BEHIND THE
Connectome

Connectomics is having a moment.

Following on the heels of genomics, proteomics, transcriptomics, metabolomics, and microbiomics, the latest "omic" to seize the spotlight is generating the kind of buzz that makes other disciplines fluorescent green with envy. As the name suggests, connectomics maps connections—specifically, the ones between the neurons in an animal's brain or nervous system.

The advent of high-throughput, computer-assisted techniques has led to an explosion of connectomic technologies and studies. The field is also amassing the sort of Big Science resources previously associated with efforts to land a man on the moon or decode the human genome. The Obama administration, for example, recently decided to pump $100 million into the Brain Research through Advancing Innovative Neurotechnologies (BRAIN) project to develop methods for recording neuronal activity on a large scale; while the European Commission is investing €1 billion in Henry Markam's Human Brain Project in Switzerland—a plan to build a computer simulation of the human brain, neuron by neuron.

At the same time, the connectomics frenzy has come under fire. Some scientists question the wisdom of devoting scarce resources to the pursuit of the human connectome when the task remains well beyond the scope of currently available technologies. As **Sebastian Seung, PhD**, author of *Connectome: How the Brain's Wiring Makes Us Who We Are*, notes in a widely watched TED talk (more than half a million views and counting), the only complete connectome we have is for the tiny nematode worm *C. elegans*, an animal with roughly 300 neurons and 7000 neuronal connections; and the job took more than a decade using conventional electron microscopy, which was used to resolve individual synaptic connections between cells. By contrast, the human brain contains roughly 100 billion neurons and 100 trillion connections. This is why Seung describes the effort to map the human connectome as "one of the greatest technological challenges of all

time," and asserts that it will "take generations to succeed." Others question the wisdom of concentrating on the development of new technologies in the absence of clearly articulated scientific goals, and on focusing so intensely on wiring diagrams that cannot by themselves explain how our brains give rise to feelings, thoughts, and perceptions.

But researchers from a wide range of backgrounds—some trained in physics or engineering, others in neuroscience or medicine—are mounting a serious effort to demonstrate the viability of connectomics. Thus far, they have focused on two broad and equally important tasks: devising faster and better methods for building connectomes; and putting connectomic data to good use. Their collective spadework constitutes the true current state of connectomics—and their success, its true promise.

Building a Connectome by Taking the Middle Road

In 1993, Francis Crick coauthored a *Nature* article ("Backwardness of Human Neuroanatomy") that lamented the lack of progress toward a "connectional map" of the human brain. Since then, not much has changed, argues **Partha Mitra, PhD**, Crick-Clay Professor of Biomathematics at Cold Spring Harbor Laboratory. In part, he says, that's because neuroanatomists have tended to beaver away in their individual labs using different paradigms and techniques. As a result, their data often doesn't integrate well, and the models of neuroanatomy and connectivity they develop stay locked inside their own brains.

In a 2009 *PLoS Computational Biology* paper, Mitra and a number of his colleagues proposed solving that problem by launching a coordinated effort to construct a mesoscale whole-brain wiring diagram for a vertebrate. That proposal led to the creation of the Mouse Brain Architecture (MBA) Project, an attempt to demonstrate the attainability of a large connectomic endeavor while also providing a testbed for developing practical neuroanatomical techniques.

As a theoretical physicist with an eye for the big

Commotion

By Alexander Gelfand

picture, Mitra sought to develop a systematic method for constructing a complete connectome using currently available technology. Taking a page from the Human Genome Project, he developed a high-throughput, semi-automated pipeline for imaging multiple mouse brains using light microscopy, a technique whose resolution lies between that of electron microscopy, which is impractical for mammalian brains, and the non-invasive yet far coarser magnetic resonance imaging (MRI) techniques used on human subjects.

Last June, Mitra released the first round of gigapixel image data collected for the MBA Project. The images can be viewed online and explored with a virtual microscope: Users can zoom in on individual neurons and their axons, the long, slender fibers that trail away

of a sphere-packing algorithm. The tracers are either absorbed into neuronal cell bodies and spread through their axons, or are taken up by axons at synapses and propagate up into the cell bodies. The brains are subsequently sectioned into 20µ-thick slices and imaged using either brightfield (i.e., white-light) or fluorescence microscopy.

