

s2100009 Technology/Biology/Mind and Body/Artificial Intelligence/Brain Simulation Jeopardy! is Elementary for Watson Brain Training: Pain But No Gain Turning on the Lights: Imaging Consciousness Monkey See, Monkey Do — A Study of Decision-Making Artificial Intelligence Beats Human Genius Researchers Watch People Think World Almanac Encyclopedia: Brain World Almanac Encyclopedia: Nervous System World Almanac Encyclopedia: Neurophysiology 20130200

A TODAY'S SCIENCE SPECIAL FEATURE

Building a Brain in a Box

by *Katya Poltorak*

- [Brain Basics](#)
- [The Mind's Mysteries](#)
- [Neural Networks](#)
- [Slings and Arrows of the Robot World](#)
- [Spaun Specialties](#)
- [Spaun the Multitasker](#)
- [Copying Quirks](#)
- [A Sampler Plate of Brain Functions](#)
- [The Next Step](#)
- [Robots of the Future](#)
- [Chris Eliasmith: Making a Model Brain](#)
- [Discussion Questions](#)
- [Journal Abstracts and Articles](#)
- [Bibliography](#)
- [Keywords](#)

In *How the Mind Works*, world-famous cognitive scientist Stephen Pinker [pinker.wjh.harvard.edu/about/index.html] talks about the "fantastically complex design behind feats of mental life we take for granted." According to Pinker, "The engineering problems that we humans solve as we see and walk and plan and make it through the day are far more challenging than landing on the moon or sequencing the human genome."

The [brain](#) controls everything we do—from tying a shoelace to laughing at a joke to solving a differential equation or writing poetry. And while we generally don't think of these things as miracles, the mechanism behind the way we process the world around us is one of the most intricate things in nature.



Laguna Design/Science Photo Library/Getty Images

Studying how the brain works is an inherently interdisciplinary project. It encompasses everything from behavior patterns to nervous system operations. A Canadian research team has come up with a simulation of the functioning brain that helps fill in some of the gaps in our knowledge.

Studying how the brain works is an inherently interdisciplinary project. It encompasses everything from behavior patterns to nervous system operations. Putting all this together gets very tricky. After all, philosophers and molecular biologists don't always see eye to eye.

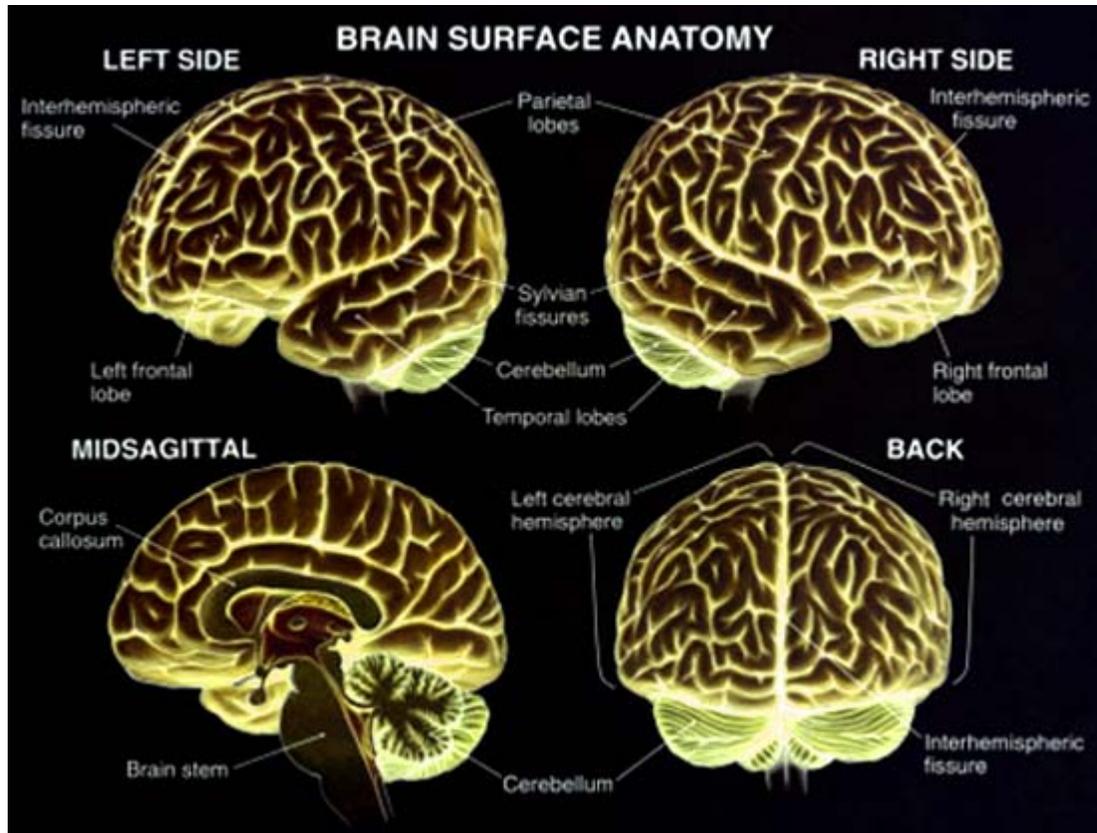
A Canadian research team led by [Chris Eliasmith](#) has come up with a simulation of the functioning brain that helps fill in some of these gaps. Known as the "Semantic Pointer Architecture Unified Network"—Spaun, for short—the model demonstrates the brain in action and shows how the activity of individual neurons gives rise to complex behavior.

Brain Basics

First, a bit of background. The [nervous system](#) is a set of structures allowing different parts of the body to communicate, with the brain, which serves as a type of supercomputer running the show. Sensory neurons pick up information about the environment—such as sounds or smells—and carry it to the brain. The brain then sends out instructions to muscles or glands via motor neurons, which initiate movement and allow us to walk, roller skate, play the piano, or do headstands in yoga class.

The largest part of the brain is called the cerebrum (or cortex). It controls the "higher functions"—thought and action—and is divided into lobes that specialize in particular types of tasks. For example, the frontal lobe is in charge of reasoning and planning, while the temporal lobe deals with memory and speech.

The other two parts of the brain are the cerebellum ("little brain") and the brainstem. Located at the base of the skull, the cerebellum regulates movement, posture and balance. Finally, the brainstem connects the cerebrum with the spinal cord and regulates "autonomic" functions, such as breathing, digestion or heart rate. Like a power strip, it also controls the "traffic" of signals between the cortex, spinal cord and peripheral nervous system.



MediVisuals/Getty Images

ABOVE: The largest part of the brain is called the cerebrum. It controls the "higher functions"—thought and action—and is divided into lobes that specialize in particular types of tasks. For example, the frontal lobe is in charge of reasoning and planning, while the temporal lobe deals with memory and speech. The other two parts of the brain are the cerebellum and the brainstem. Located at the base of the skull, the cerebellum regulates movement, posture and balance. Finally, the brainstem connects the cerebrum with the spinal cord and regulates "autonomic" functions, such as breathing. Like a power strip, it also controls the "traffic" of signals between the cortex, spinal cord and peripheral nervous system. **BELOW:** At the molecular level, this signaling mechanism is based on the interaction of nerve cells, or neurons. Each neuron has a cell body with small extensions called dendrites that surround it like tree branches. It also has an axon fiber that can be more than a meter long.



