Brain-inspired computing needs a master plan.

A. Mehonic & A.J. Kenyon*
Department of Electronic & Electrical Engineering
UCL
Torrington Place
London
WC1E 7JE
United Kingdom

* Corresponding author. E-mail address: a.kenyon@ucl.ac.uk

Preface

New computing technologies inspired by the brain promise fundamentally different ways to process information with extreme energy efficiency and the ability to handle the avalanche of unstructured and noisy data that we are generating at an ever-increasing rate. To realise this promise requires a brave and coordinated plan to bring together disparate research communities and to provide them with the funding, focus and support needed. We have done this in the past with digital technologies; we are in the process of doing it with quantum technologies; can we now do it for brain-inspired computing?

Main

Modern computing systems consume far too much energy. They are not sustainable platforms for the complex Artificial Intelligence (AI) applications that are increasingly a part of our lives. We usually don't see this, particularly in the case of cloud-based systems, as we focus on functionality – how *fast* are they; how *accurate*; how many parallel operations per second? We are so accustomed to accessing information near-instantaneously that we neglect the energy, and therefore environmental, consequences of the computing systems giving us this access. Nevertheless, each Google search has a cost: data centres currently use around 200 terawatt hours of energy per year, forecast to grow by around an order of magnitude by 2030[1]. Similarly, the astonishing achievements of high-end AI systems such as DeepMind's AlphaGo and AlphaZero, which can beat human experts at complex strategy games, require thousands of parallel processing units, each of which can consume around 200 Watts[2].

While not all data-intensive computing requires AI or Deep Learning (DL), DL is deployed so widely that we must worry about its environmental cost. We should also consider applications including the Internet of Things (IoT) and autonomous robotic agents that may not need always to be operated by computationally intense DL algorithms but must still reduce their energy consumption. The vision of the IoT cannot be achieved if the energy requirements of the myriad connected devices are too high. Recent analysis shows that increasing demand for computing power vastly outpaces improvements made through Moore's law scaling[3]. Computing power demands now double every two months (Figure 1a). Remarkable improvements have been made through a combination of smart architecture and software-hardware co-design. For example, the performance of NVIDIA GPUs has improved by the factor of 317 since 2012: far beyond what would be expected from Moore's law alone (Figure 1b) – although the power consumption of units has increased from ~25 W to around 320 W in the same period. Further impressive performance improvements have been demonstrated at the R&D stage (Figure 1b in red) and it is likely that we can achieve more [4,5]. Unfortunately, it is unlikely that conventional computing solutions alone will cope with demand over an extended period. This is especially apparent when we consider the shockingly high cost of training required for the most complex DL models (Figure 1c). We need alternative approaches.

The energy problem is largely a consequence of digital computing systems storing data separately from where they are processed. This is the classical von Neumann architecture underpinning digital computing systems. Processors spend most of their time and energy moving data. Luckily, we can improve the situation by taking inspiration from biology, which takes a different approach entirely – co-locating memory and processing, encoding information in a wholly different way or operating directly on signals, and employing massive parallelism, for example (Box 1). There is a system that achieves both energy efficiency and advanced functionality remarkably well: the brain. Recognising that we still have much to learn about how the brain operates and that our aim is not simply to emulate biological systems, we can nevertheless learn from the significant progress in neuroscience and computational neuroscience in the last few decades. We know just enough about the brain to use it as an inspiration.

Biological inspiration

Biology does not separate data storage from processing. The same elements – principally neurons and synapses – perform both functions in massively parallel and adaptable structures. The 10^{11} neurons and 10^{15} synapses contained in the typical human brain expend approximately 20 W of power, while a digital simulation of an artificial neural network of approximately the same size consumes 7.9 MW[6]. That six order of magnitude gap poses us a challenge. The brain directly processes with extreme efficiency signals that are noisy. This contrasts with the signal to data conversion and high precision computing in our conventional computer system that produces huge costs in energy and time for even the most powerful digital supercomputers. Brain-inspired, or *neuromorphic*, computing systems could therefore transform the way we process signals and data, both in terms of energy efficiency and of their capacity to handle real-world uncertainty.

This is not a new idea. The term *neuromorphic*, describing devices and systems that mimic some functions of biological neural systems, was coined in the late 1980s by Carver Mead at the California Institute of Technology[7,8]. The inspiration came from work undertaken over previous decades to model the nervous system as equivalent electrical circuits[9] and to build *analogue* electronic devices and systems to provide similar functionality (Box 1).