Then comes the tricky part. The resulting two-dimensional images must be assembled into three-dimensional stacks and registered to an anatomical reference atlas using a combination of off-the-shelf and custom software. Next, axons and cell bodies must all be identified, a step that presents a significant bottleneck. Mitra's lab is working on automating the process using machine-learning algorithms that can be trained to identify features of interest. For now, however, "somebody has to look through half a million sections," he says.
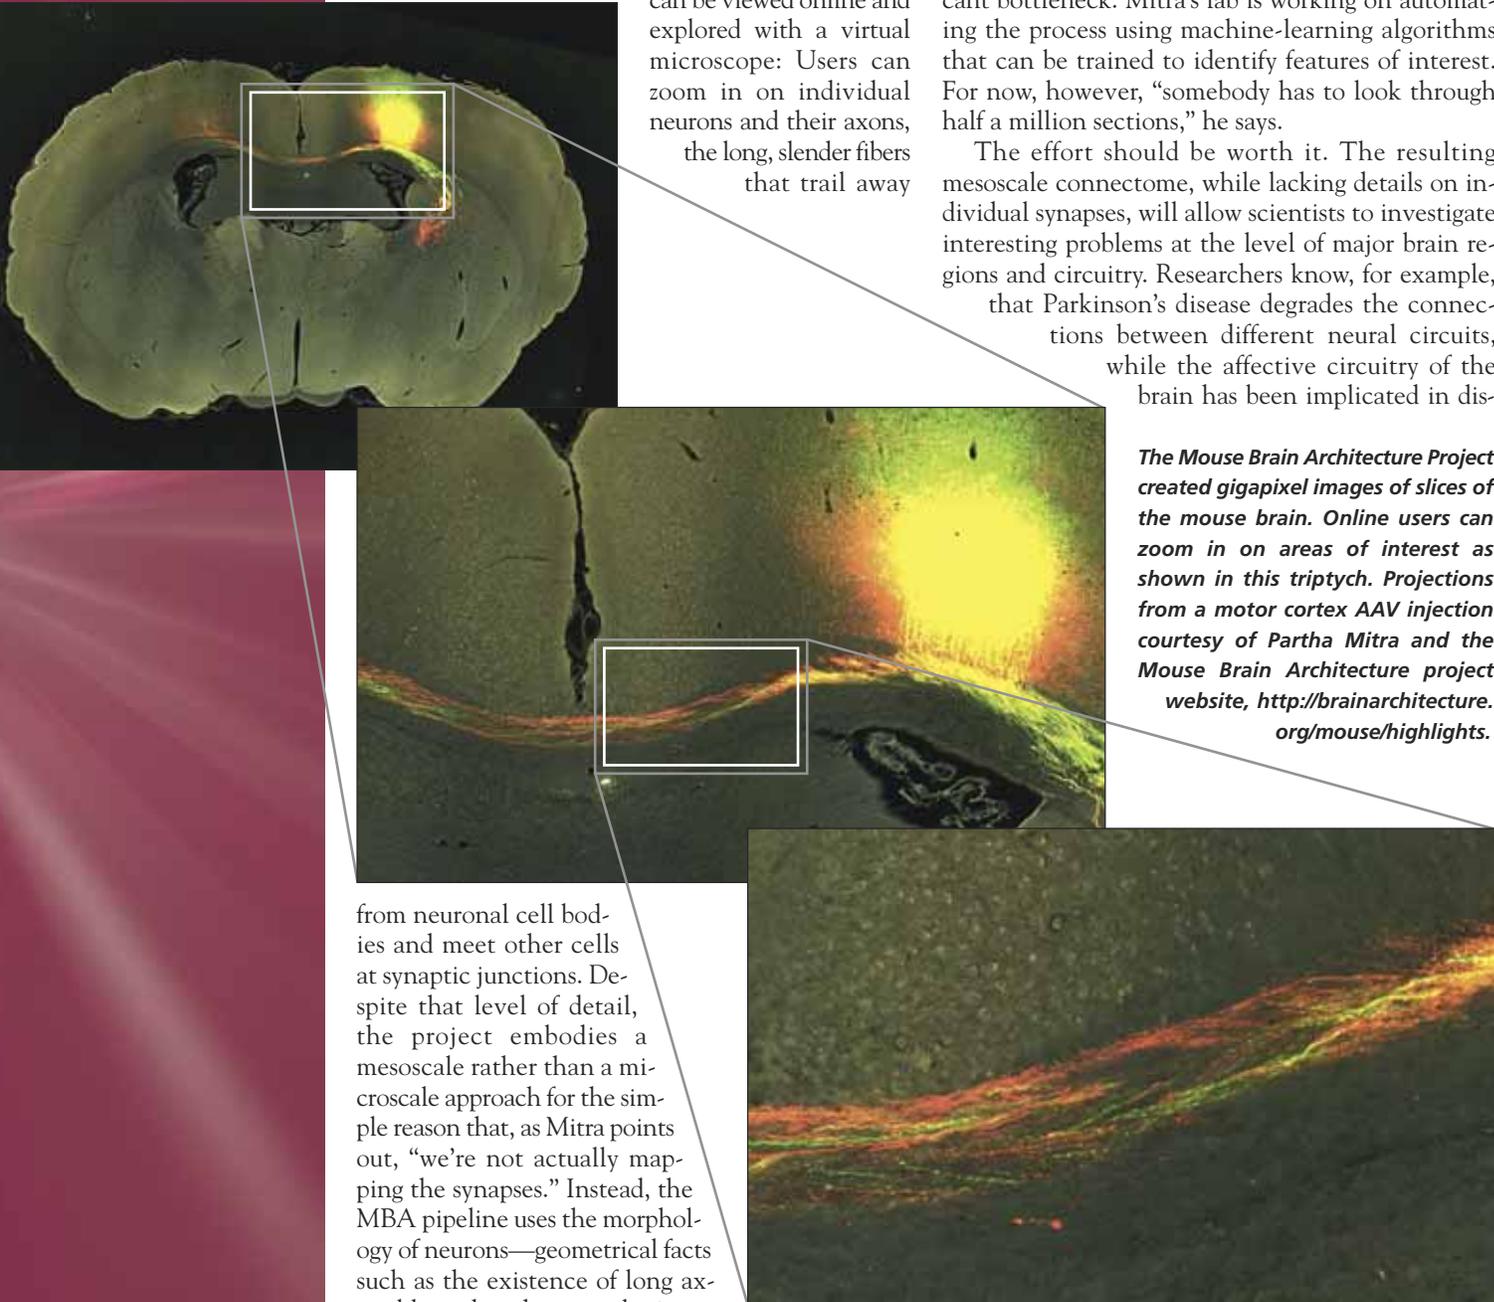
The effort should be worth it. The resulting mesoscale connectome, while lacking details on individual synapses, will allow scientists to investigate interesting problems at the level of major brain regions and circuitry. Researchers know, for example, that Parkinson's disease degrades the connections between different neural circuits, while the affective circuitry of the brain has been implicated in dis-