Left: Blue Brain; Right: Ecole Polytechnique Fédérale de Lausanne

At the molecular level, this signaling mechanism is based on the interaction of nerve cells, or neurons. Each neuron has a cell body with small extensions called dendrites that surround it like tree branches. It also has an axon fiber that can be more than a meter long.

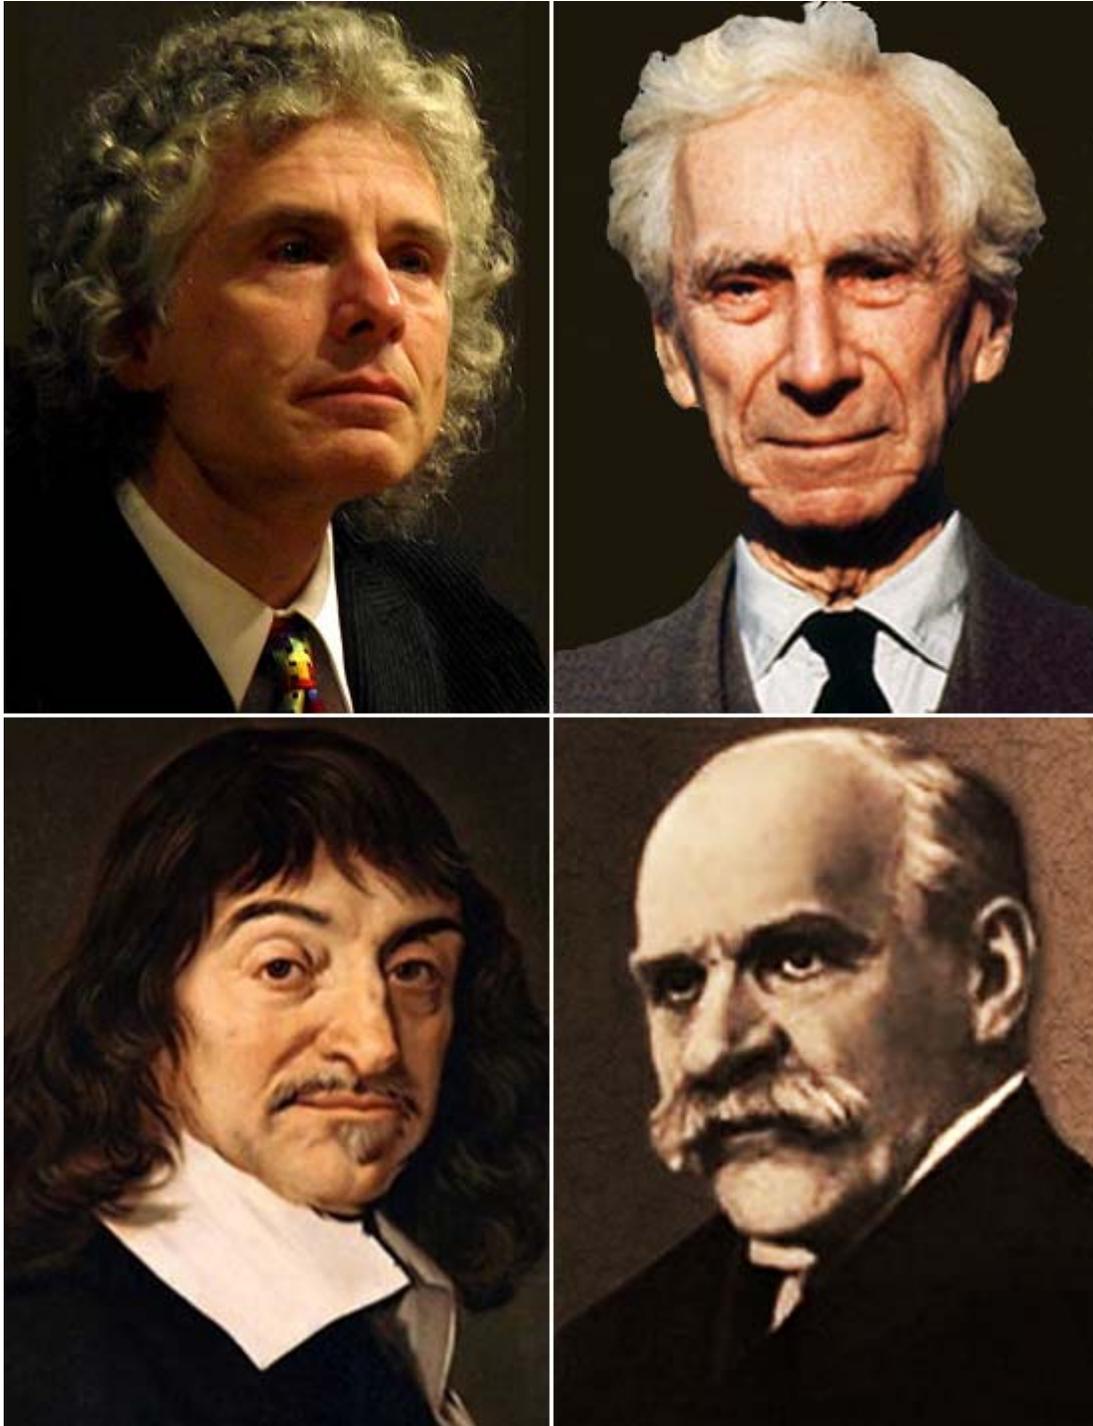
Messages travel along the axon as electrical signals caused by shifts in ion concentration along the cell membrane. At the synapse—the space where two neurons meet—the electrical signal gets "translated" into a chemical one: molecules known as neurotransmitters are released from the axon terminal of one neuron to be picked up by the dendrites of its neighbor. The "receiver" neuron then gets excited or inhibited, depending on what kind of transmitter is involved.

Since there are an almost endless number of ways for neurons to partner up, the nervous system becomes an intricate web of connections. Thus, the same basic elements can give rise to an enormous variety of patterns that give the brain its unique potential. ("Variety" is actually an understatement. The numbers we are talking about are astronomical: the cerebral cortex has about 30 billion neurons, which can create a staggering million billion synapses.)

The Mind's Mysteries

Because of this almost unimaginable potential, the brain in action is much more than the sum of its parts. Pinker compares it to the variety of effects produced by combinations of basic units in other systems: "In the same way that all books are physically just different combinations of the same seventy-five or so characters, and all movies are physically just different patterns of charges along the tracks of a videotape, the mammoth tangle of spaghetti of the brain may all look alike when examined strand by strand. The content of a book or movie lies in the pattern of ink marks or magnetic charges, and is apparent only when the piece is read or seen. Similarly, the content of brain activity lies in the patterns of connections and patterns of activity among the neurons."