A word about “data”. We use the term to describe information encoded in, say, an analogue signal or the physical response of a sensor, as well as the more standard computing-focused sense of digital data. When we refer to the brain “processing data” we describe an integrated set of signal processing tasks that do not rely on digitisation of signals in any conventional sense. We can think of brain-inspired systems operating at different levels: from analogue signal processing to working with large digital data sets. In the former case, we can avoid generating large data sets in the first place; in the latter we can greatly increase the efficiency of processing by moving away from the von Neumann model. Of course, there are good reasons why we represent data digitally for many applications: we want high precision, reliability and determinacy. However, digital abstraction discards massive amounts of information, found in the physics of transistors, for the minimum information quantum: a single bit. And we pay a significant energy cost by trading efficiency for reliability. As AI applications are often probabilistic at heart we must consider if this trade-off makes sense. The computational tasks underpinning AI applications are very compute-intensive (and therefore energy-hungry) when performed by conventional von Neumann computers. However, we might perform similar tasks much more energy-efficiently on analogue or mixed systems that use a spike-based representation of information. There has therefore been a recent resurgence in interest in neuromorphic computing, driven by the growth in AI systems and by the emergence of new devices that offer new and exciting ways to mimic some of the capabilities of biological neural systems (Box 1).

Definitions of *neuromorphic* vary considerably. Loosely speaking, the story is a hardware one: neuromorphic chips aim to integrate and utilise various useful features of the brain, including in-memory computing, spike-based information processing, fine-grained parallelism, signal processing resilient to noise and stochasticity, adaptability, learning in hardware, asynchronous communication, and analogue processing. While it is debatable how many of these need to be implemented for something to be classified as neuromorphic, this is clearly a different approach from AI implemented on mainstream computing systems. Nevertheless, we should not be lost in terminology; the main question is whether this approach is *useful*.

Approaches to neuromorphic technologies lie on a spectrum between reverse-engineering the structure and function of the brain (analysis) and living with our current lack of knowledge of the brain but taking inspiration from what we do know (synthesis). Perhaps foremost among the former approaches is the Human Brain Project, a high-profile and hugely ambitious ten-year programme funded by the European Union from 2013. The programme supported the adoption and further development of two existing neuromorphic hardware platforms – SpiNNaker (at Manchester) and BrainScaleS (at Heidelberg) – as openly accessible neuromorphic platforms. Both systems implement highly complex silicon models of brain architectures to understand better the operation of the biological brain. At the other end of the spectrum numerous groups augment the performance of digital or analogue electronics using selected biologically-inspired methods. Figure 2 summarises the range of existing neuromorphic chips, divided into four categories depending on their position on the analysis-synthesis spectrum and their technology platform. It is important to remember that neuromorphic engineering isn't just about high-level cognitive systems, but also offering energy, speed and security gains (at least by removing the need for constant communication to the Cloud) in small-scale edge devices with limited cognitive abilities.

Prospects

We do not propose that neuromorphic systems will, or should, replace conventional computing platforms. Instead, precision calculations should remain the preserve of digital computation while neuromorphic systems can process unstructured data, perform image recognition, classification of noisy and uncertain data sets, and underpin novel learning and inference systems. In autonomous and IoT-connected systems, they can provide huge energy savings over their conventional counterparts. Quantum computing is also part of this vision. A practical quantum computer, while still several years away by any estimation, would certainly revolutionise many computing tasks. However, it is unlikely that IoT-connected smart sensors, edge computing devices, or autonomous robotic systems will adopt quantum computing without depending on cloud computing. There will remain a need for low-power computing elements capable of dealing with uncertain and noisy data. We can imagine a three-way synergy between digital, neuromorphic and quantum systems.

Just as the development of semiconductor microelectronics relied on many different disciplines, including solid state physics, electronic engineering, computer science, and materials science, neuromorphic computing is profoundly cross- and inter-disciplinary. Physicists, chemists, engineers, computer scientists, biologists, neuroscientists, all play key roles. Simply getting researchers from such a diverse set of disciplines to speak a common language is challenging. In our own work we spend considerable time and effort ensuring that everyone in the room understands terminology and concepts in the same way. A case for bridging the communities of computer science (specifically AI) and neuroscience (initially computational neuroscience) is clear. After all, many concepts found in today's state-of-the-art AI systems arose in the 1970s and 80s in neuroscience though, of course, AI systems need not be completely bio-realistic. We must include other disciplines, recognising that many of the strides we have made in AI or neuroscience have been enabled by different communities – e.g. innovations in material science, nanotechnology, or electronic engineering. Further, conventional CMOS technology may not be the best fabric to efficiently implement new brain-inspired algorithms; innovations across the board are needed. Engaging these communities early reduces the risk of wasting effort on directions that have already been explored and failed, or of reinventing the wheel.

Further, we should not neglect the challenges of integrating new neuromorphic technologies at the system level. Beyond the development of brain-inspired devices and algorithms there are pressing questions around how existing, mainstream, AI systems can be replaced with functionally equivalent neuromorphic alternatives. This further emphasises the need for a fully integrated approach to brain-inspired computation.

