Neuromorphic Hardware as a Replacement for Von Neumann Architecture

Samuel Lord, Undergraduate, School of Computing and Mathematics, University of Plymouth
samuel.lord@students.plymouth.ac.uk

Abstract

The present study explores the application and usability of neuromorphic hardware to determine its viability as a replacement for von Neumann architecture as traditional hardware becomes limited in its improvements. A broad history of the development of neuromorphic hardware is explored and informs the discussion of the suitability of neuromorphic hardware to replace von Neumann systems. It is found that the hardware is currently cost-prohibitive for consumers and difficulty in processing general computing tasks make it unlikely that neuromorphic hardware will fully replace von Neumann systems,. The possibility of neuromorphic hardware as a supplementary device at both consumer and commercial levels is more likely.

Introduction

The prevalence of big data in modern business practices has facilitated the widespread use of brain-inspired, 'neuromorphic', algorithms for deep learning; learning customer behaviour and predicting trends to inform business strategies among other use cases.

As the adoption of these technologies becomes ever widespread, a demand for fast, cost efficient and end-user-friendly solutions increases. As such, the development of hardware designed specifically to emulate brain processes and run brain-inspired algorithms has increased accordingly. Informed by neuroscience research which often uses simulation and hardware emulation of the brain to better understand biological processes, "computing for neuroscience", there is an increasing trend to use the knowledge obtained in neuroscience to develop hardware to emulate the brain for computing, "neuroscience for computing".

The current study reviews the history and development of brain-inspired systems and its culmination in neuromorphic hardware and contrasts with its traditional von Neumann counterpart.

History of Neuromorphic Systems

The use of neural networks and other brain-inspired software systems, especially in the classification and pattern recognition domains, is widespread and breakthroughs such as DeepMind's AlphaGo often receive widespread media attention. Less often covered by the mass media is the application of brain research in the development of hardware emulating brain cells and systems to both develop more efficient processors and further understand the brain. Neuromorphic computing, first conceived in 1990 by Carver Mead initially referred to analogue components that imitate any biological neural system. More recently, however, the term has become more generalised to refer to artificial neural networks (including those run on von Neumann architectures), or components that are biologically inspired.

The foundations for modern neuromorphic systems was laid in 1943 with the first mathematical model of a neural network in Mcculloch and Pitts' "*A Logical Calculus of Ideas Immanent in Nervous Activity*". The neuron described in the study is the same neuron that is found in deep learning networks today.

Although it would not be published until much later (1969), in 1948 Alan Turing wrote his paper *Intelligent Machinery*. The paper described neurons arranged in a graph with 'modifier devices' that would propagate or cancel signals (Turing, 1969).

In 1958 Frank Rosenblatt discussed the creation of the perceptron algorithm in the previous year at a US Navy press conference (Mason, Stewart and Gill, 1958).
At this time the perceptron had been implemented in software but the US Navy later built a hardware implementation, the "Mark I Perceptron", for image recognition in 1960.

In 1959 Hubel and Wiesel discovered what they referred to as *simple cells* and *complex cells* through experimentation with the primary visual cortex of cats. Simple cells were neurons that responded to lines, or edges, at a particular angle whereas complex cells would respond to edges at any angle regardless of the placement in the neuron's perceptive field (Hubel and Wiesel, 1959). This study laid the groundwork for showing that the visual system could construct complex shapes from simple base detections and informed later work in pattern recognition.

In the following year, 1960, Widrow and Hoff built an adaptive linear neuron, ADELINE. ADELINE utilised electrochemical memistors which were inspired by synapses. The integration of the signals from one of the three terminals on the device modified the resistance between the other two. This allowed behaviour similar to that of a biological neuron (Widrow, 1987).

In 1971, Leon Chua imagined the "memristor" which he considered to be the "fourth fundamental circuit element" (Chua, 1971) although the truth of that statement is still being debated (Mouttet, 2012).

In 1980, Mead, Carver and Conway published "*Introduction to VLSI systems*". VLSI is very large scale integration in which an integrated circuit is constructed through the use of the combination of a large number of transistor on a single chip, for example, a microprocessor (Mead and Conway, 1980). This laid the foundations for later work in creating the first neural inspired chip.

