For a century, neuroscientists have dissected, traced, eavesdropped on, and are now compiling a seemingly endless cast of players in the nervous system. As we keep gathering more and more molecular details, how do we know when we know enough?

Reverse Engineering the Brain

By Roberta Friedman, PhD

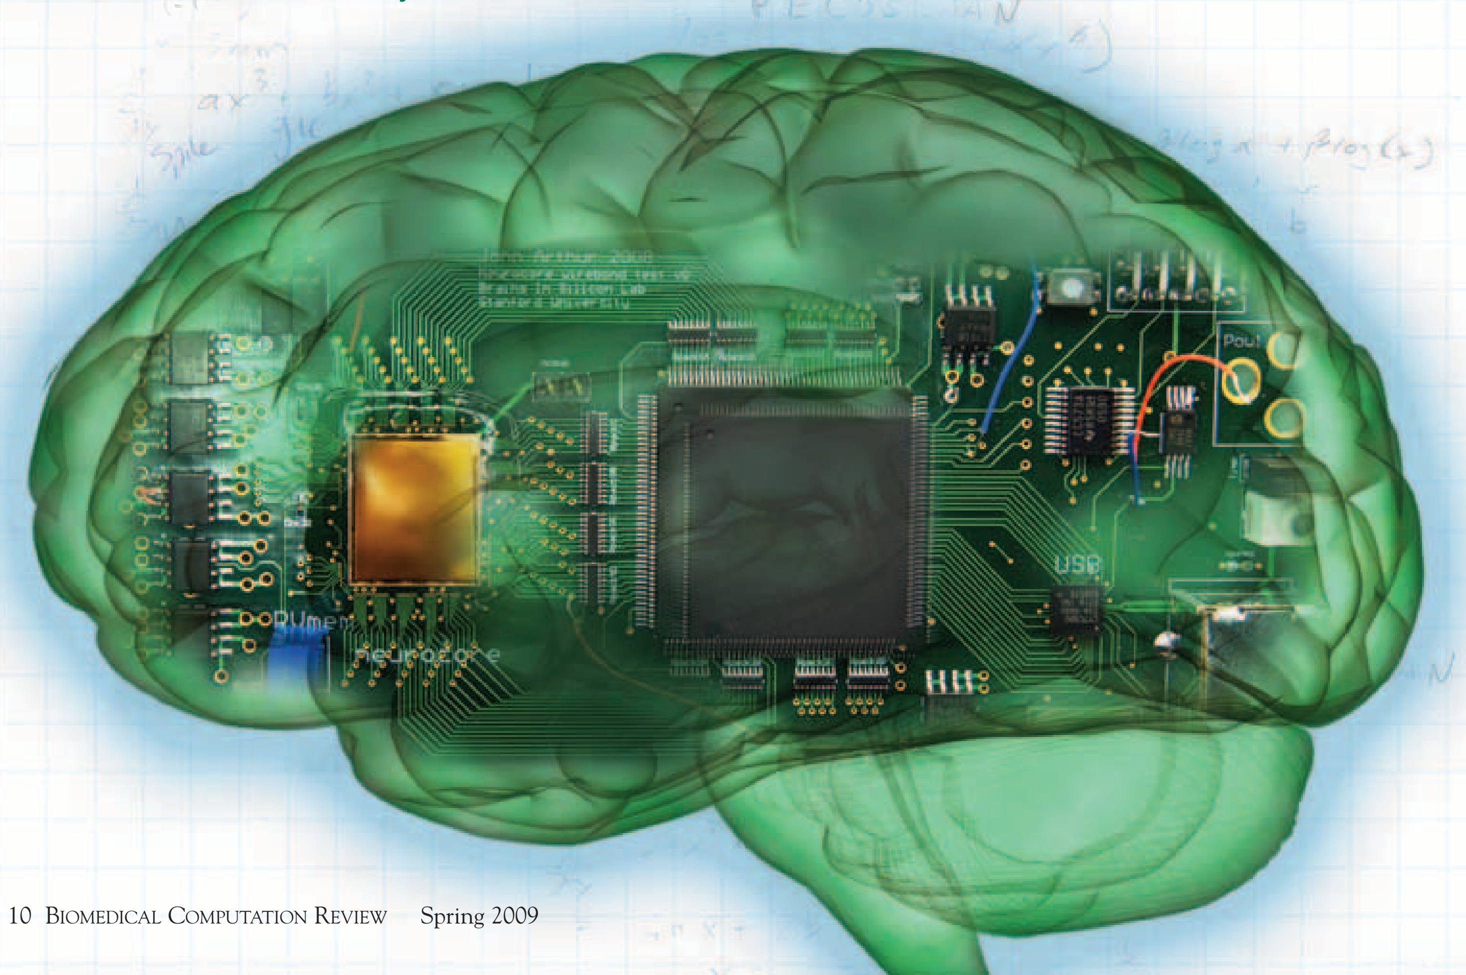

Some have decided it's time to just go ahead and create a brain *in silico*. And to a surprising extent, they've done it: Labs around the world are populated with autonomously functioning brains based on what we know so far. These simulations match what happens at the cellular level in the brain when the nerve cells, or neurons, that make up the brain pump ions and produce electrochemical activity that propagates across the synapse from one neuron to another. Robots or avatars activated by these engineered brains are directing movement, perceiving visual objects, and even responding to rewards—exhibiting behaviors associated with our "thinking" brains.

Eerily, the most recent simulations show the same oscillating rhythms seen when physicians record human brain waves using an electroencephalogram (commonly known as an EEG).

Computer simulations of the brain already allow experiments impossible to carry out with animals. "As good as modern neuroscience is—and it has been brilliant over the last two decades—we can't really sample every neuron and every synapse as they are performing a behavior," notes consciousness researcher **Gerald Edelman, MD, PhD**, director of the Neurosciences Institute and chair of neurobiology at the Scripps Research Institute in San Diego, California.

Researchers are looking to develop even more efficient simulated brains to help produce computers that can think while at the same time accelerating neuroscience. Ultimately brain simulations promise the ability to study the effect of drugs and disease and aid in the design of new therapeutic strategies.