We should point out that, despite the potential outlined above, there is as yet no compelling demonstration of a commercial neuromorphic technology. Existing systems and platforms are primarily research tools. However, this is equally true of quantum computing, which remains a longer-term prospect. We should not let this delay the development of brain-inspired computing; the need for lower power computing systems is pressing and we are tantalisingly close to achieving this with all the added functionality that comes from a radically different approach to computation. Commercial systems will surely emerge.

Seizing the opportunity

If neuromorphic computing is needed, how to achieve it? First, the technical requirements. Bringing together diverse research communities is necessary but not sufficient. Incentives, opportunities and infrastructure are needed. The neuromorphic community is a disparate one lacking the focus of quantum computing, or the clear roadmap of the semiconductor industry. Initiatives around the globe are starting to gather the required expertise, and early stage momentum is building. How can we build on this? Funding is key. Investment in neuromorphic research is nowhere near the scale of that in digital AI or quantum technologies (Box 2). While that is not surprising given the maturity of digital semiconductor technology, it is a missed opportunity. There are a few examples of medium-scale investment in neuromorphic R&D such as the IBM AI Hardware Centre's range of brain-inspired projects (including the TrueNorth chip), Intel's development of the Loihi processor, and the US Brain Initiative project, but the sums committed are well below what they should be given the promise of the technology to disrupt digital AI.

The neuromorphic community is a large and growing one, but one that lacks a focus. While there are numerous conferences, symposia and journals emerging in this space there remains much work to be done to bring the disparate communities together and to corral their efforts to persuade funding bodies and governments of the importance of this field.

The time is ripe for bold initiatives. At a national level, governments need to work with academic researchers and industry to create mission-oriented research centres to accelerate the development of neuromorphic technologies. This has worked well in areas such as quantum technologies and nanotechnology (the US National Nanotechnology Initiative demonstrates this very well[10]) and provides focus and stimulus. Such centres may be physical or virtual but must bring together the best researchers across diverse fields. Their approach must be different from that of conventional electronic technologies in which every level of abstraction (materials, devices, circuits, systems, algorithms and applications) belongs to a different community. We need holistic and concurrent design across the whole stack. It isn't enough for circuit designers to consult computational neuroscientists before designing systems; engineers and neuroscientists must work together throughout the process to ensure as full an integration of bio-inspired principles into hardware as possible. Interdisciplinary co-creation must be at the heart of our approach. Research centres must house a broad constituency of researchers.

Alongside the required physical and financial infrastructure, we need a trained workforce. Electronic engineers are rarely exposed to ideas from neuroscience, and vice-versa. Circuit designers and physicists may have a passing knowledge of neurons and synapses but are unlikely to be familiar with cutting edge computational neuroscience. There is a strong case to set up Masters courses and doctoral training programmes to develop neuromorphic engineers. UK research councils sponsor Centres for Doctoral Training (CDTs) – focused programmes supporting areas with an identified need for trained researchers. CDTs can be single- or multi-institution; there are significant benefits to institutions collaborating on these programmes by creating complementary teams across institutional boundaries. Programmes generally work closely with industry and build cohorts of highly skilled researchers in ways that more traditional doctoral programmes often do not. There is a good case to be made to develop something similar, to stimulate interaction between nascent neuromorphic engineering communities and provide the next generation of researchers and research leaders. Pioneering examples include the Groningen Cognitive Systems and Materials research programme, which aims to train tens of doctoral students specifically in materials for cognitive (AI) systems[11]; the Masters programme in Neuroengineering at the Technical University of Munich[12], ETH Zurich courses on analogue circuit design for neuromorphic engineering[13]; large-scale neural modelling at Stanford University[14], and development of visual neuromorphic systems at the Instituto de Microelectrónica de Sevilla[15]. There is scope to do much more.

Similar approaches could work at the trans-national level. As always in research, collaboration is most successful when it is the best working with the best, irrespective of borders. In such an interdisciplinary endeavour as neuromorphic computing this is critical, so international research networks and projects undoubtedly have a role to play. Early examples include the European Neurotech consortium[16], focusing on neuromorphic computing technologies, as well as the Chua Memristor Centre at the University of Dresden[17], which brings together many of the leading memristor researchers across materials, devices and algorithms. Again, much more can and must be done.

How to make this attractive to governments? Government commitment to more energy-efficient bio-inspired computing can be part of a broader large-scale decarbonisation push. This will not only address climate change but also will accelerate the emergence of new, low-carbon, industries around big data, IoT, healthcare analytics, modelling for drug and vaccine discovery, and robotics, amongst others. If existing industries rely on ever more large-scale conventional digital data analysis, they increase their energy cost while offering sub-optimal performance. We can instead create a virtuous circle in which we greatly reduce the carbon footprint of the knowledge technologies that will drive the next generation of disruptive industries and, in doing so, seed a host of new neuromorphic industries.

