

Shutterstock/jagsandrew

Smarter machines

Although today's computers can perform superhuman feats, even the best are no match for human brains at tasks like processing speech. But as **Jessamyn Fairfield** explains, a new generation of computational devices is being developed to mimic the networks of neurons inside our heads

The most impressive electronics you possess aren't in your pocket. They're in your head. Not only does your brain keep you breathing, moving, speaking and thinking enough to read a magazine about physics, but it does it all on the cheap. The brain runs on about 20 watts – less than a light bulb and far less than the megawatts used by supercomputing centres. And yet on this power the brain is able to do things that supercomputers still struggle with, such as advanced pattern recognition.


For decades, scientists have tried to develop “brain-like” code for conventional computers, aiming to emulate the unique talents of the brain. (Imagine what could be achieved by combining the computational power of supercomputers with our own cognitive abilities.) Researchers have, however, found this a very challenging quest, since the architecture of conventional computing hardware is so different from that of the brain. Instead, a community of scientists and engineers is now developing a new kind of computing hardware architecture that is physically,

and functionally, more analogous to the computers inside our heads.

Our brain has a plethora of electrically active cells, called neurons, each of which has a profusion of connections, called synapses, to other neurons. With time and experience, the synapses in the brain strengthen and weaken to allow the flow of more or fewer current pulses, which changes the strength, or “synaptic weight”, of connections. Additionally, neurons can be discarded or grown anew, and these structural changes lead to learning and memory. This ability of the brain to adapt and rewire itself, known as its “plasticity”, is why it can learn new tasks without needing an external programmer. Sensibly, the brain has evolved computational strengths that are adapted not so much to abstract computation, but to real-world problems, like not getting eaten by a tiger.

Supercomputers, in contrast, run on switches and logic gates, and can far outperform the brain on computational power alone. This powerful mathematical approach is implemented using semiconduc-

Jessamyn Fairfield is a physicist specializing in nanomaterials and neuromorphic devices at NUI Galway, Ireland, e-mail jessamyn.fairfield@nuigalway.ie



Nanowire networks

Films of nanomaterials can be used to create adaptable electronics that mimic the brain, and the connectivity paths through these materials can be imaged using scanning electron microscopy. When the metal electrode (bottom) is connected to electrical ground, nanowires with a conducting path to the electrode and the electrode itself appear dark due to electrons flowing out. In contrast, nanowires isolated from the electrode experience a build-up of electrons, causing them to appear light as more electrons are reflected into the detector. The scale bar represents 10 μm .

tors, which have not only enabled huge advances in computing, but also made it widely affordable. While the brain can rewire its own hardware, this feature is neither replicated nor missed in the main silicon computer-chip industry. The 1980s did see the invention of a hardware component that could be reprogrammed “in the field”, but the “field-programmable gate array” did not prove popular except in niche applications and lacked the full flexibility of neurons and synapses.

Although the raw computational power of supercomputers cannot be matched, Moore’s law, in which the number of transistors on a chip doubles every two years, is slowly eroding. We are perhaps close to the end of gains to be eked from traditional silicon architectures, as we approach the limit of heat dissipation in ever more miniaturized circuits. Another restriction of this technology is that the circuits for computation and memory are kept separate, requiring large throughput to move information from computation to memory and back again. This “von Neumann” bottleneck is avoided in the brain because memory is largely stored in the shape of the network, and in the way in which neurons are connected to which others, rather than as a series of bits that have to be shuffled around.

Many impressive forays into simulating the human brain have nevertheless managed to do so using conventional computer hardware. “Deep learning”, for example, is effectively a software approach incorporating the brain’s patterns of connectivity to tackle applications like computer vision, speech recognition and language processing. Deep-learning algorithms often comprise multiple layers of artificial neural networks – sets of software nodes each of which feeds into the next – that can either perform “supervised” learning, involving human guidance, or “unsupervised” learning. Such computational approaches are key to the Human Brain Project, a €1bn European Union initiative to simulate the brain. Deep learning seems to be a promising step towards achieving this goal.