HOW TO BUILD A BRAIN 101

To build a simulated brain requires a vast amount of detailed information about this complex organ, starting from its basic unit (the neuron) and building up to the complex network of connections between them that produces perception and cognition. None of this information is available from any single species. Much of the data on how individual neurons behave comes from rat studies. Observations of primates have provided data about how neurons are wired together across brain regions. And cat and human research led to an understanding of the finer, local circuitry in specific areas of the brain. Nevertheless, the basics of the nervous system are similar enough across mammals that Edelman and others have cobbled together chimeric, rudimentary brain simulations that show remarkable similarities to the real item. "We can simulate the neuronal dynamics beautifully so that you can't tell the difference between neurons in the model and real neurons," Edelman says.

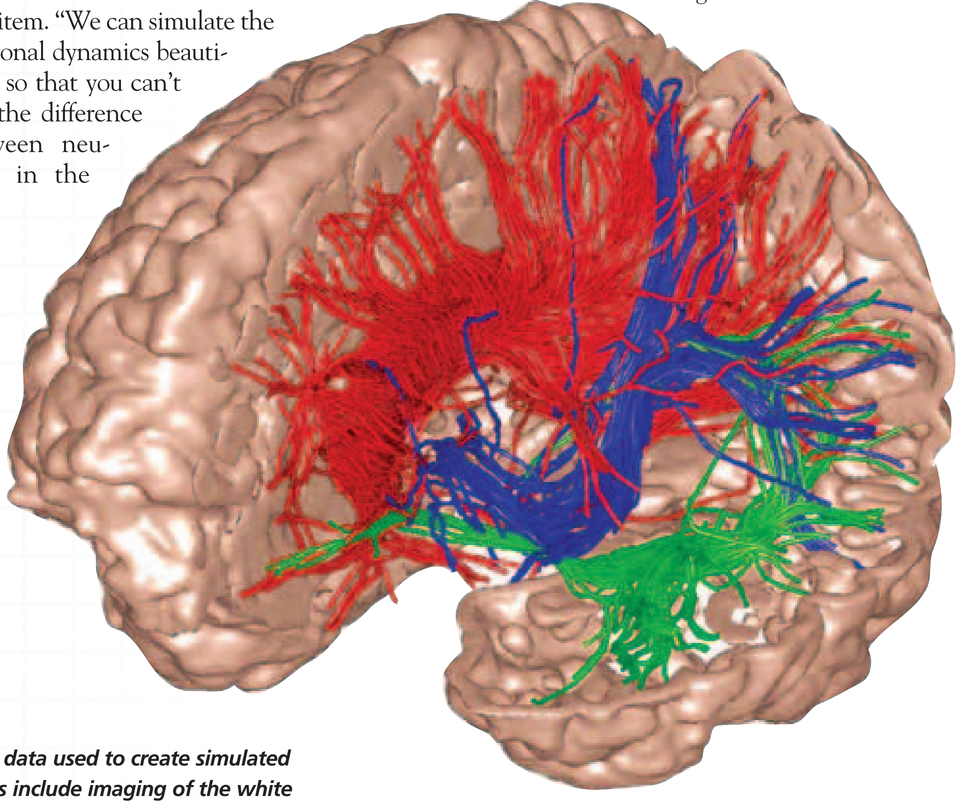

Brain data used to create simulated brains include imaging of the white matter fibers in the brain using a technique called diffusion tensor magnetic resonance imaging (DTMRI). Reprinted from Izhikevich et al., Large-scale model of mammalian thalamocortical systems,* Proceedings of the National Academy of Sciences *(2008) 105:3593-3598.

To build a simulated brain, Edelman and others start with what's known about the neuron, a cell that actively maintains a separation of charged ions across its membrane. Specific channels in the membrane allow certain ions in, and these are quickly pumped back out, or sequestered internally. But when a certain threshold of charge is reached the neuron fires a spike of current toward an adjacent neuron.

Here, at the synapse—a microscopic gap between each nerve cell—current becomes chemistry (and here is where drugs alter that chemistry). A spike wave arriving at the synapse triggers the release of neurotransmitters—to activate the next cell—provided enough inputs arrive in a very short time. Sufficient impulses strengthen the synapse. Neglected, the synaptic strength weakens and the particular connection is diminished. In the developing brain, synapses are ruthlessly pruned. This is what neuroscientists have uncovered during decades of listening in with electrodes a hundred times finer than a human hair. And this is the basic information that Edelman and others use to construct their simulated neurons.

To determine how these neurons are connected, simulators turn to microscopists and their latest technologies. Techniques from immunology have brought incredible resolution on the molecular level: cells containing particular molecules can be tagged by dye-bearing antibodies so that researchers can distinguish them from from their

> "We can simulate the neuronal dynamics beautifully so that you can't tell the difference between the model and real neurons," says Gerald Edelman.

The simulation of major brain centers and their microcircuits is able to generate its own inherent activity—similar to what is seen in real brains.

fellows and follow their links to one another. Scanning electron microscopy has been able to home in on the fine molecular scale at the synapse.

Knowing how individual neurons function and how they're connected will not make a brain work. Simulators need to know the bigger picture of brain area networks. To understand the function of brain regions, neuroscientists initially used data from scalp EEG and depth electrodes placed within the brains of living patients and animals, as well as observational reports such as from accidents that selectively damaged specific brain areas. These days computer-analyzed imaging can reveal additional details of the normal brain. Simulators employ all of these lines of evidence, and still seek more.

But none of this data could produce an engineered brain without huge advances in computer simulation. Alan Turing's idea for a calculating machine at the end of World War II laid the groundwork. Warren McCulloch and Walter Pitts set forth the initial properties of an electronic replica of a neuron in 1943. In the mid 50s, IBM researchers ran a simulation of 512 neurons.

These are the lines of investigation picked up by Edelman who entered the field of reverse brain engineering after receiving a Nobel Prize in Physiology or Medicine (for immunology research) in 1972.

ACHIEVING AUTONOMY: EDELMAN'S SIMULATED BRAIN RHYTHMS

The latest of Edelman's simulations incorporates the known circuitry from the thalamus, a central command post in the core of the brain. The thalamus drives the cortex (the brain's covering layers—also modeled) into and out of sleep and through various levels of alertness. When the thalamus no longer talks to the cortex, vegetative states result. The model also includes circuitry of the hippocampus, a seahorse shaped curl of brain tissue beneath the temples, which is crucial for long-term memory, a region attacked in Alzheimer's disease.

