ISTOCKPHOTO.COM/BLIXEN, ILLUSTRATION BY SPOOKY POOKA

HOW TO BUILD A BRAIN

It's the body's most complex organ – the tangle of cells that makes us who we are. But, says Moheb Costandi, scientists think they can now create an artificial human brain

Alan Mathison Turing was one of the leading scientific minds of his time. He cracked German codes in World War II and is widely regarded as the father of the modern computer. But he wasn't simply driven by the challenge of solving problems. It was the sudden death of his close friend Christopher Morcom, when Turing was just 17, that profoundly shifted his view of the world. It underpinned many of his greatest discoveries and led him towards his biggest challenge – to build a human brain.

The death of Morcom in 1930, an intellectual companion for Turing, led him to consider how the human brain worked. He came to the conclusion that it was simply a machine with a series of logic gates – on/off switches – that underpin all our thought processes. As such, he believed there was nothing that couldn't be replicated in a man-made machine. In other words, we could build an artificial brain.

Turing died before he achieved this goal. But now, 100 years after his birth, some scientists are daring to think they might be able to recreate the body's most complex organ. Their plan isn't to build a brain out of the 86 billion nerve cells, or neurones, that make up our brains. Instead, they plan to recreate a brain in digital form – using software, silicon and wires.

One of the most ambitious of these initiatives is the Human Brain Project, led by Professor Henry Markram of École Polytechnique Fédérale de Lausanne in Switzerland. His plan is to integrate everything that is known about the brain, from the molecular level up through to its large-scale anatomical structure (see 'Building blocks of the brain', on p34), in one working model that resides in a supercomputer. It's a huge challenge requiring big money. In March, Markram's proposal was selected as one of six finalists in a European Commission competition for €1 billion (£0.8 billion).

To build the artificial brain, models of all the processes that go on inside the real thing would need to be encoded in software and brought together so they can interact. The hope is that the resulting 'unified model' would provide insights into what makes us tick – exactly how our thoughts and behaviour come about.

"We know a lot about the brain and we have very good knowledge about its components," says the brain project's spokesman Richard Walker, "but our systematic knowledge of how they work together is weak.

"We can put information about the behaviour of individual cars into a computer and generate a simulation that will predict traffic jams. In the same way, simulating how the components of the brain behave ▸

▸ and interact could give us an overall picture of how they all work together. What the brain actually does in life will emerge from the interactions of these lower-level parts."

Mind-boggling complexity

It's not going to be an easy task. The brain's 86 billion neurones form something like one quadrillion synapses (that's one thousand billion), which are being modified continuously, countless times every second. Simulating this complexity would involve processing a huge data set, the size of which approaches that of all the stored information in the world. This would require at least an exaflop of computing power, executing roughly a trillion operations per second. Even today's most powerful supercomputers aren't up to this gargantuan task.

But Markram has a good track record in the field. In 2008, he and his colleagues used IBM's Blue Gene supercomputer to simulate a tube-shaped column of neurones in a rat brain, a column containing 10,000 nerve cells with 30 million connections, or synapses, between them. Now, they have simulated 100 such columns.

"Simulating the brain would involve processing a data set the size of all the stored information in the world"

Markram's team will generate the human brain simulation using some of their own data. Earlier this year, for example, they published details about the combination of ion channels expressed in several types of neurones in the rat brain. Data from outside sources, such as the Human Connectome Project in which the major pathways of neurones in the brain are being plotted, will be built into the simulation too.

As well as providing new insights into the workings of the brain, Walker believes the resulting simulation could lead to new treatments for diseases. "If we know that schizophrenia is due to some molecular defect, we can simulate that in the model and then do experiments that are otherwise impossible to do in the lab," he says.

And the benefits would not be limited to medicine. "The human brain is amazing. Every day, it absorbs the same amount of energy as a light bulb. It learns by itself without being programmed and can run for more than 70 years without serious defects," says Walker. "If we could imitate the brain to build machines with some of these capabilities, we could revolutionise information and communications technology."

If Markram and his colleagues hear they have been awarded the money later this year, they will bring together scientists from more than 100 international research centres working in neuroscience, robotics, genetics, and mathematics. If they don't get the money, they say they will continue with the work regardless, albeit at a slower pace.

Many researchers have criticised the project, arguing that it is ill-conceived and over-ambitious. One of the most outspoken critics is neuroscientist Professor Rodney Douglas of the Institute of Neuroinformatics in Zürich. "A description of the brain is not an explanation of how it works," he says. "It doesn't follow that a better understanding of the brain will necessarily emerge from encoding its detailed organisation in a computer. The claim that a theory of brain function would emerge from a simulation is very extravagant."