However, trying to implement software for neural behaviour onto traditionally structured electronics is like playing a symphony on only timpanis. How

much more could we do if we had electronics that mimicked the timbre and shape of the brain?

Neuromorphic engineering

While the quest to understand how we think, and hence how the brain works, is nearly as old as science, only recently have we developed electronics anywhere near capable of reproducing its circuitry. Early computational devices were based on the binary logic of the transistor, first implemented with vacuum tubes and later in silicon, where bits of information can have values of only 1 or 0. Binary logic is based on algebra first proposed by George Boole in the mid-1800s, an English mathematician working in Cork, Ireland. But in the 1960s Lotfi Zadeh at the University of California, Berkeley, proposed an analogue algebraic system, where a range of less crisply defined values are the basis for logic operations. This “fuzzy logic” is much closer to the changeable synaptic weights of the brain.

One of the first pioneers who saw the relevance of fuzzy logic to biological electronic systems was engineer Carver Mead at the California Institute of Technology, US, who in the 1980s coined the word “neuromorphic” to describe electronic devices modelled on human biology. He went on to create computer interfaces for touch, hearing and vision. While Mead and other early researchers in neuromorphic engineering found niche applications for their devices, ultimately they were limited because most widely available electronics were still based on digital logic using silicon, and the primary industrial goal was to make these silicon transistors as small and reliable as possible.

Decades later, a breakthrough result enabled a renaissance in neuromorphic engineering. In 2008 neuromorphic engineers learned of a new circuit element for their arsenal, when a team led by Stanley Williams at Hewlett Packard reported the first experimental realization of a memristor (*Nature* **453** 80). This device had first been predicted by Leon Chua in 1971 as a theoretical necessity to complete the set of basic circuit elements: resistor, capacitor and inductor. The memristor in its ideal form is a circuit element whose resistance depends on how much current has previously flowed through it, meaning it can be used as an analogue form of computer memory that is much closer to the plasticity of synapses in the brain.

Hewlett Packard’s implementation of the memristor consisted of a nanoscale film of titanium dioxide, containing positively charged oxygen vacancies, sandwiched between two electrodes. When a current is passed through the film, the oxygen vacancies move one way or the other across the film, depending on the sign of the current. This repositioning of the oxygen vacancies changes the resistance across the film, leading to the device memory.

While the ideal memristor is a passive device, requiring no power to operate, the experimental realizations of memristors have generally involved some sort of energy storage and release. Since this is required in atomic reordering or chemical reactions at the nanoscale, many researchers say that pure and



Mimicking the brain IBM TrueNorth computer chips are based on simulated neurons and synapses.

hence passive memristors are only theoretically possible, though they do say that “memristive systems” like the titanium dioxide film can be implemented as useful circuit elements.

As many researchers realized once the first results were announced, memory of past measurements is a common nanoscale feature. What had previously been seen as a flaw in the electronic behaviour of many nanomaterials could now become a strength. Current flow can often cause small changes in materials, by moving atoms or even just by changing the electron distribution in the material between mobile states and trapped states. These changes may not be noticeable in the properties of a bulk material, but when the material is thin or has a small feature size, the changes can affect material properties in a measurable way. This robust memristance in nanomaterials has thus led to many neuromorphic devices based on nanoscale phenomena and interactions.

Materials for brain-like electronics

The development of memristive devices initially continued in the vein of the thin-film-based Hewlett Packard device, but a variety of other designs soon followed. A phenomenon called resistive switching drives one increasingly popular type of memristor, most notably thanks to work led by Rainer Waser at RWTH Aachen University, Germany, and Wei Lu at the University of Michigan, US. In this device, a conductive filament across a non-conducting oxide can be formed and then strengthened or broken by electrical impulses, allowing patterns of conductivity to be defined and used for either memory or computation. It’s challenging to control the exact filament location and structure, as with many processes at the nanoscale, but these resistive switches are self-healing and responsive to electrical stimuli. The filaments themselves parallel synaptic connections between neurons, whose synaptic weight can be tuned by repeated stimuli. And since filamentary conduction has been found in a wide range of metal oxides at the

The memristor can be used as an analogue form of computer memory that is much closer to the plasticity of synapses in the brain

nanoscale, the exact structure of the device can be tuned for the environment and application.