Once enough of the brain's macro and microcircuitry is simulated, the *in silico* model is able to generate its own inherent activity—similar to what is seen in real brains. "When you stimulate the neural model, it takes off on its own and is constantly active," Edelman says. "We've never succeeded in doing this before." Moreover, oscillating waves of synchronous neural firing not explicitly built-in emerged spontaneously, the researchers reported in the March 4, 2008, *Proceedings of the National Academy of Sciences*. The researchers also were able to induce and reproduce spontaneous, low-level activity at the synapses—called miniature postsynaptic potentials or minis. The results suggest that, as a real brain develops in a fetus, minis like these might prime neurons for action.

EAVESDROPPING ON SIMULATED NEURONS: THE BLUE BRAIN PROJECT

Edelman is not alone in simulating the brain. **Henry Markram, PhD**, co-Director of École Polytechnique Fédérale de Lausanne (EPFL), in Lausanne, Switzerland, directs the data-intensive Blue Brain project.

Edelman's group relied on a top-down approach based on global network properties of the brain and mathematical formulas to reproduce known types of neuron behavior. In a complementary approach, Blue Brain focuses

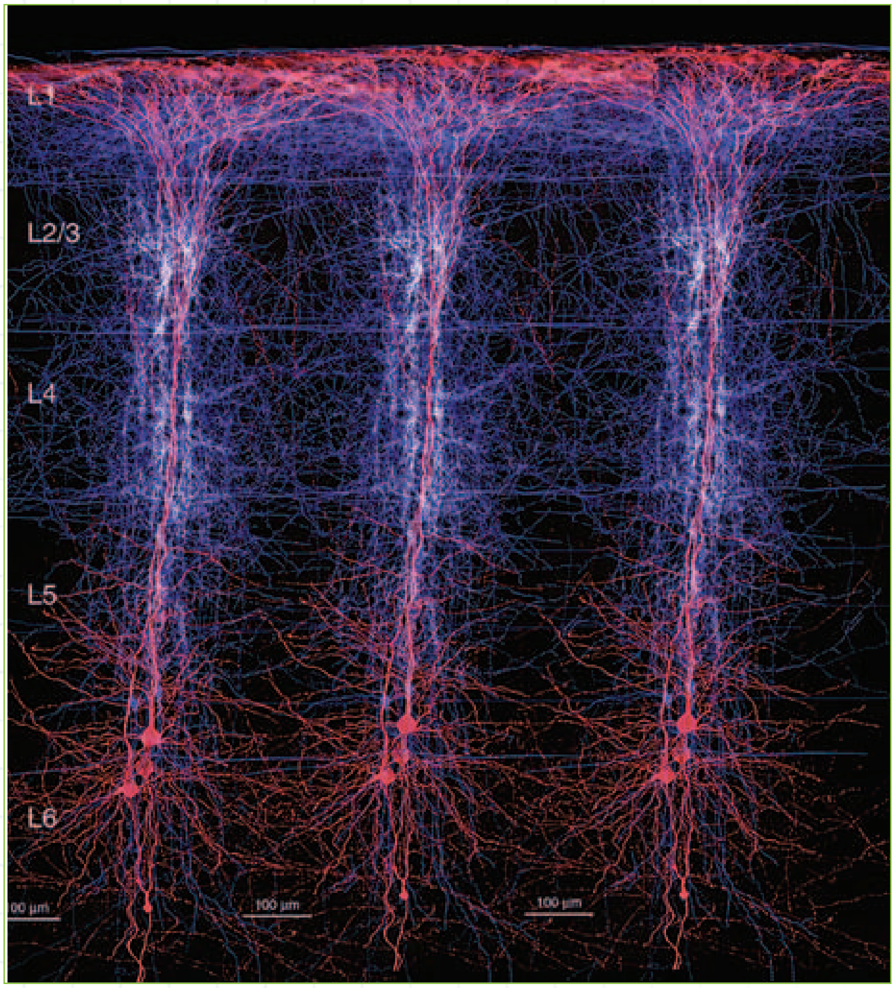

Cells in the cortex form columns. In this image the red neurons, called pyramidal cells, are revealed to be entwined by blue fibers from other, inhhibitory neurons that slow their firing. The layers of the column are indicated by the numerals to the left: L1, at the surface of the brain, through L6, the deepest cortical layer. Pyramidal cells, which receive messages along their extensively branched fibers, and send long fibers out to other brain areas or down to the spinal cord, are crucial in movement control and in cognition. They have their cell bodies in layer 5 of the cortex, and the main receiving fiber, the apical dendrite, rises up to the surface, layer 1. ©BBP/EPFL

on exact structural and molecular details to model a particular piece of the brain, building up from exact details of individual neurons, Markram says. “We are constrained by biology. There are so many theoretical ways to do it you would be lost forever. Biology has chosen a certain way and when you choose that, it becomes easier, not more difficult.”

“We are constrained by biology,” says Henry Markram. “There are so many theoretical ways to do it you would be lost forever. Biology has chosen a certain way and when you choose that, it becomes easier, not more difficult.”

Data for the Blue Brain project was gathered using a key innovation: the ability to record ion signals from many neurons at once using what’s called a multiple unit patch clamp technique. By eavesdropping on the interactions among neurons, researchers learned what synaptic currents were being generated and where.

In addition, they gathered data on gene activity within neurons—as an indicator of which discrete ion channels are present. In most neurons, a dozen or more types of these pores regulate ion flow. The Blue Brain simulation specifies which ones are present in each neuron. They also captured the precise connecting points of each neuron, by injecting dye once they were done recording the electrical activity. “The details are accurate, down to the micron,” for each contact point of each nerve fiber, adds **Phil Goodman, MD**, professor of Internal Medicine and Biomedical Engineering at the University of Nevada, Reno, who collaborated on Blue Brain. “It is a simulation, in time and space, of cells in real life.”

So far, the project has reproduced the architecture and electrical properties of a single cortical column of a two-week-old rat. The living cerebral cortex is comprised of millions of such columns, with each column consisting of a vertical stack of six layers of over 400 types of neurons. The cortex column is has a blueprint which is quite similar from mouse to man and across brain regions with only subtle variations.