In 1982, John Hopfield popularised an earlier described associative neural network, now widely referred to as Hopfield networks (Hopfield, 1982). In a Hopfield network, recurrency is added to a perceptron, with the limitation that a neuron may not connect to itself and connection must be symmetric.
In the same year, Bienenstock, Cooper and Munro described their theory of synaptic plasticity in the brain, later called the BCM plasticity rule. Unlike the earlier Hebbian learning, BCM allowed for long term depression (LTD) as well as long term potentiation (LTP) (Bienenstock, Cooper, and Munro 1982).
BCM has since been evidenced in the visual cortex and hippocampus.

In 1986 Rumelhart, Hinton and Williams showed experimentally, for the first time, that back propagation can be applied to neural networks (Rumelhart, Hinton and Williams, 1991).

In 1987 Carpenter and Grossberg introduced adaptive resonance theory; a description of aspects of brain information processing. They further explored a number of neural network models which addressed issues of pattern recognition and prediction using both supervised and unsupervised learning (Carpenter and Grossberg, 1987).

In 1989 Mead and Ismail's book "*Analog VLSI Implementation of Neural Systems"* was published containing proceedings from a symposium in May of the same year. The first neural-inspired chips were described. This included chips for maze-solving, image focusing, cochlear models and components for general purpose analogue neural computers (Mead and Ismail, 1989).
Building on Barlow's work from 1989, Bell and Sejnowski showed that the independent component of a natural scene is an edge filter. This work provided a mathematical foundation for neural plasticity which proved neurons

modified their synaptic weights to identify specific components of a scene (Bell and Sejnowski, 1989).
At the same time, Mead found that complementary metal-oxide-semiconductor (CMOS) transistors operating below their activation threshold acted in much the same way as the ion channels in neurons and as such could be used as neuron replacements in analogue circuits (Mead, 1990). This is when the term 'neuromorphic' was first coined.

In 1990 the work of Vapnik and Chervonenkis, which had started in 1960, culminated in a statistics-based explanation of computational learning.

In the following year, 1991, Mahowald and Douglas, using Mead's observations from 1989, invented silicon neurons; spiking elements in VLSI chips that did not require any digital software simulation (Mahowald and Douglas, 1991).

In 2006 Hasler, Farquhar and Gordon presented a paper outlining the creation of "*an analog circuit capable of accurately emulating large complex cells, or multiple less complex ones*" and coined the term 'field programmable neural array', or FNPA, to describe it. The FNPA contained components simulating the effects of active channels, dendrites, and synapses (Hasler, Farquhar and Gordon, 2006).

In 2008 HP claimed to have created Chua's missing memristor from his proposition in 1971 (Strukov et al, 2009). However, some controversy surrounds the claim not only due to the issues surrounding Chua's original claim but also due to the supposed lack of scientific method employed in the development of HP's memristor (Mouttet, 2012a)(Mouttet, 2012b). HP also used their memristor in neuromorphic circuits in the same year.

In 2011 researchers at MIT produced the first chip that was capable of simulating the analogue ion-based communication of neurons by utilising 400 transistors. The chip was produced using standard CMOS manufacture techniques.

In 2012 researchers at Purdue proposed several new designs for neuromorphic architectures that could operate on very low power. The designs utilised new spin devices; magneto-electronic devices that utilise the spin of electrons alongside their electrostatic charge to encode information.

In 2014 Nugent and Molter published a paper introducing Anti-Hebbian and Hebbian (AHaH) computing as a complementary system to today's von Neumann architecture with a view to reducing the impact of the von Neumann bottleneck. The paper, "*AHaH Computing–From Metastable Switches to Attractors to Machine Learning*", showed that AHaH synaptic plasticity can form "computationally complete set of logic functions" alongside extracting independent features from data streams (Nugent and Molter 2014).
Nugent continued this work in 2016 and detailed how to implement the AHaH plasticity rule in circuits composed of memristive elements.

Contemporary Neuromorphic Hardware

The culmination of the research starting in 1943, among other work, is the increasing development of modern hardware inspired by brain systems. In contrast to the current standard in computer hardware of von Neumann architecture, neuromorphic devices take inspiration from brain systems in their designs. Von Neumann systems held advantages over previous systems in that programs and associated data are stored in memory and data is exchanged between the processor and memory as a program executes. This means that the program and data do not have to be input to the computer manually as it runs. However, limitations of von Neumann systems may be rectified by neuromorphic hardware.