If this sounds a tall order, consider quantum technologies. In the UK the government has so far committed around £1 billion to a range of quantum initiatives, largely under the umbrella of the National Quantum Technologies Programme. A series of research hubs, bringing together industry and academia, translate quantum science into technologies targeted at sensors and metrology, imaging, communications, and computing. A separate National Quantum

Computing Centre builds on the work of the hubs and other researchers to deliver demonstrator hardware and software to develop a general purpose quantum computer. China has established a multi-billion dollar Chinese National Laboratory for Quantum Information Sciences, while the USA in 2018 commissioned a National Strategic Overview for Quantum Information Science[18], which resulted in a 5-year $1.2 billion investment, on top of supporting a range of national quantum research centres[19]. Thanks to this research work there has been a global rush to start up quantum technology companies. One analysis found that in 2017 and 2018 funding for private companies reached $450 million[20]. No such joined-up support exists for neuromorphic computing, despite the technology being more established than quantum, and despite its potential to disrupt existing AI technologies on a much shorter time horizon. Of the three strands of future computing in our vision, neuromorphic is woefully under-invested.

Finally, some words about what bearing the COVID-19 pandemic might have on our arguments. There is a growing consensus that the crisis has accelerated many developments already under way: for example, the move to more homeworking. While reducing commuting and travel has direct benefits – some estimates put the reduction in global CO_2 as a result of the crisis at up to 17%[21] – new ways of working have a cost. To what extent will carbon savings from reduced travel be offset by increased data centre emissions? If anything, the COVID pandemic further emphasises the need to develop low carbon computing technologies such as neuromorphic systems.

Our message about how to realise the potential of neuromorphic systems is clear: Provide targeted support for collaborative research through the establishment of research centres of excellence; provide agile funding mechanisms to enable rapid progress; provide mechanisms for close collaboration with industry to bring in commercial funding and generate new spin-outs and start-ups, similar to schemes already in place for quantum tech; develop training programmes for the next generation of neuromorphic researchers and entrepreneurs; and do all of this quickly and at scale.

Neuromorphic computing has the potential to transform our approach to AI. Thanks to the conjunction of new technologies and a massive, growing demand for efficient AI we have a timely opportunity. Bold thinking is needed, and bold initiatives to support this thinking. Will we seize the opportunity?

Acknowledgements

AJK thanks the Engineering and Physical Sciences for financial support from grants EP/K01739X/1 and EP/P013503/1. AM gratefully acknowledges financial support from the Royal Academy of Engineering in the form of a Research Fellowship (RF201617\16\9).

Author contributions

Both authors contributed equally to the manuscript and revisions.

Competing interests

The authors declare the following competing interests:
The authors are founders and directors of Intrinsic Semiconductor Technologies Ltd (www.intrinsicst.com), a spin-out company commercialising silicon oxide RRAM.

Figures

Figure 1.

Figure 1. Computational demands are increasing rapidly. (a) *The increase in computing power demands over the past four decades expressed in PetaFLOPS-days. Until 2012, computing power demand doubled every 24 months; recently this has shortened to approximately every two months. The colour legend indicates different application domains. Data taken from [3].* ***(b)*** *Improvements in AI hardware efficiency over the last five years. State-*

of-the-art solutions have driven increases in computing efficiency of over 300 times. Solutions in research and development promise further improvements. ***(c)*** *Increase since 2011 of the costs of training AI models. Such an exponential increase is clearly unsustainable. Data taken from [22].*

Figure 2.