But while we can describe the processes responsible for the brain's unique abilities, mapping structure and function is notoriously difficult. Decades ago, British philosopher [Bertrand Russell](#) showed the way in which even the simplest action—such as looking at the sky—involves a disjunction between matter and mind. After the light gets picked up by the eye, it sets off a chain reaction that ends up in the brain. And while "the whole chain of these events, from the sun to the top of [the brain] is physical" and "each step is an electrical reaction," the result is a change "wholly unlike any that led up to it, and wholly inexplicable by us..." "I see the dome of the sky and the sun in it, and a hundred other visual things beside...I perceive a picture of the world around me."



Clockwise from Top Left: Wikimedia; Bertrand Russell.org; Stephen J. Gould Archive; Wikimedia (Frans Hals portrait, 1648)

The mind's mysteries have been examined by philosophers and scientists throughout Western civilization. Among the thinkers who have wrestled with these mysteries are (clockwise from top left) Stephen Pinker, Bertrand Russell, J.B.S. Haldane and René Descartes.

What Russell was discussing is sometimes referred to as the "mind-body problem," and it has vexed thinkers as far back as the ancient Greeks. In the 17th century, philosopher [René Descartes](#) famously argued that mental and physical states are inherently separate. He came to this conclusion after setting out to make a list of things that we know for sure—a list that turned out to be rather small. Descartes found that since pretty much everything was filtered through the "thinking thing" called the mind, it seemed to be the only thing you can be completely certain about. Things that appeared real, such as the trees and grass outside your window, the voices of your neighbors, and so on, could be an illusion. Your senses could be deceiving you, or you could be dreaming. There was even a chance that some "demon" was trying to trick you by planting apparently "real" impressions in your mind. A slim chance, yes, but Descartes wanted to limit the things on his list that he was really, really sure

about.

This dualism persisted throughout Western intellectual history. Even as scientists continued to make progress in learning how the brain works, there still seemed to be a frustrating disconnect between matter and mind. Part of the paradox comes from the fact that we have to use our brain in order to study it. As [John Burden Sanderson Haldane](#) writes in *Possible Worlds*: "If my mental processes are determined wholly by the motions of atoms in my brain, I have no reason to suppose that my beliefs are true...and hence I have no reason for supposing my brain to be composed of atoms."

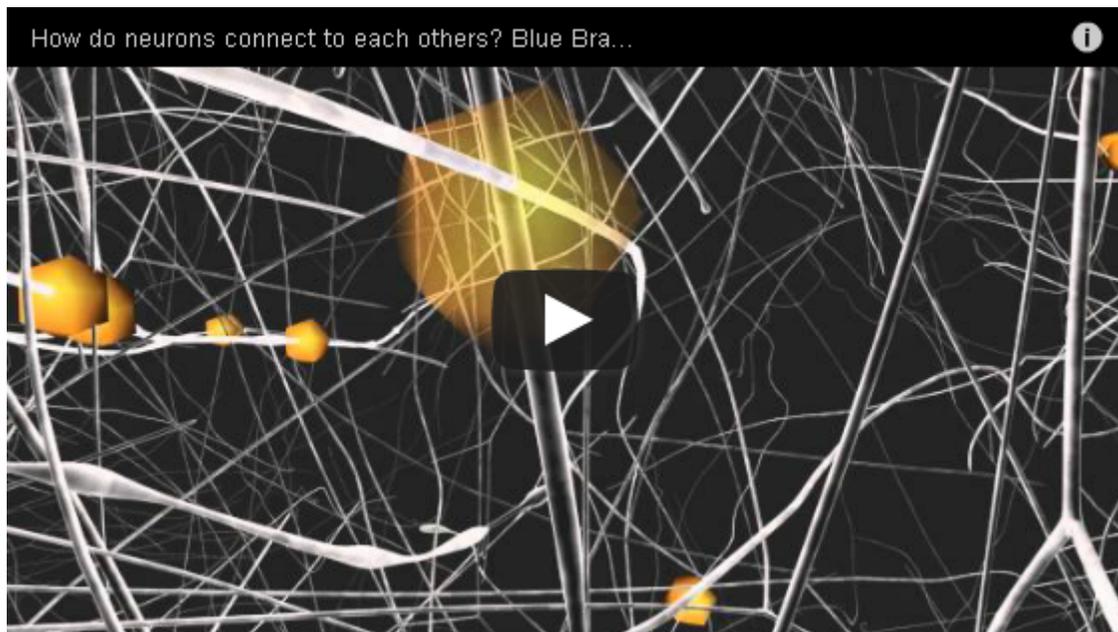
Neural Networks

Notwithstanding the profound philosophical questions that continue to be raised about the relationship between mind and body, on a more practical level scientists have made some progress in correlating mental activity with its physical underpinnings. In particular, they now believe that you can't "translate" particular brain structures into mental functions: instead, you have to look at neural activity as a whole.

Understanding the way physical processes in the brain give rise to particular functions is especially important for constructing computer simulations of the brain—a field sometimes referred to as artificial intelligence (AI). After all, you can't just throw the right "ingredients" together and expect them to self-assemble into a working brain: the key lies in the patterns these ingredients create.

Reproducing such patterns in a virtual brain is very difficult. As Eliasmith explains, "Throwing a lot of neurons together and hoping something interesting emerges doesn't seem like a plausible way of understanding something as sophisticated as the brain." It's a bit like the theory that a monkey will supposedly type up a play by [William Shakespeare](#) if given enough time to play with a typewriter. The catch is, of course, that you may have to wait forever for that to happen. [See [Monkeyssssssssss....](#), May 2003.]

University of Michigan scientist John Laird became one of the mavericks in the field of AI. In 1983, his research team developed a model of "cognitive architecture" known as "Soar" (State, Operator, and Result). More recently, a team led by neuroscientist Henry Markham came up with the Blue Brain Project—a supercomputer with microchips that are programmed to act as neurons. In late 2012, Markham and his research team reported their latest results in the *Proceedings of the National Academy of Sciences*. By comparing their simulation with actual mammalian samples, they were able to identify patterns behind synapse formation and map the physical arrangement of neurons in the brain.



EPFL News

In late 2012, neuroscientist Henry Markham reported the latest results of the Blue Brain Project—a supercomputer with microchips that are programmed to act as neurons. By comparing their simulation with actual mammalian samples, the Blue Brain Project was able to identify patterns behind synapse formation and to map the physical arrangement of neurons in the brain.

As Markham explains, "There are lots of models out there, but this is the only one that is totally biologically accurate. We began with the most basic facts about the brain and just worked from there." In a few years, Blue Brain will even make it to the big screen: a 10-part documentary about the project has been in the works since 2009 and is set to come out in 2020.