Neuroscience-based advancements for neuromorphic hardware have largely been driven by the limited speed with which simulations of the brain can be accomplished on von Neumann computers. These simulations are often orders of magnitude slower than real time. In fact, even comparing two of the most high performance architectures used by the Human Brain Foundation, SpiNNaker and BrainScaleS, shows a significant margin. SpiNNaker utilises more than half a million custom, but traditional digital, ARM processors in an asynchronous network based on spiking neurons, in an attempt to simulate the brain. In comparison, BrainScaleS uses a physical model analogue system utilising four million artificial neurons and one billion artificial synapses which allows simulation of up to ten thousand times real-time speeds (Meier, 2016).

Inspiration from brain systems is increasingly prevalent in architectures designed for online learning. Here, the brain is a good model for the device's requirements in that processing, long and short-term memory and learning are all performed efficiently in terms of power and weight requirements (research.ibm.com, 2014). Neuromorphic hardware varies between implementations but that which aims not to simulate, but more closely mimic the behaviour of neurons is analogue based.

In order to understand neuromorphic systems, it is important to understand the biological hardware used in the brain; neurons.

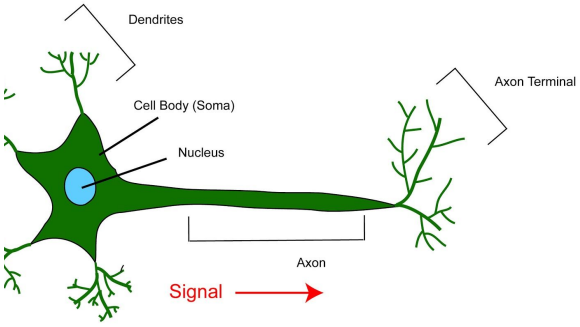

Figure 1. The structure of a Neuron (CorticalBrain, 2017)

Although there are various types of neuron, all neurons are made up of three key elements; dendrites, an axon and a soma (Fig 1). The soma is the processor of the cell. It receives inputs as action potentials from the dendrites and, if an activation threshold is reached, causes an action potential along the axon (Gerstner et al., n.d.). This process is caused by the membrane potential of neurons. Current flowing into and out of a neuron causes a change in the membrane potential. It is this potential that must reach a threshold level to cause the neuron to activate. The membrane potential of a cell is driven by concentration gradients across

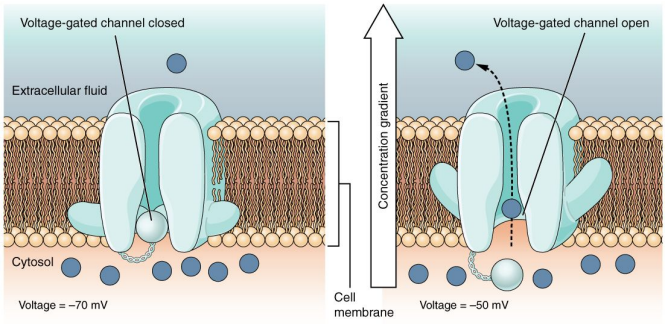

Figure 2. Voltage-gated ion channels (OERPUB, 2017)

the cell membrane, which is regulated by voltage-gated ion channels of the cell (Fig. 2). Input action potentials to the neuron cause the sodium channels of the cell open. If the activity is large enough the activation threshold is reached, and this causes further opening of sodium channels and therefore an increase in the action potential of the neuron. When a critical voltage is reached, the sodium channels close and potassium channels open which causes a counter current, thereby quickly returning the neuron to its resting state (Physiology and Maintenance - Vol. V; Neurons, Action Potentials, and Synapses, 2017). In many cases, the membrane potential settles at a value lower than the resting potential of the cell due to the hyperpolarisation caused by the potassium channels. This results in a refractory period in which the neuron cannot reactivate. The firing patterns caused by this process result in spiking action potentials being delivered via the axon to the next neuron in the network. Neuromorphic hardware attempts to mimic these processes through the use of materials that exhibit stochastic behaviour, like that seen

in cells, as opposed to the largely deterministic nature of traditional computer chips. For example, in 2015 IBM Research patented a neuromorphic synapse that used a phase-change material (International Business Machines Corporation, 2017). At the nano scale, phase-change materials display behaviour very similar to that of neuron membranes, and the manufactured neuron is, therefore, able to have an artificial counterpart to the soma, with connections to other artificial neurons forming the axon and dendrites. IBM's artificial soma receives action potentials from its dendrites and then generates an output on its axon dependant on the integration of the incoming action in the soma potential exceeding the required threshold, making it a memristor.
When the artificial neuron activates, the phase-change device fires a spike and then is reset back to a non-conductive state. The stochastic nature of the artificial neuron is achieved by the phase-change material which behaves in such a way as to return the soma to a slightly differing state to the original. IBM's artificial synapse has been shown to successfully detect structured data, namely the IBM and Watson logos, within noise unsupervised (IBM Research, 2017).