The Mouse Brain Architecture Project created gigapixel images of slices of the mouse brain. Online users can zoom in on areas of interest as shown in this triptych. Projections from a motor cortex AAV injection courtesy of Partha Mitra and the Mouse Brain Architecture project website, http://brainarchitecture. org/mouse/highlights.

from neuronal cell bodies and meet other cells at synaptic junctions. Despite that level of detail, the project embodies a mesoscale rather than a microscale approach for the simple reason that, as Mitra points out, "we're not actually mapping the synapses." Instead, the MBA pipeline uses the morphology of neurons—geometrical facts such as the existence of long axonal branches that span brain regions—to infer patterns of connectivity.

To gather those geometrical facts, Mitra's team injects mouse brains with four different tracers, some of which express fluorescent proteins, in 262 uniformly spaced sites that were chosen with the help

orders such as anxiety and depression; and many believe that conditions like autism and schizophrenia are caused by pathological patterns of neural connectivity, or "connectopathies." Mesoscale connectomes could lead to better diagnostic tools for major disor-

ders, better drug therapies, and even to a better understanding of how genetic variation influences behavior by shaping the wiring of the brain.

Building a Connectome at the Microscale

Meanwhile, Mitra's Cold Spring colleague **Anthony Zador**, **MD**, **PhD**, co-founder of the Computational and Systems Neuroscience (Cosyne) conference, is working on a microscale approach to building connectomes that would dispense with microscopes altogether.

As someone who studies attention and auditory processing, Zador wants a way to model and interrogate neural circuits *in silico* in order to streamline the lengthy and laborious process of running experiments to determine how groups of neurons feed information to one another. Doing that, however, requires a quick and inexpensive method for tracing the synaptic connections between individual cells. If Mitra's mesoscale approach to identifying axonal pathways is akin to sketching the major highways in the United States, says Zador, he's interesting in identifying "each and every street, road, and country lane."

Zador was initially inspired by the Brainbow technology invented by Harvard researchers **Jeff Lichtman**, **MD**, **PhD**, and **Joshua Sanes**, **PhD**. Brainbow uses genetically engineered neurons to express random combinations of up to four fluorescent proteins, producing brightly colored collections of cells that can be imaged using fluorescence microscopy. But the technique is limited to a palette of just a couple of hundred colors—not nearly enough to uniquely label all of the neurons in even a small sample of brain tissue—and offers limited resolution. So Zador began looking for optical alternatives to fluorescent proteins. Eventually, he realized that he could do away with them entirely and use DNA barcodes instead. "The readout for the fluorophores is a microscope. The readout for the barcodes is an Illumina," says Zador, referring to the ultrafast next-generation gene sequencing machines. Zador described his approach, called BOINC ("barcoding of individual neuronal connections"), in a *PLoS Biology* paper published last October.

With the help of a recombination enzyme that can scramble bits of genetic code, Zador instructs neurons to generate short random sequences of DNA. In theory, a sequence containing just 20 random nucleotides could uniquely label 2×10^{12} neurons, more than enough for the 100 million or so neurons in a mouse brain. Once the neurons are labeled with these randomly generated barcodes, Zador traces their connections by having a transsynaptic virus spread the barcodes from cell to cell. That has the effect of turning each neuron into a "bag of barcodes" that contains not only its own unique DNA label, but also the unique identifiers belonging to

each neuron that's connected to it. Some more genetic engineering technology is used *in vivo* to join each neuron's barcode with the barcodes from neurons to which it is connected by a synapse, creating sequences of fused barcodes that represent networks of neurons; harvesting and reading the fused barcodes with a high-throughput sequencer yields a connectivity matrix. Computation plays a role in several places: Zador and his colleagues had to develop novel algorithms to clean up the barcodes, correct for any sequencing errors, and determine the connectivity matrices. They have been running proof-of-principle experiments using neurons cultured in an incubator, and have successfully completed each step in the process in isolation. Despite some remaining technical hurdles, Zador expects to combine all of the steps within a matter of months.