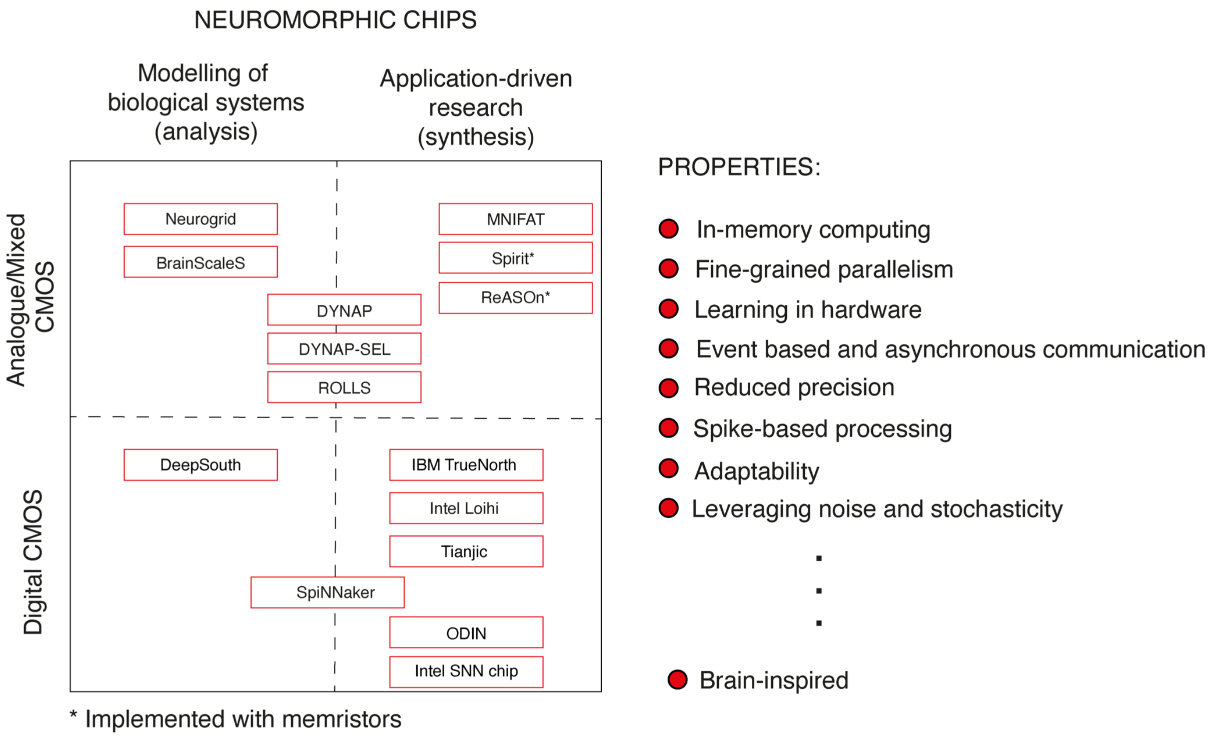

Figure 2. The landscape of neuromorphic systems. *Neuromorphic chips can be classified as either modelling biological systems or applying brain-inspired principles to novel computing applications. They may be further subdivided into those based on digital CMOS with novel architecture (for example, spikes may be simulated in the digital domain rather than implemented as analogue voltages) and those implemented using some degree of analogue circuitry. In all cases, however, they share at least some of the properties listed on the right-hand side, which distinguish them from conventional CMOS chips. Here we classify examples of recently developed neuromorphic chips. Further details of each can be found in the relevant reference: Neurogrid*[23]*, BrainSclaseS*[24]*, MNIFAT*[25]*, DYNAP*[26]*, DYNAP-SEL*[27]*, ROLLS*[28]*, Spirit*[29]*, ReASOn*[30]*, DeepSouth*[31]*, SpiNNaker*[32]*, IBM TrueNorth*[33]*, Intel Loihi*[34]*, Tianjic*[35]*, ODIN*[36]*, and the Intel SNN chip*[37]*.*