The development of neuromorphic hardware is often attributed to simulating brain systems (Humanbrainproject.eu, 2017), however just as von Neumann architecture offered advantages over its predecessor, neuromorphic systems may offer solutions to some of the limitations current computer hardware faces. Similarly, as the ways we use computers change, for example the increased use of pattern recognition in personal computing devices, a specialised chip that can process this information efficiently will be beneficial.

While there are a wide range of applications for neuromorphic chips in online supervised and unsupervised learning; from self-driving cars (Hall-Geisler, 2017) to expert systems like Watson in the health sector (Bakkar et al., 2017) to domain-specific pattern recognition, there has been criticism that neuromorphic hardware does not have the flexibility to be a feasible replacement for modern computing devices. However, the fact that neural networks can be Turing-complete (Graves, Wayne and Danihelka, 2017) and the advantages neuromorphic hardware hold with respect to power and weight efficiency, make them a strong candidate for cloud-based devices, even if they remain distinct from general consumer computing. Moreover, the use of graphical processing units (GPUs) is widespread and neuromorphic hardware may be useful as a supplementary device for von Neumann computers as a "neuromorphic processing unit" or NPU.

A Solution to Von Neumann Limitations

The fundamental differences in architecture between von Neumann and neuromorphic hardware present the opportunity for the limitations currently faced in developing more advanced von Neumann systems. These differences can be split into five high-level differences; Spiking, Plasticity, Learning and Adaptability, Connectedness, and Hardware Components. These differences may allow us to realise the benefits found in computations by the brain. For example, the brain is a factor of roughly 10 million times more power efficient than conceivable computers (Mead, 1990). Indeed, this is what is seen when considering current neuromorphic chips. For example, IBM's TrueNorth chip has a power density of 20 mW/cm^2 whereas modern, high-performance data centres operate at 1,292 mW/cm^2.

Spiking

In the brain, spiking is the change in voltage of the output of a neuron and is the method by which signals are transmitted from one neuron to the next when activated. For example, a neuron, *N*, may have a resting membrane potential of -50mV, with an activation potential of -70mV. When the summative potential from excitatory (positive) and inhibitory (negative) signals is greater than this activation threshold, *N* will fire. Further, post-activation *N* will enter a refractory period due to the depletion of the

pool of neurotransmitters at *N*'s synapse. This will typically last a few milliseconds, wherein *N* will be unable to activate. A short while after this, a large net excitatory signal may be able to reactivate *N* and later still *N* will return to its original state. The spikes produced during activation can also encode information in the timing of the spikes. Successful software implementations of spiking neurons in spiking neural networks (SNNs) have been developed (The GENESIS Simulator, 2017) (Delorme and Thorpe, 2017). In addition to the use of artificial synapses and neurons found in artificial neural networks (ANNs), in an SNN time is also used as an encoding mechanism, as in the brain. Information is encoded in the frequency of spiking from single action potentials and can be found in both software and hardware: For example, Thorpe (2012) employed an unsupervised, STDP-inspired learning rule to train a network containing hardware-based spiking retinas for image processing. This method proved effective in that during a simulation involving cars passing on a highway, within minutes the system could count the cars passing in the various lanes.
Spiking is not seen in the digital systems in use in von Neumann architectures where output has only binary values or analogue continuous signals, in contrast to the output of spiking systems where information is transmitted via action potentials and encodes information within the activation of a neuron.

Plasticity
Spiking also enables spike timing dependant plasticity (STPD). STDP is the process by which synaptic connections may be strengthened or weakened and is a key component for learning. Synapses are strengthened when an action potential to a neuron, on average, occurs immediately before the output spike of that neuron. This is Hebbian learning but more often referred to as long term potentiation (LTP) since the discovery of STDP. LTP is often simplified to *"what fires together wires together"*, although in reality LTP only occurs with causation and therefore cells firing simultaneously will not experience LTP. In the case of STDP-like methods in SNNs, there is even more similarity to the learning that occurs in the brain due to the simulation of membrane potential.
This behaviour heavily contrasts with that of chips used in general computing today, where given an input, a specific output is always given rather than a changeable output as a neuromorphic system adapts to the stimulus it receives.