Sectioning the brain before extracting the DNA for sequencing will allow Zador to identify the brain
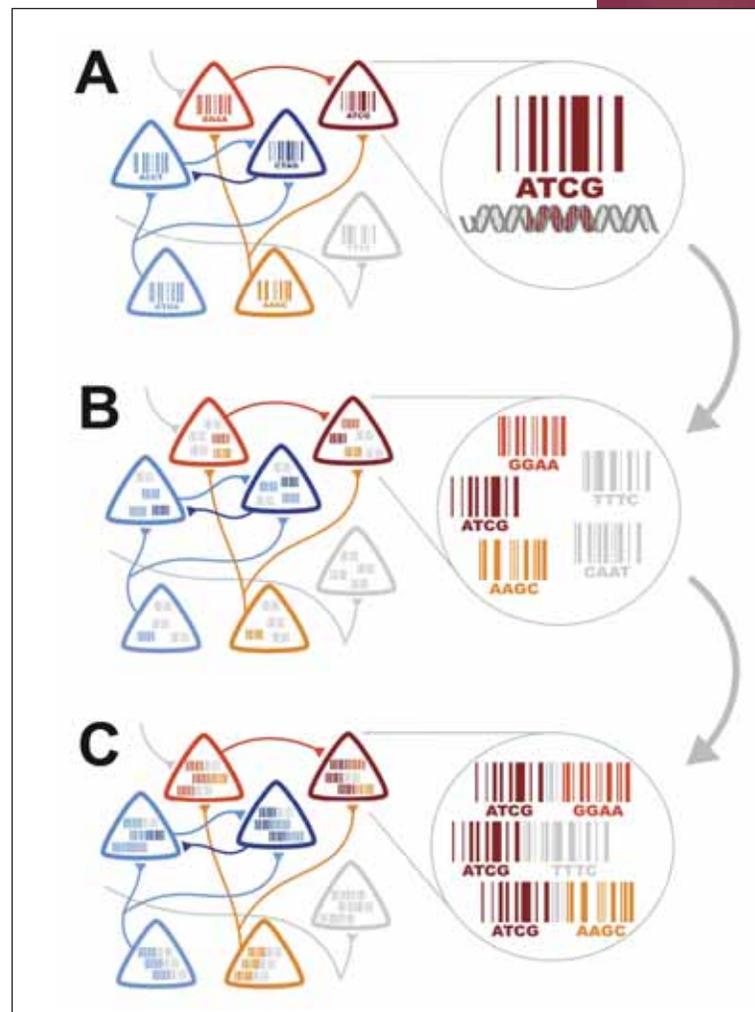
Zador and his colleagues convert neuron connectivity into a sequencing problem that can be broken down conceptually into three components—labelling neurons with unique DNA barcodes; associating barcodes from synaptically connected neurons; and joining host and neighboring (invader) barcodes into pairs for sequencing. Reprinted from Zador AM, Dubnau J, Oyibo HK, Zhan H, Cao G, et al. (2012) Sequencing the Connectome. **PLoS Biol** *10(10): e1001411. doi:10.1371/journal.pbio.1001411, (2012).*

region that each neuron comes from. And the use of genetic technology ought to allow Zador to determine the particular kinds of neurons (e.g., inhibitory, excitatory) involved, too. Inhibitory neurons, for example, express a particular enzyme that is encoded in mRNA; by tagging the appropriate mRNA in a given batch of neurons with barcodes, Zador should be able to identify the inhibitory ones. That would add a level of detail about the identity of individual neurons that would be hard to come by even using electron microscopy; and if his sequencing approach works, it would be cheaper and faster than anything currently out there. In his *PLoS* paper, Zador estimates that sequencing the connectome of a mouse cortex would cost $40,000 at current rates, "and could easily drop several orders of magnitude in a few years." Sequencing a fruit fly brain would cost one dollar, and doing *C. elegans* would be "essentially negligible." That would put neuroscientists in a position to quickly and inexpensively map brain circuits, allowing them to develop testable hypotheses and design experiments far more efficiently.

## Using Connectomes To Understand Behavior

Some researchers are already using connectomics to understand behavior. In a December 2012 paper published in *Cell*, for example, **Daniel Bumbarger, PhD**, used differences in neural connectivity to help explain the divergent feeding behaviors of *C. elegans* and its nematode cousin, *P. pacificus*. Unlike *C. elegans*, which feeds exclusively on microbes, *P. pacificus* is capable of switching into predator mode and eating other nematodes. Typically, neuroscientists have explained such behavioral differences by looking at the physiology of the neurons involved, or the neurotransmitters that modulate them. But Bumbarger, who is a postdoctoral fellow at the Max Planck Institute for Developmental Biology in Tuebingen, Ger-

many, wanted to see if those differences in feeding styles could be related to differences in patterns of synaptic connectivity. So he compared the wiring diagrams for the pharyngeal nervous systems of the two worms—a task that first required imaging the 300μm-long pharynxes of several *P. pacificus* specimens using an electron microscope, something that in itself took nearly two years of work.

What he found was striking. Despite having basically the same number and types of neurons in their pharyngeal nervous systems—a remarkable conservation of cell identity—the two nemotode species displayed a "massive rewiring of synaptic connectivity," Bumbarger says, with *P. pacificus* demonstrating much higher and more complex connectivity than *C. elegans*.

To better understand how differences in connectivity might be driving differences in behavior, Bumbarger turned to graph theory, the branch of mathematics that gave rise to network analysis. Graphs are defined as sets of nodes connected by edges, or lines; consequently, graph theory can be used to analyze the characteristics of virtually any network, including neural ones in which the nodes are brain regions or neurons, and the edges are axons or synapses. To compare the relative importance of the various neurons shared by *C. elegans* and *P. pacificus*, Bumbarger computed a variety of measures that evaluate the centrality of nodes within their networks—measures like degree centrality, for example, which counts the connections associated with a node, and PageRank centrality, which gauges the probability of stopping at it. (PageRank helps Google rate the importance of webpages.) He also developed a new tool, called focused network cen-

