-prediction errors everywhere and all the time, the brain implements approximate Bayesian inference, and the resulting Bayesian best guess is what we perceive.

To understand how dramatically this perspective shifts our intuitions about the neurological basis of perception, it is helpful to think in terms of bottom-up and top-down directions of signal flow in the brain. If we assume that perception is a direct window onto an external reality, then it is natural to think that the content of perception is carried by bottom-up signals—those that flow from the sensory surfaces inward. Top-down signals might contextualize or finesse what is perceived, but nothing more. Call this the "how things seem" view because it seems as if the world is revealing itself to us directly through our senses.

POORLY EXPOSED photograph of a dress appears blue and black to some people, white and gold to others.

The prediction machine scenario is very different. Here the heavy lifting of perception is performed by the top-down signals that convey perceptual predictions, with the bottom-up sensory flow serving only to calibrate these predictions, keeping them yoked, in some appropriate way, to their causes in the world. In this view, our perceptions come from the inside out just as much as, if not more than, from the outside in. Rather than being a passive registration of an external objective reality, perception emerges as a process of active construction—a controlled hallucination, as it has come to be known.

Why controlled hallucination? People tend to think of hallucination as a kind of false perception, in clear contrast to veridical, true-to-reality, normal perception. The prediction machine view suggests instead a continuity between hallucination and normal perception. Both depend on an interaction between top-down, brain-based predictions and bottom-up sensory data, but during hallucinations, sensory signals no longer keep these top-down predictions appropriately tied to their causes in the world. What we call hallucination, then, is just a form of uncontrolled perception, just as normal perception is a controlled form of hallucination.

This view of perception does not mean that nothing is real. Writing in the 17th century, English philosopher John Locke made an influential distinction between "primary" and "secondary" qualities. Primary qualities of an object, such as solidity and occupancy of space, exist independently of a perceiver. Secondary qualities, in contrast, exist only in relation to a perceiver—color is a good example. This distinction explains why conceiving of perception as controlled hallucination does not mean it is okay to jump in front of a bus. The bus has primary qualities of solidity and space occupancy that exist independently of our perceptual machinery and that can do us injury. It is the way in which the bus appears to us that is a controlled hallucination, not the bus itself.

TRIPPING IN THE LAB

A GROWING BODY of evidence supports the idea that perception is controlled hallucination, at least in its broad outlines. A 2015 study by Christoph Teufel of Cardiff University in Wales and his colleagues offers a striking example. In this study, the ability to recognize so-called two-tone images was evaluated in patients with early-stage psychosis who were prone to hallucinations.

Take a look at the photograph on the next page—a sample of a two-tone image. Probably all you will see is a bunch of black-and-white splotches. Now look at the image on page 17. Then have another look at the first photo; it ought to look rather different. Where previously there was a splotchy mess, there are now distinct objects, and something is happening.

What I find remarkable about this exercise is that in your second examination of the first image, the sensory signals arriving at your eyes have not changed at all from the first time you saw it. All that has changed are your brain's predictions about the causes of these sensory signals. You have acquired a new high-level perceptual expectation, and this changes what you consciously see.

If you show people many two-tone images, each followed by the full picture, they might subsequently be able to identify a good proportion of two-tone images, though not all of them. In Teufel's study, people with early-stage psychosis were better at recognizing two-tone images after having seen the full image than were healthy control subjects. In other words, being hallucination-prone went along with perceptual priors having a stronger effect on perception. This is exactly what would be expected if hallucinations in psychosis depended on an overweighting of perceptual priors so that they overwhelmed sensory-prediction errors, unmooring perceptual best guesses from their causes in the world.

Recent research has revealed more of this story. In a 2021 study, Biyu He of New York University and her colleagues had neurosurgical patients look at ambiguous images, such as a Necker cube, that constantly flip between two different appearances even though the sensory input remains the same. By analyzing

Swiked.tumblr.com

TWO-TONE IMAGE looks like a mess of black-and-white splotches, until you see the full image (*opposite page*).

the signals recorded from within the patients' brains, they discovered that information flowed more strongly in a top-down, "inside-out" direction when the perceived appearance was consistent with the patients' biases, as would be expected if perceptual predictions were strong in this case. And when the perceived appearance was inconsistent with preexisting biases, information flow was stronger in the bottom-up direction, suggesting a "prediction error" signal. This is an exciting new development in mapping the brain basis of controlled hallucinations.

In my lab we have taken a different approach to exploring the nature of perception and hallucination. Rather than looking into the brain directly, we decided to simulate the influence of overactive perceptual priors using a unique virtual-reality setup masterminded by VR guru Keisuke Suzuki. We call it, with tongue firmly in cheek, the "hallucination machine."

Using a 360-degree camera, we first recorded panoramic video footage of a busy square in the University of Sussex campus on a Tuesday at lunchtime. We then processed the footage through an algorithm based on Google's AI program DeepDream to generate a simulated hallucination. What happens is that the algorithm takes a so-called neural network—one of the workhorses of AI—and runs it backward. The network we used had been trained to recognize objects in images, so if you run it backward, updating the network's input instead of its output, the network effectively projects what it "thinks" is there onto and into the image. Its predictions overwhelm the sensory inputs, tipping the balance of perceptual best guessing toward these predictions. Our particular network was good at classifying different breeds of dogs, so the video became unusually suffused with dog presences.

Many people who have viewed the processed footage through the VR headset have commented that the experience is rather reminiscent not of the hallucinations of psychosis but of the exuberant phenomenology of psychedelic trips.

More recently, we have implemented the hallucination machine in different ways to simulate different kinds of altered visual experience. By extending our algorithm to include two coupled neural networks—a "discriminator network" much like the one in our original study and a "generator network" that has been trained to reproduce ("generate") its input image—we have been able to model different types of hallucination. For example, we have modeled the complex hallucinatory experiences reported by people with Parkinson's disease and some forms of dementia; the patterned, geometric hallucinations that occur after the loss of foveal vision, as happens in Charles Bonnet syndrome; and a range of psychedeliclike hallucinations. We hope that by understanding hallucinations better, we will be able to understand normal experience better, too, because predictive perception is at the root of all our perceptual experience.

THE PERCEPTION OF REALITY

ALTHOUGH THE HALLUCINATION machine is undoubtedly trippy, people who experience it are fully aware that what they are experiencing is not real. Indeed, despite rapid advances in VR technology and computer graphics, no current VR setup delivers an experience sufficiently convincing to be indistinguishable from reality.

This is the challenge we took up when designing a new "substitutional reality" setup at Sussex—the one we were working on when Pope Francis convened the retreat with Salva Kiir and Riek Machar. Our aim was to create a system in which volunteers would experience an environment as being real—and believe it to be real—when in fact it was not real.

The basic idea is simple. We again prerecorded some panoramic video footage, this time of the interior of our VR lab rather than of an outside campus scene. People coming to the lab are invited to sit on a stool in the middle of the room and to put on a VR headset that has a camera attached to the front. They are encouraged to look around the room and to see the room as it actually is, via the camera. At some point, without telling them, we switch the feed so that the headset displays not the live real-world scene but rather the prerecorded panoramic video. Most people in this situation continue to experience what they are seeing as real even though it is now a fake prerecording. (This is actually very tricky to pull off in practice—it requires careful color balancing and alignment to avoid people noticing any difference that would tip them off to the shift.)

I find this result fascinating because it shows that it is possible to have people experience an unreal environment as being fully real. This demonstration alone opens new frontiers for VR research: we can test the limits of what people will experience, and believe, to be real. It also allows us to investigate how experiencing things as being real can affect other aspects of perception. In one of our most recent experiments, we aimed to find out whether people are worse at detecting unexpected changes in the room when they believe that what they are experiencing is real. If things do turn out this way (the study was heavily delayed by

the global pandemic), that finding would support the idea that the perception of things as being real itself acts as a high-level prior that can substantively shape our perceptual best guesses, affecting the contents of what we perceive.

PERCEPTUAL SHIFT: Viewing this photograph changes what one consciously sees in the two-tone image (*opposite page*).

THE REALITY OF REALITY

THE IDEA THAT THE WORLD of our experience might not be real is an enduring trope of philosophy and science fiction, as well as of late-night pub discussions. Neo in *The Matrix* takes the red pill, and Morpheus shows him how what he thought was real is an elaborate simulation, while the real Neo lies prone in a human body farm, a brain-in-a-vat power source for a dystopian AI. Philosopher Nick Bostrom of the University of Oxford has famously argued, based largely on statistics, that we are likely to be living inside a computer simulation created in a posthuman age. I disagree with this argument in part because it assumes that consciousness can be simulated—I do not think that this is a safe assumption—but it is thought-provoking nonetheless.

Although these chunky metaphysical topics are fun to chew on, they are probably impossible to resolve. Instead what we have been exploring throughout this article is the relation between appearance and reality in our conscious perceptions, where part of this appearance is the appearance of being real itself.

The central idea here is that perception is a process of active interpretation geared toward adaptive interaction with the world via the body rather than a re-creation of the world within the mind. The contents of our perceptual worlds are controlled hallucinations, brain-based best guesses about the ultimately unknowable causes of sensory signals. For most of us, most of the time, these hallucinations are experienced as real. As Canadian rapper and science communicator Baba Brinkman suggested to me, when we agree about our hallucinations, maybe that's what we call reality.

But we do not always agree, and we do not always experience things as real. People with dissociative psychiatric conditions such as derealization or depersonalization syndrome report that their perceptual worlds, even their own selves, lack a sense of reality. Some kinds of hallucination, various psychedelic hallucinations among them, combine a sense of unreality with perceptual vividness, as does lucid dreaming. People with synesthesia consistently have additional sensory experiences, such as perceiving colors when viewing black letters, which they recognize as not real. Even with normal perception, if you look directly at the sun you will experience the subsequent retinal afterimage as not being real. There are many such ways in which we experience our perceptions as not fully real.

What this means to me is that the property of realness that attends most of our perceptions should not be taken for granted. It is another aspect of the way our brain settles on its Bayesian best guesses about its sensory causes. One might thus ask what purpose it serves. Perhaps the answer is that a perceptual best guess that includes the property of being real is usually more fit for purpose—that is, better able to guide behavior—than one that does not. We will behave more appropriately with respect to a coffee cup, an approaching bus or our partner's mental state when we experience it as really existing.

But there is a trade-off. As illustrated by the dress illusion, when we experience things as being real, we are less able to appreciate that our perceptual worlds may differ from those of others. (A popular explanation for the differing perceptions of the garment holds that people who spend most of their waking hours in daylight see it as white and gold; night owls, who are mainly exposed to artificial light, see it as blue and black.) And even if these differences start out small, they can become entrenched and reinforced as we proceed to harvest information differently, selecting sensory data that are best aligned with our individual emerging models of the world and then updating our perceptual models based on these biased data. We are all familiar with this process from the echo chambers of social media and the newspapers we choose to read. I am suggesting that the same principles apply also at a deeper level, underneath our sociopolitical beliefs, right down to the fabric of our perceptual realities. They may even apply to our perception of being a self—the experience of being me or of being you—because the experience of being a self is itself a perception.

This is why understanding the constructive, creative mechanisms of perception has unexpected social relevance. Perhaps once we can appreciate the diversity of experienced realities scattered among the billions of perceiving brains on this planet, we will find new platforms on which to build a shared understanding and a better future—whether between sides in a civil war, followers of different political parties, or two people sharing a house and faced with washing the dishes. SA

Anil K. Seth is a professor of cognitive and computational neuroscience at the University of Sussex in England. His research focuses on the biological basis of consciousness, and he is author of *Being You—A New Science of Consciousness* (Dutton, 2021).

Richard Armstrong/EyeEm/Getty Images (*opposite page* and *this page*)

The Brain Electric

Electrodes that stimulate brain tissue reveal the topography of conscious experience

By Christof Koch

Illustration by Zara Picken

Consider the following experiences:

- You're headed toward a storm that's a couple of miles away, and you've got to get across a hill. You ask yourself: "How am I going to get over that, through that?"
- You see little white dots on a black background, as if looking up at the stars at night.
- You look down at yourself lying in bed from above but see only your legs and lower trunk.

These may seem like idiosyncratic events drawn from the vast universe of perceptions, sensations, memories, thoughts and dreams that make up our daily stream of consciousness. In fact, each one was evoked by directly stimulating the brain with an electrode. As American poet Walt Whitman intuited in his poem "I Sing the Body Electric," these anecdotes illustrate the intimate relationship between the body and its animating soul. The brain and the conscious mind are as inexorably linked as the two sides of a coin.

Recent clinical studies have uncovered some of the laws and regularities of conscious activity, findings that have occasionally proved to be paradoxical. They show that brain areas involved in conscious perception have little to do with thinking, planning and other higher cognitive functions. Neuroengineers have been working to turn these insights into technologies to replace lost cognitive function

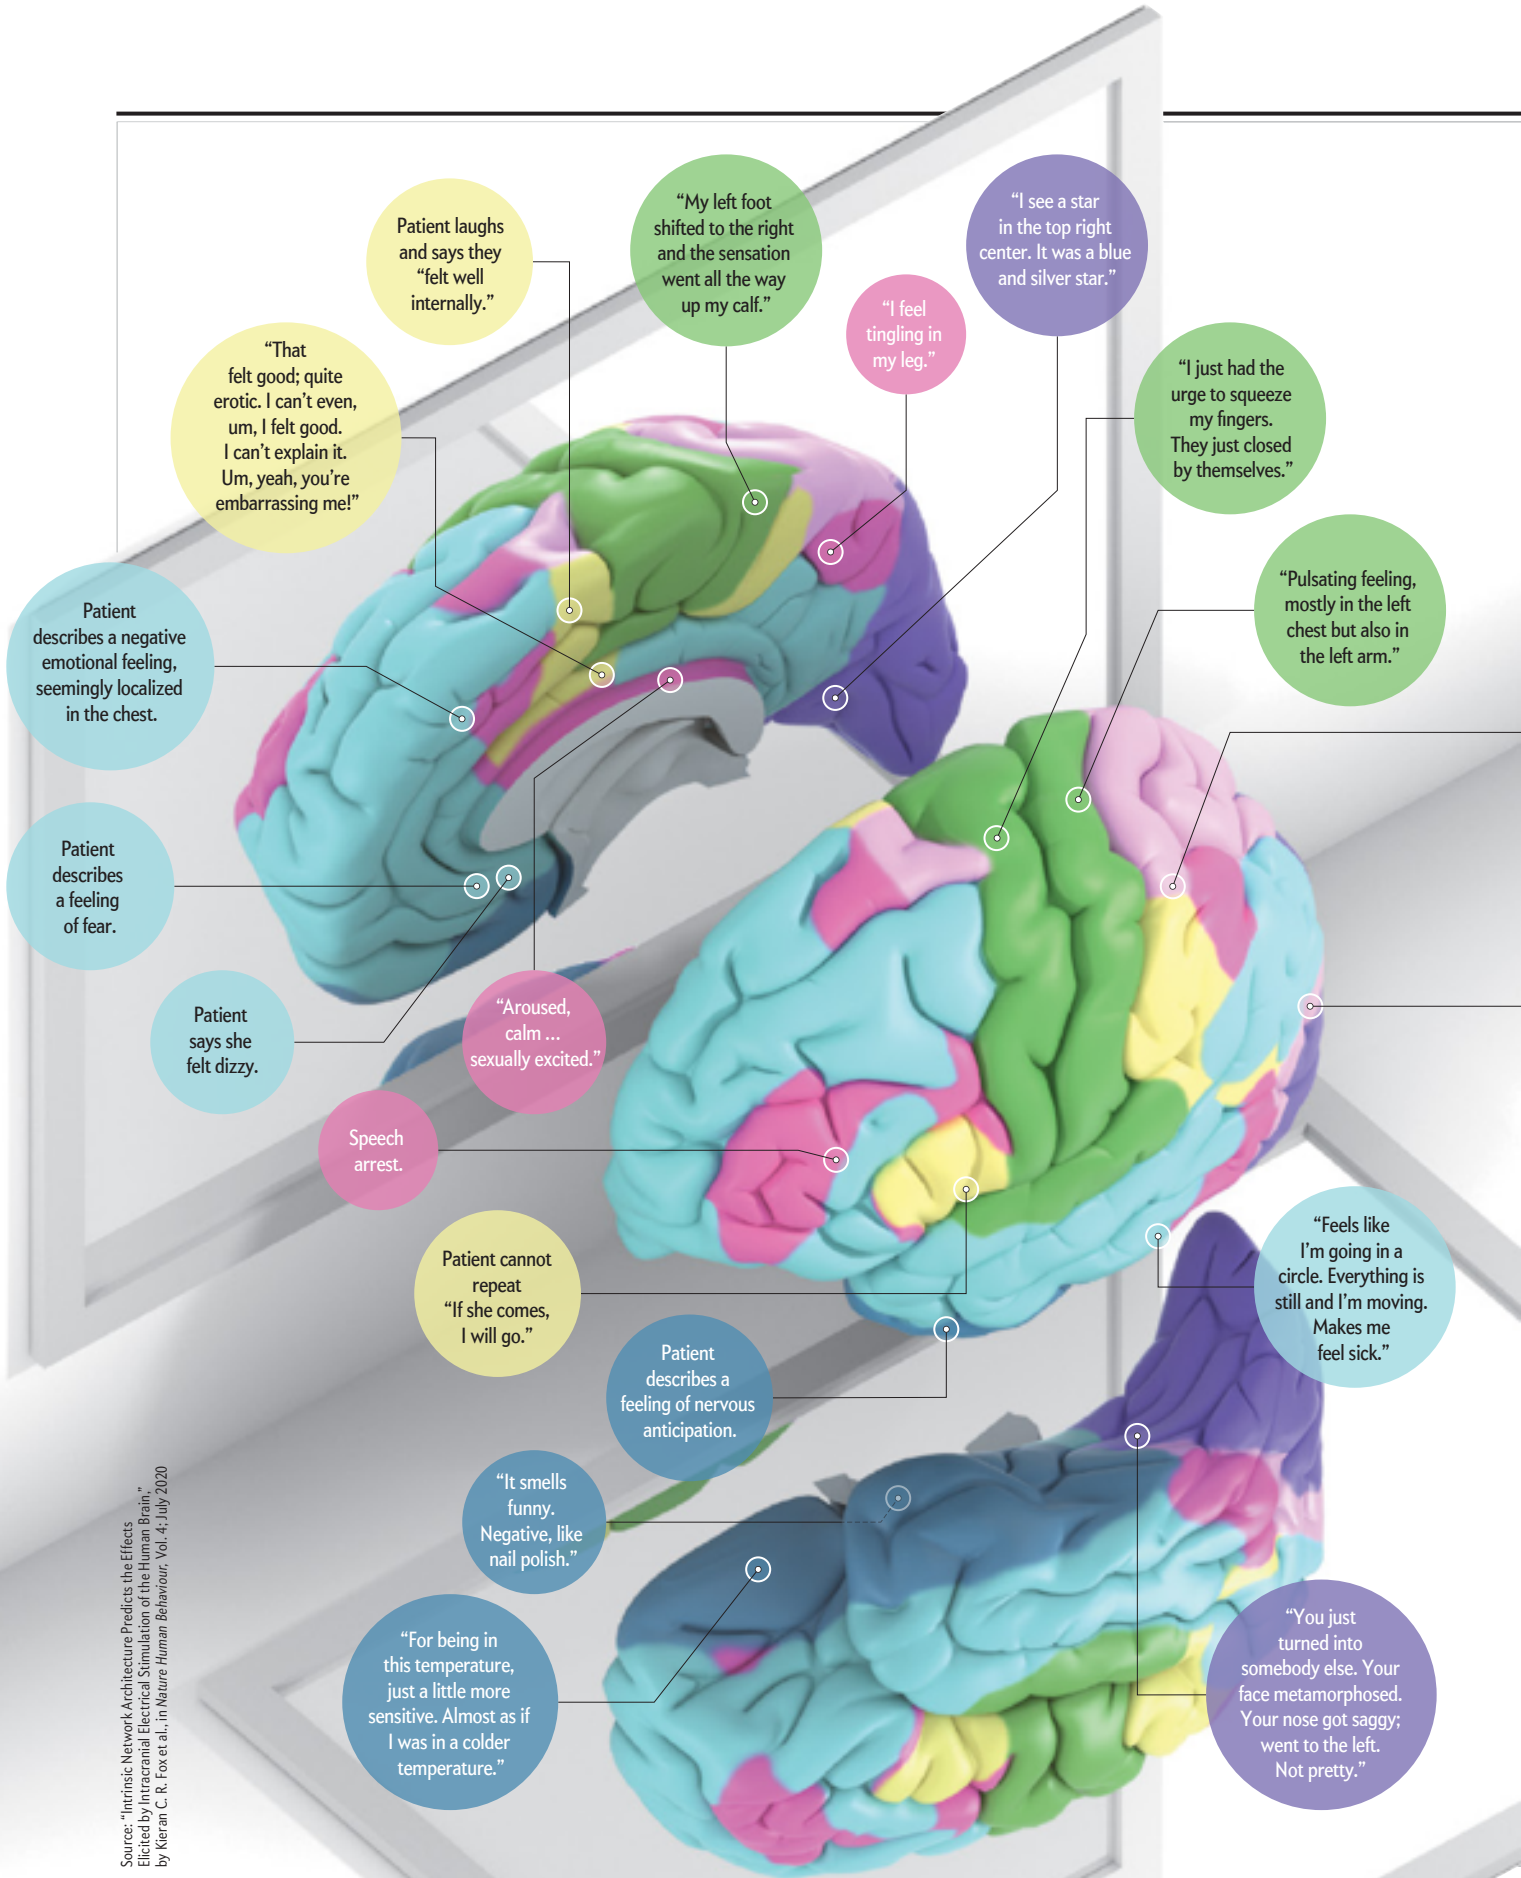

Source: "Intrinsic Network Architecture Predicts the Effects Elicited by Intracranial Electrical Stimulation of the Human Brain," by Kieran C. R. Fox et al., in *Nature Human Behaviour*, Vol. 4, July 2020

Where Experiences Live in the Brain

An atlas published in the summer of 2020 compiled the verbal reports of people with epilepsy whose cortical areas were stimulated with electrodes during surgery. What they felt and perceived varied depending on which brain region was stimulated. All of the 1,537 locations in these 67 patients where current was applied were mapped onto a digital brain model, a simplified version of which is depicted here. When stimulated at these sites, patients talked about their experiences.

"Just really couldn't move (my fingers) too much; lost the motion. The hand felt a little tight, but the thumb was out of commission."

"I felt like my arms were moving but they weren't. I felt side-to-side movements, like floating in the air."

Colors of Cognition

Stimulation was applied to varied brain networks and regions.

- Somatomotor
- Visual
- Dorsal attention
- Salience
- Frontoparietal
- Limbic
- Default

Points of Stimulation

Researchers placed dozens of electrodes onto the surface of the cerebral cortex accessed through an opening in the skull.

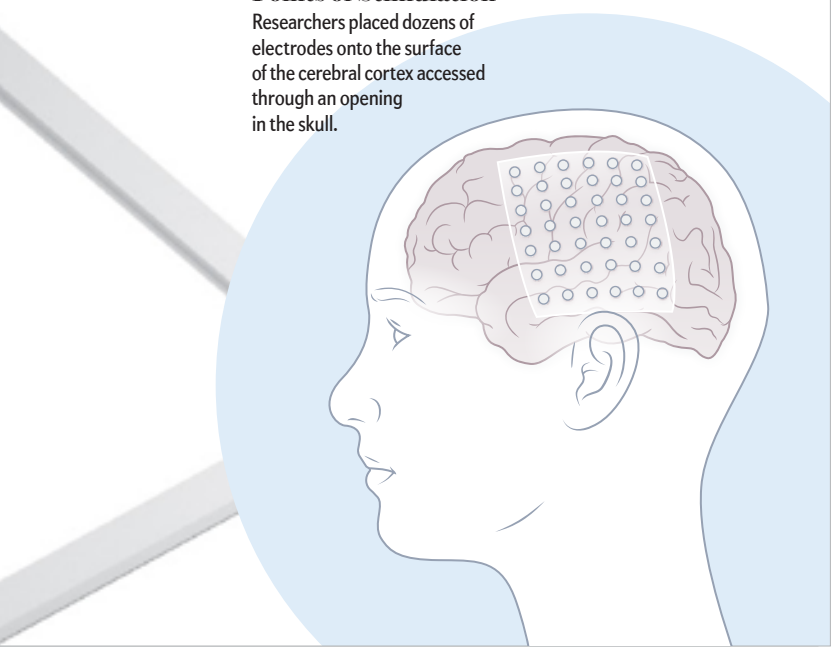

Graphic by AXS Biomedical Animation Studio

and, in the more distant future, to enhance sensory, cognitive or memory capacities. For example, a brain-machine interface can now provide completely blind people with limited abilities to perceive light. These tools, however, also reveal the difficulties of fully restoring sight or hearing. They underline even more the snags in the way of sci-fi-like enhancements that would enable access to the brain as if it were a computer storage drive.

ANIMAL ELECTRICITY

NERVOUS SYSTEMS operate on the flow of electric currents through ultradense and hyperconnected networks of switching elements. Countless physicians and scientists have studied this problem over the past two and a half centuries, beginning with Italian physician Luigi Galvani, who in the late 18th century connected a freshly killed frog to a long metal wire. By pointing the wire toward the sky during a thunderstorm, he made the frog's leg jump and twitch with each flash of lightning. Galvani's investigations revealed that nerve fibers transmitted "animal electricity," which is no different in kind from the "atmospheric electricity" that Benjamin Franklin discovered with his kite experiments in Philadelphia in 1752. In 1802 Galvani's nephew Giovanni Aldini electrically stimulated the exposed brain of a decapitated prisoner during a public event. A jaw quivered. An eye opened. The spectacle may have helped to inspire Mary Shelley to write the classic 1818 gothic novel *Frankenstein*.

Subsequent animal studies demonstrated that exciting particular brain regions triggered movements in specific muscles and limbs. These investigations led to the discovery of the motor cortex in the 1870s. In 1874 American physician Robert Bartholow performed the first direct brain stimulation of a conscious patient—a pioneering act clouded in ethical controversy because it caused the patient pain and probably hastened her death. Intracranial electrical stimulation (iES) was refined over the following decades. It became part of the neurosurgeon's toolbox thanks to the ground-breaking work of Wilder Penfield of the Montreal Neurological Institute, who between the 1930s and the 1950s used iES to map cortical areas that process motor or sensory functions.

In some people with epilepsy, drugs fail to adequately control the number or severity of seizures. Neurosurgery becomes an option if those seizures originate in a delimited neighborhood in the cortex—the outermost layer of the brain involved in perception, motor control, speech, reasoning, and so on—or in closely related structures, such as the hippocampus. Uncontrolled hyperexcitability starts because of local faulty wiring. It can grow and eventually engulf the rest of the brain. How much tissue to remove is a dilemma: cut too little, and seizures may continue; cut too much, and the patient may lose the ability to speak, see or walk. Surgeons must avoid areas of the cortex that are crucial for everyday behavior, such as the primary auditory, visual, somatosensory and motor cortices and the regions controlling understanding and producing speech, areas known as the eloquent cortex.

iES is brought in as a means to look for tissue that needs preserving. Neurosurgeons implant disk-shaped electrodes inside the skull, underneath the tough, leatherlike membrane known as the dura mater. Alternatively, they may insert needlelike electrodes into the brain's gray matter to probe its function. Once the surgeons have identified the focal point of the seizure and removed the electrodes, they extract this tissue in a follow-up operation, and the patient usually becomes seizure-free.

A different use for iES is chronic electrical stimulation, in which the electrodes are left permanently in place. Gentle pulses of current sent through the electrodes can control the tremors and rigidity of Parkinson's disease (a technique called deep-brain stimulation) or reduce the incidence and severity of seizures. Pilot clinical experiments have evaluated the use of such implanted electrodes as a visual prosthetic device to enable people with vision impairments to navigate and as a therapy for obsessive-compulsive disorder and depression.

Patients reported electrode-evoked experiences such as seeing distorted faces reminiscent of paintings by Salvador Dalí.

HOT OR NOT

IN JULY 2020 *Nature Human Behaviour* published an atlas highlighting locations across the cortex that, when aroused with electrodes, evoked conscious experiences, such as the storm and the disconnected body mentioned earlier. Led by Josef Parvizi, a professor of neurology at the Stanford University School of Medicine, the clinical team collected data from 67 people with epilepsy. The researchers recorded electrical activity from more than 1,500 sites in the cortex, primarily with subdural electrodes. They mapped the recordings from those sites to spots on a digital brain model so they could compare data from different brains (the pattern of ridges and valleys that give the organ the look of an oversized walnut differs from person to person). The team looked for "responsive" electrodes that triggered some visual or tactile sensation, muscle twitching or disrupted speech. If the patient did not feel anything when stimulated, that electrode was marked as nonresponsive.

Patients reported a range of electrode-evoked subjective experiences: briefly flashing points akin to stars of light; distorted faces like those in the paintings of Salvador Dalí; bodily feelings such as tingling, tickling, burning, pulsing and so-called out-of-body experiences; fear, unease, sexual arousal, merriment; the desire to move a limb; the will to persevere in the face of some great but unrecognized challenge. Mere tickling of neural tissue with a bit of electric current was enough to evoke these feelings. During sham stimulation (no current applied), patients did not feel anything.

Although iES is safe and effective, it is also crude. The low-impedance electrodes are six to 10 square millimeters in area and deliver up to 10 milliamperes of electric current between adjacent electrodes—enough to modulate the excitability of a million or more nerve cells. Still, effects induced by iES can be quite localized. Responsiveness can change from all to none within millimeters or across a sulcus (a groove on the cortical surface).

The Parvizi team found that electrodes in the dedicated sensory and motor areas were far more likely to be responsive than those in areas of the cortex that process higher cognitive functions. Half to two thirds of electrodes above visual and tactile (somatosensory) cortex areas triggered some conscious perception; in regions of the lateral and anteromedial prefrontal cortex, which are involved with higher thought processes, at most one in five electrodes did so. Put differently, electrodes in the back of the cortex—in areas responsible for sensory experiences—were more likely to be active than those toward the front, which consists of regions of the cortex important for cognitive activity such as thinking, planning, moral reasoning, decision-making and intelligence.

Despite their importance for thinking, these regions have little to do with consciousness. Indeed, for the past century neurosurgeons have observed that so long as the eloquent cortex is spared, massive regions of the prefrontal cortex can be ablated without causing obvious deficits in the daily stream of consciousness of these patients. These regions of noneloquent cortex can modulate consciousness, but they are, by and large, not where conscious experience appears to originate. That privilege belongs to more posterior regions—the parietal, temporal and occipital lobes. Why the physical substrate of our mental experiences should be in the back rather than in the front of the brain remains a mystery.

TO SEE OR NOT TO SEE

APPLYING iES to the visual cortex triggers optical phenomena known as phosphenes, brief flashes that resemble lightning striking a darkened plain. This observation is the source of a long-standing dream of a prosthetic device able to restore some vision to people who are blind. Millions worldwide live with deficits in both eyes from retinitis pigmentosa, age-related macular degeneration, glaucoma, infection, cancer or trauma.

Doctors, scientists and engineers started pursuing visual prosthetics in the 1960s but have been able to harness the appropriate technology to help blind people only fairly recently. One prominent example is a device known as Orion, developed by Second Sight Medical Products in Los Angeles. A tiny camera, mounted on glasses, converts images into pulses and transmits them wirelessly to fire 60 electrodes sitting on the visual cortex. The people who have had this experimental device implanted into their brain perceive clouds of dots that allow them to navigate. "It's still a blast every time I turn it on," one study participant reports. "After seeing nothing to all of a sudden seeing little flickers of

light move around and figuring out that they mean something. It's just amazing to have some form of functional vision again." Orion significantly improves the quality of life for people who previously lived in complete darkness. It enables them to safely cross the street or locate a doorway. But it does not allow them to regain the ability to recognize figures, shapes or letters.

A team at the University of California, Los Angeles, and the Baylor College of Medicine led by neurosurgeon Daniel Yoshor did accomplish this feat, as described in the journal *Cell*. They stimulated nearby locations in the visual cortex to trigger phosphenes that appear close together, demonstrating that the external visual environment is mapped in a regular fashion onto the surface of the visual cortex. This observation has led to the erroneous belief that individual phosphenes are like pixels on a computer display—that is, if you were to simultaneously stimulate a series of points on the cortical surface in the shape of a cross, the subject should see points forming a cross. This does not happen, however.

Stimulating more than one location yields unpredictable results. In one participant, simultaneous stimulation of five electrodes, each one associated with one discrete phosphene, triggered the illumination of two large phosphenes that did not coalesce into a letter or any other coherent form. If the researcher staggered activation of the electrodes in time, however, the subject could identify shapes. The staggering reflected the delay required to trace the shape of a letter, as if the researcher were outlining a letter into the hand of the subject or onto a piece of paper. In this more dynamic manner, the subject with the implant whose vision was blocked could identify a stimulus by tracing out a Z, N, V and W, rapidly distinguishing upward from downward motion or discriminating sequences of letters.

Seeing the shape of a single letter is not quite the same as seeing a glorious sunset over Homer's wine-dark sea, but it represents progress. Why staggering stimulation in time improves perception is not clear and reveals our ignorance concerning functioning cortical circuits.

WHAT LIES AHEAD

TECHNOLOGICAL PROGRESS in so-called brain-machine interfaces is proceeding at a rapid pace. Elon Musk's company Neuralink released in April 2021 an impressive video showcasing a monkey playing the computer game Pong without any controller. This was achieved with two small chips implanted into the left and right motor cortices of the animal. Each chip has 1,024 hairlike electrodes that record the chattering of individual neurons. Collectively they convey the monkey's intention to quickly move the paddle up or down the screen to return the ball to the opposite side. Everything was done wirelessly; no electronics or dangling wires were protruding from the monkey's head. Many assume that surgeons will soon routinely replace or bypass faulty biological components—defective eyes or ears, failing memories—with superior electronic substitutes. Such optimism neglects the fact that all of this requires trepanation of the skull. In general, turning scientific insights into actual therapeutics is done in decades rather than in years. I am pretty confident that such enhancements will not occur in my lifetime (I'm now 66).

The "easiest" hurdles to overcome on the way to such a utopian (or perhaps dystopian) future are technological ones—reliably, quickly, and delicately reading and writing the brain electric. Neuralink's device represents the best of currently available technology and will certainly improve in future iterations. But we still have a long way to go before we can identify which of the 50,000 or more neurons in any quinoa-size bit of brain matter are involved in any given perception or action. Only when that happens will it be possible to limit electrical stimulation to just those neurons and not the output cables of nearby cells. That Parvizi and his colleagues failed to elicit conscious perceptions in more than half of all stimulated sites shows we lack tools capable of reliably eliciting *any* sensations through electrical stimulation, let alone being able to evoke any highly specific one.

Even more challenging are surgical and regulatory hurdles that demand that prosthetic devices can be routinely and safely implanted by drilling through the hard skull into the gray matter underneath while minimizing the risk of infections, bleeding and seizures. Furthermore, the electronics has to function for years inside warm, wet and salty biological tissue—hardly an optimal operating regime. You don't want your prosthetic device to corrode or freeze up in the equivalent of the blue screen of death. For this reason, neural implants will remain a matter of last resort for those with severe sensory or motor impairments. As neuroprosthetic devices move through clinical trials, they will help people with visual impairments see and those with physical disabilities perform actions such as steering a wheelchair with their thoughts, like the mind-Pong-playing monkey. For everyone else, the benefits of highly invasive brain surgery are unlikely to outweigh the costs.

But the true Annapurna ahead involves understanding how three pounds of excitable brain matter is responsible for seeing, moving and suffering. Yes, the physical substrate of heaven and hell is rooted in bioelectric signals that obey natural laws. But that tells us precious little about how a trillion electrical signals spiking each second, streaming over networks of tens of billions of heterogeneous cells, constitute a sight, sound or emotion.

Intracranial brain stimulation highlights the daily miracle of the brain's water changing into the wine of consciousness. The question remains, though: What is it about the brain, the most complex piece of active matter in the known universe, that turns the activity of 86 billion neurons into the feeling of life itself? SA

Christof Koch is chief scientist of MindScope at the Allen Institute in Seattle and of the Tiny Blue Dot Foundation in Santa Monica, as well as author of *The Feeling of Life Itself—Why Consciousness Is Widespread but Can't be Computed.* He serves on *Scientific American*'s board of advisers.

Hidden Consciousness

Some patients who appear to be in a coma, and can't speak or move, may be aware of what is happening around them

By Jan Claassen and Brian L. Edlow • Photographs by Kholood Eid

A MEDICAL TEAM SURROUNDED MARIA MAZURKEVICH'S HOSPITAL bed, all eyes on her as she did … nothing. Mazurkevich was 30 years old and had been admitted to New York–Presbyterian Hospital at Columbia University on a blisteringly hot July day in New York City. A few days earlier, at home, she had suddenly fallen unconscious. She had suffered a ruptured blood vessel in her brain, and the bleeding area was putting tremendous pressure on critical brain regions. The nurses and physicians at the hospital's neurological intensive care unit were looking for any sign that Mazurkevich could hear them. She was on a mechanical ventilator to help her breathe, and her vital signs were stable. But she showed no signs of consciousness.

Mazurkevich's parents, also at her bedside, asked, "Can we talk to our daughter? Does she hear us?" She didn't appear to be aware of anything. One of us (Claassen) was on her medical team, and when he asked Mazurkevich to open her eyes, hold up two fingers or wiggle her toes, she remained motionless. Her eyes did not follow visual cues. Yet her loved ones still thought she was "in there."

She was. The team members gave her an EEG—placing sensors on her head to monitor her brain's electrical activity—while they asked her to "keep opening and closing your right hand." Then they asked her to "stop opening and closing your right hand." Even though her hands themselves didn't move, her brain's activity patterns differed between the two commands. These brain reactions clearly indicated that she was aware of the requests and that those requests were different. After about a week, her body began to follow her brain. Slowly, with minuscule responses, Mazurkevich started to wake up. Within a year she recovered fully without major limitations to her physical or cognitive abilities. She now works as a pharmacy technician.

Mazurkevich had "covert consciousness," a state in which the brain reacts to the outside world with some comprehension, although the body does not respond. As many as 15 to 20 percent of patients who appear to be in a coma or other unresponsive state show these inner signs of awareness when evaluated with advanced brain-imaging methods or sophisticated monitoring of electrical activity. Many of these techniques have only recently been refined. These methods are altering our understanding of coma and other disorders of consciousness. Moreover, people whose covert consciousness is detected early have a greater chance of a full conscious and functional recovery, indicated by our studies at Columbia. These discoveries, which would have startled most neurologists and neuroscientists a few decades ago, highlight the importance of recognizing this hidden conscious state and developing ways to communicate with people who are in it.

THE STANDARD DEFINITION OF A COMATOSE PATIENT is someone who is unconscious, is unable to be awakened, and has no signs of awareness or the ability to interact with the environment. Patients in a coma caused by severe brain injury may look indistinguishable from someone in a deep sleep, except that most comatose patients cannot breathe on their own and need support from a ventilator, with a tube inserted into their airway.

Some people think comas are easy to recover from or—conversely—a living death. Both are mistakes. Popular depictions in movies and elsewhere may be partly responsible for this. Uma Thurman as the Bride in *Kill Bill: Volume 1* awakens abruptly from a prolonged comatose state, appears well nourished despite not having any feeding tubes and regains full physical strength within hours. The reality is far more challenging, with frequent medical complications, physical deterioration and a long road of small steps forward with many steps backward. Patients who survive coma after severe brain injury typically require feeding tubes for nutrition, tracheostomies that allow them to breathe through

AFTER A BRAIN INJURY, Maria Mazurkevich seemed to be in a coma. But brain tests showed she was aware, and she recovered in a year.

a tube in the neck, and weeks to months of rehabilitation. Recovery is variable and unpredictable, even in those who, like Mazurkevich, ultimately return to independence. Overly pessimistic views of coma patients are also inaccurate because people may assume that all such patients are destined to die without emerging from their coma or to live with severe disability. Recovery of consciousness, communication and functional independence is quite possible in some patients, even after a prolonged time.

Views about coma and consciousness have changed in the medical profession over time. In the 1960s neurologists and neurosurgeons noted that some comatose patients opened their eyes but showed no interaction with the environment. Many of these people remained in this state until death, leading some clinicians to believe that consciousness, once lost in this way, was impossible to recover.

Immobile patients with covert consciousness can deliberately alter their brain patterns when told to move parts of their bodies.

Yet in the 1990s reports of patients in a "permanent" vegetative state who returned to consciousness began to surface in the medical literature. In a vegetative state, unlike coma, people's eyes may open and shut, but they still do not react in any deliberate manner. The reports of recovery from this condition pushed the fields of neurocritical care and rehabilitation medicine to develop more fine-tuned classifications such as the minimally conscious state. It is characterized by nonverbal responses, as when patients track objects with their eyes or intermittently follow commands. A patient's prognosis, physicians learned, was related to these states. For instance, someone who moved from a vegetative to a minimally conscious state had a greater chance of further recovery.

Detecting and predicting recovery of consciousness early on, in the intensive care unit, is often a matter of life or death. Families typically make decisions about continuing or stopping life-sustaining therapy within 10 to 14 days of the injury—the time when surgical procedures become necessary to support longer-term breathing and feeding. And a diagnosis of covert consciousness could affect clinical decisions about goals of care, pain management, bedside behavior of clinicians and family members, and management of depression and anxiety.

SO WHAT DOES COVERT CONSCIOUSNESS LOOK LIKE TO clinicians and to the patient's family? One can get some idea through the lens of locked-in syndrome, in which people may have normal or near-normal cognition but are unable to control most motor movements. Locked-in patients illustrate the limitations of judging awareness, thinking abilities, and emotions purely based on motor function. The term "locked in" was coined in 1966 by neurologists Fred Plum and Jerome Posner in their monograph *The Diagnosis of Stupor and Coma*. They refer to the description of M. Noirtier De Villefort as "a corpse with living eyes" in Alexandre Dumas's classic *The Count of Monte Cristo* (1844–1846). In clinical practice, locked-in patients do not move their extremities, but many can reliably move their eyes up and down in response to verbal commands. Some can blink or show other subtle facial movements.

The experience of living in a locked-in state was poignantly illustrated by Jean-Dominique Bauby, an editor at *Elle* magazine who, in 1995, suffered a stroke that blocked signals traveling from the motor cortex in his brain to his spinal cord and limbs. Without the ability to speak or move his extremities, he began to communicate with his speech therapist using eye movements and wrote a memoir, *The Diving Bell and the Butterfly* (1997). This book captured the fear, frustration and hope that individuals with locked-in syndrome may experience. Remarkably, some people in a locked-in state report a meaningful quality of life.

With covert consciousness, the lack of outward movement is complete, even more so than with locked-in patients. But this does not mean the absence of inner life. In 2006 neuroscientist Adrian M. Owen, now at Western University in Ontario, and his colleagues examined a young woman who had experienced a severe traumatic brain injury and was believed to be in a vegetative state. The health-care team assessed her with a type of imaging scan called functional MRI, which traces blood flow through the brain to reveal active areas. During this scan the clinicians asked her to imagine playing tennis and to imagine walking through the rooms of her house. To the surprise of Owen and his colleagues, the woman showed activation within her brain comparable to that seen in healthy volunteers. What's more, the brain-activation patterns for the tennis task were distinct from the patterns in the walking task, indicating that she could deliberately change her brain activity.

Covert consciousness was subsequently identified in patients around the world, with varying types of brain injuries. In 2017 it was detected in seemingly unaware patients who had just been admitted to the intensive care unit at Massachusetts General Hospital with severe brain injuries, indicating that the covert phenomenon can occur in people who were hurt very recently, not only after patients have been "out" for weeks. To diagnose the covert state, clinicians give patients different behavioral tasks, such as asking them to open and close their hands or imagine swimming, while recording their brain reactions with an EEG or functional MRI. These responses have been reproduced by multiple research groups worldwide despite differences in methodology. Patients with covert consciousness can deliberately alter their brain patterns when told to move parts of their bodies or to envision an activity. But outwardly, in terms of body movements, they show no signs of following any prompt.

This state of being in which cognitive function exceeds motor expression is still poorly understood, and both the EEG and functional MRI techniques have limitations. The methods may not detect intentional brain activity in some patients who later regain consciousness. Both techniques may also be confounded by sedative medications, which are required for safety or comfort in most critically ill patients. Furthermore, functional MRI requires a specialized imaging room, and moving unstable patients from the intensive care unit to the MRI scanner may put them at risk. Yet another problem is that the MRI provides only a snapshot of a patient's level of consciousness during a short period because it cannot easily be repeated. An EEG can be done frequently at the patient's bedside—capturing snapshots

at different times—but the method has its own shortcomings. Its readings can be altered by electrical noise created by other machines in intensive care rooms, which can cause the test to reflect artifacts instead of reality.

Both methods need improvements, but the evidence for their usefulness is strong enough for them to have been endorsed for the diagnosis of covert consciousness in clinical guidelines in the U.S. (2018) and Europe (2020). The early detection of covert consciousness, soon after a patient's injury, predicts behavioral recovery of consciousness, long-term functional recovery and the speed of that recovery, as shown by the research that our group published in 2019 (and confirmed more recently, in 2022). Building on the momentum of these studies, scientists came together in 2019 to launch the Curing Coma Campaign, an international collaboration led by the Neurocritical Care Society to direct medical resources and public attention to the condition, with the goal of developing new therapies that promote recovery of consciousness.

Neurologists are trying to devise a test that can identify which patients are likely to be in a state of covert consciousness and thus should undergo advanced EEG and functional MRI assessments. Laboratories around the world are working to develop such screening methods, but progress has been slow because the structural and functional mechanisms that underlie covert consciousness are uncertain, so clinicians do not know exactly what to look for. Recent studies suggest that brain injuries disconnecting the thalamus—a region that relays movement signals and sensory information between the body and brain—from the cerebral cortex, which is responsible for higher-level cognitive functioning, may be responsible for the condition. Yet it is likely that not a single type of lesion but rather various combinations of lesions in several locations could cause motor dysfunction while allowing covert consciousness. Further complicating clinical efforts to detect covert consciousness is that patients with severe brain injuries often have fluctuating levels of consciousness. Such swings mean that a single assessment could miss important signs; perhaps patients need to be tested multiple times.

BUILDING ON RECENT DISCOVERIES ABOUT THE PRESENCE of covert consciousness, investigators are trying to reconnect and communicate with these patients using brain-computer interfaces. These devices typically record the brain's electrical activity while the patient is asked to move the cursor of a mouse on a computer screen. The computer "learns" to identify the physiological signals that correlate with the patient's attempts to move the cursor, left, right, up or down. Once training is completed, those brain patterns allow the patient to take control over the cursor. Patients can use it to select letters and spell out words.

Brain-computer interfaces would be ideal to provide covertly conscious patients a channel for communication with the outer world. But tremendous challenges must be overcome, particularly for acutely brain-injured patients. The capacity for sustained attention in these patients may be compromised, and prolonged training is often not feasible. Moreover, the hectic, noisy intensive care environment is not ideal for these purposes. For example, even though Mazurkevich had covert consciousness that was associated with a very good recovery, she was unable to activate a brain-computer interface to communicate with the health-care team or her family.

Communication might be possible using functional MRI, too. A few years ago Martin Monti, a cognitive psychologist at the University of California, Los Angeles, used the method to investigate the presence of covert consciousness in a group of behaviorally unresponsive patients. He wanted to see if he could train them to reliably answer "yes" or "no" to questions by using different functional MRI activation patterns. This required enormous technological coordination as the imaging data needed to be analyzed in real time. As Owen did in 2006, Monti asked patients to imagine playing tennis or imagine walking through their apartment. The difference was that he wasn't simply looking for brain activation; he wanted to see if they understood questions well enough to answer them. He told them to think about tennis if the answer to a given question was "yes" and think about walking through their home if the answer was "no." Monti identified one patient in the group who reliably communicated with him using this strategy, creating one pattern of brain activity for yes answers and another pattern for no answers. Although there are questions about whether this approach can be scaled up for wider use, his study suggested that communication with patients in a state of covert consciousness is possible.

To further improve communication, reliable tools to identify patients with covert consciousness need to be at the bedside. A number of groups are investigating advanced EEG technology because it can more easily be integrated into the clinical routine of an intensive care unit. And with brain-computer interfaces, the accuracy of the algorithm that decodes the patient's attempts to control the computer might be enhanced by using additional biological signals, such as heart rate, along with brain activity.

Beyond the urgent matter of caring for critically ill patients, diagnosis and exploration of covert consciousness have the potential to teach us about the human mind. In covert consciousness, the very foundation of our experience as humans, our consciousness, is dissociated from our behavior. What is the inner mental life of the covertly conscious patient? Detecting covert consciousness fundamentally affects our conceptualization of an individual's personhood and autonomy. Brain-computer interfaces have not yet allowed in-depth conversations, and to date patients with covert consciousness who recovered the ability to communicate and were interviewed later did not remember the experience of being covertly conscious. Mazurkevich, for instance, does not recall any aspect of her time in the intensive care unit when she appeared to be comatose. So the experience is still largely a mystery.

There is no mystery, however, about the ethical imperative that physicians now have to search for consciousness in patients who appear unresponsive, using all available technologies and resources. Increasing access to these technologies and resources is a fundamental goal, and challenge, for the medical community, spearheaded by the Curing Coma Campaign. With those tools, we can look forward to a future in which all covertly conscious people are given a way to speak for themselves. SA

Jan Claassen is an associate professor of neurology at Columbia University and a critical care neurologist at New York–Presbyterian Hospital at Columbia University, where he is chief of critical care and hospitalist neurology and head of the disorders of consciousness laboratory.

Brian L. Edlow is an associate professor of neurology at Harvard Medical School and a critical care neurologist at Massachusetts General Hospital, where he directs the Laboratory for Neuroimaging of Coma and Consciousness and is associate director of the Center for Neurotechnology and Neurorecovery.

Creating Our Sense of Self

One brain region helps people maintain a consistent identity

By Robert Martone

WE ARE ALL TIME TRAVELERS. EVERY DAY we experience new things as we travel forward through time. As we do, the countless connections among the nerve cells in our brain are recalibrated to accommodate these experiences. It's as if we reassemble ourselves daily, maintaining a mental construct of ourselves in physical time, and the glue that holds together our core identity is memory.

