

Zurich Open Repository and Archive University of Zurich University Library Strickhofstrasse 39 CH-8057 Zurich www.zora.uzh.ch

Year: 2022

2022 roadmap on neuromorphic computing and engineering

Christensen, Dennis V ; Dittmann, Regina ; Linares-Barranco, Bernabe ; Sebastian, Abu ; Le Gallo, Manuel ; Redaelli, Andrea ; Slesazeck, Stefan ; Mikolajick, Thomas ; Spiga, Sabina ; Menzel, Stephan ; Valov, Ilia ; Milano, Gianluca ; Ricciardi, Carlo ; Liang, Shi-Jun ; Miao, Feng ; Lanza, Mario ; Quill, Tyler J ; Keene, Scott T ; Salleo, Alberto ; Grollier, Julie ; Marković, Danijela ; Mizrahi, Alice ; Yao, Peng ; Yang, J Joshua ; Indiveri, Giacomo ; Strachan, John Paul ; Datta, Suman ; Vianello, Elisa ; Valentian, Alexandre ; Feldmann, Johannes ; et al

Abstract: Modern computation based on von Neumann architecture is now a mature cutting-edge science. In the von Neumann architecture, processing and memory units are implemented as separate blocks interchanging data intensively and continuously. This data transfer is responsible for a large part of the power consumption. The next generation computer technology is expected to solve problems at the exascale with 10¹⁸ calculations each second. Even though these future computers will be incredibly powerful, if they are based on von Neumann type architectures, they will consume between 20 and 30 megawatts of power and will not have intrinsic physically built-in capabilities to learn or deal with complex data as our brain does. These needs can be addressed by neuromorphic computing systems which are inspired by the biological concepts of the human brain. This new generation of computers has the potential to be used for the storage and processing of large amounts of digital information with much lower power consumption than conventional processors. Among their potential future applications, an important niche is moving the control from data centers to edge devices. The aim of this roadmap is to present a snapshot of the present state of neuromorphic technology and provide an opinion on the challenges and opportunities that the future holds in the major areas of neuromorphic technology, namely materials, devices, neuromorphic circuits, neuromorphic algorithms, applications, and ethics. The roadmap is a collection of perspectives where leading researchers in the neuromorphic community provide their own view about the current state and the future challenges for each research area. We hope that this roadmap will be a useful resource by providing a concise yet comprehensive introduction to readers outside this field, for those who are just entering the field, as well as providing future perspectives for those who are well established in the neuromorphic computing community.

DOI: https://doi.org/10.1088/2634-4386/ac4a83

Posted at the Zurich Open Repository and Archive, University of Zurich ZORA URL: https://doi.org/10.5167/uzh-231198 Journal Article Published Version



The following work is licensed under a Creative Commons: Attribution 4.0 International (CC BY 4.0) License.

Originally published at:

Christensen, Dennis V; Dittmann, Regina; Linares-Barranco, Bernabe; Sebastian, Abu; Le Gallo, Manuel; Redaelli, Andrea; Slesazeck, Stefan; Mikolajick, Thomas; Spiga, Sabina; Menzel, Stephan; Valov, Ilia; Milano, Gianluca; Ricciardi, Carlo; Liang, Shi-Jun; Miao, Feng; Lanza, Mario; Quill, Tyler J; Keene, Scott T; Salleo, Alberto; Grollier, Julie; Marković, Danijela; Mizrahi, Alice; Yao, Peng; Yang, J Joshua; Indiveri, Giacomo; Strachan, John Paul; Datta, Suman; Vianello, Elisa; Valentian, Alexandre; Feldmann, Johannes; et al (2022). 2022 roadmap on neuromorphic computing and engineering. Neuromorphic Computing and Engineering, 2(2):022501. DOI: https://doi.org/10.1088/2634-4386/ac4a83

ROADMAP · OPEN ACCESS

2022 roadmap on neuromorphic computing and engineering

To cite this article: Dennis V Christensen et al 2022 Neuromorph. Comput. Eng. 2 022501

View the article online for updates and enhancements.

You may also like

- <u>Roadmap on emerging hardware and</u> <u>technology for machine learning</u> Karl Berggren, Qiangfei Xia, Konstantin K Likharev et al.
- <u>The 2022 applied physics by pioneering</u> women: a roadmap Begoña Abad, Kirstin Alberi, Katherine E Ayers et al.
- <u>Emerging memory technologies for</u> <u>neuromorphic computing</u> Chul-Heung Kim, Suhwan Lim, Sung Yun Woo et al.

NEUROMORPHIC

Computing and Engineering

ROADMAP



OPEN ACCESS

RECEIVED

REVISED

PUBLISHED

2022 roadmap on neuromorphic computing and engineering

Dennis V Christensen¹, Regina Dittmann², Bernabe Linares-Barranco³, Abu Sebastian⁴^(b), Manuel Le Gallo⁴^(b), Andrea Redaelli⁵, Stefan Slesazeck⁶^(b), Thomas 14 May 2021 Mikolajick^{6,7}, Sabina Spiga⁸, Stephan Menzel⁹, Ilia Valov⁹, Gianluca Milano¹⁰, Carlo Ricciardi¹¹, Shi-Jun Liang¹², Feng Miao¹², 23 October 2021 Mario Lanza¹³, Tyler J Quill¹⁴, Scott T Keene¹⁵, Alberto Salleo¹⁴, Julie Grollier¹⁶, ACCEPTED FOR PUBLICATION Danijela Marković¹⁶, Alice Mizrahi¹⁶, Peng Yao¹⁷, J Joshua Yang¹⁷, 12 January 2022 Giacomo Indiveri¹⁸, John Paul Strachan¹⁹, Suman Datta²⁰, Elisa Vianello²¹, 20 May 2022 Alexandre Valentian²², Johannes Feldmann¹, Xuan Li²³, Wolfram H P Pernice^{24,25}, Harish Bhaskaran²³, Steve Furber²⁶, Emre Neftci²⁷, Franz Scherr²⁸, Original content from Wolfgang Maass²⁸, Srikanth Ramaswamy²⁹^(D), Jonathan Tapson³⁰ this work may be used Priyadarshini Panda³¹, Youngeun Kim³¹, Gouhei Tanaka³², Simon Thorpe³³, under the terms of the Creative Commons Chiara Bartolozzi³⁴, Thomas A Cleland³⁵, Christoph Posch³⁶, ShihChii Liu¹⁸, Attribution 4.0 licence Gabriella Panuccio³⁷, Mufti Mahmud³⁸, Arnab Neelim Mazumder³⁹, Any further distribution Morteza Hosseini³⁹, Tinoosh Mohsenin³⁹, Elisa Donati¹⁸, Silvia Tolu¹⁴, of this work must maintain attribution to Roberto Galeazzi⁴⁰, Martin Ejsing Christensen⁴¹, Sune Holm⁴², Daniele Ielmini⁴³ and the author(s) and the N Pryds¹ title of the work, journal citation and DOI. Department of Energy Conversion and Storage, Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark Peter Gruenberg Institute 7, Forschungszentrum Juelich GmbH, 52425 Juelich, Germany and JARA-FIT, RWTH Aachen University, **(i)** 52056 Aachen, Germany Instituto de Microelectrónica de Sevilla (IMSE-CNM), CSIC and Universidad de Sevilla, 41092 Seville, Spain IBM Research - Zurich, Switzerland STMicroelectronics, Agrate, Italy 6 NaMLab gGmbH, 01187 Dresden, Germany ⁷ Institute of Semiconductors and Microsystems, TU Dresden; Dresden, Germany CNR-IMM, Unit of Agrate Brianza, Via C. Olivetti 2, Agrate Brianza (MB), Italy FZ Juelich (PGI-7), Juelich, Germany ¹⁰ Advanced Materials Metrology and Life Science Division, INRiM (Istituto Nazionale di Ricerca Metrologica), Torino, Italy ¹¹ Department of Applied Science and Technology, Politecnico di Torino, Torino, Italy ¹² National Laboratory of Solid State Microstructures, School of Physics, Collaborative Innovation Center of Advanced Microstructures, Nanjing University, Nanjing, People's Republic of China ¹³ Physical Sciences and Engineering Division, King Abdullah University of Science and Technology (KAUST), 23955-6900 Thuwal, Saudi Arabia ¹⁴ Department of Materials Science and Engineering, Stanford University, Stanford, CA 94305, United States of America ¹⁵ Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, United Kingdom ¹⁶ Unité Mixte de Physique, CNRS, Thales, Université Paris-Saclay, 91767 Palaiseau, France ¹⁷ Electrical and Computer Engineering Department, University of Southern California, Los Angeles, CA, United States of America ¹⁸ Institute of Neuroinformatics, University of Zurich and ETH Zurich, Switzerland ¹⁹ Hewlett Packard Laboratories, Hewlett Packard Enterprise, San Jose, CA, United States of America ²⁰ Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN, United States of America ²¹ CEA, LETI, Université Grenoble Alpes, Grenoble, France

- ²² CEA, LIST, Université Grenoble Alpes, Grenoble, France
- ²³ Department of Materials, University of Oxford, Parks Road, OX1 3PH Oxford, United Kingdom
- 24 Institute of Physics, University of Münster, Heisenbergstr. 11, 48149 Münster, Heisenbergstr., Germany
- ²⁵ Center for Soft Nanoscience, University of Münster, 48149 Münster, Germany
- ²⁶ The University of Manchester, United Kingdom
- ²⁷ Department of Cognitive Sciences, University of California, Irvine, Irvine, CA, United States of America
- 28 Institute of Theoretical Computer Science, Graz University of Technology, Graz, Austria
- ²⁹ École Polytechnique Fédérale de Lausanne, Geneva, Switzerland
- ³⁰ School of Electrical and Data Engineering, University of Technology, Sydney, Australia
- ³¹ Department of Electrical Engineering, New Haven, Yale University, United States of America
- ³² International Research Center for Neurointelligence (IRCN), The University of Tokyo, 7-3-1 Hongo Bunkyo-ku, Tokyo 113-0033, Iapan
- ³³ CerCo, Université Toulouse 3, CNRS, CHU Purpan, Pavillon Baudot, 31059 Toulouse, France
- ³⁴ Event Driven Perception for Robotics, Italian Institute of Technology, iCub Facility, Genoa, Italy
- ³⁵ Dept. of Psychology, Cornell University, Ithaca, NY, United States of America
- 36 Prophesee, Paris, France
- ³⁷ Enhanced Regenerative Medicine, Istituto Italiano di Tecnologia, Italy
- ³⁸ Department of Computer Science and Medical Technologies Innovation Facility, Nottingham Trent University, United Kingdom
- ³⁹ University of Maryland, Baltimore County, Catonsville, United States of America
- ⁴⁰ Technical University of Denmark, Denmark

- ⁴¹ The Danish Council on Ethics, Denmark
- ⁴² Department of Food and Resource Economics, University of Copenhagen, Denmark
- ⁴³ Dipartimento di Elettronica, Informazione e Bioingegneria Politecnico di Milano and IU.NET, 20133 Milano, Italy
- * Author to whom any correspondence should be addressed.

E-mail: nipr@dtu.dk

Keywords: neuromorphic computation, spiking neural networks, robotics, memristor, convolutional neural networks, self-driving cars, deep learning

Abstract

Modern computation based on von Neumann architecture is now a mature cutting-edge science. In the von Neumann architecture, processing and memory units are implemented as separate blocks interchanging data intensively and continuously. This data transfer is responsible for a large part of the power consumption. The next generation computer technology is expected to solve problems at the exascale with 10¹⁸ calculations each second. Even though these future computers will be incredibly powerful, if they are based on von Neumann type architectures, they will consume between 20 and 30 megawatts of power and will not have intrinsic physically built-in capabilities to learn or deal with complex data as our brain does. These needs can be addressed by neuromorphic computing systems which are inspired by the biological concepts of the human brain. This new generation of computers has the potential to be used for the storage and processing of large amounts of digital information with much lower power consumption than conventional processors. Among their potential future applications, an important niche is moving the control from data centers to edge devices. The aim of this roadmap is to present a snapshot of the present state of neuromorphic technology and provide an opinion on the challenges and opportunities that the future holds in the major areas of neuromorphic technology, namely materials, devices, neuromorphic circuits, neuromorphic algorithms, applications, and ethics. The roadmap is a collection of perspectives where leading researchers in the neuromorphic community provide their own view about the current state and the future challenges for each research area. We hope that this roadmap will be a useful resource by providing a concise yet comprehensive introduction to readers outside this field, for those who are just entering the field, as well as providing future perspectives for those who are well established in the neuromorphic computing community.