Learning and Adaptability
A key application of neuromorphic hardware is live learning. For example, in biological systems synaptic scaling during sleep allows for the advantages of STDP through the down-scaling of synaptic strength, and therefore efficacy, without the harmful over-excitation of neurons. Whereas LTP and LTD are a synaptic process, synaptic scaling is a cellular homeostatic process through which the incoming synapses to a neuron are normalised. This reduces the energy requirements of synapses whilst maintaining the learned structure. The normalisation of inputs also allows for the continuation of selectivity over time by reducing the chances of LTP occurring from spurious action potentials from strengthened synapses. Due to the fact that all synapses are equally dampened, the relative strengths remain the same and the effects of STDP are maintained without the resulting long term damage to neurons.
The traditional von Neumann architecture has no learning and is, therefore, unable to adapt to dynamically changing inputs nor able to self-program in the way brains and neuromorphic systems are able to.

Connectedness
In the brain neurons are connected by extended structures called axons which may connect a neuron to thousands of other neurons via their dendrites. This is in contrast to the tens (or less) of connections seen between components on a modern chip. The density of connections within the brain, and to a lesser extent neuromorphic chips, allows for the creation of power-efficient, dense neural networks at a hardware level, and

enables the previously discussed benefits of plasticity, learning and adaptability. Although the density of the brain has not yet been matched, the TrueNorth system contained 1,000,000 neurons and 256 synapses and FACETS, a neuromorphic system in development at the University of Heidelberg contains 180,000 neurons with $4x10^7$ synapses per wafer.

Components

The fundamental components of traditional chips and the brain vary significantly. Where the brain uses the described neurons and axon connections, traditional chips use transistors and separated memory that describes von Neumann architecture (Fig. 3). These systems are now reaching fundamental, atomic physical limits past which neuromorphic chips may offer alternative solutions to the techniques of transistor miniaturisation seen in hardware production in past years.

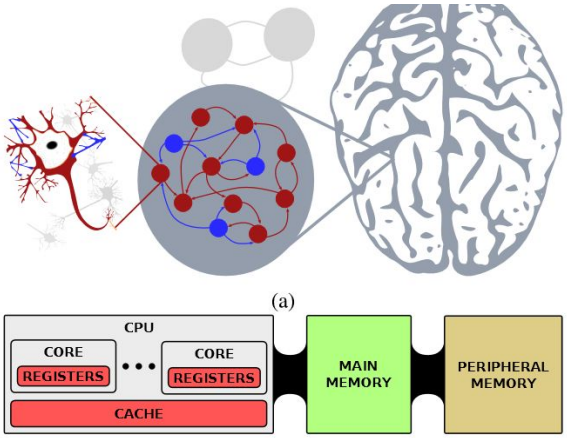

Figure 3. Comparison between von Neumann and Neuromorphic architectures (Indiveri and Liu, 2015)

Moore's Law

Moore's Law, named after Intel co-founder Gordon Moore's observations in 1965, states that the number of transistor on a dense integrated circuit will double roughly every two years. The law is stated to break down around 2025 (IEEE Spectrum: Technology, Engineering, and Science News, 2017), although some have argued both that the law has not been applicable since 2016 and that it acted as a self-fulfilling prophecy due to the industry using it as a target to meet (Bright, 2017).

The difficulty in continuing to match Moore's law is twofold. Firstly it is due to the fact that the development of highly dense VLSI chips, that are in widespread use today, are reaching hard, physical limits. For example, the effects of quantum tunnelling, whereby electrons 'pass through' barriers (Fig. 4), become increasingly problematic for transistors as they approach smaller sizes. As soon as any internal barrier within the transistor reaches 1 nm or less, current will continue to flow when the transistor is off (Bright, 2017). Whilst debate continues in the industry on how to overcome, or utilise, these effects neuromorphic alternatives have arisen.

Secondly, economics has a large part to play in the direction computing will take as costs continue to increase to develop smaller transistors. This may offer an opportunity for neuromorphic chips to become more widespread as the cost to switch becomes lessened. This is not to say that industry is ignoring neuromorphic technology. Indeed, as recently as September 2017 Intel unveiled a new neuromorphic chip designed for use in

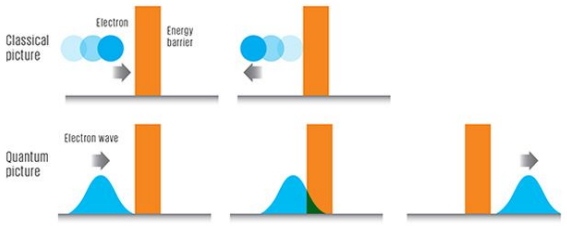

Figure 4. The effects of quantum tunneling (IEEE Spectrum, 2017)

machine learning which is stated to be 1,000 times more power efficient than traditional chips.