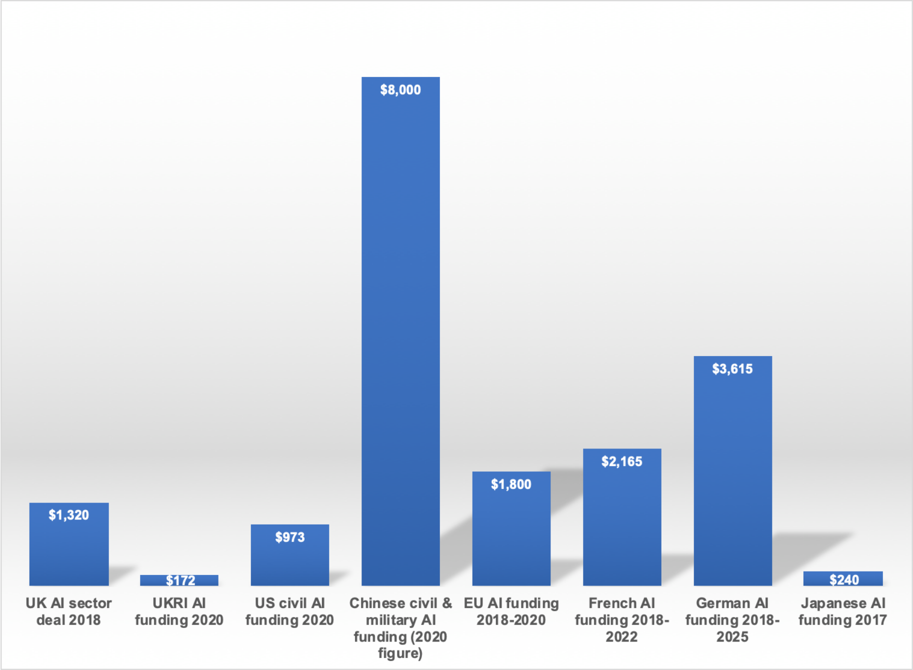

***Box 2 figure**: **Public investment in AI.** A comparison of recent global public research funding of digital AI technologies. Figures are in US dollar equivalent (2021 exchange rate) and are expressed as millions of dollars. While some are in-year snapshots (eg UKRI funding committed for 2020), some have no specified period (the UK AI sector deal, for example), and others are for multi-year programmes, the figure illustrates the scale of public funding in digital technologies. Disruption of the AI ecosystem by the development of efficient neuromorphic technologies would put much of this investment at risk.*

BOX 1 – What do we mean by "neuromorphic" systems?

Taking inspiration from the brain allows us to approach information processing fundamentally differently to the way existing conventional computing systems work. Different brain-inspired ("neuromorphic") platforms use combinations of different approaches: analogue data processing, asynchronous communication, massively parallel information processing or spiking-based information representation. These properties distinguish them from von Neumann computers.

The term *neuromorphic* encompasses at least three broad communities of researchers, distinguished by whether their aim is to emulate neural function (i.e. reverse-engineer the brain), simulate neural networks (i.e. develop new computational approaches), or engineer new classes of electronic device.

Neuromorphic engineering looks at how the brain uses the physics of biological synapses and neurons to "compute". Neuromorphic engineers work to emulate the functions of biological neurons and synapses by harnessing the physics of analogue electronics – such as carrier tunnelling, charge retention on silicon floating gates, and the exponential dependence of various device or material properties on field – to define elementary operations to underpin audio or video processing or smart sensors, for example. Transistors are used as analogue circuit elements with rich dynamic behaviour rather than binary switches. More details can be found in [38] and related work.

Neuromorphic computing looks to biology to inspire new ways to process data. This could be considered as the computing science of neuromorphic systems. Research looks to simulate the structure and/or operation of biological neural networks., which may mean co-locating storage and computing, as the brain does; or perhaps adopting wholly different ways of computing based on voltage spikes modelling the action potentials of biological systems.

Underpinning everything are the devices and materials needed to implement bio-inspired functions. Here, recent developments promise new electronic and photonic devices whose properties we can tailor to mimic biological elements such as synapses and neurons. These *neuromorphic devices* could provide exciting new technologies to expand the capabilities of neuromorphic engineering and computing.

Foremost amongst these new devices are *memristors*: electronic devices whose resistance is a function of their history. Their complex dynamic electrical response means they can be used as digital memory elements, as variable weights in artificial synapses, as cognitive processing elements, optical sensors, and devices that mimic biological neurons[39]. They may embody some of the functionality of biological dendrites[40] and their dynamic response can generate oscillatory behaviour similar to that of the brain – controversially, operating on the edge of chaos[41,42]. They may also be linked with biological neurons in a single system[43]. They do all of this while expending very little energy.

BOX 2 – The AI funding landscape

Investment in “conventional” digital AI is booming, fuelled by the need to process ever-increasing volumes of data, and the development of hardware to support existing compute- and memory-intensive algorithms. The UK government announced in April 2018 a £950 million “sector deal” in digital AI, in addition to existing research council support. France announced a €1.8 billion government investment in AI from 2018 to 2022[44], Germany committed €3 billion from 2018 to 2025, while Japan spent ¥26 trillion in 2017. US government funding of civil AI technologies was $973 million in 2020[45]; figures are harder to come by for US military AI funding, as non-AI projects are often included in published analysis. China is estimated to be investing up to $8 billion in both civil and military AI and is constructing a $2.1 billion AI research park near Beijing[46], while the European Commission committed €1.5 billion in the period 2018-2020[47]. Commercial investment dwarfs this. In the USA one estimate puts the total investment in AI companies in 2019 at $19.5 billion[48], and global investment is predicted to be around $98 billion by 2023[49]. Such sums must be considered at risk if our current hardware systems cannot support potentially disruptive neuromorphic algorithms and architectures. If neuromorphic technologies offer anything like the efficiency savings and enhanced performance they promise, smart money will hedge its bets on novel technologies and architectures alongside digital systems.

Comparable figures are not available for neuromorphic technologies, as they currently lack focus and government-level visibility. Research funding is therefore piecemeal and at project, rather than strategic, level. While there have been various estimates published – for example, that the global neuromorphic chip market will grow from $111 million in 2019 to $366 million in 2025[50], the safest conclusion to draw is that funding of neuromorphic systems lags way behind that of digital AI or of quantum (of which more below).