The Blue Brain researchers can probe the simulated cortical column with simulated electrodes. As in Edelman’s lab, once a few stimulations are fed in, the simulation keeps going with its own intrinsic activity. For example, if thalamic fibers arriving at a deep layer of the cortical column are stimulated, the activity spreads, and finally the most superficial layer lights up. Markram notes that laboratory experiments failed to make this observation because they failed to listen in at the superficial layer. Thus, the simulation has already generated observations that could easily be missed in the lab, suggesting how simulations can guide brain research.

In the next six months, the Blue Brain project plans to publish “key insights never seen before in the neocortical column,” Markram says. “By the end of the year we will publish the entire circuit with the blueprint. It’s like the genome map—it’s a comprehensive description of the neocortical column.”

“It took 15 years to get the data for this small piece of brain,” Markram says. “Every week the model becomes more biological,” he adds. “It’s very much like a real little bit of tissue.” And now that they’ve built one cortical column, building another is a simple task. “We can (now) push a button and build an unlimited amount of neurons automatically.” The goal is to build up the brain from this discrete piece.

“We can (now) push a button and build an unlimited amount of neurons automatically,” Markram says.

There’s still a need for more data about brain anatomy, Markram says. Some neuroanatomists are working toward a map to locate every single neuron in the human brain. This so-called “connectome,” says Markram, will undoubtedly help the next generation of brain simulation. **Javier DeFelipe, PhD**, from the Cajal Institute in Madrid has joined the project to provide Blue Brain with data at the electron microscopic level. “Blue Brain is hungry for data,” Markram says.

A POCKET-SIZED SIMULATED BRAIN: NEUROGRID CHIPS

To simulate the human brain, to really know how we think, is not a research problem many can take on. Electricity alone for a supercomputer to simulate a million neurons eats through $200,000 a month, restricting brain simulations to the very few able to get that kind of funding. “This is something we want to change,” says **Kwabena Boahen, PhD**, associate professor of bioengineering and the principal investigator for Brains in Silicon at Stanford University. “If we can create a tool to allow a lot of people to play at this scale, as a community, we will progress faster.”

To that end, Boahen and members of his lab have developed the Neurogrid chip with funding from the NIH Director's Pioneer Award Program. No bigger than a fingernail, 16 of these chips will be assembled in an iPod-sized device that can do what a supercomputer does—simulate a million neurons—at only $40,000. The Neurogrid chips have been received from the silicon foundry and should allow the group to emulate a million neurons in the cortex in real time at a thousandth of the cost of supercomputing. “Everybody can play now,” says Boahen. “Not just IBM.”

The Neurogrid chip works the same way the brain does, Boahen says. Its circuitry is analog because that is the way neurons compute: They sum their inputs continuously, not discretely. It is only past a certain threshold that the process