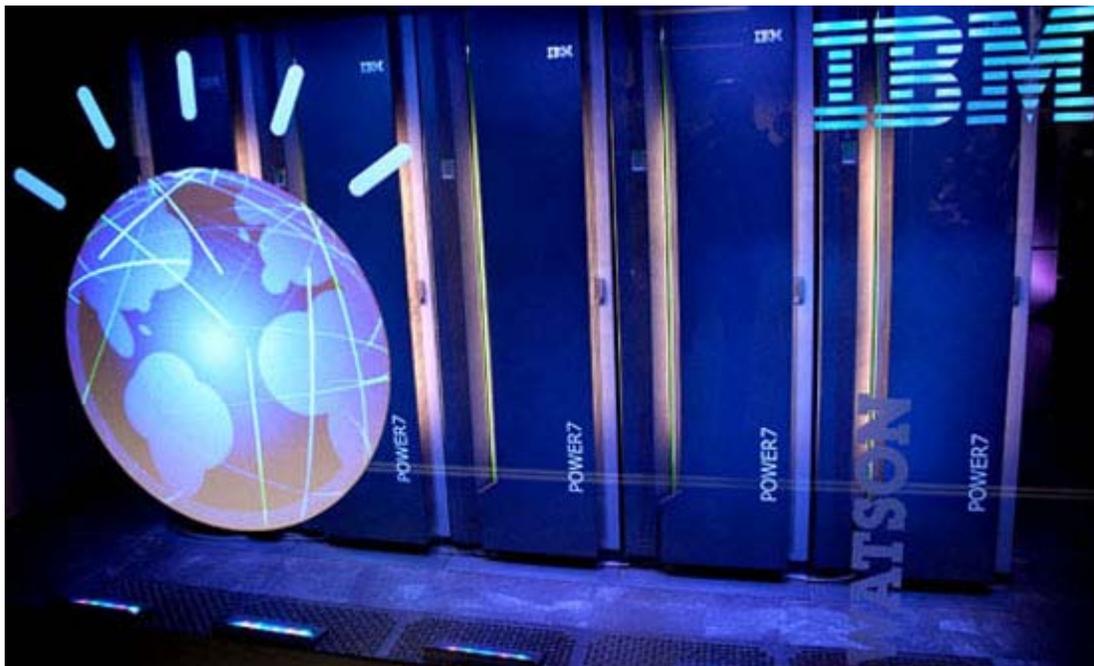
The development of Blue Brain also paved the way for an even more ambitious Human Brain Project—a long-term endeavor

aimed at modeling the entire brain. This "mother-of-all simulations" incorporates one million virtual neurons and illuminates the complexity of the connections between them.

Slings and Arrows of the Robot World

While many of these brain simulations are massive in scope and provide a detailed picture of neural networks, they don't show the relationship of physical signals to complex functions performed by the different parts of the brain. Referring to Blue Brain and SyNAPSE (another large-scale simulation), Eliasmith points out that "these artificial brains don't actually do anything. They don't remember, they don't recognize objects. They sit there and generate complex voltage patterns, but those complex voltage patterns aren't led to behavior."

The simulations that do perform some type of action are often limited to one particular task. For example, IBM's "Watson" is constructed to answer questions posed in natural language. Having access to massive amounts of information, including the entire Wikipedia online encyclopedia, makes Watson a whiz at trivia games. When pitted against real people on the TV quiz show "Jeopardy," the unusual contestant was able to beat all the top previous human contestants and won an unprecedented one million dollars in prize money. [See [Jeopardy! is Elementary for Watson](#), March 2011.]



Ben Hider/Getty Images

While many brain simulations are massive in scope and provide a detailed picture of neural networks, they don't show the relationship of physical signals to complex functions performed by the different parts of the brain. For example, IBM's "Watson" is constructed to answer questions posed in natural language. Having access to massive amounts of information makes Watson a whiz at trivia games and very little else. Meanwhile, "robot cars" won a DARPA-sponsored race across California's Mojave Desert. However, the same vehicle could not navigate through California traffic.



DARPA

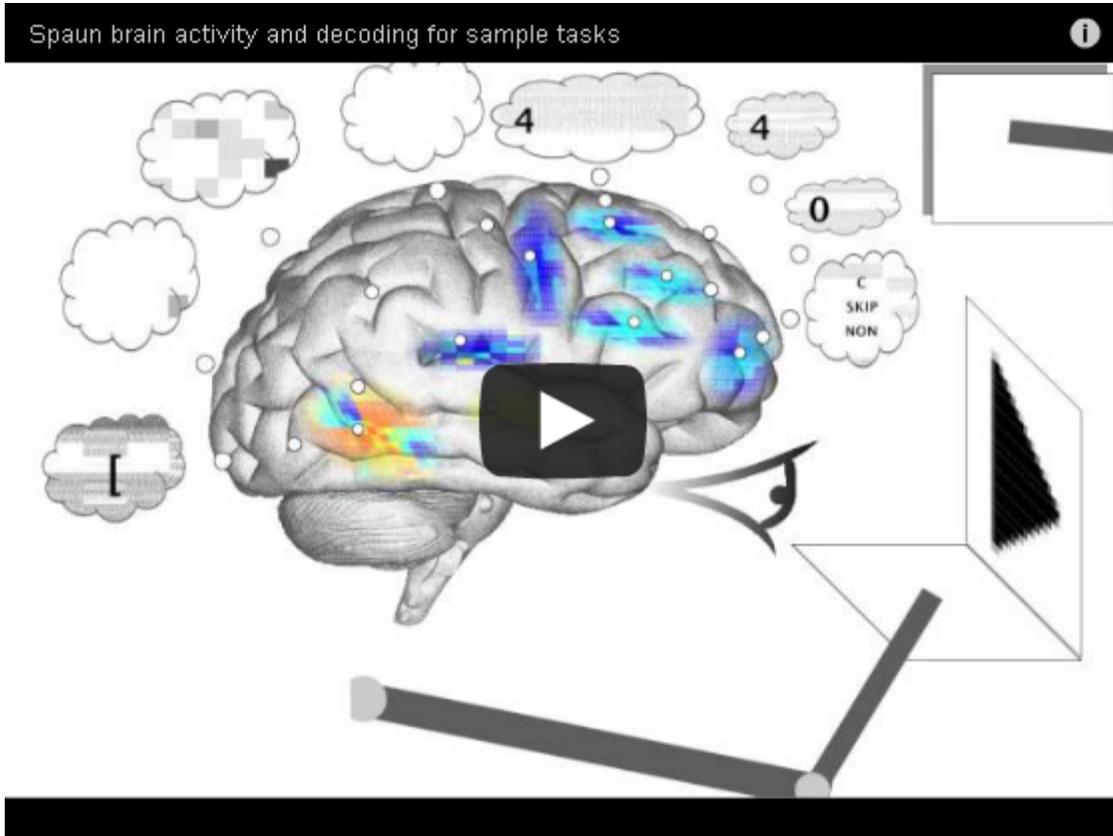
Part of the challenge lies in the brain's capacity for innovation: simulations often lack the brain's ability to come up with creative approaches to the task at hand. For example, driverless "robot cars" won a race across California's Mojave Desert sponsored by the Defense Advanced Research Projects Agency (DARPA). However, when trying to navigate through traffic, the same vehicles couldn't deal with obstacles in their path: they ran into one another when trying to share a lane and could be thrown off course by something as simple as a rock.