Contents

Introduction	
Section 1. Materials and devices	6
1. Phase-change memory devices	6
2. Ferroelectric devices	9
3. Valence change memory	11
4. Electrochemical metallization cells	16
5. Nanowire networks	19
6. 2D materials	22
7. Organic materials	25
Section 2. Neuromorphic circuits	29
8. Spintronics	29
9. Deep learning	32
10. Spiking neural networks	
11. Emerging hardware approaches for optimization	37
12. Enabling technologies for the future heterogeneous neuromorphic accelerators	41
13. Photonics	44
14. Large-scale neuromorphic computing platforms	48
Section 3. Neuromorphic algorithms	50
15. Learning in spiking neural networks	
16. Learning-to-learn for neuromorphic hardware	
17. Computational neuroscience	
18. Stochastic computing	
19. Convolutional spiking neural networks	

20. Reservoir computing	63
21. Computing with spikes	66
Section 4. Applications	69
22. Robotics	69
23. Self-driving cars	
24. Olfaction and chemosensation	
25. Event vision sensors	
26. Neuromorphic audition	
27. Biohybrid systems for brain repair	
28. Embedded devices for neuromorphic time-series assessment	
29. Electromyography processing using wearable neuromorphic technologies	
30. Collaborative autonomous systems	
Section 5. Ethics	
31. The ethics of developing neuromorphic technology	
Data availability statement	
References	

Introduction

N Pryds¹, Dennis V Christensen¹, Bernabe Linares-Barranco², Daniele Ielmini³ and Regina Dittmann⁴

- ¹Technical University of Denmark
- ²Instituto de Microelectrónica de Sevilla, CSIC and University of Seville
- ³Politecnico di Milano and IU.NET

⁴Forschungszentrum Jülich GmbH

Computers have become essential to all aspects of modern life and are omnipresent all over the globe. Today, data-intensive applications have placed a high demand on hardware performance, in terms of short access latency, high capacity, large bandwidth, low cost, and ability to execute artificial intelligence (AI) tasks. However, the ever-growing pressure for big data creates additional challenges: on the one hand, *energy consumption* has become a remarkable challenge, due to the rapid development of sophisticated algorithms and architectures. Currently, about 5%–15% of the world's energy is spent in some form of data manipulation, such as transmission or processing [1], and this fraction is expected to rapidly increase due to the exponential increase of data generated by ubiquitous sensors in the era of internet of things. On the other hand, data processing is increasingly limited by the *memory bandwidth* due to the von-Neumann's architecture with physical separation between processing and memory units. While the von Neumann computer architecture has made an incredible contribution to the world of science and technology for decades, its performance is largely inefficient due to the relatively slow and energy demanding data movement.

Conventional von Neumann computers based on complementary metal oxide semiconductor (CMOS) technology do not possess the intrinsic capabilities to learn or deal with complex data as the human brain does. To address the limits of digital computers, there are significant research efforts worldwide in developing profoundly different approaches inspired by biological principles. One of these approaches is the development of *neuromorphic systems*, namely computing systems mimicking the type of information processing in the human brain.

The term 'neuromorphic' was originally coined in the 1990s by Carver Mead to refer to mixed signal analog/digital very large scale integration computing systems that take inspiration from the neuro-biological architectures of the brain [2]. 'Neuromorphic engineering' emerged as an interdisciplinary research field that focused on building electronic neural processing systems to directly 'emulate' the bio-physics of real neurons and synapses [3]. More recently, the definition of the term neuromorphic has been extended in two additional directions [4]. Firstly, the term neuromorphic was used to describe spike-based processing systems engineered to explore large-scale computational neuroscience models. Secondly, neuromorphic computing comprises dedicated electronic neural architectures that implement neuron and synapse circuits. Note that this concept is distinct from AI machine learning approaches which are based on pure software algorithms developed to minimize the recognition error in pattern recognition tasks [5]. However, a precise definition of neuromorphic computing is somewhat debated. It can range from very strict high-fidelity mimicking of neuroscience principles where very detailed synaptic chemical dynamics are mandatory, to very vague high-level loosely brain-inspired principles, such as the simple vector (input) times matrix (synapses) multiplication. In general, as of today, there is a wide consensus that neuromorphic computing should at least encompass some

time-, event-, or data-driven computation. In this sense, systems like spiking neural networks (SNN), sometimes referred to as the third generation of neural networks [6], are strongly representative. However, there is an important cross-fertilization between the technologies required to develop efficient SNNs and those for more traditional non-SNN, referred to as artificial neural networks (ANN), which are typically more timestep-driven. While the former definition of neuromorphic computing is more plausible, in this roadmap we aim at broadening the scope to emphasize the cross-fertilization between ANN and SNN.

Nature is a vital inspiration for the advancement to a more sustainable computing scenario, where neuromorphic systems display much lower power consumption than conventional processors, due to the integration of non-volatile memory and analog/digital processing circuits as well as the dynamic learning capabilities in the context of complex data. Building ANNs that mimic a biological counterpart is one of the remaining challenges in computing. If the fundamental technical issues are solved in the next few years, the neuromorphic computing market is projected to rise from \$0.2 billion in 2025 to \$22 billion in 2035 [7] as neuromorphic computers with ultra-low power consumption and high speed advance and drive demands for neuromorphic devices.

In line with these increasingly pressing issues, the general aim of the *roadmap on neuromorphic computing and engineering* is to provide an overview of the different fields of research and development that contribute to the advancement of the field, to assess the potential applications of neuromorphic technology in cutting edge technologies and to highlight the necessary advances required to reach these. The roadmap addresses:

- Neuromorphic materials and devices
- Neuromorphic circuits
- Neuromorphic algorithms
- Applications
- Ethics

Neuromorphic materials and devices: To advance the field of neuromorphic computing and engineering, the exploration of novel materials and devices will be of key relevance in order to improve the power efficiency and scalability of state-of-the-art CMOS solutions in a disruptive manner [4, 8]. Memristive devices, which can change their conductance in response to electrical pulses [9–11], are promising candidates to act as energy- and space-efficient hardware representation for synapses and neurons in neuromorphic circuits. Memristive devices have originally been proposed as binary non-volatile random-access memory and research in this field has been mainly driven by the search for higher performance in solid-state drive technologies (e.g., flash replacement) or storage class memory [12]. However, thanks to their analog tunability and complex switching dynamics, memristive devices also enable novel computing functions such as analog computing or the realisation of brain-inspired learning rules. A large variety of different physical phenomena has been reported to exhibit memristive behaviour, including electronic effects, ionic effects as well as structural or ferroic ordering effects. The material classes range from magnetic alloys, metal oxides, chalcogenides to 2D van de Waals materials or organic materials. Within this roadmap, we cover a broad range of materials and phenomena with different maturity levels with respect to their use in neuromorphic circuits. We consider emerging memory devices that are already commercially available as binary non-volatile memory such as phase-change memory (PCM), magnetic random-access memory, ferroelectric memory as well as redox-based resistive random-access memory and review their prospects for neuromorphic computing and engineering. We complement it with nanowire networks, 2D materials, and organic materials that are less mature but may offer extended functionalities and new opportunities for flexible electronics or 3D integration.

Neuromorphic circuits: Neuromorphic devices can be integrated with conventional CMOS transistors to develop fully functional neuromorphic circuits. A key element in neuromorphic circuits is their non-von Neumann architecture, for instance consisting of multiple cores each implementing distributed computing and memory. Both SSNs, adopting spikes to represent, exchange and compute data in analogy to action potentials in the brain, as well as circuits that are only loosely inspired by the brain, such as ANNs, are generally included in the roster of neuromorphic circuits, thus will be covered in this roadmap. Regardless of the specific learning and processing algorithm, a key processing element in neuromorphic circuits is the neural network, including several synapses and neurons. Given the central role of the neural network, a significant research effort is currently aimed at technological solutions to realize dense, fast, and energy-efficient neural networks by inmemory computing [13]. For instance, a memory array can accelerate the matrix-vector multiplication (MVM) [14]. This is a common feature of many neuromorphic circuits, including spiking and non-spiking networks, and takes advantage of Ohm's and Kirchhoff's laws to implement multiplication and summation in the network. The MVM crosspoint circuit allows for the straightforward hardware implementation of synaptic layers with high density, high real-time processing speed, and high energy efficiency, although the accuracy is challenged by stochastic variations in memristive devices in particular, and analog computing in general. An additional circuit challenge is the mixed analog-digital computation, which results in the need for large and energyhungry

analog-digital converter circuits at the interface between the analog crosspoint array and the digital system. Finally, neuromorphic circuits seem to take the most benefit from hybrid integration, combining front-end CMOS technology with novel memory devices that can implement MVM and neuro-biological functions, such as spike integration, short-term memory, and synaptic plasticity [15]. Hybrid integration may also need to extend, in the long term, to alternative nanotechnology concepts, such as bottom-up nanowire networks [16], and alternative computing concepts, such as photonic [17] and even quantum computing [18], within a single system or even a single chip with 3D integration. In this scenario, a roadmap for the development and assessment of each of these individual innovative concepts is essential.

Neuromorphic algorithms: A fundamental challenge in neuromorphic engineering for real application systems is to train them directly in the spiking domain in order to be more energy-efficient, more precise, and also be able to continuously learn and update the knowledge on the portable devices themselves without relying on heavy cloud computing servers. Spiking data tend to be sparse with some stochasticity and embedded noise, interacting with non-ideal non-linear synapses and neurons. Biology knows how to use all this to its advantage to efficiently acquire knowledge from the surrounding environment. In this sense, computational neuroscience can be a key ingredient to inspire neuromorphic engineering, and learn from this discipline how brains perform computations at a variety of scales, from small neurons ensembles, mesoscale aggregations, up to full tissues, brain regions and the complete brain interacting with peripheral sensors and motor actuators. On the other hand, fundamental questions arise on how information is encoded in the brain using nervous spikes. Obviously, to maximize energy efficiency for both processing and communication, the brain maximizes information per unit spike [19]. This means unravelling the information encoding and processing by exploiting spatio-temporal signal processing to maximize information while minimizing energy, speed, and resources.

Applications: The realm of applications for neuromorphic computing and engineering continues to grow at a steady rate, although remaining within the boundaries of research and development. While it is becoming clear that many applications are well suited to neuromorphic computing and engineering, it is also important to identify new potential applications to further understand how neuromorphic materials and hardware can address them. The roadmap includes some of these emerging applications as examples of biologically-inspired computing approaches for implementation in robots, autonomous transport capability or in perception engineering where the applications are based on integration with sensory modalities of humans.

Ethics: While the future development and application of neuromorphic systems offer possibilities beyond the state of the art, the progress should also be addressed from an ethical point of view where, e.g., lack of transparency in complex neuromorphic systems and autonomous decision making can be a concern. The roadmap thus ends with a final section addressing some of the key ethical questions that may arise in the wake of advancements in neuromorphic computation.

We hope that this roadmap represents an overview and updated picture of the current state-of-the-art as well as being the future projection in these exciting research areas. Each contribution, written by leading researchers in their topic, provides the current state of the field, the open challenges, and a future perspective. This should guide the expected transition towards efficient neuromorphic computations and highlight the opportunities for societal impact in multiple fields.

Acknowledgements

DVC and NP acknowledge the funding from Novo Nordic Foundation Challenge Program for the BioMag project (Grant No. NNF21OC0066526), Villum Fonden, for the NEED project (00027993), Danish Council for Independent Research Technology and Production Sciences for the DFF Research Project 3 (Grant No. 00069B), the European Union's Horizon 2020, Future and Emerging Technologies (FET) programme (Grant No. 801267) and Danish Council for Independent Research Technology and Production Sciences for the DFF-Research Project 2 (Grant No. 48293). RD acknowledges funding from the German Science foundation within the SFB 917 'Nanoswitches', by the Helmholtz Association Initiative and Networking Fund under Project Number SO-092 (Advanced Computing Architectures, ACA), the Federal Ministry of Education and Research (project NEUROTEC Grant No. 16ES1133K) and the Marie Sklodowska-Curie H2020 European Training Network, 'Materials for neuromorphic circuits' (MANIC), grant Agreement No. 861153. BLB acknowledges funding from the European Union's Horizon 2020 (Grants 824164, 871371, 871501, and 899559). DI acknowledges funding from the European Union's Horizon 2020 (Grants 824164, 899559 and 101007321).

Section 1. Materials and devices

1. Phase-change memory devices

Abu Sebastian¹, Manuel Le Gallo¹ and Andrea Redaelli²

¹IBM Research - Zurich, Switzerland ²ST Microelectronics, Italy

1.1. Status

PCM exploits the behaviour of certain phase-change materials, typically compounds of Ge, Sb and Te, that can be switched reversibly between amorphous and crystalline phases of different electrical resistivity [20]. A PCM device consists of a certain nanometric volume of such phase change material sandwiched between two electrodes (figure 1).

In recent years, PCM devices are being explored for brain-inspired or neuromorphic computing mostly by exploiting the physical attributes of these devices to perform certain associated computational primitives inplace in the memory itself [13, 21]. One of the key properties of PCM that enables such immemory computing (IMC) is simply the ability to store two levels of resistance/conductance values in a non-volatile manner and to reversibly switch from one level to the other (binary storage capability). This property facilitates in-memory logical operations enabled through the interaction between the voltage and resistance state variables [21]. Applications of in-memory logic include database query [22] and hyper-dimensional computing [23].

Another key property of PCM that enables IMC is its ability to achieve not just two levels but a continuum of resistance values (analogue storage capability) [20]. This is typically achieved by creating intermediate phase configurations through the application of partial RESET pulses. The analogue storage capability facilitates the realization of MVM operations in O(1) time complexity by exploiting Kirchhoff's circuit laws. The most prominent application for this is DNN inference [24]. It is possible to map each synaptic layer of a DNN to a crossbar array of PCM devices. There is a widening industrial interest in this application owing to the promise of significantly improved latency and energy consumption with respect to existing solutions. This in-memory MVM operations also enable non-neuromorphic applications such as linear-solvers and compressed sensing recovery [21].

The third key property that enables IMC is the accumulative property arising from the crystallization kinetics. This property can be utilized to implement DNN training [25, 26]. It is also the central property that is exploited for realizing local learning rules like spike-timing-dependent plasticity in SNN [27, 28]. In both cases, the accumulative property is exploited to implement the synaptic weight update in an efficient manner. It has also been exploited to emulate neuronal dynamics [29].

Note that, PCM is at a very high maturity level of development with products already on the market and a well-established roadmap for scaling. This fact, together with the ease of embedding PCM on logic platforms (embedded PCM) [30] make this technology of unique interest for neuromorphic computing and IMC in general.

1.2. Current and future challenges

PCM devices have several attractive properties such as the ability to operate them at timescales on the order of tens of nanoseconds. The cycling endurance is orders of magnitude higher for PCM compared to other non-volatile memory devices such as flash memory. The retention time can also be tuned relatively easily with the appropriate choice of materials, although the retention time associated with the intermediate phase configurations could be substantially lower than that of the full amorphous state.