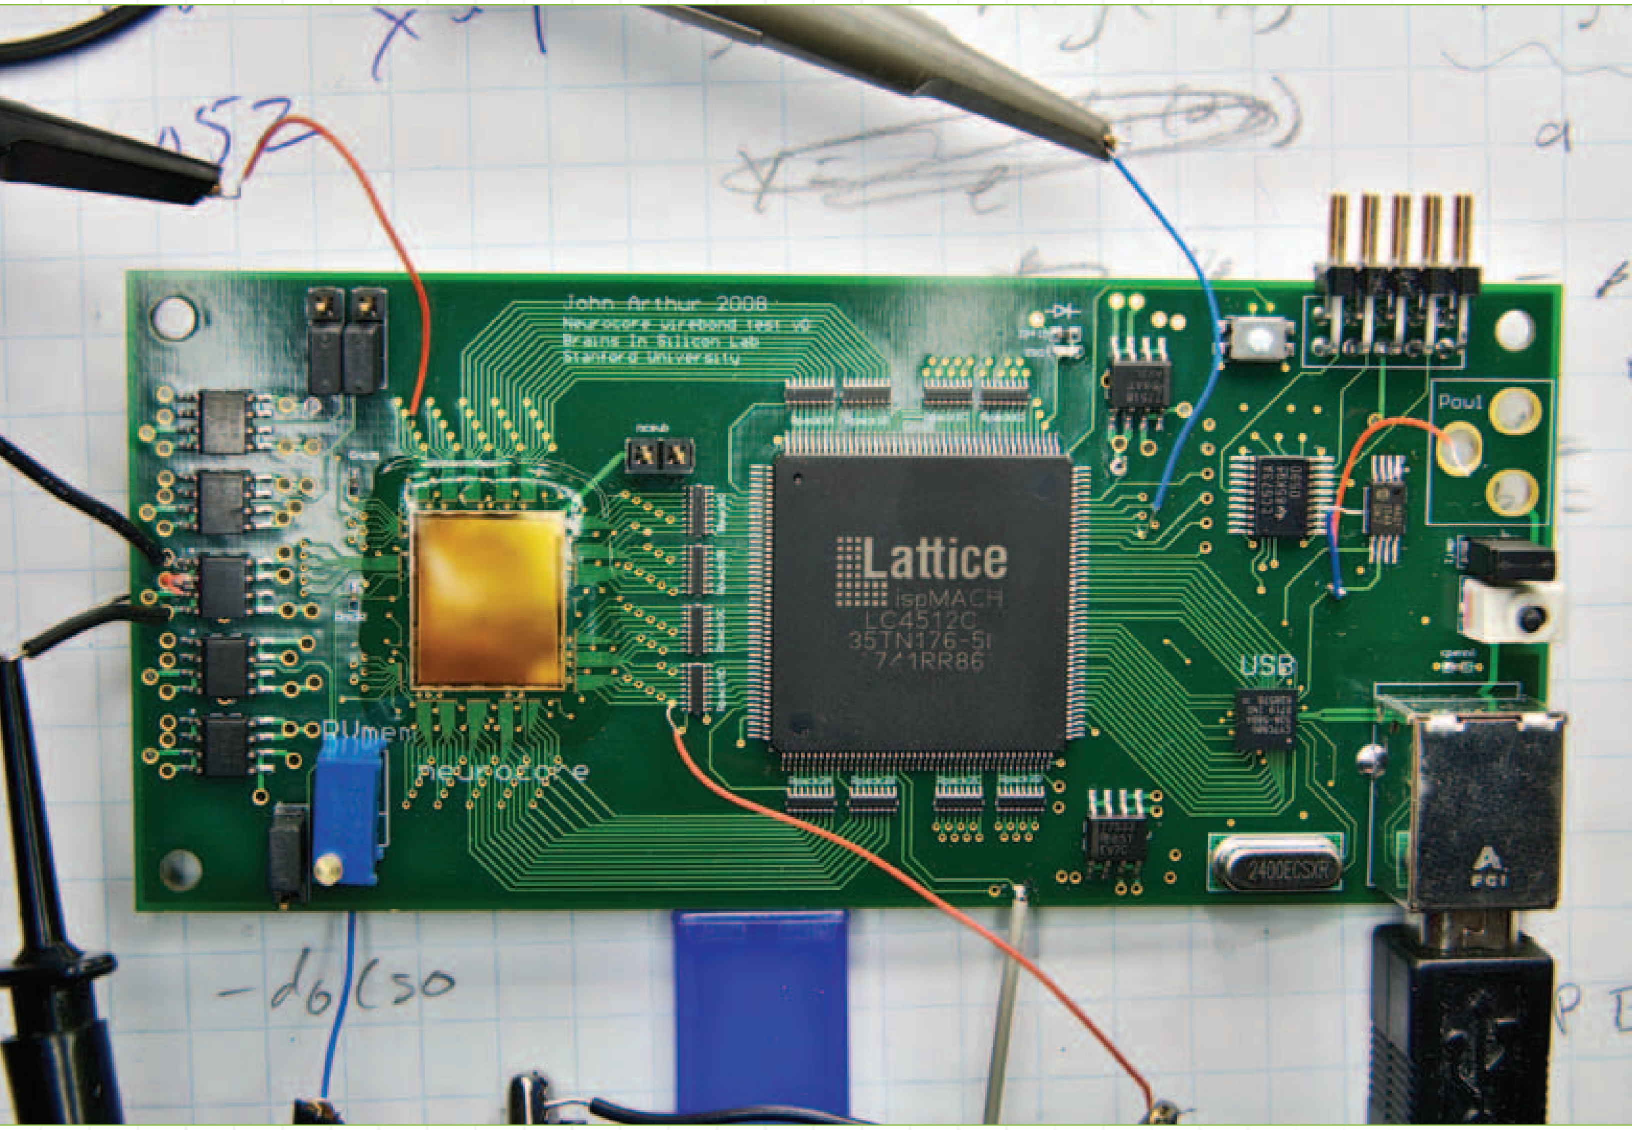

A Neurogrid chip (Neurocore) mounted on a test printed circuit board. Each Neurocore has 65,536 programmable neurons in 162 mm^2 of silicon. Sixteen Neurocores connected together will form the first hardware system with over one million model neurons operating in real-time, while consuming less than 10 Watts and taking up less space than a paperback book. Courtesy of Rodrigo Alvarez and Kwabena Boahen, Brains in Silicon, Stanford University.

becomes digital, generating a spike of electrical activity—all or nothing (the spike is like a one; lack of spike, a zero).

"Instead of using transistors as switches, I can build a capacitor and sum currents and get the same voltage on the capacitor that a neuron makes," Boahen says. With one transistor and a capacitor, he says, you can solve a differential equation that would take a thousand transistors in the traditional arrangement in a computer.

Dharmendra S. Modha, PhD, manager of cognitive computing at IBM's Almaden Research Center in San Jose, and a collaborator of Boahen's, says "Neurogrid is a genuine technical breakthrough. It has the potential to transform computational neuroscience."

Modha cites the mouse cortex model that his team has created as a prime example. Their simulation shows the oscillations present in living brains just as Edelman's do and runs "in near real time" on a 4096 processor BlueGene/L supercomputer with a terabyte of memory. Modha explains that even so, that was still seven to ten times slower than the action in the rodent brains.

Obviously, the requirements for brain simulation outstrip the available hardware unless alternatives such as Neurogrid or others are achieved. "The

> "If we can create a tool to allow a lot of people to play at this scale, as a community, we will progress faster," says Kwabena Boahen, who, with colleagues, has developed the Neorogrid chip.

brain that Mother Nature has created is enormously complex," Modha says. "Any attempt to emulate it is always a radical simplification."

Edelman, too, is looking for ways to simulate the brain using less computing power. He can currently simulate 10 million neurons and half a billion synapses. But the human cortex has at least 3,000 times that many neurons and almost a million times more connections. He says his group has designed and built their own completely new computer architecture in order to be able to add regional microcircuit details into their generic cortex simulated so far. Their simulations to-date have used a Beowulf cluster of 64 interactive processors. "Beowulf is seven feet high and 250 to 300 pounds," Edelman says. The new system—which hasn't yet been described in a published paper—"is about 10 inches by three inches and weighs a few pounds. It can be stuffed inside a brain-based device and is more powerful."

Markram is also starting to feel constrained. "Our BlueGene supercomputer is only just enough to launch this project. It is enough to simulate about 50,000 fully complex neurons close to real-time. Much more power will be needed to go beyond this."

SIMULATED BRAINS IN THE REAL WORLD: THROW IT A BONE

Simulated brains on computers may be interesting research, but like real brains, they are best understood by how they respond to the real world. To test simulated brains in real world settings, some researchers, such as Edelman, use robot-like devices; others use computer avatars; and still others, with a focus on computer vision, struggle to achieve object recognition.

Edelman emphasizes that real world interactions have shaped brain evolution. He has formulated a theory he calls neural Darwinism focused on the role of reward as a driver of brain evolution. "The brain is embodied, and the body and brain are embedded in the real world environment," Edelman says. "And that environment, enormously rich, provides the reward that drives real brains to make choices."

Edelman has tested this theory using "embodied" devices run by a brain-based network. These brain-based devices, called "the Darwin series," are fitted with cameras and microphones that serve as their eyes and ears, and they can sense conductance ("taste") between their grippers.