Spaun Specialties

In an effort to move past these challenges, the scientists from the Waterloo Centre for Theoretical Neuroscience, in Waterloo, Ontario, presented a simulation that links the brain's physical apparatus to its behavioral functions. Eliasmith claimed, "Until now, the race has been who could get a human-sized brain simulation running, regardless of what behaviors and functions such simulation exhibits." However, "from now on, the race is more [about] who can get the most biological functions and animal-like behaviors. So far, Spaun is the winner."

In late 2012, the researchers published a paper that explained the details of their new virtual brain. Built to be able to interact with its environment, Spaun has an "eye" that can see, a two-joint "arm" that can draw, and a relatively modest 2.5 million virtual neurons. The neurons are grouped into "modules" that represent the prefrontal cortex, basal ganglia, thalamus, and other structures of the brain—areas in charge of processing images, orchestrating movements, and storing information in short-term memory.

The key feature of this model has to do with the so-called "semantic pointer" units of the simulated network. Instead of representing specific bits of information, semantic pointers can be put together in different ways to represent complex systems of connections. In this way, a semantic pointer works as an "address" rather than a "value."



CTNWaterloo

Built to be able to interact with its environment, Spaun has an "eye" that can see, a two-joint "arm" that can draw, and a relatively modest 2.5 million virtual neurons. Spaun's virtual sensory neurons pick up visual information, which consists of letters, shapes and numbers. The visual data gets sent to the "brain" to be stored in its memory. The brain then processes the input and sends a new signal to virtual motor neurons, allowing Spaun to use its arm in order to produce a written response to the data.

Input

0	1	2	3	4	5	6	7	8	9
0	1	2	3	4	5	6	7	8	9

Output

0	1	2	3	4
5	6	7	8	9

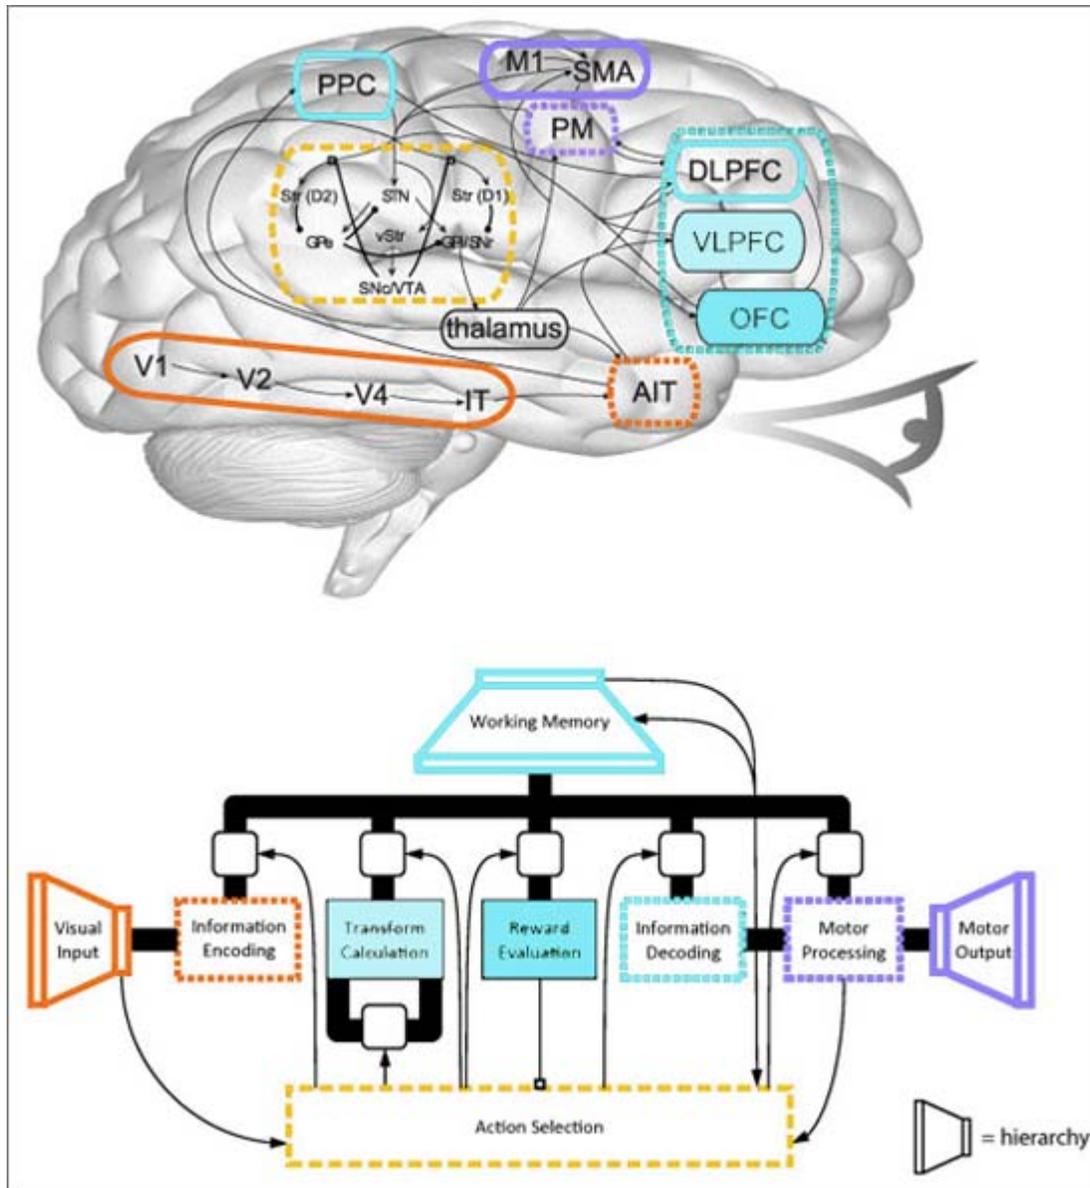
Putting this setup to work, Spaun's virtual sensory neurons pick up visual information, which consists of letters, shapes and numbers. The visual data gets sent to the "brain" to be stored in its memory. The brain then processes the input and sends a new signal to virtual motor neurons, allowing Spaun to use its arm in order to produce a written response to the data.

The signals that dash across the virtual nervous system consist of voltage spikes and simulated neurotransmitters. "It's all in a machine," Eliasmith explains. "But we're actually simulating all those voltages and currents down to the level of things you can measure in real cells." What results is a network of connections representing the intricate patterns in a real brain.

Spaun the Multitasker

There are eight tasks that Spaun is built to handle. The simplest perceptual problems require the virtual brain to recognize numbers in different styles of handwriting and copy them with the mechanical arm. The slightly more difficult working-memory questions involve sequence recall and reinforcement learning. Finally, the most challenging complex cognitive tasks call for Spaun to complete number sequences by figuring out the rule governing them.