However, there are also several device-level challenges as shown in figure 2. One of the key challenges associated with the use of PCM for in-memory logic operations is the wide distribution of the SET states. These distributions could detrimentally impact the evaluation of logical operations. The central challenge associated with in-memory MVM operations is the limited precision arising from the 1/*f* noise as well as conductance drift. Drift is attributed to the structural relaxation of the melt-quenched amorphous phase [31]. Temperature-induced conductance variations could also pose challenges. One additional challenge is related to the stoichiometric stability during cycling where ion migration effects can occur [32]. Moreover, the accumulative behaviour in PCM is highly nonlinear and stochastic. While one could exploit this intrinsic stochasticity to realize stochastically firing neurons and for stochastic computing, this behaviour is detrimental for applications such as DNN training in which the conductance must be precisely modulated.

PCM-based IMC has the potential for ultra-high compute density since PCM devices can be scaled to nanoscale dimensions. However, it is not straightforward to fabricate such devices in a large array due to fabrication challenges such as etch damage and deposition of materials in high-aspect ratio pores [33]. The





Figure 1. Key physical attributes that enable neuromorphic computing. (a) Non-volatile binary storage facilitates in-memory logical operations relevant for applications such as hyper-dimensional computing. (b) Analog storage enables efficient matrix-vector multiply (MVM) operations that are key to applications such as deep neural network (DNN) inference. (c) The accumulative behaviour facilitates applications such as DNN training and emulation of neuronal and synaptic dynamics in SNN.



is detrimental for applications such as in-memory logic. (b) The drift and noise associated with analogue conductance values results in imprecise matrixvector multiply operations. (c) The nonlinear and stochastic accumulative behaviour result in imprecise synaptic weight updates.

integration density is also limited by the access device, which could be a selector in the backend-of-the-line (BEOL) or front-end bipolar junction transistors (BJT) or metal-oxide-semiconductor field effect transistors (MOSFET). The threshold voltage must be overcome when SET operations are performed, so the access device must be able to manage voltages at least as high as the threshold voltage. While MOSFET selector size is mainly determined by the PCM RESET current, the BJT and BEOL selectors can guarantee a minimum cell size of 4F², leading to very high density [34]. However, BEOL selector-based arrays have some drawbacks in terms of precise current control, while the management of parasitic drops is more complex for BJT-based arrays [35].

1.3. Advances in science and technology to meet challenges

A promising solution towards addressing the PCM nonidealities such as 1/f noise and drift is that of projected phase-change memory (projected PCM) [36, 37]. In these devices, there is a non-insulating projection segment in parallel to the phase-change material segment. By exploiting the highly nonlinear I-V characteristics of phase-change materials, one could ensure that during the SET/RESET process, the projection segment has minor impact on the operation of the device. An increase in the reset current is anyway expected and some work should be done on material engineering side to compensate for that. However, during read, the device

conductance is mostly determined by the projection segment that appears parallel to the amorphous phasechange segment. Recently, it was shown that it is possible to achieve remarkably high precision in-memory scalar multiplication (equivalent to 8 bit fixed point arithmetic) using projected PCM devices [38]. These projected PCM devices also facilitate array-level temperature compensation schemes. Alternate multi-layered PCM devices have also been proposed that exhibit substantially lower drift [39].

There is a perennial focus on trying to reduce the RESET current via scaling the switchable volume of the PCM device. Either by shrinking the overall dimension of the device in a confined geometry or by scaling the bottom electrode dimensions of a mushroom-type device. The exploration of new material classes such as single elemental antimony could help with the scaling challenge [40].

The limited endurance and various other non-idealities associated with the accumulative behaviour such as limited dynamic range, nonlinearity and stochasticity can be partially circumvented with multiPCM synaptic architectures. Recently, a multi-PCM synaptic architecture was proposed that employs an efficient counterbased arbitration scheme [41]. However, to improve the accumulation behaviour at the device level, more research is required on the effect of device geometries as well as the randomness associated with crystal growth.

Besides conventional electrical PCM devices, photonic memory devices based on phase-change materials, which can be written, erased, and accessed optically, are rapidly bridging a gap towards allphotonic chip-scale information processing. By integrating phase-change materials onto an integrated photonics chip, the analogue multiplication of an incoming optical signal by a scalar value encoded in the state of the phase change material was achieved [42]. It was also shown that by exploiting wavelength division multiplexing, it is possible to perform convolution operations in a single time step [43]. This creates opportunities to design phase-change materials that undergo faster phase transitions and have a higher optical contrast between the crystalline and amorphous phases [44].

1.4. Concluding remarks

The non-volatile binary storage, analogue storage and accumulative behaviour associated with PCM devices can be exploited to perform in-memory computing (IMC). Compared to other non-volatile memory technologies, the key advantages of PCM are the well understood device physics, volumetric switching and easy embeddability in a CMOS platform. However, there are several device and fabrication-level challenges that need be overcome to enable PCM-based IMC and this is an active area of research.

It will also be rather interesting to see how PCM-based neuromorphic computing will eventually be commercialized. Prior to true IMC, a hybrid architecture where PCM memory chips are used to store synaptic weights in a non-volatile manner while the computing is performed in a stacked logic chip is likely to be considered as an option by the industry. Despite the tight interconnect between the stacked chips, data transfer will remain a bottleneck for this approach. A better solution could be PCM directly embedded with the logic itself BEOL without any interconnect bottleneck and eventually we could foresee full-fledged non-von Neumann accelerator chips where the embedded PCM is also used for analogue IMC.

Acknowledgements

This work was supported in part by the European Research Council through the European Union's Horizon 2020 Research and Innovation Programme under Grant No. 682675.

2. Ferroelectric devices

Stefan Slesazeck¹ and Thomas Mikolajick^{1,2}

¹NaMLab gGmbH

²Institute of Semiconductors and Microsystems, Germany

2.1. Status

Ferroelectricity was firstly discovered in 1920 by Valasek in Rochelle salt [45] and describes the ability of a non-centrosymmetric crystalline material to exhibit a permanent and switchable electrical polarization due to the formation of stable electric dipoles. Historically, the term ferroelectricity stems from the analogous behavior with the magnetization hysteresis of ferromagnets when plotting the ferroelectric polarization versus the electrical field. Regions of opposing polarization are called domains. The polarization direction of such domains can be switched typically by 180° but, based on the crystal structure, other angles are also possible. Since the discovery of the stable ferroelectric barium titanate in 1943, ferroelectrics have found application in capacitors in electronics industry. Already, in the 1950s, ferroelectric capacitor (FeCAP) based memories (FeRAM) have been proposed [46], where the information is stored as polarization state of the ferroelectric material. Read and write operation are performed by applying an electric field larger than the coercive field E_C . The destructive read operation after readout. Thanks to the development of mature processing techniques for ferroelectric lead zirconium tantalate FeRAMs, these have been commercially available since the early 1990s [47]. However, the need for a sufficiently large capacitor together with the limited thin-film manufacturability of perovskite materials has so far restricted their use to niche applications [48].

The ferroelectric field effect transistors (FeFET) that was proposed in 1957 [49] features a FeCAP as gate insulator, modulating the transistor's threshold voltage that can be sensed non-destructively by measuring the drain-source current. Perovskite based FeFET memory arrays with up to 64 kBit have been demonstrated [50]. However, due to difficulties in the technological implementation, limited scalability and data retention issues, no commercial devices became available.

The ferroelectric tunneling junction (FTJ) was proposed by Esaki *et al* in 1970 s as a 'polar switch' [51] and was firstly demonstrated in 2009 using a BaTiO₃ ferroelectric layer [52]. The FTJ features a ferroelectric layer sandwiched between two electrodes, thus modifying the tunneling electro-resistance. A polarization-dependent current is measured non-destructively when applying electrical fields smaller than $E_{\rm C}$.

Since the fortuitous discovery of ferroelectricity in hafnium oxide (HfO₂) in 2008 and its first publication in 2011 [53] the well-established and CMOS-compatible fluorite-structure material has been extensively studied and has recently gained a lot of interest in the field of nonvolatile memories and beyond von-Neumann computing [54, 55] (figure 3).


2.2. Current and future challenges

Very encouraging electrical results of fully front-end-of-line integrated FeFET devices featuring switching speeds <50 ns at <5 V pulse voltage have been reported recently based on >1 Mbit memory arrays [56]. The ability of fine-grained co-integration of FeFET memory devices together with CMOS logic transistors paves the way for the realization of braininspired architectures to overcome the limitations of the von-Neumann bottleneck, which restricts the data transfer due to limited memory and data bus bandwidth [57]. However, one of the main challenges for the FeFET devices and therefore a topic of intense research is the formation of ferroelectric HfO₂-based thin films featuring a uniform polarization behavior at nano-scale as an important prerequisite for the realization of small scaled devices with feature sizes <100 nm.

Another important challenge for many application cases is the limited cycling endurance of silicon-based FeFETs that is typically in the range of 10^5 cycles. This value is mainly dictated by the breakdown of the dielectric SiO₂ interfacial layer that forms between the Si channel and the ferroelectric gate insulator.

FeCAPs have been successfully integrated into the back-end-of-line (BEOL) of modern CMOS technologies and operation of a HfO₂-based FeRAM memory array at 2.5 V and 14 ns switching pulses was successfully demonstrated [58]. At this point the main challenge is the decrease of the ferroelectric layer thickness well below 10 nm to allow scaling of 3D capacitors towards the 10 nm node. Moreover, phenomenon such as the so called 'wake-up effect' with increasing of P_r for low cycle counts as well as the 'fatigue effect' resulting in a reduction of P_r at high cycle counts due to oxygen vacancy redistribution [59] and defect generation have to be tackled. That is especially important for fine-grained circuit implementations where the switching properties of single ferroelectric devices impact the designed operation point of analogue circuits.

One of the most interesting benefits of FTJ devices is the small current density making them very attractive for applications requiring massive parallel operations such as analogue matrixvector-multiplications in




larger cross-bar structures [60]. However, increasing the ratio between the on-current density and the selfcapacitance of FTJ devices turns out to be one of the main challenges to increasing the reading speed for these devices. The tunneling current densities depend strongly on the thickness of the ferroelectric layer and the composition of the multi-layer stacks. The formation of very thin ferroelectric layers is hindered by unintentional formation of interfacial dead layers towards the electrodes and increasing leakage currents due to defects and grain-boundaries in the poly-crystalline thin films.

2.3. Advances in science and technology to meet challenges

Although ferroelectricity in hafnium oxide has been extensively studied for over one decade now, there are still many open questions in understanding the formation of the ferroelectric *Pca2*₁ phase and regarding the interaction with material layers such as electrodes, dielectric tunneling barriers in multi-layer FTJs or interfacial layers in FeFETs. Moreover, the interplay between charge trapping phenomenon and ferroelectric switching mechanisms [61], the trade-off between switching speed and voltage of the nucleation limited switching and its impact on device reliability or the different behavior of abrupt single domain switching [55] and smooth polarization transitions in negative capacitance devices that were observed in the very similar material stacks are still not completely understood. However, that knowledge will be an important ingredient for proper optimization of material stacks as well as electrical device operation conditions.

On the materials side, the stabilization of the ferroelectric orthorhombic $Pca2_1$ phase in crystallized HfO₂ thin films has to be optimized further. Adding dopants, changing oxygen vacancy densities or inducing stress by suitable material stack and electrode engineering are typical measures. In most cases a poly-crystalline material layer is attained consisting of a mixture of different crystalline ferroelectric and non-ferroelectric phase



fractions. Moreover, ferroelectric grains that differ in size or orientation of the polarization axis, electronically active defects as well as grain size dependent surface energy effects give rise to the formation of ferroelectric domains that possess different electrical properties in terms of coercive field $E_{\rm C}$ (typical values $\sim 1 \text{ MV cm}^{-1}$) or remnant polarization $P_{\rm r}$ (typical values $10-40 \,\mu {\rm C cm}^{-2}$) with impact on the device-to-device variability and the gradual switching properties that are important especially for analog synaptic devices. Some drawbacks of the poly-crystallinity of ferroelectric HfO₂- and ZrO₂-based thin films could be tackled by the development of epitaxial growth of monocrystalline ferroelectric layers [62] where domains might extend over a larger area. The case of FTJs in particular demonstrate the effect of domain wall motion that might allow a more gradual and analogue switching behavior even in small scaled devices. The utilization of an anti-ferroelectric hysteretic switching that was demonstrated in ZrO₂ thin films bears the potential to overcome some limitations that are related to the high coercive field of ferroelectric HfO₂, such as operation voltages being larger than the typical core voltages in modern CMOS technologies or the limited cycling endurance [63].

Finally, in addition to the very encouraging results adopting ferroelectric HfO₂, in 2019 another promising material was realized. AlScN is a semiconductor processing compatible and already utilized piezoelectric material that was made ferroelectric [64] (figure 4).

2.4. Concluding remarks

The discovery of ferroelectricity in hafnium oxide has led to a resumption in the research on ferroelectric memory devices, since hafnium oxide is a well-established and fully CMOS compatible material in both front end of line and back end of line processing. Besides the expected prospective realization of densely integrated non-volatile and ultra-low-power ferroelectric memories in near future, this development directly leads to the adoption of the trinity of ferroelectric memory devices—FeCAP, FeFET and FTJ—for beyond von Neumann computing. While in the memory application the important topic of reliability on the array level is yet to be solved, for neuromorphic applications the linear switching to many different states, especially in scaled down devices, is a topic that needs further attention. Moreover, very specific properties of the different ferroelectric device types demand for the development of new circuit architectures that facilitate a proper device operation taking into account the existing non-idealities. A thorough design technology co-optimization will be the key to fully exploit their potential in neuromorphic and edge computing. Finally, large scale demonstrations of ferroelectrics based neuromorphic circuits need to be investigated to identify all possible issues.