Darwin VII ran on a brain simulation that included elements of the reward circuitry of the mammalian brain. The knee-high device started out randomly picking up little blocks placed in its roaming zone. One set had stripes, the others, spots. One pattern had high conductance, the other, low conductance. High conductance was arbitrarily rewarded, strengthening the appropriate connections in the brain-based network. This eventually led the Darwin to pick up only this one type of block and ignore the other.

As in the brain, strengthening and weakening of synapses determines if neurons next in line will fire, Edelman says. "Just like synapses act in real brains, the next one won't necessarily fire unless enough stimulation occurs."

Experience changed the synaptic strength. In other words, the Darwin learns.

This reinforced behavior is exactly equivalent to how mammals learn to choose what to eat based on taste.

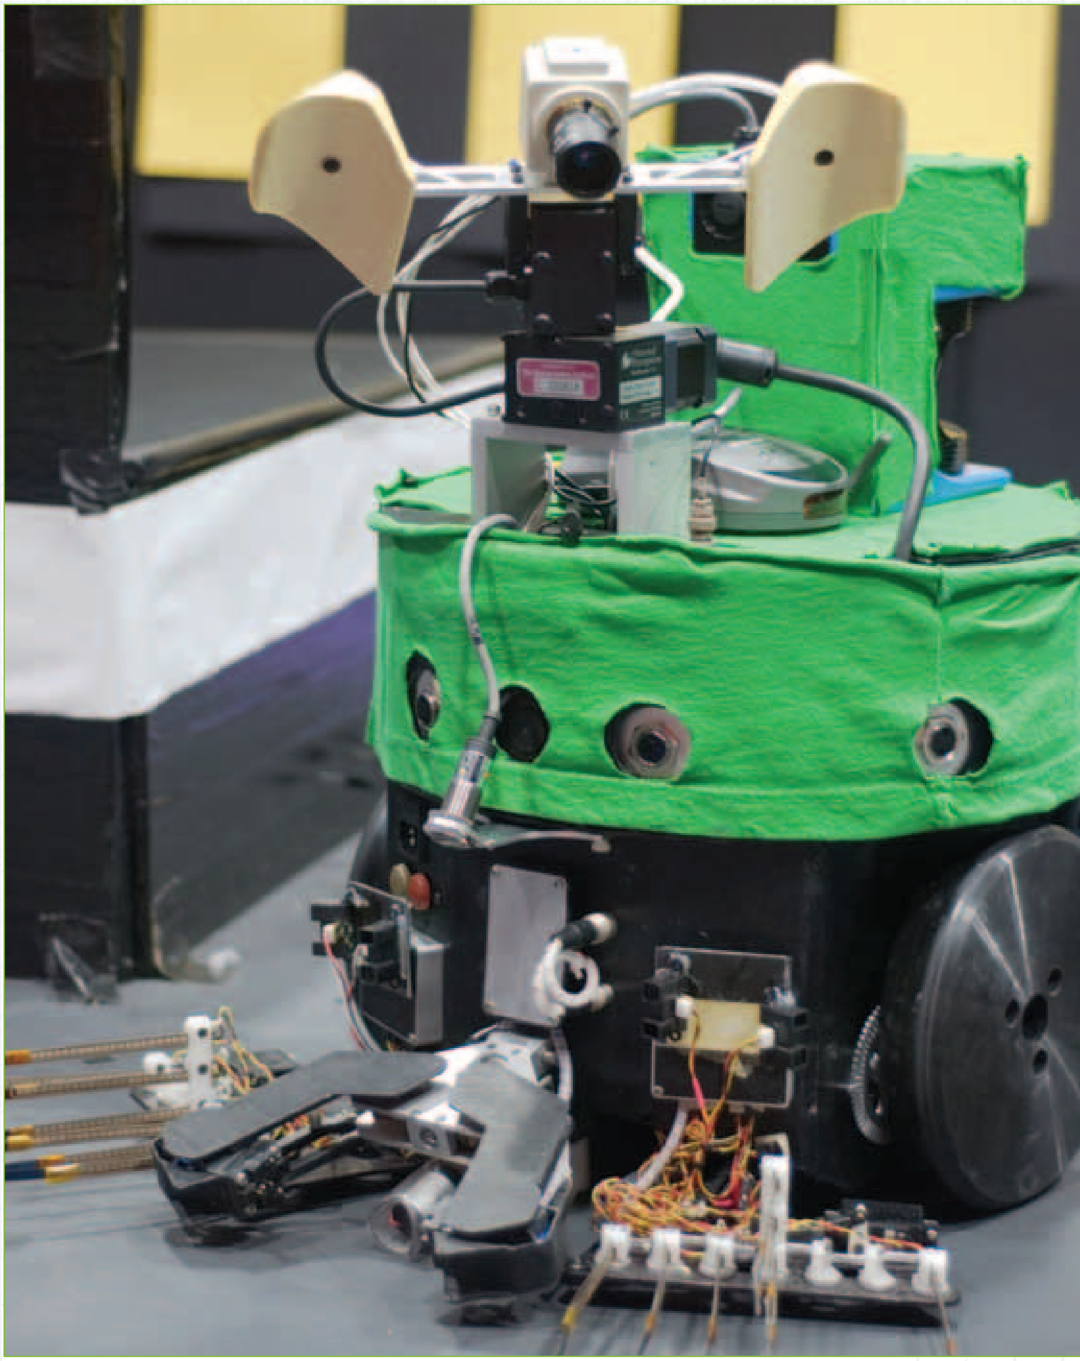

As Darwin XI learned to navigate mazes, its hippocampus exhibited responses similar to those seen in rats engaged in the same task. Courtesy of Jason Fleischer/The Neurosciences Institute.

After all, taste is a random function of the chemicals in food detected by the olfactory system. The Darwin's sensing of the conductance was equivalent to the mammal's ability to taste food.

"The world is not a coded piece of tape. It can't be explicitly contained in an algorithm," Edelman says. A brain-based device, with a value system, learns by making mistakes. "Hook that to a Turing machine and what you will get is not artificial intelligence, but an entirely new machine," he says—for example, an aerial drone that could decide on its own about threats.

> "The brain is embodied, and the body and brain are embedded in the real world environment," Edelman says. "And that environment, enormously rich, provides the reward that drives real brains to make choices."