Bumbarger and his colleagues compared the synaptic connectivity of the two nematode species C. elegans *(A—based on previous work by others) and* P. pacificus *(B), both shown here in a two-dimensional representation. Nodes indicate neurons (blue), muscle cells (red), and other network outputs (yellow). Edges curve clockwise from the presynaptic to the postsynaptic node and are colored the same as their postsynaptic partners, with edge width indicating connection weight, or strength, according to multiplicity of synapses. Bumbarger and his colleagues also mapped differences in PageRank centrality onto the* P. pacificus *network (C). Node size is proportional to magnitude of the difference in PageRank between* C. elegans *and* P. pacificus. *Orange nodes have a higher centrality in* P. pacificus, *whereas blue nodes have a higher centrality in* C. elegans. *Nodes with connections to anterior pharynx output cells (red edges), including those nodes proposed to control predatory feeding, have a higher PageRank in* P. pacificus *than in* C. elegans. *Nodes with connections to posterior pharynx outputs (blue edges) have a higher PageRank in* C. elegans *than in* P. pacificus. *Reprinted with permission from Bumbarger, DJ et al., System-wide Rewiring Underlies Behavioral Differences in Predatory and Bacterial-Feeding Nematodes,* Cell *152:109-119 (2013).*
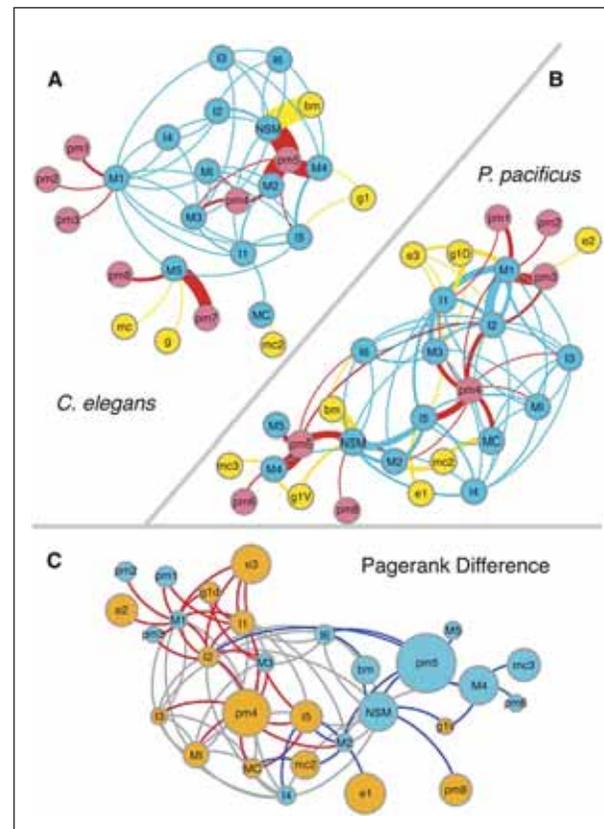
trality, to determine which parts of each network were most important to particular nodes.

Among other things, Bumbarger found that there was a general shift in network focus between the two species, with more going on in the anterior portion of *P. pacificus'* pharyngeal network than in the posterior portion—precisely the opposite of *C. elegans.* He also found that while information tended to follow the shortest path across *C. elegans'* pharyngeal network, information flow in *P. pacificus* was more indirect, suggesting more complex processing that could correlate with its more diverse feeding behaviors. And there were significant differences in connectivity and information flow associated with the two neurons that play the largest role in regulating feeding behavior in both worms.

It's hard to know exactly what all this means in terms of function, but Bumbarger's findings point the way toward experiments that could help explain how differences in connectivity and network architecture affect behavior. He'd now like to do laser ablation experiments on both species—blasting away at those two neurons, for example—to see what, if any, changes in feeding behavior ensue.

Simulating a Human Connectome: Spaun

Worms are one thing, people another. And the amount of time it took Bumbarger to map just a tiny piece of *P. pacificus'* total connectome—let alone begin to understand its functional significance— gives some indication of why even the most ardent advocates of mapping the human connectome see it as a very long-term goal. But that hasn't stopped some researchers from taking the data that's already out there and using it to model human behavior.

Chris Eliasmith, **PhD**, a theoretical neuroscientist at the University of Waterloo in Ontario, Canada and author of the forthcoming book, *How to Build a Brain*, has developed a large-scale computational model of the human brain made up of two and a half million virtual neurons. Known as Spaun (for Semantic Pointer Architecture Unified Network), the model, which is described in a 2012 *Science* paper,