Acknowledgements

This work was financially supported out of the State budget approved by the delegates of the Saxon State Parliament.

3. Valence change memory

Sabina Spiga¹ and Stephan Menzel²

¹CNR-IMM, Italy ²FZ Juelich (PGI-7), Germany

3.1. Status


Resistive random access memories (RRAMs), also named memristive devices, change their resistance state upon electrical stimuli. They can store and compute information at the same time, thus enabling in-memory and brain-inspired computing [13, 65]. RRAM devices relying on oxygen ion migration effects and subsequent valence changes are named valence change memory (VCM) [66]. They have been proposed to implement electronic synapses in hardware neural networks, due to the ability to adapt their strength (conductance) in an analogue fashion as a function of incoming electrical pulses (synaptic plasticity), leading to long-term (short-term) potentiation and depression. In addition, learning rules such as spike-time or spike-rate dependent plasticity, paired-pulse facilitation or the voltage threshold—based plasticity have been demonstrated; the stochasticity of the switching process has been exploited for stochastic update rules [67-69]. Most of the VCM devices are based on a two-terminal configuration, and the switching geometry involves either confined filamentary, or interfacial regions (figure 5(A)). Filamentary VCMs are today the most advanced in terms of integration and scaling. Their switching mechanism relies on the creation and rupture of conductive filaments (CF), formed by a localized concentration of defects, shorting the two electrodes. The modulation/control of the CF diameter and/or CF dissolution can lead to two or multiple stable resistance states [70, 71]. Prototypes of neuromorphic chips have been recently shown, integrating HfO_x and TaO_x -based filamentary-VCM as synaptic nodes in combination with CMOS neurons [72-74]. In interfacial VCM devices, the conductance scales with the junction area of the device, and the mechanism is related to a homogenous oxygen ion movement through the oxides, either at the electrode/oxide or oxide/oxide interface. Reference material systems are based on complex oxides, such as bismuth ferrite [75] and praseodymium calcium manganite [76]; or bilayers stacks, e.g. TiO₂/TaO₂ [77] and a-Si/TiO₂ [78]. Finally, three-terminal VCM redox transistors have been recently studied (figure 5(A) right), where the switching mechanism is related to the control of the oxygen vacancy concentration in the bulk of the transistor channel [79, 80]. While interfacial and redox-transistor devices are today at low technological readiness, and most of the studies are reported at single device level, they promise future advancement in neuromorphic computing in terms of analogue control, higher resistance values, improved reliability, reduced stochasticity with respect to filamentary devices [81]. To design neuromorphic circuits including VCM devices, compact models are requested. For filamentary devices compact models including variability are available [81, 82], but lacking for interfacial VCM and redox-based transistors.

3.2. Current and future challenges

VCM devices have been developed in the last 15 years mainly for storage applications, but for neuromorphic applications the required properties differ. In general, desirable properties of memories for neural networks include (i) analogue behaviour or controllable multilevel states, (ii) compatibility with learning rules supporting also online learning, (iii) tuneable short-term and long-term stability of the weights to implement various dynamics and timescales in synaptic and neuronal circuits [67–69]. A significant debate still refers to the linear/non-linear and symmetric/asymmetric conductance update of experimental devices, synaptic resolution (number of resistance levels), and how to exploit or mitigate these features (figures 5(B) and (C)).

Filamentary devices are the most mature type of VCMs. Nevertheless, many issues are pending: e.g. control of multi-level operation, device variability, intrinsic stochasticity, program and read disturbs, and the still too low resistance level range for neuromorphic circuits [83]. Moreover, the understanding/modelling of their switching mechanism is still under debate. Whereas first models including switching variability and read noise are available [81, 82], retention modelling, and the modelling of volatile effects and device failures are current challenges. First hybrid CMOS-VCM chips have been developed demonstrating inference application, but so far they do not support on-chip learning [72–74].

Interfacial VCM devices show in general less variability, less (no) read instability and a very analogue tuning of the conductance states, which can leads to a more deterministic and linear conductance update compared to filamentary devices [76]. Still these properties are not characterized on a high statistical basis. The retention, especially for thin oxide devices, is lower than for filamentary devices, which may be still compatible with some applications. As the conductance scales with area, the achievable high resistance levels promise a low power operation. Typical devices, however, have a large area or thick switching oxides, and scaling them to the nanoscale is an open issue. Moreover, devices showing a large resistance modulation require high switching voltages, not easily compatible with scaled CMOS nodes. The fabrication and characterization of interfacial



strategies or materials engineering.

VCM arrays needs to be further addressed. Simulation models for interfacial VCM are not available yet and need to be developed.

Redox-based VCM transistors have been only shown on a single device level [79, 80]. Thus, reliable statistical data on cycle-to-cycle variability, device-to-device variability and stability of the programmed states is not available yet. Moreover, the trade-off between switching speed and voltage has not been studied in detail. Another challenge is the understanding of the switching mechanism and the development of suitable models for circuit design.

The open challenges for all three types of VCM devices are summarized in table 1.

	Filamentary	Interfacial	Redox transistor
Single device	Various materials systems and scaling down to mm scale, low V operation	Various materials, scaling at nm scales, still overall high V and thick oxide	Recently studied, promising but still at initial development stage
Array	Demonstrated as 1T-1R and 1S-1R. Total capacity up to 1–2 Mb in recent works	Published in few works, up to few Kb, mainly 1R array only	To be addressed
Monolithic CMOS integration	Demonstrated, down to the 2× techn. nodes, and in 3D architectures	Partial demonstration, mainly for oxide bilayer and not for perovskites	To be addressed
Compact models	Available, retention and variability modelling to be optimized	Almost not availalbe	Almost not available
Binary	Available, with good endurance $(>10^6-10^9)$ and retention $(>years)$	Possible, but lack of statistical data. Endurance and long retention to be optimized	Very new devices, mostly proposed for multilevel applications, lack of statistical data
Analog multilevel memory	Hard to control, multilevel demonstrated in array using program/verify algorithms	Promising, less variability, high <i>R</i> , but reduced dynamic range and mainly data for single devices	Promising, high <i>R</i> value, deterministic control of <i>G</i> update, but shown for single devices
Stochastic memory	Proposed in some works, to be further validated	_	_
Long term memory (LTM)	Yes, retention at high T and for 6–10 years. Depends on <i>R</i> levels	Possible, lack of statistic data on array, single device retention up to years for some material stacks	Possible. Few studies. Lack of statistical data
Short term memory (STM)	Usually difficult to achieve controlled decay	Possible, to be further addressed and optimized	Possible, to be further addressed. Few studies

Table 1. Summary of status and open challenges of the three types of VCM devices.

3.3. Advances in science and technology to meet challenges

The current challenges for VCM-type devices push the research in various but connected directions, which span from material, to theory, devices and architecture. A better understanding of material properties and microscopic switching mechanisms is definitely required. However, the key step is to demonstrate the device integration in complex circuits and hybrid CMOS-VCM hardware neuromorphic chips. While VCMs are not ideal devices, many issues can be solved or mitigated at circuit level still taking advantage of their properties in term of power, density, and dynamic properties.

In this context, **filamentary** VCM devices are the most mature technology, but their deployment into neuromorphic computing hardware is still at its infancy. A comprehensive compact model, depicting complete dynamics including retention effects, e.g. to accurately simulate online learning, is required for the development of optimized circuits. On the material level, the biggest issues are read noise and switching variability. Due to the inherent Joule heating effect, the transition time of the conductance switching is very short and depends strongly on the device state [84]. This makes it hard to control the conductance update. Future research could explore very fast pulses in the range of the transition time to update the cell conductance, or use thermal engineering of the device stacks to increase the transition time. Finally, to achieve low power operation, resistance state values should be moved to the M Ω regime.

For interfacial and redox-transistor VCM devices, one of the next important steps is to shift from single device research to large arrays, possibly co-integrated with CMOS. This step enables to collect a large amount of data, which is required for modelling and demonstrating robust neuromorphic functions. It would be highly

desirable to identify a reference material system with a robust switching mechanism supported by a comprehensive understanding and modelling from underlying physics to compact and circuits modelling. Indeed, the modelling of these devices are still at its infancy. One open question for both devices is the trade-off between data retention and switching speed. In contrast to the filamentary devices, the velocity of the ions are probably not accelerated by Joule heating. Thus, the voltage needs to be increased more than in filamentary devices, to operate the devices at fast speed [85]. This might limit the application of these device to a certain time domain as the CMOS might not be able to provide the required voltage. By using thinner device layers or material engineering this issue could be addressed.

3.4. Concluding remarks

The VCM device technologies can integrate novel functionalities in hardware as key elements of the synaptic nodes in neural networks, i.e. to store the synaptic weight. Moreover, they can enable new learning algorithms that enable bio-plausible functions over multiple timescales. At the moment, it is still not clear which can be the best 'final' VCM material system and/or VCM device type, having each of them advantages and disadvantages. The missing 'killer' system, with consolidated properties/understanding/easy manufacturing, prevents to concentrate the efforts of the scientific community in single direction to bring VCM device to industrial real applications beyond a niche market. While filamentary VCMs are already been implemented in neuromorphic computing hardware, interfacial VCM or redox transistor can open new perspectives in the long term. To this end, there is an urgent request to further develop VCM devices enhancing new properties through a combined synergetic development based on materials design, physical and electrical characterizations and multiscale modelling to support the microscopic understanding of the link between the device structure and the electrical characteristics. Moreover, the device development targeting braininspired computing systems can only go hand-in-hand with theory and architectures design in a holistic view.

Acknowledgements

This work was partially supported by the Horizon 2020 European projects MeM-Scales (Grant No. 871371), MNEMOSENE (Grant No. 780215), and NEUROTECH (Grant No. 824103); in part by the Deutsche Forschungsgemeinschaft (SFB 917); in part by the Helmholtz Association Initiative and Networking Fund under Project Number SO-092 (Advanced Computing Architectures, ACA) and in part by the Federal Ministry of Education and Research (BMBF, Germany) in the project NEUROTEC (Project Numbers 16ES1134 and 16ES1133K).

4. Electrochemical metallization cells

Ilia Valov

Research Centre Juelich

4.1. Status


Electrochemical metallization memories were introduced in nanoelectronics with the intention to be used as memory, optical, programmable resistor/capacitor devices, sensors and for crossbar arrays and rudimentary neuromorphic circuits by Kozicki *et al* [86, 87] under the name programmable metallization cells. These types of devices are also termed conductive bridging random access memories or atomic switches [88]. The principles of operation of these two electrode devices using thin layers as ion transporting media are schematically shown in figure 6. As electrochemically active electrodes Ag, Cu, Fe or Ni are mostly used and for counter electrodes Pt, Ru, Pd, TiN or W are preferred. Electrochemical reactions at the electrodes and ionic transport within the device are triggered by internal [89] or applied voltage causing the formation of metallic filament (bridge) short-circuiting the electrodes and defining the low resistance state (LRS). Voltage of opposite polarity is used to dissolve the filament, returning the resistance to high ohmic state (HRS). LRS and HRS are used to define Boolean 1 and 0, respectively.


Apart from the prospective for a paradigm shift in computing and information technology offered by memrsitive devices in general [8], ECMs provide particular advantages compared to other redox-based resistive memories. They operate at low voltages (\sim 0.2 V to \sim 1 V) and currents (from nA to μ A range) allowing for low power consumption. A huge spectrum of materials can be used as solid electrolytes, ionic conductors, mixed conductors, semiconductors, macroscopic insulators and even high-k materials such as SiO₂, HfO₂, Ta_2O_5 etc, predominantly in amorphous but also in crystalline state [90]. The spectrum of these materials includes 1D and 2D materials but also different polymers, bioinspired/bio-compatible materials, proteins and other organic and composite materials [91, 92]. The metallic filament can vary in thickness and may either completely bridge the device, or be only partially dissolved providing multilevel to analog behaviour. Very thin filaments are extremely unstable and dissolve fast (down to 10^{10} s) [93]. The devices are stable against radiation/cosmic rays, high energy particles and electromagnetic waves and can operate over a large temperature range [94, 95]. Due to these properties, ECMs can be implemented in various environments, systems and technologies. The typical applications are as selector devices, volatile, non-volatile digital and analog memories, transparent and flexible devices, sensors, artificial neurons and synapses [96–98]. The devices can combine more functions and are thought of as basic units for the fields of autonomous systems, beyond von Neumann computing and AI. Further development in the field is essential to realise the full potential of this technology.


4.2. Current and future challenges

Despite the apparent simplicity and ease of operation, ECM cells are complex nanoscale systems, relying on redox reactions and ion transport at extreme conditions [100]. Despite low absolute voltages and currents, the devices are exposed to electric fields of up to 10^8 V cm⁻¹ and current densities of up to $\sim 10^{10}$ A cm⁻². There is no other example in the entire field of electrochemical applications even approaching these conditions. Small device volume, harsh and strongly non-equilibrium conditions make the understanding of fundamental processes and their control extremely challenging. The latter results in less precise (or missing) control over the functionalities and reliable operation. Indeed, maybe the most serious disadvantage of ECMs is the large variability in switching voltages, currents and resistive states. Additional problems are fluctuations and drift of the resistance states, as well their chemically and/or physically determined instabilities (figure 7).