Dennard Scaling

Related to Moore's law is Dennard scaling which states that the performance per watt of computing increases exponentially year on year. However, around 2006 this effect began to break down, with the performance increase slowing in relation to the number of transistors in integrated circuits. This is due to the current leakage caused by quantum tunnelling in smaller transistors (IEEE Spectrum: Technology,

Engineering, and Science News, 2017). Current leakage also results in increased heat on the chip and therefore increases in cycle frequency that previously was used to increase the computing power is no longer possible.

Von Neumann Bottleneck

The von Neumann bottleneck is a throughput limitation of traditional computer architectures whereby the limiting factor for a computer's speed becomes the speed at which memory can be transferred between the CPU and memory. This is due to the architecture of von Neumann systems where the processor and memory are separate. Where memory storage density has seen dramatic and consistent improvement in recent years, memory bandwidth has not and therefore the von Neumann bottleneck becomes an ever more prevalent issue. The apposition of processing and memory in many neuromorphic architectures reduce or completely eliminate the issue of the von Neumann bottleneck as there is no latency when fetching data from memory as the processing of data is inextricably linked to its storage (Fig. 3) (Indiveri and Liu, 2015). These systems take inspiration from the fact that neurons are the base element in both processing and storing information in the brain. This means that the processing of data is inextricably tied to the storage of that same data thereby making data throughput for processing many times more efficient than that found in traditional hardware.

These differences between architectures may cause significant difficulties in the adoption of neuromorphic hardware despite the potential benefits it offers. This is due to the present difficulty in writing non-trivial general programs for neuromorphic hardware and therefore the lack of support for general use computing it provides.

Conclusion

The present study has shown neuromorphic hardware to be a promising technology for hardware-based neural network processing. However, the significant difference between neuromorphic and von Neumann devices likely means that a dramatic shift to neuromorphic hardware is unlikely, especially in consumer devices. In large-scale data centres where a significant amount of ANN or SNN processing is being performed, it may be cost-efficient, due to the power efficiency of neuromorphic devices, to use neuromorphic hardware as a GPU-like peripheral device in conjunction with von Neumann systems. This would allow data centres to streamline and cheapen neural processing whilst maintaining support for software developed for traditional hardware. Furthermore, as costs reduce and if the use of brain-inspired processing continues to rise, it may also be beneficial to consumers to have 'NPUs' for local neural processing, especially in the cases where weight and power efficiency are a key factor.

Future Work

The present study has shown that a significant limiting factor in the adoption of neuromorphic hardware as general computing devices is the difficulty in writing general code for them. It may, therefore, be worthwhile to explore the feasibility of automated compilation of existing code for neuromorphic hardware, and the performance impact this has when compared with von Neumann instances of the same program.

References

Bakkar, N., Kovalik, T., Lorenzini, I., Spangler, S., Lacoste, A., Sponaugle, K., Ferrante, P., Argentinis, E., Sattler, R. and Bowser, R. (2017). Artificial intelligence in neurodegenerative disease research: use of IBM Watson to identify additional RNA‐binding proteins altered in amyotrophic lateral sclerosis. *Acta Neuropathol*. [online] Available at: https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=HLW03040USEN& [Accessed 14 Dec. 2017].

Bell, A. and Sejnowski, T. (1997). The "independent components" of natural scenes are edge filters. *Vision Research*, 37(23), pp.3327-3338.

Bichler, O., Querlioz, D., Thorpe, S., Bourgoin, J. and Gamrat, C. (2012). Extraction of temporally correlated features from dynamic vision sensors with spike-timing-dependent plasticity. *Neural Networks*, 32, pp.339-348.

Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982).Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. J. Neurosci., 2, 32–48.

Bright, P. (2017). *Moore's law really is dead this time*. [online] Ars Technica. Available at: https://arstechnica.com/information-technology/2016/02/moores-law-really-is-dead-this-time/ [Accessed 14 Dec. 2017].