References

---

[1] Jones, N. How to stop data centres from gobbling up the world’s electricity. *Nature* **561**, 163 (2018).

[2] Wu, K.J. Google’s New AI Is a Master of Games, but How Does It Compare to the Human Mind? https://www.smithsonianmag.com/innovation/google-ai-deepminds-alphazero-games-chess-and-go-180970981/ accessed 17th August 2021

[3] Amodei, D. Hernandez, D. AI and Compute, https://openai.com/blog/ai-and-compute/ (accessed: September 2020).

[4] Venkatesan, R. *et al.* A 0.11 PJ/OP, 0.32-128 Tops, Scalable Multi-Chip-Module-Based Deep Neural Network Accelerator Designed with A High-Productivity vlsi Methodology. in *2019 IEEE Hot Chips 31 Symposium (HCS)* 1–24 (IEEE, 2019). doi:10.1109/HOTCHIPS.2019.8875657.

[5] Venkatesan, R. *et al.* MAGNet: A Modular Accelerator Generator for Neural Networks. in *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)* 1–8 (IEEE, 2019). doi:10.1109/ICCAD45719.2019.8942127.

[6] Wong, T.M. Preissl, R. Datta, P. Flickner, M. Singh, R. Esser *et al.* 10^{14}, *IBM Research Report RJ10502* (ALM1211-004) (2012). **The power consumption of this simulation of the brain puts ththat of conventional digital systems into context.**

[7] Mead, C.A. *Analog VLSI and Neural Systems*, Addison-Wesley, Reading MA (1989).

[8] Mead, C.A. Author Correction: How we created neuromorphic engineering. *Nat Electron* **3**, 579–579 (2020).

[9] Hodgkin A.L. and Huxley, A.F. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.,* **117**, 500 (1952). **This seminal work developed equivalent electrical circuits and circuit models for the neural membrane. More complex models followed, but this remains the clearest and an excellent starting point.**

[10] https://www.nano.gov Accessed 18th August 2021

[11] https://www.rug.nl/research/fse/cognitive-systems-and-materials/about/ Accessed 9th November 2020.

---

[12] https://www.tum.de/en/studies/degree-programs/detail/detail/StudyCourse/neuroengineering-master-of-science-msc/ Accessed 18th August 2021

[13] http://www.vvz.ethz.ch/Vorlesungsverzeichnis/lerneinheit.view?lerneinheitId=132789&semkez=2019W&ansicht=KATALOGDATEN&lang=en Accessed 9th November 2020.

[14] https://web.stanford.edu/group/brainsinsilicon/courses.html Accessed 9th November 2020

[15] http://www2.imse-cnm.csic.es/neuromorphs/ Accessed 9th November 2020

[16] https://neurotechai.eu Accessed 18th August 2021

[17] https://cmc-dresden.org/members/ Accessed 9th November 2020.

[18] https://www.whitehouse.gov/wp-content/uploads/2018/09/National-Strategic-Overview-for-Quantum-Information-Science.pdf Accessed 9th November 2020.

[19] P. Smith-Goodson, *Forbes*, (2019). https://www.forbes.com/sites/moorinsights/2019/10/10/quantum-usa-vs-quantum-china-the-worlds-most-important-technology-race/#371aad5172de Accessed 9th November 2020.

[20] Gibney, E. Quantum gold rush: the private funding pouring into quantum start-ups. *Nature* **574**, 22 (2019).

[21] Le Quéré, C., Jackson, R.B., Jones, M.W. et al. Temporary reduction in daily global CO_2 emissions during the COVID-19 forced confinement. *Nat. Clim. Change,* **10**, 647 (2020).

[22] https://research.ark-invest.com/hubfs/1_Download_Files_ARK-Invest/White_Papers/ARK–Invest_BigIdeas_2021.pdf. Accessed 27th April 2021

[23] Benjamin, B. V. *et al.* Neurogrid: A Mixed-Analog-Digital Multichip System for Large-Scale Neural Simulations. *Proc. IEEE* **102**, 699–716 (2014).

[24] Schmitt, S. *et al.* Neuromorphic hardware in the loop: Training a deep spiking network on the BrainScaleS wafer-scale system. *2017 International Joint Conference on Neural Networks (IJCNN)* (2017). doi:10.1109/ijcnn.2017.7966125

[25] Lichtsteiner, P., Posch, C. & Delbruck, T. A 128 x 128 120 dB 15 μs Latency Asynchronous Temporal Contrast Vision Sensor. *IEEE J. Solid-State Circuits* **43**, 566–576 (2008).