Several notable issues should be taken in consideration: (1) missing unequivocal experimental value about what part of the applied current is carried by ions and by electrons. Whereas in macroscopic systems these numbers are constant, in nanoscale ECMs they may vary depending on the conditions and charge concentration. (2) The charge/ion concentration may vary with time. Due to the small volume, it is easy to enrich or deplete the film with mobile ions (acting as donors/acceptors) during the operation cycles, resulting in deviation of the switching voltages and currents and finally to failures. (3) Again, due to the small volume, even a low number of foreign atoms/ions (impurities) will cause considerable changes in the electronic properties. Impurities or dopants as well as the matrix significantly alter the characteristics due to effects on the switching layer [102, 103] or on the electrodes [104]. (4) Effects of protons and oxygen; both can be incorporated either during device preparation (e.g. lithography steps, or deposition technique e.g. ALD etc) or from the environment [105], even if a capping layer is used. Many devices even cannot operate without the presence of protons and many electrode materials such as Cu, Mo, W or TiN etc can be partially or even are fully oxidized by environmental factors. (5) Interfacial interactions are commonly occurring at the electrode/solid electrolyte interface. The thickness of these interfacial layers can sometimes even exceed the thickness of the switching







layer and inhibit or support reliable operation [101]. All these effects have their origin in the nanosize of the devices and highly non-equilibrium operating conditions.

4.3. Advances in science and technology to meet challenges

Addressing the challenges and issues that still limit the implementation of ECM devices in the praxis, should be considered on different levels. On a fundamental level, an in-depth understanding of the nanoscale processes and rate-limiting steps that determine the resistive switching mechanism is essential. To overcome the current limitations the theory should be further improved to account not only quantitatively but also qualitatively for the fundamental differences in thermodynamics and kinetics on the nanoscale compared to the macroscale. The scientific equipment needs to be improved to address the demand on sufficient mass and charge sensitivity as well as lateral and vertical resolution. Focus should be set on *in situ* and *in operando* techniques under real conditions enhanced by high time and imaging resolutions.

On a materials level, efforts should be made to understand and effectively use the relation between physical and chemical material properties, such as chemical composition, non-stoichiometry, purity, doping, density, thickness and mechanical properties and device performance and functionalities. A more narrow selection from the vast sea of ECM materials should be made on which systematic research should be performed. The final task to be achieved by a selective materials research approach is establishing a universal materials treasure map.

On a device/circuit/technology level, common problems such as the sneak path problem still need to be addressed. Limitation of interactions between devices and high-density integration (also within CMOS) needs to be further improved. The control during the deposition of layer materials should be adjusted to avoid layer intermixing, contaminations and incorporation of impurities. In many cases, deposition of thin films of non-oxidized elements or components with higher affinity to oxygen such as W, Mo, TiN or oxygen-free containing chalcogenides is possible only after special pre-care. The technological processes must be adapted and regularly controlled to ensure high quality and defined chemical composition. Additional efforts should also be made to integrate devices utilizing different functionalities and allowing for higher degree of complexity. The internal electromotive force should be further explored and utilized in respect to autonomous systems and applications in space technologies and medicine should also be further developed.

These issues are in fact highly interrelated and closely depend on each other. Most important on the current stage of development of ECM devices is to understand and control the relation between material properties, physical processes and device performance and functionalities. This knowledge will result in improved reliability of the devices and advanced technology.

4.4. Concluding remarks

ECM devices have been intensively developed in the last 20 years however, still not reaching their full potential. Opportunities for various applications in the fields of nanoelectronics, nanoionics, magnetics, optics, sensorics etc and the prospectsfor implementation as basic units in neuromorphic computing, big data processing, autonomous systems and AI are impeded by insufficient control of the nanoscale processes and incomplete knowledge on the relation between material properties, fundamental processes and devices characteristics and functionalities. To achieve these tasks, not only existing theory but also the scientific equipment and characterization techniques should be further improved allowing a direct insight in the complex nanoscale phenomena. Interacting and complementing fundamental and applied research is the key to address these issues in order to deploy the advantages and opportunities offered by the electrochemical metallization of cells into modern information and communications technologies.

5. Nanowire networks

Gianluca Milano¹ and Carlo Ricciardi²

¹INRiM (Istituto Nazionale di Ricerca Metrologica), Italy ²Politecnico di Torino, Italy

5.1. Status

The human brain is a complex network of about 10¹¹ neurons connected by 10¹⁴ synapses, anatomically organized over multiple scales of space, and functionally interacting over multiple scales of time [106]. Synaptic plasticity, i.e. the ability of synaptic connections to strengthen or weaken over time depending on external stimulation, is at the root of information processing and memory capabilities of neuronal circuits. As building blocks for the realization of artificial neurons and synapses, memristive devices organized in large crossbar arrays with a top-down approach have been recently proposed [107]. Despite the state-of-art of this rapidly growing technology demonstrated hardware implementation of supervised and unsupervised learning paradigms in ANN, the rigid top-down and grid-like architecture of crossbar arrays fails in emulating the topology, connectivity and adaptability of biological neural networks, where the principle of self-organization governs both structure and functions [106]. Inspired by biological systems (figure 8(a)), more biologically plausible nanoarchitectures based on self-organized memristive nanowire (NW) networks have been proposed [16, 108–112] (figures 8(b) and (c)). Here, the main goal is to focus on the emergent behaviour of the system arising from complexity rather than on learning schemes that require addressing of single ANN elements. Indeed, in this case main players are not individual nano objects but their interactions [113]. In this framework, the cross-talk in between individual devices, that represents an unwanted source of sneak currents in conventional crossbar architectures, here represents an essential component for the network emerging behaviour needed for the implementation of unconventional computing paradigms. NW networks can be fabricated by randomly dispersing NWs with a metallic core and an insulating shell layer on a substrate by a low-cost drop casting technique that does not require nanolithography or cleanroom facilities. The obtained NW network topology shows small-world architecture similarly to biological systems [114]. Both single NW junctions and single NWs show memristive behaviour due to the formation/rupture of a metallic filament across the insulating shell layer and to breakdown events followed by electromigration effects in the formed nanogap, respectively (figures 8(e) and (h)) [16]. Emerging network-wide memristive dynamics were observed to arise from the mutual electrochemical interaction in between NWs, where the information is encoded in 'winnertakes-all' conductivity pathways that depend on the spatial location and temporal sequence of stimulation [115–117]. By exploiting these dynamics, NW networks in multiterminal configuration can exhibit homosynaptic, heterosynaptic and structural plasticity with spatiotemporal processing of input signals [16]. Also, nanonetworks have been reported to exhibit fingerprints of self-organized criticality similarly to our brain [108, 118, 119], a feature that is considered responsible for optimization of information transfer and processing in biological circuits. Because of both topological structure and functionalities, NW networks are considered as very promising platforms for hardware realization of biologically plausible intelligent systems.

5.2. Current and future challenges

Current and future challenges for hardware implementation of neuromorphic computing in the bottom-up NW network will need integrated theoretical and experimental multidisciplinary approaches involving material physics, electronics engineering, neuroscience and network science (an overview of the roadmap is shown in figure 9). In NW networks, unconventional computing paradigms that emphasize the network as a whole rather than the role of single elements need to be developed. In this framework, great attention has recently been devoted to the reservoir computing (RC) paradigm where a complex network of nonlinear elements is exploited to map input signals into a higher dimensional feature space that is then analysed by means of a readout function. In this framework, nano-networks have been proposed [120] and experimentally exploited as 'physical' reservoirs for *in materia* implementation of the RC paradigm [121–123]. However, fundamental research is needed to address remaining challenges. The design and fabrication of multiterminal memristive NW networks able to process multiple spatio-temporal inputs with nonlinear dynamics, fading memory (short-term memory) and echo-state properties minimizing energy dissipation are needed. Importantly, these NW networks have to operate at low voltages and currents to be implemented with conventional electronics. These represent challenges from the material science point of view, since to achieve this goal NWs have to be optimized in terms of core-shell structures for tailoring ionic dynamics underlying resistive switching mechanism. Also, a fully-hardware RC system requires hardware implementation of the readout function for processing outputs of the NW network physical reservoir. Despite the neural network readout can be implemented by means of crossbar arrays of ReRAM devices to realize a fully-memristive architecture as demonstrated in





Figure 8. Bio-inspired memristive NW networks. (a) Biological neural networks where synaptic connections between neurons are represented by bright fluorescent boutons (image of primary mouse hippocampal neurons); (b) self-organizing memristive Ag NW networks realized by drop-casting (scale bar, 500 nm). Adapted from [16] under the terms of Creative Commons Attribution 4.0 License, copyright 2020, Wiley-VCH. (c) Atomic switch network of Ag wires. Adapted from [112], copyright 2013, IOP Publishing. (d) and (e) Single NW junction device where the memristive mechanism rely on the formation/rupture of a metallic conductive filament in between metallic cores of intersecting NWs under the action of an applied electric field and (f) and (g) single NW device where the switching mechanism, after the formation of a nanogap along the NW due to an electrical breakdown, is related to the electromigration of metal ions across this gap. Adapted from [16] under the terms of Creative Commons Attribution 4.0 License, copyright 2020, Wiley-VCH.



reference [121], the software/hardware for interfacing the NW network with the ReRAM readout represents a challenge from the electronic engineering point of view. To fully investigate the computing capabilities of these self-organized systems, modelling of the emergent behaviour is required for understanding the interplay in between network topology and functionalities. This relationship can be explored with a complex network approach by means of graph theory metrics. Current challenges in understanding and modelling the emergent behaviour of NW networks rely on the experimental investigation of resistive switching mechanism in single network elements, including a statistical analysis of inherent stochastic switching features of individual memristive elements. Also, combined experiment and modelling are essential to investigate hallmarks of criticality including short and longrange correlations among network elements, power-law distributions of events and

avalanche effects by means of an information theory approach. Despite scale-free networks operating near the critical point similarly to the cortical tissue are expected to enhance information processing, understanding how critical phenomena affect computational capabilities of self-organized NW networks still remain an open challenge.

5.3. Advances in science and technology to meet challenges

Understanding dynamics from the nanoscale, at the single NW/NW junction level, to the macroscale where a collective behaviour emerges is a key requirement for implementing neuromorphic-type of data processing in NW networks. At the nanoscale, scanning probe microscopy techniques can be employed to assess local network dynamics. In particular, Conductive Atomic Force Microscopy (C-AFM), that provides information on the local NW network conductivity, can be exploited not only as a tool to investigate changes of conductivity after switching events, but also for locally manipulating the electrical connectivity at the single NW/NW junction level [124]. Scanning thermal microscopy can be employed to locally measure the network temperature with spatial resolution <10 nm, well below the resolution of the conventional lock-in thermography [116], providing information about nanoscale current pathways across the sample. At the macroscale, advances in electrical characterization techniques are required for analysing the spatial distribution of electrical properties across the network and their evolution over time upon stimulation. In this framework, one-probe electrical mapping can be adopted for spatially visualizing voltage equipotential lines across the network [125], even if this scanning technique does not allow an analysis of the network evolution over time. In contrast, non-scanning electrical resistance tomography (ERT) have been recently demonstrated as a versatile tool for mapping the network conductivity over time at the macroscale (\sim cm²) [126]. Thus, ERT can allow *in situ* direct visualization of the formation and spontaneous relaxation of conductive pathways, providing quantitative information on the conductivity and morphology of conductive pathways in relation with the spatio-temporal location of stimulation. Advancements in the synthesis of core-shell NWs are required for engineering the insulating shell layer surrounding the metallic inner core that acts as a solid electrolyte. Taking into advantage of the possibility of producing conformal thin films with control of thickness and composition at the atomic level, atomic layer deposition (ALD) represents one of the most promising techniques for the realization of metal-oxide shell layers. Also, alternative bottom-up nanopatterning techniques such as direct self-assembly (DSA) of block copolymers (BCPs) can be explored for the fabrication of selforganizing NW networks with the possibility of controlling correlation lengths and degree of order [127]. This approach can allow a statistical control of network topology. Customized characterization techniques, from the nanoscale to the macroscale, coupled with a proper engineering of NW structure/materials and network topology, will ultimately enable the control of network dynamics needed for efficient computing implementations.

5.4. Concluding remarks

Self-organized NW networks can provide a new paradigm for the realization of neuromorphic hardware. The concept of nanoarchitecture, where the mutual interaction among a huge number of nano parts causes new functionalities to emerge, resembles our brain, where an emergent behaviour arises from the synaptic interactions among a huge number of neurons. Besides RC that represents one of the most promising computing paradigms to be implemented on these nanoarchitectures, unconventional computing frameworks able to process sensor inputs from the environment can be explored for online adapting of robot behavior. In perspective, more complex network dynamics can be explored by realizing computing nanoarchitectures composed of multiple interconnected networks or by stimulating networks with heterogeneous stimuli. In this scenario, NW networks that can learn and adapt when externally stimulated—thus mimicking the processes of experience-dependent synaptic plasticity that shapes connectivity of our nervous system—would not only represent a breakthrough platform for neuro-inspired computing but could also facilitate the understanding of information processing in our brain, where structure and functionalities are intrinsically related.

Acknowledgements

This work was supported by the European project MEMQuD, code 20FUN06. This project (EMPIR 20FUN06 MEMQuD) received funding from the EMPIR programme co-financed by the participating states and from the European Union's Horizon 2020 research and innovation programme.

6. 2D materials

Shi-Jun Liang¹, Feng Miao¹ and Mario Lanza²

¹Nanjing University, China

²King Abdullah University of Science and Technology (KAUST), Saudi Arabia

6.1. Status

With more and more deployed edge devices, huge volumes of data are being generated each day and are waiting for real-time analysis. To process these raw data, these data have to be collected and stored, which are accomplished in sensors, memory unit and computing units, respectively. This usually gives rise to large delay and high energy consumption, which becomes severe with an explosive growth in data generation. Computing in sensory or memory devices allows for reducing latency and power consumption associated with data transfer [128] and is promising for real-time analysis. Functional diversity and performances of these two distinct computing paradigms are largely determined by the type of functional materials. Two dimensional (2D) materials represent a novel class of materials and show many promising properties, such as atomically thin geometry, exceptional mechanical flexibility and strength, etc. Stacking distinct 2D materials on top of each other enables creation of diverse van der Waals (vdW) heterostructures with different combinations and stacking orders, not only retaining the properties of dividual 2D components but also exhibiting additional intriguing properties beyond those of individual 2D materials.