C. D. Schuman, T. E. Potok, R. M. Patton, J. D. Birdwell, M. E. Dean, G. S. Rose and J. S. Plank, "A Survey of Neuromorphic Computing and Neural Networks in Hardware," arXiv:1705.06963, 2017.

Carpenter, G.A. & Grossberg, S. (2003), Adaptive Resonance Theory, In Michael A. Arbib (Ed.), The Handbook of Brain Theory and Neural Networks, Second Edition (pp. 87-90). Cambridge, MA: MIT Press.

Carver Mead and Lynn Conway, Introduction to VLSI Systems (Addison-Wesley, Reading, MA, 1980).

Chua L (1971) Memristor—the missing circuit element. IEEE Transactions on Circuit Theory 18: 507–519.
Citri, A. and Malenka, R. (2007). Synaptic Plasticity: Multiple Forms, Functions, and Mechanisms. *Neuropsychopharmacology*, 33(1), pp.18-41.
cortex. J. Physiol. 148, 574-591.

CorticalBrain (2017). Neuron of Human Brain. [image] Available at: http://corticalbrain.com/neuroanatomy/home-content/Brain-Structure-and-Neurons-neurochemicals-and-neurocircuits.php [Accessed 14 Dec. 2017].

D. E. Rumelhart, G. E. Hinton, and R. J. Williams, Learning internal representations by error propagation, in Parallel Distributed Processing: Explorations in the Microstructure of Cognition, D. E. Rumelhart (AAAI-91), July 1991, pp. 762-767.

Delorme and Thorpe (2017). *SpikeNET*.

De Pittà, M., Brunel, N. and Volterra, A. (2016). Astrocytes: Orchestrating synaptic plasticity?. *Neuroscience*, 323, pp.43-61.

Farquah, E., Gordon, C. and Hasler, P. (2017). A Field Programmable Neural Array. In: *ISCAS 2006*. IEEE.

Gerstner, W., Kistler, W., Naud, R. and Paninski, L. (n.d.). *Neuronal dynamics*. Cambridge University Press, pp.3.0.

Graves, A., Wayne, G. and Danihelka, I. (2017). *Neural Turing Machines*. Google DeepMind.

Hall-Geisler, K. (2017). *Cortica teaches autonomous vehicles with unsupervised learning*. [online] TechCrunch. Available at: https://techcrunch.com/2017/07/13/cortica-teaches-autonomous-vehicles-with-unsupervised-learning/ [Accessed 14 Dec. 2017].

Hasler, P., Farquhar, E. and Gordon, C. (2017). Building Large Networks of Biological Neurons. In: *EMBS Annual International Conference*. [online] IEEE, pp.6548-6551. Available at: http://ieeexplore.ieee.org/document/4030595/ [Accessed 14 Dec. 2017].

HBP Neuromorphic Computing Platform Guidebook. (2017). [ebook] Human Brain Project. Available at: https://electronicvisions.github.io/hbp-sp9-guidebook/ [Accessed 11 Dec. 2017].

Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational abilities. Proceedings of the National Academy of Sciences, 79(8), pp.2554-2558.

Hubel, D. H. & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate
Humanbrainproject.eu. (2017). *Brain Simulation*. [online] Available at: https://www.humanbrainproject.eu/en/brain-simulation/ [Accessed 11 Dec. 2017].

IBM Research (2017). *All-memristive neuromorphic computing with level-tuned neurons*. [video] Available at:

https://www.youtube.com/watch?v=hXeO8Kzz3bo [Accessed 13 Dec. 2017].

IEEE Spectrum: Technology, Engineering, and Science News. (2017). *The Tunneling Transistor*. [online] Available at: https://spectrum.ieee.org/semiconductors/devices/the-tunneling-transistor [Accessed 14 Dec. 2017].

IEEE Spectrum: Technology, Engineering, and Science News. (2017). *Transistor Wars*. [online] Available at: https://spectrum.ieee.org/semiconductors/devices/transistor-wars [Accessed 14 Dec. 2017].

IEEE Spectrum (2017). *Back or Through*. [image] Available at: https://spectrum.ieee.org/semiconductors/devices/the-tunneling-transistor [Accessed 14 Dec. 2017].

Indiveri, G. and Liu, S. (2015). Memory and Information Processing in Neuromorphic Systems. *Proceedings of the IEEE*, 103(8), pp.1379-1397.

International Business Machines Corporation (2017). *Neuromorphic synapses*. US 20160125287 A1.