Our travels are not limited to physical time. We also experience mental time travel. We visit the past through our memories and then journey into the future by imagining what tomorrow or next year might bring. When we do so, we think of ourselves as we are now, remember who we once were and envision how we will be.

A study published in 2021 in the journal *Social Cognitive and Affective Neuroscience* (*SCAN*) explores how one particular brain region helps to knit together memories of the present and future self. When people sustain an injury to this area, it leads to an impaired sense of identity. The region—called the ventral medial prefrontal cortex (vmPFC)—may produce a fundamental model of oneself and place it in mental time. When the region does so, this study suggests, it may be the source of our sense of self.

Psychologists have long noticed that a person's mind handles information about oneself differently from other details. Memories that reference the self are easier to recall than other forms of memory. They benefit from what researchers have called a self-reference effect, in which information related to oneself is privileged and more salient in our thoughts. Self-related memories are distinct from both episodic memory, the category of recollections that pertains to specific events and experiences, and semantic memory, which connects to more general knowledge, such as the color of grass and the characteristics of the seasons.

Self-reference effects, then, are a way to investigate how our sense of self emerges from the workings of the brain—something that multiple research groups have studied intensely. For example, previous research employed functional magnetic resonance imaging (fMRI), a method that uses blood flow and oxygen consumption in specific brain areas as a measure of neural activity, to identify regions that were activated by self-reference. These studies identified the medial prefrontal cortex (mPFC) as a brain region related to self-thought.

This area, the mPFC, can be further divided into upper and lower regions (called dorsal and ventral, respectively), and it turns out that each one makes different contributions to self-related thought. The dorsal section plays a role in distinguishing self from other and appears to be task-related, whereas the ventral section, the vmPFC, contributes more to emotional processing.

In the *SCAN* study, the researchers used the self-reference effect to assess memories of present and future selves among people who had brain lesions to the vmPFC. The scientists worked with seven people who had lesions to this area and then compared them with a control group made up of eight people with injuries to other parts of the brain, as well as 23 healthy individuals without brain injuries. By comparing these groups, the scientists could investigate whether brain lesions in general or those to the vmPFC specifically might affect self-reference. All people in the study underwent a thorough neuropsychological evaluation, which confirmed that they were within normal ranges for a variety of cognitive assessments, including measures of verbal fluency and spatial short-term memory. The researchers then asked the participants to list adjectives to describe themselves and a well-known celebrity, both in the present and 10 years in the future. Later, the participants had to recall these same traits.

The researchers discovered that people in their control group could recall more adjectives linked to themselves in the present and future than adjectives linked to the celebrity. In other words, scientists found that the self-reference effect extends to both the future and the present self. Although there was some variation in the group—people with brain injuries to areas other than the vmPFC were somewhat less able to recall details about their future self when compared with healthy participants—the self-reference effect still held true.

Results were distinctly different, however, for the participants with injuries to the vmPFC. People with lesions in this area had little or no ability to recall references to the self, regardless of the context of time. Their identification of adjectives for celebrities in the present or future was also significantly impaired when compared with the rest of the participants' responses. In addition, people with vmPFC lesions had less confidence about an individual's abil-

ity to possess traits than other people in the study. All of this evidence points to a central role for the vmPFC in the formation and maintenance of identity.

These findings are intriguing for several reasons. Brain lesions can help us understand the normal function of the region involved. Lesions of the vmPFC are associated with altered personality, blunted emotions, and a number of changes in emotional and executive function. Injury to this area is most often associated with confabulations: false memories that people recite to listeners with great confidence. Although it may be tempting for someone to view confabulations as deliberate or creative falsehoods, people who tell them actually are unaware that their stories are false. Instead it is possible their confusion could stem from misfunctioning memory retrieval and monitoring mechanisms.

More broadly, the study helps us understand how self-related memories—recollections key to maintaining our core sense of identity—depend on the function of the vmPFC. But what about our past selves? Curiously, in previous studies that asked people to consider their past selves, there was no more activation of the mPFC than when considering someone else. Our past selves seem foreign to us, as if they were individuals apart from us.

One idea that scientists have put forward to understand this distinction is that perhaps we are not very kind in our judgments of our past selves. Instead we may be rather critical and harshly judgmental of our previous behavior, emotions and personal traits. In these situations, we may use our past primarily to construct a more positive self-image in the present. Put another way, because we may recognize flaws in our past self's behavior, we tend to distance ourselves from the person we once were.

Bringing the present and future into the spotlight, then, is central to understanding the way our brain and thoughts shape our current identities. In many ways, it makes sense that the mPFC is key in this process of recalling present details and imagining future ones that build on our memories. The prefrontal cortex, including the mPFC and its subdivisions, forms a network in the brain that is involved in future planning.

That network also includes the hippocampus, a brain structure that is central to episodic memory formation and that can track moments as sequential events in time. In earlier work, researchers found that manipulating the activity of the hippocampus alters creative and future imaginings, which suggests an important role for brain structures supporting memory in imagining the future. In fact, although we often think of memory as the brain's accurate and dispassionate recording device, some scholars have characterized it as a form of imagination.

Future thought is a vital component of being human. Its importance in our culture is embodied in the mythological figure and pre-Olympian god Prometheus (whose name means "fore-thinker"), patron of the arts and sciences. According to Greek legend, he shaped humans out of clay and bestowed them with fire and the skills of craftsmanship. These are acts that illustrate the power of imagining a novel future. Although there is debate as to whether thinking about the future is an exclusively human feature—birds such as Western Scrub-Jays, for example, appear to anticipate and plan for future food needs—it is clear that future thought has played a significant role in human evolution. This ability may have contributed to the development of language, and it has a key part in human interactions, where the vmPFC is central to evaluating and taking advantage of social context.

Now, thanks to this research, we have a better idea than ever about the way a small region within our brains is able to build and hold this core ability to maintain our identity. SA

Robert Martone is a research scientist with expertise in neurodegeneration. He spends his free time kayaking and translating Renaissance Italian literature.

Illustration by Zoe Liu

Constructing the World from Inside Out

The brain probes your physical surroundings to select just the information needed to survive and flourish

By György Buzsáki

Illustration by Stefania Infante

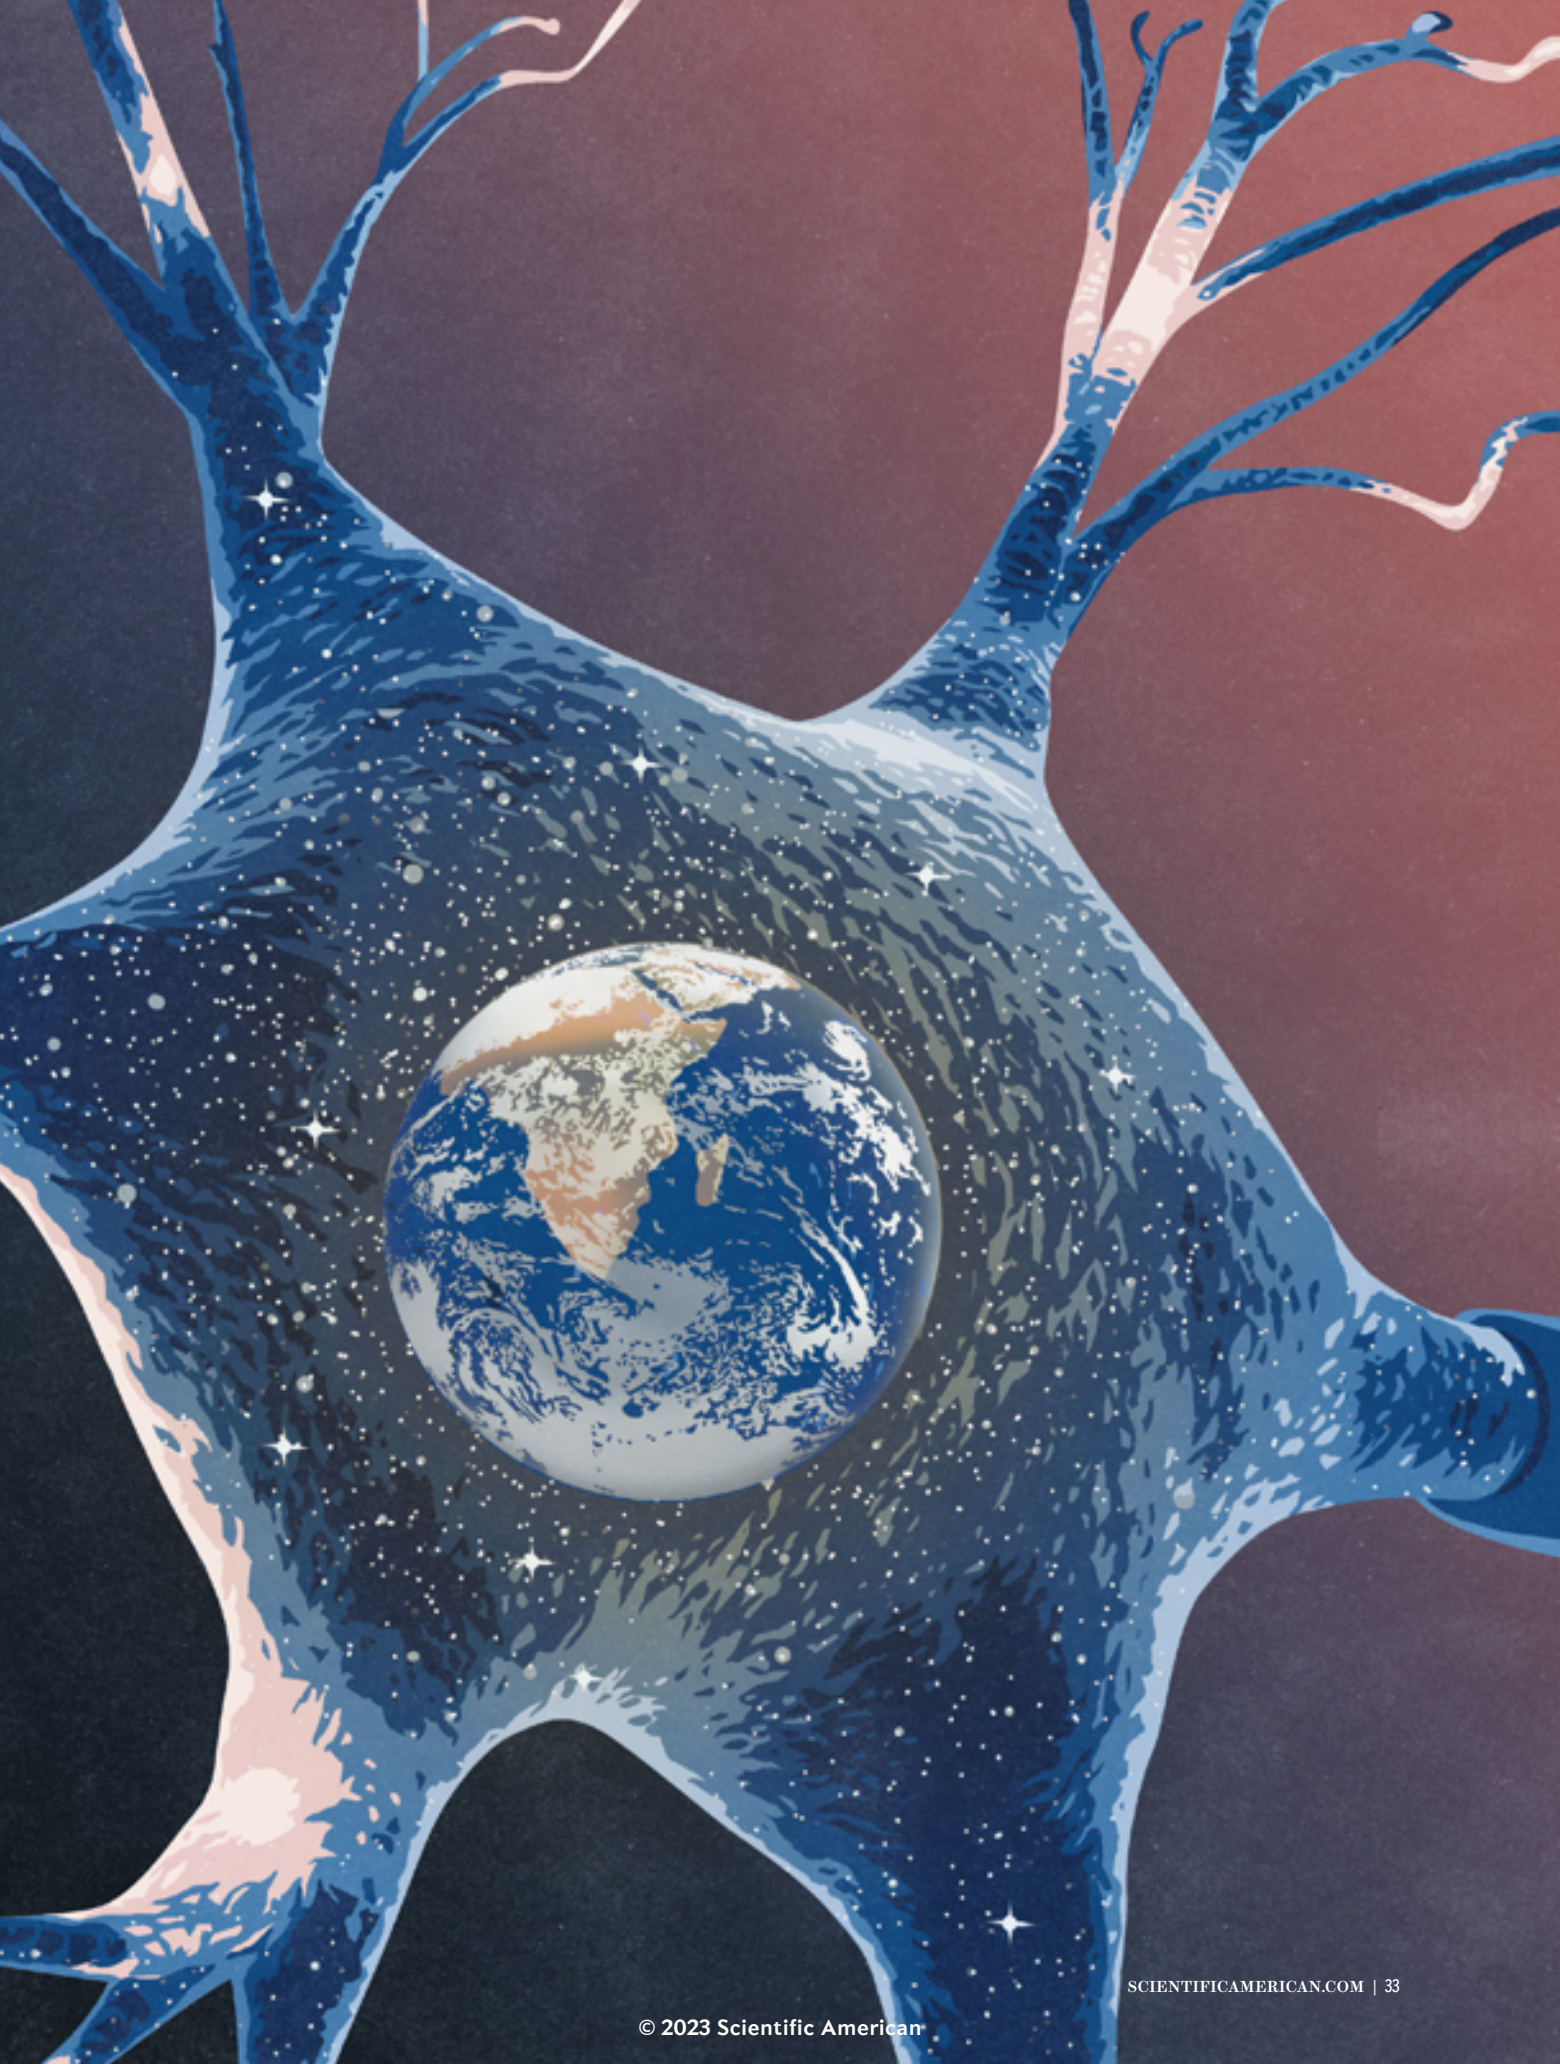

AS A YOUNG COURSE INSTRUCTOR IN SEMINARS FOR MEDICAL students, I faithfully taught neurophysiology by the book, enthusiastically explaining how the brain perceives the world and controls the body. Sensory stimuli from the eyes, ears, and such are converted to electrical signals and then transmitted to the relevant parts of the sensory cortex that process these inputs and induce perception. To initiate a movement, impulses from the motor cortex instruct the spinal cord neurons to produce muscular contraction.

Most students were happy with my textbook explanations of the brain's input-output mechanisms. Yet a minority—the clever ones—always asked a series of awkward questions. "Where in the brain does perception occur?" "What initiates a finger movement before cells in the motor cortex fire?" I would always dispatch their queries with a simple answer: "That all happens in the neocortex." Then I would skillfully change the subject or use a few obscure Latin terms that my students did not really understand but that seemed scientific enough so that my authoritative-sounding accounts temporarily satisfied them.

Like other young researchers, I began my investigation of the brain without worrying much about whether this perception-action theoretical framework was right or wrong. I was happy for many years with my own progress and the spectacular discoveries that gradually evolved into what became known in the 1960s as the field of "neuroscience." Yet my inability to give satisfactory answers to the legitimate questions of my smartest students has haunted me ever since. I had to wrestle with the difficulty of trying to explain something that I didn't really understand.

Over the years I realized that this frustration was not uniquely my own. Many of my colleagues, whether they admitted it or not, felt the same way. There was a bright side, though, because these frustrations energized my career. They nudged me to develop a perspective that provides an alternative description of how the brain interacts with the outside world.

The challenge for me and other neuroscientists involves the weighty question of what, exactly, is the mind. Ever since the time of Aristotle, thinkers have assumed that the soul or the mind is initially a blank slate, a tabula rasa on which experiences are painted. This view has influenced thinking in Christian and Persian philosophies, British empiricism and Marxist doctrine. In the past century it has also permeated psychology and cognitive science. This "outside-in" view portrays the mind as a tool for learning about the true nature of the world. The alternative view—one that has defined my research—asserts that the primary preoccupation of brain networks is to maintain their own internal dynamics and perpetually generate myriad nonsensical patterns of neural activity. When a seemingly random action offers a benefit to an organism's survival, the neuronal pattern leading to that action gains meaning. When an infant utters "te-te," the parent happily offers the baby "teddy," so the sound "te-te" acquires the meaning of the teddy bear. Recent progress in neuroscience has lent support to this framework.

DOES THE BRAIN "REPRESENT" THE WORLD?

NEUROSCIENCE INHERITED the blank slate framework millennia after early thinkers gave names like tabula rasa to mental operations. Even today we still search for neural mechanisms that might relate to their dreamed-up ideas. The dominance of the outside-in framework is illustrated by the outstanding discoveries of the legendary scientific duo David Hubel and Torsten Wiesel, who introduced single-neuronal recordings to study the visual system and were awarded the Nobel Prize in Physiology or Medicine in 1981. In their signature experiments, they recorded neural activity in animals while showing them images of various shapes. Moving lines, edges, light or dark areas,

and other physical qualities elicited firing in different sets of neurons. The assumption was that neuronal computation starts with simple patterns that are synthesized into more complex ones. These features are then bound together somewhere in the brain to represent an object. No active participation is needed. The brain automatically performs this exercise.

The outside-in framework presumes that the brain's fundamental function is to perceive "signals" from the world and correctly interpret them. But if this assumption is true, an additional operation is needed to respond to these signals. Wedged between perceptual inputs and outputs resides a hypothetical central processor—which takes in sensory representations from the environment and makes decisions about what to do with them to perform the correct action.

So what exactly is the central processor in this outside-in paradigm? This poorly understood and speculative entity goes by various names—free will, homunculus, decision maker, executive function, intervening variables or simply just a "black box." It all depends on the experimenter's philosophical inclination and whether the mental operation in question is applied to the human brain, brains of other animals, or computer models. Yet all these concepts refer to the same thing.

An implicit practical implication of the outside-in framework is that the next frontier for progress in contemporary neuroscience should be to find where the putative central processor resides in the brain and systematically elaborate the neuronal mechanisms of decision-making. Indeed, the physiology of decision-making has become one of the most popular focuses in contemporary neuroscience. Higher-order brain regions, such as the prefrontal cortex, have been postulated as the place where "all things come together" and "all outputs are initiated." When we look more closely, however, the outside-in framework does not hold together.

This approach cannot explain how photons falling on the retina are transformed into a recollection of a summer outing. The outside-in framework requires the artificial insertion of a human experimenter who observes this event [*see box at right*]. The experimenter-in-the-middle is needed because even if neurons change their firing patterns when receptors on sensory organs are stimulated—by light or sound, for instance—these changes do not intrinsically "represent" anything that can be absorbed and integrated by the brain. The neurons in the visual cortex that respond to the image of, say, a rose have no clue. They do not "see" the appearance of a flower. They simply generate electrical oscillations in response to inputs from other parts of the brain, including those arriving along multiple complex pathways from the retina.

In other words, neurons in sensory cortical areas and even in the hypothetical central processor cannot "see" events that happen in the world. There is no interpreter in the brain to assign meaning to these changes in neuronal firing patterns. Without a magi-

Outside In vs. Inside Out

The idea of the brain as a blank slate onto which experience is written has existed since ancient times—and persists today in modified form. Some neuroscientists have begun to question this theory because it requires a hard-to-justify assumption about the way we perceive and process events from the outside world—in particular, the need to involve a hypothetical "interpreter" to explain what is happening.

Outside-In Framework

A stimulus—the image of a flower—reaches the eyes, and the brain responds by causing neurons to fire. This theory is plausible only with the involvement of an "experimenter" to observe and establish a relation between the flower and the neuronal responses it induces. Absent the experimenter, neurons in the sensory cortex do not "see" the flower.

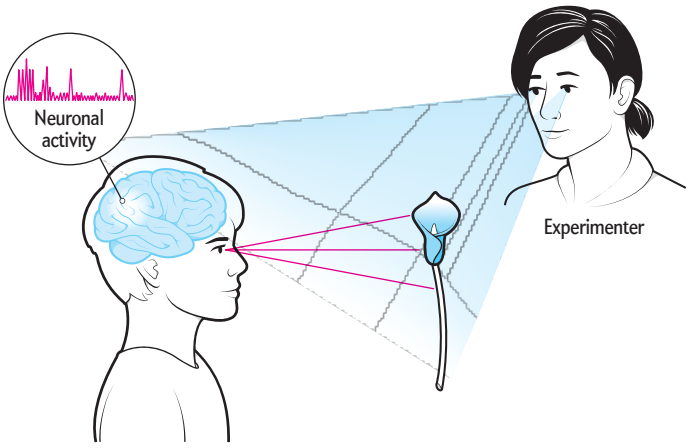

Inside-Out Framework

The alternative, inside-out theory does away with the experimenter. It presumes instead that we come to understand the external world by taking actions—moving a flower, for instance—to learn about an object. To accomplish this task, inputs from action-initiating neurons combine with sensory inputs to provide an understanding of the object's size, shape and other attributes. A meaningful picture arises, allowing the neurons to "see" the flower.

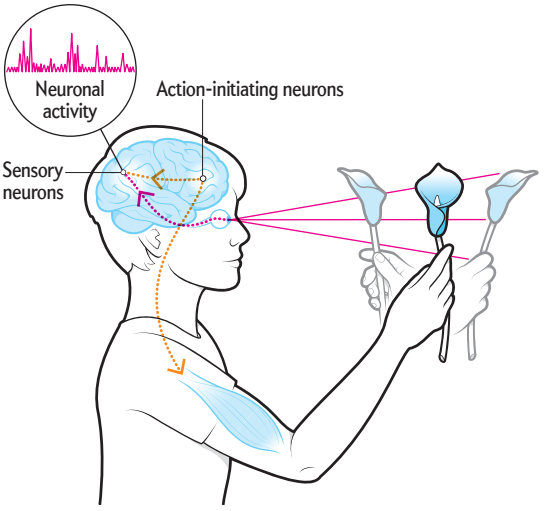

Graphic by Brown Bird Design

cal homunculus watching the activities of all the neurons in the brain with the omniscience of the experimenter, the neurons that take this all in are unaware of the events that caused these changes in their firing patterns. Fluctuations in neuronal activity are meaningful only for the scientist who is in the privileged position of observing both events in the brain and events in the outside world and then comparing the two perspectives.

PERCEPTION IS WHAT WE DO

BECAUSE NEURONS have no direct access to the outside world, they need a way to compare or "ground" their firing patterns to something else. The term "grounding" refers to the ability of the brain's circuits to assign meaning to changes in neuronal firing patterns that result from sensory inputs. They accomplish this task by relating this activity to something else. The "dah-dah-dit" Morse code pattern becomes meaningful only when it has previously been linked to the letter "G." In the brain, the only available source of a second opinion appears when we initiate some action.

We learn that sticks that look bent in water are not broken by moving them. Similarly, the distances between two trees and two mountain peaks may appear identical, but by moving around and shifting our perspective we learn the difference.

The outside-in framework follows a chain of events from perception to decision to action. In this model, neurons in dedicated sensory areas are "driven" by environmental signals and thus cannot relate their activity to something else. But the brain is not a serial processing unit; it does not proceed one by one through each of these steps. Instead any action a person takes involves the brain's motor areas informing the rest of the cerebral cortex about the action initiated—a message known as a corollary discharge.

Neuronal circuits that initiate an action dedicate themselves to two tasks. The first is to send a command to the muscles that control the eyes and other bodily sensors (the fingers and tongue, among others). These circuits orient bodily sensors in the optimal direction for in-depth investigation of the source of an input and enhance the brain's ability to identify the nature and location of initially ambiguous incoming signals from the senses.

The second task of these same action circuits involves sending notifications—the corollary discharges—to sensory and higher-order brain areas. Think of them as registered mail receipts. Neurons that initiate eye movement also notify visual sensory areas of the cortex about what is happening and disambiguate whether, say, a flower is moving in the wind or being handled by the person observing it.

This corollary message provides the second opinion sensory circuits need for grounding—a confirmation that "my own action is the agent of change." Similar corollary messages are sent to the rest of the brain when a person takes actions to investigate the flower and its relationship to that person and other objects. Without such exploration, stimuli from the flower alone—the photons arriving on the retina connected to an inexperienced brain—would never become signals that furnish a meaningful description of the flower's size and shape. Perception then can be defined as what we *do*—not what we passively take in through our senses.

You can demonstrate a simple version of the corollary discharge mechanism. Cover one of your eyes with one hand and press against the side of the other eye gently with the tip of your finger at about three times per second while reading this text. You will see immediately that the page is moving back and forth. By comparison, when you move your eyes naturally to read or look around the room, nothing else seems to move. This constancy occurs because neurons that initiate eye movements to scan sentences also send a corollary signal to the visual system to indicate whether the world or the eyeball is moving, thus stabilizing the perception of your surroundings.

LEARNING BY MATCHING

THE CONTRAST between outside-in and inside-out approaches becomes most striking when used to explain the mechanisms of learning. A tacit assumption of the blank slate model is that the complexity of the brain grows with the amount of experience. As we learn, the interactions of brain circuits should become increasingly more elaborate. In the inside-out framework, however, experience is not the main source of the brain's complexity.

Instead the brain organizes itself into a vast repertoire of preformed patterns of firing known as neuronal trajectories. This self-organized brain model can be likened to a dictionary filled initially with nonsensical words. New experience does not change the way these networks function—their overall activity level, for instance. Learning takes place, rather, through a process of matching the preexisting neuronal trajectories to events in the world.

To understand the matching process, we need to examine the advantages and constraints brain dynamics impose on experience. In the basic version, models of blank slate neuronal networks assume a collection of largely similar, randomly connected neurons. The presumption is that brain circuits are highly plastic and that any arbitrary input can alter the activity of neuronal circuits.

We can see the fallacy of this approach by considering an example from the field of artificial intelligence. Classical AI research—particularly the branch known as connectionism, the basis for artificial neural networks—adheres to the outside-in, tabula rasa model. This prevailing view was perhaps most explicitly promoted in the 20th century by Alan Turing, the great pioneer of mind modeling: "Presumably the child brain is something like a notebook as one buys it from the stationer's," he wrote.

Artificial neural networks built to "write" inputs

Imagining the Road Ahead

An experiment demonstrates that distinct sets of neurons fire—each set in a different order—when a rat is planning whether to take the left or right route to receive a reward.

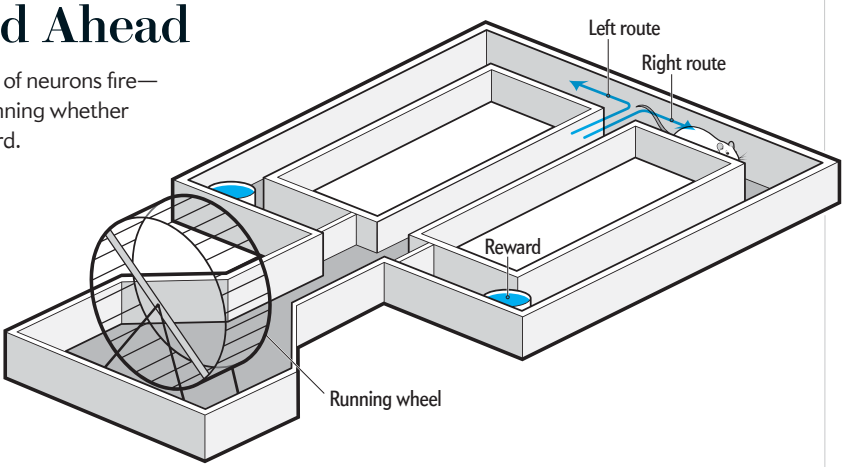

Experimental Setup

A running wheel is located at the entrance to a maze with two route options, both of which lead to a reward. The rat is free to choose a path through the maze after a 15-second run on the wheel. Neuronal firing patterns are recorded during both maze and wheel activity.

Results

Neuronal activity while the rat was running in the wheel predicted the direction it would take in the maze many seconds later, as if the animal were imagining the path to come. The "left trials" panel represents a sequence of neuronal firing that differs from the one for the right trials. When the "left" pattern occurred while the rat was in the wheel, it took the left route in the maze moments later.

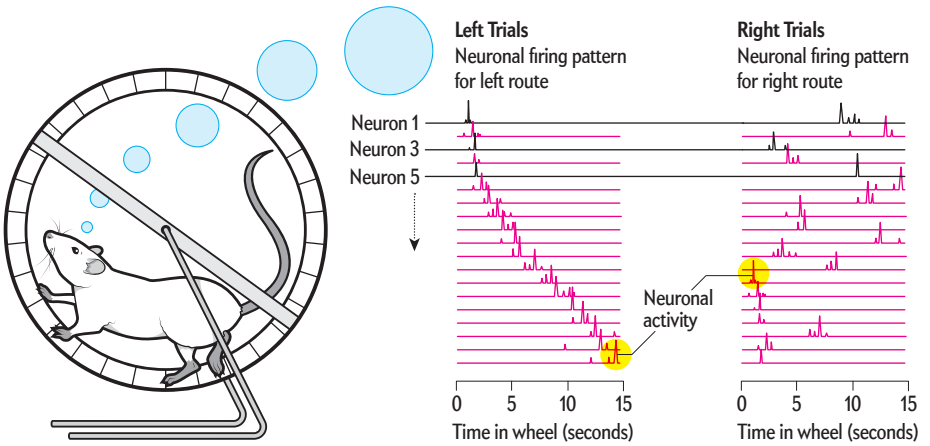

onto a neural circuit often fail because each new input inevitably modifies the circuit's connections and dynamics. The circuit is said to exhibit plasticity. But there is a pitfall. While constantly adjusting the connections in its networks when learning, the AI system, at an unpredictable point, can erase all stored memories—a bug known as catastrophic interference, an event a real brain never experiences.

The inside-out model, in contrast, suggests that self-organized brain networks should resist such perturbations. Yet they should also exhibit plasticity selectively when needed. The way the brain strikes this balance relates to vast differences in the connection strength of different groups of neurons. Connections among neurons exist on a continuum. Most neurons are only weakly connected to others, but a smaller subset retains robust links. The strongly connected minority is always on the alert. It fires rapidly, shares information readily within its own group, and stubbornly resists any modifications to the neurons' circuitry. Because of the multitude of connections and their high communication speeds, these elite subnetworks, sometimes described as a "rich club," remain well informed about neuronal events throughout the brain.

The hard-working rich club makes up roughly 20 percent of the overall population of neurons, but it is in charge of nearly half of the brain's activity. In contrast to the rich club, most of the brain's neurons—the neural "poor club"—tend to fire slowly and are weakly connected to other neurons. But they are also highly plastic and able to physically alter the connection points between neurons, known as synapses.

Both rich and poor clubs are key for maintaining brain dynamics. Members of the ever ready rich club fire similarly in response to diverse experiences. They offer quick, good-enough solutions under most conditions. We make good guesses about the unknown not because we remember it but because our brains always make a surmise about a new, unfamiliar event. Nothing is completely novel to the brain because it always relates the new to the old. It generalizes. Even an inexperienced brain has a vast reservoir of neuronal trajectories at the ready, offering opportunities to match events in the world to preexisting brain patterns without requiring substantial reconfiguring of connections. A brain that remakes itself constantly would be unable to adapt

Graphic by Brown Bird Design

Rehearsal and Playback

A group of neurons fire before, during and after a rat runs a lap on an elevated track. Neurons firing rapidly at the beginning and end of the run (*insets*) are the same ones active during the run, constituting either a rehearsal or a playback (the prerun sequence in reverse) of the rat's trajectory. These early and late events are known as sharp wave ripples and enable a mental process that selects and remembers an optimal path.

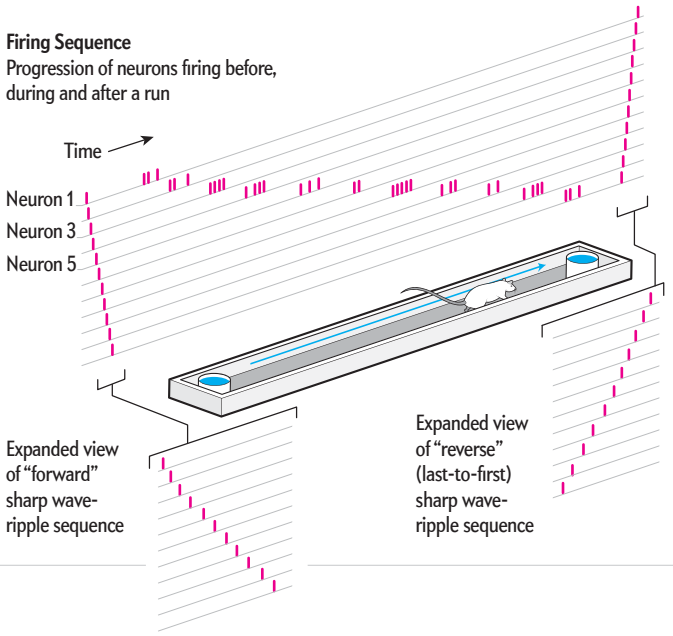

quickly to fast-changing events in the outside world.

But there also is a critical role for the plastic, slow-firing-rate neurons. These neurons come into play when something of importance to the organism is detected and needs to be recorded for future reference. They then go on to mobilize their vast reserve to capture subtle differences between one thing and another by changing the strength of some connections to other neurons. Children learn the meaning of the word "dog" after seeing various kinds of canines. When a youngster sees a sheep for the first time, they may say "dog." Only when the distinction matters—understanding the difference between a pet and livestock—will they learn to differentiate.

COGNITION AS INTERNALIZED ACTION

AS AN EXPERIMENTER, I did not set out to build a theory in opposition to the outside-in framework. Only decades after I started my work studying the self-organization of brain circuits and the rhythmic firing of neuronal populations in the hippocampus did I realize that the brain is more occupied with itself than with what is happening around it. This realization led to a whole new research agenda for my lab. Our experiments, along with findings from other groups, revealed that neurons devote most of their activity to sustaining the brain's perpetually varying internal states rather than being controlled by stimuli impinging on our senses.

During the course of natural selection, organisms adapt to the ecological niches in which they live and learn to predict the likely outcomes of their actions in those niches. As brain complexity increases, more intricate connections and neuronal computations insert themselves between motor outputs and sensory inputs. This investment enables the prediction of planned actions in more complex and changing environments and at lengthy timescales far in the future. More sophisticated brains also organize themselves to allow computations to continue when sensory inputs vanish temporarily and an animal's actions come to a halt. When you close your eyes, you still know where you are because a great deal of what defines "seeing" is rooted in brain activity. This disengaged mode of neuronal activity provides access to an internalized virtual world of vicarious or imagined experience and serves as a gateway to a variety of cognitive processes.

Let me offer an example of such a disengaged mode of brain operation from our work on the brain's temporal lobe, an area that includes the hippocampus, the nearby entorhinal cortex, and related structures involved with multiple aspects of navigation (the tracking of direction, speed, distance traveled, environmental boundaries, and so on).

Our research builds on leading theories of the functions of the hippocampal system, such as the spectacular Nobel-winning discovery by John O'Keefe of University College London. O'Keefe found that firing of an animal's hippocampal neurons during navigation correlates with that animal's spatial location. For that reason, these neurons are known as place cells.

When a rat walks through a maze, distinct assemblies of place cells become active in a sequential chain corresponding to where it is on its journey. From that observation, one can tentatively conclude that continually changing sensory inputs from the environment exercise control over the firing of neurons, in line with the outside-in model.

Yet other experiments, including in humans, show that these same networks are used for our internal worlds that keep track of personal memories, engage in planning and imagine future actions. If cognition is approached from an inside-out perspective, it becomes clear that navigation through either a physical space or a landscape that exists only in the imagination is processed by identical neural mechanisms.

Fifteen years ago my lab set out to explore the mechanisms of spatial navigation and memory in the hippocampus to contrast the outside-in and inside-out frameworks. In 2008 Eva Pastalkova, then a postdoctoral fellow, and I trained rats to alternate between the left and right arms of a maze to find water. At the beginning of each traversal of the maze, the rat was required to run in a wheel for 15 seconds, which helped to ensure that memory alone of the maze routes, and not environmental and body-derived cues, allowed it to choose

Graphic by Brown Bird Design

a particular arm of the maze. We reasoned that if hippocampal neurons "represent" places in the maze corridors and the wheel, as predicted by O'Keefe's spatial navigation theory, a few neurons should fire continuously at each spot whether the rat is in the corridors or the wheel. In contrast, if the neurons' firing is generated by internal brain mechanisms that can support both navigation and memory, the duration of neuronal firing should be similar at all locations, including inside the wheel.

The findings of these experiments defied outside-in explanations. Not a single neuron among the hundreds recorded fired continuously throughout the wheel running. Instead many neurons fired transiently one after the other in a continuous sequence.

Obviously these neurons could not be called place cells, because the animal's body was not displaced while at the single location of the running wheel. Moreover, the firing patterns of individual neurons in this neuronal trajectory could not be distinguished from neurons active when the rat was traversing the arms of the maze.

When we sorted individual trials according to the rat's future choice of left or right arms, the neuronal trajectories were uniquely different. The distinct trajectories eliminated the possibility that these neuronal sequences arose from counting steps, estimating muscular effort or some other undetected feedback stimuli from the body. Also, the unique neuronal trajectories allowed us to predict the animal's maze arm choice from the moment it entered the wheel and throughout wheel running, a period in which the rat had to keep in mind the previously visited arm. The animals needed to correctly choose the alternate maze arm each time to get their rewards [*see box on page 37*].

These experiments lead us to the idea that the neuronal algorithms that we can use to walk to the supermarket govern internalized mental travel. Disengaged navigation takes us through the series of events that make up personal recollections, known as episodic memories.

In truth, episodic memories are more than recollections of past events. They also let us look ahead to plan for the future. They function as a kind of "search engine" that allows us to probe both past and future. This realization also presages a broadening in nomenclature. These experiments show that progressions of place cell activity are internally generated as preconfigured sequences selected for each maze corridor. Same mechanism, multiple designations—so they can be termed place cells, memory cells or planning cells, depending on the circumstance.

Further support for the importance of disengaged circuit operations comes from "offline" brain activity when an animal is milling around doing nothing, consuming a reward or just sleeping. As a rat rests in the home cage after a maze exploration, its hippocampus generates brief, self-organized neuronal trajectories. These sharp wave ripples, as they are known, occur in 100-millisecond time windows and reactivate the same neurons that were firing during several seconds of maze running, recapitulating the neuronal sequences that occurred during maze traversals. Sharp wave-ripple sequences help to form our long-term memories and are essential to normal brain functioning. In fact, alteration of sharp wave-ripple events by experimental manipulations or disease results in serious memory impairment [*see box on opposite page*].

Clever experiments performed in human subjects and in animals over the past decade show that the time-compressed ripple events constitute an internalized trial-and-error process that subconsciously creates real or fictive alternatives for making decisions about an optimal strategy, constructing novel inferences and planning ahead for future actions without having to immediately test them by undertaking a real exploit. In this sense, our thoughts and plans are deferred actions, and disengaged brain activity is an active, essential brain operation. In contrast, the outside-in theory does not make any attempt to assign a role to the disengaged brain when it is at rest or even in the midst of sleep.

THE MEANING OF INSIDE OUT

IN ADDITION to its theoretical implications, the inside-out approach has a number of practical applications. It may help in the search to find better diagnostic tools for brain disease. Current terminology often fails to describe accurately the underlying biological mechanisms of mental and neurological illnesses. Psychiatrists are aware of the problem but have been hindered by limited understanding of pathological mechanisms and their relation to symptoms and drug responses.

The inside-out theory should also be considered as an alternative to some of the most prevalent connectionist models for conducting AI research. A substitute for them might build models that maintain their own self-organized activity and that learn by "matching" rather than by continual adjustments to their circuitry. Machines constructed this way could disengage their operations from the inputs of electronic sensors and create novel forms of computation that resemble internal cognitive processes.

In real brains, neural processes that operate through disengagement from the senses go hand in hand with mechanisms that promote interactions with the surrounding world. All brains, simple or complex, use the same basic principles. Disengaged neural activity, calibrated simultaneously by outside experience, is the essence of cognition. I wish I had had this knowledge when my smart medical students asked their legitimate questions that I brushed off too quickly. SA

György Buzsáki is Biggs Professor of Neuroscience at New York University Grossman School of Medicine. He studies memory formation and how brain rhythms affect cognition. He was a co-recipient of the 2011 Brain Prize from the Lundbeck Foundation. Buzsáki is author most recently of *The Brain from Inside Out* (Oxford University Press, 2019).

BORN TO COUNT

As Plato anticipated, babies do math

By Jacob Beck and Sam Clarke

Photograph by Jamie Chung

HOW MANY GUMBALLS are in the jar? This popular guessing game involves our innate number sense. (Find the answer in the fold on page 47.)

IMAGINE HOSTING A PARTY. YOU ARRANGE SNACKS, CURATE A PLAYLIST AND PLACE a variety of beers in the refrigerator. Your first guest shows up, adding a six-pack before taking one bottle for himself. You watch your next guest arrive and contribute a few more beers, minus one for herself. Ready for a drink, you open the fridge and are surprised to find only eight beers remaining. You haven't been consciously counting the beers, but you know there should be more, so you start poking around. Sure enough, in the crisper drawer, behind a rotting head of romaine, are several bottles.

How did you know to look for the missing beer? It's not like you were standing guard at the refrigerator, tallying how many bottles went in and out. Rather you were using what cognitive scientists call your number sense, a part of the mind that unconsciously solves simple math problems. While you were immersed in conversation with guests, your number sense was keeping tabs on how many beers were in the fridge.

For a long time scientists, mathematicians and philosophers have debated whether this number sense comes preinstalled or is learned over time. Plato was among the first in the Western tradition to propose that humans have innate mathematical abilities. In Plato's dialogue *Meno,* Socrates coaxes the Pythagorean theorem out of an uneducated boy by asking him a series of simple questions. Socrates's takeaway is that the boy had innate knowledge of the Pythagorean theorem all along; the questioning just helped him express it.

In the 17th century John Locke rejected this idea, insisting that the human mind begins as a tabula rasa, or blank slate, with almost all knowledge acquired through experience. This view, known as empiricism, in contrast to Plato's nativism, was later further developed by John Stuart Mill, who argued that we learn two plus three is five by seeing many examples where it holds true: two apples and three apples make five apples, two beers and three beers make five beers, and so on.

In short, empiricism dominated philosophy and psychology until the second half of the 20th century, when nativist-friendly thinkers such as Noam Chomsky swung the pendulum back toward Plato. Chomsky focused on language, proposing that children are born with an innate language instinct that enables them to quickly acquire their first language with little in the way of explicit instruction.

Others then extended Chomsky's hypothesis to mathematics. In the late 1970s cognitive scientists C. R. Gallistel and Rochel Gelman argued that children learn to count by mapping the number words in their language onto an innate system of preverbal counting that humans share with many other animals. In his landmark book *The Number Sense,* first published in 1997, French neuroscientist Stanislas Dehaene drew attention to the converging evidence for this preverbal system, helping researchers from diverse disciplines—animal cognition, developmental psychology, cognitive psychology, neuroscience, education—realize they were all studying the same thing.

In our 2021 paper in the journal *Behavioral and Brain Sciences,* we argued that there is no longer a serious alternative to the view that humans and many nonhuman animals have evolved a capacity to process numbers. Whereas Plato proposed that we have innate mathematical knowledge, or a capacity to think about numbers, we argue that we have innate mathematical perception—an ability to see or sense numbers. When you opened the fridge, it's not that you saw the beer bottles and made an inference about their number in the way that you saw Heineken labels

and inferred that someone brought a pale lager from the Netherlands. Rather you saw their number much the way you perceived their shape and color.

But not everyone agrees with this idea, and a new wave of empiricism has emerged over the past decade. Critics who reject the existence of an ability to innately sense numbers highlight a broader and important scientific challenge: How could we ever know the contents of an infant's or a nonhuman animal's mind? As philosophers of cognitive science, we supplement thousands of years of philosophical thinking about this issue by drawing on a mountain of experimental evidence that simply was not available to past thinkers.

EMERGING EVIDENCE

IMAGINE YOU SEE TWO collections of dots flash on a computer screen in quick succession. There's no time to count them, but if you're like the thousands of people who have done this exercise in studies, you'll be able to tell which grouping had more dots if the two are sufficiently different. Although you might struggle to distinguish 50 from 51 dots, you'll be pretty good at recognizing that 40 dots are less numerous than 50. Is this ability innate or the product of years of mathematical education?

In 2004 a team of French researchers, led by Dehaene and Pierre Pica, took that question deep into the Brazilian Amazon. With a solar-powered laptop, Pica performed the same dot-flashing experiment with people from isolated Indigenous villages. People from this Indigenous group had the same ability to distinguish between sufficiently different numbers of dots, even though they had limited or no formal mathematical training and spoke a language in which precise number words went no higher than five.

Around the same time, a different group of researchers, including developmental psychologists Elizabeth S. Spelke and Hilary Barth, then both at Harvard University, used a modified dot-flash experiment to show that five-year-olds in Massachusetts also had this ability. One possible explanation is that the children weren't really tracking the number of dots but rather were focusing on some other aspect, such as the total area the dots covered on the screen or the total perimeter of the cluster. When one collection of dots was replaced with a rapid sequence of audible tones, however, the children determined which quantity was greater—that is, whether they had heard more tones or seen more dots—equally as well as in the dots-only experiment. The children couldn't have used surface area or perimeter for that comparison, because the tones didn't have those features. Nor do dots have loudness or pitch. The children weren't using sound duration, either: the tones were presented sequentially over variable durations, whereas the dots were presented all at once for a fixed duration. It seems that the five-year-olds really did have a sense for the number of dots and tones.

Barth and her colleagues proceeded to show that these numerical abilities support basic forms of arithmetic. In another experiment, the five-year-olds saw two collections of blue dots move behind an opaque block, one after the other. From that point on, none of the blue dots were visible. Then some red dots appeared beside the block. The children were asked whether there were more blue dots or red dots in total. They answered correctly, indicating that they could add the two groups of blue dots together even though they couldn't see them anymore and then compare their total to that of the red dots. In 2021 Chuyan Qu and her colleagues in Elizabeth Brannon's laboratory at the University of Pennsylvania took this further, showing that children as young as five can perform approximate multiplication—an operation that isn't taught until third grade in the U.S.

It's easy to wonder whether these five-year-olds had learned something about numbers from their adult caretakers who had learned math in school. But

Number Sense in Infants

Fei Xu and Elizabeth S. Spelke repeatedly presented six-month-old infants with different displays containing either 8 or 16 dots. They then presented the infants with a new display. When the new display contained a new number of dots (8 when they had initially seen 16 or 16 when they had initially seen 8), the infants looked longer, an indication that they noticed something different. But when the new display contained the same number of dots, they did not. Six-month-olds showed the same behavior for other dots in a 2:1 ratio, such as 4 vs. 8 and 16 vs. 32. By nine months they also showed sensitivity to dots in a 3:2 ratio, such as 8 vs. 12, 4 vs. 6, and 16 vs. 24. Because the displays controlled for confounding variables such as brightness, density and area, these results suggest that infants have an innate number sense that, like visual-spatial acuity, becomes more precise with age.