A different approach is being taken by Dr Dharmendra Modha who works for IBM in the US and runs the global ▸

BUILDING BLOCKS OF THE BRAIN

The key components that would need to be incorporated into an artificial brain

Cortical columns

These are vertical structures in the cerebral cortex, the outer layer of the brain. It plays a key role in memory, thought and language. The columns contain groups of neurones specialised to perform a particular function. There are about half a million of them, each containing about 60,000 neurones.

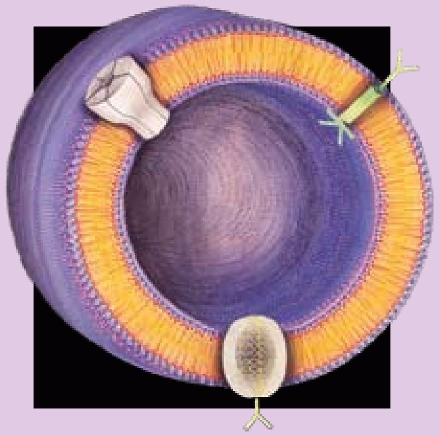

Ion channels

These pores span the nerve cell membrane and control the electrical activity of neurones, turning them on or off. As such, their behaviour is crucial to how the brain behaves. They exert their control by allowing electrically-charged atoms (or ions) of calcium, sodium and potassium to enter or leave the cell.

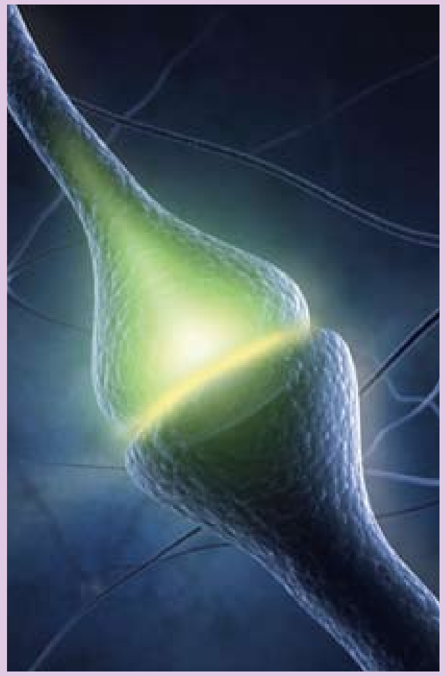

Synapses

The junctions between neurones where communication between different nerve cells takes place, either electrically or chemically. Each neurone can have many synapses, with the result that the human brain contains something like one quadrillion junctions. Learning and memory are widely believed to involve the strengthening of synapses in distributed networks of neurones.

White matter tracts

These are large bundles of nerve fibres, or axons, that form long-range connections between different parts of the brain. The fibres are insulated by a fatty tissue, myelin, which increases the velocity at which nervous impulses travel and gives the tissue its white appearance. Recent research shows that the brain's white matter tracts are arranged in a highly ordered three-dimensional grid.

Neurones and glia

These are the two main cell types within the brain. It is the interconnected networks of neurones that encode and process information and about 86 billion of them would be needed to be included in an artificial human brain. Glial cells have been thought of as nothing more than support cells, but we now know that certain types make important contributions to information processing.

ADVANCES IN NEUROSCIENCE

Key steps towards building an artificial brain

1852

German physician Hermann von Helmholtz measures the speed of nerve impulses. The conventional wisdom at the time was that impulses are propagated instantaneously. Helmholtz used an induction coil to send electrical surges through the nerves of frogs, showing that they conduct electricity at a speed slow enough to be measured.

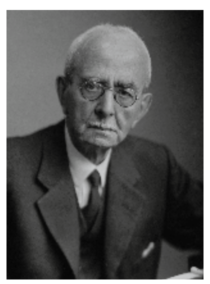

1897

English physiologist Charles Sherrington coins the term 'synapse' to describe the junctions between neurones following experiments on spinal reflexes in which he noticed a short delay in the transmission of impulses. It was not until the 1950s, when powerful electron microscopes were invented, that synapses were observed directly.

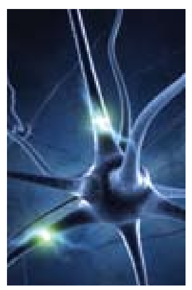

1906

Using a new dye, Spanish neuroscientist Santiago Ramón y Cajal determines that neurones are the brain's functional units. In the 19th century an ongoing debate had some scientists arguing that the brain consisted of distinct cells, others that it was a network of tissues. Cajal settled the argument.