Though one might think the Darwin device hovers on the brink of consciousness, a lot still separates these simulations from actual brain processes. Phil Goodman emphasizes the role of intention and emotion in mammalian brain action. He embodies simulated brain circuitry through projection of a virtual device, an avatar, similar to what video gamers are used to seeing and controlling.

One of his avatars is a dog with pre-programmed behaviors: It starts out lying down, gives a threatening bark while sitting up, or engages in panting and tail-wagging while standing. Sensors allow the simulated brain that is steering the avatar to see and hear. So much of communication of emotions is subliminal that Goodman says, "if (an avatar) is to be social, it needs to interact with our own bodies." So his model incorporates aspects of the emotional processing regions of the brain, the so called limbic system.

A supercomputer runs programs that process sensory input, producing probabilities of neuronal firing, which in turn trigger behavior. A stranger's posture and actions elicit the appropriate reaction of the projected canine. Upright posture with a raised arm will trigger the threat behavior. The happy reception is elicited by crouching with soothing words—and petting on a touch pad.

FEEDING THE WORLD INTO THE BRAIN AND BACK AGAIN

The brain remodels itself in response to perceptions through its sense organs. Thus, simulators need to tackle these brain accessories as well in order to recreate cognition.

Object recognition is vital for a virtual or a material brain-based device such as the Darwin series or Goodman's avatars. Yet it has been one of the most challenging tasks for artificial intelligence. Goodman uses fairly primitive visual processing in his model, but **Thomas Serre, PhD**, a postdoc working with **Tomaso Poggio, PhD**, at the Massachusetts Institute of Technology, has recreated in a machine the ability to perceive objects when flashed at the threshold of human visual perception. Remarkably, the simulation performs as well as people (as described in a News Byte in the Summer 2007 issue of *Biomedical Computation Review* http://biomedicalcomputationreview.org/3/3/4.pdf).

Serre's experiment was limited, however, to the brain's response to an image flashed for less than 150 milliseconds. Thus, it provides just a skeleton of a complete theory of vision, Serre says. He's now working on what happens beyond the first 150 milliseconds of visual processing—"when you move your eyes and shift attention."

The visual system involves a complex of more than 30 brain areas propagating signals from the retina through the visual cortex to the region of motor cortex that controls how the person (or the simulator) responds. Living brain also contains back projections, echoing all the way back to the primary visual area that receives the initial signals

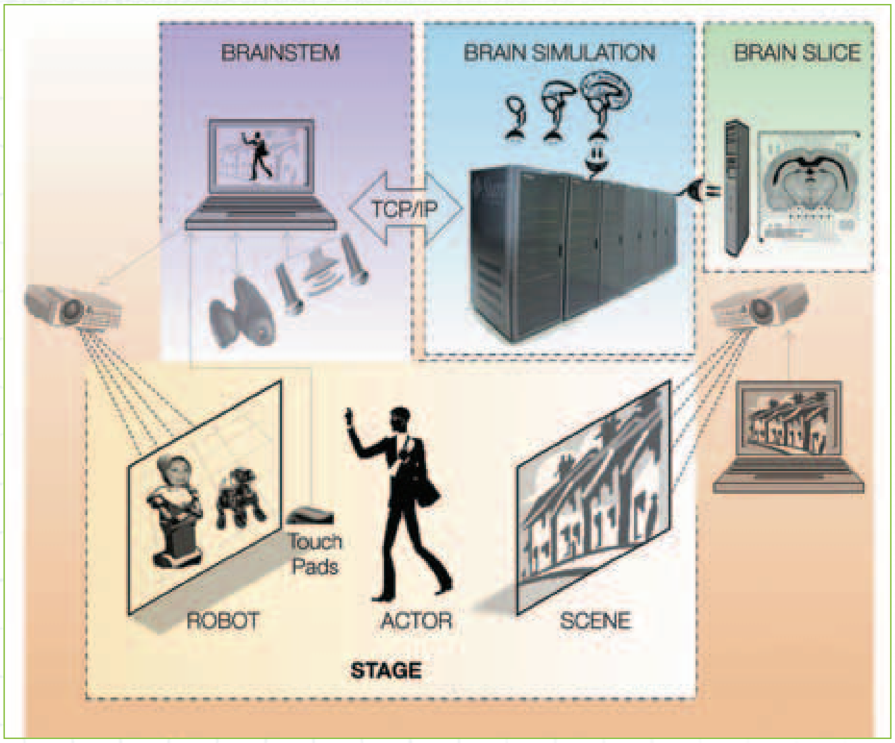

Schematic of a virtual neurorobotic system. By creating an avatar of a robot, Goodman and his colleagues avoid the complex engineering of the physical robot. The virtual robot can still respond to environmental stimuli provided through a mouse pad, microphone and camera. Reprinted from Goodman et al., Virtual Neurorobotics (VNR) to Accelerate Development of Plausible Neuromorphic Brain Architectures, **Frontiers in Neurorobotics, *(2008) 2:123:128.***