In these screenshots from movies of Spaun processing an input image, the spiking neural networks in the model are mapped to the corresponding anatomical areas. For example, the highest level of visual hierarchy lies at the back of the brain (in the inferotemporal cortex); the motor areas in the middle; and executive control in the front. When shown the number two, the visual area responds; and when prompted by a question mark to write the number, the motor area kicks in, recognizing not only the number but details such as the numeral two's loop (or lack thereof), demonstrating its ability to capture this subtle visual difference. Screenshots taken from http://nengo.ca/build-a-brain/severaltasks, courtesy of Chris Eliasmith.
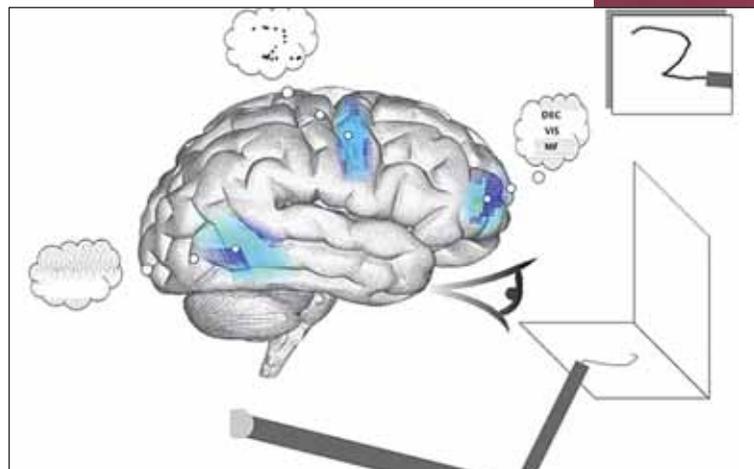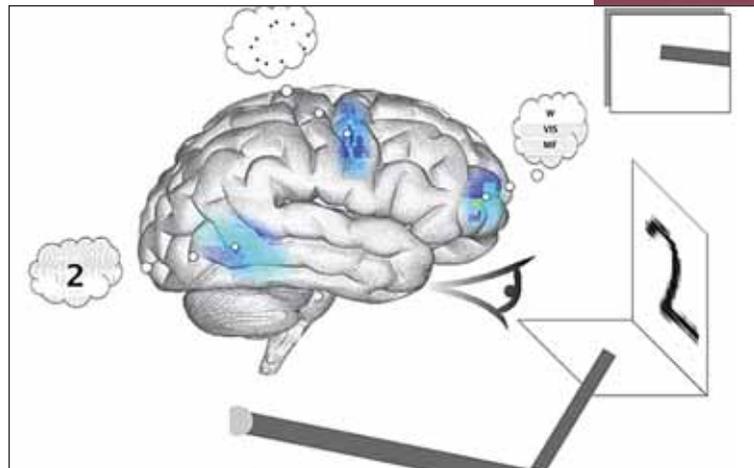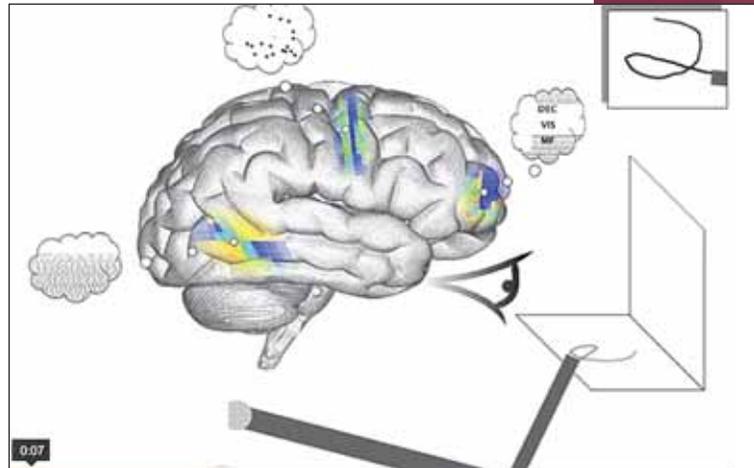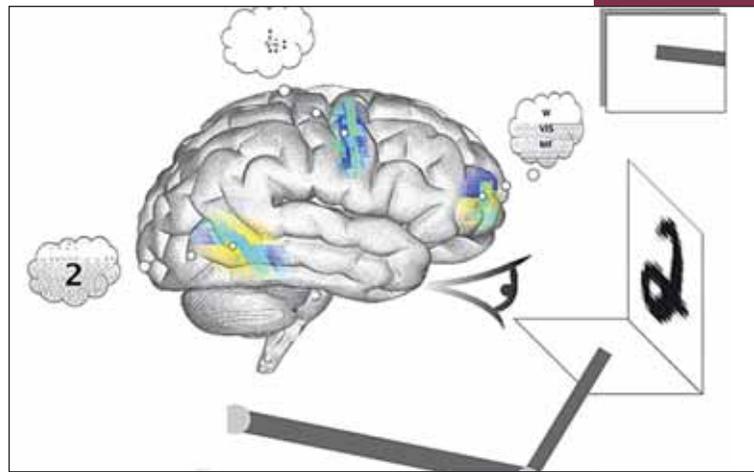
can see with a simulated eye and write with a simulated arm. Eliasmith and his colleagues used as much neuroanatomical and biological data as they could to build Spaun; its simulated neurons are grouped into 20 anatomical structures (primary visual cortex, primary motor cortex, and so on) that are wired together in a realistic manner, mimicking functional brain areas that communicate with one another to reproduce a variety of cognitive behaviors.

Spaun was designed to perform eight basic tasks, including one that involves viewing a sequence of digits displayed on a screen, remembering them, and writing them down. Spaun can do it, but just like a real human being, it's better at remembering the items at the beginning and end of the list. It also performs about as well as most people would on a reasoning task that resembles the kinds of problems included on a common IQ test.

The fact that Spaun can do these things almost as well as people can, while also making the same kinds of mistakes they do, lends credence to the assumptions about brain wiring and function upon which it is based. For example, Spaun's ability to handle a diverse set of tasks is made possible by the way in which its virtual basal ganglia—a group of neurons that are associated with functions such as motor control and procedural learning—route information through simulated synaptic connections to different portions of its virtual cortex depending on the job at hand. To some degree, the system is even capable of changing its connection weights, or the strength of the connections between its neurons, a property that is believed to play a key role in memory formation, information processing, and behavior.