2D materials and vdW heterostructures have recently shown great potential on achieving insensor computing and IMC, as shown in figure 10. There has intense interest in exploring unique properties of 2D materials and their vdW heterostructures for designing computational sensing devices. For example, photovoltaic properties of gate-tuneable p–n homojunction based on ambipolar material WSe₂ were exploited for ultrafast vision sensor capable of processing images within 50 ns [129]. Employing gate-tuneable optoelectronic response of WSe₂/h-BN vdW heterostructure can emulate the hierarchical architecture and biological functionalities of human retina to design reconfigurable retinomorphic sensor array [130].


2D materials and their associated vdW heterostructures were also introduced for IMC devices and circuits to improve the switching characteristics and offering additional functionalities. Several switching mechanisms such as conductive filament [131], chargingdischarging [132–134], grain boundary migration [135], ionic intercalation [136, 137], lattice phase transition [138], etc, have been reported in 2D materials-based planar and vertical devices. With strict limitation in available space and the number of references, only a few representative works are mentioned in this roadmap. Interested readers are encouraged to refer to previous review article [139]. Based on superior thermal stability and atomically-sharp interface of graphene/MoS_{2-x}O_x/graphene vdW heterostructure, a robust memristive device was reported to exhibit endurance of 10^7 at room temperature and stable switching performance in a recordhigh operating temperature of 340 °C [140]. Different from oxide-based memristive devices, metal/2D material/metal vertical devices with layered-structure feature of switching medium were used to mimic high-performance electronic synapses with good energy efficiency [141], which holds promise for modelling artificial neural network in a high-density memristive crossbar array [142]. Reducing the thickness of switching medium down to monolayer allows for fabrication of thinnest resistive switching devices with featuring the conductive-point resistive switching mechanism [143, 144].

6.2. Current and future challenges

In these prototype demonstrations of in-sensor computing, the fabricated device arrays are limited due to the challenge in large area synthesis of 2D materials and vdW heterostructures. However, all these works no doubt show that the unique properties of 2D materials and vdW heterostructures can be utilized to achieve ultralow-latency and reconfigurable in-sensor computing [129, 130]. To eventually realize practical applications of in-sensor computing, innovation is demanded to address issues associated with materials, device physics, array size and controlling peripheral circuit. The challenges that must be overcome in the future include growing large-area single crystal materials, exploiting suitable sensing device structures and mechanisms to handle sensory information, fabricating large-scale computational sensory device arrays with good uniformity, high yield and reliability, as well as designing peripheral circuits that efficiently control programmable operations of in-sensor computing arrays.

In contrast to in-sensor computing, many distinct operating mechanisms were already explored to realize 2D materials-based memristive devices for IMC. From structural point of view, planar or vertical metal/insulator/metal (MIM) configurations of isolated devices with relatively large area are mainly studied. Vertical MIM devices with small lateral area would enable high integration density and would be considered by the industry.

IOP Publishing



However, crossbar study of MIM vertical devices is limited due to the difficulty in synthesis of large area 2D materials with controllable thickness and high quality vdW heterostructure with controllable interface. From electrical point of view, most 2D materials based MIM devices cannot achieve endurances of larger than 10⁶ cycles and stability study of the resistive states was not always demonstrated in multilevel resistive switching devices reported so far. Besides, a unified criteria for yield and variability has been not yet established, which leads to a challenge in evaluating the maturity of 2D materials technology for circuit- and system-level applications of IMC. Clearly stating yield-pass criteria and variability windows of memristive devices is especially important in 2D materials given the large number of local defects intrinsic to scalable synthesis methods as well as other extrinsic defects introduced during integration.

6.3. Advances in science and technology to meet challenges

2D materials for in-sensor computing and IMC are emerging research fields and are still in their infancy. The family of 2D materials database is rapidly expanding and a large number of family members already reported in experiments are available for computational sensory and memory devices. In particular, some of air-stable 2D single-crystal materials such as graphene, h-BN and MoS₂, etc can be synthesized directly on metal wafers and transferred to target wafer substrate in a reliable approach [145]. Besides these advances in materials growth,

recent advances in device physics and arrays as well as peripheral circuits would offer unprecedented opportunities to realize devices arrays on wafer scale with 2D materials that are suitable for in-sensor and IMC applications.

For practical applications of in-sensor computing, further exploration of novel device physics related to 2D materials is required. For example, in the case of vision sensor, a few distinct types of visual information (i.e. orientation, colour, etc) have to be sensed and processed simultaneously with low power consumption and low latency. Notably, significant progresses have been achieved in anisotropic optoelectronics based on low-lattice symmetry 2D materials and bandgap engineering by electrical field and quantum confinement in 2D materials. This would facilitate the device design with new mechanism that enables to sense and process visual information related to orientation, colour and others. Recently, a 32×32 optoelectronic machine vision array has been fabricated with a large-area monolayer MoS₂ synthesized by metal-organic chemical vapor deposition to propel the functional complexity in visual processing to an unprecedented level [146]. Together with the advance in industrial foundry synthesis of large-area ambipolar WS₂ directly on dielectric by plasma enhanced ALD, the promising demonstrations of in-sensor computing should be extended to a larger scale array to benchmark against the performance of conventional materialbased technology.

Traditionally, IMC is usually implemented in 1T1R crossbar array to avoid the sneak path issue. Similarly, 2D material based resistive switching devices should be organized in such way. To that end, radically new growth processes are desired to achieve all 2D materials 1T1R integrated circuit applications. Furthermore, to fabricate a large-scale crossbar array with high yield and low variance, it is required to spatially engineer the precise atomic vacancy patterns on the surface of wafer-scale single crystal 2D semiconductors or insulators, in particular for monolayer form.

Beyond individual 2D materials, vdW heterostructures by stacking 2D materials with distinct electronic properties can retain the properties of each component and exhibit additional properties inaccessible in individual 2D materials. With the breakthrough in material synthesis and fabrication of large-scale integrated arrays as well as peripheral circuits, use of 2D vdW heterostructures in in-sensor and IMC would provide a disruptive technology to solve the challenges of traditional electronics based on von Neumann architecture.

6.4. Concluding remarks

In conclusion, more exploration of 2D vdW heterostructures and continued efforts in exploiting novel device physics will offer more possibilities for energy-efficiency and low-latency in-sensor and IMC, respectively. For example, vdW heterostructure device with high photoresponsivity over broadband spectrum is expected to deal with visual information in an ultra-wide dynamic range close or beyond human retina. By sensing visual information encoded with other degrees of freedom, e.g. multi-coloured wavelengths, polarization, phase, etc, it would further enhance the information processing capability of insensor computing chip. The unique properties of 2D vdW heterostructures are not limited to in-sensor computing applications, but also shows promise in IMC. Exploring spin–orbit torque and ferroelectric polarization in vdW heterostructure for energy-efficient IMC would be a case in point. The current and future challenges of achieving computing in sensory and non-volatile devices mainly arise from materials synthesis, device physics and array integration. However, all recent advances have indicated that 2D vdW heterostructures can provide numerous opportunities for exploration and hold promise for innovation in material growth, device physics, array integration and peripheral circuit for desirable in-sensor and IMC devices as well as their fusion for real-time and high energy-efficiency data analysis applications [147].

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (61625402, 62034004, 61921005, 61974176), and the Collaborative Innovation Center of Advanced Microstructures. FM would like to acknowledge the support from AIQ foundation.

7. Organic materials

*Tyler J Quill*¹, *Scott T Keene*² *and Alberto Salleo*¹

¹Stanford University, United States of America ²University of Cambridge, United Kingdom

7.1. Status

Organic semiconductors (OSCs) have emerged as candidate materials for artificial synaptic devices owing to their low switching energies, wide-range of tunability, and facile ion-migration due to the large free volume within the material. OSCs emulate neuroplasticity at the single unit level with a wide range of synaptic switching mechanisms demonstrated for both two-terminal devices, which utilize filament formation [148], charge trapping [149], and ion migration, as well as three-terminal transistor-like architectures such as ion-gated electrochemical transistors [150] and charge trapping transistors. In most cases, the resistive switching of polymers is either via metal-ion migration to form conductive pathways (figures 11(a) and (b)) or by reversible doping where the oxidation state of the OSC is modulated via charge trapping on defect sites (such as implanted nanoparticles), redox reactions (e.g. protonation/deprotonation), or ion intercalation (figures 11(c) and (d)).

The ability to tailor the properties of OSCs makes them a particularly promising class of materials for neuromorphic devices since both chemical and microstructural control over the materials can dramatically influence device performance (figure 11(e)). Side-chain engineering of OSCs can enhance ionic mobility in the materials, enabling relatively high-speed device operation [151], whereas modification of chemical moieties on the polymer backbone can be used to tune of energy levels and electronic conductivity [152]. The crystallinity and microstructure of these materials allow for yet another degree of freedom which can be exploited to further optimize them to emulate synaptic behavior [153]. Lastly, the relatively low-cost and solution processability makes OSCs particularly attractive where large-area or printable devices are desired, such as when interfacing with biological systems.


Thus far, OSC neuromorphic devices have demonstrated a variety of synaptic functionality, including the representation of synaptic weight as electrical resistance [150], excitatory postsynaptic potential (EPSC), global connectivity [154], and pulse shaping [155]. This broad functionality makes OSCs promising for applications ranging from high-performance computing to biological interfacing of neuromorphic systems. Recently, three-terminal electrochemical devices with low switching energy have been demonstrated which can overcome several challenges associated with parallel operation of a hardware neural network in a crossbar architecture [156], showing the promise for organic materials in neuromorphic engineering. In this work, however, we will discuss the general challenges and outlook for using OSCs in neuromorphic computing without focusing on any single device, application, or architecture.

7.2. Current and future challenges

Speed. Increasing the speed of organic devices has long been a central goal in materials engineering. In OSCs, device speed can be limited by electronic mobilities, ionic mobilities, defects, or stray capacitances. Recent advances in side-chain engineering of mixed ionic/electronic conducting OSCs have notably improved the speed at which organic devices operate, but their speeds still lag behind their inorganic counterparts. Furthermore, the electronic mobilities of OSCs are typically lower than their inorganic/crystalline counterparts, also limiting speed.

Density. Patterning of OSCs presents a fundamental challenge for increasing device density due to the incompatibility between OSCs and many of the solvents and photon wavelengths used in photolithography. Consequently, outside of additive manufacturing methods such as printing, the most widespread methods of patterning OSCs rely on either sacrificial hard masks to protect the OSC from solvents [157], or the use of orthogonal photoresist solvents which do not damage the OSC [158]. Additionally, deposition of highly uniform OSC films is challenging due to complex microstructures, and nontraditional fabrication techniques are required to enable vertical architectures to reduce the individual device footprint (three-terminal) and interconnect complexity.

Integration. An additional challenge revolves around incorporating organic neuromorphic devices/systems with traditional digital systems while avoiding damage during back-end-of-the-line (BEOL) processing. The electronic properties of OSCs typically degrade at elevated temperatures (typically >150 °C) due to phase transitions (e.g. T_g , T_m) or temperature-induced morphological changes (e.g. thermal expansion, backbone twisting). This temperature sensitivity is problematic due to required processing temperatures of ~400 °C used to anneal Cu interconnects in inorganic device stacks. Although crossbar architectures offer a straightforward method of replicating the vector-matrix multiplication (VMM) desired in ANNs, the need to control sneak



currents often requires an access device, increasing the complexity of the array and providing an additional integration challenge.

Environmental and electronic stability. A final remaining challenge for OSCs is to achieve long-term device stability and resistance state retention. Interfaces of OSCs and dielectrics are susceptible to formation of traps resulting from exposure to oxygen or moisture, leading to irreversible changes in device performance. Additionally, because of the inherently low switching energy found in many organic neuromorphic devices, 'SET' OSCs are susceptible to leakage due to parasitic reactions with the surrounding atmosphere [159]. Finally, both the charge transport and doping reactions in OSCs must be stable at the typical operating temperatures of computers (\sim 85 °C) without suffering from changes in morphology due to thermal annealing.

7.3. Advances in science and technology to meet challenges

Speed. To improve organic device speed, it is first essential to identify the rate limitations. For example, in ion-gated devices, both the ion and electronic mobilities can dictate the switching speed. Spyropoulos *et al* have shown that for ion-gated devices, presence of the electrolyte ions within the OSC can improve device response time [162]. Additionally, strain-relaxation following the insertion or removal of species may also limit the time to reach a stable resistance state [163]. Once the fundamental speed limits are identified, polymers and devices can be engineered to optimize for the critical parameters (e.g. ion mobility, parasitic capacitances, strain response). In a recent example, selection of the electrolyte and OSC materials allowed for electrochemical neuromorphic devices to operate with 200 ns write operations and <1 μ s read-write cycles (figures 12(a) and (b)) [160].

Density. Novel nanofabrication processes which can accurately define OSCs in vertical architectures can reduce device footprints, increasing device density. Strategies such as utilizing metal contacts as hard masks




Figure 12. State-of-the-art organic neuromorphic devices. (a) Analog resistance tuning of an electrochemical neuromorphic device under ± 2 V 200 ns write pulses (gray shaded area), followed by 100 ns write-read delay and ± 0.3 V 500 ns readout (orange shaded area). The horizontal dashed lines are a guide to the eye to represent tunable conductance states. (b) The volumetric scaling of electrochemical doping enables channel conductance of devices to be tuned with increasingly lower write energies and shorter write pulses as device sizes are reduced. (c) Cross-sectional schematic of fabrication procedure of densely packed ion-gel-gated vertical P3HT synapses and (d) optical microscopy images of a crossbar array. (e) A non-volatile ionic floating gate (IFG) memory consisting of a filament forming access device (green) attached to a PEDOT: PSS organic synapse (blue). (f) Schematic of parallelprogrammable neuromorphic array using IFG memory divided into a two-layer neural network, as indicated by orange and green. Analog network inputs V_i^R are applied across the source-drain rows, while programming inputs V_i^W and V_j^W are applied along the gate row and drain column, respectively. Adapted from reference [160], AAAS, (a) and (b); reproduced from reference [161], Springer Nature Ltd, (c) and (d); adapted from reference [156], AAAS, (e) and (f).

can enable nanopatterning of OSC channels with resolutions limited by conventional lithographic techniques [164], but defining gate and electrolyte geometries with similar precision for complete three-terminal devices introduces additional complexity. Choi *et al* recently demonstrated vertical three-terminal electrochemical neuromorphic devices which reduced the single cell footprint to *ca* 100 μ m by 100 μ m in a crossbar architecture using photo-crosslinked P3HT as the channel material (figures 12(c) and (d)). In principle, this cell could be reduced significantly using the same general technique with the use of advanced photolithography.