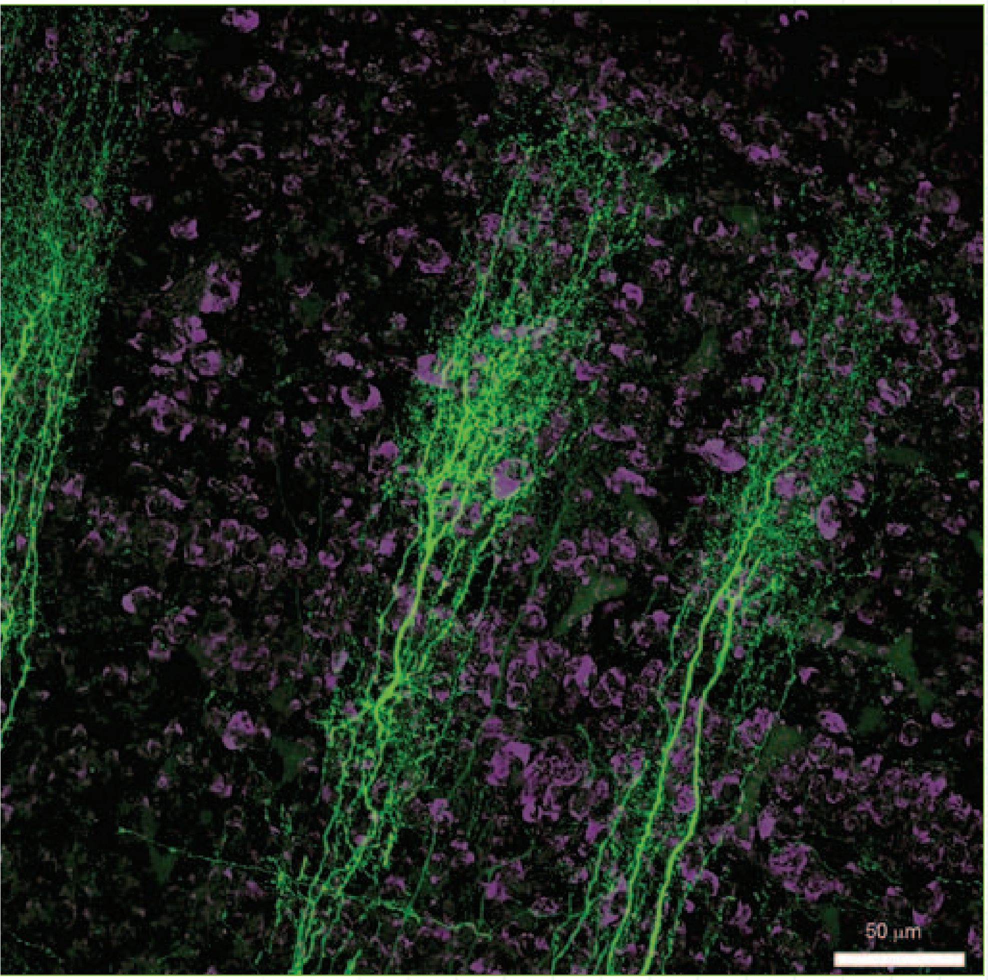

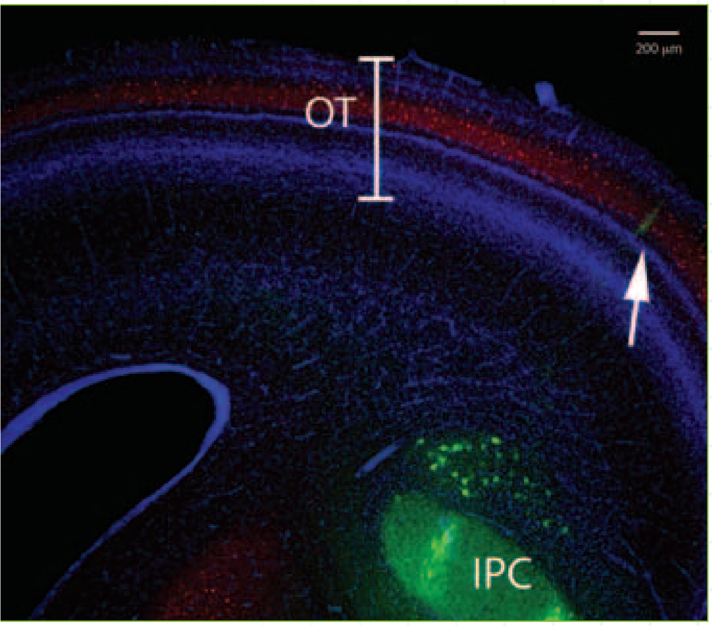

The optic tectum (OT) is a brain structure important for gaze in birds and still present in mammalian brains. Boahen's collaboration is building a model of the OT on a silicon microchip while parallel efforts by Eric Knudsen's team attempt to uncover its biological properties in living brain slices. Below left, we see a cross-section of the bird brain through both the OT and an area that connects with it, the isthmic nucleus (ipc, stained green). The arrow points to a green line marking the location of the close-up image shown at left. The close-up shows nerve fibers in the OT (stained purple) with cell fibers (axons) in green arriving from the ipc. Images courtesy of Dr. Alex Goddard, postdoctoral researcher in the Knudsen laboratory.

from the retina. Vision researchers suspect these back connections may be the way that the visual system can pick out a target object from complex scenes. "By adding back projections to the model, and allowing one shift in attention, to one part of the image, we are (now) able to mimic the next level of performance of a human observer when the image is left just 30 ms longer on the screen, just enough for people to shift their attention once," Serre says.

Boahen at Stanford heads a team working on recreating the basics of different parts of the perceiving brain. Much of the circuitry they plan to model will include back projections. Boahen agrees that feedback likely mediates attention, as competing firing is suppressed. As with other brain simulations, his also shows synchrony, the living rhythms of the brain, including gamma waves with attention.

To find out what the oscillations mean for visual attention, team member **Sridharan Devarajan** and Stanford neurobiologist **Eric Knudsen, PhD**, are working to understand the wiring in a barn owl's tectum, the brain area that controls gaze. Other collaborators are working with Boahen to simulate the wiring of the frontal eye fields in monkeys, an area that allows primates to gauge where attention is needed. This brain area evolved in the social setting of primate life, allowing monkeys to suppress a direct gaze at a superior monkey while still attending to what needs to be watched—covertly. These brain regions feed forward as well as back to higher and more basic levels of visual processing in the brain. Thus, simulations of this area will help researchers to understand the role of feedback circuits in perception.

COMPUTER CONSCIOUSNESS

Where are the eavesdroppers and engineers going with all this? Better business machines may be IBM's mantra. Modha's favorite saying is that the mind arises from the wetware of the brain. "The quickest and cheapest way to engineer mind-like intelligence into machines is to reverse engineer the structure, function, and dynamics of the brain" with its low power consumption and compact size, Modha says. "This is our quest."

Some may be scared by this quest. Others eagerly await the emergence of machine intelligence. Eventually, brain scientists hope to simulate the effect of strokes, tumors, or neurological disorders such as Alzheimer's or Parkinson's disease to understand how they derail brain dynamics.

Edelman states frankly his intention: to craft a conscious artifact. "Philosophers have owned the field of consciousness research from time immemorial. What could be more romantic, remarkable or valuable," Edelman says, "than to take on their quest? Right now, you might say, I am going for broke." □