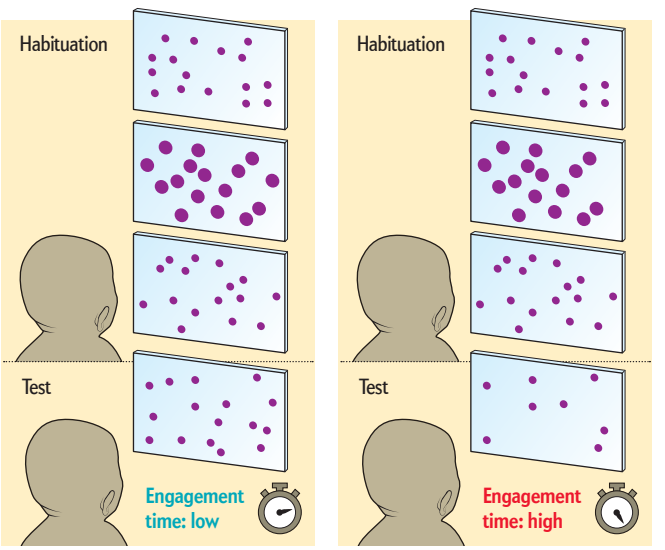

Source: "Large Number Discrimination in 6-Month-Old Infants," by Fei Xu and Elizabeth S. Spelke, in *Cognition*, Vol. 74; January 2000 (*reference*)

Graphics by Jen Christiansen

similar results have been found in a wide range of animal species. Wolves consider the size of their pack before deciding to hunt, preferring a group of two to six to attack an elk but a group of at least nine to take on a bison. Rats learn to press a lever a certain number of times in exchange for food. Ducks take account of how many morsels of food each of two people is throwing into a pond before deciding whom to approach. This behavior suggests that the number sense is evolutionarily ancient, similar to the ability to see colors or to feel warmth or cold.

These examples don't quite get at the question of whether the number sense is innate, though. Perhaps all they suggest is that formal schooling isn't necessary for humans or animals to learn to count. The ideal subjects for testing an innate number sense are newborn infants because they haven't had time to learn much of anything. Of course, they can't talk, so we can't ask them which of two collections contains more. They don't even crawl or reach. They are, however, capable of a simpler action: *looking*. By measuring where infants look and how much time they spend gazing there, scientists have discovered a window into their minds.

In a 2009 collaboration between Spelke and a team of researchers in France led by Véronique Izard and Arlette Streri, newborns at a hospital in Paris—all younger than five days old—listened for two minutes to auditory sequences containing either four sounds ("tuuu tuuu tuuu tuuu") or 12 sounds. The researchers then presented the infants with a visual display containing four or 12 objects. Infants are known to like looking at familiar things, such as their mother's face. Izard and her colleagues reasoned that if the infants extracted numbers from the auditory stimuli, they would prefer to look at a display containing a matching number of items—four objects right after they had heard sequences containing four sounds or 12 objects right after hearing sequences containing 12 sounds. Thus, they should look longer at the display when it numerically matches the sounds than when it does not.

And that's exactly what Izard and her colleagues found. (Just as a smile doesn't always mean the same thing, nor does longer looking: Here newborns looked longer at a match in number, indicating coincidence, whereas the six-month-olds described in the box on the preceding page looked longer at a change in number, reflecting surprise at the unexpected. In either case, reliable differences in looking behavior show that infants are sensitive to number.)

Critics such as Tali Leibovich-Raveh of Haifa University and Avishai Henik of Ben-Gurion University of the Negev in Israel raised concerns about overinterpreting these results, given that newborns have poor eyesight. But the fact that this result held for 15 of the

Preschool Mathematicians

Hilary Barth and her colleagues presented preschoolers with a collection of blue dots moving onto a screen and then being covered up by an opaque block. A collection of red dots then appeared, and the children were asked if there were more blue dots behind the barrier or red dots on the screen. They answered correctly even though they had not yet studied math.

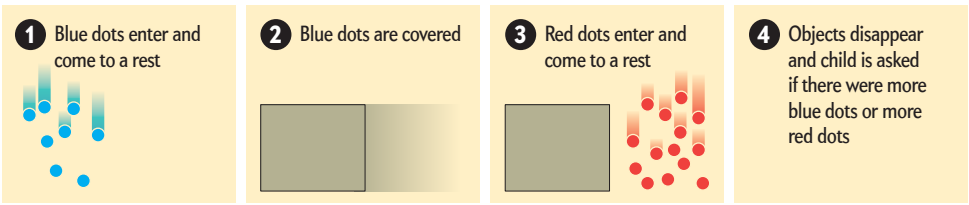

They also succeeded when they had to add two collections of blue dots and compare them with a third collection of red dots.

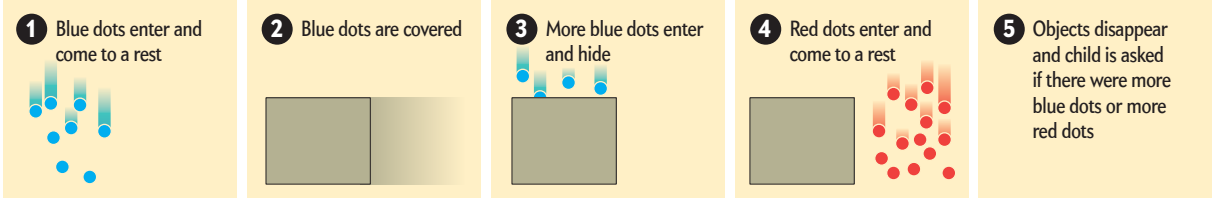

Interestingly, follow-up experiments found that their performance was just as good when the collection of red dots was indicated by a sequence of beeps. This is striking because it suggests that the preschoolers were really tracking number. For example, they couldn't have simply compared the total surface area of the collections, because the beeps didn't have a surface area. Conversely, they couldn't have simply compared the loudness of the collections, because the blue dots didn't make a noise.

Source: "Abstract Number and Arithmetic in Preschool Children," by Hilary Barth et al., in *PNAS*, Vol. 102, September 19, 2005 *(reference)*

Seeing Numbers, Part 1

Stare at the red cross for 30 seconds and then quickly turn the page. Does the collection of dots on the left now appear more or less numerous than the one on the right?

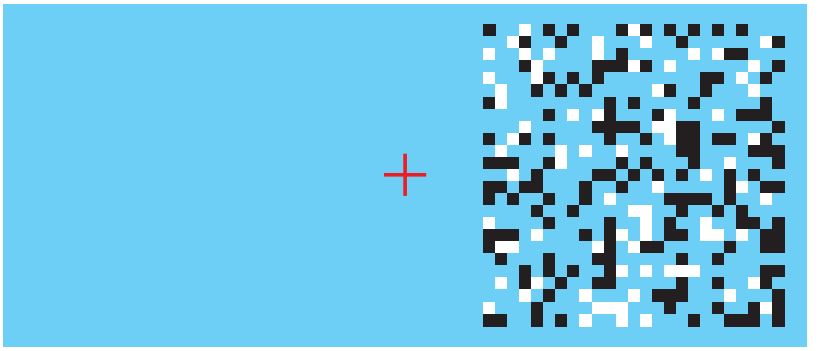

16 infants tested who didn't succumb to sleep or fussiness is certainly suggestive.

NUMBERS VS. NUMERALS

RECALL OUR SUGGESTION that when you opened your fridge at the party, you saw how many beers were present in much the way you saw their shape and color. This wasn't thoughtlessly worded: You didn't see the beer bottles and then judge their number. Rather the number sense enabled you to see the number like you see colors and shapes.

To clarify this idea, it's first important to distinguish numbers from numerals. Numerals are symbols used to refer to numbers. The numerals "7" and "VII" are distinct, but both refer to the same number. The claim that you see numbers should not be confused with a claim that you see numerals. Just as seeing the word "red" is distinct from seeing the color red, seeing the numeral "7" is not the same as seeing the number seven.

Moreover, just as seeing the size of a beer bottle does not involve a symbol like "12 oz" popping up in your visual field, seeing the number of beers in the fridge does not involve seeing a numeral such as "7." When you see the size of a beer bottle, it looks a certain way—a way that would change if the bottle got bigger or smaller. You can tell by looking that one bottle is bigger than another. Correspondingly, when you see how many beers there are, the beers look a certain way—a way that would change if there were more or fewer of them. Thus, you can tell, just by looking, whether there are more beers over here or over there.

Of course, even once numbers are distinguished from numerals, the concept of seeing numbers may still seem puzzling. After all, numbers are abstract. You can't point at them, because they aren't located in space, and your eyes certainly can't detect any light being reflected off them.

But the idea that you see numbers is not so different from the idea that you see shapes. Although you can see the sail of a boat as triangular, you cannot see a pure triangle on its own, independent of any physical objects. Likewise, although you can see some beers as being about seven in number, you cannot see the number seven all by itself. You can see shapes and numbers but only as attributes of an object or collection of objects that reflects light into your eyes.

So how can we tell when something is seen? If two lines appear on your COVID test, you might say that you can "see" that you have COVID. But that's loose talk. You certainly see the lines, but you merely judge that you have COVID. How can we draw this distinction scientifically?

There are many answers to this question, but one of the most helpful appeals to what is called perceptual adaptation. An example is the way a person's eyes eventually get used to the sun when they go picnicking on a sunny day. When that person later heads indoors, the bathroom appears dimly lit even when all the lights are on. After someone's eyes adapt to bright light, a normally lit room looks dark.

Whereas Plato proposed that we have a capacity to think about numbers, we argue that we have an ability to see or sense numbers.

Adaptation is a hallmark of perception. If you can perceive something, you can probably adapt to it—including its brightness, color, orientation, shape and motion. Thus, if numbers are perceived, people should adapt to numbers, too. This is precisely what vision researchers David Burr of the University of Florence and John Ross of the University of Western Australia reported in a paper published in 2008.

Burr and Ross showed that if a person stares at an array of lots of dots, a later array containing a middling number of dots will appear less numerous than it would have otherwise. For instance, they found that after someone stared at 400 dots for 30 seconds, they saw a group of 100 dots as if it had just 30 dots. Therefore, in much the way that our eyes get used to the sun, they can get used to large numbers, leading to striking visual effects.

Source: "A Visual Sense of Number," by David Burr and John Ross, in *Current Biology*, Vol. 18; 2008 *(reference)*

Other researchers, among them Frank Durgin of Swarthmore College, questioned whether the adaptation was to number as opposed to texture density (how frequently items appear in a given region of space). As a display of dots increases in number while the area it covers stays the same, it also increases in texture density. But a 2020 study by vision scientists Kevin DeSimone, Minjung Kim and Richard F. Murray teased these possibilities apart and showed that observers adapt to number independently of texture density. Strange as it sounds, humans see numbers.

NATURAL-BORN ENUMERATORS

DESPITE ABUNDANT EVIDENCE, contemporary empiricists—folks who follow in the tradition of Locke and Mill and believe that all mathematical knowledge is acquired through experience—remain skeptical about the existence of the number sense. After all, the ability to do arithmetic is traditionally considered a hard-won cultural achievement. Now we're supposed to believe babies do math?

Psychologists do have a checkered history when it comes to overinterpreting the numerical abilities of nonhuman animals. Undergraduate psychology majors are sternly warned of the "Clever Hans effect," named after a horse that was prematurely credited with sophisticated arithmetic abilities (not to mention the ability to tell time and spell long words in German). It was later revealed that he was simply responding to subtle cues in his trainer's behavior. Researchers nowadays take great care to avoid inadvertently cueing their subjects, but that doesn't resolve everything.

Rafael Núñez of the University of California, San Diego, argues, for instance, that the number sense simply couldn't represent numbers, because numbers are precise: 30 is exactly one more than 29 and exactly one less than 31. The number sense, in contrast, is imprecise: if you see 30 dots flashed on a screen, you'll have a rough idea of how many there are, but you won't know there are exactly 30. Núñez concludes that whatever the number sense is representing, it cannot be number. As he put it in a 2017 article in *Trends in Cognitive Sciences,* "a basic competence involving, say, the number 'eight,' should require that the quantity is treated as being categorically different from 'seven,' and not merely treated as often—or highly likely to be—different from it."

In our 2021 *Behavioral and Brain Sciences* article, we reply that such concerns are misplaced because any quantity can be represented imprecisely. You can represent someone's height precisely as 1.9 meters, but you can also represent it imprecisely as almost two meters. Similarly, you can represent the number of coins in your pocket precisely as five, but you can also represent it imprecisely as a few. You're representing height and number, respectively. All that changes is how you represent those quantities—precisely or imprecisely. Consequently, it's hard to see why the imprecision of the number sense should be taken to suggest that it's representing some attribute other than number.

This may seem like an issue of semantics, but it has a substantive implication. If we follow Núñez in supposing that the number sense doesn't represent numbers, then we need to say what it represents instead. And no one seems to have any good ideas about that. In many studies of the number sense, other variables—density, area, duration, height, weight, volume, brightness, and so on—have been controlled for.

Another reason to think the number sense concerns numbers (as opposed to height, weight, volume or other quantities) comes from late 19th-century German philosopher and logician Gottlob Frege. In his work on the foundations of arithmetic, Frege noted that numbers are unique in that they presuppose a way of describing the stuff they quantify. Imagine you point to a deck of cards and ask, "How many?" There's no single correct answer. We first need to decide whether we're counting the number of decks

Newborn Numeracy

Véronique Izard and her colleagues spent two minutes familiarizing newborn infants to sequences of four or 12 phonemes. The infants were then shown a display containing four or 12 seen items. Fifteen out of the 16 newborn infants who completed the study looked longer when the seen items matched the heard phonemes in number.

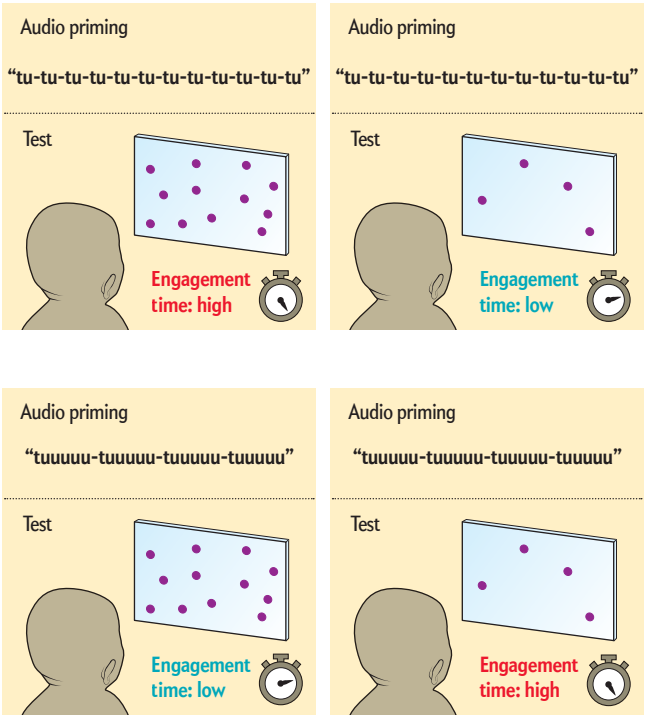

Source: "Newborn Infants Perceive Abstract Numbers," by Véronique Izard et al., in *PNAS*, Vol. 106; June 23, 2009 (*reference*)

Seeing Numbers, Part 2

If you kept your eyes fixated on the cross, you probably found that the left-hand array looked more numerous when you first turned the page. Your eyes adapted to the large number of dots on the right in the previous image, giving rise to a striking visual aftereffect.

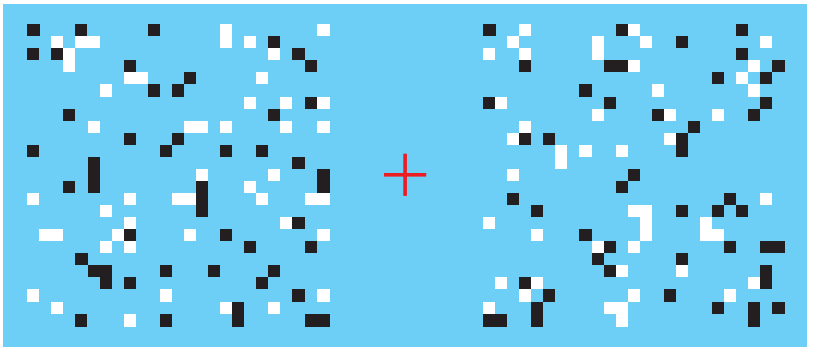

(one) or the number of cards (52), even though the 52 cards and the deck are one and the same.

Frege observed that other quantities aren't like this. If we want to know how much the cards weigh, we toss them on a scale and read off our answer. It won't make a difference to their weight if we think of them as a single deck or as a collection of cards. The same goes for their volume. The cards take up the same amount of space whether we describe them as one deck or 52 cards. (Of course, if we remove one card from the deck, it will have a different weight and volume than the full deck does. But then we're changing what we're describing, not just how we're describing it.) If the number sense is sensitive to how stuff is described, we can guess that it truly represents numbers and not other quantities.

This is precisely what we find when we apply Frege's insights. Work by a team of researchers led by Steven Franconeri of Northwestern University gives a vivid illustration. In a 2009 study, they presented subjects with a sequence of two screens containing circles and thin lines. Similar to many of the abovementioned experiments, subjects were asked which screen had more circles. They were also told to ignore the thin lines entirely. But when lines happened to connect two circles, effectively turning the pair of circles into a single dumbbell-shaped object, the subjects underestimated the number of circles on the screen. It seems they couldn't help but see a dumbbell as one object, even when they were trying to ignore the connecting lines and focus only on the circles.

The observers were not simply tracking some other quantity such as the total surface area of the items or the total number of pixels on the screen. After all, whether two circles and a line are connected to make a dumbbell doesn't affect the total surface area or pixel number. What it would, and seemingly did, affect is the perceived number of items in the display. So just as describing something as a deck of cards versus a collection of individual cards influences how you count it, whether you visually interpret something as a single dumbbell or a pair of circles influences how many items you seem to see—exactly as Frege would have predicted for a visual system that tracks number.

None of this is to deny that the mathematical abilities endowed by the number sense differ dramatically from the mature mathematical abilities most human adults possess. If you ask children for exactly 15 jelly beans, only ones who have learned to count in language can honor your request. But that's no reason to suppose their number sense isn't representing number. Just as children can perceive and discriminate distances long before they can think about them precisely, they have the ability to represent numbers before they learn to count in language and think of numbers precisely.

On their own, these innate mathematical abilities to perceive, add, subtract and operate over numbers are limited. But to understand how an infant develops into an Einstein, we must not underestimate babies' initial grasp of the world. To learn, we need something of substance to build on, and the number sense provides infants with part of the foundation from which new numerical abilities can arise—the ability to track coins and create monied economies, to develop modern mathematics or, more prosaically, to find those missing beers in the back of the fridge. SA

To understand how an infant develops into an Einstein, we must not underestimate babies' initial grasp of the world.

Jacob Beck is York Research Chair in the Philosophy of Visual Perception and an associate professor in the department of philosophy at York University in Toronto.

Sam Clarke is a MindCORE research fellow at the University of Pennsylvania, where he divides his time between philosophy and psychology. Beginning in July 2023, he will be an assistant professor of philosophy at the University of Southern California.

The number of gumballs in the jar on page 41 is 296.

age of opportunity

A refined understanding of the adolescent brain could lead to improvements in education and mental health

By Lydia Denworth

Illustrations by Alison Seiffer

HERE IS A PARABLE FOR OUR TIME: THERE ONCE WAS AN ADULT who wanted to encourage eighth graders to eat healthier food. The adult designed a lesson plan full of nutritional information—why fruit and vegetables are good for you, why junk food is bad for you, and so on. A similar approach had worked with younger children. But the eighth graders declared the intervention—and, if we're being honest, the adult—boring. They carried on eating junk food, some of them in greater quantities than they had before.

Versions of that story play out in real life all the time, although the age of the adolescents varies, and the goal could be anything from reducing bullying or depression to increasing engagement with math. With discouraging regularity, researchers find that what works with younger children is no longer effective with adolescents. Eighth grade seems to be the inflection point.

If we thought more carefully about what it is to be an eighth grader, however, down to the level of changes in the brain, our parable could have a happier ending. Thirteen-year-olds are concerned with status and respect—these kids do not want to feel patronized by adults. In a study published in 2019 in *Nature Human Behaviour,* instead of nutritional information, researchers showed more than 300 eighth graders in Texas investigative reports revealing that food company executives use unhealthy ingredients, target young adolescents in their marketing, and won't let their own children eat their products. The students were outraged and began to see healthy eating as a way of taking a stand against being manipulated. For the next three months the students made healthier snack purchases in the cafeteria. And in a follow-up study, the researchers found that the students, especially boys, with higher levels of testosterone (a marker of pubertal maturation in both boys and girls) were most likely to respond well to the intervention.

Over the past 15 years neuroscience has dramatically changed our understanding of the structural and functional changes in the brain during adolescence, which runs from around the age of 10 all the way into the mid-20s. It is a time of rapid brain growth and neuronal fine-tuning when young people are especially sensitive to social cues and rewards. More recent research has focused on how the adolescent brain interacts with the social environment. It shows that social context and acceptance strongly influence behavior. Adolescence might even constitute a sensitive period for social and emotional learning, a window of time when the brain is uniquely primed by neurochemical changes to make use of social cues for learning.

A growing group of researchers and clinicians see these neuroscientific findings as a chance to do things differently. When a young brain is looking for experience, teachers, parents and other influential adults should seek to capitalize on the richness of learning and stave off negative experiences such as smoking or drug use. This was a central idea in the 2019 National Academies of Sciences, Engineering, and Medicine report on the promise of adolescence, which called for investments in programs and interventions that use the brain's capacity to change during adolescence to promote beneficial shifts in young people's life trajectories.

A sensitive period for social and emotional processing also suggests that certain phases of adolescence may be more opportune than others for certain approaches. Early adolescence in particular—from roughly age nine

to 11—could be an opportunity to launch kids on a positive path by buttressing their sense of self and motivation to learn. The nutrition experiment shows the benefits of fine-tuning interventions for middle adolescents, who have been through puberty. And no one wants to suggest that it's ever too late to help young people in trouble, especially given that the most serious behavioral and health problems of adolescence tend to occur at 16 and beyond.

To meaningfully compare the results of which interventions work best at age 10 or 14 or 18 requires extensive longitudinal studies, which have not yet been done. Even so, the advances in developmental science appear poised to lead to wiser, more effective approaches to supporting young people's education and physical and mental health. These new methods emphasize adolescents' concern with status and respect, their evolving sense of self in relation to the wider world, and their need to contribute and find purpose. Similar ideas already underpin the growing interest in social and emotional learning among educators. Rather than focusing on the storminess of the teenage years, these ideas offer a sunnier view of adolescence as a window of opportunity.

RETHINKING ADOLESCENCE

FOR DECADES much of the research on adolescence focused on its dark side. Although those years are the physically healthiest period in life, when strength, speed, reaction time, reasoning abilities and immune function all improve or peak, adolescence also brings alarming increases in rates of accidents, suicide, homicide, depression, alcohol and substance use, violence, reckless behaviors, eating disorders, obesity and sexually transmitted disease compared with the rates for younger children.

But a different interpretation of adolescence emerged in the 2000s, stemming from two important new findings. Neuroscientists showed that puberty ushers in a period of exuberant neuronal growth followed by a pruning of neural connections that is second only to the similar process that occurs in the first three years of life. They also showed that the maturation of the adolescent brain is not linear. The limbic system, a collection of brain areas that are sensitive to emotion, reward, novelty, threat and peer expectations, undergoes a growth spurt while the brain areas responsible for reasoning, judgment and executive function continue their slow, steady march toward adulthood. The resulting imbalance in the developmental forces helps to explain adolescent impulsivity, risk taking, and sensitivity to social reward and learning. In an evolutionary sense, much of adolescents' behavior pushes them to leave the safety of family to explore the larger social world—a step on the way to becoming independent adults.

Another line of research, from the human connectome project, shows that adult brains vary in their patterns of neural connections throughout the brain, whereas children's connectomes are less distinctive. Those differentiated patterns of connection emerge in adolescence—between the ages of 10 and 16, just when social values and cognition are developing quickly. And the changes in the connectome data show up on average a year to a year and a half earlier in girls than in boys, just like puberty does, which suggests that the two things are intertwined.

The idea that adolescence might constitute a sensitive period for social and emotional processing was put forward in 2014 by neuroscientists Sarah-Jayne Blakemore and Kathryn Mills, now at the University of Cambridge and the University of Oregon, respectively. Previous research had assumed that social-cognitive abilities such as theory of mind were mature by the middle of childhood, but Blakemore and Mills laid out the many continuing changes across adolescence in social cognition and the network of brain regions governing social behavior.

Sensitive, or critical, periods are windows of time when the brain is primed to make specific neural connections that depend on the input received. They are timed to when important information is available and most useful for development. For sensory processing such as vision and hearing, such periods are well defined with an opening, peak and closing. A brain deprived of sight or sound early in development will never be able to see or hear normally. Likewise, a sensitive period for language acquisition explains why people who learn a foreign language after puberty typically have an accent. Sensitive periods for social learning have been harder to pin down.

During adolescence the brain may be uniquely primed by neurochemical changes to make use of social cues for learning.

Animal research has identified some versions of sensitive periods for social learning. Songbirds can delay the closing of the sensitive period for vocal learning if they need more time to learn their songs, which usually happens in adolescence. "It's a gorgeous example of a sensitive period for learning that has social function," says Linda Wilbrecht of the University of California, Berkeley, who has studied sensitive periods in songbirds, mice and humans.

Neuroscientist Gül Dölen and her colleagues at Johns Hopkins University identified an adolescent critical period in mice for something called social conditioned place preference (social CPP). The research followed up on an observation by the late Estonian neuroscientist Jaak Panksepp. He presented mice with two different kinds of bedding—on one, the mice were alone; on the other, they were with friends. When the mice subsequently had a choice of bedding, adolescents, in particular, showed a preference for the bedding that carried a memory of friends.

Dölen ran similar experiments with around 900

mice at 14 different ages and mapped out exactly when this preference for place occurs. Triggered by changes in oxytocin that lead to increased synaptic plasticity, it peaks 42 days after birth (roughly age 14 in humans), when the mice become sexually mature. "It's a really important stage of their lives when they're leaving the nest and trying to create their own groups," Dölen says. "[In] that window of time, when they're really sensitive to what other members of their group are doing, when they're learning from their group, when they're forming attachments to the group—that's when that peaks." It seems the brain is suddenly alert to and rewarded by information that it had previously ignored. "There's information flowing by us all the time," Wilbrecht says. "Once puberty and hormones pass through the circuit, suddenly those cues have meaning. They don't have salience until you shift into the adolescent phase."

PRIMED FOR LEARNING

THESE WINDOWS of rapid change create both learning opportunities and vulnerabilities. What adolescents are learning is all-important. "The adolescent brain is primed for social and emotional learning, to explore, to interact, to take chances so they can learn, but it all depends on what we do to give them scaffolded opportunities in order to learn," says psychologist Andrew Fuligni of the University of California, Los Angeles. Harmful experiences may lead to negative spirals from which it's hard to recover. Research has shown that earlier experimentation with alcohol and drugs makes an adolescent more likely to become addicted.

Research shows that adolescents have a need to contribute to society, and doing so can safeguard against anxiety and depression.

"When your brain is undergoing rapid reorganization, that's probably not the best time to introduce external chemicals," says developmental psychologist Anthony Burrow of Cornell University. "Your body and brain are paying attention in a slightly different way. [Your brain is] going to organize itself around what you've done to it at that particular moment."

Protective factors in the adolescent's environment could support positive trajectories. What do protective factors look like? They include supportive relationships with family and caretakers and access to resources such as scaffolded opportunities to learn in positive ways. They also include some elements that have previously been underappreciated. Fuligni's research shows that adolescents have a need to contribute to society, and doing so makes them feel valued and can safeguard against anxiety and depression. "Part of what the brain is designed to do during the teenage years is to learn how to contribute to the social world," Fuligni says.

This need is particularly significant in adolescence, he argues, because it's a time when the social world is expanding and young people are becoming capable of "making contributions of consequence." These contributions can occur within peer groups, in the family, or at a larger societal level. It's no accident that recent social protest movements for gun control and against structural racism have been spearheaded in large part by young people.

The specifics of what today's adolescents are learning—and what they are not—may bear on the alarming rises in depression, anxiety and suicidal ideation at that age compared with earlier ages (as well as with previous generations). Some of the information they encounter about mental health may be amplifying their problems, says psychologist Nicholas Allen of the University of Oregon. He points to the controversial Netflix series *13 Reasons Why,* which when it first aired in 2017 included a detailed depiction of a character's suicide and which research suggests was associated with an increase in adolescent suicides. "Whether it's a supportive, solution-oriented discussion or whether it's a ruminative, hopeless discussion will have a big effect," Allen says. "Too often adolescents who are tending toward depression, anxiety or suicidal ideation have a tendency to ruminate, and they find friends—both online and offline—who feed that tendency rather than help the teenagers move beyond it."

EFFECTIVE INTERVENTIONS

THERE IS STILL DEBATE about how best to use the new neuroscientific knowledge to help adolescents. "We've learned an enormous amount about the brain, but the application of that knowledge is not straightforward," Allen says.

A big question is when to intervene. One argument for zeroing in on early adolescence is to act preemptively. Because so many of the problems of adolescence occur in the mid- to late teenage years, many interventions target that time. "If you're a developmentalist, that is too late," says Ronald Dahl, a pediatrician and developmental scientist and founder of the Center for the Developing Adolescent at U.C. Berkeley. "Smaller, more subtle, positive interventions earlier are probably a much more promising way to improve population health." The logic of that idea first struck Dahl when he was still practicing as a pediatrician. At conferences, he started mentioning the importance of reaching kids early and found educators nodding their heads. They introduced Dahl to the idea of the fifth grade slump and the eighth grade cliff, a phenomenon in which children's disengagement with education starts slowly with a dip in grades and participation around fifth grade, when most students are 10, and accelerates so that those same students are failing three years later.

The neuroscience also suggests that acting early could make sense. "What we're increasingly learning is that there's another node of new plasticity around the time puberty starts," Dahl said at a conference in

early 2020. "We talk about this as a high-stakes pivotal transition in terms of patterns that are beginning to be shaped." In a study in Tanzania, Dahl and his colleagues succeeded both quantitatively and qualitatively in reducing ideas of gender inequality among 10- and 11-year-olds with a series of technology lessons at which girls were as likely to shine as boys.

Others are wary of focusing too much on any one phase. They emphasize that what neuroscience contributes to the discussion is a reminder of what to prioritize. "What is the thing at this stage of life that is most plastic, that is open for input? That tells you where the risk is, but it also tells you where the opportunity is," Allen says. "What the brain science says is that you should be looking in this area: social and emotional learning."

It is not surprising then that those interventions that look most promising take into account adolescents' desire for status and respect, as well as their need to contribute and find a sense of purpose. According to Fuligni, the most successful volunteer programs give adolescents a say in what to work on and a chance to reflect on the work, and the projects also feel meaningful. Meaning seems to matter in other efforts, too. In a study of early adolescents participating in a 4-H program, Burrow found that those who were asked to write about their sense of purpose before engaging in an educational activity were more likely to engage with the activity and find it important and interesting. "Purpose is a pretty powerful form of identity capital because it's not just an answer to the question of who you are, but it's an answer to the question of who you're going to be and the direction you're heading in," Burrow says. "It's got legs."

Psychologist David Yeager of the University of Texas at Austin has explored how best to frame messages to teenagers and studied whether their effectiveness interacts with pubertal maturation, a sign that the neurochemical changes are playing a role. "You should be able to show that if you communicated respectfully to teenagers in a way that felt authentic and supported their autonomy and independence, you should have bigger effects for adolescents, especially if they're more mature in terms of their puberty," he says.

His research bears that out. One series of experiments showed that the framing of a request to take medicine predicted different rates of compliance and that those rates varied with testosterone levels. Some 18- and 19-year-olds came into the lab and were given instructions in a condescending way: *I'm the expert, I know what's good for you, take this.* Another group of young adults were given instructions in a more respectful manner: *Let me explain the reasons this medicine can be useful.*

For ethical reasons, the medicine in question was actually a spoonful of Vegemite, a notoriously strong-tasting condiment. Asked respectfully, people were twice as likely to take the Vegemite. Furthermore, participants with higher testosterone levels were significantly less likely to take the medicine in the disrespectful condition and more likely to comply in the respectful condition. When Yeager and his colleagues manipulated testosterone levels with a nasal inhaler, they found that doing so made individuals with naturally low testosterone levels behave just like those with naturally high testosterone levels.

While the medicine study was a nice test of how respectfulness might matter, Yeager says that the 2019 nutrition study informing eighth graders about unsavory food industry practices, which he helped lead, is even more promising. "That's the first direct evidence that these pubertal hormones sensitize you to status and respect and therefore change the way you respond to health messages," he says. "And not just how you respond in the moment but the way you internalize them and continue to keep acting on them after the treatment is over."

In other words, now we know more about what causes adolescents to put up a wall and resist attempts to change their habits, beliefs and ways of coping. That same knowledge offers ways to break down that wall. "It's only recently that we know how to work with those sensitivities and not against them," Yeager says. "I'd like it to be a wake-up call for adults who work with kids." SA

Lydia Denworth is a Brooklyn, N.Y.–based science writer, a contributing editor for *Scientific American,* and author of *Friendship: The Evolution, Biology, and Extraordinary Power of Life's Fundamental Bond* (W. W. Norton, 2020).

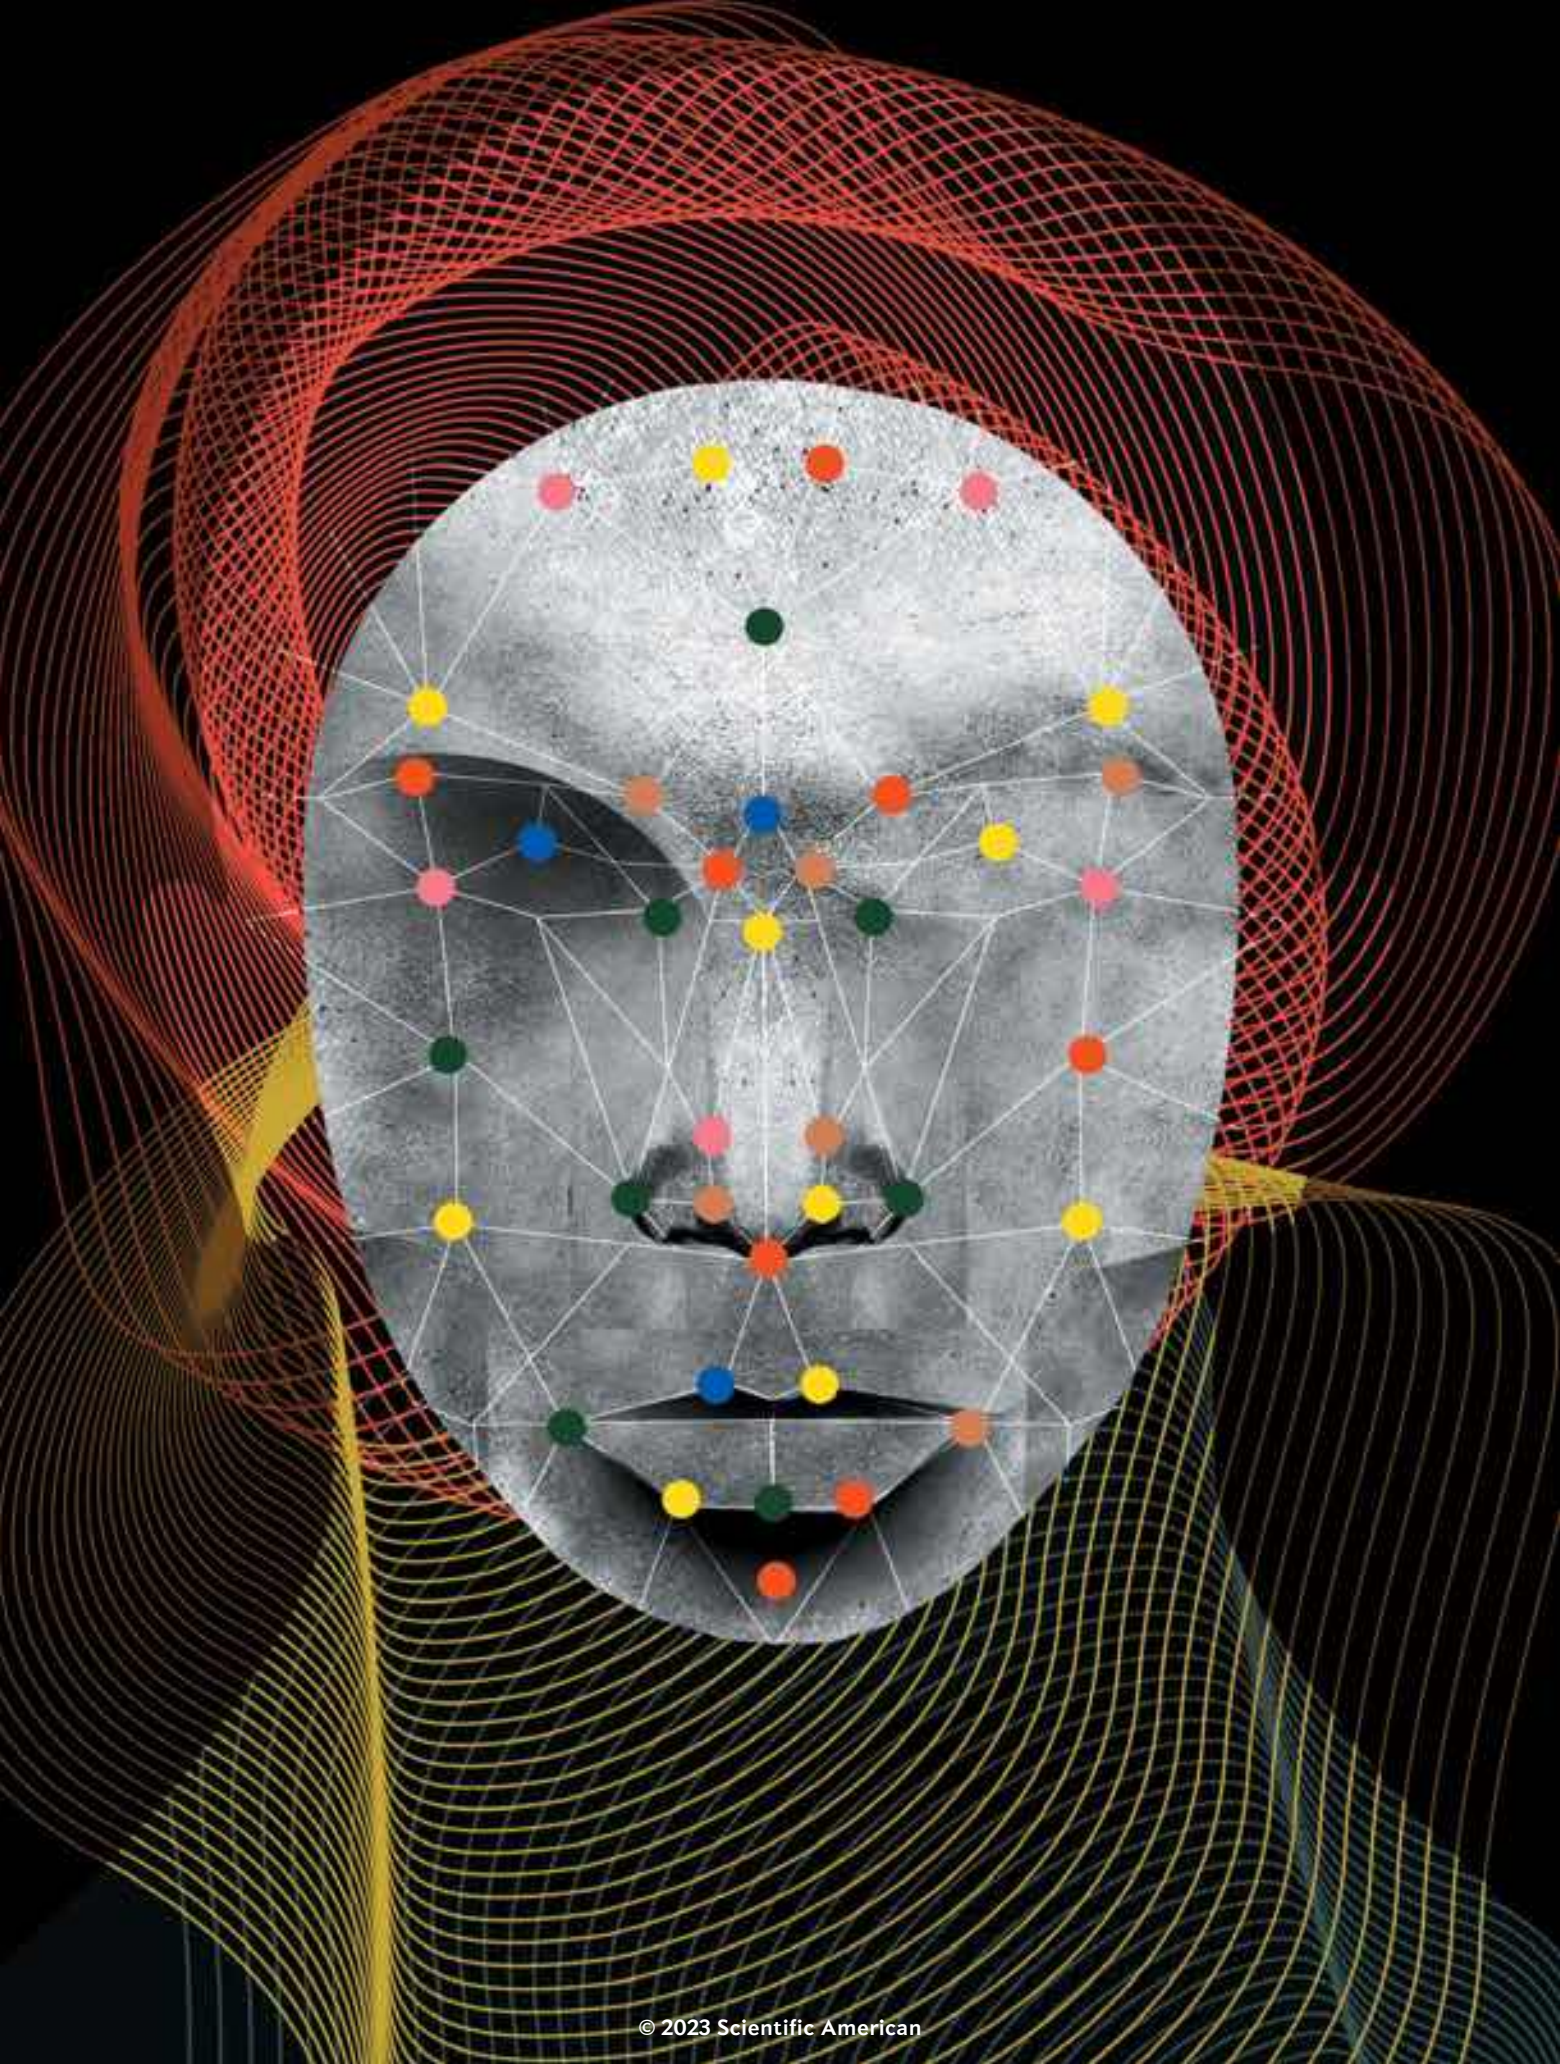

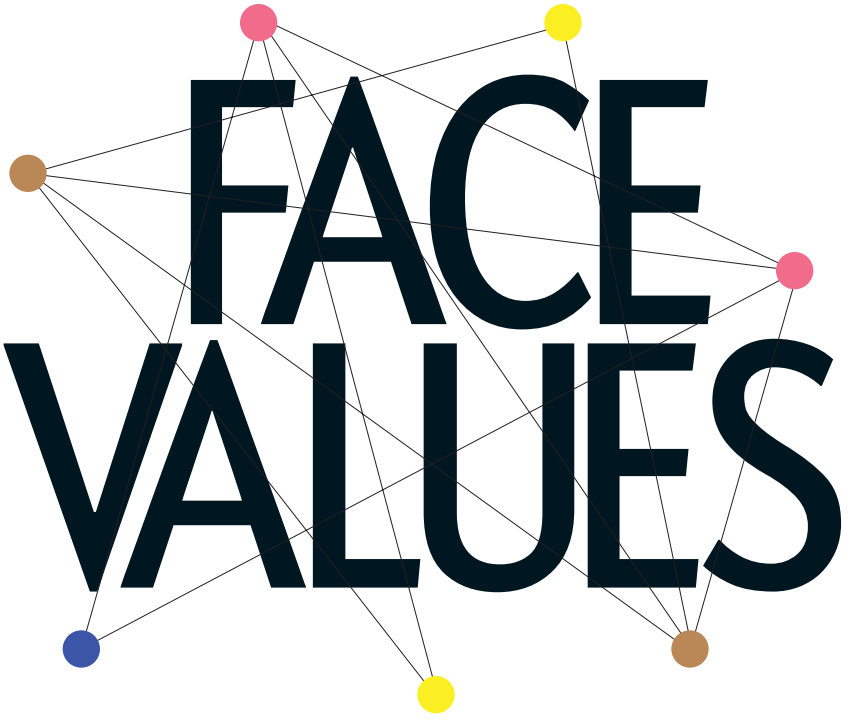

Brain regions that process faces reveal deep insights into the neural mechanisms of vision

By Doris Y. Tsao

Illustration by Brian Stauffer

WHEN I WAS IN HIGH SCHOOL, I LEARNED ONE DAY ABOUT THE DENSITY of curves in an introductory course on calculus. A simple pair of differential equations that model the interactions of predators and prey can give rise to an infinite number of closed curves—picture concentric circles, one nested within another, like a bull's-eye. What is more, the density of these curves varies depending on their location.

This last fact seemed so strange to me. I could easily imagine a finite set of curves coming close together or pulling apart. But how could an infinity of curves be denser in one region and less dense in another? I soon learned that there are different types of infinity with paradoxical qualities, such as Hilbert's Hotel (where the rooms are always fully booked but new guests can always be accommodated) and the Banach-Tarski apple (which can be split into five pieces and rearranged to make two apples with the same volume as the original). I spent hours poring over these mathematical proofs. Ultimately they struck me as symbolic magic of no real consequence, but the seed of interest had taken root.

Later, as an undergraduate at the California Institute of Technology, I learned about the experiments of David Hubel and Torsten Wiesel and their landmark discovery of how a region in the brain called the primary visual cortex extracts edges from the images relayed from the eyes. I realized that what had mystified me back in high school was the act of trying to *imagine* different densities of infinity. Unlike the mathematical tricks I had studied in high school, the edges that Hubel and Wiesel described are processed by neurons, so they actually exist in the brain. I came to recognize that visual neuroscience was a way to understand how this neural activity gives rise to the conscious perception of a curve.

The sense of excitement this realization triggered is hard to describe. I believe at each stage in life one has a duty. And the duty of a college student is to dream, to find the thing that captures one's heart and seems worth devoting a whole life to. Indeed, this is the single most important step in science: to find the right problem. I was captivated by the challenge of understanding vision and embarked on a quest to learn how patterns of electrical activity in the brain are able to encode perceptions of visual objects—not just lines and curves but even objects as hard to define as faces. Accomplishing this objective required pinpointing the specific brain regions dedicated to facial recognition and deciphering their underlying neural code—the means by which a pattern of electrical impulses allows us to identify people around us.

FACE PATCHES DO ACT AS AN ASSEMBLY LINE TO SOLVE ONE OF THE BIG CHALLENGES OF VISION: HOW TO RECOGNIZE THINGS AROUND US DESPITE CHANGES IN THE WAY THEY LOOK.

The journey of discovery began in graduate school at Harvard University, where I studied stereopsis, the mechanism by which depth perception arises from differences between the images in the two eyes. One day I came across a paper by neuroscientist Nancy Kanwisher, now at the Massachusetts Institute of Technology, and her colleagues, reporting the discovery of an area in the human brain that responded much more strongly to pictures of faces than to images of any other object when a person was inside a functional magnetic resonance imaging (fMRI) brain scanner. The paper seemed bizarre. I was used to the brain being made of parts with names like "basal ganglia" and "orbitofrontal cortex" that had some vague purpose one could only begin to fathom. The concept of an area specifically devoted to processing faces seemed all too comprehensible and therefore impossible. Anyone could make a reasonable conjecture about the function of a face area—it should probably represent all the different faces that we know and something about their expression and gender.

As a graduate student, I had used fMRI on monkeys to identify areas activated by the perception of three-dimensionality in images. I decided to show pictures of faces and other objects to a monkey. When I compared activation in the monkey's brain in response to faces with activation for other objects, I found several areas that lit up selectively for faces in the temporal lobe (the area underneath the temple)—specifically in a region called the inferotemporal (IT) cortex. Charles Gross, a pioneer in the field of object vision, had discovered face-selective neurons in the IT cortex of macaques in the early 1970s. But he had reported that these cells were randomly scattered throughout the IT cortex. Our fMRI results provided the first indication that face cells might be concentrated in defined regions.

FACE PATCHES

AFTER PUBLISHING MY WORK, I was invited to give a talk describing the fMRI study as a candidate for a faculty position at Caltech, but I was not offered the job. Many people were skeptical of the value of fMRI, which measures local blood flow, the brain's plumbing. They argued that showing increased blood flow to a brain area when a subject is looking at faces falls far short of clarifying what neurons in the area are actually encoding because the relation between blood flow and electrical activity is unclear. Perhaps by chance these face patches simply contained a slightly larger number of neurons responsive to faces, like icebergs randomly clustered at sea.

Because I had done the imaging experiment in a monkey, I could directly address this concern by inserting an electrode into an fMRI-identified face area and asking, What images drive single neurons in this region most strongly? I performed this experiment together with Winrich Freiwald, then a postdoctoral fellow in Margaret Livingstone's laboratory at Harvard, where I was a graduate student. We presented faces and other objects to a monkey while amplifying the electrical activity of individual neurons recorded by the electrode. To monitor responses in real time, we converted the neurons' electrical signals to an audio signal that we could hear with a loudspeaker in the lab.

This experiment revealed an astonishing result: almost every single cell in the area identified through fMRI was dedicated to processing faces. I can recall the excitement of our first recording, hearing the "pop" of cell after cell responding strongly to faces and very little to other objects. We sensed we were on to something important, a piece of cortex that could reveal the brain's high-level code for visual objects. Marge remarked on the face patches: "You've found a golden egg."

I also remember feeling surprised during that first experiment. I had expected the face area would contain cells that responded selectively to specific individuals, analogous to orientation-selective cells in the primary visual cortex that each respond to a specific edge orientation. In fact, a number of well-publicized studies had suggested that single neurons can be remarkably selec-

tive for the faces of familiar people—responding, say, only to Jennifer Aniston. Contrary to my expectation, each cell seemed to fire vigorously for almost any face.

I plugged madly away at Photoshop during these early experiments and found that the cells responded not just to faces of humans and monkeys but even to highly simplified cartoon faces.