The model is far from complete. "It's large-scale in a way, but it's 40,000 times smaller than the

The practical benefits of a (relatively) large-scale model of a functioning brain are already apparent.

brain," Eliasmith says. And while it is capable of enacting very small variations on the routing schemes supplied by Eliasmith and his colleagues, Spaun will have to figure out how to rewire itself more substantially in order to learn new tasks. Still, the practical benefits of a (relatively) large-scale model of a functioning brain are already apparent.

On the one hand, Spaun should help neuroscientists figure out why specific connections matter, and how neural anatomy and physiology underwrite behavior. That, says Eliasmith, will help them understand how our brains relate to who we are and what we do. On the other hand, while Spaun can replicate normal cognitive behavior, it can also be used to model the cognitive decline associated with aging, or the damage inflicted by diseases like Parkinson's and Alzheimer's.

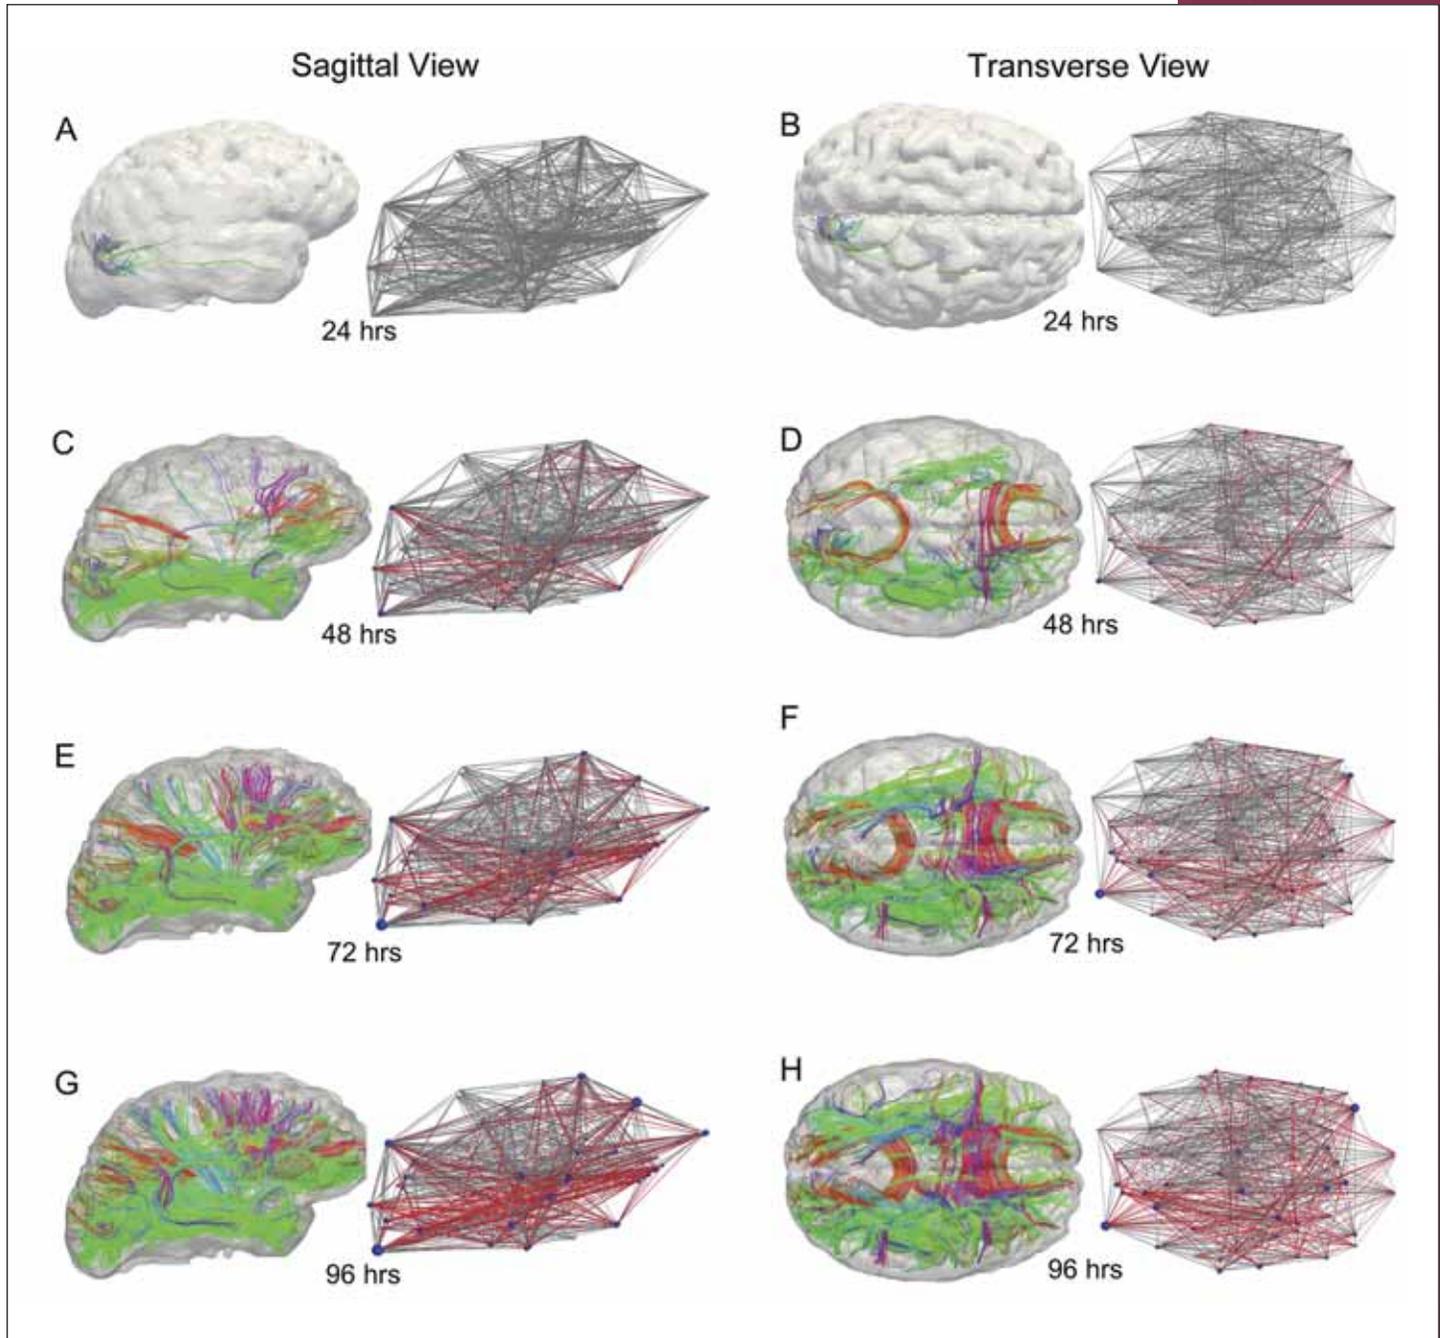## Broken Connectomes: Understanding Brain Trauma