Integration. Advancements in non-traditional chip manufacturing (BEOL alternatives) [165] are necessary for seamless integration of OSCs with silicon technology. Sneak currents in neuromorphic arrays can be avoided by using filament-forming access devices coupled to three-terminal memories, as shown by Fuller *et al* (figures 12(e) and (f)) [156]. Increasing the temperature stability of OSCs also helps enable complete integration with conventional BEOL processing. Recently, Gumyusenge *et al* demonstrated that nanoconfined OSCs

in high-performance polymer blends exhibit robust temperature-independent mobilities up to 220 °C [166], a notable step towards integration.

Environmental and electronic stability. Although the stability of OSCs presents challenges for developing neuromorphic devices, recent design strategies provide promise. The modularity of OSCs enables tuning of both molecular orbital energies as well as morphology. For example, engineering OSCs with high ionization potentials can eliminate cross-reactions with moisture or oxygen [152]. Further optimization of OSC crystallinity [153] and encapsulation methods [159], which shield devices from the ambient atmosphere, could further improve stability.

7.4. Concluding remarks

Organic materials have rapidly grown into a promising class of materials for neuromorphic systems and could be harbingers of other unconventional semiconductors for these applications. While there are great challenges facing organics before they are suitable for commercial neuromorphic computing systems, including significant improvements to speed, density, integration, and stability, there are no fundamental barriers preventing OSCs from satisfying these metrics.

Owing to their biocompatibility and softer mechanical properties, organics are of interest for direct connections between biological systems and neuromorphic computers, such as in brain-machine interfaces and adaptive prosthetic devices. Inspired by the biomimicking nature of neuromorphic systems, there is a strong push towards direct integration with prosthetics to match the low power computation already found in the human brain. These systems and devices would form direct interfaces with tissue, repairing augmenting functionality or act as 'smart system' for wearable electronics. Recently, organic neuromorphic devices have been tuned to match either the output signals for 'talking' to [155] or by responding to neurotransmitter signals for 'listening' to [167] the biological domain. We postulate that organic materials will shine as neuromorphic devices in bioelectronic interfaces due to the relative maturity of the materials class in the bioelectronic space.

Acknowledgements

AS and STK acknowledge support from the National Science Foundation and the Semiconductor Research Corporation (Award NSF E2CDA #1507826). TJQ acknowledges support from the National Science Foundation Graduate Research Fellowship Program under Grant DGE-1656518.

Section 2. Neuromorphic circuits

8. Spintronics

Julie Grollier, Danijela Marković and Alice Mizrahi

Université Paris-Saclay, France

8.1. Status

Spintronics, or spin electronics, manipulates the spin of electrons in addition to their charge. This brings multiple interesting features for neuromorphic computing: the non-volatile memory provided by nanomagnets and the non-linear dynamics of magnetization induced by fields or currents [168]. These two aspects allow the same materials to be used to mimic the essential operations of synapses and neurons. Important experimental results have thus been obtained in recent years.


Synapses. The first way to realize spintronic synapses is to store the weights in digital spin torquemagnetic random access memories (ST-MRAMs) [169]. Gigabit devices from the latter are now commercially available in several large foundries. They consist of magnetic tunnel junctions (MTJ), formed by an ultra-thin (~1 nm) insulator sandwiched between magnetic layers, integrated in the CMOS process. The main advantage of ST-MRAMs over their competitors is their endurance, which is more than two orders of magnitude higher, a very important factor for the chips dedicated to learning, that will require very many read/write cycles. Indeed, the resistance change mechanism comes from a reversal of magnetization by current pulses of the order of nanoseconds and a hundred millivolts, a purely electronic phenomenon that does not require the movement of ions or atoms in a nanostructure as in ReRAMs or PCMs. Moreover, they are non-volatile, retaining information even when the power is switched off. Associative memories integrating ST-MRAMs (figure 13(a)) have enabled significant gains in power consumption, with only 600 μ W per recognition operation, i.e. a 91.2% reduction compared to a twin chip using conventional static random access memory [169].

The second way to realize spintronic synapses is to directly imitate a synapse with a magnetic tunnel junction. In this case, the junction acts as a memristor device, which takes as input a current and multiplies it by its resistance, which thus plays the role of the synaptic weight. The stability of magnetization in MTJ allows them to retain the value of the weight. Since magnetization is naturally bistable, MTJ are very good candidates for neural networks with binary weights [170]. It is also possible to modify the materials or geometry so that the magnetization changes orientation via non-uniform states. This has allowed to experimentally realize analog synapses (figure 13(b)) [171–173], as well as to train a small neural network with magnetic multi-state synapses (figure 13(c)) [174].

Neurons. In most neural network algorithms, neurons simply apply a non-linear function to the realvalued synaptic inputs they receive. The characteristics of the nonlinear dynamics of spintronics can be exploited to mimic biology more closely, which could lead to increased computing functionalities such as local and unsupervised learning. Biological neurons transform the voltage on their membrane into electrical spike trains, with a mean frequency that is non-linearly dependent on the voltage. MTJ transform DC inputs into an oscillating voltage with a frequency that depends non-linearly on the injected current. This property can be used to imitate neurons. In stable junctions such as those used for ST-MRAMs, the spin torque can induce oscillations between about ten MHz and ten GHz depending on the materials and geometry. These oscillations have been used with a single device to recognize pronounced digits with a time-multiplexed reservoir [175]. Four coupled spintronic nano-oscillators were also trained to recognize vowels via their synchronization patterns to RF inputs (figure 14(a)) [176]. In unstable junctions, thermal fluctuations may be sufficient to induce tele-graphic voltage behavior, allowing the mimicking of stochastic neurons with minimal energy consumption. Neuromorphic tasks have been performed by small experimental systems composed of such junctions, using neural networks [177, 178] or probabilistic algorithms [179].


8.2. Current and future challenges


The challenge is now to create large-scale spintronic neuromorphic circuits capable of solving useful tasks at low energy consumption, especially for embedded AI. This requires the integration of layers of synapses and spintronic neurons interfaced in deep networks. The major disadvantage of spintronics compared to other technologies is the low resistance ratio between the OFF and ON states of the MTJ, of the order of 2–3 compared to more than 10 in other resistive memory technologies. These small variations in resistance make the reading of the state of the junctions more complex, and have so far prevented the development of magnetic tunnel junction crossbar arrays as realized in other resistive technologies. It is therefore necessary to continue the effort on the material side to push the resistance variations towards their theoretical R_{off}/R_{on} value >100



Roadmap

Figure 13. Spintronic synapses. (a) Schematic of an associative memory circuit with ST-MRAM cell, reproduced from [169]. (b) R–I hysteresis loop of a spintronic memristor based on current-induced domain wall displacement in a magnetic tunnel junction, reproduced from [171]. (c) R–I hysteresis loop of a spintronic memristor exploiting spin–orbit torques in a ferromagnetic/antiferromagnetic bilayer, reproduced from [173].





[180, 181]. On the CMOS design side, the development of low-power circuits allowing efficient reading of the state of the junctions, such as sense-amplifier, is crucial. As for all technologies, device reliability and scaling is a challenge, especially in analog implementations. The first demonstrations will certainly rely on binarization of resistance values for the inference phase and implementation of hardware binary neural networks, before end-to-end on-chip learning solutions are developed.

Combining ionic and spintronic effects will be one of the keys to efficient learning of neuromorphic chips. It was recently demonstrated that strong magnetoelectric effects enable control of magnetic dynamics by the electric field created by the interface, more efficiently than previous methods [182].

A critical challenge for the development of hardware neural networks is to achieve a high density of connections. Spintronics offers several opportunities to tackle this issue. Long-range connections can be implemented via spin currents and magnetic waves or by physically moving magnetic textures such as skyrmions and solitons [168, 183]. Furthermore, the multilayer nature of spintronic devices allows them to naturally stack in three dimensions, opening the path to vertical communication [184]. Spintronic neuromorphic chips will be able to receive as inputs fast signals compatible with digital electronics (classical binary junctions), radio-frequency inputs (GHz oscillator), as well as inputs varying at the speed of the living world, thanks to superparamagnetic junctions or magneto-electric effects that can operate at timescales between seconds and milliseconds. There is active research on developing spintronic devices for on-chip communication (using their capability to emit and receive microwaves), magnetic sensing (with promising biomedical applications) and energy harvesting, all of which could benefit neuromorphic chips [168].

Taking full advantage of the dynamical behavior of spintronic devices will require the development of dedicated learning algorithms, inspired by advances in both machine learning and computational neuroscience. The fact that the behavior of spintronic devices relies on purely physical phenomena that can be predictively described and integrated into neural network programming libraries is a key enabler for this task [185].

8.3. Advances in science and technology to meet challenges

Spintronics is undergoing promising new developments from which neuromorphic chips could benefit. Antiferromagnetic materials and interfaces with optics bring the possibility of information processing and transmission at THz speed [186]. Spin/charge conversions are increasingly efficient thanks to new materials such as topological insulators.

Finally, the multifunctionality of spintronics also makes it possible to train complex physical systems that do not exactly reproduce the synapse/neuron structure, such as, for example, arrays of spin wave transmitters/receivers, or fixed or mobile magnetic particles, such as skyrmions and domain walls [187]. Micromagnetic simulations with predictive power, coupled with gradient descent, have modeled learning tasks [185]. Experimental demonstrations with these complex physical systems remain to be carried out.

8.4. Concluding remarks

In the short term, neuromorphic spintronics should see the commercialization of AI chips storing synaptic weights into current and future generations of ST-MRAMs. This should be followed by the development hardware neuron circuits leveraging the dynamical properties of MTJ to implement synapses and neurons for inference and learning. In the longer term, more exotic materials and textures offer the fascinating prospect of in-materio computation based on complex physical effects.

9. Deep learning

Peng Yao and J Joshua Yang

University of Southern California, United States of America

9.1. Status

The development of deep learning (DL) has brought AI to the spotlight of broad research communities. The brain-inspired neural network models with different structures and configurations have made significant progress in a variety of complex tasks [5]. However, in conventional von Neumann architecture, the physically separated computing unit and memory unit require frequent data shuttling between them, which results in considerable power consumption and latency cost. One promising approach to tackle this issue is to realize IMC paradigm where each underlying device component functions as memory and computation elements simultaneously. Non-volatile devices based on resistive switching phenomena [13, 188], such as redox memristor, phase change, magnetic and ferroelectric devices, could support such computing system and show greatly improved performance in data centric computation tasks.


Analogue resistive-switching memory based IMC is promising to bring orders of magnitudes improvement in energy efficiency compared to the conventional von Neumann hardware. The devices are assembled in a crossbar structure to conduct VMM operations, where the input vectors are encoded as voltage amplitude, pulse widths, pulse numbers, or sequential pulses with different significances, and the matrix elements are mapped to tunable cell conductance where each cell is often represented in the differential form of a pair of devices. Thanking to Ohm's law for multiplication and Kirchhoff's current law for accumulation, the dense crossbar could conduct multiplication-accumulation (MAC) fully in parallel and the computation occurs at the data location. Since VMM calculation accounts for the majority of computation during inference and training of deep learning algorithms, this IMC paradigm could help the hardware to meet stringent requests of low power dissipation and high computing throughput.

Major progresses have been made in this area, spanned from device optimization to system demonstration [13, 21, 188]. The oxide-memristor devices have been scaled down to 2 nm in an array [189] and 3D stacked architecture has been fabricated in laboratory to enhance the network connectivity [190]. In addition, various DNN models, including perceptron [191, 192], multiple layer perceptron (MLPs) [193], long short term memory (LSTM) [194] based recurrent neural networks (RNNs), and convolutional neural networks (CNNs) [74], have been demonstrated based on nonvolatile resistive-switching crossbars or macro circuits. These demonstrations have covered the typical learning algorithms for supervised learning, unsupervised learning and reinforcement learning. More recently, a multiple-array based memristor system [74] and some monolithically integrated memristor chips have been demonstrated [195, 196], and it is encouraging to see that this kind of IMC system could achieve an accuracy comparable to software results and reach >10 TOPS/W energy efficiency using 8 bit input precision [74]. However, despite the fast development of hardware prototypes and demonstrations, a monolithically integrated IMC chip with large and tiled crossbars (shown in figure 15) for practical and sophisticated DL models (e.g. ResNET50) is still under-explored, and the accomplished tasks are limited to relatively small dataset (e.g. MNIST, CIFAR10) rather than handling large workloads (e.g. ImageNet).