Observing this phenomenon, I decided to create cartoon faces with 19 different features that seemed pertinent to defining the identity of a face, including inter-eye distance, face aspect ratio and mouth height, among other characteristics. We then went on to alter these values—the inter-eye distance, for instance, varied from almost cyclopean to just inside the face boundary. Individual cells responded to most faces but interestingly did not always exhibit the exact same rate of firing with all faces. Instead there was a systematic variation in their response: when we plotted the firing of cells for the different cartoon features, we found a pattern in which there was a minimal response to one feature extreme, such as the smallest inter-eye distance, and a maximal response to the opposite extreme—the largest eye separation—with intermediate responses to feature values in the middle. The response as a function of the value for each feature looked like a ramp, a line slanted up or down.

Where Are the Face Detectors?

A set of six nodes in the inferotemporal cortex of both brain hemispheres specializes in identifying faces. These "face patches" function as an assembly line: in the middle lateral and middle fundus patches, one neuron might become active when faces look straight ahead; another might turn on for a face looking to the right. At the end of the assembly line, in the anterior medial patch, varying views are stitched together. Neurons in this patch are active in response to the face of a specific individual, no matter if the view is from the front or side. Responses from a face patch of one monkey are generated for faces but not objects (*red areas in* A) and for the same individual, such as the dark-haired man, from varying angles (*red areas in* B).

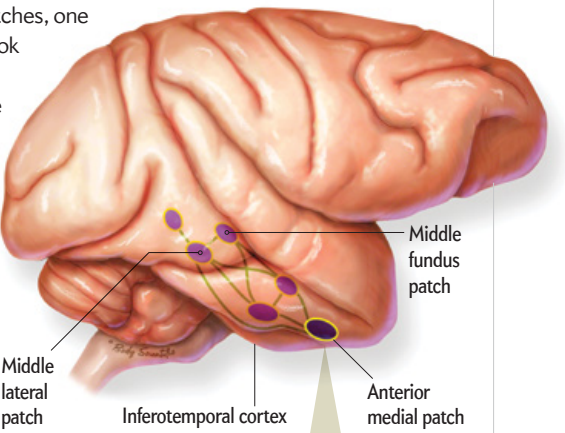

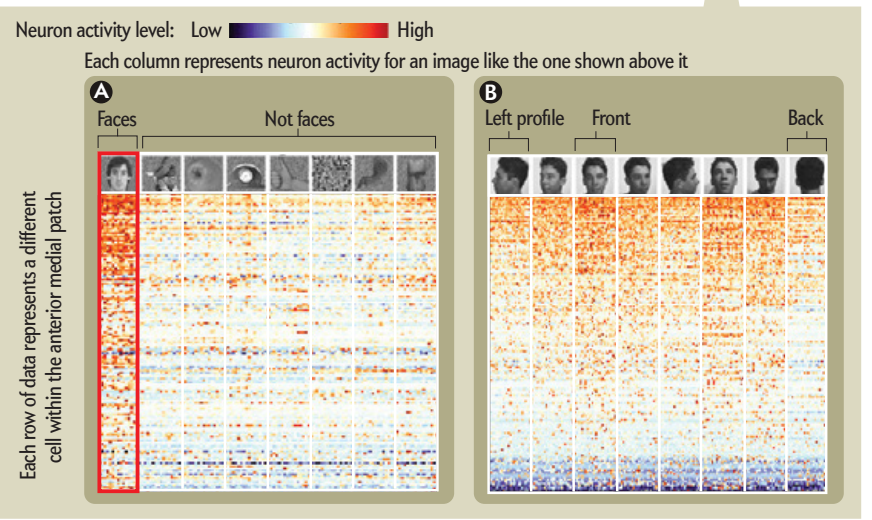

Once again, I was invited to give a job talk at Caltech. Returning, I had more to offer than just fMRI images. With the addition of the new results from single-cell recordings, it was clear to everyone that these face patches were real and likely to have an important role in facial recognition. Furthermore, understanding their underlying neural processes seemed like an effective way to gain traction on the general problem of how the brain represents visual objects. This time I was offered the job.

CONTRAST IS KEY

AT CALTECH, my colleagues and I dug deeper into the question of how these cells detect faces. We took inspiration from a paper by Pawan Sinha, a vision and computational neuroscientist at M.I.T., that suggested faces could be discerned on the basis of specific contrast relations between different regions of the face—whether the forehead region is brighter than the mouth region, for example. Sinha suggested a clever way to determine *which* contrast relations can be used to recognize a face: they should be the ones that are immune to changes in lighting. For example, "left eye darker than nose" is a useful feature for detecting a face because it does not matter if a face is photographed with lighting from above, left, right or below: the left eye is *always* darker than the nose (check for yourself).

From a theoretical standpoint, this idea provides a simple, elegant computational mechanism for facial recognition, and we wondered whether face cells might be using it. When we measured the response of cells to faces in which different regions varied in brightness, we found that cells often had a significant preference for a particular contrast feature in an image.

To our astonishment, almost all the cells were wholly consistent in their contrast preferences—just a single cell was found that preferred the opposite polarity. Moreover, the preferred features were precisely those identified by Sinha as being invulnerable to lighting

From "Functional Compartmentalization and Viewpoint Generalization within the Macaque Face-Processing System," by Winrich A. Freiwald and Doris Y. Tsao, in *Science*, Vol. 330; November 5, 2010 (*data grids* and *photo insets*)

Graphic by Body Scientific (brain)

Shape + Appearance = Face

Identifying the face patches was only a first step. It then became necessary to explore what happens in the neurons within each patch, setting off a search for the brain's coding scheme for faces. To derive quantitative measures for faces, the Tsao laboratory at Caltech came up with 25 features for shape and 25 for appearance that could be used by each neuron in a face patch—a 50-dimensional face space. The shape features can be thought of as those defining the skeleton—how wide the head is or the distance between the eyes. The appearance features specify the face's surface texture (complexion, eye or hair color, and so on).

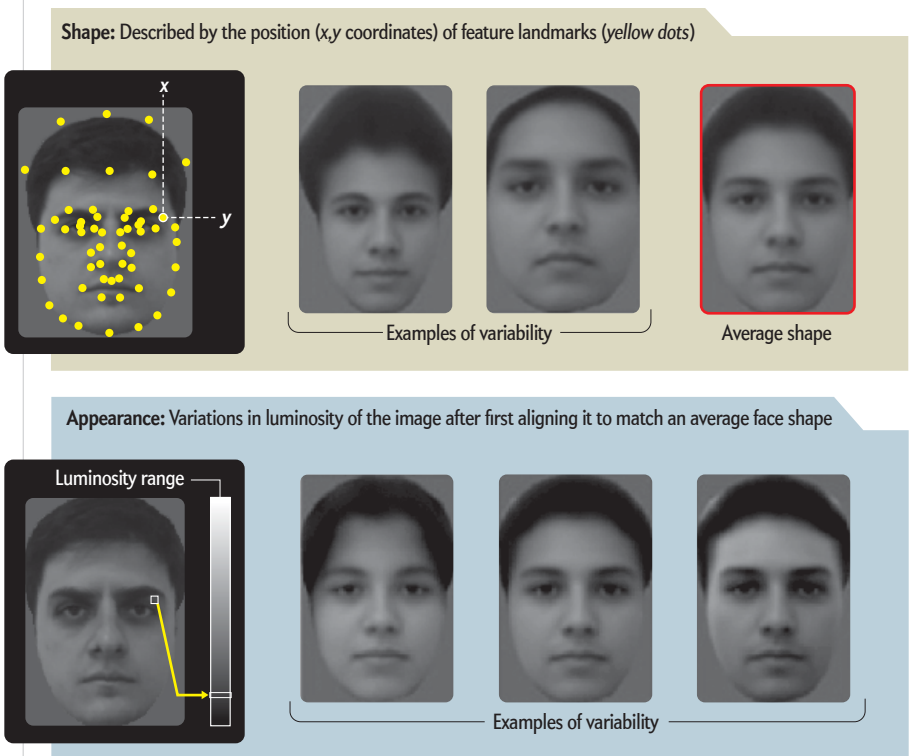

changes. The experiment thus confirmed that face cells use contrast relations to detect faces.

More broadly, the result confirmed that these cells truly were face cells. At talks, skeptics would ask, How do you know? You can't test every possible stimulus. How can you be sure it's a face cell and not a pomegranate cell or a lawn mower cell? This result nailed it for me. The precise match between the way cells reacted to changes in contrast between different parts of the face and Sinha's computational prediction was uncanny.

Our initial experiments had revealed two nearby cortical patches that lit up for faces. But after further scanning (with the help of a contrast agent that increased severalfold the robustness of the signal), it became clear that there are in fact six face patches in each of the brain's two hemispheres (making a dozen golden eggs total). They are distributed along the entire length of the temporal lobe. These six patches, moreover, are not randomly scattered throughout the IT cortex. They are located in similar locations across hemispheres in each animal. Work by our group and others has found that a similar pattern of multiple face patches spanning the IT cortex exists in humans and other primates such as marmosets.

This observation of a stereotyped pattern suggested that the patches might constitute a kind of assembly line for processing faces. If so, one would expect the six patches to be connected to one another and each patch to serve a distinct function.

To explore the neural connections among patches, we electrically stimulated different patches with tiny amounts of current—a technique called microstimulation—while the monkey was inside an fMRI scanner. The goal was to find out what other parts of the brain light up when a particular face patch is stimulated. We discovered that whenever we stimulated one face patch, the other patches would light up, but the surrounding cortex would not, indicating that, indeed, the face patches are strongly interconnected. Furthermore, we found that each patch performs a different function. We presented pictures of 25 people, each at eight different head orientations, to monkeys and recorded responses from cells in three regions: the middle lateral and middle fundus patches (ML/MF), the anterior lateral patch (AL) and the anterior medial patch (AM).

We found striking differences among these regions. In ML/MF, cells responded selectively to specific views. For example, one cell might prefer faces looking straight ahead, whereas another might opt for faces looking to the left. In AL, cells were less view-specific. One class of cells responded to faces looking up, down and straight ahead; another responded to faces looking to the left or right. In AM, cells responded to specific individuals regardless of whether the view of the face was frontal or in profile. Thus, at the end of the network in AM, view-specific representations were successfully stitched into a view-invariant one.

Apparently face patches do act as an assembly line to solve one of the big challenges of vision: how to recognize things around us despite changes in the way they look. A car can have any make and color, appear at any viewing angle and distance, and be partially obscured by closer objects such as trees or other cars. Recognizing an object despite these visual transformations is called the invariance problem, and it became clear to us that a major function of the face-patch network is to overcome this impediment.

From "The Code for Facial Identity in the Primate Brain," by Le Chang and Doris Y. Tsao, in *Cell*, Vol. 169, No. 6, June 1, 2017 (*face images*)

Graphics by Jen Christiansen

Given the great sensitivity of cells in face patches to changes in facial identity, one might expect that altering these cells' responses should modify an animal's perception of facial identity. Neuroscientists Josef Parvizi and Kalanit Grill-Spector of Stanford University had electrically stimulated a face-patch area in human subjects who had electrodes implanted in their brains for the purpose of identifying the source of epileptic seizures and found that stimulation distorted the subjects' perception of a face.

We wondered whether we would find the same effect in monkeys when we stimulated their face patches. Would doing so alter the perception only of faces, or would it affect that of other objects as well? The boundary between a face and a nonface object is fluid—one can see a face in a cloud or an electrical outlet if prompted. We wanted to use electrical microstimulation as a tool to delineate precisely what constitutes a face for a face patch. We trained monkeys to report whether two sequentially presented faces were the same or different. Consistent with the earlier results in humans, we found that microstimulation of face patches strongly distorted perception so that the animal would always signal two identical faces as being different.

Interestingly, microstimulation had no effect on the perception of many nonface objects, but it did significantly affect responses to a few objects whose shape is consistent with a face—apples, for one. But why does this stimulation influence the perception of an apple?

One possibility is that the face patches are typically used to represent not just faces but also other round objects like apples. Another hypothesis is that face patches are not normally used to represent these objects, but stimulation induces an apple to appear face-like. It remains unclear whether face patches are useful for detecting any nonface objects.

CRACKING THE CODE

UNCOVERING the organization of the face-patch system and properties of the cells within was a major accomplishment. But my dream when we first began recording from face patches was to achieve something more. I had intuited that these cells would allow us to crack the neural code for facial identity. That means understanding how individual neurons process faces at a level of detail that would let us predict a cell's response to any given face or decode the identity of an arbitrary face based only on neural activity.

The central challenge was to figure out a way to describe faces quantitatively with high precision. Le Chang, then a postdoc in my lab, had the brilliant insight to adopt a technique from the field of computer vision called the active appearance model. In this approach, a face has two sets of descriptors, one for shape and another for appearance. Think of the shape features as those defined by the skeleton—how wide the head is or the distance between the eyes. The appearance features define the surface texture of the face (complexion, eye or hair color, and so on).

To generate these shape and appearance descriptors for faces, we started with a large database of face images. For each face, we placed a set of markers on key features. The spatial locations of these markers described

The Face Code, at Last

Having 50 coordinates that describe shape and appearance allows for a description of neurons' firing in response to a particular face—a description that functions as a code that can be visualized geometrically. In this code, each face cell receives inputs for a face in the form of the 50 coordinates, or dimensions. The neuron then fires with a particular intensity in response to a certain face (*red outlines*), along what is called the preferred axis. The intensity increases steadily (monotonically) along the preferred axis. Furthermore, the response is the same for every face on an axis at right angles to the preferred axis, even though those faces may look very different. This axis model of facial coding differs from a previous exemplar model that suggests that each neuron fires with maximum intensity to a single most preferred face.

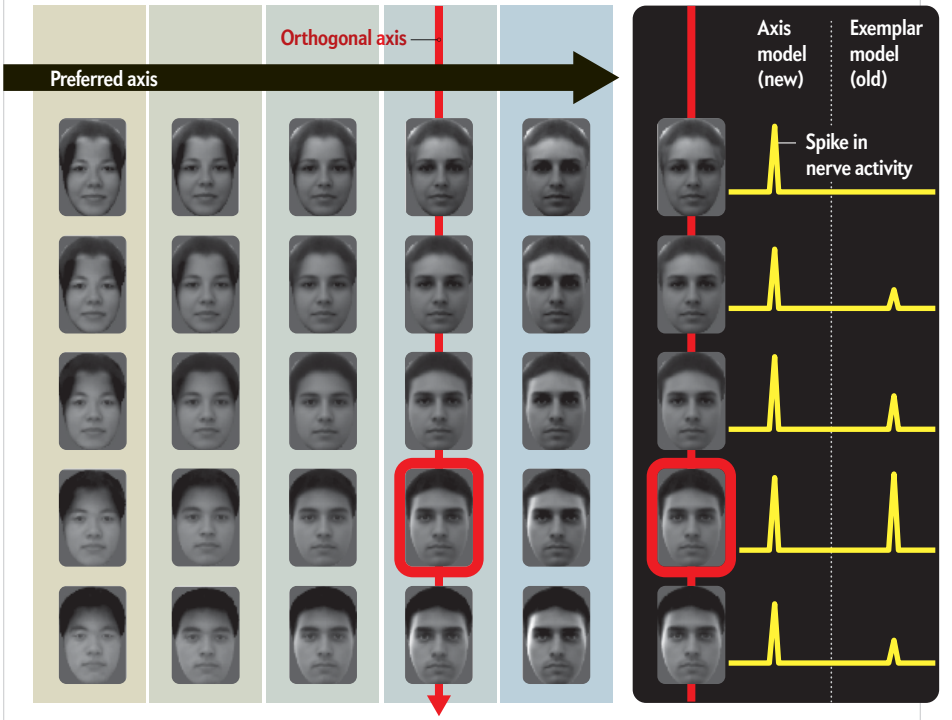

From "The Code for Facial Identity in the Primate Brain," by Le Chang and Doris Y. Tsao, in *Cell*, Vol. 169, No. 6, June 1, 2017 (*face grid*)

Pictures Worth 205 Neurons

For a given face, we can predict how a cell will respond by taking a weighted sum of all 50 face coordinates. To predict what face the monkey saw from neuronal activity, this entire process can be reversed: by knowing the response of 205 face cells, it is possible to predict the 50 coordinates defining the exact facial features—and make a highly accurate reconstruction of a given face.

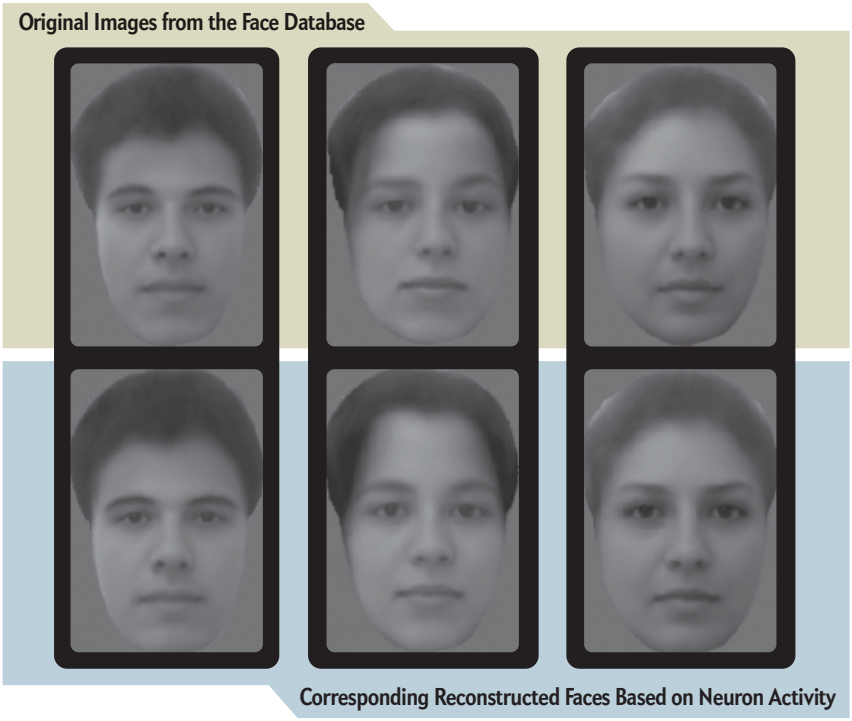

the shape of the face. From these varied shapes, we calculated an average face. We then morphed each face image in the database so its key features exactly matched those of the average face. The resulting images constituted the appearance of the faces independent of shape.

We then performed principal components analysis independently on the shape and appearance descriptors across the entire set of faces. This is a mathematical technique that finds the dimensions that vary the most in a complex data set.

By taking the top 25 principal components for shape and the top 25 for appearance, we created a 50-dimensional face space. This space is similar to our familiar 3-D space, but each point represents a face rather than a spatial location, and it comprises much more than just three dimensions. For 3-D space, any point can be described by three coordinates (x,y,z). For a 50-D face space, any point can be described by 50 coordinates.

In our experiment, we randomly drew 2,000 faces and presented them to a monkey while recording cells from two face patches. We found that almost every cell showed graded responses—resembling a ramp slanting up or down—to a subset of the 50 features, consistent with my earlier experiments with cartoon faces. But we had a new insight about why this is important. If a face cell has ramp-shaped tuning to different features, its response can be roughly approximated by a simple weighted sum of the facial features, with weights determined by the slopes of the ramp-shaped tuning functions. In other words:

response of face cells = weight matrix × 50 face features

We can then simply invert this equation to convert it to a form that lets us *predict* the face being shown from face cell responses:

50 face features = (1/weight matrix) × response of face cells

At first, this equation seemed impossibly simple to us. To test it, we used responses to all but one of the 2,000 faces to learn the weight matrix and then tried to predict the 50 face features of the excluded face. Astonishingly, the prediction turned out to be almost indistinguishable from the actual face.

A WIN-WIN BET

AT A MEETING in Ascona, Switzerland, I presented our findings on face reconstruction using neural activity. After my talk, Rodrigo Quian Quiroga of the University of Leicester in England, who with his colleagues discovered the so-called Jennifer Aniston cell in the human medial temporal lobe in 2005, asked me how my cells related to his concept that single neurons react to the faces of specific people. The Jennifer Aniston cell, also known as a grandmother cell, is a putative type of neuron that switches on in response to the face of a recognizable person—a celebrity or a close relative.

I told Rodrigo I thought our cells could be the building blocks for his cells, without thinking very deeply about how this would work. That night, sleepless from jet lag, I recognized a major difference between our face cells and his. I had described in my talk how our face cells computed their response to weighted sums of different face features. In the middle of the night, I realized this computation is the same as a mathematical operation known as the dot product, whose geometric representation is the projection of a vector onto an axis (like the sun projecting the shadow of a flagpole onto the ground).

Remembering my high school linear algebra, I realized this implied that we should be able to construct a large "null space" of faces for each cell—a series of faces of varying identity that lie on an axis perpendicular to the axis of projection. Moreover, all these faces

Doris Y. Tsao (*face images*)

would cause the cell to fire in exactly the same way.

And this, in turn, would suggest cells in face patches are fundamentally different from grandmother cells. It would demolish the vague intuition everyone shared about face cells—that they should be tuned to specific faces.

I was the first person in the meeting's breakfast hall at 5 A.M. the next morning and hoped to find Rodrigo so I could tell him about this counterintuitive prediction. Amazingly, when he finally showed up, he told me he had the exact same idea. So we made a bet, and Rodrigo allowed the terms to be framed in a way that would be win-win for me. If each cell really turned out to have the same response to different faces, then I would send Rodrigo an expensive bottle of wine. If, on the other hand, the prediction did not pan out, he would send me solace wine.

In search of an answer back in our lab at Caltech, Le Chang first mapped the preferred axis for a given cell using responses to the 2,000 faces. Then he generated, while still recording from the same cell, a range of faces that could all be placed on an axis perpendicular to the cell's preferred axis. Remarkably, all these faces elicited exactly the same response in the cell. The next week Rodrigo received an exquisite bottle of Cabernet.

The finding proved that face cells are not encoding the identities of specific individuals in the IT cortex. Instead they are performing an axis projection, a much more abstract computation.

An analogy can be made to color. Colors can be coded by specific names, such as periwinkle, celandine and azure. Alternatively, one can code colors by particular combinations of three simple numbers that represent the amount of red, green and blue that make up that color. In the latter scheme, a color cell performing a projection onto the red axis would fire electrical impulses, or spikes, proportional to the amount of red in any color. Such a cell would fire at the same intensity for a brown or yellow color containing the same amount of red mixed in with other colors. Face cells use the same scheme, but instead of just three axes, there are 50. And instead of each axis coding the amount of red, green or blue, each axis codes the amount of deviation of the shape or appearance of any given face from an average face.

It would seem then that the Jennifer Aniston cells do not exist, at least not in the IT cortex. But single neurons responding selectively to specific familiar individuals may still be at work in a part of the brain that processes the output of face cells. Memory storage regions—the hippocampus and surrounding areas—may contain cells that help a person recognize someone from past experience, akin to the grandmother cells.

Facial recognition in the IT cortex thus rests on a set of about 50 numbers in total that represent the measurement of a face along a set of axes. And the discovery of this extremely simple code for face identity has major implications for our understanding of visual object representation. It is possible that all of the IT cortex might be organized along the same principles governing the face-patch system, with clusters of neurons encoding different sets of axes to represent an object. We are now conducting experiments to test this idea.

NEURAL ROSETTA STONE

IF YOU EVER GO to the British Museum, you will see an amazing artifact, the Rosetta stone, on which the same decree of Memphis is engraved in three different languages: Egyptian hieroglyphics, Demotic and ancient Greek. Because philologists knew ancient Greek, they could use the Rosetta stone to help decipher Egyptian hieroglyphics and Demotic. Similarly, faces, face patches and the IT cortex form a neural Rosetta stone—one that is still being deciphered. By showing pictures of faces to monkeys, we discovered face patches and learned how cells within these patches detect and identify faces. In turn, understanding coding principles in the face-patch network may one day lead to insight into the organization of the entire IT cortex, revealing the secret to how object identity more generally is encoded. Perhaps the IT cortex contains additional networks specialized for processing other types of objects—a whirring factory with multiple assembly lines.

UNDERSTANDING CODING PRINCIPLES IN THE FACE-PATCH NETWORK MAY ONE DAY LEAD TO INSIGHT INTO THE ORGANIZATION OF THE ENTIRE INFERO-TEMPORAL CORTEX, REVEALING THE SECRET TO HOW OBJECT IDENTITY MORE GENERALLY IS ENCODED.

I also hope that knowing the code for facial identity can help fulfill my college dream of discovering how we imagine curves. Now that we understand face patches, we can begin to train animals to imagine faces and explore how neural activity is shaped by the purely internal act of imagination. Lots of new questions arise. Does imagination reactivate the code for the imagined face in the face patches? Does it bring back even earlier representations of contours and shading that provide inputs to the face-patch system? We now have the tools to probe these questions and better understand how the brain sees objects, imagined or real.

Because almost all the brain's core behaviors—consciousness, visual memory, decision-making, language—require object interactions, a deep understanding of object perception will help us gain insight into the entire brain, not just the visual cortex. We are only starting to solve the enigma of the face. SA

Doris Y. Tsao is a professor of biology at the University of California, Berkeley's Helen Wills Neuroscience Institute and an investigator of the Howard Hughes Medical Institute. In 2018 she was named a MacArthur Fellow.

THE BRAIN'S SOCIAL ROAD MAPS

Neural circuits that track our whereabouts in space and time may also play vital roles in determining how we relate to other people

By Matthew Schafer and Daniela Schiller

Illustration by Richard Borge

WE ARE OFTEN TOLD THAT THERE ARE NO SHORTCUTS IN LIFE. But the brain—even the brain of a rat—is wired in a way that completely ignores this kind of advice. The organ, in fact, epitomizes a shortcut-finding machine.

The first indication that the brain has a knack for finding alternative routes was described in 1948 by Edward Tolman of the University of California, Berkeley. Tolman performed a curious experiment in which a hungry rat ran across an unpainted circular table into a dark, narrow corridor. The rat turned left, then right, and then took another right and scurried to the far end of a well-lit narrow strip, where, finally, a cup of food awaited it. There were no choices to be made. The rat had to follow the one available winding path, and so it did, time and time again, for four days.

On the fifth day, as the rat once again ran straight across the table into the corridor, it hit a wall—the path was blocked. The animal went back to the table and started looking for alternatives. Overnight, the circular table had turned into a sunburst arena. Instead of one track, there were now 18 radial paths to explore, all branching off from the sides of the table. After venturing out a few inches on a few different paths, the rat finally chose to run all the way down path number six, the one leading directly to the food.

Taking the path straight to the food cup without prior experience may seem trivial, but from the perspective of behavioral psychologists at the time, the rat's navigational accomplishment was a remarkable feat. The main school of animal learning in that era believed that maze behavior in a rat is a matter of simple stimulus-response associations. When stimuli in the environment reliably produce a successful response, neural connections that represent this association get strengthened.

In this view, the brain operates like a telephone switchboard that maintains only reliable connections between incoming calls from our sense organs and outgoing messages to the muscles. But the behavioral switchboard was unable to explain the ability to correctly choose a shortcut right off the bat without having first experienced that specific path. Shortcuts and many other intriguing observations along these lines lent support to a rival school of thought promulgated by theorists who believe that in the course of learning, a map gets established in a rat's brain. Tolman—a proponent of that school—coined the term: the cognitive map.

According to Tolman, the brain does more than just learn the direct associations among stimuli. Indeed, such associations are often brittle, rendered outdated by changes in the environment. As psychologists have learned in the decades since Tolman's work, the brain also builds, stores and uses mental maps. These models of the world enable us to navigate our surroundings, despite complex, changing environments—affording the flexibility to use shortcuts or detours as needed. The hungry rat in Tolman's experiment must have remembered the location of the food, inferred the angle to it and chosen the route most likely to bring it to its goal. Quite simply, it must have built a model of the environment.

Such model building or mapmaking extends to more than physical space. Mental maps may exist at the core of many of our most "human" capacities, including memory, imagination, inferences, abstract reasoning and even the dynamics of social interactions. Researchers have begun to explore whether mental maps document how close or distant one individual is to another and where that individual resides in a group's social hierarchy. How does the brain, in fact, create the maps that allow us to make our way about the world?

A SPATIAL MAP

THE FIRST HINTS of a neural basis for mental maps came in the 1970s. While studying a brain region called the hippocampus in rodents, John O'Keefe of University College London, along with his student Jonathan Dostrovsky, discovered a particular class of neurons that becomes active when mice occupy specific locations in their environment. Some of these neurons fired when the animal was in one location, and others switched on when it moved to the next spot on the path along which it traveled, as if the cells were specialized to track *where* the animal was in space. By linking sequences of these "place cells" together, researchers were able to reconstruct an animal's navigational trajectory. Work over the intervening decades confirmed the existence of place cells in other animals, including humans, and clarified many of their properties. Along the way, a host of cell types surfaced, each uniquely contributing to the brain's encoding of spatial representations.

In the nearby entorhinal cortex, a region connected to the hippocampus, a research team led by Edvard Moser and May-Britt Moser, former postdoctoral visiting fellows in O'Keefe's laboratory, discovered neurons highly similar to place cells. These cells also fired when an animal was in specific locations. But unlike place cells, each of these newly discovered cells spiked in multiple, regular locations. When mapped onto the animal's position, the activity patterns of these "grid cells" resembled highly regular, equilateral triangles. Like a spatial metric, these cells fired when an animal passed over the vertices of the triangles. The discovery of these cell types sparked excitement because of the emerging picture of how the brain controls navigation. Place cells

Giving Way to the Abstract

Maps simplify the world by reducing an overwhelming amount of sensory and cognitive data into a format that can be used for navigating physical space, pointing to shortcuts and detours to reach a destination faster. The organization of such maps—built on the activity of cells dedicated to tracking both space and time—scales in the abstraction of what they represent: from the recognition of another individual along the way to even a complex space that denotes social power and closeness to others.

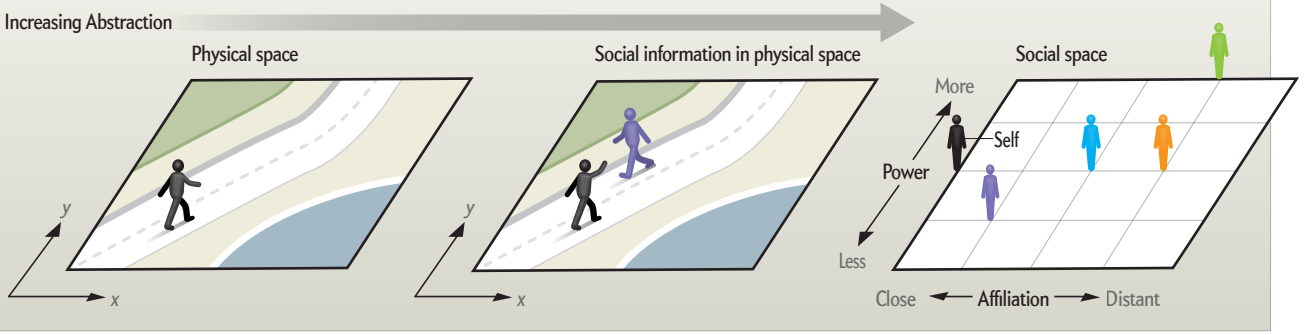

and grid cells could provide a means to locate oneself in space and determine distance and direction. These navigational tools are crucial for building mental maps. (O'Keefe and the Mosers received the 2014 Nobel Prize in Medicine or Physiology for their work on place and grid cells.)

A wide variety of information is useful for creating such a map, and the hippocampus-entorhinal system encodes much of it. Discovering the location of a physical goal is one example: as an animal navigates toward an objective, some hippocampal neurons fire depending on the direction and distance to reach it. The cells increase their firing rate as the animal approaches the goal.

Other cells also enter the picture. A dedicated population of "reward" cells encodes reward locations across different environments, providing a signal to guide an animal's navigation (think of an "X" marking the spot of treasure on a pirate's map). Other cells track speed and direction and in doing so act like internal speedometers and compasses that compute an animal's progress as it travels through the environment. Specific cells that signal the locations of landmarks in the surroundings serve as references to correct errors in the animal's trajectory. A map must also have edges: cells that fire more as the animal approaches the map's perimeter.

For humans, the importance of such an abundance of cell types seems obvious: the brain is responsible for knowing the location of home and work, walls and dead ends, a favorite shop or the corner store. It is still a mystery as to how all of this information is drawn together into a coherent map, but these cells appear to provide the parts list for the elements of neural mapmaking.

This hippocampal-entorhinal system is more than a mapmaker, though, and the maps are more than a way to locate oneself in space. These maps also are used for active planning. When a rat comes to a junction in a familiar maze, it will pause while place cell firing sequences that relate to the different options are activated, as if the animal is contemplating the choices.

Humans engage similar processes. Research in participants navigating virtual environments while their brains were scanned with functional magnetic resonance imaging shows that the hippocampus becomes active in ways consistent with spatial planning, such as considering and planning routes.

Shaping of plans also occurs during sleep. Sequences of place cell activity can be reactivated during sleep to replay the past or simulate the future. Without the ability to simulate new behaviors, we would have to explore a multitude of real-world options before deciding on what action to take. We would be constant empiricists, able to act only on direct observations. Instead offline simulations give us the ability to envision possibilities without directly experiencing them.

MENTAL TIME TRAVEL

TIME AND SPACE are inextricably linked. It is difficult to talk about time without borrowing a spatial metaphor: time "passes" as we "move" through it. We look "forward" to the future and "back" on our past. The same hippocampal-entorhinal system tracks movement through time. Work done largely in the lab of the late Howard Eichenbaum of Boston University revealed neurons in the hippocampal-entorhinal system that encode the time course of an animal's experience. Time cells fire at successive moments but do not track time in a simple clocklike fashion. Instead they mark temporal context—stretching or shrinking their firing durations if the length of a task changes, for example. Some time cells encode space as well. In the brain, in fact, physical and temporal space may be bound together.

The discovery of the crucial importance of these brain areas in space and time was not totally surprising. Psychologists had long suspected it to be the case. In 1953 Henry Molaison underwent bilateral hippocampal resection surgery to reduce extreme, life-disrupting epileptic seizures. The surgery was successful at quelling the seizures. But Molaison—known for decades only as H.M.—became one of the most renowned cases in the history of the brain sciences.

Molaison could remember most experiences from before his surgery—people he knew and recollections from culture and politics. But his ability to form such explicit memories postsurgery was practically nonexistent. Even so, certain types of learning and memory remained untouched: he could still learn some new skills with enough practice. But his recollections of new people, facts and events were immediately lost.

Source: "Navigating Social Space," by Matthew Schafer and Daniela Schiller, in *Neuron*, Vol. 100; October 24, 2018

Graphics by Jen Christiansen

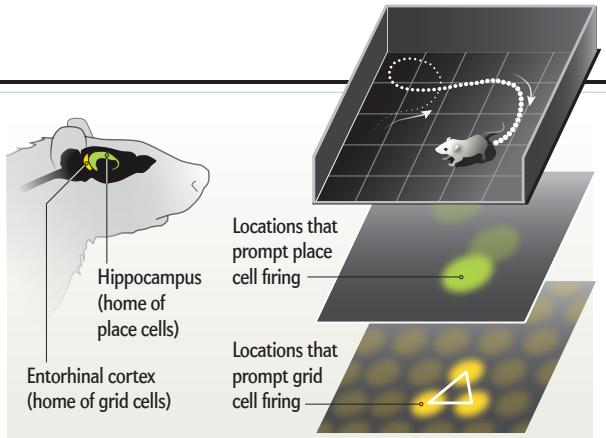

Cognitive Cartography Is Physical *and* Social

The brain forms the idea of friend or foe by stitching together diverse social characteristics from memories that track one's whereabouts. The recollections, research suggests, can then be used to place an individual within a social hierarchy that elucidates, say, where one stands in relation to others.

PLACE AND GRID CELLS

Place cells pinpoint the animal's whereabouts, each cell firing when a particular spot on a mental map is reached. A grid cell activates when an animal passes over the vertices of triangles superimposed on a mental map. The triangles' pattern of activation helps the animal compute the direction and distance traveled along a route.

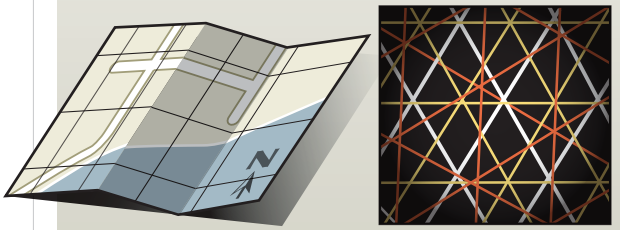

MAKING THE LEAP TO SOCIAL MAPS

Go right at the corner and continue to your destination. Building a map of physical surroundings is the work of place and grid cells. But the brain may also use these cells for constructing maps for social milieus: locating an acquaintance who grows closer but loses power in a relationship.

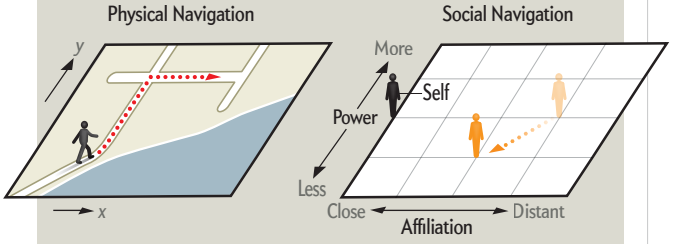

From observing Molaison, neuroscientists discerned that the hippocampus was essential in forming the episodic memories that record facts and events. Research on the role of the hippocampus in episodic memory exploded, largely in parallel to studies on its maplike functions.

The discoveries about the roles of the hippocampus and entorhinal cortex in spatial navigation and episodic memory were significant for at least a couple of reasons. The work in spatial navigation in rodents marked the first time that a higher-order cognitive function—something beyond basic sensory processes—mapped onto clear neural correlates. H.M. showed us that there were multiple types of memory supported by at least partially different neural systems, with the hippocampus playing a central role in the formation and storage of new episodic memories. These discoveries hinted that mechanisms of spatial and temporal navigation might underlie episodic memory. This synthesis is perhaps best explained by the theoretical construct proposed decades earlier by Tolman; both episodic memory and spatial navigation might reflect the brain's formation and use of cognitive maps.

Maps are not accurate portraits of the world in all of its complexity. Rather they are representations of relations—distances and directions between locations and what exists where. Maps reduce a dizzying amount of real-world information into a simple, easily readable format that is useful for effective, flexible navigation. The cell types mentioned earlier (place cells, grid cells and border cells, among others) may piece together such related elements into a mental map, which other brain regions can then read out to guide "navigation," amounting to adaptive decision-making. Mapping allows relations to be inferred, even when they have not been experienced. It also allows for mental shortcuts that go beyond the purview of the spatial and temporal domains. In fact, reasoning using abstract concepts may depend on some of these same neural foundations.

In one example of this line of work, researchers Alexandra Constantinescu, Jill O'Reilly and Timothy Behrens, all then at the University of Oxford, asked participants to learn associations of different symbols with images of "stick" birds with various neck and leg lengths. A bird with a long neck but short legs, for example, might be linked with the image of a bell, whereas a bird with a short neck and long legs might be connected to a teddy bear. These linkages created a two-dimensional association space. Despite neuroimaging being too crude to detect actual grid cells in the human brain, imaging conducted during the learned-association testing nonetheless revealed a gridlike pattern of activation within the entorhinal cortex.

This finding builds on earlier work by Christian Doeller of the Max Planck Institute for Human Cognitive and Brain Science in Leipzig, Germany, and Neil Burgess of University College London that first showed an entorhinal gridlike representation in humans navigating a virtual maze. For both physical and abstract relations, the gridlike organization is highly efficient. It makes the linkages of places or concepts more predictable, enhancing how quickly inferences can be made about these relations. As in physical space, this organization of information allows for the inference of shortcuts—relations between ideas or perhaps analogies, stereotypes and even some aspects of creativity itself could depend on such inferences.

Sources: "Scientific Background: The Brain's Navigational Place and Grid Cell System," by Ole Kiehn and Hans Forssberg, with illustrations by Mattias Karlen. Nobelprize.org; "Navigating Social Space," by Matthew Schafer and Daniela Schiller, in *Neuron*, Vol. 100; October 24, 2018

PEOPLE MAPS

THE PROGRESSION from the physical to the abstract carries over into the way the brain represents social relationships. Various bits of knowledge about another person are distilled into the concept of that individual. When we see a photograph of someone or hear or see that person's name, the same hippocampal cells will fire, regardless of the sensory details of the stimulus (for example, the famous "Jennifer Aniston neuron" described by Itzhak Fried of the University of California, Los Angeles, and his colleagues). These hippocampal cells are responsible for representing concepts of specific individuals.

Other hippocampal cells track the physical locations of others and are called social place cells. In an experiment by David Omer of the Hebrew University of Jerusalem, Nachum Ulanovsky of the Weizmann Institute of Science in Rehovot, Israel, and their colleagues, bats observed other bats navigating a simple maze to reach a reward. The task of an observer bat was to simply watch and learn from a navigating bat, enabling it to subsequently navigate the same route to get the same reward. When the observer bat watched, hippocampal cells fired corresponding to the location of the other bat.

Neural circuitry within specific subregions of the hippocampus (in particular, areas called CA1 and CA2) contributes to such social memories. Artificial stimulation or inactivation of these hippocampal areas enhances or diminishes an animal's ability to recognize other animals. In humans, hippocampal injury often spares memory for specific, individual faces, but the relation between this cardinal identifier of another person and that individual's behavior may be lost. That observation suggests that the hippocampus does not simply record a face or some other personal detail but rather ties together diverse social characteristics.

Hippocampal activity also tracks social hierarchies: the demands of a boss and a co-worker, for instance, may be valued differently and confer different social standings. Common metaphors illustrate the spatial dimensions of a hierarchy: a person may try to gain status to "climb the social ladder" or "look down" at someone below them. Other factors are also critical. Biological relatedness, common group goals, the remembered history of favors and slights—all determine social proximity or distance. Human relationships can be conceived of as geometric coordinates in social space that are defined by the dimensions of hierarchy and affiliation.

Work in our lab has explored these ideas in recent years. Our results suggest that, as with other spaces, the hippocampus organizes social information into a maplike format. To test this hypothesis, we put individuals in a choose-your-own-adventure game in which they interacted with cartoon characters and made decisions while their brains were scanned.

In the game, players had just moved to a new town and needed to interact with the fictional characters to secure a job and a place to stay. Participants made decisions on how to deal with a given character. Players could request that others perform favors to demonstrate their power, or they could submit to demands made on them. In a subsequent interaction, they could decide whether to make a gesture of attachment—giving a hug or remaining at a distance.

Using these decisions, we plotted each character at certain coordinates on a map representing their movement along the dimensions of power and affiliation. In each interaction, we drew a line or vector from the participant to the character. In this way, we charted the evolving relations as trajectories through social space and computed information about the angles and lengths of the social vectors.

We searched for neural signals that tracked this information by correlating a participant's brain activity with the angle and length of the vectors for each decision. Activity in the hippocampus tracked the angle of the characters to the participant. The degree to which hippocampal activity captured these social coordinates also reflected the participants' self-reported social skills. These findings suggest that the hippocampus monitors social dynamics as it does physical locations by encoding relations between points in multidimensional space. Indeed, it may be that along any arbitrary dimension in which we can order information, whether physical or abstract, the hippocampus-entorhinal system plays a part.

Many questions about the brain's social maps still remain unanswered. How does this system interact with other brain regions? For example, in our role-playing study, we found that the posterior cingulate cortex, a region also involved in representing spatial information, tracked the length of social vectors—functioning in effect as a measuring stick of "social distance." Further, a gridlike signal was found in brain regions that are interconnected with and tend to co-activate with the hippocampal-entorhinal system, suggesting they form a network of brain regions with common functional properties.

As research accumulates, questions of clinical importance arise as well. Can flawed mapping processes explain psychiatric dysfunction? Another possibility is that insights garnered from this brain architecture could inform artificial-intelligence development. Well-organized internal models of the world might be key to building more intelligent machines.

That the same mapping system may underlie navigation through space and time, reasoning, memory, imagination and even social dynamics suggests that our ability to construct models of the world might be what makes us such adaptive learners. The world is full of both physical and abstract relations. Road maps of city streets and mental maps of interrelated concepts help us make sense of the world by extracting, organizing and storing related information. A new coffee shop on a familiar street can be easily placed within an existing spatial map. Fresh concepts can be related to older ideas. And a new acquaintance can reshape our social space.

Maps let us simulate possibilities and make predictions, all within the safety of our own heads. The mental shortcuts we can so readily conjure up might have their basis in the same system that allows us to figure out a detour around a traffic jam. We have just begun to discover the varied properties and capacities of this system. Mental maps do more than help us find shortcuts through physical space—they enable us to navigate life itself. SA

Matthew Schafer is pursuing a doctorate in neuroscience at the Icahn School of Medicine at Mount Sinai, focusing on the neural mechanisms of social cognition in the human brain.

Daniela Schiller is a professor of both neuroscience and psychiatry at the Icahn School of Medicine at Mount Sinai. She researches the neural mechanisms underlying emotional control needed to adapt to constantly changing environments.

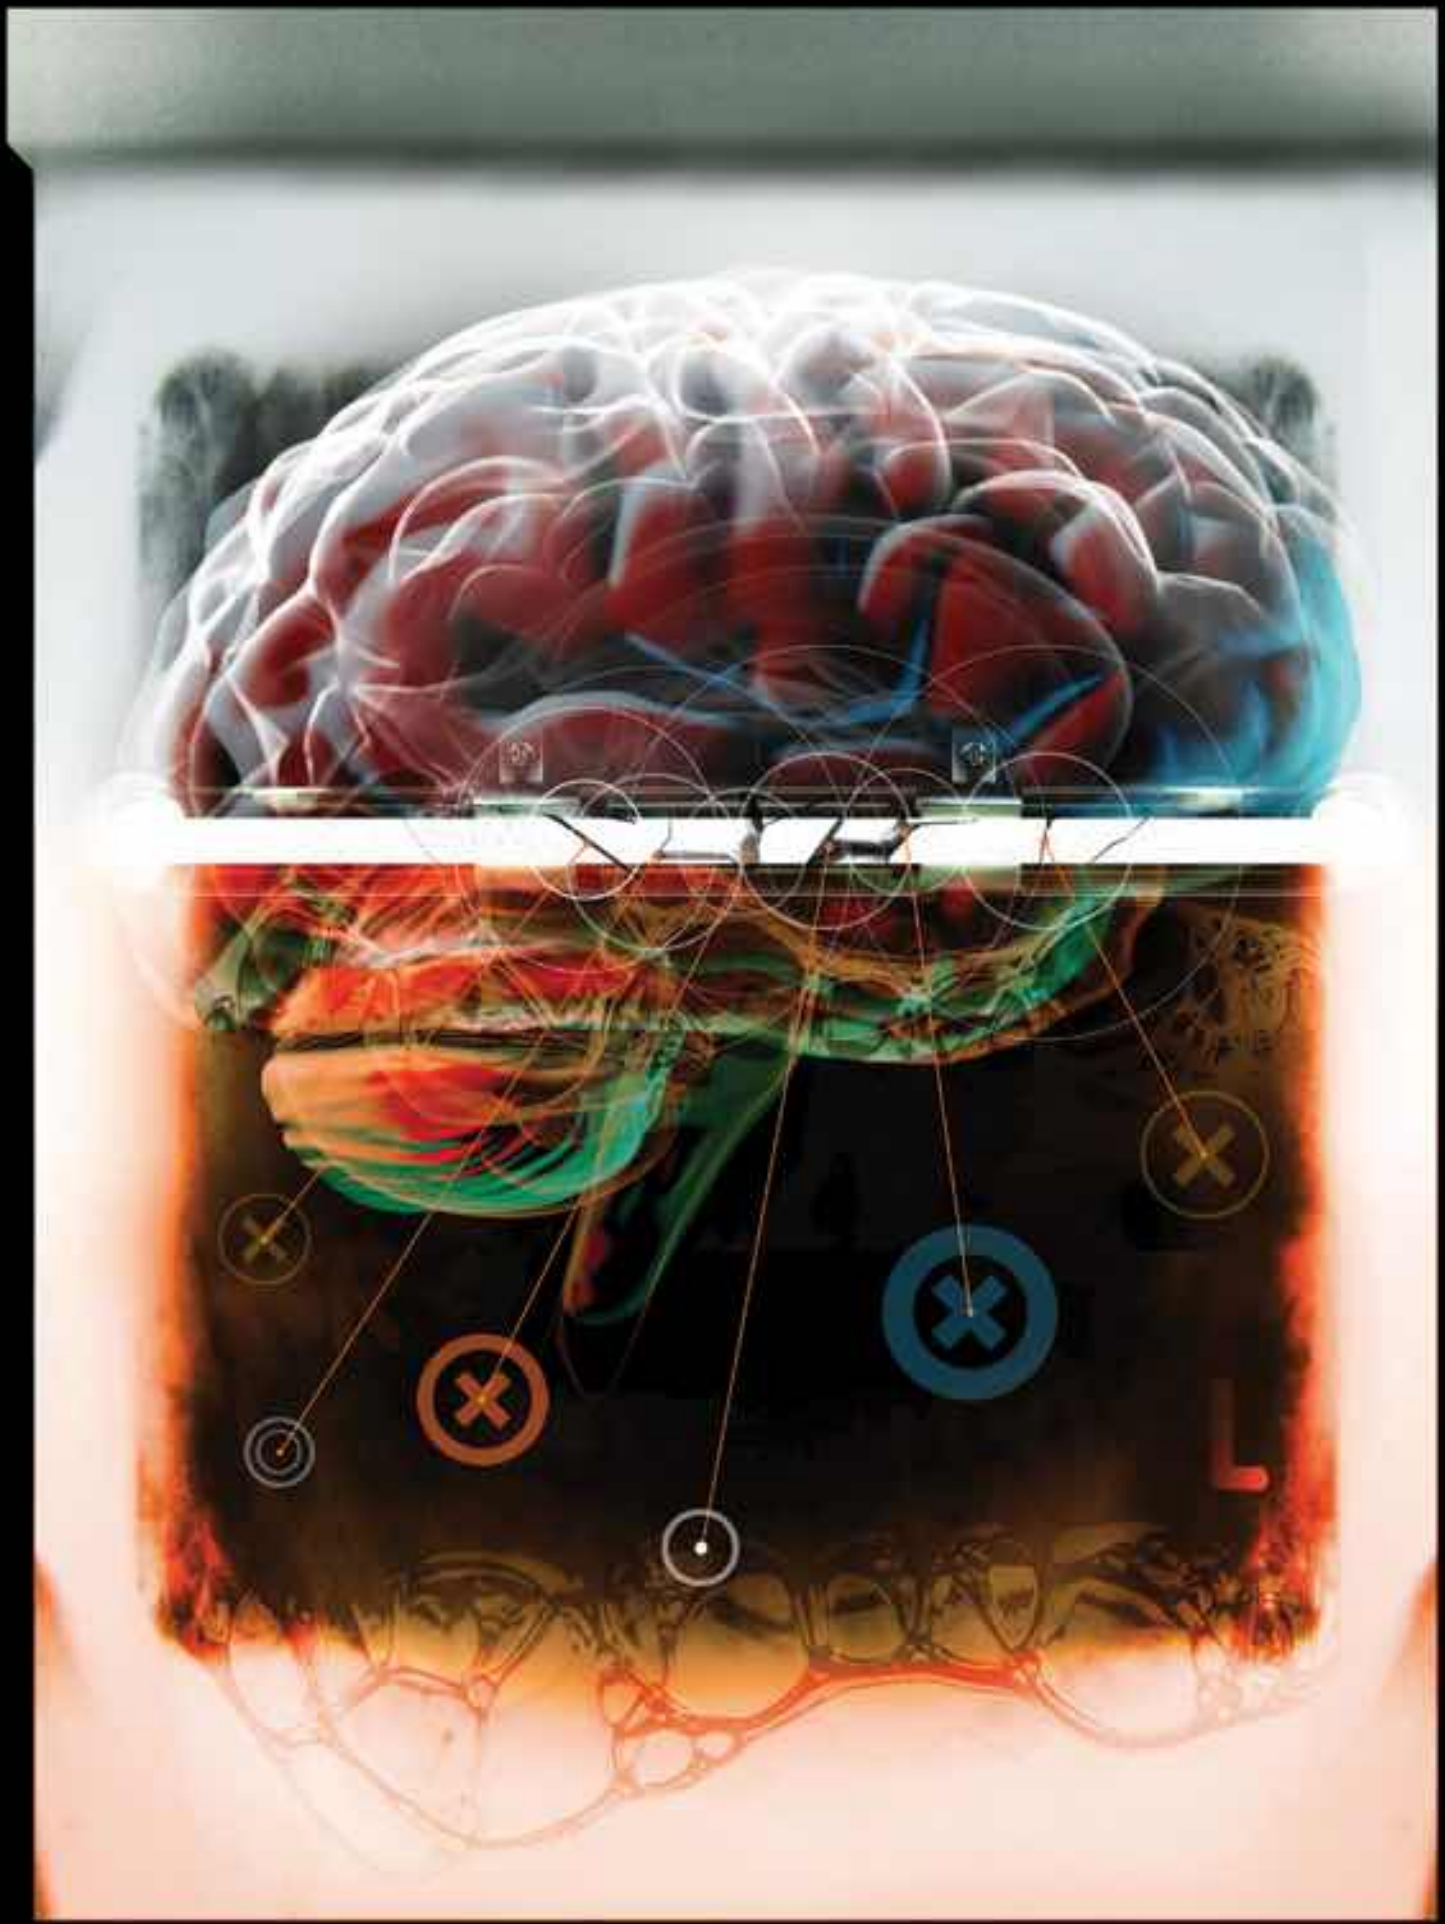

HOLES IN THE SHIELD

Leaks in a protective filter called the blood-brain barrier may lead to Alzheimer's and other dementias. Reversing the effects makes aging animal brains look young and healthy

By Daniela Kaufer and Alon Friedman

Illustration by Viktor Koen

IT WAS THE MIDDLE OF THE NIGHT IN JERUSALEM, AND WE WERE WATCHING MICE SWIM. The year was 1994, and the two of us were crouching over a pool of cold water in a laboratory at the Hebrew University. The room was chilly, our hunched backs ached, and we had been repeating this routine over many nights, so we were tired and uncomfortable. So were the mice. Mice really dislike swimming, especially in cold water—but we wanted to stress them out.