In order for the signals involved in these operations to travel through the simulation effectively, input and output information has to be "packed up" in certain ways. First, the visual image gets "compressed," or whittled down to eliminate anything irrelevant or redundant. This happens as the information passes through a series of layers known as "Boltzmann machines," which correspond to the areas of the brain that make up the "ventral visual stream." The reverse happens in the response phase: the simple command generated by the virtual brain gets "expanded" into the many basic components of the complex arm movement.



Courtesy of Chris Eliasmith/University of Waterloo

In order for the appropriate signals to travel through the simulation effectively, input and output information has to be "packed up" in certain ways. First, the visual image gets "compressed," or whittled down to eliminate anything irrelevant or redundant. This happens as the information passes through a series of layers known as "Boltzmann machines," which correspond to the areas of the brain that make up the "ventral visual stream." The reverse happens in the response phase: the simple command generated by the virtual brain gets "expanded" into the many basic components of the complex arm movement.

The "cognitive machinery" involved in tasks Spaun performs depends on two main features: working memory and action selection. Working memory is based in the prefrontal cortex and allows Spaun to store information it picks up from the environment. It also links new data to information that is already stored, making it possible to "infer relations between past and present stimuli."

Action selection allows Spaun to switch from one task to another while using the same "machinery." This process is controlled by the basal ganglia—a brain structure that serves as a switchboard that shuttles information between different parts of the cortex and determines which step should come next. As Eliasmith suggests, "You can think of the basal ganglia as controlling the flow of information through [the] cortex, in order to solve different tasks."

Copying Quirks

This ability to switch gears is an essential aspect of human brain functioning, and is crucial to understanding how the brain's physical processes generate actions. According to Eliasmith, "Human cognition isn't interesting because we can recognize visual patterns...move our arms in an integrated way...[or] solve a particular task or puzzle. It is interesting because we can do all of this with the same brain, in any order, and at any time."

And while Spaun may be less powerful than some previous brain simulations, its ability to use the same "raw material" to create different types of connections to meet the demands of particular problems is an important step in bridging the "brain-behavior" gap. It was especially interesting for the scientists to see that the simulated version of this process developed quirks similar to those manifested by real brains. For example, Spaun hesitates before answering questions and has trouble remembering large sets of numbers. Just like the human brain, it's also better at recalling items at the beginning and end of a list, compared to items in the middle.

"That's the kind of thing we couldn't program in," Eliasmith points out. "We weren't surprised that it could do tasks...but we were often surprised that subtle features like the time it took or the errors it made were the same as for humans."

A Sampler Plate of Brain Functions

Of course, any virtual brain, no matter how sophisticated, can't outsmart the one each one of us already has—at least not yet. So, understandably, Spaun has several limitations. For example, it's nowhere near as fast as the real brain: each second of simulation requires two hours of computation. Moreover, the virtual brain can't control its own input by switching its attention or changing the position of the eye.

Finally, even though Spaun can successfully crank out answers to questions it's "ready" for, it will be thrown off guard by any unexpected deviations. Unlike the human brain, the simulation is not as "adaptive and cannot acquire new skills or learn new tasks. And since there are still many gaps in our understanding of learning scientists can't reproduce the learning process from scratch.

As a result, Spaun reproduces only a subset of brain functions and lacks some of the details of an actual brain. "People are enormously more complicated," Eliasmith says, "[Spaun] is nowhere near as intricate or sophisticated as human brains." Instead, you can think of Spaun as "a sampler plate of brain functions." As Eliasmith explains, "There are some fairly subtle details of human behavior that the model does capture. It's definitely not on the same scale as a human brain....It gives a flavor of a lot of different things brains can do." [For more details, see nengo.ca/build-a-brain/.]

The Next Step

The Waterloo scientists plan to address some of these limitations in the future. A more complex model, one that has more plasticity and ability to tailor responses in innovative ways, will soon be in the works. It should be able to learn new tasks and process instructions at a more complex level.

At some point, Eliasmith hopes, direct instructions might not be necessary at all: "We would just tell it if it is doing a good job or a bad job. Eventually it would discover its own strategy for accomplishing its own task." It would also be a lot faster, eventually working in real time. "Instead of us giving it strategies for performing these tasks, it would be able to discover strategies based on experience, just like humans do."

Eliasmith describes some of the possibilities of brain simulation in *How to Build a Brain*, a book set to come out in early 2013. At the moment, however, there are still many gaps in the whole picture. As neuroscientist Sean Hill points out, we have a long way to go when it comes to understanding the brain: "Right now, we're in a crisis in neuroscience. There's a lot of wonderful data being gathered but we don't have a place where we can put those experimental results together and understand their implications."

Robots of the Future

Still, while AI's future remains unclear, the ability to construct more sophisticated brain simulations is likely to open a number of doors. For example, scientists can use such simulations to build better robots—ones more closely modeled on our own mental machinery. Such robots, in turn, should be easier for us to interact with since they would "think" like we do. "We can try to discover the algorithms used by biology, and maybe understand the principle behind them, to build better artificial agents," Eliasmith suggests. "Those kinds of features are important in a way because if we're interacting with an agent and it has a kind of memory that we're familiar with, it'll [be] more natural to interact with."



Left: Victor Habbick Visions/Getty Images; Right: Hanson Robotics

Scientists can use more sophisticated brain simulations to build better robots—ones more closely modeled on our own mental machinery. Such robots, in turn, should be easier for us to interact with since they would "think" like we do.

The new simulations will also be more flexible, thanks to their ability to switch between tasks the way real brains do. As a result, "robots of the future" will be able to perform more complex jobs. For example, Eliasmith suggests that one day the postal system might turn to robots for certain tasks: "A robot that is able to navigate through a city and deliver a package from one place to another. I think that kind of thing will be within reach in the next 10 years."

Accurate brain simulations also have important applications in medical science. They provide a useful platform for research on neurological conditions and the process of aging. By intentionally altering, or even "damaging," parts of the model brain (part of the model or changing how neurotransmitters work), scientists may/should be able to determine how structural changes affect the brain's cognitive functions—and how particular treatments can help restore those functions to health. In fact, this project is already under way: the Waterloo research team submitted a paper in which they describe an "aging" Spaun losing neurons at the same rate as humans do. It turned out that the loss of virtual neurons led to a similar type of cognitive decline.

As neuroscientist Michael Harper explains, "Currently, Spaun only lives inside a simulated environment deep inside a computer, but Eliasmith and team hope this model will be a huge step for understanding how the brain reacts to different drugs and treatments as well as how the brain ages and learns. The average human brain is comprised of nearly 100 billion neurons, working together to perform the tasks of our everyday life."

Chris Eliasmith: Making a Model Brain

Chris Eliasmith is director of the Centre for Theoretical Neuroscience at the University of Waterloo in Ontario, Canada and holds the Canada Research Chair in Theoretical Neuroscience. He currently is appointed to the department of philosophy and the department of systems design engineering. In addition, he has a cross-appointment to the school of computer science.