1921

Pharmacologist Otto Loewi discovers the first known neurotransmitter. He took two frog hearts, one with its vagus nerve attached. When this was slowed with an electrical current, the second slowed too. Loewi surmised that a chemical, we now know to be acetylcholine, was released by the first heart and affected the other.

1944

Alan Turing tells a colleague Donald Bayley of his desire to build a brain. The impetus is the death of his friend Christopher Morcom in 1930. It is an event that leads him to consider the mind and its relationship to the brain. Turing is convinced a machine can be created that can mimic anything the brain can do.

BLUE BRAIN PROJECT, SCIENCE PHOTO LIBRARY X3, THINKSTOCK X3, UNIVERSITY OF MANCHESTER, CORBIS X2

CAN WE REALLY BUILD AN ARTIFICIAL HUMAN BRAIN?

The prospect of reconstructing the most complex organ divides opinion

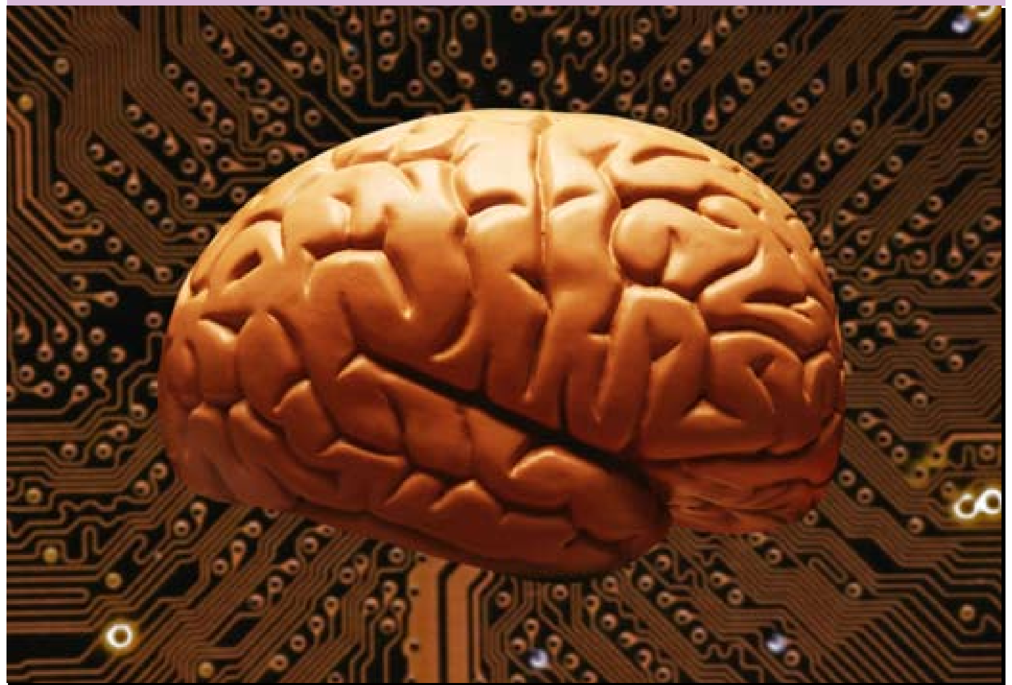

YES

Professor Henry Markram
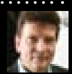
Leader of the Human Brain Project

"Supercomputers are becoming powerful enough to simulate the brain with biological detail and accuracy. They can help us construct a multi-level view of the brain showing how all its parts work together. The volume of data on the healthy and diseased brain is increasing exponentially: big science initiatives are screening the brain at various levels and theoretical neuroscience, mathematics and statistics have reached a point where they can begin to analyse it all. But understanding the brain is not a one-man show. It will require many disciplines coming together, with hundreds of Einsteins sitting around a table to solve the problems systematically."

NO

Professor Rodney Douglas
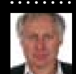
Neuroscientist at the Institute of Neuroinformatics, Zürich

"Claims that we can learn something new from a simulation miss an important principle. That is that simulation involves understanding a functional system from observed data, then creating an abstract model before converting it into a code that is optimised for computer hardware. Your computer then runs this simulation, but not the physical processes themselves. You'll only retrieve the consequences of what you have input from the simulation; not the behaviour of the original system. In other words, it never rains inside a weather simulation because the physics of computation has nothing to do with the physics of weather."