We humans were on the night shift because both of us had other things to do during the day. Kaufer was working on a doctorate in molecular neurobiology, and Friedman was an Israel Defense Forces physician and was often on call. What brought us together with the mice every evening was an attempt to understand a medical mystery: Gulf War syndrome. After the conflict ended in 1991, there were an increasing number of reports of soldiers from the U.S.-led coalition who were afflicted with chronic fatigue, muscle pain, sleep problems and cognitive deterioration, and those soldiers were hospitalized at higher rates than nondeployed veterans. Some doctors suspected that pyridostigmine, a drug that had been given to soldiers to protect them from chemical weapons, could cause these ailments if it made it into their brains.

There was a big problem with this theory, however: pyridostigmine in the bloodstream was not supposed to reach the brain. Blood vessels that course through this vital organ have walls made of specialized cells, packed very closely and with abilities to selectively control what can get into the brain. They form a shield that keeps toxins, pathogens such as bacteria and viruses, and most drugs safely within the vessels. This structure is called the blood-brain barrier, or BBB for short, and the drug should not have been able to pass through it.

That is, if the barrier was intact. We were testing the hypothesis that the physical and mental stress of combat might somehow trigger leaks in the shield. The swimming mice were our way of testing whether stress led to damage. When the swim session was done, we pulled each mouse from the pool and injected a drop of blue dye into one of its veins. Then we waited as the dye passed through its body, gradually turning the mouse blue. If the BBB was intact, the brain should have remained its normal pinkish-white color. We eu-

thanized the mice so we could take a look at their brains under a dissecting microscope. Over several nights we had tried various lengths of swim time, but we did not see any brain changes.

But on this night, after two dips in slightly colder water, things looked different: the brains had a strong blue tint! Lab work is usually tedious, and success is often subtle, but this time we were jumping up and down and hugging each other, giddy with excitement. Our weird experiment had worked. Stressful situations could make the BBB spring leaks. With the help of our mentor, neuroscientist Hermona Soreq, we went on to show that the effect occurred with pyridostigmine and altered brain cell activity. We published these results in 1996 in *Nature Medicine* and in 1998 in *Nature*.

A quarter of a century later we can say that looking at these blue brains turned out to be a defining moment for both our careers, as well as the beginning of a lifelong friendship and scientific collaboration. Discovering the telltale blue tinge was the first step on a path that, over many years, led us to probe more and more deeply into the connection between other brain diseases and flaws in the organ's protective shell. Today pyridostigmine penetration is a leading hypothesis for the cause of Gulf War syndrome, although there are other candidates. And our investigations have linked BBB damage—caused by aging in addition to injury or acute stress—to several more familiar illnesses: Alzheimer's and related dementias, epilepsy and traumatic brain injury. In two papers published in 2019 in *Science Translational Medicine,* we demonstrated that as people get older, this shield loses integrity and starts leaking, allowing blood proteins into the brain that normally do not get there. These proteins in turn activate a cascade of events among brain cells that can produce some of the most notable and widely seen changes associated with aging and illness: inflammation, abnormal neuron activity and cognitive impairment.

The cause-and-effect connection looks especially strong because stopping the reactions set off by these leaks actually reverses signs of disease, at least in rodents. In older mice, we can abolish the inflammatory fog with a targeted drug that protects brain cells from being irritated by blood proteins or by making genetic modifications that prevent those cells from releasing inflammatory molecules. Within days of the treatment the aged brains of these mice started to function more like young brains. Abnormal electrical activity subsided. Markers of inflammation dropped to low levels. When placed in mazes, the animals made their way through as quickly and accurately as young mice did. We cannot try the same experimental brain modifications in humans; it is not ethical. But we have been able to use magnetic resonance imaging techniques, electroencephalography recordings and analysis of postmortem brain specimens to compare the brains of aged people and of individuals with Alzheimer's with those of young and healthy people. The images show excessive and progressive BBB disruption and leaking with aging and in those with the disease, as well as other features of the chemical cascade.

We do not know whether a damaged barrier is the only cause of Alzheimer's or other brain illnesses. It could play a contributing role along with other causes, including genetics and a variety of cellular problems that have been seen in aging brains. Or it could be collateral damage. And experiments in mice often do not pan out in people. But right now the long-standing dominant theory for Alzheimer's—that it is triggered by a buildup of a protein called beta-amyloid in the brain—is looking less convincing than ever. Many people have high levels of beta-amyloid in their brain but show no decline in mental function. Moreover, numerous experiments have reduced levels of this protein in the brain, yet the disease and associated mental decline remained unaffected. Drugs that target beta-amyloid have failed to help. Given that there are now more than 55 million people worldwide with dementia and almost 10 million new cases diagnosed every year, according to the World Health Organization, many scientists say it is high time to consider alternative explanations and therapeutic directions. If flaws in the brain's protective shield start a chain of events that leads to disease—a chain that experiments suggest can be blocked to restore brain health—it is a path of investigation worth pursuing.

GAPS IN THE WALL

WITH "BARRIER" IN ITS NAME, the BBB sounds like a wall around the brain, but it is really more like a distributed filter. Our body's control center gets 20 percent of the oxygen-rich blood pumped out by the heart, delivered by an intricate mesh of blood vessels. They look different from vessels in the rest of the body, with walls made of tightly packed cells with particular transport systems that form a semipermeable filter. Networks of brain cells need a carefully controlled environment to function, so this filter lets molecules such as oxygen and glucose diffuse through but blocks blood proteins, certain ions, immune system cells and pathogens. This protective mesh extends throughout most areas of the brain, from the outer layers of the cortex, where higher-order cognition occurs, to deep places such as the hippocampus, which regulates memory storage. Problems with the filter can therefore lead to all kinds of neurological difficulties.

Back in the 1990s, as we were completing our initial work on Gulf War syndrome, we knew that other researchers had noted BBB damage in patients with some brain disorders, including Alzheimer's. But we did not know whether this problem was a cause or an effect or how leaks in the shield get started and what they might do to alter brain function. We did, however, want to find out.

After our time working in Jerusalem, Kaufer went to Stanford University for her postdoctoral fellowship, and Friedman continued his medical training in Israel, specializing in neurosurgery. But time and distance did not let us forget. On a vacation together with our families, sailing between Greek islands, we caught up. Kaufer was learning more about how stress affects the brains of mice in work at Stanford. Friedman, in his own practice, was reaffirming the early observations from other researchers who saw flawed BBBs in many patients suffering from very different neurological conditions. Just what was the damaged barrier doing?

We began to figure out the answer to this question in the mid-2000s, when Friedman got the chance to work in Berlin with the late neuroscientist Uwe Heinemann of the Institute for Neurophysiology, part of the Charité Center for Basic Medicine. Heinemann opened his lab to the next key experiment. We wanted to observe brain function directly after the BBB started to malfunction, so we gave rats a chemical that essentially poked holes in the barrier and then dissected their brains at various later times. We kept the brain slices alive in nourishing fluid and used an electrode to record the electrical signals that the cells used to communicate with one another in the presence of a leaky barrier.

The first few days were boring. The neurons were giving off signals one after another in staccato, irregular patterns, "talking" as if nothing unusual had happened. We almost decided to give up. Then, on the fifth day, the cells' chatter patterns changed. More and more

neurons started to pulse together in synchrony. After a full week we nudged them with a small signal from an electrode, mimicking a brief electrical message within the cerebral cortex. This nudge produced a storm of cells firing together, similar to what is observed in people and animals with epilepsy.

We think what happened with these cells is analogous to generating a Twitter storm. Imagine that you created a Twitter account today and tweeted some sensational statement. You would probably get a very small response because you would not have many followers. If in the next few days you built a bigger network of followers and tweeted again, however, the same statement would be likely to be retweeted, recruiting more followers who would also retweet it, eventually leading to a storm of tweets on the social media platform. Similarly, when we disrupted the BBB, neurons in the brain were not discombobulated right away, but after they had spent a week building a new network of connections, a small jolt prompted a bigger electrical response. We now have recorded these electrical responses in the brains of people diagnosed with Alzheimer's and with epilepsy, and we termed those events "paroxysmal slow-wave events."

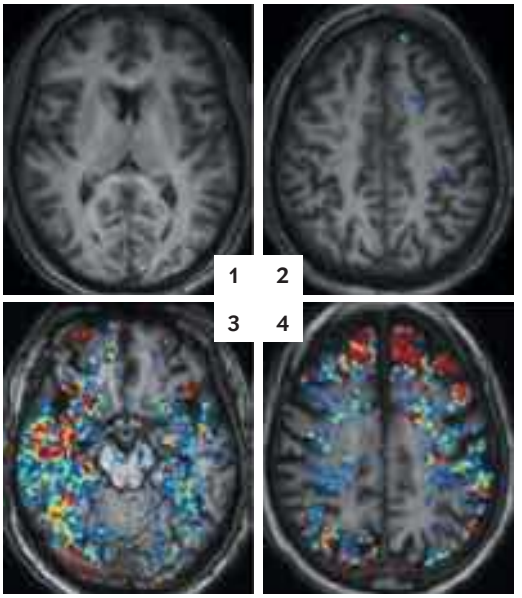

AGING BARRICADE: Brain scans, highlighting a colored tracer molecule in the blood, show more leaks in the blood-brain barrier as people age. A 30-year-old looks clear (*1*). At age 42, blue spots indicate small seeps (*2*). By age 65, red and yellow spots show bigger flows (*3*). At age 76, the pattern continues (*4*).

This storm happened only after we mimicked a BBB leak. Without one, our brain slices were untroubled by any electrical tempests. So we hypothesized that there was some element from the blood that was reaching these neurons to trigger the brain reaction. We tested this theory in a young, healthy rat with a normal BBB by injecting blood directly into its brain—bypassing the barrier—and monitoring electrical activity. It took several days, but again the storm built and exploded. Clearly, it had something to do with the blood. But blood is a complex fluid containing many different kinds of cells and proteins, so we set off on a painstaking filter-and-trap expedition. We used various filters to capture blood components of different sizes, then repeated the injection to see whether the storm recurred. It was a process of elimination. Eventually we found one blood protein that created the disturbances: albumin.

THE START OF TROUBLE

WE WERE NOT THRILLED with our catch. Albumin is very common and is involved in many bodily functions, so it was hard to isolate what it was doing in this situation. We would have preferred a rarer component. But albumin was what we got, so we dug in. Kaufer moved to the University of California, Berkeley, to run her own lab, and Friedman started his, first at Ben-Gurion University of the Negev in Israel and later at Dalhousie University in Nova Scotia. We planned a joint, long-distance series of experiments over several years to delineate the steps from BBB disruption and albumin leakage to the appearance of neurological disorders.

The first thing we learned was that when albumin gets into the brain, it appears to stimulate astrocytes, key brain cells that provide structural, functional and chemical support for neurons and their connections. When albumin contacts an astrocyte, it binds to receptors on the cell that under normal conditions bind a signaling molecule called transforming growth factor beta (TGFβ). Among other things, TGFβ provokes inflammation, starting a cascade of molecules called cytokines, and activates astrocytes and sentinel cells called microglia. Normally, this mechanism is the brain's way of limiting damage by destroying malfunctioning cells in a targeted assault. But when albumin seeps in, it leads to a positive-feedback loop of overproduction of TGFβ and other cytokines, setting in motion a chain of events. Recent evidence describes "senescent" cells—cells that remain chemically active after their cell-proliferation functions have stopped—that detect tissue damage and turn on a genetic program that amplifies the signal, broadcasting it to neighboring cells. We showed in a paper published this year in *Aging Cell* that albumin that seeps into the brain induces astrocytes to become senescent. Eventually lots of brain cells get damaged, key neural connections are modified and the function of these circuits deteriorates.

These findings tracked the steps from a leak in the BBB (induced by brain injury, for example) to neural dysfunction and the possible development of epilepsy. We became curious about whether this offered a possible explanation for why brain function deteriorates with age in some individuals, and indeed, subsequent experiments showed how this sequence of events plays out in aging mice. The animals typically live a bit more than two years on average. We allowed a colony of mice to age peacefully and looked inside their brains at various points. Albumin, we saw, was not in the brain at all in younger mice, but it began to show up in middle age. The effect was modest at first, but there was a clear decline in the integrity of the barrier, and it got worse in some mice as they got older. The affected mice also had more trouble learning and remembering their way through mazes than did their younger and relatively albumin-free counterparts.

When the albumin showed up, other experiments revealed, TGFβ receptors started to get active. We stained the mouse brains with antibodies that recognized an inflammatory protein by-product of TGFβ-receptor activation, and then we used green fluorescence to localize astrocytes that had albumin and TGFβ-receptor hyperactivity. The inflammatory signal always started after albumin appeared in astrocytes and increased with a greater degree of albumin leakage. We saw that albumin and the inflammation it caused were especially abundant in the hippocampus, a brain area that is a key component in memory regulation.

Within the past decade we have been able to provide good evidence that this same process happens in people. We used tracer molecules to tag signals of barrier leakage in people. With MRI, we could see how the brain changes its signal when the tracer shows up minutes after being injected. Similar to mice, some middle-aged

Alon Friedman and Daniela Kaufer

Mind the Gaps

In the brain, blood vessel walls are made of closely packed cells that form a protective filter called the blood-brain barrier (BBB). The filter allows crucial nutrients to get from the blood to brain cells, but it blocks pathogens and certain blood proteins that can harm delicate neurons. Research has shown that this filter leaks more as people age. In the brain's memory areas such as the hippocampus, the leaks start a cascade of damaging events that can lead to Alzheimer's and other dementias.

A Healthy Barrier

Young, disease-free brains have a well-functioning BBB that lets vital molecules, including oxygen and glucose, move through it. These substances nourish neurons as well as specialized cells called astrocytes that support neuron connections and help them transmit signals to one another.

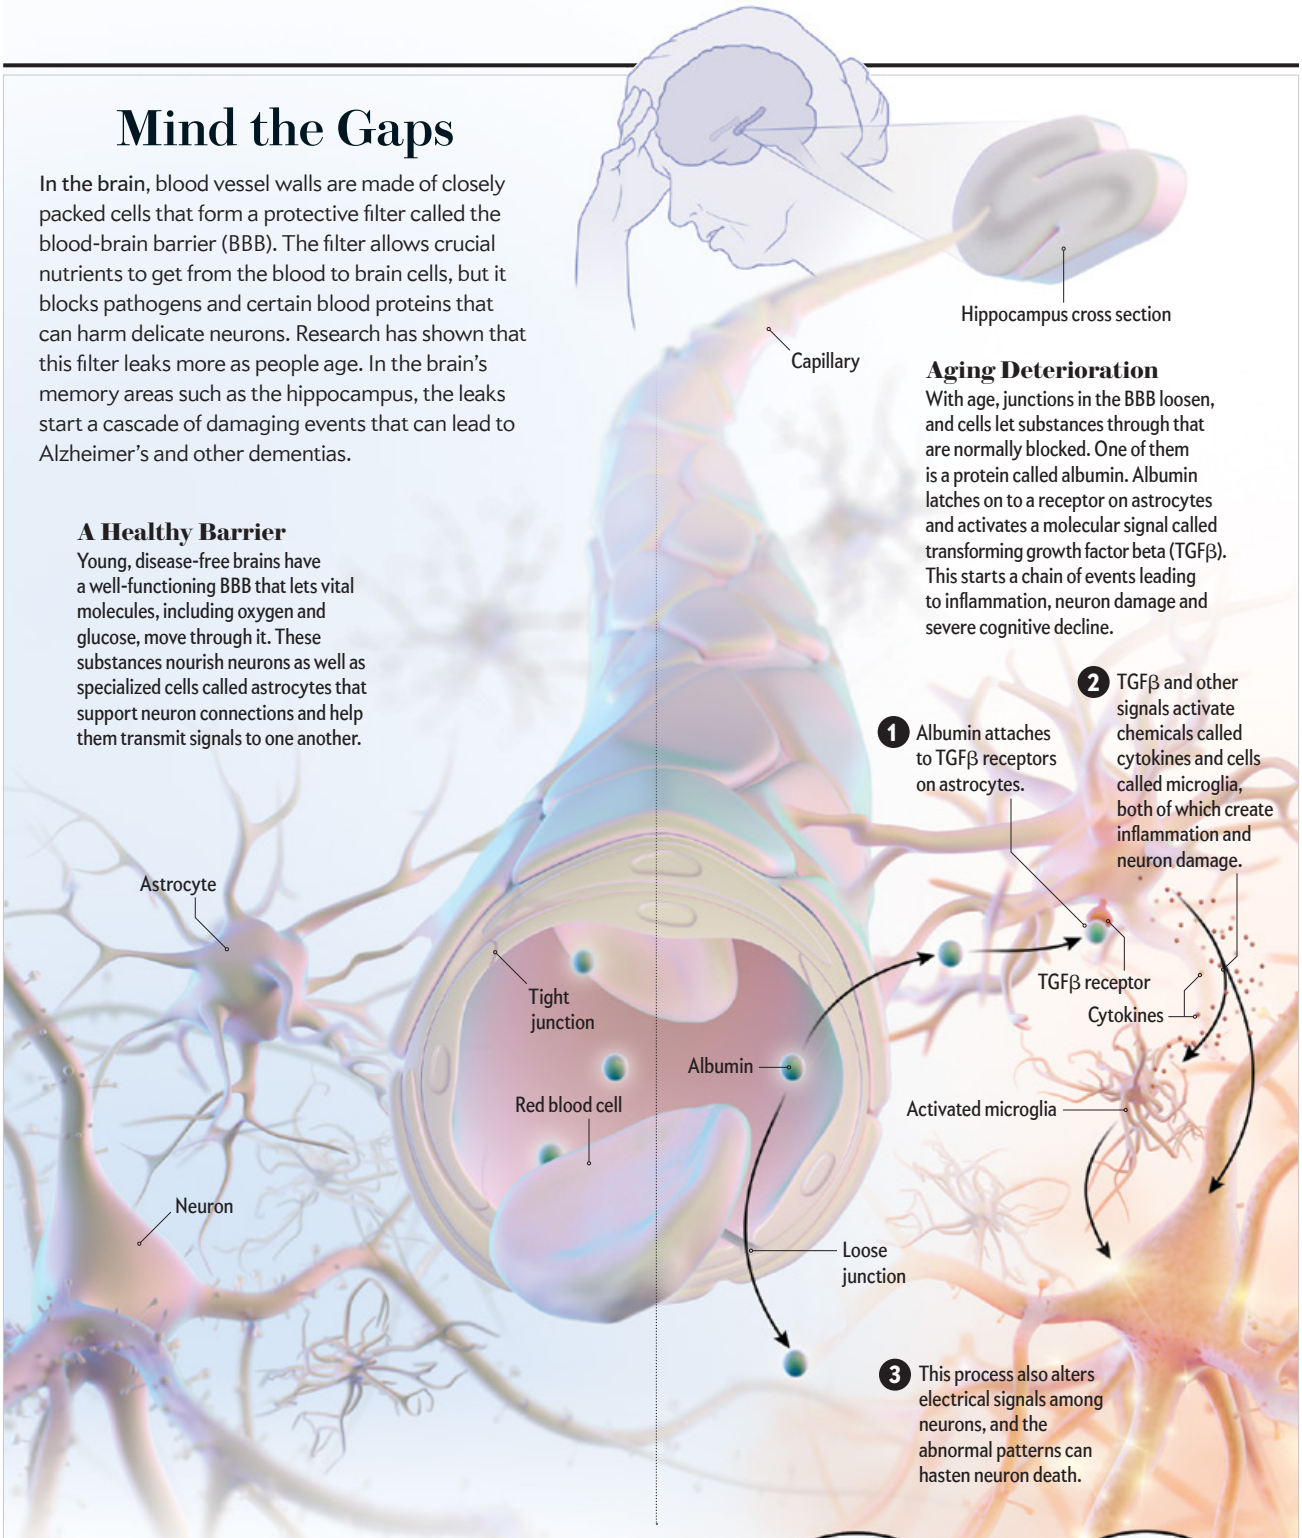

Aging Deterioration

With age, junctions in the BBB loosen, and cells let substances through that are normally blocked. One of them is a protein called albumin. Albumin latches on to a receptor on astrocytes and activates a molecular signal called transforming growth factor beta (TGFβ). This starts a chain of events leading to inflammation, neuron damage and severe cognitive decline.

1. Albumin attaches to TGFβ receptors on astrocytes.
2. TGFβ and other signals activate chemicals called cytokines and cells called microglia, both of which create inflammation and neuron damage.
3. This process also alters electrical signals among neurons, and the abnormal patterns can hasten neuron death.

Plugging the Leaks

Because albumin that gets through the barrier appears to start a sequence that ends in disease, scientists would like to stop this cascade. One approach would be to plug the initial breaches in the shield, but researchers have not yet demonstrated an effective way to do so **A**. Scientists have interrupted the sequence later on, when albumin contacts astrocytes. In mice, an experimental drug called IPW blocks TGFβ receptors on those cells, so albumin cannot latch on **B**. Researchers also have genetically altered animals so their astrocytes do not produce these receptors. Both approaches have made aged, damaged mouse brains look healthy again.

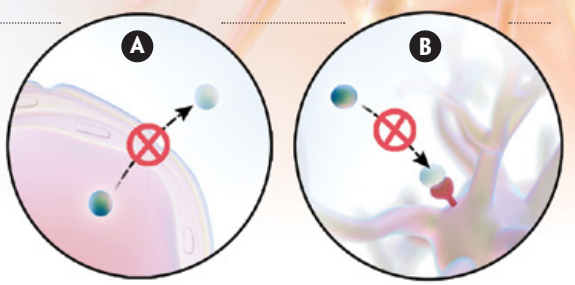

Graphic by Now Medical Studios

people already showed the tracer leaking into their brain at that point, and the older people were, the leakier their barrier became. Other researchers, such as Berislav V. Zlokovic of the University of Southern California Keck School of Medicine and his colleagues, used slightly different imaging methods to show similar age-related deterioration in barrier integrity in the hippocampi of living people with cognitive impairment. In our work, we added autopsies of a separate group of people and showed that heightened albumin levels accompanied greater amounts of TGFβ, always in astrocytes. These concentrations were higher in older people and in individuals who had died of Alzheimer's than in those without the disease.

BRAIN REJUVENATION

NOW WE KNEW that there was a correlation between BBB dysfunction in aging mice and the process in humans. But demonstrating correlation does not prove causation. First, we proved that this mechanism is sufficient to activate an aging program in the brain: When we infused albumin into the intact healthy brains of young mice for one week, these brains functioned like aged brains, with abnormal neuronal function, higher susceptibility to seizures and reduced learning capacity (the mice had a hard time learning an escape route in a water maze task). Furthermore, activation of this mechanism was essential for the aging response: we reversed the deterioration in mice by blocking the albumin-TGFβ cascade that came after the leaks. We developed a group of mice in which we genetically cut out the portion of DNA that tells astrocytes to produce TGFβ receptors, eliminating that feature from the cells. When the mice were still relatively young, we implanted a tiny pump in their brains that injected albumin. We did the same thing to a group of young, normal mice. Then we put both groups into a tricky water maze. (Watching mice swim seems to be a recurring theme with us.) The mice with receptors (and albumin in their brains) had a lot of trouble. But the animals without receptors swam the maze like young, healthy mice—speedily and accurately—and when we changed the maze configuration, they learned the new route, too. When we looked at their brains, we saw low levels of inflammation, senescent astrocytes and abnormal electrical activity.

This was really very encouraging. But for people, the option of knocking out a gene for a brain feature will not be an available therapy any day soon. There is, however, another form of treatment. Barry Hart, a medicinal chemist at Innovation Pathways, a start-up drug company in Palo Alto, Calif., had designed an anticancer drug that specifically blocked the activity of the TGFβ receptor. Such growth factors play a role in tumor progression, so blocking them could be therapeutic. Hart contacted us and suggested that we try the drug, called IPW, on our mice.

When we gave the drug to middle-aged mice—the ones that were starting to show albumin leakage—we learned that it made their brains look young again. TGFβ activity dropped to levels seen in youthful mice, markers of inflammation went way down, and abnormal electrical activity and seizure susceptibility diminished.

But the big surprise came when we tested actual behavior and cognition. We set up another maze, and this time we ran older mice through it. Some of the aged animals were treated with IPW, and some were not. We did not predict a lot of improvement, because we thought irreversible damage had already been done. (Our mice without the TGFβ gene had been spared the long months of deterioration inflicted by the inflammatory cascade, but these animals had not.) Within days, however, the treated mice were almost as good at learning the maze as rodents half their age. The untreated mice just shambled along as usual. Moreover, the mice that got IPW showed no sign of the "Twitter storm" effect that we typically see in humans with Alzheimer's or epilepsy and not much evidence of inflammation. It was as if an inflammatory fog had lifted, allowing the brain to regain its youthful abilities. These, along with the studies of human brains, are the results we published in two papers in 2019 in *Science Translational Medicine*.

The maze outcome was so unexpected, even to us, because, like most people, we had considered aging damage as a one-way trip—deterioration that cannot be undone. That is probably the case for major brain trouble, such as the havoc that occurs in individuals with Parkinson's disease or in advanced Alzheimer's once clumps of beta-amyloid have accumulated to such a degree that they kill off swaths of neurons and other cells. But this work may indicate that in the absence of a lot of cell death, the aging brain has a hidden capacity to rebound from some types of insults.

And our findings have implications for acute injuries as well, not just gradual deterioration. Treating rodents with IPW after concussions or traumatic brain injury with this drug alleviated the inflammation, seizures and cognitive decline that developed in the placebo-treated animals.

FIXING THE DAMAGE

THE WORLD POPULATION IS AGING, and the number of people with dementia and Alzheimer's is on the rise. Neuroscientists have a poor understanding of the early triggers of the transition from a young, healthy brain to an old, dysfunctional one. Alzheimer's and other neurological diseases of aging are complex. Defects in the way the brain disposes of aberrant proteins may play a role in how these illnesses get started, or the trigger could be the impairment of electrical signals among neurons, to name just two possibilities.

Now a leaky BBB has to be considered a strong contender as well, although it might not be the only cause or the only route to treatment. This barrier-breach theory provides a remarkably intuitive and straightforward new model to understand why the brain declines with age. And it is a model that gives us optimism: the results of our work strongly hint that the aging brain retains a capacity for reshaping and restoring itself, an ability that may be chronically suppressed, but not irreversibly lost, by persistent leakiness and the ensuing chain of events we have traced.

Our next step is to look for strategies and therapies to reduce barrier leakage. In the past, pharmaceutical research into the barrier focused on ways to increase permeability, not limit it, to get more drugs across it to treat brain tumors or infections. Our results show that it is time to flip the question: Can we come up with ways to stop the shield from degrading, stop harmful substances from getting across, or at least interrupt the fall of molecular dominoes if they do? There is a chance to do a lot of good for a lot of people if we can figure these things out. SA

Daniela Kaufer is Class of 1943 Memorial Chair, a professor in the department of integrative biology, an investigator at the Helen Wills Neuroscience Institute and associate dean of biological sciences at the University of California, Berkeley.

Alon Friedman is a professor in the department of medical neuroscience and William Dennis Chair in Epilepsy Research at Dalhousie University in Nova Scotia. He is also a professor in the department of cognitive and brain sciences and Dr. Helena Rachmanska-Putzman Chair in Neurology at Ben-Gurion University of the Negev in Israel.

THE BRAIN AND LONG COVID

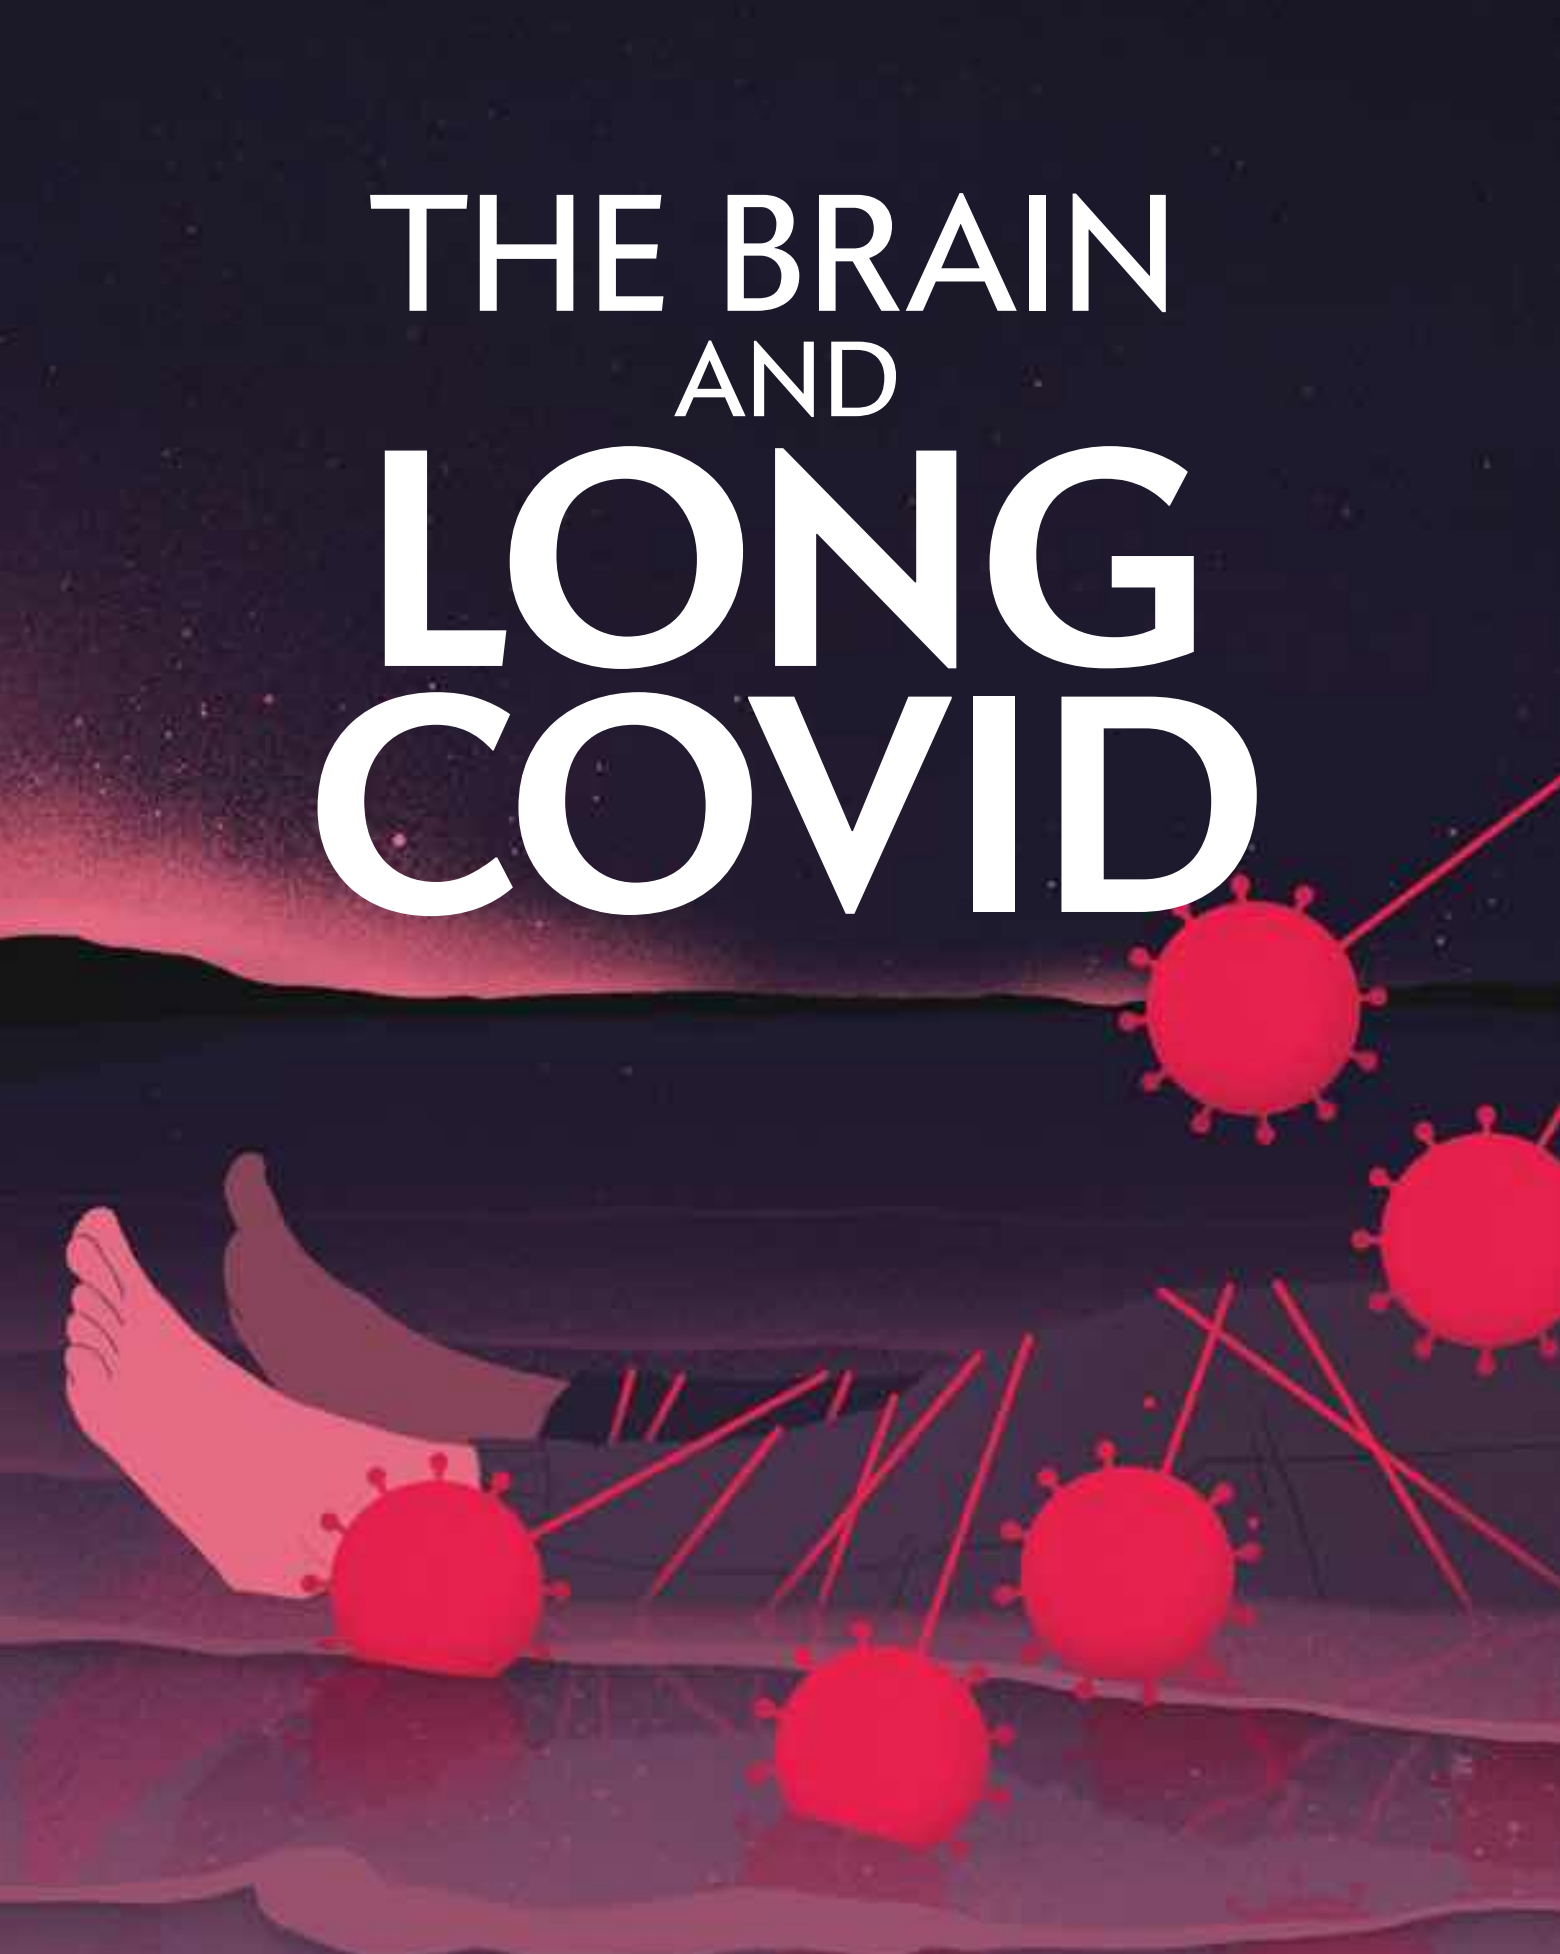

Millions of people are still suffering long after infection. Now researchers are finding neurological causes for their symptoms

By Stephani Sutherland

Illustration by Stephanie Shafer

TARA GHORMLEY HAS ALWAYS BEEN AN OVERACHIEVER. SHE FINISHED AT THE TOP of her class in high school, graduated summa cum laude from college and earned top honors in veterinary school. She went on to complete a rigorous training program and build a successful career as a veterinary internal medicine specialist. But in March 2020 she got infected with the SARS-CoV-2 virus—just the 24th case in the small, coastal central California town she lived in at the time, near the site of an early outbreak in the COVID pandemic. "I could have done without being first at this," she says.

Almost three years after apparently clearing the virus from her body, Ghormley is still suffering. She gets exhausted quickly, her heartbeat suddenly races, and she goes through periods where she can't concentrate or think clearly. Ghormley and her husband, who have relocated to a Los Angeles suburb, once spent their free time visiting their "happiest place on Earth"—Disneyland—but her health prevented that for more than a year. She still spends most of her days off resting in the dark or going to her many doctors' appointments. Her early infection and ongoing symptoms make her one of the first people in the country with "long COVID," a condition where symptoms persist for at least three months after the infection and can last for years. The syndrome is known by medical professionals as postacute sequelae of COVID-19, or PASC.

People with long COVID have symptoms such as pain, extreme fatigue and "brain fog," or difficulty concentrating or remembering things. As of March 2023, the syndrome was estimated to affect more than 15 million adults in the U.S., and a 2022 report found that it had forced between two million and four million Americans out of the workforce. Long COVID often arises in otherwise healthy young people, and it can follow even a mild initial infection. The risk appears at least slightly higher in people who were hospitalized for COVID and in older adults (who end up in the hospital more often). Women and those at socioeconomic disadvantage also face higher risk, as do people who smoke, are obese, or have any of an array of health conditions, particularly autoimmune disease. Vaccination appears to reduce the danger but does not entirely prevent long COVID.

The most common, persistent and disabling symptoms of long COVID are neurological. Some are easily recognized as brain- or nerve-related: many people experience cognitive dysfunction in the form of difficulty with memory, attention, sleep and mood. Others may seem rooted more in the body than the brain, such as pain and postexertional malaise (PEM), a kind of "energy crash" that people experience after even mild exercise. But those, too, result from nerve dysfunction, often in the autonomic nervous system, which directs our bodies to breathe and digest food and generally runs our organs on autopilot. This so-called dysautonomia can lead to dizziness, a racing heart, high or low blood pressure, and gut disturbances, sometimes leaving people unable to work or even function independently.

The SARS-CoV-2 virus is new, but postviral syndromes are not. Research on other viruses, and on neurological damage from the human immunodeficiency virus (HIV) in particular, is guiding work on long COVID. And the recognition that the syndrome may cause its many effects through the brain and the nervous system is beginning to shape approaches to medical treatment. "I now think of COVID as a neurological disease as much as I think of it as a pulmonary disease, and that's definitely true in long COVID," says William Pittman, a physician at UCLA Health in Los Angeles, who treats Ghormley and many similar patients.

Although 15 million current U.S. sufferers is a reasonable estimate of the condition's toll, there are other, more dire assessments. A meta-analysis of 41 studies conducted in 2021 concluded that worldwide, 43 percent of people infected with SARS-CoV-2 may develop long COVID, with about 30 percent—translating to approximately 30 million people—affected in the U.S. Some studies have offered more conservative numbers. A June 2022 survey reported by the U.S. National Center for Health Statistics found that among adults who had had COVID, one in five was experiencing long COVID three months later; the U.K. Office for National Statistics put the estimate at one in 10. Even if only a small share of infections result in long COVID, experts say, they will add up to millions more people affected—and potentially disabled.

Most of the first recognized cases of long COVID were in patients who needed extended respiratory therapy or who had obvious organ damage that caused lasting symptoms. People reporting neurological symptoms were often overlooked or dis-

"EVERYTHING FELL APART FOR ME," says Tara Ghormley, who has been struggling with long COVID since 2020.

missed as traumatized by their initial illness and hospitalization. But as 2020 came to an end, says Helen Lavretsky, a psychiatrist at the University of California, Los Angeles, "we started getting to a place of sorting through what was really going on … and it became very evident at that time that neuropsychiatric symptoms were quite prevalent," most commonly fatigue, malaise, brain fog, smell loss and post-traumatic stress disorder, as well as cognitive problems and even psychosis.

Ghormley was in her late 30s and relatively healthy when she caught the virus, but she had underlying conditions—including rheumatoid arthritis and asthma—that put her at risk for severe COVID. She spent several days at home, struggling to breathe, and then she went to the hospital, where her blood pressure soared and her blood glucose dropped precipitously. She mostly recovered from this acute phase within a few weeks, but, she says, "I never really got better."

Soon after coming home from the hospital, Ghormley developed what her husband called "goldfish brain." "I'd put something down and have no idea where I put it," she recalls. "It kept happening over and over. I was thinking, 'This is getting weird.' My husband said I was not remembering anything. I'd try to talk, and I knew what I wanted to say, but I couldn't think of the word."

She also experienced tremors, dramatic mood swings and painful hypersensitivity to sounds. "My husband opening a paper bag felt like knives stabbing me in the ear," she recounts. Any exertion—physical or mental—left her exhausted and in pain. The changes were jarring to Ghormley, who prided herself on her sharp mind. "The thing that bothered me the most was that I was really having trouble thinking, speaking, remembering—trying to complete a task and having no idea what it was. Suddenly I had quite profound neurological deficits. Everything fell apart for me at that time. That was horribly traumatic … it kind of broke me. I didn't feel like me."

ROOTS OF DYSFUNCTION

AS A VETERINARY INTERNIST, Ghormley says, it's her job to problem solve when mysterious symptoms arise, including her own. "I was actively trying to find reasons and find what I could do." She theorized that some of her neurological symptoms might be the result of thrombotic events, blood clots that can cause ministrokes. Several early studies showed that COVID attacks endothelial cells, which line blood vessels. That can lead to clotting and oxygen deprivation in multiple organs, including the brain. Even subtle disruption of endothelial cells in the brain could contribute to cognitive dysfunction.

One study found that in people with neurological COVID symptoms, the immune system seems to be activated specifically in the central nervous system, creating inflammation. But brain inflammation is probably not caused by the virus infecting that organ directly. Avindra Nath, who has long studied postviral neurological syndromes at the National Institutes of Health, found something similar in an autopsy study of people who died of COVID. "When you look at the COVID brain, you don't actually find [huge amounts of virus, but] we found a lot of immune activation," he says, particularly around blood vessels. The examinations suggested that immune cells called macrophages had

Photograph by Ewan Burns

been stirred up. "Macrophages are not that precise in their attack," Nath says. "They come and start chewing things up; they produce all kinds of free radicals, cytokines. It's almost like blanket bombing—it ends up causing a lot of damage. And they're very hard to shut down, so they persist for a long time. These are the unwelcome guests" that may be causing persistent inflammation in the brain.

Determining which patients have ongoing inflammation could help inform treatments. Early research identified markers that often are elevated in people with the condition, says Troy Torgerson, an immunologist at the Allen Institute in Seattle. Three cell-signaling molecules—tumor necrosis factor alpha, interleukin 6 and interferon beta—stood out in long COVID patients. But this pattern wasn't found in absolutely everyone. "We're trying to sort through long COVID patients and say, 'This would be a good group to take to trials of an anti-inflammatory drug, whereas this group may need to focus more on rehabilitation,'" Torgerson says. He led a study (currently released as a preprint, without formal scientific review by a journal) in which his team measured proteins from the blood of 55 patients. The researchers found that a subset had persistent inflammation. Among those people, they saw a distinct immune pathway linked to a lasting response to infection. "One subset of patients does appear to have an ongoing response to some virus," Torgerson says.

Isolated pockets of SARS-CoV-2 or even pieces of viral proteins may remain in the body well after the initial infection and continue to elicit an immune attack. The first solid evidence for "viral persistence" outside the lungs came in 2021 from researchers in Singapore who found viral proteins throughout the gut in five patients who had recovered from COVID as much as six months earlier. A study conducted at the University of California, San Francisco, found evidence for viral particles in the brains of people with long COVID. Scientists collected exosomes, or tiny packets of cellular material, released specifically from cells of the central nervous system. The exosomes contained pieces of viral proteins as well as mitochondrial proteins, which may indicate an immune attack on those vital cellular organelles. Amounts of such suspicious proteins were higher in patients with neuropsychiatric symptoms than in those without them.

The virus could linger in the brain for months, according to research conducted at the NIH and reported in *Nature* in December 2022. The autopsy study of 44 people who died of COVID found rampant inflammation mainly in the respiratory tract, but viral RNA was detected throughout the body, even in the brain, as long as 230 days after infection. Two other studies, both published last year in the *Proceedings of the National Academy of Sciences USA,* showed evidence that SARS-CoV-2 may infect astrocytes, a type of neural support cell, gaining entrance via neurons in the skin lining the nose.