Eliasmith earned his bachelor's degree from the University of Waterloo in systems design engineering in 1994 and completed his master's degree in philosophy a year later. In 2000 Eliasmith received his Ph.D. in philosophy in the philosophy-neuroscience-psychology program at Washington University in St. Louis, Missouri. He completed one-year as a postdoctoral researcher in computational neuroscience at Washington University Medical School before returning to the philosophy department at Waterloo in 2001. A licensed engineer since 2008, Eliasmith considers his research interdisciplinary.

Eliasmith is the author of two books on neural architecture, as well as numerous articles, reviews and lectures on theoretical neuroscience, theoretical psychology, philosophy and machine intelligence. He is the recipient of numerous awards and grants and is active in professional organizations and committees.

Eliasmith's research interests lie in the intersection of theoretical and computational neuroscience, and the philosophy of mind, language and cognitive science. He is focused on large-scale simulations of specific behaviors, computational cognitive neuroscience involving memory, emotions and decision-making, neural coding, computational modeling and neurobiological explanations.

Below are Eliasmith's January 29, 2013 responses to questions posed to him by Today's Science.



Courtesy of Chris Eliasmith/University of Waterloo

"To me, theoretical neuroscience is like theoretical physics, but for your brain instead of the universe. No one would deny the contributions of theoretical physics, but the same isn't yet true of theoretical neuroscience."

Q. When did you realize you wanted to become a scientist?

A. I consider myself more of an "academic" — someone who wants to have a career generating and evaluating knowledge — than a scientist per se (though lots of the knowledge I'm interested in is scientific). I'm currently appointed to the department of philosophy and the department of systems design engineering, and I have a cross-appointment to computer science. I decided I wanted to pursue this path after discovering that I had had the most fun engaging with new ideas — a realization I made upon completing my undergraduate degree.

Q. How did you choose your field?

I took many philosophy and psychology courses during my undergraduate studies in engineering. I was fascinated by the confluence of these fields. This made me pursue the question of how the brain works from several perspectives at once. I found my current place at the university by pursuing that question.

Q. Are there particular scientists, whether you know them in person or not, that you find inspiring?

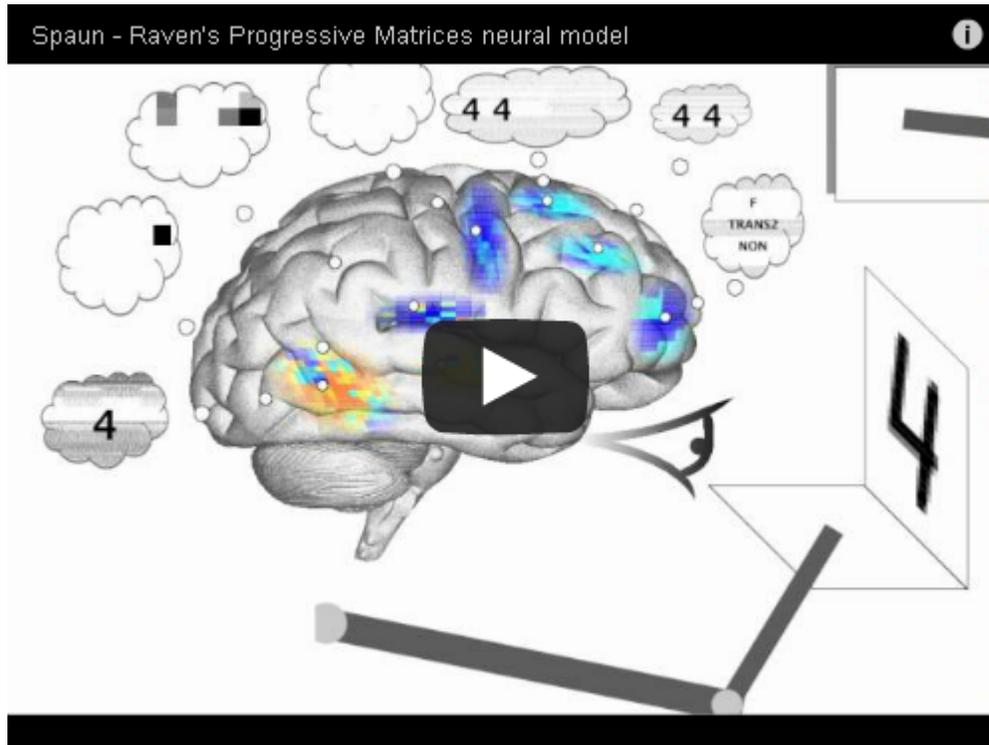
A. I'm most inspired by broad-minded generalists who don't shy away from either technical or conceptual problems. Several thinkers who are now called "philosophers" were like this (e.g., [René] [Descartes](#), [Aristotle](#), [Gottfried Wilhelm von] [Leibniz](#)), as were some who were typically called "scientists" (e.g., [Albert] [Einstein](#), [John] [von Neumann](#), [Norbert] [Wiener](#), [Alan] [Turing](#)).

Q. What do you think is the biggest misconception about your profession?

A. Maybe that it isn't yet mature enough to contribute to our understanding of how the brain works. To me, theoretical neuroscience is like theoretical physics, but for your brain instead of the universe. No one would deny the contributions of theoretical physics, but the same isn't yet true of theoretical neuroscience.

Q. You built your model brain using simulated neurons. Was the level of interconnectivity of your simulated neurons comparable to that of neurons in the brain? Did your simulation allow new connections to be made between neurons?

A. Our model has only about 20 of the 1,000 or so areas identified in the brain. However, all connections in our model exist in the real brain, and the degree of connectivity between neurons is similar. Yes, there are tasks (in particular, reinforcement learning) that the model brain performs by changing these connections between neurons. However, many of the tasks do not require such changes, even though learning occurs (e.g., in the rapid-variable-creation and Raven's tasks [nonverbal group tests typically used in educational settings]).



CTNWaterloo

"...there are tasks (in particular, reinforcement learning) that the model brain performs by changing these connections between neurons. However, many of the tasks do not require such changes, even though learning occurs (e.g., in the rapid-variable-creation and Raven's tasks [nonverbal group tests typically used in educational settings])."

Q. This is a very broad question, I know, but how would you characterize the biggest differences between your brain simulation and a conventionally designed computer intended to perform tasks similar to those your simulation could do?

A. There are two main differences. One is that the processing in our brain model is highly parallel, with all 2.5 million neurons constantly processing information. The second is that the representations used in our model are quite unlike the "symbols" used in traditional computers.

Without getting into details, the representations we use are analog and highly distributed, whereas a computer uses more digital and less distributed representations.

Q. What do you think are realistic goals for artificial intelligence in, say, the next 10 years? — That is, what kinds of tasks do you think machines/robots will be able to do in 2023 that they cannot do now? What, if any, are the goals that strike you as more long-term and difficult?

A. I think we will have fairly sophisticated robots that are quite physically and cognitively flexible. They will likely be able to do things like deliver a package anywhere in a city, or provide some assistance in the home. We will also likely see the continued improvement of "disembodied" AI, like Siri and Watson. This will become more integrated into our online experiences. However, I'm uncertain about how well we will solve the problem of building agents that can carry on an informative discussion, or be good teachers, etc. This requires

sophisticated social and emotional processing as well as cognitive, perceptual and motor skills.