SyNAPSE project. Instead of building a brain using software, his team plans to build its components using hardware – reproducing the organ's structures in silicon. And where Markram and his team's primary aim is to understand how the brain works, Modha is more interested in mimicking the brain to build truly intelligent machines that can program themselves and learn from experience.

Hard-wired brain

Modha's approach, known as neuromorphic computing, was pioneered by Professor Carver Mead of the California Institute of Technology. Working with PhD student Misha Mahowald, he developed the first neuromorphic device in the late 1980s – a silicon retina in which electrical circuits mimic the function of the eye's rod and cone cells.

Neuromorphic computing promises to provide supercomputing performance at an affordable price by overcoming the constraints of conventional computers. Developed in the 1940s by early computer scientist John Von Neumann and others, conventional computer architecture separates the central processing unit from the memory, so that information has to be shuttled back and forth between the two through another component called a bus. Packets of information are processed one at a time and sent to the memory individually. This creates a processing bottleneck that uses lots of energy, slowing the computer.

Crucially, neuromorphic chips integrate the computational units and the memory components on the same silicon circuit and process multiple pieces of information in parallel, more closely resembling the brain – increasing processing speed and reducing the energy required. "Von Neumann architectures are very good at calculations and repetitive tasks, but the brain is fundamentally different," says Ton Engbersen, a senior researcher at IBM's Zurich lab who is involved with the SyNAPSE project. "It's very good at interpreting what we see and all other kinds of what we call 'unstructured data'."

Last year, Modha's global team of scientists unveiled the most sophisticated neuromorphic chip yet: a 'neurosynaptic core' consisting of 256 neurones and more than 260,000 synapses replicated in silicon. The 'weight', or strength, of each synapse can be pre-programmed into the simulation before it is run, enabling the chip to perform navigation and pattern-recognition tasks.

"I believe there's an algorithm for how the brain learns, although we don't know that yet," says Engbersen. "We are trying to build the next generation of computers that can process information and learn in the same way as the brain." In the long run, the IBM team is planning to develop a human-scale system containing one quadrillion synapses distributed across multiple neurosynaptic core chips.

But while most neuroscientists agree that the mind is an emergent property of the brain, it's not clear whether something like human intelligence would emerge from one of these artificial brains.

> *"I believe there's an algorithm for how the brain learns, although we don't know that yet"*


Dr Dharmendra Modha is trying to build a brain in silicon with the SyNAPSE project

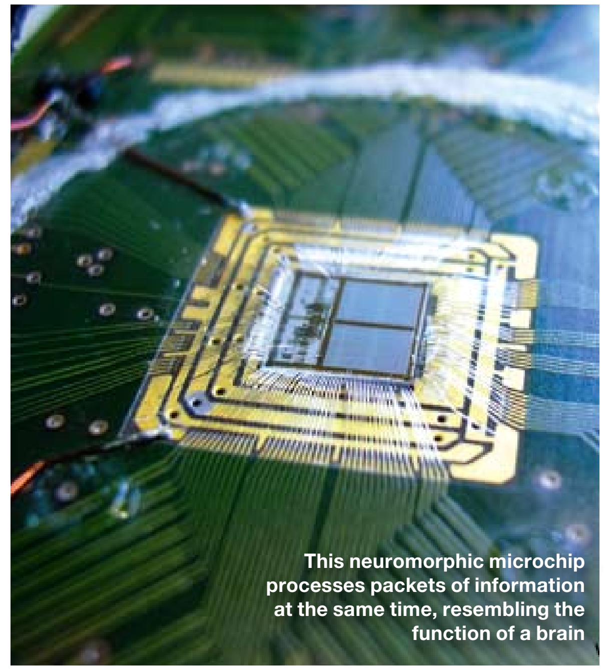
This neuromorphic microchip processes packets of information at the same time, resembling the function of a brain

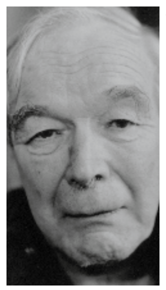

1952 British physiologists Alan Hodgkin and Andrew Huxley set out to discover what happens when a neurone is stimulated. By impaling giant nerve fibres, or axons, from squid with fine wires called microelectrodes, they determine that impulses are generated by the flow of ions, or charged atoms, across the nerve cell membrane.

1959 Neurobiologists David Hubel and Torsten Wiesel publish a research paper describing the visual system of cats. They found that individual neurones in the primary visual cortex are sensitive to movement in specific directions. The pair go on to show that the visual cortex is organised in alternating columns devoted to the left and right eye.