Researchers are examining inflammatory signals in patients with long COVID in increasingly fine detail. A small study led by Joanna Hellmuth, a neurologist at U.C.S.F., found that patients with cognitive symptoms had immune-related abnormalities in their cerebrospinal fluid, whereas none of the patients without cognitive symptoms did. At the 2022 meeting of the Society for Neuroscience, Hellmuth reported that she had looked at more specific immune markers in people with cognitive symptoms and found that some patients had an elevated level of VEGF-C, a marker of endothelial dysfunction. Higher VEGF-C concentrations are associated with higher levels of immune cells getting into the brain, she says, and "they're not doing their normal function of maintaining the blood-brain barrier; they're distracted and perhaps activated." Although the studies are small, Hellmuth adds, they reveal "real biological distinctions and inflammation in the brain. This is not a psychological or psychosomatic disorder; this is a neuroimmune disorder."

What keeps the immune system in attack mode? According to Torgerson, "one option is that you've developed autoimmunity," in which antibodies produced by the immune system to fight the virus also mark a person's own cells for immune attack. The response to the virus "turns the autoimmunity on, and that doesn't get better even when the virus goes away," he says. Several studies have found evidence of autoimmune components called autoantibodies that interact with nerve cells in people with long COVID.

Clues about the inflammatory processes at work could point toward treatments for neurological symptoms. "If it's a macrophage-mediated inflammatory process … intravenous immunoglobulin could make a difference [to] dampen the macrophages," Nath says. The treatment, referred to as IVIg, contains a cocktail of proteins and antibodies that can mitigate an overactive immune response.

IVIg can also be used to block autoantibodies. And a therapy called rituximab that targets antibody-producing B cells provides "a time-tested therapy for a lot of autoantibody-mediated syndromes," Nath says. Another strategy is to use corticosteroids to dampen immune activity altogether, although those drugs can be used for only a limited time. "That's a sledgehammer approach, and you can see if it makes a difference. At least it gives you an idea that, yes, it's an immune-mediated phenomenon, and now we need to find a better way to target it," Nath says.

If the virus does hang around in some form, antiviral medications could potentially clear it, which might help resolve neurological symptoms. That's the hope of scientists running a clinical trial of Paxlovid, Pfizer's antiviral drug for acute COVID.

A CHRONIC FATIGUE CONNECTION?

POSTVIRAL SYNDROMES have been documented for more than a century, arising after infection with viruses from HIV to the flu. Epstein-Barr virus, which causes mononucleosis, is one of several viruses linked to a condition called myalgic encephalomyelitis/chronic fatigue syndrome (ME/CFS), which is estimated to affect at least one and a half million people in the U.S. ME/CFS bears striking resemblances to long COVID, with symptoms such as immune system dysregulation, fatigue and cognitive dysfunction. "One of the patterns we see is patients who definitely meet the criteria for ME/CFS. This is something we are seeing and treating all the time" in long COVID patients, Pittman says. ME/CFS can be severe, with some people losing mobility and becoming bedbound.

Nath, who also studies ME/CFS, says that "we think mechanistically they are going to be related." Researchers suspect that ME/CFS, like some cases of long COVID, could be autoimmune in nature, with autoantibodies keeping the immune system activated. ME/CFS has been difficult to study because it often arises long after a mild infection, making it hard to identify a viral trigger. But with long COVID, Nath says, "the advantage is that we know exactly what started the process, and you can catch cas-

How SARS-CoV-2 Can Harm the Brain and Nerves

Researchers have found evidence that the COVID-causing virus, SARS-CoV-2, can reach the brain and other parts of the central nervous system. This contact may lead to persistent and devastating symptoms of long COVID, which—more and more scientists say—appears to be a neurological disease. Cognitive symptoms include difficulty thinking and remembering things. And physical ailments, such as pain, extreme fatigue and a racing heartbeat, are tied to problems with the autonomic nervous system, which ordinarily runs our bodies on autopilot.

Into the Brain

Genetic material from the virus, and viral proteins, has been found in cells that line passages deep within the nose. Neurons project into this lining, and the virus can travel through them into brain areas that control breathing and the heart. It can also infect astrocytes, a crucial neural support cell.

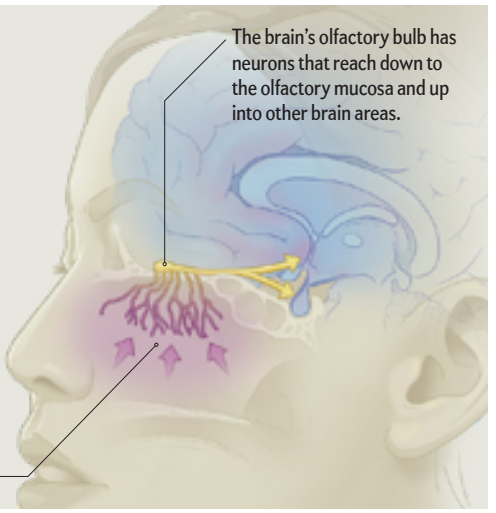

Immune System Abnormalities

Studies of long COVID patients with cognitive problems found signs that immune system cells from blood vessel walls had moved into the brain. These cells are not supposed to be in that organ and can cause damaging inflammation there. Patients without cognitive difficulties had lower levels of this unusual immune activity.

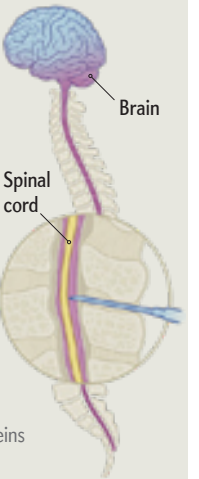

Cerebrospinal fluid collected from patients' spinal cords via a lumbar puncture (spinal tap) included proteins associated with inflammation.

Lingering Virus

In COVID patients with neuropsychiatric symptoms, proteins specific to SARS-CoV-2 appeared in small packets of cellular material that came from their neurons and astrocytes as long as three months after initial infection. This indicates the virus persists in the central nervous system for a long time. Another study found genetic material from the virus in a patient's brain almost eight months after symptoms began.

Macrophage Attack

Brains of people who have died from COVID show signs of an assault from macrophages, a type of immune system cell that reacts to invaders such as viruses. The cells surround and destroy the interlopers. But macrophages also damage nearby tissue, especially around brain blood vessels, says Avindra Nath, a neurologist at the National Institutes of Health.

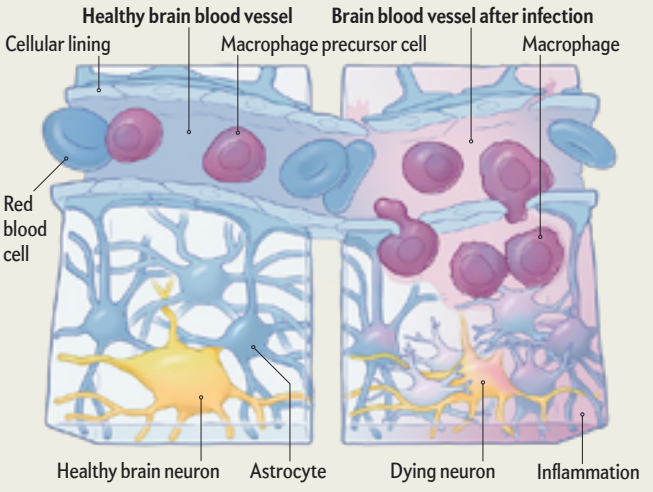

Sources: "Neurovascular Injury with Complement Activation and Inflammation in COVID-19," by Myoung-Hwa Lee et al., in *Brain*, Vol. 145; July 2022 (*blood vessel reference*); "Olfactory Transmucosal SARS-CoV-2 Invasion as a Port of Central Nervous System Entry in Individuals with COVID-19," by Jenny Meinhardt et al., in *Nature Neuroscience*, Vol. 24; February 2021 (*nasal passage reference*)

es early in the [development of] ME/CFS-like symptoms." In people who have had ME/CFS for years, "it's done damage, and it's hard to reverse that." Nath speculates that for long COVID, if doctors could study people early in the illness, they would have a better chance of reversing the process.

Torgerson hopes that researchers will ultimately come to better understand ME/CFS because of COVID. "COVID has been more carefully studied with better technology in the time we've had it than any other infectious disease ever. I think we'll learn things that will be applicable to other inflammatory diseases driven by infection followed by an autoimmune process."

TEAM TREATMENT

GHORMLEY, AFTER MONTHS of illness, sought care at UCLA Health's long COVID clinic, among the country's few comprehensive, multidisciplinary programs for people with this syndrome. Even though her symptoms are rooted in nervous system dysfunction, she needed an array of medical specialists to treat them. The clinic grew out of a program aimed at coordinating care for medically complex COVID patients, says its director Nisha Viswanathan, an internist and primary care physician. In following up with COVID patients after several months, she realized that "we had a group of patients who still had symptoms. There was no

Graphic by Now Medical Studios

understanding around the condition; we were just trying to see what we could offer them." Viswanathan and others convened a biweekly meeting of UCLA Health doctors in pulmonology, cardiology, neurology, psychiatry and other specialties to discuss individual cases and overall trends.

At UCLA Health, Pittman coordinates Ghormley's treatment. He says the interdisciplinary team is crucial to getting patients the best possible care. "Oftentimes there are so many symptoms," and some patients have seen multiple specialists before arriving, but not necessarily the right ones. As long COVID primary care providers, he says, "we do the initial testing and get them to the right person." For Ghormley, that list of providers includes Pittman, along with a neurologist, a pulmonologist, a cardiologist, a psychiatrist, a trauma counselor, a rheumatologist and a gynecologist.

The team approach has also been critical for doctors trying to understand a brand-new disease, Pittman says. "It's been a very interesting journey from knowing almost nothing to knowing a little bit now, and we're learning more every day, every week, every month," he says. The term "long COVID" "is an umbrella, and I think there are multiple diseases under that umbrella." Although each long COVID patient is unique, Pittman says, "we start to see patterns developing. And with Ghormley, we saw a pattern of dysautonomia, which we see frequently."

Dysautonomia impairs the autonomic nervous system, a network of nerves that branch out from the brain or spinal cord and extend through the body, controlling unconscious functions such as heartbeat, breathing, sweating and blood vessel dilation. For Ghormley, like many people with long COVID, dysautonomia takes the form of postural orthostatic tachycardia syndrome, or POTS. The syndrome encompasses a collection of symptoms that include a racing heart rate—particularly on standing—and fatigue, and it can cause bowel and bladder irregularities. POTS can also be a component of the exhaustion that comes with PEM. Although the symptoms may seem to affect the body, they stem from nervous system dysfunction.

Ghormley's dysautonomia led her to see cardiologist Megha Agarwal at a UCLA clinic near her home. Many physicians are not familiar with POTS, but Agarwal is particularly attuned to it, having seen it in some of her patients before COVID hit. "There's dysregulation of the nervous system, and so many things can cause it: some cancer therapies, viruses, autoimmune conditions." Agarwal recognized POTS in Ghormley in the fall of 2020, when very little was known about long COVID. Now she believes "POTS is really what long-haul COVID is causing" in many patients. Luckily, Agarwal says, there are medical interventions that can help.

Tachycardia—the T in POTS—causes the heartbeat to speed up, contributing to exhaustion and fatigue in addition to stressing the heart itself. Drugs called beta-blockers (for the beta-adrenergic receptors they shut off in the heart) can lower the heart rate and improve symptoms. "When heart rate is controlled, not only does the

COVID from the Brain to the Heart

One common affliction of people with long COVID is a heart condition called postural orthostatic tachycardia syndrome, or POTS. When a person stands up or even sits up after lying down, their heart starts racing. It accelerates by 30 beats per minute or more. This makes it hard to breathe or think ("brain fog"), and it leads to exhaustion, headaches and other symptoms. What's happening is that the brain and nervous system are losing control of the heart, something that is ordinarily managed unconsciously and automatically.

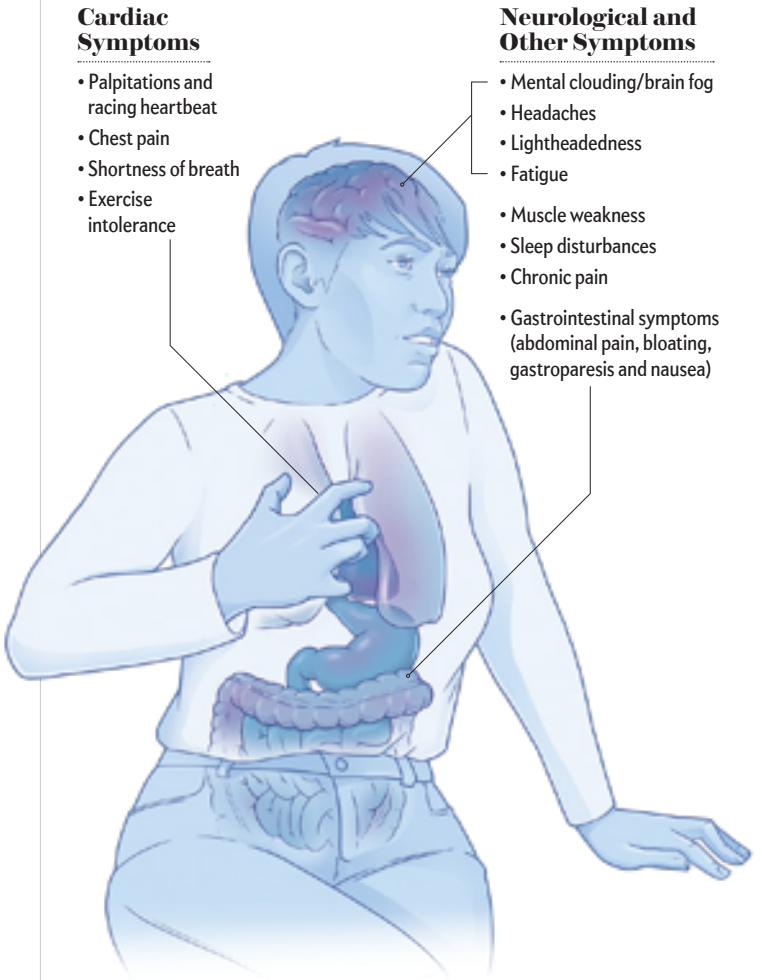

Long COVID Connections

Medical scientists have several theories about the ways that the COVID-causing virus can lead to POTS. One idea is that when a person is infected, fever, sweating and spending a long time in bed can reduce blood volume. The heart can also lose strength. To compensate and get more blood around the body, the brain gets the heart to beat faster. A second notion is that the virus could directly infect and damage nerves that set heart rhythm and send signals for blood vessel contraction. A third theory is that infection triggers an overreaction by the immune system, whose cells harm nerves in their zeal to attack the virus; some of those nerves affect heartbeat control. It is possible that some or all of these mechanisms may overlap.

Sources: "Postural Orthostatic Tachycardia Syndrome as a Sequela of COVID-19," by Cameron K. Ormiston et al., in *Heart Rhythm*, Vol. 19; November 2022; "Long COVID-19 and Postural Orthostatic Tachycardia Syndrome—Is Dysautonomia to Be Blamed?" by Karan R. Chadda et al., in *Frontiers in Cardiovascular Medicine*; March 2022 (*references*)

Graphic by Now Medical Studios

pump improve," Agarwal says, "[but people's] energy improves, their fatigue is gone, and sometimes there's better mental clarity." For some patients like Ghormley, beta-blockers are not enough, so Agarwal adds a medication called ivabradine. "It's a bit off-label, but it's currently being aggressively studied" for POTS. For Ghormley, the combination led to real improvements, "so now she doesn't feel like she ran the Boston Marathon when all she did was sit down and stand up at work or take a shower," Agarwal says.

Among Ghormley's toughest symptoms is her brain fog, a catchall term for a slew of cognitive problems that make it hard for her to function. For days when Ghormley works, her psychiatrist prescribes Adderall, a stimulant used to treat attention deficit hyperactivity disorder that helps her concentrate and stay focused. That has "helped immensely," Ghormley says.

Ghormley credits her doctors and Agarwal in particular with doing the detective work to dig into her symptoms. "Nobody knew anything about it, but everyone listened to me," Ghormley says. Perhaps because she was a professional from a medical field, no one "brushed me aside."

That's unusual for people with long COVID, many of them women, who are often dismissed by physicians who doubt their complaints are real. "Patients just don't feel heard," Viswanathan says. "I had a patient who told me everything, and after, I just said, 'This must be so hard for you. I want you to know that everything you're feeling is real, and I've seen so many patients like you.' And she started crying. She said, 'No one has told me that. I can't tell you the number of times I was told it was in my head.'"

In addition to drugs, other types of therapies, including physical therapy, can help improve some symptoms. But people who experience PEM face a particular challenge when using movement therapies. Pittman says the exertion can make these patients feel worse. "We don't want patients to go to not moving at all, but sometimes the type of movement they're doing may be flaring their symptoms." He notes that often PEM strikes young, previously healthy people who will say, "'I need to push myself,' and then they go way too far and get worse. Our job is to try to find that middle ground and then make that consistent over time, so they're not getting further deconditioned but they don't have the PEM, which has been shown to set them back."

THE LONG HAUL

SOME PATIENTS, Pittman says, "have the expectation that they're going to come in, and within a month they're going to be back to normal. And resetting those expectations can be really challenging. You have to be really empathetic because people's lives have completely changed." But sometimes patients' quality of life can improve noticeably when they are able to adjust to a new normal. Still, he says, "patients have so many questions, and I can't lead them down a physiological pathway. I can tell them there's neuroinflammation, maybe there's autoimmunity, but we still don't have the answers. Sometimes it's really tough for us to accept and for the patient to accept that we just have to try our best."

A number of people, Viswanathan says, benefit from reducing various treatments they have accumulated. Some people become so desperate that they will try anything from supplements to off-label medications to untested potions from the Internet. Stopping those sometimes leads to improved symptoms, she says.

Psychological care and support groups can help. Lavretsky adds that "lifestyle choices can play a huge role in improvement," particularly better sleep habits and the use of breathing exercises to control anxiety. She tells people their bodies can heal themselves if the patients and clinicians find the right tools.

Whether that's true for everyone remains to be seen, Viswanathan says. "We see many patients who have gotten better with time. I have patients whose symptoms have disappeared in the course of a year, or they disappear and occasionally flare up again." But for some, she says, "it could last many years."

"We're going to be addressing this for probably decades," Viswanathan says. "COVID is not going to go away so much as we're just going to get used to living with it, but part of [that] means that people will continue to develop long COVID."

Vaccination appears to reduce the risk of long COVID. But a study published in May 2022 in *Nature Medicine* suggests the protection, though real, is not as good as one might hope. The survey of electronic health records from the U.S. Department of Veterans Affairs looked at the relatively small portion of vaccinated people who subsequently became infected. They developed long COVID only 15 percent less often than unvaccinated people. "These patients can have symptoms for one to two years or longer, and so every month you're racking up more patients. Even if it's 15 percent less, the total population of patients is still growing and exploding," Pittman says. The best way to avoid getting long COVID, experts all agree, is to avoid getting COVID at all.

> An interdisciplinary medical team is crucial to getting patients the best care because often there are so many symptoms.

The syndrome is still mired in a lot of medical uncertainty. Patients might have one or a combination of the problems investigated so far: Long COVID might be caused by viral particles that persist in the brain or other parts of the nervous system. Or it might be an autoimmune disorder that lasts long after the virus has disappeared. Maybe overactive immune cells continue to perturb the nervous system and nearby blood vessels. Fortunately, the increasing ability to recognize specific problems is helping clinicians hone treatments that give patients the best chance of recovery.

Although Ghormley says her care has dramatically improved her symptoms and allowed her to "do some normal things again," she continues to experience flare-ups that make it impossible for her to work for weeks at a time. One day last year she skipped a dose of her heart medication and made a Target run in the southern California heat. "I got home and basically collapsed in the hallway," she says, and even months later things were "out of whack. If I try to move around, my legs give out." Most frustrating—and scary—to Ghormley is the unpredictability of her symptoms. "They have changed so much; some are manageable, some debilitating. One thing will get better, and another thing comes back. I'm always hopeful that it's going to get better, but I just don't know." SA

Stephani Sutherland is a neuroscientist and science journalist based in southern California. She wrote about the causes of autoimmune diseases in our September 2021 issue. Follow her on Twitter @SutherlandPhD

A Talking Cure for Psychosis

Cognitive-behavioral therapy helps to treat symptoms for which drugs are ineffective

By Matthew M. Kurtz

Schizophrenia is a mental disorder that, at its worst, ravages the totality of everyday life. It is hard to imagine what people with the severest forms of the ailment experience as anything but biologically driven, a direct consequence of aberrant chemical and electrical activity occurring deep within the brain.

As a neuropsychologist, I have often seen convincing evidence of schizophrenia's biological underpinnings in my dealings with patients. To illustrate what I mean, I will describe "Billy"—a composite profile derived from various patients I encountered in my work at the Institute of Living

in Hartford, Conn. Billy exhibits the detachment from reality and emotional agony brought on by a psychotic episode, symptoms experienced by many people with schizophrenia.

Think of Billy as a 35-year-old man who lopes from one corner of a psychiatric hospital to another, gazing down at his feet and repeating, "Billy likes model trains. Billy likes model trains. Billy likes model trains." When a clinician on the unit asks Billy about his treatment goals, he replies, "Saturn is going to crash into Mother Earth." Billy was recently flooded with anxiety as he became convinced that the Blob, the smothering gelatinous substance from the 1958 science-fiction movie of the same name, was about to engulf his neighborhood.

Billy's severe psychotic episodes point to a brain in disequilibrium and, correspondingly, to the need for drugs and treatments to alter this pathology as a basic standard of care. Findings from decades of research and clinical practice support the crucial role of antipsychotic medications, which interfere with the transmission of the brain neurotransmitter dopamine, in the attenuation of delusions, hallucinations and other symptoms that are so apparent in Billy's case.

Yet key features of schizophrenia, such as reduced spoken communication and inappropriate expression of emotions, remain entirely untouched by pharmacological interventions. Similarly, these medications make little difference in the social disability that characterizes the disorder: the chronic unemployment, social withdrawal and isolation, high suicide rates, and abbreviated life spans that typically coincide with a diagnosis. None of these medications has been found to help prevent the onset of the disorder in people at high risk because of family history and mild psychological symptoms that are not severe enough to meet criteria for a diagnosis of schizophrenia.

Research in the psychological sciences and related disciplines is now broadening understanding of both the emergence of schizophrenia and its treatment. These newer approaches focus on psychosocial stress and the patients' own belief systems. Methodologically rigorous, large-scale, population-based studies are delving more deeply into environmental factors linked to disease onset. Other research has shown that talk-based psychotherapy may be able to constrain the aberrant beliefs associated with schizophrenia.

Medical professionals are starting to pair therapy with methods to strengthen thinking skills. For the past 20 years my laboratory has studied ways to measure and improve concentration, memory and problem-solving. The results of these studies show that impairment in these thinking skills, even beyond other, more visible symptoms, often stands in the way of functional improvement for people with schizophrenia and related illnesses. Such work has led to a much more nuanced view of the disorder that highlights psychological factors and complements the biological models that have dominated the field.

THE PATHOLOGICAL PARENT

FOR MUCH of the 20th century and as recently as 50 years ago, cases such as Billy's were viewed in many professional quarters as resulting from a disorder of the mind that was rooted in pathological parenting styles and influenced by cultural milieus. Sigmund Freud's inordinate influence on psychiatry, particularly in the U.S., led professionals in that field to view illnesses through a lens of environmental factors centered on unresolved family trauma, with talk therapy being the key to relief. But this approach yielded scant success.

The first psychological models of schizophrenia declined in favor for a number of reasons, including stubbornly high psychiatric hospitalization rates and persistently poor outcomes, even among patients with access to the most intensive psychological care and therapy. Charismatic health-care providers developed treatments based on their own theories and supported their claims with isolated case studies rather than exacting scientific data. Many psychotherapists of this era resisted randomized, controlled trials. Solid evidence that the prevailing treatments worked for most patients never materialized.

A watershed event occurred in 1952 with the publication of a clinical trial of chlorpromazine at St. Anne's Hospital in Paris, heralding the arrival of a new class of drugs known as antipsychotic medications. These pharmaceuticals moderated irrational, often paranoid beliefs and hallucinations for many people with schizophrenia, stabilizing them in the hospital and, in many cases, allowing them to reenter the community for the first time in years. The psychological approach to treatment became less dominant as therapists embraced pharmacological therapies targeting the brain.

Three additional factors played a pivotal role in reinforcing the view of schizophrenia as a neuroscientific entity. First there was a new focus on psychiatric illnesses as diseases with consistent signs and symptoms, just like other medical illnesses, which meant they could be studied through rigorous biological analysis. This medical model made it much more likely that patients with similar patterns of symptoms would reliably receive the same psychiatric diagnosis.

The second factor was the emergence of highly detailed imaging technologies that researchers used to look first at the brain's structure and later at the function of its various regions. By the early 2000s it was becoming clear that people with schizophrenia had reductions in brain activity and tissue volume across a broad range of neural systems, particularly in the frontal and side (temporal) brain lobes. Researchers replicated these findings and discerned these changes even in patients experiencing their first episode of schizophrenia, before they had received any antipsychotic therapy.

Third, the mapping of the human genome in the early 2000s and the development of cheaper technology for identifying genetic variants raised the possibili-

ty of determining which genes put people at the greatest risk for schizophrenia. Studies have identified more than 100 locations on DNA that confer increased susceptibility to the disorder. If researchers could use this genetic analysis to identify aberrant protein synthesis, new drugs could be formulated to interfere with this process.

A NEW ERA OF PSYCHOLOGY

SCIENTISTS HAVE MADE undeniable advancements in the neurological and genetic understanding of schizophrenia. But over the past 20 years a growing body of work has suggested yet another revised view of the disorder. Much of this research comes from academic precincts neighboring neuroscience—not just psychology but epidemiology and anthropology. This fresh perspective goes beyond the physiology of schizophrenia to encompass personal belief systems, social interactions and the destructiveness of psychological stress. It also emphasizes the importance of the environment in which a patient lives in explaining the origins of symptoms, not just for schizophrenia but for related psychoses, such as bipolar disorder with accompanying psychotic symptoms.

The field's return to a psychological emphasis is a product, in part, of neuroscience's failure to deliver clear answers about schizophrenia despite its promise for identifying the neural underpinnings of psychiatric conditions. Hundreds of informative structural and functional neuroimaging studies have identified locations in the brains of affected people where there is diminished tissue volume or aberrant activity. But huge genome-wide association studies—which recruit thousands of patients to pinpoint genetic variants that may place people at higher risk for schizophrenia—have failed to define the causes of the illness. So far none of these findings has led to the development of drug treatments that meaningfully alter the course of the disorder.

Evidence from recent studies has helped bolster renewed interest in the psychological underpinnings of schizophrenia. These are rigorous, well-designed, large scientific studies in which researchers carefully quantified patients' experiences, using measures that are consistent over time and that have been validated with other forms of evidence. New psychological therapies undergo testing to minimize biases that might affect whether a treatment is judged effective. Markers of patients' improvement in these studies are standardized and made objective, and study participants and their evaluators often do not know whether participants are in a treatment or a control group. This type of study design helps to ensure that participants don't appear to improve simply because they or their evaluators believe they should be getting better.

Such studies have produced a consensus that adverse experiences and environments contribute in important ways to the development of schizophrenia. For example, rates of psychosis are dramatically higher in some minority immigrant communities compared with those of the native-born populations in their adopted countries. The United Kingdom Aetiology and Ethnicity Study in Schizophrenia and Other Psychoses (AESOP) followed patients who sought clinical treatment starting with their first episode of psychosis. The investigators used census data to obtain estimates of incidence rates, and diagnoses were based on chart notes and a standardized interview and were analyzed by psychiatrists who never learned the patients' ethnicity.

AESOP confirmed previous findings that African-Caribbean and Black African people living in the U.K. were diagnosed with schizophrenia at a rate as much as five to 10 times higher than that of white Britons. A study of the incidence of psychosis in a broader range of immigrant groups—from the Middle East, North Africa, China, Vietnam and Japan—showed that their likelihood of developing psychosis was almost three times greater than that of white Britons. (People from these regions were combined into one study group because the size of each ethnic group alone was too small to be statistically meaningful.) Some aspect of the immigrant experience or a person's minority status, or a combination of these two factors, appeared to be contributing to elevated rates of schizophrenia. More important, these rates were typically much higher than those in the migrants' countries of origin.

Rigorous scientific studies have brought a renewed focus on schizophrenia's psychological roots.

The authors found that psychological and social stresses correlated with the increased incidence of schizophrenia in minority ethnic groups. Separation from a parent during childhood was associated with rates of diagnosis two to three times higher than among people whose families had remained intact. A variety of other markers of social disadvantage, including unemployment, living alone, being single, a lack of formal education and limited social networks, all increased the likelihood of schizophrenia onset in various ethnic groups. White Britons showed similar links between psychosocial stressors and the likelihood of schizophrenia, but immigrant minority groups experienced those stressors more frequently.

Social discrimination may also increase a person's chances of developing schizophrenia. A study from the Netherlands looked at all non-Western immigrants seeking services for a first episode of psychosis in The Hague between 2000 and 2005. The researchers studied minority groups from Morocco, the Antil-

les, Suriname and Turkey, among other regions, and interviewed members of these groups about the perceived levels of discrimination they encountered. Moroccans, the ethnic group that experienced the most discrimination, showed the highest incidence of psychosis, whereas ethnic groups reporting less bias (Turks, Surinamese, and others) had lower rates.

The European Network of National Schizophrenia Networks Studying Gene-Environment Interactions (EU-GEI) took a detailed look at the contribution of immigration in the emergence of psychosis. Using a data set of more than 200 migrants and 200 control participants matched for a variety of variables, including family history of psychosis, the authors defined indicators of social disadvantage for each phase of immigration. In the premigration phase, indicators consisted of parental social class, type of employment and whether the participant lived with their family of origin. In the active migration process, indicators consisted of age, whether the person was detained at any point during immigration and whether they had any plans to return to their country of origin. For the postmigratory phase, the measures included employment during the previous five years, long-term relationships and family structure.

> Social disadvantages and adversity in the migrant experience doubled the chances of developing psychosis.

The study found that among first-generation migrants, social disadvantages and adversities during the premigration and postmigration phases doubled a person's chances of developing psychosis even when other risk factors such as cannabis use and age were statistically controlled for. Mismatches between the expectations people held before leaving their native countries and their actual postmigration achievements were also associated with an increased likelihood of psychosis. The risk of illness increased with the number of cumulative adversities. These findings all suggest that providing psychological support to immigrants might alter the onset of schizophrenia in those who are facing high levels of social adversity.

HEARING VOICES

PEOPLE FROM DIFFERENT cultures experience psychosis in distinct ways. Psychological anthropologist T. M. Luhrmann and her collaborators have shown that the emotional tone of auditory hallucinations may vary widely across cultures, suggesting that what one hears may be influenced by cultural expectations. Those in subcultures that have encountered the violence of war or other social upheaval may experience loud or disruptive hallucinations, whereas those with dense family ties may have more benign symptoms.

Luhrmann's research team conducted structured interviews of people diagnosed with schizophrenia, most of them ill for years, from the U.S., Ghana and India. The researchers asked people how many voices they heard, how often they heard them, what their opinion of these voices was and what they believed was the source. In general, the results support the idea that schizophrenia is a biologically based condition that manifests similarly across cultures. Participants in all three countries heard "good" and "bad" voices. Some reported having back-and-forth conversations with the voices they heard, and others thought the voices came from God. At every site, at least some participants disliked their voices and viewed them as an intrusion on their daily mental life.

Many of the symptoms were similar across groups, but the interpretation and consequent emotional tone of these hallucinatory experiences diverged substantially across cultures. Participants in the U.S. were much more likely to use an unadorned clinical diagnostic label—"I am a schizophrenic"—to describe their lives. They tended to report violent imagery associated with their voices more frequently than participants from India or Ghana.

In the other two countries, people were more likely to maintain close relationships with their voices and less apt to describe them as the expression of a mind violated by auditory hallucinations. In Ghana, study participants insisted that voices spoken by an invisible person were controlled by God and that at times evil voices were entirely absent. Rarely did people there describe voices as an intrusion on their everyday mental life. Participants from India often experienced their voices as those of family members. They said the voices conveyed a mix of agreeable and unpleasant utterances but did not have stressful or harsh overtones.

COGNITIVE-BEHAVIORAL THERAPY

IF CULTURAL ATTITUDES exert such a profound effect on the experience of symptoms, that insight holds promise for psychotherapies. It raises the possibility that talking to people with schizophrenia and offering them strategies for changing the way they think about their symptoms can reduce the distress those symptoms cause.

An important question is whether altering beliefs around symptoms can improve people's ability to function in society. A growing body of scientific literature suggests such a goal is achievable. A form of cognitive-behavioral therapy (CBT) has been developed specifically to treat psychosis. It focuses on detrimental thinking—intrusive thoughts that crop up such as "Why even try? I always fail."

The aim of CBT is to help people deal with their

emotional and behavioral responses to psychological experiences that cause distress. Atypical or disproportionate responses, which are often among the most debilitating of schizophrenia's symptoms, can make it difficult to carry out daily tasks. Patients undergoing CBT are taught to think about their symptoms in a new way. They might tell themselves, "The voices in my head don't have to make me anxious; it is the way I think about them that makes me anxious."

People with schizophrenia commonly believe their voices are all-knowing, all-powerful and uncontrollable. CBT can generate alternative explanations for these auditory hallucinations. It can begin a process of interrogating and weakening unhelpful beliefs. A clinician might suggest that a voice a client hears could be that of a family member as opposed to a deity or the devil. The doctor might frame this for the patient in a simple question and statement: "Are we certain that the voice you hear is not your father? Many of the statements the voice makes seem to be similar to ones you have attributed to your father in the past." Such reconsideration of a voice's meaning can lead to a significant decrease in the distress associated with the hallucination.

Other strategies include behavioral tasks to show that voices are not in fact uncontrollable. A therapist might lead a client in an activity—walking outside or listening to music on headphones—to help the person quiet the constant chatter, gain mastery over their symptoms, and disrupt their beliefs about the voices being an inevitable, eternal intrusive presence. The client also might try simply ignoring the stream of commands issued by their inner voices. This can undermine their belief that the commands from hallucinations must be followed or terrible consequences will ensue. When the client discovers that ignoring voices does not produce some feared catastrophe, the realization supports the counterargument that their voices are not all-powerful.

Research provides evidence that this suite of interventions may be effective even for people with the most severe symptoms. In one of the most remarkable demonstrations of the benefits of therapy for psychosis to date, CBT pioneer Aaron Beck, in some of the most influential work he conducted before his death in 2021 at the age of 100, worked with Paul Grant and their colleagues at the University of Pennsylvania to evaluate the impact of a modified approach to CBT that addresses the needs of low-functioning people with schizophrenia. Their study was published in 2012 in *Archives of General Psychiatry*.

The patients they worked with had moderate to high levels of what are labeled negative symptoms of schizophrenia: low motivation, diminished pleasure in life, near absence of spoken language, and reduced emotional expressiveness to the point that they maintained a "wooden" expression during social interactions. Among the most disabling, these symptoms are also the hardest to treat with medication and are disproportionately represented in the most persistently ill patients. People with a high intensity of negative symptoms also typically have the most elevated levels of cognitive distortions and biases. To date, there is no pharmaceutical treatment for negative symptoms.

In the study, participants were randomly assigned to either a control group, in which patients were given standard treatment, including prescribed drugs, or a test group, in which they received CBT in addition to the standard therapy. The CBT was intended to help clients establish long-term goals (seeking independent housing, relationships or a job) as well as intermediate- and short-term goals (calling a friend that day).

Therapy can quiet the train of negative thoughts—beliefs such as "taking even a small risk is foolish because the loss is likely to be a disaster" and "making new friends is not worth the energy it takes." Participants in the CBT group also took part in exercises, games, role-playing and community outings designed to instill belief in their own abilities. The benefits of this therapy persisted for months after treatment ended. Clients assigned to CBT had meaningful improvements in functioning—better motivation and reduced delusions and hallucinations—compared with patients who received only standard treatment.

From what researchers have learned in recent years, adverse experiences increase the likelihood that someone will develop schizophrenia. In addition, the cultural context in which people experience symptoms may affect their ability to come to terms with those symptoms. All these findings support the argument that key aspects of schizophrenia are rooted in the psychology of stress and trauma and in attitudes and biases that are shaped by the persistent lingering of a patient's mental anguish. Treatments designed to address negative biases and societal discrimination and stigma can improve symptoms and functioning in people with schizophrenia, which further highlights the key role that psychology is starting to play in understanding and treating the disorder.

None of these findings throws into question the changes in brain structure that accompany schizophrenia or the genes currently implicated in the disorder. What they suggest is that if methods of prevention and treatment for schizophrenia are to progress, increased public health focus on mitigation of damaging social experiences, along with therapies focused on psychological beliefs and attitudes, is critical.

Psychological therapies need to be prioritized by both practitioners and federal funding agencies and placed on more equal footing with gene and brain-imaging studies. Psychoactive medications will take a person with schizophrenia only so far in adapting to the personal struggles their condition brings. That is why an interaction with a therapist able to question their ideas and basic beliefs is also essential to make peace with the din of voices in their head. SA

Matthew M. Kurtz has professorships in psychology and in neuroscience and behavior at Wesleyan University.

WHEN DREAMS FORESHADOW BRAIN DISEASE

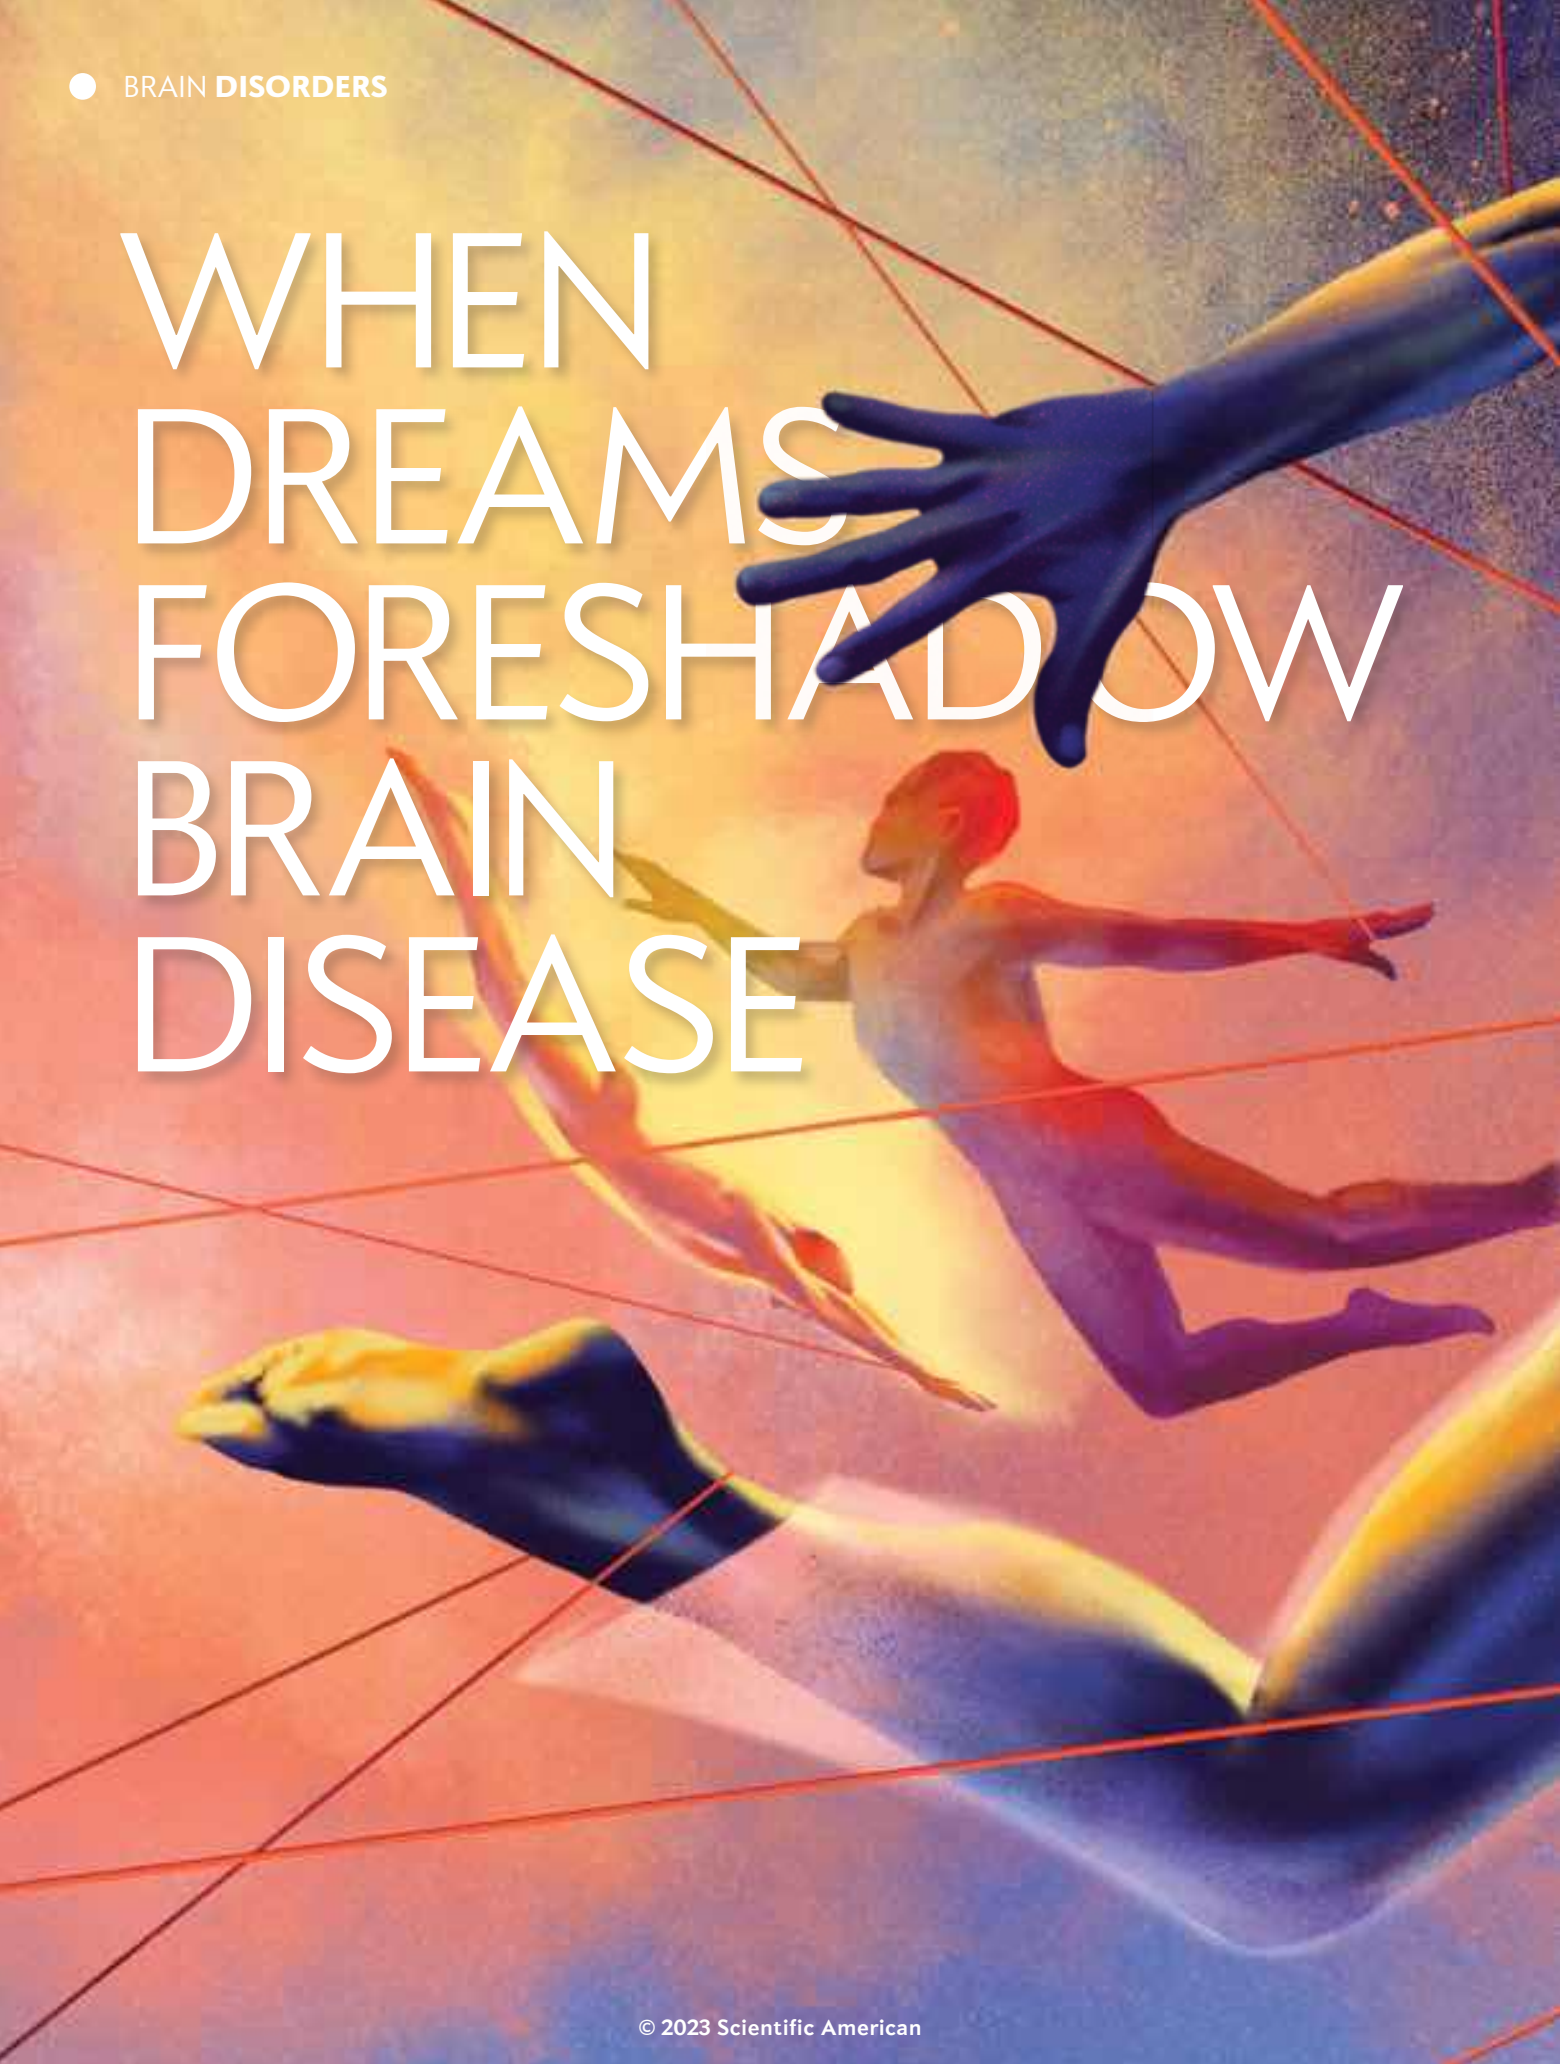

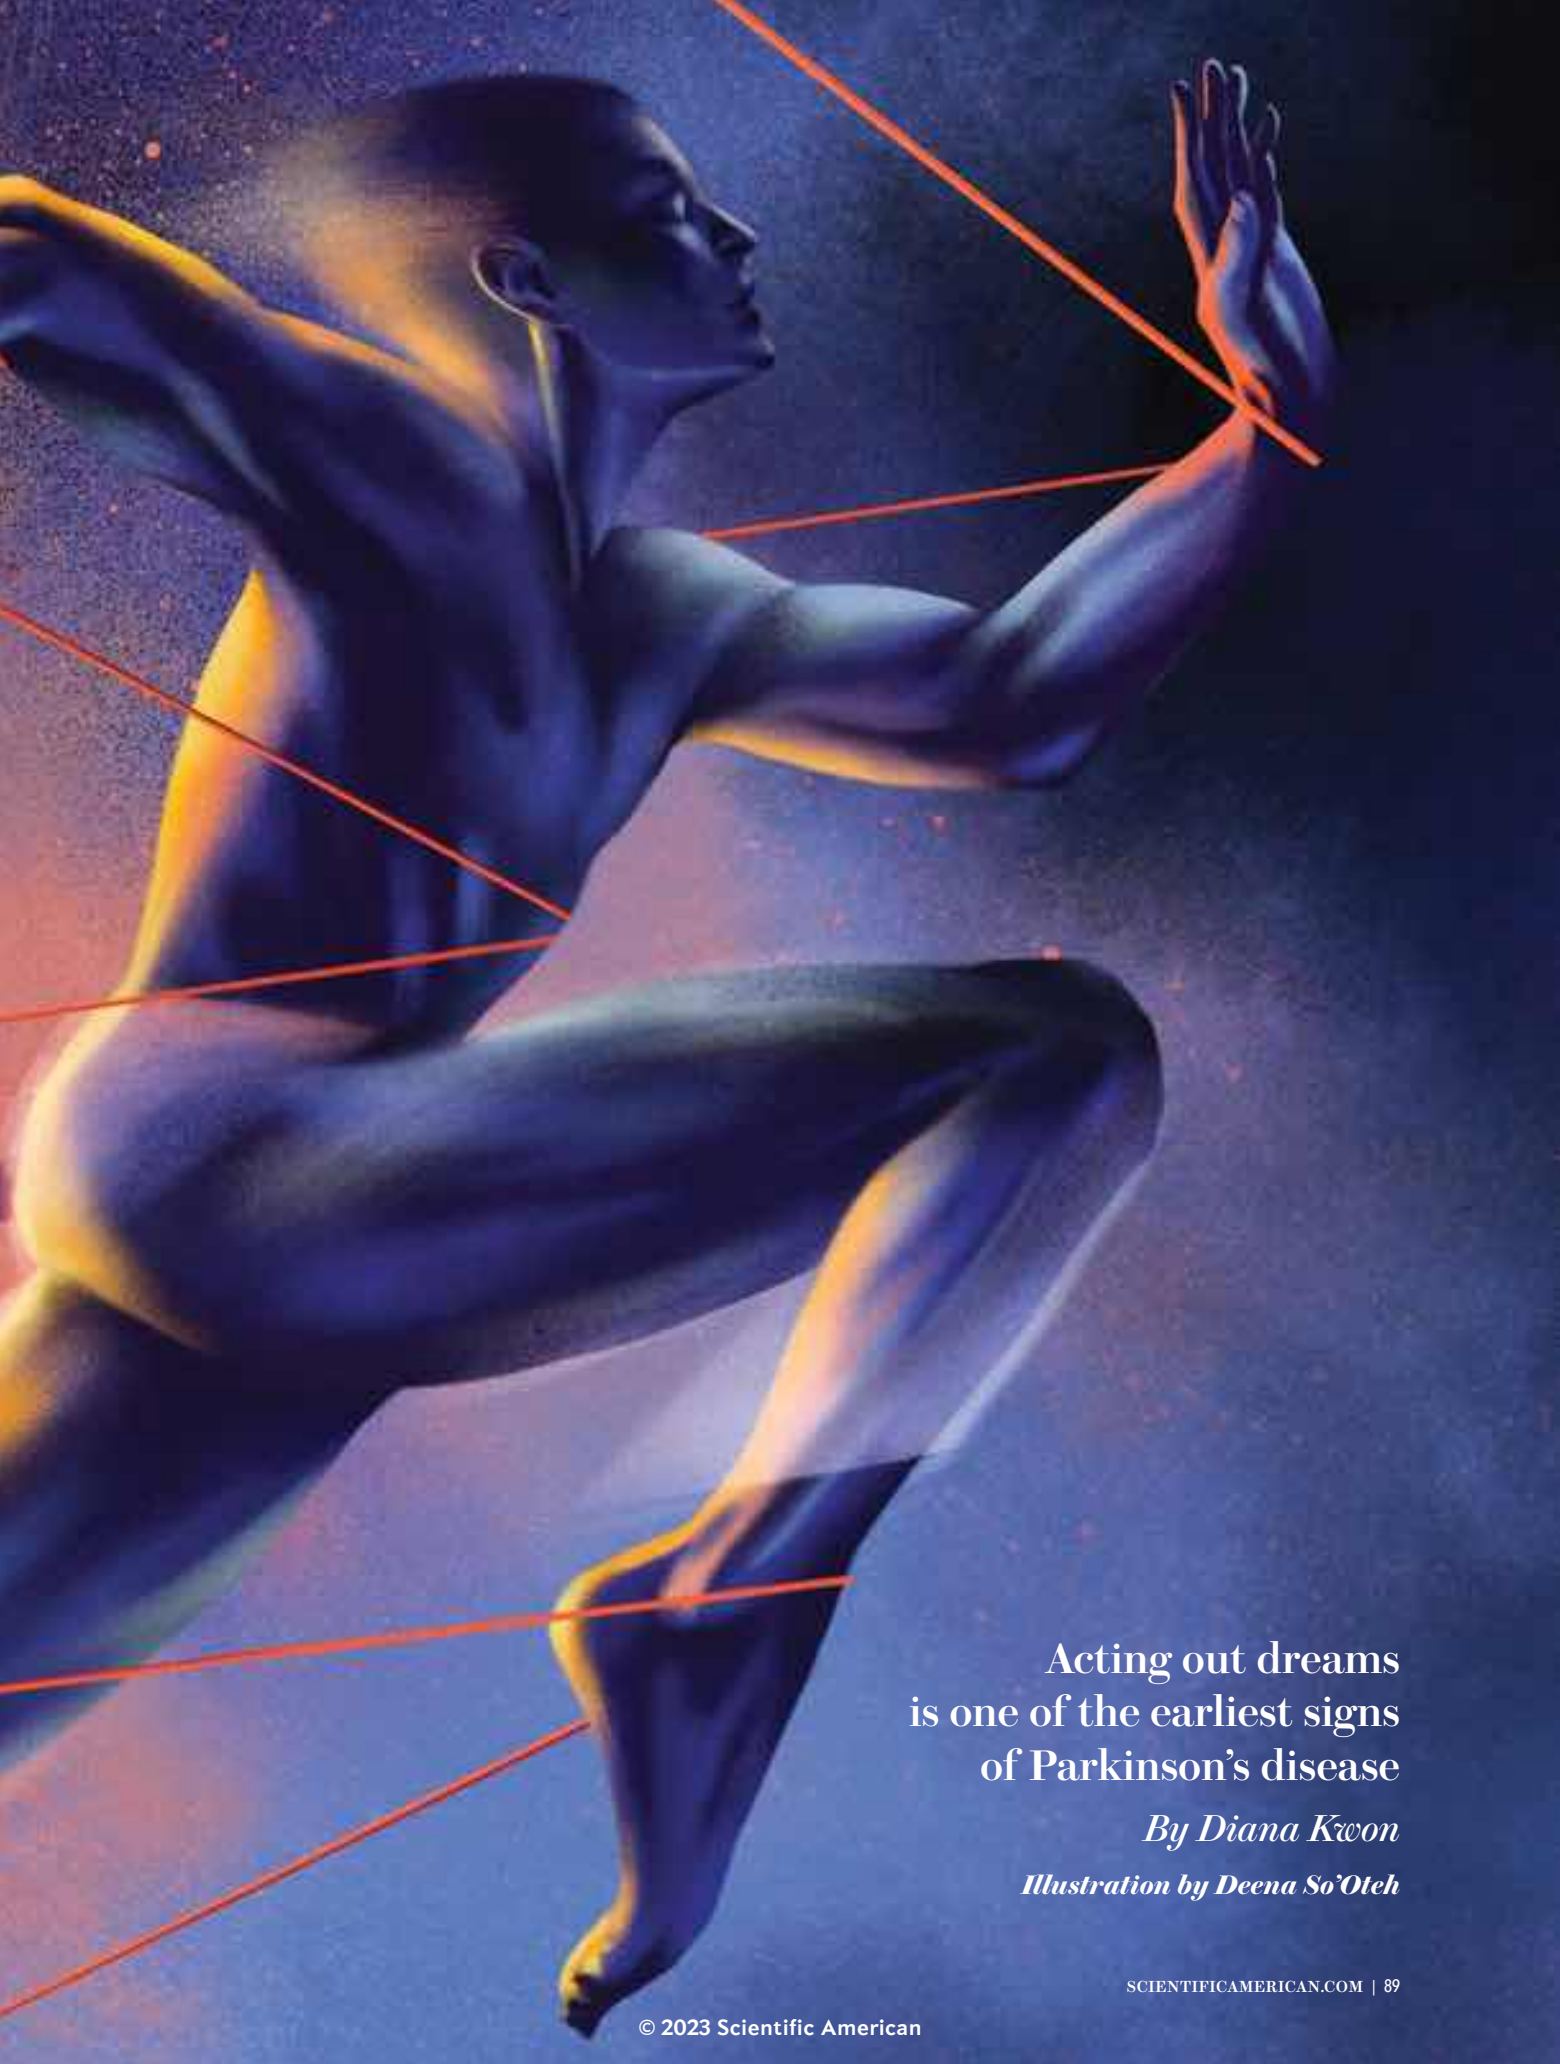

Acting out dreams is one of the earliest signs of Parkinson's disease

By Diana Kwon

Illustration by Deena So'Oteh

ALAN ALDA WAS RUNNING FOR HIS LIFE. THE ACTOR, BEST known for his role on the television series *M*A*S*H,* wasn't on a set. This threat was real—or at least it felt that way. So when he saw a bag of potatoes in front of him, he grabbed it and threw it at his attacker. Suddenly, the scene shifted. He was in his bedroom, having lurched out of sleep, and the sack of potatoes was a pillow he'd just chucked at his wife.