Q. Do you think the bigger payoff of your research — and that of others working in the field — will be in enhancing our artificial intelligence capabilities, or in achieving a better understanding of how our own brains work?

A. I'm hoping that it will provide big payoffs in both areas. At the moment, I'm not sure I could pick one over the other. If anything, I think that the fact that we are linking these two aspects of brain function will be what allows us to get mutually reinforcing payoffs in both areas.

Q. Where do you spend most of your workday? Who are the people you work with?

A. I spend most of my time in my office, reading, writing, and conversing via email or in person. I work most closely with my students and postdocs, since we share many of the same immediate goals. However, I also talk a lot to others in the Centre here [the Waterloo Centre for Theoretical Neuroscience]. It is a great facility with lots of exciting things going on, so I like to keep up to date with all that's happening nearby.

Q. What do you find most rewarding about your job? What do you find most challenging about your job?

A. The most rewarding aspect is building and designing things that actually work. A close second is coming up with good ways to explain those artifacts — but this conceptual work is very dependent on success in building something interesting in the first place. I find it challenging to balance administration with research. I'm not a big fan of the former, and the trend seems to be towards more complex and time-wasting administrative rules, which can be frustrating.

Q. What has been the most exciting development in your field in the last 20 years? What do you think will be the most exciting development in your field in the next 20 years?

A. This is a hard question for me to answer because I participate in so many fields! In some ways, I think the convergence of different ideas is quite exciting.

Seeing work in probability theory and machine learning collide with our neuroscientific understanding of perception is quite exciting. The same type of thing has happened in reinforcement learning and motor control. I think this is exciting because over the next 20 years, I hope this will develop into a unified theory of much of brain functioning.

Q. How does the research in your field affect our daily lives?

A. Research in machine learning is constantly being introduced into our daily lives, through ever more sophisticated "services" on the web, and on our devices. Research in neuroscience is leading to a better understanding of the effects of drugs, new ways of helping people with brain or spinal damage (e.g., with robotic prosthetics), and improving our ability to treat diseases including Parkinson's, Alzheimer's, Huntington's and many more.

Q. For young people interested in pursuing a career in science, what are some helpful things to do in school? What are some helpful things to do outside of school?

A. The most helpful thing to do in school is to be a sponge! Learning everything you can about every field by fully participating in learning.

This doesn't mean getting high marks on everything, but rather working on something until you actually understand it. I'd also recommend leaning towards technical subjects — math, science, computer programming — as these are the main tools of science. Outside of school should really be no different — build things, design things, figure out how stuff works, even if it seems obvious — be a sponge.

Discussion Questions

How does the brain-simulation approach discussed in this story differ from the building of conventional computers, and how is it similar? Can you think of ways in which AI researchers could try to simulate the conditions under which the human brain evolved so as to "evolve" some form of AI?

Journal Abstracts and Articles

"A Large-Scale Model of the Functioning Brain." www.sciencemag.org/content/338/6111/1202.

Bibliography

The Blue Brain Project. bluebrain.epfl.ch/.

"Blue Brain' Project Accurately Predicts Connections Between Neurons." *ScienceDaily* (September 17, 2012) [accessed December 11, 2012]: www.sciencedaily.com/releases/2012/09/120917152043.htm.

Edelman, Gerald M. and Giulio Tononi. *A Universe of Consciousness: How Matter Becomes Imagination*. Basic Books, 2000.

Edwards, Lin. "Spaun, the New Human Brain Simulator, Can Carry Out Tasks." *Phys.org* (November 30, 2012) [accessed December 7, 2012]: phys.org/news/2012-11-spaun-human-brain-simulator-tasks.html.

Eliasmith, Chris, Terrence C. Stewart, Xuan Choo, Trevor Bekolay, Travis DeWolf, Charlie Tang and Daniel Rasmussen. "A Large-Scale Model of the Functioning Brain." *Science*, November 30, 2012, pages 1202-1205.

Fildes, Jonathan. "Artificial Brain '10 Years Away.'" *BBC News* (July 22, 2009) [accessed December 11, 2012]: news.bbc.co.uk/2/hi/8164060.stm.

Givetash, Linda. "Brain Simulator Could Provide Insight to Damage, Diseases." *The Record* (November 30, 2012) [accessed December 8, 2012]: www.therecord.com/news/local/article/846117--brain-simulator-could-provide-insight-to-damage-diseases.

Harper, Michael. "Spaun Brain Model Created Inside a Computer." *Redorbit* (November 30, 2012) [accessed December 7, 2012]: www.redorbit.com/news/technology/1112739456/human-brain-computer-model-spaun-eliasmith-113012/.

The Human Brain Project. www.humanbrainproject.eu/vision.html.

Lehrer, Jonah. "Can a Thinking, Remembering, Decision-Making, Biologically Accurate Brain Be Built From a Supercomputer?" *Seed* (January 11, 2012) [accessed January 15, 2012]: seedmagazine.com/content/article/out_of_the_blue/?utm_source=SB-bottom&utm_medium=linklist&utm_content=magazine&utm_campaign=internal%252Bblinkshare.

Machens, Christian K. "Building the Human Brain." *Science*, November 30, 2012, pages 1156-1157.

Malein, Flora. "Artificial Brain 'Spaun' Passes IQ Tests." *TG Daily* (December 3, 2012) [accessed December 11, 2012]: www.tgdaily.com/general-sciences-features/67819-artificial-brain-spaun-passes-iq-tests.

Markram, Henry. "A Countdown to a Digital Simulation of Every Last Neuron in the Human Brain." *Scientific American* (June 11, 2012) [accessed December 12, 2012]: www.scientificamerican.com/article.cfm?id=human-brain-project-digital-simulation-neuron.

Pinker, Steven. *How the Mind Works*. W.W. Norton, 1997.

Sanders, Laura. "Simulated Brain Mimics Human Quirks." *ScienceNews* (November 29, 2012) [accessed December 7, 2012]: www.canada.com/technology/Canadian+scientists+create+functioning+virtual+brain/7628440/story.html.

Seung, Sebastian. "Another Perspective on Massive Brain Simulations." *Scientific American* (June 11, 2012) [accessed December 12, 2012]: www.scientificamerican.com/article.cfm?id=massive-brain-simulators-seung-conntectome&page=4.

Waldrop, Mitchell M. "Computer Modelling: Brain in a Box." *Nature* (February 22, 2012) [accessed December 11, 2012]: www.nature.com/news/computer-modelling-brain-in-a-box-1.10066.

Yong, Ed. "Simulated Brain Scores Top Test Marks." *Nature News* (November 29, 2012) [accessed December 7, 2012]: www.nature.com/news/simulated-brain-scores-top-test-marks-1.11914.

Keywords

brain simulation, artificial intelligence, Blue Brain Project, Human Brain Project, "Spaun," computer modeling, Waterloo Centre for Theoretical Neuroscience, Henry Markham, Chris Eliasmith