Damage is precisely what **Reuben Kraft**, **PhD**, assistant professor of mechanical engineering and member of the Institute for Cyberscience at Penn State University, wants to understand—and prevent. While working at the Army Research Laboratory, Kraft began trying to model traumatic brain injury (TBI), the signature injury suffered by American troops in Iraq and Afghanistan. Studies have linked TBI to chronic traumatic encephalopathy, a progressive neurodegenerative disease that affects both soldiers who are subjected to blast and concussion, and athletes such as boxers and football players who suffer repeated head injuries. Colleagues in the Translational Neuroscience Branch introduced him to connectomics, and Kraft, enticed by the combination of imaging and network analysis, was off and running. In a *PLoS Computational Biology* paper published last August, he and his collaborators used magnetic resonance imaging, biomechanical modeling, and graph theory to examine how a blow to the head might affect an individual's neural network.

Kraft used a technique known as finite element modeling to turn standard MRI scans taken from a graduate student at the University of California, Santa Barbara, into a three-dimensional model of a human head that included everything from skull and skin to brain tissue and cerebrospinal fluid. He then combined that model with a low-resolution connectome constructed via diffusion tensor imaging (DTI), a form of MRI that traces the approximate location of bundles of axons by analyzing the movement of water molecules along the fibers.

DTI provides far less detail than microscopy, but the technology is well suited to the kind of macroscale study that Kraft wanted to perform. By the time they were done, Kraft and his colleagues had a 3-D model kitted out with a connectome in which the nodes represented anatomical regions of the brain rather than individual neurons, and the edges represented the axonal highways that linked them together. By throwing in experimentally based models that predict how cells die off in the hours and days following an initial insult, Kraft was able to predict how the connectome would evolve over time after an impact to the head, with connections between brain regions degrading or disappearing completely as cells expired. He then analyzed the local and global effects on the network using graph theoretical measures of efficiency and connectivity. The system proved to be surprisingly robust: Even in the worst-

case scenarios, with many connections lost, none of the brain regions wound up being completely disconnected. Moreover, Kraft suspects that the extreme

Sagittal View   **Transverse View**

A   24 hrs   B   24 hrs

C   48 hrs   D   48 hrs

E   72 hrs   F   72 hrs

G   96 hrs   H   96 hrs

Finite element simulations coupled with cellular death predictions are used to specify injuries to white matter and subsequent damage over time. Damaged edges are shown in red, and node size increases as connections are lost. The predicted evolution of damage is shown for 24 (a and b), 48 (c and d), 72 (e and f), and 96 hours (g and h). Reprinted from Kraft RH, Mckee PJ, Dagro AM, Grafton ST, Combining the Finite Element Method with Structural Connectome-based Analysis for Modeling Neurotrauma: Connectome Neurotrauma Mechanics. PLoS Comput Biol 8(8): e1002619. doi:10.1371/journal.pcbi.1002619 (2012).

damage predicted at the outer limits of the model doesn't actually occur in nature, suggesting the existence of a protective or regenerative mechanism that has yet to be defined mathematically.

Kraft is already working on a follow-up paper that will use the model to examine network response and degradation in the face of blast injuries, which differ mechanically from physical impacts. Ultimately, he'd like to see this kind of biomechanical/connectomic model used to predict the likelihood that a particular person might be susceptible to long-term neurodegenerative disease following traumatic brain injury of any kind, whether sustained on the battlefield or the gridiron. That information could form the basis of a "pro-

team, or wearing a particular kind of helmet.

It's the kind of idea—practical, useful, maybe even possible—that makes a proposal to map the human connectome seem like something that even the most skeptical critic could support. □