Acting out dreams marks a disorder that occurs during the rapid eye movement (REM) phase of sleep. Called RBD, for REM sleep behavior disorder, it affects an estimated 0.5 to 1.25 percent of the general population and is more commonly reported in older adults, particularly men. Apart from being hazardous to dreamers and their partners, RBD may foreshadow neurodegenerative disease, primarily synucleinopathies—conditions in which the protein α-synuclein (or alpha-synuclein) forms toxic clumps in the brain.

Not all nocturnal behaviors are RBD. Sleepwalking and sleep talking, which occur more often during childhood and adolescence, take place during non-REM sleep. This difference is clearly distinguishable in a sleep laboratory, where clinicians can monitor stages of sleep to see when a person moves. Nor is RBD always associated with a synucleinopathy: it can also be triggered by certain drugs such as antidepressants or caused by other underlying conditions such as narcolepsy or a brain stem tumor.

When RBD occurs in the absence of these alternative explanations, the chance of future disease is high. Some epidemiological studies suggest that enacted dreaming predicts a more than 80 percent chance of developing a neurodegenerative disease within the patient's lifetime. It may also be the first sign of neurodegenerative disease, which on average shows up within 10 to 15 years after onset of the dream disorder.

One of the most common RBD-linked ailments is Parkinson's disease, characterized mainly by progressive loss of motor control. Another is Lewy body dementia, in which small clusters of α-synuclein called Lewy bodies build up in the brain, disrupting movement and cognition. A third type of synucleinopathy, multiple system atrophy, interferes with both movement and involuntary functions such as digestion. RBD is one of the strongest harbingers of future synucleinopathy, more predictive than other early markers such as chronic constipation and a diminished sense of smell.

Descriptions of dream enactment by people with Parkinson's are as old as recognition of the disease itself. In James Parkinson's original description, "An Essay on the Shaking Palsy," published in 1817, he wrote: "Tremulous motions of the limbs occur during sleep, and augment until they awaken the patient, and frequently with much agitation and alarm." But despite similar reports over the next two centuries, the connection between dreams and disease remained obscure—so much so that Alda had to convince his neurologist to do a brain scan for Parkinson's

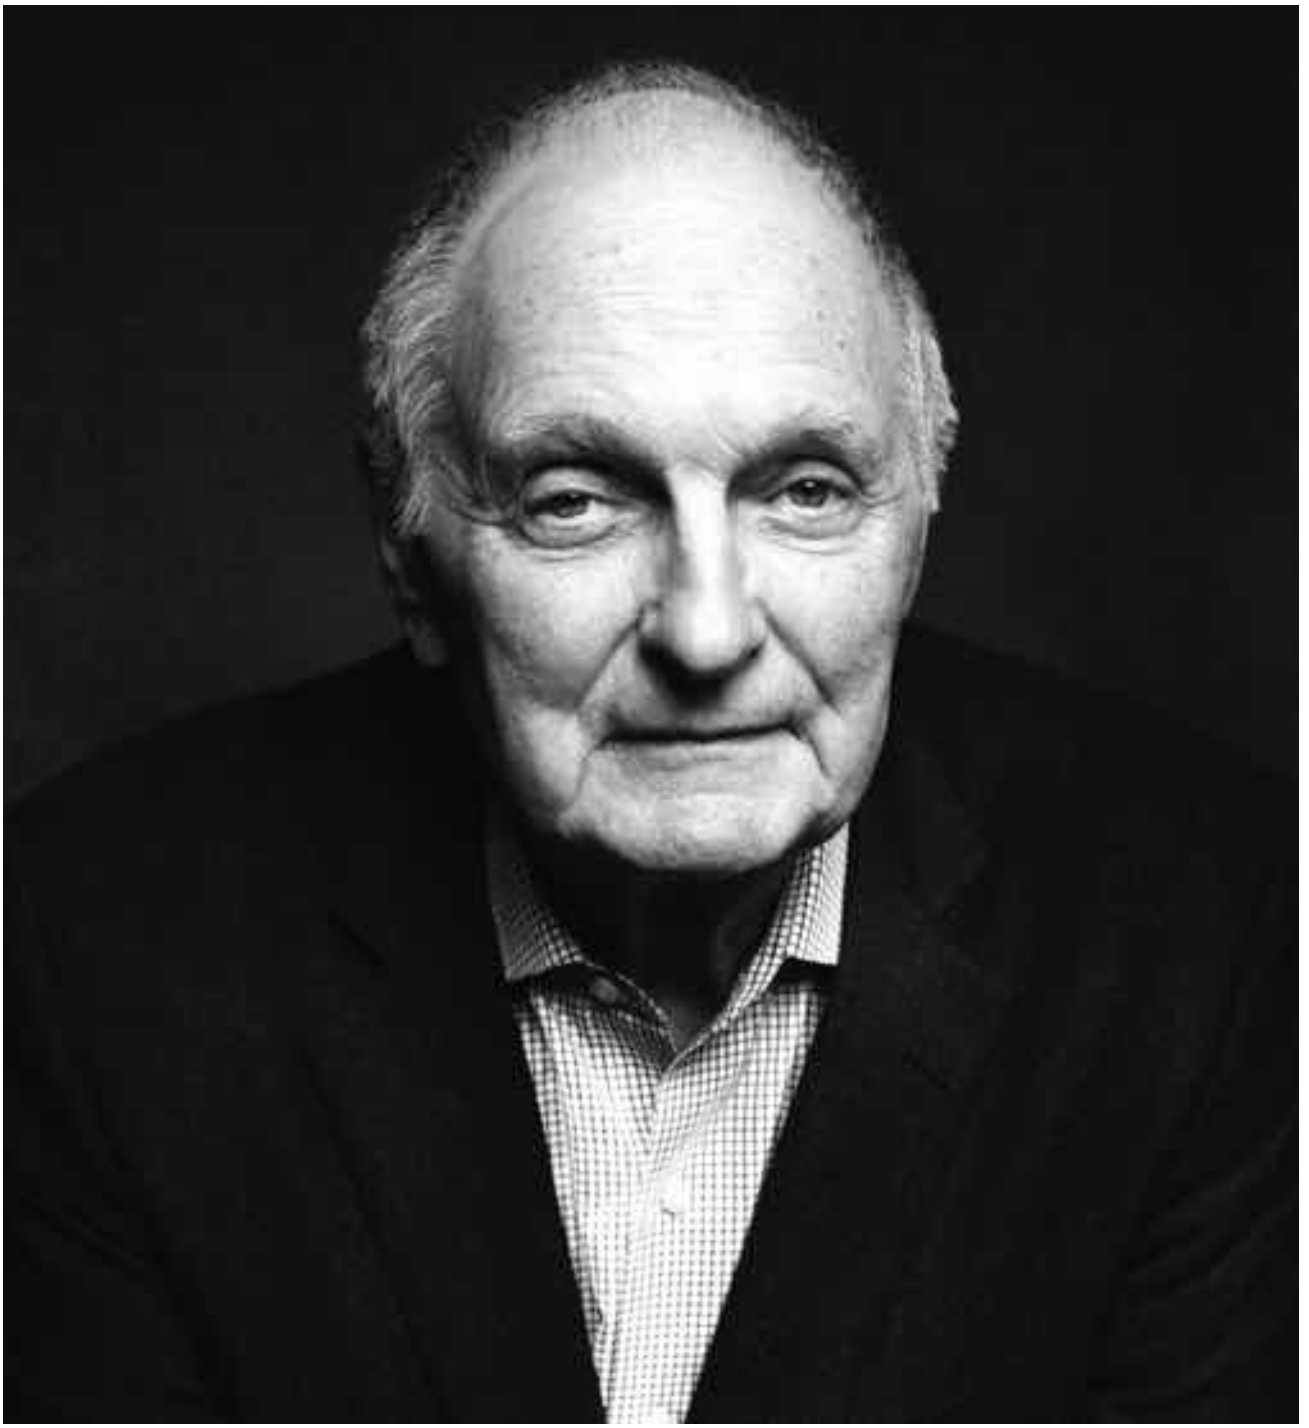

Jesse Dittmar/Redux

ACTOR ALAN ALDA helps to raise awareness of Parkinson's and its early symptoms to give people a head start on dealing with the disease.

after he read about the link in a 2015 news article.

Those scans confirmed Alda's suspicion: he had Parkinson's. He shared his experience with the public "because I thought anybody who has any symptom, even if it's not one of the usual ones, could get a head start on dealing with the progressive nature of the disease," he says. "The sooner you attack it, I think, the better chance you have to hold off the symptoms."

In recent years awareness of RBD and an understanding of how it relates to synucleinopathies have grown. Studying this link is giving researchers ideas for early intervention. These advances contribute to a growing appreciation of the so-called prodromal phase of Parkinson's and other neurodegenerative disorders—when preliminary signs appear, but a definitive diagnosis has not yet been made. Among the early clues for Parkinson's, "RBD is special," says Daniella Berg, a neurologist at the University Hospital Schleswig-Holstein in Germany. "It's the strongest clinical prodromal marker we have."

Dream Cycles

When we fall asleep, the brain begins to cycle through different stages, marked by characteristic patterns of activity. Brain waves during rapid eye movement (REM) sleep look rather like those in awake brains, representing vivid dreams. Muscles are normally immobilized to prevent injury from acting dreams out. But in RBD, or REM sleep behavior disorder, this sleep-time paralysis is lifted, likely by damage to the brain stem.

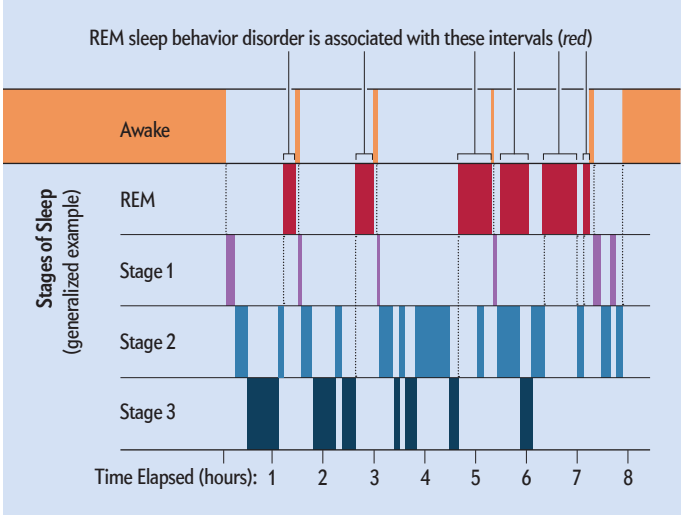

LIFTING THE BRAKE

RAY MERRELL, a 66-year-old living in New Jersey, started acting out his dreams around 15 years ago. His dreamscapes became action-packed, like "something you'd watch on TV," Merrell says. He often found himself either being chased by or chasing a person, animal or something else. In the real world, Merrell was flailing, kicking and jumping out of bed. Some of his violent nighttime behaviors injured him or his wife.

In people with RBD, the brakes that normally immobilize them during REM sleep—the stage of sleep most closely linked with dreaming—are lifted. (Dreaming also occurs in non-REM sleep, but dreams during REM are longer, more vivid and more bizarre.)

In the 1950s and 1960s French neuroscientist Michel Jouvet conducted a series of experiments that revealed just how chaotic unrestricted movements during REM sleep could be. By lesioning parts of the brain stem in cats, Jouvet inhibited the muscle paralysis that occurs in many species during REM sleep. Cats that had gone through the procedure acted normally when awake, but when asleep they became unusually active, exhibiting intermittent bursts of activity such as prowling, swatting, biting, playing and grooming. Despite this remarkably awakelike behavior, the cats remained fast asleep. Jouvet observed that the cats' sleeping actions often were unlike their waking habits. Felines that were "always very friendly when awake," he wrote, behaved aggressively during REM sleep.

In the late 1980s Carlos Schenck, a psychiatrist at the University of Minnesota, and his colleagues published the first case reports of RBD. Patients described having violent dreams and aggressive sleep behaviors that contrasted sharply with their nonviolent nature while awake—echoing Jouvet's documentation of otherwise friendly felines that turned belligerent during sleep. One patient, for example, said he had a dream about a motorcyclist trying to ram him on the highway. He turned to kick the bike away—and woke to his wife saying, "What in heavens are you doing to me?" because he was "kicking the hell out of her." Another said he dreamed of breaking a deer's neck and woke up with his arms wrapped around his spouse's head.

To test whether these bizarre behaviors may reflect damage to the brain stem, as in Jouvet's cats, Schenck and his colleagues kept track of such patients to see whether they might develop a brain disease. In 1996 they reported that in a group of 29 RBD patients, all of whom were male and age 50 or older, 11 had developed neurodegenerative disease an average of 13 years after the onset of their RBD. By 2013, 21 of them, or more than 80 percent, had developed a neurodegenerative condition—the most common of which was Parkinson's.

Subsequent studies confirmed this link. Of 1,280 patients across 24 centers around the world, 74 percent of people with RBD were diagnosed with a neurodegenerative disease within 12 years. Sometimes RBD shows up decades before other neurological symptoms, although the average lag appears to be about 10 years. When dream enactment occurs alongside other early signs of synucleinopathies, people tend to develop a neurodegenerative disease more rapidly.

Many researchers expressed skepticism about this link early on, says Bradley Boeve, a professor of neurology at the Mayo Clinic in Rochester, Minn. "We would get reviewer comments back saying that this is hogwash," he says. But the connection between RBD and synucleinopathy has become well accepted: "I think that's pretty much gospel now."

Some scientists suspect RBD results from an aggregation of synuclein and associated neurodegeneration in areas of the brain stem that immobilize us during REM sleep. In its normal, benign, form, the protein is involved in the functioning of neurons, but when "misfolded" into an atypical configuration, it can form toxic clumps. Autopsies have shown that more than 90 percent of people with RBD die with signs of synuclein buildup in their brains. There are no established methods to probe for synuclein clusters in the brains of living patients, but scientists have looked for the toxin in other parts of the body. Alejandro Iranzo, a neurologist at the Hospital Clinic Barcelona in Spain, and his colleagues were able to detect misfolded synuclein in the cere-

Source: "Across the Consciousness Continuum—From Unresponsive Wakefulness to Sleep," by Christine Blume et al., in *Frontiers in Human Neuroscience*, Vol. 9; March 2015 (*reference*)

Graphic by Jen Christiansen

brospinal fluid of 90 percent of patients with RBD.

As an early manifestation of Parkinson's and related diseases, RBD can help scientists trace the ways in which toxic synuclein spreads throughout the body and brain. Evidence is mounting that at least in some patients, pathology may begin in the gut and spread up through lower brain structures such as the brain stem to the higher regions influencing movement and cognition. One likely pathway is the vagus nerve, a bundle of nerve fibers connecting all the major organs with the brain. Alpha-synuclein clumps injected into the guts of mice can spread to the brain via the vagus—and in humans, at least one epidemiological study has shown that cutting the vagus, a procedure sometimes used to treat chronic stomach ulcers, decreases the risk for Parkinson's later in life.

Some researchers hold that Parkinson's has two subtypes: gut first and brain first. RBD is highly predictive of later Parkinson's, says Per Borghammer, a professor of clinical medicine at Aarhus University in Denmark, but the converse is not true: only about a third of people with Parkinson's get RBD before developing motor symptoms. People with RBD have gut-first Parkinson's, Borghammer posits, and generally experience symptoms such as constipation long before motor and cognitive decline. But in the two thirds of patients who are brain first, RBD may emerge later than problems with movement—or never appear.

THE DREAM THEATER

DOES DAMAGE to the brain stem also affect the content of dreams and the actions of dreamers? Sleep researcher Isabelle Arnulf, a professor of neurology at Sorbonne University in Paris, developed a keen interest in the dream-time behaviors of her Parkinson's patients after noticing an unusual pattern: although these people struggled with movement while awake, their spouses often reported that they had no trouble moving while asleep. One particularly memorable patient, according to Arnulf, had been dreaming of crocodiles in the sleep lab when he lifted a heavy bedside table above his head and loudly shouted, "Crocodile! Crocodile!" to an empty room. When awake, he struggled to lift objects and to speak.

Intrigued by such observations, Arnulf and her colleagues began compiling the behaviors people exhibited during REM sleep. This collection, which has grown over the past decade and a half to include hundreds of hours of footage of dream-enacting sleepers and hundreds of dream reports, has enabled Arnulf to uncover unexpected features of RBD dreams and insights into some fundamental questions about how—and why—we dream.

Merrell, Alda and many other people with RBD often have dreams in which they face danger. In one study led by Arnulf, researchers found that among people with RBD, 60 percent reported dreams involving some kind of threat, and 75 percent confronted their attacker instead of running away. People who report more frequent distressing dreams are also at greater risk of developing Parkinson's. "It's textbook for people with RBD to have violent dreams where they are on the defensive," says Yo-El Ju, a professor of neurology at Washington University School of Medicine in St. Louis. But whether this is attributable to recall bias—people tending to remember more violent dreams because they are more memorable—remains an open question, she adds.

Arnulf's group also found that a significant proportion of RBD dreams are nonviolent. In one study, 18 percent of patients flew, sang, danced, laughed, lectured or enacted other peaceable activities. In another study with 52 RBD patients, the researchers looked at subtle changes in facial expressions during sleep. Half the people smiled and a third laughed during mainly REM sleep, suggesting that RBD dreams may be more positive than previously described. Arnulf hypothesizes that violent dreams may be reported more often because aggressive behaviors are more likely to wake up the dreamer or their spouse. "I'm pretty convinced that in RBD patients, it's just that the window is open on dreaming, but their dreams are not different from ours," Arnulf says.

In people with the dream behavior disorder called RBD, the brakes that normally immobilize them during REM sleep are lifted.

The finding that RBD patients display a range of emotions while dreaming led Arnulf to believe that what researchers learn about their dreams may apply to the broader population. Her team discovered, for example, that a small percentage of people with RBD were never able to recall their dreams despite acting out dreamlike behaviors while asleep—suggesting that self-described nondreamers may, in fact, dream.

One mystery of RBD is whether people are acting out their dreams or whether their movements are modifying their dream narratives, says Birgit Högl, a professor of neurology and sleep medicine at the Medical University of Innsbruck in Austria. As for the question that originally intrigued Arnulf—why the impaired movement characteristic of Parkinson's seems to disappear during sleep in some patients—work by other groups has helped suggest an answer. Neurologist and psychiatrist Geert Mayer, formerly at Hephata Clinic in Germany, and his colleagues revealed in a 2015 study that the basal ganglia, movement-related structures near the base of the brain where neurodegeneration occurs in people with Parkinson's, were silent during dream enactments in RBD patients. But other brain regions involved in producing

Where Does Parkinson's Begin?

For decades neurologists saw Parkinson's primarily as a disease caused by the progressive loss of neurons in the substantia nigra, a brain region involved in movement. In the early 2000s neuroanatomist Heiko Braak of the University of Ulm in Germany and his colleagues proposed that the ailment may begin in the gastrointestinal tract, an idea that has gained support in recent years. Some researchers hold that Parkinson's has two subtypes, one originating in the gut (*left*) and the other in the brain (*right*). In this view, enacted dreaming precedes Parkinson's when the disease travels from the gut to the brain.

In the Gut

This model posits that toxic clumps of a protein called alpha-synuclein begin aggregating in the gut, then travel through the highways of nerve fibers that help to control heart rate, digestion, and other bodily functions. Eventually they spread to the brain stem, leading to RBD, and years later reach higher brain regions, affecting movement, cognition and other processes.

In the Brain

In this model, alpha-synuclein builds up in the amygdala, substantia nigra, and other parts of the brain, affecting emotion, movement and cognition, and then descends through the brain stem. So RBD develops after the onset of more characteristic symptoms of Parkinson's.

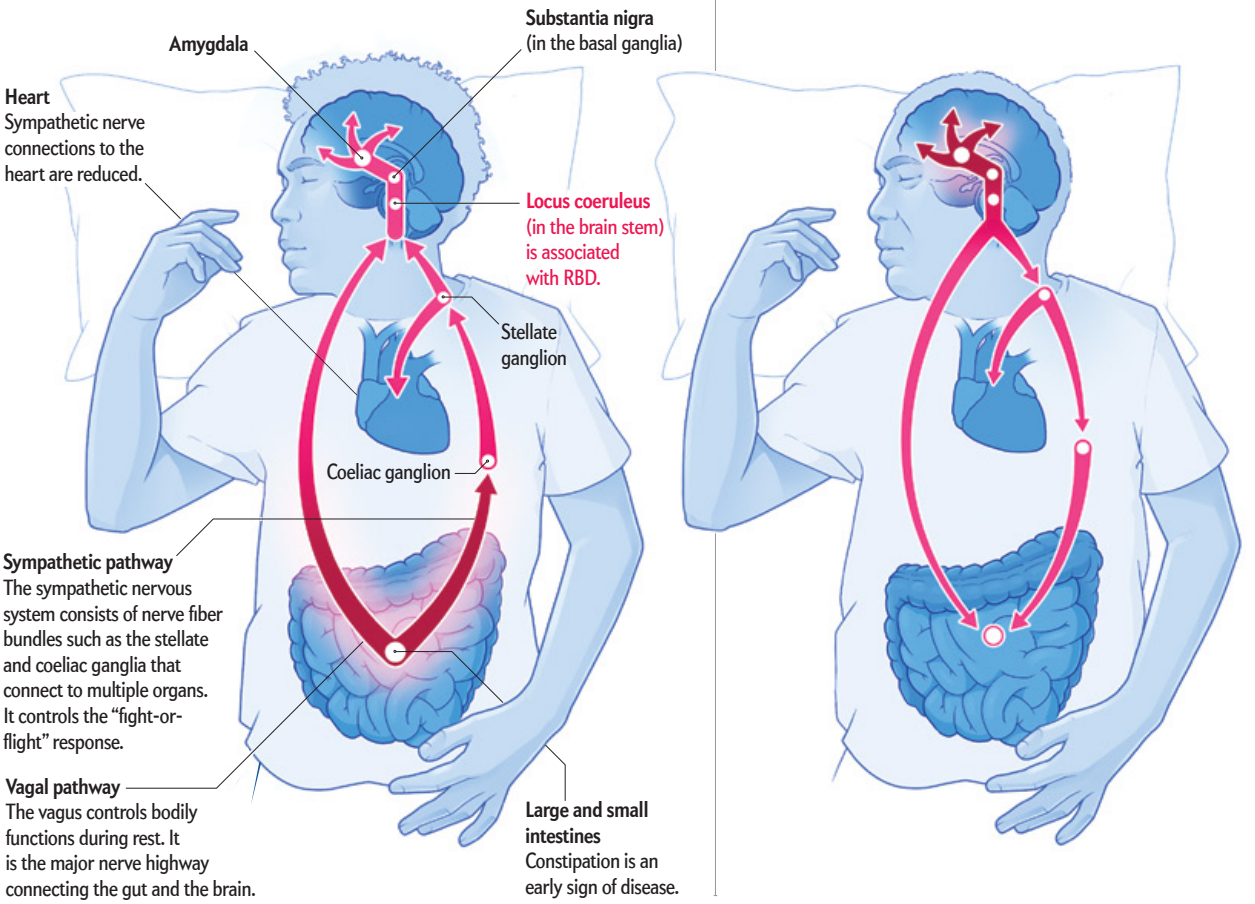

movement, such as the motor cortex, were active.

Findings such as these suggest that in people with RBD, movement is generated through a motor circuit that bypasses the basal ganglia. "This sort of shows that whatever's going on in Parkinson's disease in terms of your movement doesn't apply to you when you're asleep," says Ronald Postuma, a professor of neurology at McGill University. It also raises a tantalizing possibility for therapy: "What if you could mimic whatever that motor state is when a person is asleep but keep them otherwise awake?"

EARLY INTERVENTION

MERRELL HAD BEEN enacting dreams for several years before he realized it might be a sign of a bigger problem. It began during a rough patch at work, and he had dismissed the occasional sleep outbursts as deriving from job-related stress. One evening, mid-dream, Merrell threw himself into a corner of a nightstand, breaking his skin but narrowly missing his breastbone. The close call with a very serious injury "really got me thinking that I better look into this," he says.

Source: "Prodromal Parkinson Disease Subtypes—Key to Understanding Heterogeneity," by Daniela Berg et al., in *Nature Reviews Neurology*, Vol. 17; April 2021 (*reference*)

Graphic by Now Medical Studios

When Merrell was diagnosed with RBD in 2011, his doctor briefly mentioned the risk of developing other conditions down the line but "didn't give me assurances or any other advice," Merrell recalls. But when he began researching the condition online, he discovered many studies on RBD patients who developed a neurodegenerative disease in later life. "The more I searched," he adds, "the more I realized, wow, this has some pretty significant implications."

The available treatments for Parkinson's and other synucleinopathies can currently only manage symptoms. They're unable to slow or stop the underlying neurodegeneration. "The worst news I have to give as a sleep doctor is to tell someone that they have RBD," Ju says.

But several new therapeutics for Parkinson's and other synucleinopathies are being developed, and many neurologists believe early intervention could be crucial. "The Parkinson's disease field, in particular, is full of failed treatment trials," Ju says. "By the time people have the disease, it's probably too late to intervene—too many cells have died." Going back and testing these seemingly failed medications in RBD patients may prove more successful, she adds, because as a much earlier stage of disease, RBD provides a window where treatments are more likely to be effective. "A lot of people are viewing RBD as similar to high cholesterol," Boeve says. "If you have high lipid levels, they increase your risk for heart disease and stroke. If you can alter that pathophysiological process, you can reduce the risk or delay the onset."

Ju, Postuma and Boeve are co-leaders of the North American Prodromal Synucleinopathy (NAPS) Consortium, which launched in 2018. The NAPS investigators aim to pinpoint clinical and biological markers through various means, including brain scans, genetic screens, and tests of blood and cerebrospinal fluid. The researchers hope these markers will eventually indicate how and when a person with RBD will develop a neurodegenerative disease later in life—and which disease they will end up with. Ideally, such biomarkers would help scientists identify RBD patients for investigative therapies that target α-synuclein years before debilitating symptoms appear. The ultimate goal of NAPS, Ju says, "is essentially to prepare for clinical trials for protective treatments."

In 2021 NAPS received a $35-million grant from the National Institutes of Health for this work, which will be carried out across eight sites in the U.S. and one in Canada. In a parallel effort, Högl, along with other researchers in Europe, is gathering a similar cohort of patients from multiple institutions across the continent for future clinical studies. Wolfgang Oertel, a neurologist at Philipps University of Marburg in Germany, who is involved in the European effort, is optimistic about the future for people with RBD. He expects that of the dozens of potentially disease-modifying Parkinson's drugs currently in clinical trials, at least a few will be available soon. "I tell my patients, 'You've come at the right moment,'" Oertel says. "You will be one of the first to get the right drugs."

Högl has also been involved in another active area of investigation: finding ways to better characterize RBD. Working with Ambra Stefani of the Medical University of Innsbruck and other colleagues, she has been gathering measurements of muscle activity during sleep in people with RBD. They hope that this work will not only help to streamline the diagnosis of RBD but also help doctors to detect the sleep disorder even earlier, in so-called prodromal RBD, where overt dream enactments might not occur, or in people who may have RBD but exhibit only small, difficult-to-detect movements. Their work suggests that the elaborate, violent behaviors seen in RBD are "just the tip of the iceberg," Högl says. They may occur on one night but not another. Minor muscle jerks in the hands or elsewhere, in contrast, appear to be much more frequent—and a more stable sign because they occur hundreds of times during the night, she adds.

Often appearing as a much earlier stage of Parkinson's disease, RBD provides a chance to intervene when treatments are more likely to be effective.

For now there is no cure for RBD or Parkinson's—but that doesn't mean there is nothing patients can do. A growing body of evidence indicates that moderate to intense exercise helps to improve both motor and cognitive symptoms of Parkinson's, and many neurologists already recommend such physical activity to their patients with RBD. "The evidence suggests that the benefits of exercise are more than just symptomatic," says Michael Howell, a neurologist at the University of Minnesota. "It appears that this actually is helping to protect brain cells."

Both Alda and Merrell have taken that advice to heart. In addition to medications, Alda has taken up exercise-based therapy for Parkinson's. Merrell, too, has integrated regular physical activity into his routine, hiking for several miles every other day. He's gotten involved in clinical research and is one of the NAPS participants. This contribution helps Merrell feel empowered—he hopes to aid the discovery of effective neuroprotective therapeutics. "Somebody always had stepped up in other illnesses or conditions that allowed for clinical trials and the therapies that we have today," Merrell says. "I just happened to be queued up for this—and I accept that challenge." SA

Diana Kwon is a freelance journalist who covers health and the life sciences. She is based in Berlin.

Mental Illness and Dementia

Why do psychiatric conditions multiply the risk of cognitive decline?

By Claudia Wallis

AGE IS THE SINGLE BIGGEST FACTOR for dementia, with a person's risk doubling about every five years after age 65. But many things influence those odds for a given individual. Genetic vulnerability is a contributor, as are so-called modifiable risk factors such as smoking, cardiovascular disease, social isolation, and impaired hearing and vision. Certain mental conditions, particularly depression and schizophrenia, have also been linked to dementia. But because depression can itself be a sign of cognitive decline, the causality has been a bit muddy. In 2022 an analysis of data from New Zealand provided some of the most convincing evidence to date linking many kinds of mental illness with dementia. The study raises important questions about the reasons for this increased risk and what could be done to reduce it.

The study looked at the health records of 1.7 million New Zealanders born between 1928 and 1967 covering a 30-year period ending in mid-2018. It found that those with a diagnosed mental disorder—such as anxiety disorders, depression or bipolar disorder—had four times the rate of ultimately developing dementia compared with people without such a diagnosis. For those with a psychosis such as schizophrenia, it was six times the rate. Among people who developed dementia, those with a psychiatric disorder were affected 5.6 years earlier, on average.

The study did not examine biological, social or other reasons for the increased risk, but research on dementia points to several possible explanations. "There might be shared genetic risk factors," suggests psychologist Leah Richmond-Rakerd of the University of Michigan, lead author of the study. Other research in recent years has found some overlap in genetic markers associated with Alzheimer's disease and those linked to bipolar disorder and to major depression. Long-term use of psychiatric medications could also be playing a role in dementia, but Richmond-Rakerd and her co-authors do not think it is a major contributor.

They suspect that a more significant risk factor is the chronic stress associated with having a psychiatric disorder, which may degrade brain health over time. Both studies in animals and human autopsy studies have linked chronic stress to a loss of neural connections in the hippocampus, the brain's memory center, which is where Alzheimer's takes a heavy toll. Evidence suggests that stress drives inflammation and immune dysregulation in the body and brain, impacting brain connectivity, says Harvard University neurologist and dementia researcher Steven Arnold. "If you have fewer connections and synapses to begin with because of stress, then you can't afford to lose as many with aging before it starts to show up as what we might call dementia." In other words, people with mental illnesses may have less "cognitive reserve"—brainpower that is sufficiently robust to withstand normal aging without obvious losses of function.

Vulnerability in this population may also be linked to their finding it more difficult to lead healthy lives, physically and socially, Richmond-Rakerd says. "They

Illustration by Fatinha Ramos

might exercise less, or drink alcohol excessively, or have trouble staying socially connected"—all of which increase the risk for dementia. People with certain psychiatric conditions tend to have higher-than-average rates of smoking and fewer years of education, which are also risk factors.

Could a more holistic approach to treating mental illness mitigate the risk for dementia? Researchers tend to think so. In 2020 the British-based *Lancet* Commission on dementia prevention estimated that four in 10 cases could be prevented or delayed if society did a better job of addressing 12 modifiable risk factors, including such psychosocial contributors as depression, poor social support and low education level. Progress on some of these factors may explain why dementia rates have already fallen 15 percent per decade for the past 30 years in high-income countries. "We think there are two main reasons: better cardiovascular risk factor control and a big increase in education level," says Kenneth Langa of the University of Michigan, associate director of the Health and Retirement Study, one of the major efforts tracking these trends.

In an ideal world, Langa and other researchers say, efforts to prevent dementia would begin in childhood with strong investments in education and the inculcation of healthy habits. Such efforts would be incorporated into the treatment of depression and other mental illnesses that often emerge in the teen and early adult years. Sadly, we do not live in that ideal world; mental illness continues to be stigmatized and undertreated. But given the high costs to society and the personal tolls exacted by both mental illness and dementia, it's hard to imagine a wiser investment. SA

Claudia Wallis is an award-winning science journalist whose work has appeared in the *New York Times, Time, Fortune* and the *New Republic*. She was science editor at *Time* and managing editor of *Scientific American Mind*.

Aneurysm in a Dish

Scientists operate on a 3-D-printed model of a ballooning blood vessel

By Tanya Lewis

BRAIN ANEURYSMS, WHICH AFFECT AS MANY as one in 50 people, occur when a blood vessel wall weakens and bulges, setting the scene for a potentially deadly rupture. Scientists have created a 3-D-printed aneurysm model in the laboratory and "operated" on it: they inserted a device to seal it off and prevent it from bursting. Such models could be tailored to replicate an individual patient's blood vessel, letting doctors try different treatments and find the best solution.

To treat an aneurysm, brain surgeons sometimes operate to install a metal clip on the ballooning vessel that prevents the pooling of blood. A less invasive method involves inserting tiny metal coils into the aneurysm via a catheter to induce a blood clot that seals it off. Most treatment devices are tested in animals, whose blood vessels do not perfectly resemble those of humans. And previous lab-dish aneurysms could not mimic the properties of living blood vessels. "We thought maybe there could be a better way of testing those [treatment] devices," says Lindy Jang, who led the research when she was a biomedical engineering graduate student at Texas A&M University. The study was published in 2020 in *Biofabrication*.

Jang and her colleagues 3-D-printed an aneurysm structure that had a water-based gel and populated it with human cells that line the brain's blood vessels. The researchers operated on the aneurysm, injecting platinum coils into the bulging vessel. Finally, they filled the blood vessel with plasma (the liquid component of blood), which then formed a clot that sealed off the bulge.

"We're trying to streamline the treatment of aneurysms and take the guesswork out," says William "Rick" Hynes, a study coauthor and biofabrication research engineer at Lawrence Livermore National Laboratory, who performed the surgery. "The goal is to use these devices to validate models so someone could take a 3-D scan, recreate it in the simulation, then try adding [blood] flow and determine if they need to treat the aneurysm or leave it alone."

"I think it's really significant," Matthew Gounis, a biomedical engineer at the University of Massachusetts Medical School, says of the new model. Other groups have developed aneurysm models, but this one is exciting because it better replicates a human blood vessel by adding living cells, says Gounis, who was not involved in the study. Surgeons could practice on such models before operating on a real patient, he says: "If you have a particularly challenging case, you can print out the case, and you can basically practice before you get to the patient, in their anatomy." SA

Tanya Lewis is a senior editor covering health and medicine for *Scientific American*.

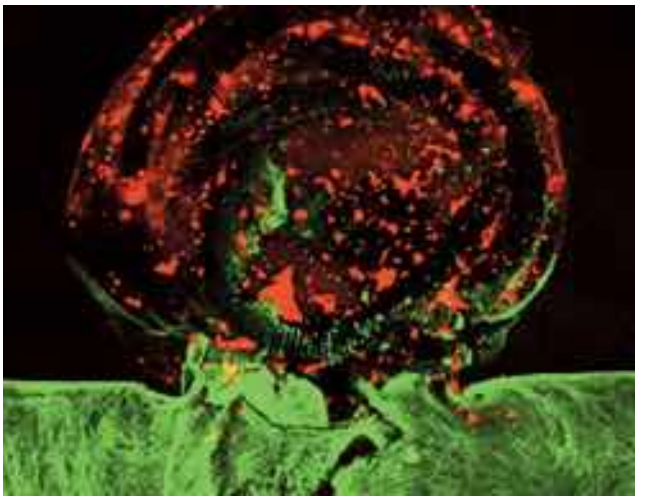

Elisa Wasson/Lawrence Livermore Laboratory

AN INDUCED BLOOD CLOT (*red balloon*) is visible in a 3-D-printed model of a brain aneurysm.

Deep Sleep Gives Your Brain a Deep Clean

Slow-wave activity during dreamless slumber helps to wash out neural detritus

By Simon Makin

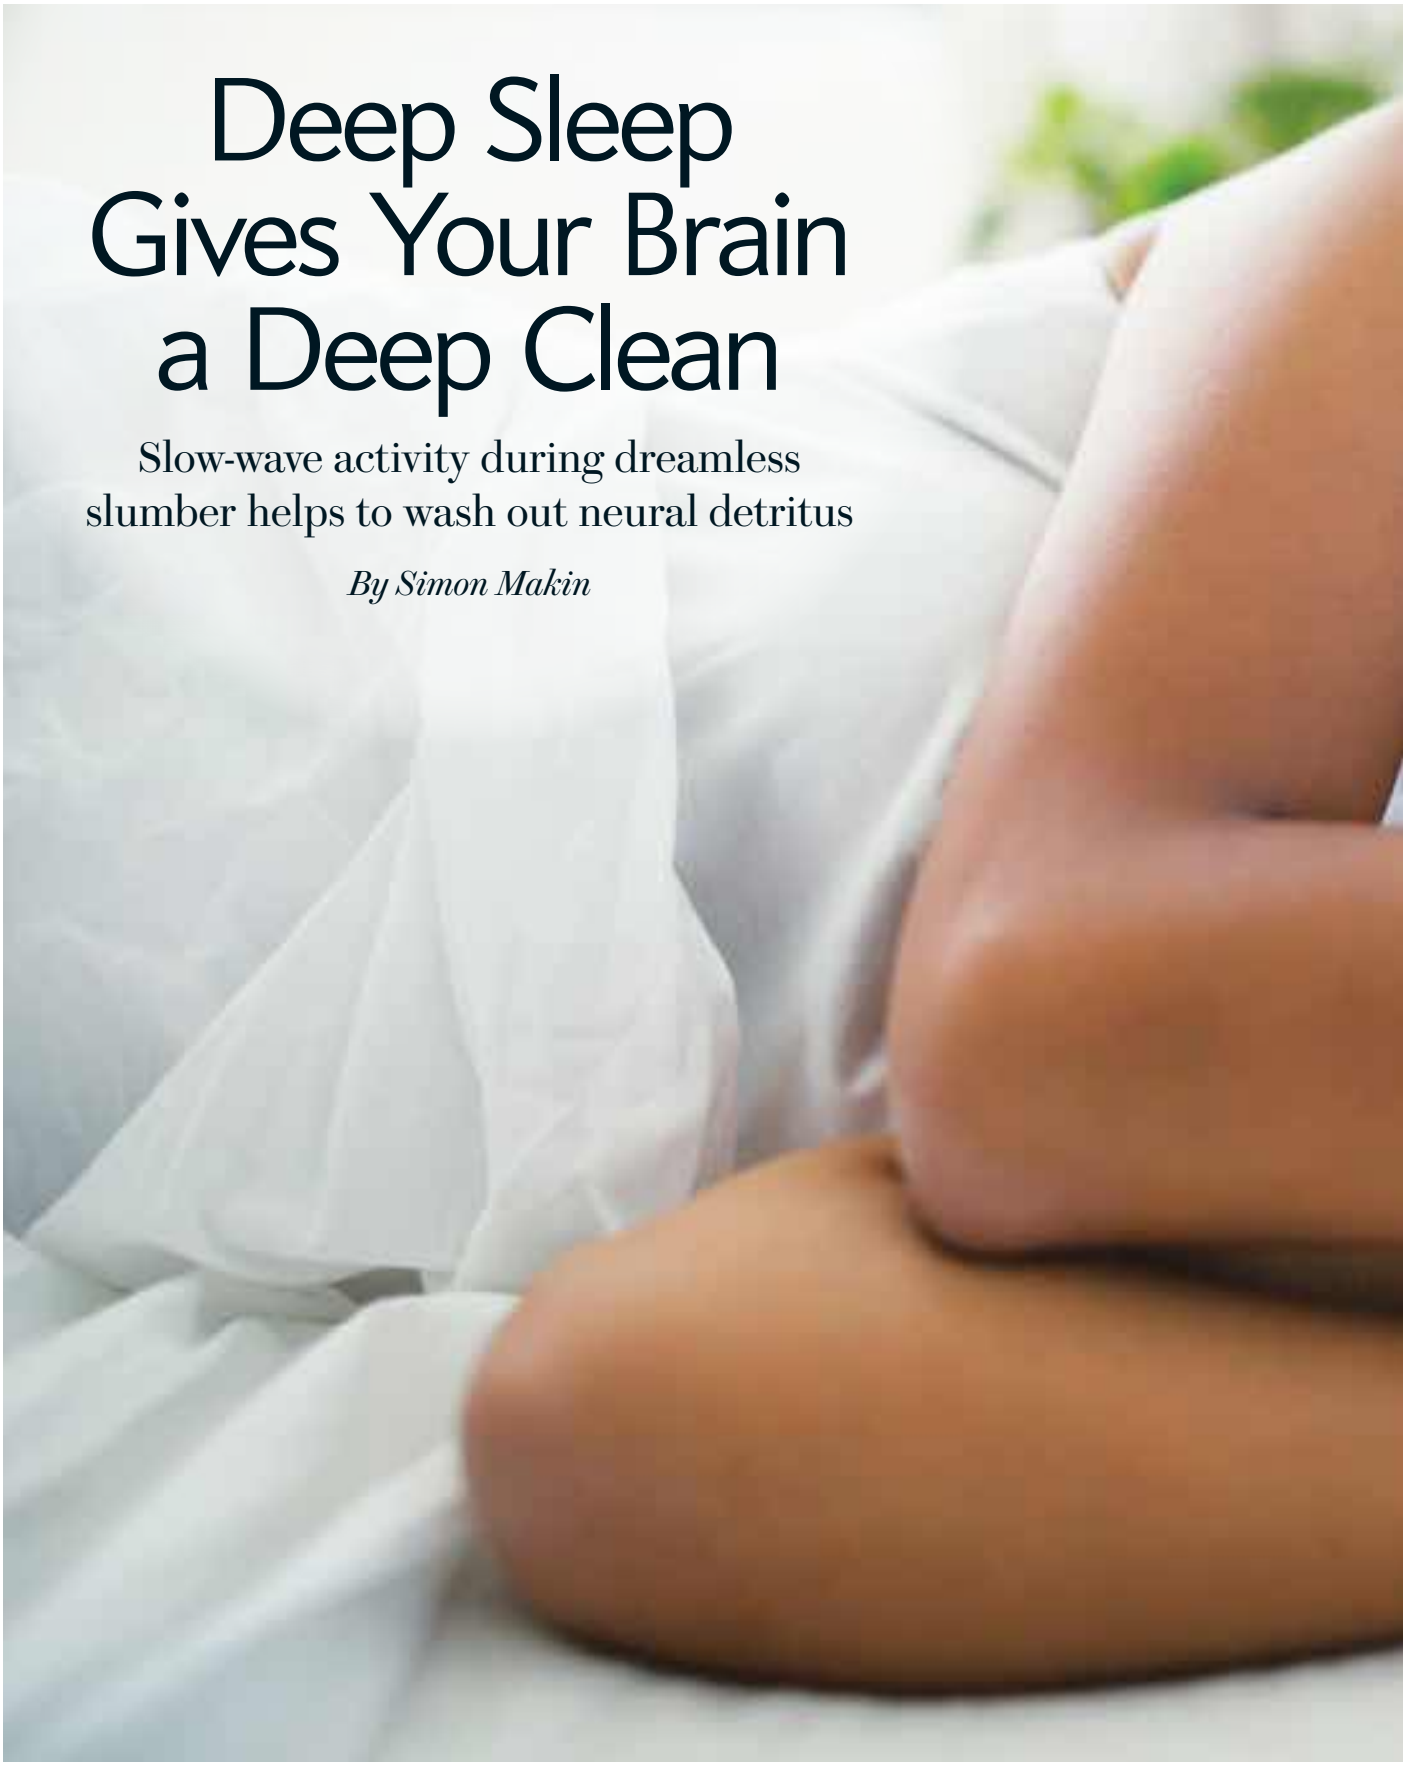

Westend61/Getty Images

WHY SLEEP HAS RESTORATIVE—OR DAMAGING—EFFECTS ON COGNITION AND BRAIN HEALTH

has been an enduring mystery in biology. Researchers think cerebrospinal fluid (CSF) may flush toxic waste out, "cleaning" the brain, and studies have shown that garbage clearance is hugely improved during sleep. Scientists were not sure exactly how all this works, however, or why it should be so enhanced during sleep.