Irizarry-Valle, Y. and Parker, A. (2015). An Astrocyte Neuromorphic Circuit That Influences Neuronal Phase Synchrony. *IEEE Transactions on Biomedical Circuits and Systems*, 9(2), pp.175-187.

J. Seo, et al., "A 45nm CMOS neuromorphic chip with a scalable architecture for learning in networks of spiking neurons," in Custom Integrated Circuits Conference (CICC), 2011 IEEE. IEEE, 2011, pp. 1–4.

Maass, W. (1997). Networks of spiking neurons: The third generation of neural network models. *Neural Networks*, 10(9).

Mahowald, M. and Douglas, R. (1991). A silicon neuron. *Nature*, 354(6354), pp.515-518.

Mason, H., Stewart, D. and Gill, B. (1958). Rival. *The New Yorker*, [online] p.44. Available at: https://www.newyorker.com/magazine/1958/12/06/rival-2 [Accessed 12 Dec. 2017].

Mead, C. (1990). Neuromorphic electronic systems. *Proceedings of the IEEE*, 78(10), pp.1629-1636.

Mead, C. and Ismail, M. (1989). *Analog VLSI Implementation of Neural Systems*. Boston, MA:Springer US.

Meier, K. (2016). *Neuromorphic Computing - Concepts, Achievements & Challenges*.

Mouttet, Blaise. 2012b. The memristor and the scientific method. arXiv.org. [Online]. Available at: https://arxiv.org/pdf/1611.04456.pdf. Accessed 13th December 2017.

Mouttet, Blaise. Memistors, Memristors, and Memresistors [Online]. Version 25. 2012a. Available from: https://memresistor.wordpress.com/article/memistors-memristors-memresistors-and-23zgknsxnlchu-7/.

Nugent, M. and Molter, T. (2014). AHaH Computing–From Metastable Switches to Attractors to Machine Learning. *PLoS ONE*, 9(2), p.e85175.

OERPUB (2017). *Voltage-Gated Channels*. [image] Available at: http://oerpub.github.io/epubjs-demo-book/content/m46526.xhtml [Accessed 14 Dec. 2017].

Pais-Vieira, M., Chiuffa, G., Lebedev, M., Yadav, A. and Nicolelis, M. (2015). Building an organic computing device with multiple interconnected brains. *Scientific Reports*, 5(1).

Physiology and Maintenance - Vol. V; Neurons, Action Potentials, and Synapses. (2017). [ebook] Encyclopedia of Life Support Systems, p.4. Available at: http://www.eolss.net/sample-chapters/c03/e6-54-09-04.pdf [Accessed 14 Dec. 2017].

Rasmussen, N. (2017). Calculating Space and Power Density Requirements for Data Centers. [online] Schneider Electric, p.3. Available at: http://www.apc.com/salestools/NRAN-8FL6LW/NRAN-8FL6LW_R0_EN.pdf [Accessed 11 Dec. 2017].

Research.ibm.com. (2014). *IBM Research - Brain Power: The origins of a brain-inspired chip*. [online] Available at:

http://www.research.ibm.com/cognitive-computing/brainpower [Accessed 14 Dec. 2017].

Rovere G., Ning Q., Bartolozzi C., Indiveri G. (2014). Ultra low leakage synaptic scaling circuits for implementing homeostatic plasticity in neuromorphic architectures, in International Symposium on Circuits and Systems, (ISCAS) (Melbourne VIC: IEEE;), 2073–2076.

Strukov, D., Snider, G., Stewart, D. and Williams, R. (2009). The missing memristor found. Nature, 453, pp.80-83.

The GENESIS Simulator. (2017). genesis-sim.org.

Thorpe, S. (2012). Spike-Based Image Processing: Can We Reproduce Biological Vision in Hardware?. In: *European Conference on Computer Vision*. pp.516-521.

Turing, A.M. (1948), 'Intelligent Machinery', National Physical Laboratory Report, in Meltzer and Michie (1969).

Widrow B (1987) The original adaptive neural net broom-balancer. In: Proc. 1987 IEEE International Symposium on Circuits and Systems. volume 2, pp. 351–357.

Wijekoon, J. and Dudek, P. (2008). Compact silicon neuron circuit with spiking and bursting behaviour. *Neural Networks*, 21(2-3), pp.524-534.

Wijekoon, J. H. B. & Dudek, P., (2006). Simple analogue VLSI circuit of a cortical neuron. In IEEE int. conf. on electronics, circuits and systems (pp. 1344–1347).