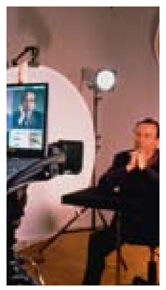

1988 US computer scientist Carver Mead develops the concept of neuromorphic computing by describing very-large-scale integrated (VLSI) systems mimicking the architecture of the brain. He subsequently works with PhD student Misha Mahowald to develop the first neuromorphic microchip, a silicon replica of three cell layers in the retina.

1990 Japanese researcher Seiji Ogawa and colleagues develop functional magnetic resonance imaging (fMRI), a technique that measures activity in the brain by detecting blood flow. fMRI is now widely used and provides information about how connected networks in the brain participate.

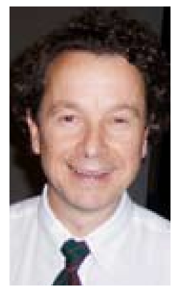

1998 Rod MacKinnon, a neurobiologist in the US, describes the three-dimensional structure of an ion channel in a neurone for the first time. MacKinnon and his colleagues isolate the potassium channel from bacteria and discover what it looks like using a technique called X-ray crystallography.

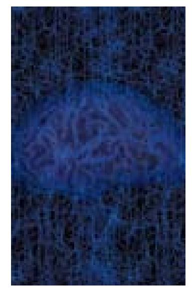

2012 Neuroscientist Henry Markram proposes the Human Brain Project, which aims to collate all existing knowledge about the brain, including data from the Allen Brain Institute's gene expression atlas. The aim is to use the data to generate a supercomputer simulation of the human brain.

GETTY X2, SYNAPSE PROJECT, BLUE BRAIN PROJECT X2, HANNES ROST, SCIENCE PHOTO LIBRARY X2, THINKSTOCK

Henry Markram gives a TED talk on his work with the human brain

▶ Dr Kwabena Boahen, principal investigator at Stanford University's Brains in Silicon research group, believes it would. "How else could intelligence emerge? It has to come out of stuff we can build. The physics of how neurones work is very similar to the physics of transistors. As we learn more about the brain, we can replicate its functions." Turing would have undoubtedly agreed.

Boahen and his colleagues have built a neuromorphic supercomputer that rivals the performance of IBM's Blue Gene while consuming 100,000 times less energy. They are using it to run a real-time simulation of a million neurones interconnected with something in the order of a billion synapses. The aim is to understand how the brain decides where to focus attention and how we choose what to do – how cognition happens through the actions of neurones.

"The brain is not made out of some alien technology," says Boahen, "and there's nothing going on in there that we can't replicate in silicon. We know a lot about how individual neurones and synapses work. It's a matter of figuring out how complex phenomena emerge when lots of them are hooked up to each other."

Melanie Mitchell, Professor of Computer Science at Portland State University in the US, agrees that the brain could eventually be replicated, but argues that our knowledge of how it functions is woefully inadequate. "In principle, we can replicate anything the brain does," she says. "But in practice it's going to be very difficult. We don't yet know enough about the brain to say what its core algorithms are."

"There's nothing going on in the brain that we can't replicate in silicon"

Markram's team used the Blue Gene supercomputer to simulate part of a rat's brain

Fathoming the brain

New research is revealing just how little we understand about our most complex organ. After British physiologists Alan Hodgkin and Andrew Huxley discovered that nervous impulses are generated by the flow of ions, or charged atoms, in 1952, we came to view neurones as computational units that function as binary switches, existing in either an 'on' or 'off' state. In recent years, however, neuroscientists have come to realise that neurones are more complex. Rather than being binary switches, they exist in many different states. And the nervous impulses that travel along them can move in both directions, not just one.

No-one – not even Markram and Modha – believe that building an artificial brain is going to be easy. To a large degree, the human brain is a black box filled with many smaller black boxes: after a century of investigation, the workings of some of these boxes remains a mystery. But while the prospects of completely achieving Turing's ultimate aim still seem slim, just like Turing, the lessons we can learn as we strive to achieve this goal could be revolutionary. ■

Moheb Costandi is a molecular and developmental neurobiologist turned science writer

FIND OUT MORE

http://bit.ly/henrymarkram
TED talk by Henry Markram on simulating a brain

http://bit.ly/kwabenaboahen
TED talk by Kwabena Boahen on building a brain

TED/JAMES DUNCAN DAVIDSON, BLUE BRAIN PROJECT