One aspect of sleep that is well understood is how the slow electrical oscillations (or "slow waves") that characterize deep, non-REM sleep contribute to memory consolidation, the process whereby new memories are transferred into long-term storage. A 2019 study, from a team led by neuroscientist Laura Lewis of Boston University, gives insight into what drives CSF flow through the brain, suggesting that the same slow waves that coordinate memory consolidation drive oscillations in blood flow and CSF in the brain.

The work has implications for understanding the relations between sleep disturbance and psychiatric and neurodegenerative conditions, and it may even point to new approaches to diagnosis and treatment. "We've discovered there are really large waves of CSF that appear in the brain only during sleep," Lewis says. "This effect is really striking, and we're also interested in what it means for maintaining brain health, especially in disorders such as Alzheimer's disease."

In the study, published in *Science,* the team set out to investigate how the dynamics of CSF flow change during sleep and how such changes might relate to alterations in brain blood flow and electri-

A Symphony in Two Movements

Dramatic differences characterize two key sleep phases. The slow waves of deep sleep dominate the early part of the night. During slow-wave sleep, some memories spontaneously reactivate. Interventions that promote this process can ensure that memories are retained. Rapid eye movement (REM) sleep prevails in the latter part of a night's slumber, but how it interacts with memory remains controversial.

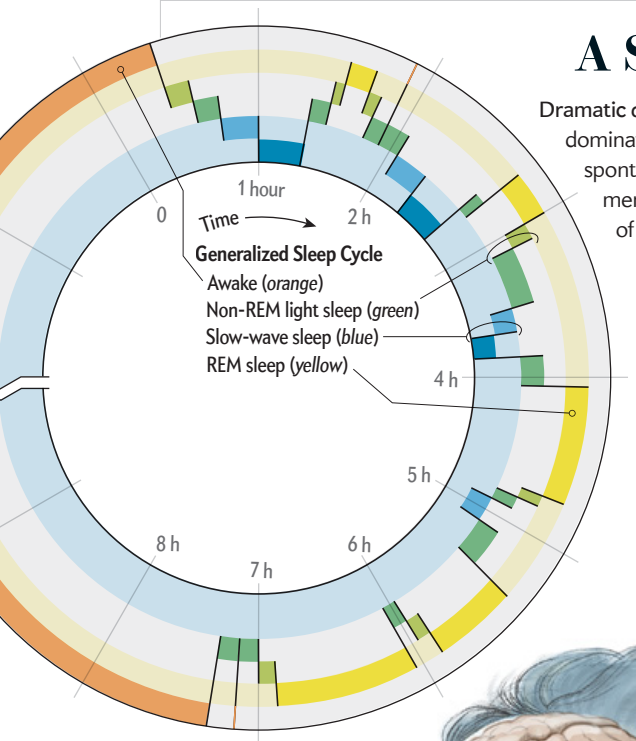

Harmonizing Brain Waves

Brain oscillations during sleep appear to play a role in strengthening new memories. A key event is the "up" phase of a slow oscillation that coordinates the activity of other brain rhythms. The ascending part of a slow oscillation in the cortex synchronizes with sleep spindles in the thalamus. The spindles coordinate the activity of sharp-wave ripples in the hippocampus. Ripples tend to coincide with a spindle trough.

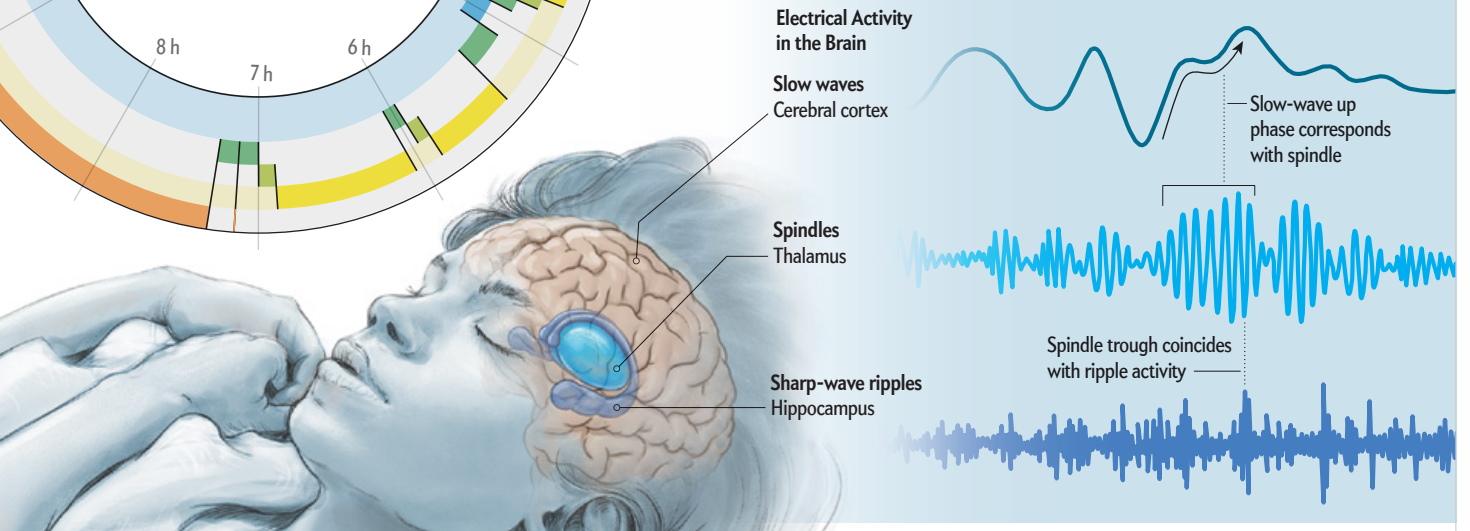

Graphic by Mesa Schumacher

cal activity. "We know sleep is really important for brain health, and waste clearance is probably a key reason why. What was less clear is: Why is this changed during sleep?" Lewis says. "That led us to ask what was happening in the CSF."

The researchers used electroencephalography (EEG) to monitor the brain waves of 13 sleeping healthy adults, while also using a cutting-edge, "accelerated" functional magnetic resonance imaging technique to capture faster changes than standard fMRI can manage. That allowed them to measure both blood-oxygenation changes (which indicate blood flowing to electrically active, oxygen-hungry regions) and CSF flows. The latter was possible only because of a flaw in this method that means any newly arriving fluid (not just oxygenated blood) lights up in the image. "We realized we could take advantage of this to measure CSF flow at the same time as blood oxygenation," Lewis says. "That was critical because it turns out these things are coupled to each other in a way we never would have seen if we didn't measure blood, CSF and electrical activity simultaneously."

What the team found was that the slow waves seen in non-REM sleep occur in lockstep with changes in both blood flow and CSF. Just because things occur together doesn't necessarily mean one causes the other, but the team also built a computer model incorporating information about the physics linking these processes, which predicted that slow waves would have just these kinds of effects on blood and CSF. What seems to be happening is that as brain activity alters blood flow, the volume of blood in the brain decreases, and because the brain is a closed vessel, CSF flows in to fill the space. "It's very convincing," says neurologist Maiken Nedergaard of the University of Rochester, who was not involved with the research. "It also really makes sense: electrical activity drives blood flow changes that then drive CSF changes."

The team measured this CSF inflow going into the fourth ventricle, one of four fluid-filled cavities involved in producing CSF (by filtering blood plasma) and circulating it around the brain. As CSF usually flows out of the fourth ventricle, this suggests a "pulsatile" flow, like a wave. This pushes CSF around the ventricles and into spaces between membranes surrounding the brain and spinal cord, called the meninges, where it mixes with "interstitial fluid" within the brain to carry away toxic waste products.

Because slow waves are important for memory consolidation, this links two disparate functions of sleep. "What's exciting about this is it's combining features of brain function that people don't normally think of as connected," Nedergaard says. It isn't obvious things had to be this way, Lewis says, but it may represent an example of nature being efficient. "It's a matter of nature not dividing tasks between higher level and lower level, like how you run a company, where you have a boss making decisions and cleaning people coming in," Nedergaard says. "In biology, it's everybody contributing, as it makes more sense."

The findings have implications for neurodegenerative diseases, which might be caused by buildup of toxic proteins in the brain, such as beta-amyloid in Alzheimer's disease. Previous research has shown that beta-amyloid is cleared more efficiently during sleep, which is often disrupted in patients. Disturbances in slow-wave sleep also often accompany aging, which may be linked to cognitive decline. "We know that people with Alzheimer's have fewer slow waves, so we may find they also have fewer CSF waves," Lewis says. "We have to do these studies now in older adults and patient populations to understand what this might mean for those disorders." Sleep disturbance is also a feature of many psychiatric disorders, from depression to schizophrenia. "Different electrical signatures of sleep are disrupted in different psychiatric conditions," she says. "So this will be very interesting to follow up on in a multitude of disorders."

> "We know sleep is really important for brain health, and waste clearance is probably a key reason why."
> —Laura Lewis *Boston University*

The team has talked about the possibility of modulating brain activity to try to nail down whether electrical oscillations truly do cause the changes they observed in CSF flow. "It would be great to find the right collaborator and do a study in mice where we manipulate neural activity, then watch the downstream consequences," Lewis says. "We're also thinking about ways to safely and noninvasively manipulate neural oscillations in humans." It may be possible to use electromagnetic stimulation to influence brain waves as a treatment for brain disorders. Researchers have seen encouraging results of this approach in mice, and these findings may help explain why.

Another potential application may come from assessing whether changes in CSF flows can serve as a diagnostic marker for some of these conditions. "It gives us a ton of interesting new biology to explore and understand since it seems like things the brain is doing during sleep are related to each other in surprising ways," Lewis says. "Maybe the most important take-home message is that sleep is a serious thing," Nedergaard says. "You really need to sleep to keep a healthy brain because it links electrical activity to a practical housekeeping function." SA

Simon Makin is a freelance science journalist based in the U.K. His work has appeared in *New Scientist, The Economist, Scientific American* and *Nature,* among others. He covers the life sciences and specializes in neuroscience, psychology and mental health.

Exercise Pill

Isolated proteins might one day produce health gains without the exertional pain

By Emily Willingham

THE DRUMBEAT OF EXERCISE'S BRAIN BENEFITS MAY sound familiar. Most of us know that getting our move on can mean a boost to mental and neurological health. But what if, through these biochemical processes, we could get all of that brain gain without going through the exercise pain? Mouse experiments have already demonstrated the feasibility of such a shortcut. And there is a hint that the results in rodents could hold for humans as well.

When plasma from well-exercised mice is injected into their idle counterparts, the sedentary rodents have improved memory and reduced brain inflammation. The blood of Olympic athletes is not about to be transfused into the arms of sofa spuds—at least not yet. But people with mild cognitive impairment who exercise for six months show increases in a key protein identified in the runner-mouse plasma. The same protein may be able to whisper its chemical message across the notoriously choosy blood-brain barrier and trigger anti-inflammatory processes in the brain.

These findings, published on December 8, 2021, in *Nature,* offer new details of how exercise benefits the brain and how molecules boosted by physical activity communicate across the organ's strict gatekeeper. The results also hint at a surprising role for the liver and anticlotting systems in these effects and possibly point the way to a futuristic scenario of exercise in a pill—or perhaps a plasma injection.

"Puzzle pieces are coming together," says Saul Villeda, who is an associate professor in the department of anatomy at the University of California, San Francisco, about these hints of multisystem involvement in exercise's effects on the brain. Villeda, who was not involved in the *Nature* study, and his colleagues previously identified a protein in exercised-mouse plasma that refreshed neurons in the aging mouse brain. "We're starting to identify factors in the blood that can target different facets of decline or pathology, and this one really highlights blood factors affecting inflammation in the brain," he says. "The word that keeps popping into my head is 'convergence.'"

On the path to convergence, behavioral scientist Zurine De Miguel, now an assistant professor at California State University, Monterey Bay, and her colleagues at Stanford University and the Department of Veterans Affairs Palo Alto Health Care System first had to let mice exercise. The animals ran their hearts out for 28 days, and then their plasma was transferred to mice that had not touched a running wheel during that time. The recipient animals showed improvements in learning and memory after they had received the "runner plasma." Their brains, in turn, revved up genes that produced proteins that facilitated memory and learning and showed a dampened inflammatory response. When the researchers deliberately induced brain inflammation in the animals, the runner-mouse plasma dialed back that response, too.

The team next looked at what the runner plasma contained. They found increased levels of anticlotting proteins, including one called clusterin, which helps to clear cells of debris. Homing in on this protein, the investigators tested the effects of stripping it from the runner plasma. Brains of sedentary mice receiving clusterin-free plasma showed much less anti-inflammatory activity.

The group also found that clusterin readily attached to the cells that form the blood-brain barrier. When they mimicked the effects of physical activity by injecting the protein into the circulation of mice genetically modified to have neurodegenerative disease, the animals' brain inflammation also declined.

Finally, the researchers wanted to see whether exercise causes clusterin elevation in people. They measured the protein in 20 veterans with mild cognitive impairment before and after six months of structured physical activity and found that their levels increased.

De Miguel notes that in her and her colleagues' study, results differed somewhat between male and female mice. Although anticlotting-protein profiles were similar in the two groups, the females showed more variability. The hormones they make can affect anticlotting factors,

Maridav/Alamy Stock Photo

De Miguel says, and the possibility that some female mice were in a sexually receptive stage during the study might explain this greater variation.

The experiment illustrates a growing recognition of the brain's dependence on assistance from outside the neural no-fly zone. The liver and heart are the most likely sources of clusterin, the authors say. The results implicate both organs as sources of beneficial molecules resulting from physical exercise, De Miguel says. "They all seem to be cross-talking to the brain," she adds.

Villeda says that his group's work with runner plasma in aging mice also implicates the liver. The organ produces an enzyme linked to cognitive improvements in the animals, and the same enzyme was also increased in the blood of older active people. This connection to the liver "was surprising to us because it wasn't usually what you focus on when you think about exercise," he says. With this finding, Villeda adds, "these mechanisms are starting to converge and come into a similar space."

Although physical activity is closely linked to good health, it may be possible to overdo exercise. There are hints that some people who engage often in highly strenuous physical activity may have increased risk for amyotrophic lateral sclerosis. "There is some information out there that says too much exercise can impair some of your immune response and make you susceptible to opportunistic infections," De Miguel says.

How will runner plasma be used as a therapy if these effects in mice are also confirmed in humans? "I have more hope now than when I started my lab because it was difficult to think about identifying all of these factors," Villeda says. "But now we have candidates, and when you have those, you can start thinking about small molecule development."

CERTAIN PROTEINS could boost brain function without exercise.

De Miguel says that a possible first step might be testing which exercise protocols trigger the biggest increases in proteins that carry a brain benefit. As with mice, someone in need of the brain-boosting power of physical exercise could simply receive an injection of runner plasma, getting a runner's gain without the accompanying pain. SA

Emily Willingham is a science writer and author based in Texas. Her latest book is *The Tailored Brain* (Basic Books, 2021).

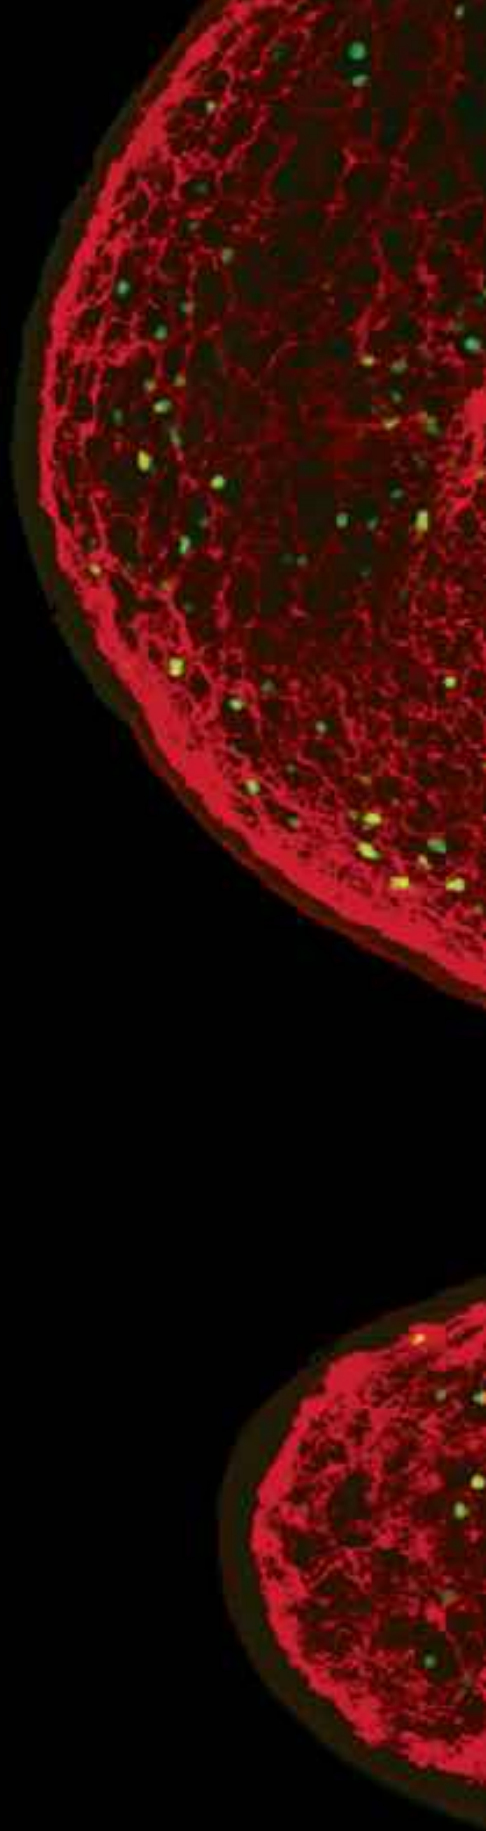

Brain over Body

Scientists are deciphering how the brain choreographs immune responses, hoping to find treatments for a range of diseases

By Diana Kwon

NEURONAL CELLS (*red*) in the gut interface with cells of the immune system (*green*).

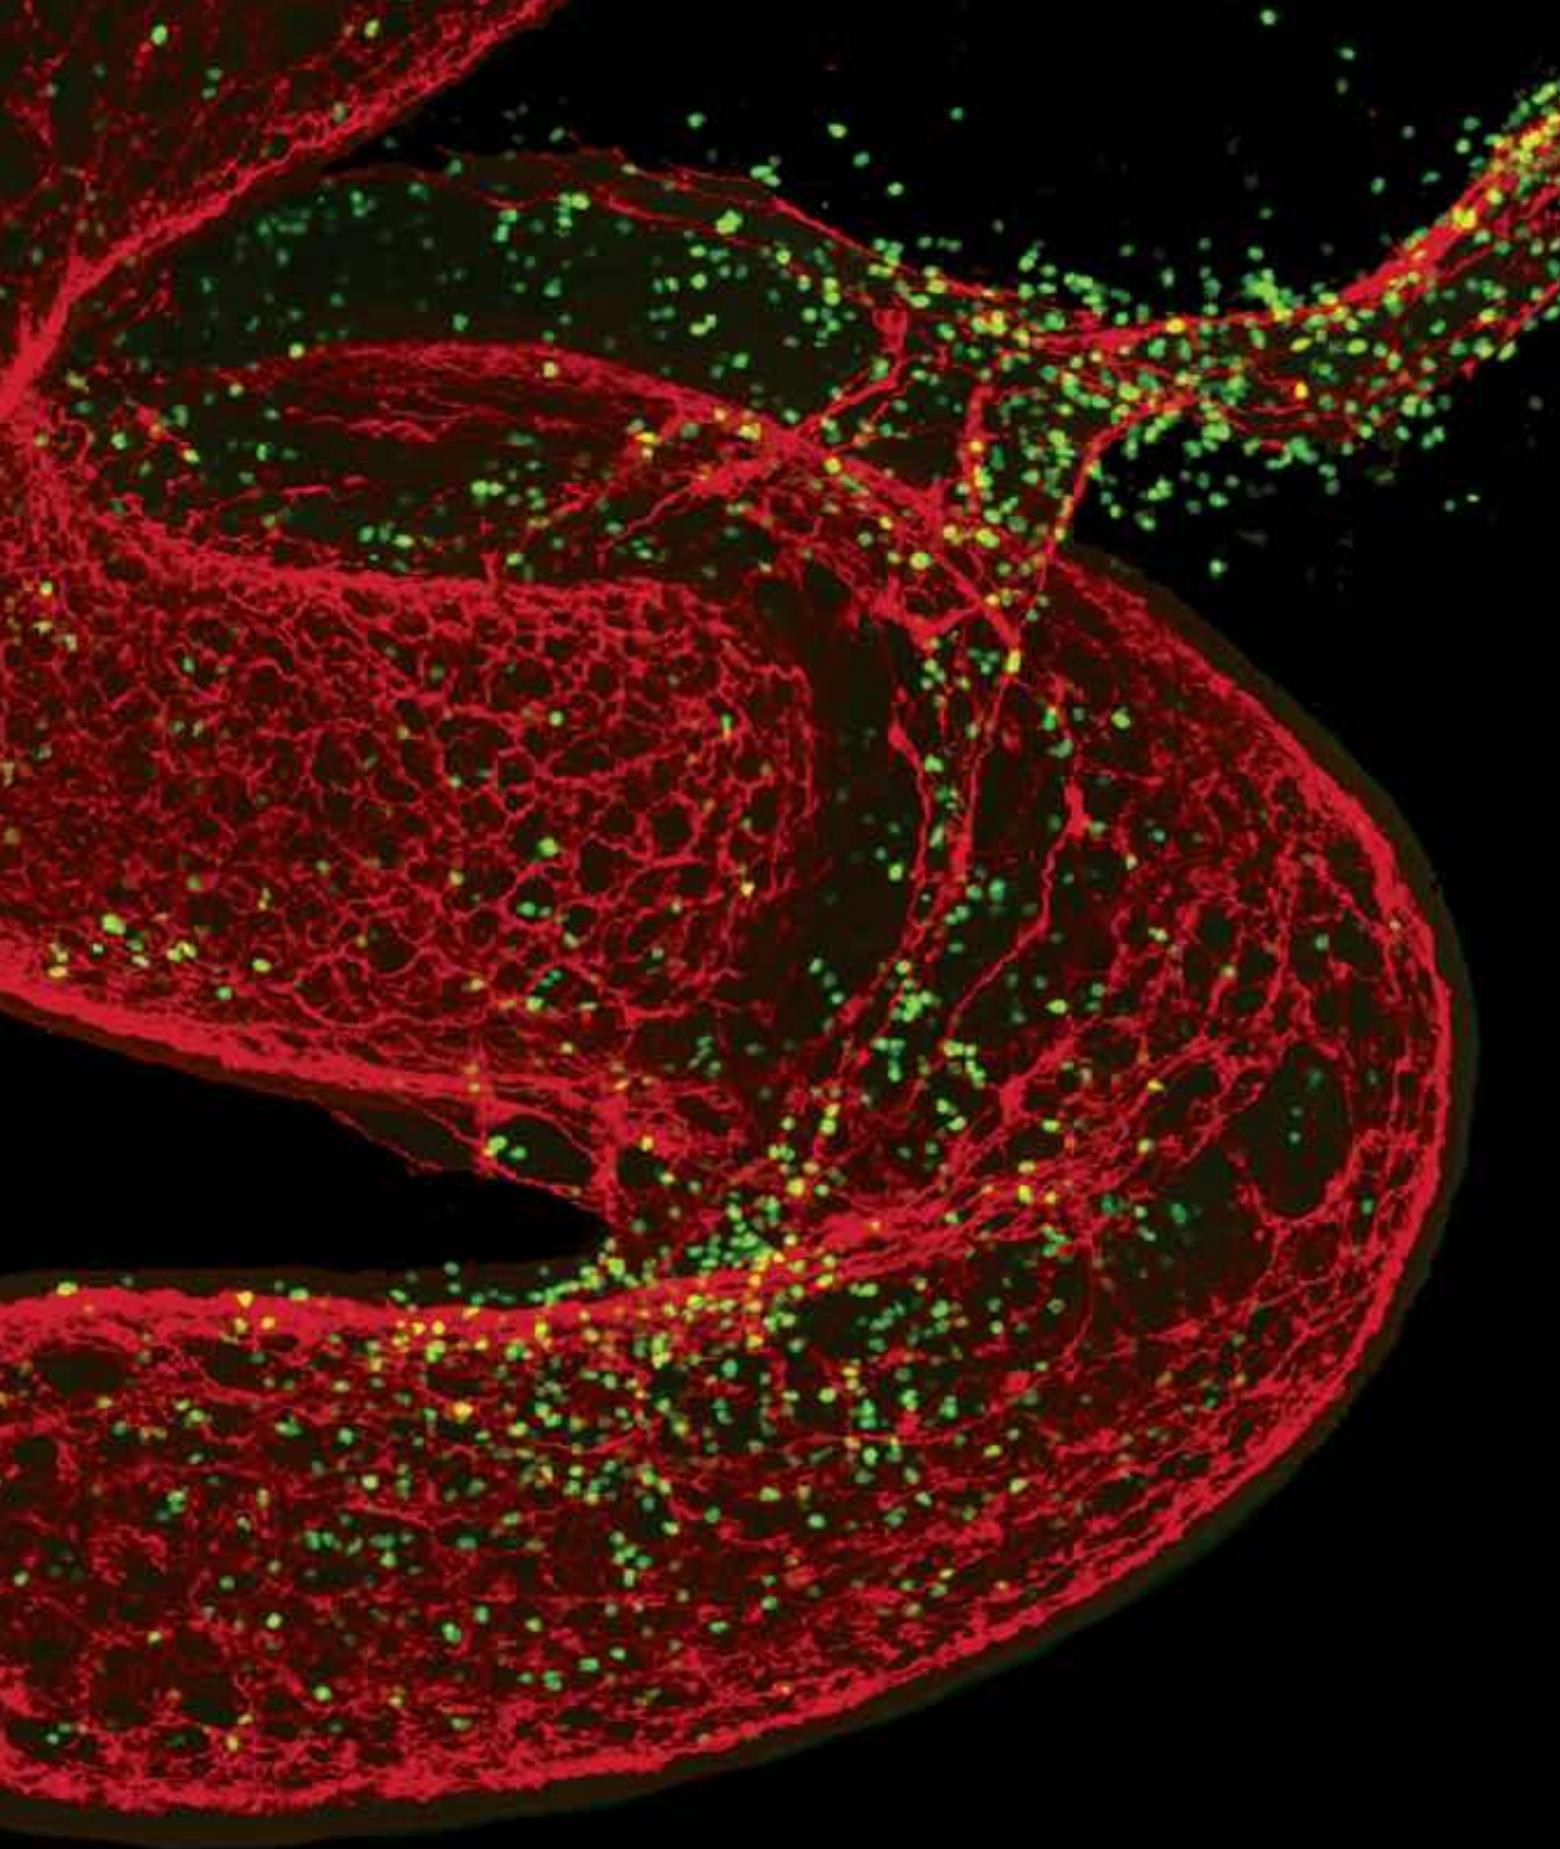

HUNDREDS OF SCIENTISTS AROUND THE WORLD ARE LOOKING for ways to treat heart attacks. But few started where Hedva Haykin did: in the brain. Haykin, a doctoral student at the Technion–Israel Institute of Technology in Haifa, wants to know whether stimulating a region of the brain involved in positive emotion and motivation can influence how the heart heals.

In late 2022, in a small, windowless room, she pulled out slides from a thin black box one by one. On them were slices of hearts, no bigger than pumpkin seeds, from mice that had experienced heart attacks. Under a microscope some of the samples were clearly marred by scars left in the aftermath of the infarction. Others showed mere speckles of damage among streaks of healthy, red-stained cells.

The difference in the hearts' appearance originated in the brain, Haykin explains. The healthier-looking samples came from mice that had received stimulation of a brain area involved in positive emotion and motivation. Those marked with scars were from unstimulated mice.

"In the beginning we were sure that it was too good to be true," Haykin says. It was only after repeating the experiment several times, she adds, that she was able to accept that the effect she was seeing was real.

Haykin, alongside her supervisors at the Technion—Asya Rolls, a neuroimmunologist, and Lior Gepstein, a cardiologist—is trying to work out exactly how this happens. On the basis of their experiments so far, which have not yet been published, activation of this reward center in the brain—called the ventral tegmental area (VTA)—seems to trigger immune changes that contribute to the reduction of scar tissue.

This study has its roots in decades of research pointing to the contribution of a person's psychological state to their heart health. In a well-known condition known as broken heart syndrome, an extremely stressful event can generate the symptoms of a heart attack—and can, in rare cases, be fatal. Conversely, studies have suggested that a positive mindset can lead to better outcomes in those with cardiovascular disease. But the mechanisms behind these links remain elusive.

Rolls is used to being surprised by the results in her laboratory, where the main focus is on how the brain directs the immune response and how this connection influences health and disease. Although Rolls can barely contain her excitement as she discusses her group's eclectic mix of ongoing studies, she's also cautious. Because of the often unexpected nature of her team's discoveries, she never lets herself believe an experiment's results until they have been repeated multiple times—a policy that Haykin and others in her group have adopted. "You need to convince yourself all the time with this stuff," Rolls says.

For Rolls, the implications of this work are broad. She wants to provide an explanation for a phenomenon that many clinicians and researchers are aware of: mental states can have a profound effect on how ill we get—and how well we recover. In Rolls's view, working out the mechanism for this influence could enable physicians to tap into the power of the mind over the body. Understanding the link could help scientists boost the placebo effect, destroy cancers, enhance responses to vaccination and even reevaluate illnesses that for centuries have been dismissed as psychologically driven, she says. "I think we're ready to say that psychosomatic [conditions] can be treated differently."

It wasn't until the late 1990s that neuroscience researchers began drawing connections to the body's master conductor, the brain.

Lara Santos/Veiga-Fernandes Lab (*preceding pages*)

NEUROSCIENTIST Talma Hendler readies a participant for a brain scan. The results will inform a study of whether learning to control brain activity can improve a person's immune response to a vaccine.

Sagol Brain Institute/Tel Aviv Sourasky Medical Center

She is part of a growing group of scientists who are mapping out the brain's control over the body's immune responses. There are multiple lines of communication between the nervous and immune systems—from small local circuits in organs such as the skin to longer-range routes beginning in the brain—with roles in a wide variety of diseases from autoimmunity to cancer. This field "has really exploded over the last several years," says Filip Swirski, an immunologist at the Icahn School of Medicine at Mount Sinai in New York City.

Some parts of the system—such as the vagus nerve, a huge highway of some 100,000 nerve fibers that connects the brain to the heart, lungs, gastrointestinal tract and other major organs—have inspired treatments for several autoimmune diseases that are currently being tested in clinical trials. Other studies investigating ways to recruit the brain itself—which some think could provide powerful therapies—are still nascent. Rolls, for one, has just begun examining whether the pathways her team has found in mice are also present in humans, and she has launched a start-up company to try to develop treatments based on the group's findings.

Although these developments are encouraging to researchers, much is still a mystery. "We often have a black box between the brain and the effect we see in the periphery," says Henrique Veiga-Fernandes, a neuroimmunologist at the Champalimaud Center for the Unknown in Lisbon. "If we want to use it in the therapeutic context, we actually need to understand the mechanism."

A TALE OF TWO SYSTEMS

FOR MORE THAN A CENTURY scientists have been finding hints of a close-knit relation between the nervous and immune systems. In the late 1800s and early 1900s, for example, scientists demonstrated that cutting nerves to the skin could curb some hallmarks of inflammation.

It wasn't until the late 1990s that researchers in this field began drawing connections to the body's master conductor, the brain. Neurosurgeon Kevin Tracey, then at North Shore University Hospital in Manhasset, N.Y., and his colleagues found something unexpected while investigating whether an experimental anti-inflammatory drug could help tame brain inflammation caused by stroke.

When delivered into the brains of rodents that had experienced strokes, the drug had the expected effect: it reduced neuroinflammation. As a control, the team injected the drug into the brains of animals that had inflammation throughout their bodies, thinking it would work exclusively in the brain. To their surprise, it also worked in the body. "This was a real head-scratcher," says Tracey, now president and chief executive of the Feinstein Institutes for Medical Research in Manhasset.

After months of trying to determine the path of the drug from brain to body, the researchers decided

NEUROSCIENTISTS Catherine Dulac (*right*) and Jessica Osterhout look at images of neurons in the hypothalamus that control symptoms of sickness, such as fever and loss of appetite.

Kris Snibbe/Harvard Staff Photographer

to cut the vagus nerve. With that nerve snipped, the anti-inflammatory effect of the brain-administered drug disappeared.

Inspired by this discovery, Tracey's group and others have continued to explore other ways in which the vagus nerve—and the rest of the nervous system—directs immune responses. A driving force for these developments, Swirski says, has been the advent of scientific tools that enable scientists to begin to chart the interactions between the nervous and immune systems in an unprecedented way.

Some researchers are focusing on particular body systems. For instance, a team led by Andreas Habenicht, a cardiologist at LMU Munich in Germany, reported in 2022 that the interaction between immune cells and nerves in the outermost layer of artery walls modulated the progression of atherosclerosis, an inflammatory disease in which vessels become clogged with cholesterol and other substances.

Meanwhile Veiga-Fernandes and his group have documented clusters of neuronal and immune cells in various tissues and discovered how they work together to sense damage and mobilize immune reactions. His team is now looking at how these little switchboards can be controlled by the brain.

The brain itself is also beginning to give up its secrets. Neuroscientist Catherine Dulac and her team at Harvard University have pinpointed neurons in an area called the hypothalamus that control symptoms that include fever, warmth-seeking and loss of appetite in response to infection. "Most people probably assume that when you feel sick, it's because the bacteria or viruses are messing up your body," she says. But her group demonstrated that activating these neurons could generate symptoms of sickness even in the absence of a pathogen. An open question, Dulac adds, is whether these hypothalamic neurons can be activated by triggers other than pathogens, such as chronic inflammation.

Just above the hypothalamus sits a region called the insula, which is involved in processing emotion and bodily sensations. In a 2021 study, one of Rolls's doctoral students, Tamar Koren, found that neurons in the insula store memories of past bouts of gut inflammation—and that stimulating those brain cells reactivated the immune response.

Rolls, Koren and their colleagues suspect that such a reaction might prime the body to fight potential threats. But these reactions could also backfire and start up in the absence of the original trigger. This could be the case for certain conditions, such as irritable bowel syndrome, that can be exacerbated by negative psychological states.

MIND OVER MATTER

MANY SCIENTISTS hope to pin down how such mental states influence immune responses.

Rolls and Fahed Hakim, a pediatrician and director of Nazareth Hospital EMMS in Israel, were inspired to investigate this question after coming across a 1989 study reporting that among women with breast cancer, those who underwent supportive group therapy and self-hypnosis in addition to routine cancer care survived longer than those who received only the standard treatment. Several other studies have documented a similar link between survival and the mental states of people with cancer.

To test the link, Rolls, Hakim and their team zoomed in on the VTA—the same region they targeted in the heart attack study and in a previous experiment looking at bacterial infection. This time they focused on

mice with lung and skin tumors. Activating neurons in the VTA noticeably shrank the cancers. It turned out that VTA activation subdued cells in the bone marrow that would usually repress immune activity, freeing the immune system to fight the cancer.

Clinicians have known about the effect of positive thinking on disease progression for a long time, Hakim says. But this evidence has been largely anecdotal or correlational, so being able to identify a pathway through which such an effect occurs—and manipulate it experimentally in animals—makes it much more real, he says.

Negative mental states can also influence the body's immune response. In a 2022 study, Swirski and his team identified specific brain circuits that mobilize immune cells in the bodies of mice during acute stress. The researchers found two pathways, one originating in the motor cortex that directed immune cells to the site of injury or infection and another beginning in the hypothalamus—a key responder in times of stress—that reduced the number of immune cells circulating in the blood. Swirski's group is currently investigating the role of stress-mediated circuits in chronic inflammatory diseases.

Neuroscientist Jeremy Borniger of Cold Spring Harbor Laboratory in New York and his colleagues have also found that activating neurons in the mouse hypothalamus can generate an immune response. They are now examining how manipulation of these cells can alter tumor growth.

Some groups are hoping to replicate their findings in humans. Swirski's team, for instance, plans to use tools such as virtual reality to alter people's stress levels and see how that changes their immune response.

Koren and Rolls are working with Talma Hendler, a neuroscientist and psychiatrist at Tel Aviv University in Israel, to see whether boosting the reward system in people's brains before they receive a vaccine can improve their immune response. Rather than stimulating the brain directly, they are using a method called neurofeedback. In this approach, individuals learn to observe and control their own brain activity, which researchers measure using methods such as functional magnetic resonance imaging.

THE ROAD TO THE CLINIC

OVER THE YEARS Rolls has chatted with her good friend Tehila Ben-Moshe about her research. Ben-Moshe is chief executive of Biond Biologics, an Israel-based biopharmaceutical company that focuses on using immune cells to target cancer. During one such discussion in 2022, Ben-Moshe realized that Rolls's brain-stimulation experiments were acting on some of the same immune cells her company was trying to target, and she immediately saw the therapeutic potential. "When I saw Asya's data, I couldn't believe what I saw," Ben-Moshe says. "The question then became, How can I translate what she's doing with mice into patients?" The two are working on launching a company.

Ben-Moshe and Rolls hope to harness existing brain-stimulation technologies—for example, transcranial magnetic stimulation, which uses magnetic pulses to alter brain activity, or focused ultrasound, which uses sound waves—to modulate the immune systems of people with cancer, autoimmune diseases or other conditions. As a first step, their team has been reaching out to companies that have developed such technologies. Before starting their clinical trials, Ben-Moshe and Rolls want to examine blood samples from past trials that used these techniques so they can see whether there are signs of immune-system alterations before and after treatment.

Potential therapies targeting the vagus nerve are nearer the clinic. A company co-founded by Tracey—SetPoint Medical in Valencia, Calif.—is testing pill-sized nerve stimulators implanted in the vagus nerve in the neck in people with autoimmune diseases such as Crohn's disease, multiple sclerosis and rheumatoid arthritis. The rheumatoid arthritis trial is furthest along—the team has shown in a small trial in Europe that its device can reduce disease severity. The technique is currently undergoing a randomized, sham-controlled trial (in which the control group receives an implant but no active stimulation) in 250 patients in various centers across the U.S.

> "You can call something psychosomatic, but in the end, it's somatic. How long can we ignore what is there?"
> —Asya Rolls *Technion–Israel Institute of Technology*

Rolls's hope is that this work will ultimately help physicians to understand, and act on, the mind-body connections they see in their practices. The need is clear: when Rolls put out a call to speak to psychologists from the hospital where her lab is based, the meeting room was packed. People from departments ranging from dermatology to oncology were eager to share their stories. Many clinicians pass people with seemingly psychosomatic issues on to psychologists, saying there is nothing physically wrong, one attendee said. This can be distressing for the person seeking treatment. Even being able to simply tell people that there is a brain-immune connection that is responsible for their symptoms can make an enormous difference.

It's time that both researchers and clinicians take the link between psychology and physiology seriously, Rolls says. "You can call something psychosomatic, but in the end, it's somatic. How long can we ignore what is there?" SA

Diana Kwon is a freelance science journalist based in Berlin.

Growing Older Means Less Stress

Good news: people report that worries and anxiety drop away as they age

By Daisy Yuhas

NO ONE IS A STRANGER TO STRESS. DECADES OF RESEARCH make it clear that major life events, such as the death of a loved one or the start of a new job, can take a lot of our energy and attention. But recently scientists have begun to reveal how smaller daily stressors can create substantial difficulties. David Almeida, a developmental psychologist at Pennsylvania State University, has been following the stressors of daily life in a group of more than 3,000 adults since 1995. Almeida spoke with *Scientific American* contributor Daisy Yuhas to describe a benefit of aging that he has uncovered: stress levels go down and coping skills may actually go up after people pass their 20s. Still, upsetting national or global events can increase our distress.

[*An edited transcript of the interview follows.*]

You've been tracking people's daily experiences for two decades. How has that shifted your perspective as a psychologist?

In my work, I try to characterize a day in the life of an individual. I look at how people use their time, how they experience stressors and positive events, their mood and their physical symptoms. I chart how this changes from day to day—the ebb and flow of daily experiences. So even though I'm a psychologist, my unit of analysis is a day, not a person.

The more I've dug into this work, the more I've begun to see that people actually differ from themselves day to day as much as you differ from somebody else. Our identity isn't just who we are based on the average of our experiences. Our identities may be in the range in our behavior, the extent to which we're going up and down with our experiences.

How do you track daily stressors?

We ask people to answer a series of structured questions at the end of every day. Originally we used telephone calls, and now we use web-based approaches. We ask about how they spent their time, their mood, their physical symptoms, who they interacted with, and then we ask a lot of questions about the types of stressors they experienced that day. For some studies, we also collect a sample of saliva, which lets us determine the amount of stress hormones in the body.

With that method, we've worked with a large group of people. I want to acknowledge that the wonderful participants in the National Study of Daily Experiences—which is part of a large-scale investigation called Midlife in the United States—have shared their lives with me for the past 20 years. It's been a privilege to follow them.

You recently published findings from an analysis of 2,845 adults—ages 22 to 77 at the start—over 20 years. In that work, you found that people seem less stressed as they grow older. Can you unpack that?

Yes, finally some good news about daily stress! It seems to get a little bit better. We find that younger people report more exposure to stressful events—things they find challenging, upsetting or disruptive—than older people do. So people in their 20s may report stressors on at least 40 to 45 percent of days, but by the time they're in their 70s, that goes down to maybe 20 to 25 percent of days.

In addition, we looked at how much distress people experience or the way they *respond* to stress. Here we see the same type of pattern, with young adults having higher distress on days with stressors than older people. But around 55 years old, that age advantage—where your response to stress gets better with age—starts to taper off and plateau.

Why is there an age advantage in dealing with stress?

I think three reasons could contribute and work together. One has to do with the social roles people inhabit. When you're young, these roles could include being a parent of a young child, starting a job, getting into new relationships. New roles are stressful, as are role conflicts that happen when you have multiple roles going on at once.

A second reason could be that as we grow older, we realize we have only so much life left and want to make the most of it—so we are very motivated to enjoy it.

The third reason, which I am most interested in, is that just by virtue of experiences, opportunities and past stressors, we learn how to deal with them and become more expert in coping with daily stressors as we get older.

Does that explain why research suggests older people are happier than younger ones?

As you grow older, you can list all these things you shouldn't be looking forward to: poorer physical health, loss of friends, being sick and cognitive decline. These are not things you would expect to be related to increased happiness. But we see over and over that as people grow older, they have *increased* life satisfaction.

That said, there is a point when this pattern stops. Much later in life—in someone's 80s or 90s—I think we're seeing a time where things are really tough, and there's a decrease in life satisfaction.

How do things such as economic and political uncertainty affect our day-to-day stress?

We were able to study the effects of the 2008 recession and the postrecession period. From our data, it's fairly clear that compared with people in 1995, adults in 2010 had more stressful daily lives and were more upset by their experiences. Our hypothesis is that this reflects historical changes such as the recession and the use of technologies that have changed social interaction. From that, we can speculate on how economic downturn and other changes may affect us. In future work, we hope to see what the pandemic has done—it's possible that we won't see much of an age advantage, for example, in this period.

But what really surprised us from our analysis of the 2008 recession was that this difference in stress seems to be concentrated among people who were in midlife. I'd have thought that younger adults just starting their careers and older adults in retirement would have been worse off, but no: it was adults in their mid-40s through mid-60s who reported higher levels of psychological distress. I think that has to do with the social roles of a midlife adult. They are worried about their kids but also their parents.

On a practical note, should we be trying to remove all stressors from our daily lives?

There's something that might actually be good about having some daily stress. People who report having no stress in their lives—you think they are lucky, happy people. But they also report fewer positive things in their lives. They have fewer people in their lives and perform worse on cognitive tests.

It's the *reactivity* to stress—how you respond to it—that really matters to your health and well-being. It's not the number of stressors but actually your emotional responses that can, for example, raise your risk of cardiovascular disease, increase inflammation and contribute to dying earlier.

How should we manage our responses?

There are things people can do, such as eating well and getting enough sleep. But remember that not everyone can do them. It's not all about individual choice.

We've found that minoritized groups—by race, ethnicity and sexual orientation—have higher stress reactivity. They don't always have the resources to cope with daily stressors. For instance, when your body is experiencing stress, it wants to mobilize energy. So getting up and taking a walk is the best way to stem that emotional response. But many people cannot just get up in the middle of their workday and take a walk somewhere.

We need to start talking about how to provide resources to empower people so they can take care of themselves. SA

Daisy Yuhas edits the *Scientific American* column Mind Matters. She is a freelance science journalist and editor based in Austin, Tex.

Illustration by Dilek Baykara

The Secrets of Thirst

Your body usually knows how much water you need, without arbitrary targets

By Claudia Wallis

SERIOUS QUESTION: How much water does the average adult need to drink every day? You've probably heard the usual answer: eight 8-ounce glasses, sometimes stated as 8 × 8. But there is not much science behind this ubiquitous recommendation. A 2002 research review found essentially no reliable studies. Any truly serious answer to the how-much question will begin with some version of "it depends." Are you in a hot location? Are you exerting yourself? Are you in good health? How big are you? Do you eat a lot of salty foods, or do you mostly load up on fruits and vegetables?

We do need water every day, but the average person gets it from many sources: tea or coffee, soft drinks (which often include sugar that you don't need), and food. "We typically get about 20 percent of our fluid requirements from solid foods and about 80 percent from beverages," says Brenda M. Davy, professor of human nutrition at Virginia Polytechnic Institute and State University. To maintain a healthy balance of water, minerals and salts, health authorities say adults should drink about a liter (34 ounces) of liquid for every 1,000 kilocalories consumed. That works out to be a little over eight cups for someone who takes in 2,000 daily calories—a possible source of the 8 × 8 notion. But most Americans achieve this level of hydration from a variety of foods and drinks, with about a third coming from plain water, according to a 2013 study of nearly 16,000 U.S. adults.

Natural thirst mechanisms are the reason that most of us do not need to be overly concerned about hydration. The adult body is roughly 60 percent water—closer to 80 percent in the lungs and kidneys—and it carefully controls the concentration of water. We are all familiar with the sensory aspect of this regulation: the dry throat and urgent alert of thirst. But recently neuroscientists have gained other remarkable insights into how thirst is monitored in the body and controlled in the brain.

Researchers have known since the 1950s that a pea-size structure in the brain's hypothalamus controls thirst. In a series of experiments in which he infused salt into the brains of goats, Swedish physiologist Bengt Andersson showed that a region called the subfornical organ (SFO) monitors the concentration of water and salts in blood and triggers the urge to drink. The SFO plays the same role in people. But Andersson's ideas failed to fully explain how humans experience thirst. For instance, when we gulp a drink, we feel almost instantly satisfied, yet it takes 10 to 15 minutes for a liquid to make it from our mouth through the digestive tract and into the bloodstream. "Something in the brain is saying that your blood may not have changed conditions yet but that you drank enough water so you can stop feeling thirsty," explains neuroscientist Christopher Zimmerman of Princeton University.

In a series of elegant experiments with mice, Zimmerman and his associates measured the activity of neurons in the SFO. "We saw that their activity changed very fast when the mouse drank water or drank salt water and when it ate food," Zimmerman says. The researchers showed that signals converged on the SFO from several places. "You get a signal from the blood that tells your current state of hydration, a signal from the mouth that tells you how much fluid you drank, and a signal from the gut that tells you what was consumed—was it water? Was it something else?" The SFO neurons, he explains, "add these signals together" and then transmit the urge to drink or stop drinking.

The big takeaway of Zimmerman's work is that for the most part you can trust your thirst system to tell you when you need to drink, as opposed to following some arbitrary advice. But there are exceptions. Because the system's sensitivity may decline with age, older adults may need to set reminders to drink—the 2013 study found that, on average, people older than 70 failed to get adequate hydration. People with certain conditions, including kidney stones and diarrhea, also need extra water. And research by Davy and others indicates that middle-aged and older people who are trying to lose weight or maintain weight loss consume fewer calories if they fill up with 16 ounces of water before meals.

Other parts of the brain—the ones used in planning—should help with hydration on hot days and when exercising. Thirsty or not, Zimmerman says, he drinks water before going for a run: "My thirst neurons don't know I'm about to run 10 miles." SA

Claudia Wallis is an award-winning science journalist whose work has appeared in the *New York Times*, *Time*, *Fortune* and the *New Republic*. She was science editor at *Time* and managing editor of *Scientific American Mind*.

Illustration by Fatinha Ramos