[26] Moradi, S., Qiao, N., Stefanini, F. & Indiveri, G. A Scalable Multicore Architecture With Heterogeneous Memory Structures for Dynamic Neuromorphic Asynchronous Processors (DYNAPs). *IEEE Trans. Biomed. Circuits Syst.* **12**, 106–122 (2018).

[27] Thakur, C. S. *et al.* Large-Scale Neuromorphic Spiking Array Processors: A Quest to Mimic the Brain. *Frontiers in Neuroscience* **12,** (2018).

[28] Qiao, N. *et al.* A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128K synapses. *Front. Neurosci.* **9**, (2015).

[29] Valentian, A. *et al.* Fully Integrated Spiking Neural Network with Analog Neurons and RRAM Synapses. in *2019 IEEE International Electron Devices Meeting (IEDM)* 14.3.1-14.3.4 (IEEE, 2019). doi:10.1109/IEDM19573.2019.8993431.

[30] Resistive Array of Synapses with ONline learning (ReASOn) developed by NeuRAM3 project, https://cordis.europa.eu/project/id/687299 (2021)

[31] Wang, R. *et al.* Neuromorphic Hardware Architecture Using the Neural Engineering Framework for Pattern Recognition. *IEEE Trans. Biomed. Circuits Syst.* **11**, 574–584 (2017).

[32] Furber, S. B., Galluppi, F., Temple, S. & Plana, L. A. The SpiNNaker Project. *Proc. IEEE* **102**, 652–665 (2014). **An example of a large-scale neuromorphic system as a model for the brain.**

[33] Merolla, P. A. *et al.* A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* **345**, 668–673 (2014).

[34] Davies, M. *et al.* Loihi: A Neuromorphic Manycore Processor with On-Chip Learning. *IEEE Micro* **38,** 82–99 (2018).

[35] Pei, J., Deng, L., Song, S. *et al.* Towards artificial general intelligence with hybrid Tianjic chip architecture. *Nature* **572,** 106–111 (2019).

[36] Frenkel, C., Lefebvre, M., Legat, J.-D. & Bol, D. A 0.086-mm2 12.7-pJ/SOP 64k-synapse 256-neuron online-learning digital spiking neuromorphic processor in 28-nm CMOS. *IEEE Transactions on Biomedical Circuits and Systems* 1–1 (2018).

[37] Chen, G. K., Kumar, R., Sumbul, H. E., Knag, P. C. & Krishnamurthy, R. K. A 4096-Neuron 1M-Synapse 3.8-pJ/SOP Spiking Neural Network With On-Chip STDP Learning and Sparse Weights in 10-nm FinFET CMOS. *IEEE Journal of Solid-State Circuits* **54,** 992–1002 (2019).

[38] Indiveri, G. *et al.* Neuromorphic Silicon Neuron Circuits. *Front. Neurosci.* **5**, (2011).

[39] Mehonic, A. *et al.* Memristors—From In-Memory Computing, Deep Learning Acceleration, and Spiking Neural Networks to the Future of Neuromorphic and Bio-Inspired Computing. *Advanced Intelligent Systems* 2000085 (2020) doi:10.1002/aisy.202000085. **A review of the promise of memristors across a range of applications, including spike-based neuromorphic systems.**

---

[40] Li, X. *et al.* Power-efficient neural network with artificial dendrites. *Nat. Nanotechnol.* **15**, 776–782 (2020).

[41] Chua, L. Memristor, Hodgkin-Huxley, and Edge of Chaos. *Nanotechnology*, **24**, 383001 (2013).

[42] Kumar, S., Strachan, J. P. & Williams, R. S. Chaotic dynamics in nanoscale NbO$_2$ Mott memristors for analogue computing. *Nature* **548**, 318–321 (2017).

[43] Serb, A. *et al.* Memristive synapses connect brain and silicon spiking neurons. *Sci Rep* **10**, 2590 (2020).

[44] Rosemain, M. Rose, M. *Reuters*, https://www.reuters.com/article/us-france-tech-idUSKBN1H51XP (2018) Accessed 9th November 2020.

[45] Castellanos, S. *Wall St Journal*, https://www.wsj.com/articles/executives-say-1-billion-for-ai-research-isnt-enough-11568153863 (2019).

[46] Larson, C. China’s AI Imperative. *Science*, **359**, 628 (2018).

[47] https://ec.europa.eu/digital-single-market/en/artificial-intelligence Accessed 9th November 2020.

[48] https://www.statista.com/statistics/672712/ai-funding-united-states/ Accessed 9th November 2020.

[49] https://www.idc.com/getdoc.jsp?containerId=IDC_P33198 Accessed 9th November 2020.

[50] https://www.marketwatch.com/press-release/neuromorphic-chip-market-size-demand-outlook-trends-revenue-future-growth-opportunities-by-2020-2025-2021-02-09. Accessed 27th April 2021