Another approach to memristors is to combine nanomaterials with polymers, either by coating the nanomaterials with a polymer or by mixing the two. As shown by researchers led by Dominique Vuillaume at the University of Lille, France, polymers in conjunction with conducting nanoparticle or nanowire cores can change their resistance, which is a key feature of memristive devices. Some conducting polymers can also be doped and de-doped, enabling them to act as standalone neuromorphic components when submerged in an electrolyte. Such polymer-based memristors are of special interest for brain-computer interfaces because polymers are both mechanically and electrically closer to the brain, and have more parallel channels for conduction than inorganic electronics. Recent research led by George Malliaras at the École des Mines de Saint-Étienne, France, used derivatives of poly(3,4-ethylenedioxythiophene) (PEDOT) in electrolyte-controlled transistors. These devices have shown both plasticity and another feature of synapses called timing dependence, in which the memristive function depends on how closely spaced together in time the electrical pulses arrive.

These building blocks of neurons and synapses, both the polymer and solid-state implementations, can then be combined into networks that exhibit memory, as shown by researchers led by John Boland at Trinity College Dublin, Ireland. These networks can not only reproduce tunable synaptic weights, but they can also show the same timing dependence as actual synapses do. My own published work with both the Boland and Malliaras labs investigated neuromorphic behaviour in individual nanowires, nanowire networks and polymer electronics, showing that brain-like function can be created in a very broad array of materials and device types. Piece by piece, researchers are reproducing the key components of synaptic plasticity in the brain.

Even the US government has taken an interest in neuromorphic computing, funding the SynAPSE initiative through the US Defense Advanced Research Projects Agency to the tune of more than \$100m. IBM and HRL Laboratories in the US have been the main beneficiaries, with their researchers working to make artificial neurons and synapses in silicon that are compatible with industrial fabrication processes. In 2014 IBM released its TrueNorth chip, which simulates a million neurons and 256 million programmable synapses, while consuming only one ten-thousandth of the energy of a comparable traditional microchip. Computationally, the



Thousands of patents cite


IOP Publishing content*.

Find out how our database can spark innovation and creativity in your R&D department.

Strengthen your research today.

Ask us about a free trial at
iopcorporate@iop.org

* Based on a search using the Google Patents database



Laboratory of Neuro Imaging and Martinos Center for Biomedical Imaging, Consortium of the Human Connectome Project – www.humanconnectomeproject.org/ David Shattuck and Paul M Thompson

Graphic navigation The Human Connectome Project aims to provide an unparalleled compilation of neural data, such as this high-angular resolution diffusion image of the axial view of a brain. The tensor can be visualized as an ellipsoid in 3D space, showing fluid mappings and brain connectivity.

value of memristive networks is already being shown for conventional tasks such as prime factorization, and novel neuromorphic computing paradigms are also emerging.

The path forward

While this is all exciting research, many important unanswered questions still lie beneath the surface. Which brain features are actually necessary for computing, memory or – dare we ask – consciousness? Which parts of the connectivity of the brain are critical, and which are remnants of our evolution that we no longer need? Does the brain run on electrical pulses merely out of biological necessity, or for computational reasons we have yet to realize and exploit? And even if we can reproduce the behaviour of synaptic spikes in artificial neuromorphic devices, will we gain anything from it without a way to crack the neural code that determines how these spikes are translated into movement, feeling or memory?

To find a path forward, neuromorphic engineers must talk to neuroscientists who have long approached these questions from the other side, examining fully formed and functioning brains and attempting to tease out the principles underneath. Those looking to build a brain from the ground up would be foolish not to seek out this knowledge. The amazing work coming out of the Human Connectome Project in the US, for example, has shed new light on how different regions of the brain are interconnected, light which could illuminate a new path for neuromorphic device connectivity. Only then can the scientific community get closer to the tantalizing dual prizes, of computers that can truly grow and learn, and a better understanding of the computers at work inside our very own heads.