9.2. Current and future challenges

One fundamental challenge of the IMC system originates from the device non-ideal characteristics and process issue towards large-scale heterogeneous integration. No matter what the specific switching mechanism is, the analogue-switching non-volatile memory device show some inherent device-wise and cycle-wise variance, reading fluctuations, state drift or stuck, or limited on/off ratio. Although the rich dynamics could be explored for neuromorphic systems in attempt to resemble biological phenomenon more faithfully, these uncontrolled non-ideal behaviors would inevitably affect the accuracy of the DL model and deteriorate the system performance in most cases. The high-level system researches often unrealistically assume ideal device features, resulting in some significant discrepancy from practical applications. Moreover, the forward propagation and backward propagation pose stringent requirement for linear I-V behavior to conduct accurate MAC operations, while the update procedure demands the device showing state-independent conductance tuning curves with sufficient linearity and symmetry with respect to the programming pulses to insure training convergence and efficiency. In spite of the undergoing advances in device optimization and exploration, ideal device is still missing. Thus, the IMC systems need various compensating strategies, which can substantially degrade system efficiency.

In terms of heterogeneous integration, the advanced foundry process for large-scale fabrication of analogue-switching devices is absent, regardless of the specific device category. Although NOR flash is




commercially available, it is large operation voltages and slow speeds together with its limited endurance and scalability make it at best an interim solution, to be replaced by emerging devices. Oxide memristor is promising in dense integration given the demonstration of 2 nm feature size and 3D stacking ability at lab. However, only 130 nm analog-switching technology [195, 196] and 22 nm digital-switching technology [197] in foundry have been reported. Many other kinds of devices require back-end process with high temperature, complex layer deposition or special handling process, which present obstacles for them to be monolithically integrated with the mainstream CMOS technology. The absence of high-uniformity and high-yield process in mainstream foundries for large-scale and small-footprint integration of analogue-switching devices has been slowing down the development of IMC circuits.

Errors in analog IMC and inefficiency of periphery circuits also imposes serious challenges for practical hardware. Analog computing directly utilizing physical laws is superior in energy efficiency, whereas it only suits for low-precision tasks so far. Although DL algorithms put loose constraints on parameter precisions (such as 4 bit weights for regular inference tasks), state-of-theart models still demands accurate digitalized value representations. However, the conductance states of analog devices always follow a certain distribution and deviate from the target mapping values, which would bring in weight representing errors. In addition, at the array/crossbar level, the parasitic effects along the metal wires would lead to inevitable IR drop and result in inaccurate programming and computing. This effect becomes more severe if the array size is increased for higher performance. Such systematic errors may be mitigated through some algorithm and architecture codesign, such as compensations in the mapping algorithms. The periphery circuits would also introduce computing errors due to the voltage loss on analogue switches, transistor mismatch, unfixed clamping voltage and environmental fluctuations. All these together would substantially lower analogue computing accuracy and prevent IMC system from realistic applications if not appropriately addressed.

To take the full advantage of IMC features, all necessary functional blocks should be integrated monolithically with device crossbars (as shown in figure 15), including buffers, interfacial circuits (mainly ADCs, converting the accumulated current to digital signals), routing units, control logic and digital processing. These circuits are expected to match the device operating requirements, such as programming voltage and driving currents. In such a complete on-chip system with tiled crossbars, the auxiliary periphery circuits might consume much more power, area and latency than the analogdomain VMM calculation. Although the IMC paradigm eliminate the movement of DL weights, it still needs data flowing between different layers and requires on-chip memory accessing. Meanwhile, the parallel MAC calculations desire multiple ADCs locating at the end of each column to carry out fast conversations frequently. According to the profile of a designing instance, the ADCs account for the majority of power and area overhead (shown in figure 16) [74]. Exploiting a larger crossbar to conduct VMM is beneficial to boost system performance by amortizing the periphery circuit overhead in the whole system, which, however, would lead to larger parasitic capacitance and resistance, higher dynamic range of the output current and lower device utilization ratio. The inefficiency of periphery circuits, especially the ADCs, is becoming the system bottleneck of IMC hardware, where innovations are needed in the co-design of device and architecture.

9.3. Advance in science and technology to meet challenges

Extensive multi-disciplinary efforts are needed in order to advance the development of IMC-based deep learning accelerators, as co-designs of device, circuit, architecture and algorithm are required to build practical IMC system [198].



First of all, researches in material engineering and device optimization should be conducted either based on present analogue-switching non-volatile devices or for the exploration of novel devices, aiming at enhanced reliability, improved programming linearity and symmetry while maintaining high switching speed, low programming power and intensive scaling potential. In addition, stable stack process for large-scale heterogeneous integration of highly uniform crossbar is needed for practical applications. The development of 3D process could drive the device density to next level and bring in extra dimension to explore more efficient system. Even more importantly, 3D structures enable the massive connectivity and low-loss communications required for complex neural networks.

Second, at the macro-circuit level, there is plenty of room to optimize the crossbar structure and the periphery circuits. For example, basic two-transistor-2-memristor (2T2M) configuration [195] could be utilized as a signed-weight unit to construct IMC arrays, where *in situ* subtraction operation is conducted in analog domain with the differential current being accumulated subsequently. Such configuration reduces total flowing currents to mitigate IR drop effect, which makes it available to build larger crossbar. Apart from this, encoding input signal by pulse width or low voltage amplitude range might bypass the nonlinear current–voltage characteristic issue, at the expense of increasing system latency or circuitry complexity. On the other hand, novel periphery circuitry design customized for IMC is required, including fast, low-power ADC and high-throughput routing scheme with little on-chip memory. For example, time-domain interfaces could be used to replace conventional ADC-based interfaces [199]. Furthermore, some emerging devices with rich nonlinearities [96] could potentially replace circuitry blocks directly, such as implementation of device-wise ReLU function [200, 201].

Finally, system-level innovations are critical to expedite the development of IMC hardware. From architecture perspective, time division multiplexing of ADCs and replicating same weights to different crossbars are key technologies in order to optimize the system dataflow and boost the computing parallelism. In addition, despite the difficulties in data storing and transmission in analogue domain, interfacing, transferring and processing the information in analogue format is intriguing due to the potential of huge efficiency benefits. From the algorithmic point of view, configuring and optimizing the DL models to fit IMC device features and reduce hardware cost is demanded. On-chip learning, hardware-aware learning and hybrid learning are some representative works to mitigate device non-ideal characteristics and computing errors.

9.4. Concluding remarks

IMC based on analogue-switching non-volatile device shows exceptional superiority regarding computing throughput and energy efficiency than the conventional von Neumann hardware, suited for dealing with data centric problems and brain-inspired deep learning algorithms. In spite of the significant advancements in device explorations and system demonstrations, device non-ideal behaviors, difficulties in large-scale heterogeneous integration, inaccuracies of analog computing and inefficiency of periphery circuits pose great challenges to promoting the IMC technologies for practical application. Monolithic integration of a complete system that unleash the full potential of the IMC features with tiled crossbar architecture and smooth dataflow is still missing. Consequently, extensive co-design efforts from device optimization, circuitry design, architecture exploration and algorithm tailoring are consistently needed. With the utilization of more emerging devices and advanced 3D integration process, the IMC promises bright future of deep learning hardware.

10. Spiking neural networks

Giacomo Indiveri

University of Zurich and ETH Zurich, Switzerland

10.1. Status

The design of neuromorphic circuits for implementing SNN represents one of the main activities of Neuromorphic Computing and Engineering. Currently, these activities can be divided into two main classes: (i) the design of large-scale general-purpose spiking neural network *simulation* platforms using digital circuits and advanced complementary metal-oxide semiconductor (CMOS) fabrication processes [202-204], and (ii) the design of analog biophysically realistic synaptic and neural processing circuits for the realtime emu*lation* of neural dynamics applied to specific sensory-motor online processing tasks [3, 205–210]. This latter effort pursues the original goal of Neuromorphic Engineering, set forth over thirty years ago by Carver Mead and colleagues [211, 212], to use the physics of electronic devices for understanding the principles of computation used by neural processing systems. While the strategy of building artificial neural processing systems using CMOS technologies to physically emulate cortical structures and neural processing systems was mainly restricted to academic investigations for basic research in the past, the recent advent of emerging memory technologies based on memristive devices spurred renewed interest in this approach, also for applied research and practical applications. One of the main reasons is that the analog and mixedsignal analog/digital neuromorphic processing architectures that implement adaptation, learning, and homeostatic mechanisms are, by construction, robust to device variability [4, 213]. This is a very appealing feature that enables the exploitation of the intricate physics of nanoscale memristive devices, which have a high degree of variability, for carrying out complex sensory processing, pattern recognition, and computing tasks. Another appealing feature of these mixed-signal neuromorphic computing architectures, that enables a perfect symbiosis with memristive devices, is their 'IMC' nature: these architectures are typically implemented as large crossbar arrays of synapse circuits that represent at the same time the site of memory and of computation. The synapses in each row of these arrays are connected to integrate-and-fire (I & F) soma circuits, located on the side of the array. The soma circuits sum spatially all the weighted currents produced by the synapses, integrate them over time, and produce an output pulse (spikes) when the integrated signal crosses a set threshold. In turn the synapses are typically stimulated with input spikes (e.g., arriving from other soma circuits in the network), and convert the digital pulse into a weighted analog current [3, 213]. Depending on the complexity of the synapse and soma circuits, it is possible to design systems that can exhibit complex temporal dynamics, for example to create spatiotemporal filters matched to the signals and patterns of interest, or to implement adaptive and learning mechanisms that can be used to 'train' the network to carry out specific tasks.

10.2. Current and future challenges

ANN can achieve impressive performance in solving an incredible amount of problems, thanks to their *learning* abilities. The backbone of learning in neural networks simulated on standard computing technologies is the 'backpropagation through time' (BPTT) algorithm [214]. So an important challenge for SNN is to understand how to implement learning rules as powerful as BPTT, but using plasticity mechanisms that only have access to local signals, due to their 'IMC' nature. This goal is particularly challenging for hardware SNN, because their synaptic circuits have limited resolution, and because the devices used to store their weights are often affected by cycle-to-cycle and device-to-device variability.

In addition, neuromorphic spiking neural architectures do not have access to external memory blocks where to store or retrieve information, like in von Neumann architectures. On one hand this has the advantage of removing the infamous von Neumann memory bottleneck [215, 216], which refers to the problem of having to transfer information from storage areas to computing ones at very high rates, and which accounts for the vast majority of the exceedingly large power consumption figures of standard computers. On the other however, this introduces the problem that spiking neural computing systems cannot arbitrarily choose the resolution of the data they need to store, or the times at which they access it during the course of their computations.

This adds an even more complicated challenge of understanding how to manage memory and time in such architectures: SNN can be seen as non-linear filters that process information online, as data is flowing through them. To carry out real-time computation on a sensory input stream, these networks must retain a short-term memory trace of their recent inputs. Without learning, there are fundamental limits on the lifetimes of these memory traces that depend on both the network size and the longest time-scales supported by the elements used in the network [217]. So an important requirement for enabling the construction of hardware SNN that can be deployed in a wide range of realworld applications is to develop volatile memristive technologies that have a large distribution of time scales, ranging from micro-seconds to hours and days [4].

To exploit the features of neuromorphic spiking hardware to their fullest extent, there are therefore two tightly interlinked critical challenges that need to be addressed in conjunction: (i) the development of a radically different theory of computation that combines the use of fading memory traces and non-linear dynamics with local spike-based learning mechanisms, and (ii) the development of both volatile and nonvolatile memory technologies, compatible with CMOS analog circuits, that support the theories developed.

10.3. Advances in science and technology to meet challenges

The challenges that the neuromorphic computing and engineering field faces, for understanding how to best design signal processing and online learning mechanisms in hardware SNN for solving complex problems might seem insurmountable. However, biological brains are an existence proof that robust and stable computation can be achieved, using a computing substrate that is analog, and that uses inhomogeneous and imprecise signal processing elements. So understanding how animal brains, even small insect brains with fewer than one million neurons, manage to achieve this tasks will be key for making progress also in this domain. Specifically, the advances in science that are required to meet these challenges need a strong interdisciplinary approach. Advances in theoretical and computational neuroscience will provide a core component. But these will need to be complemented with notions and results from multiple sub-fields of electrical engineering, such as information theory, signal processing, and control theory, as well as other disciplines such as mathematics, computer science, and robotics.

In parallel, it is clear that the technology used today in conventional computing systems is not ideally suited for building brain-inspired neuromorphic hardware. Emerging nanoscale memristive memory technologies represent a promising development that can provide solid-state electronic elements able to emulate different properties of biological synapses and neurons. Important advances in technology required to meet the challenges outlined above can indeed be provided by the development of both volatile and non-volatile memristive devices with characteristics that are compatible with the specifications provided by the theory. However, to mass-produce neuromorphic computing systems at scale, it is important that these device are compatible with standard CMOS fabrication processes. As one of the most important features of existing CMOS circuits used for implementing spiking neural network is their ability to perform computations using extremely small amounts of power [3], it will be important that the memristive devices designed to be co-integrated with these circuits do not require large amounts of currents to change conductance.

As the development of neuromorphic circuits and memristive devices for building hardware SNN to carry out computation is a very recent phenomenon, there is a unique opportunity for making concrete progress toward meeting the challenges faced by following, a co-design approach that drives both advances in science and in technology together.

10.4. Concluding remarks

Implementing SNN computing systems with analog CMOS circuits and memristive devices is hard. Besides few examples of proof-of-concept systems that have been applied to very specific tasks, such as sensory processing or spatio-temporal pattern recognition for bio-medical applications [96, 218–222], no general purpose solution exists yet. More importantly, no well established formal methodology exists for automatically designing or programming them.

However we are witnessing incredible progress being made independently in ANN, in machine learning, in neuroscience, and in memory technology developments. In addition, there is a large demand for the development of novel low-power computing technologies for applications 'at the edge', i.e., applications that need to process data measured locally, without connecting to remote servers on the internet, often with low latency and in compact packages.

So brain-inspired approaches for building such technologies are extremely promising, and the potential of research and development in memristive/CMOS SNN computing systems is extremely high.

Acknowledgements

This paper is supported in part by the European Union's Horizon 2020 ERC project NeuroAgents (Grant No. 724295), and in part by the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 871371 (project MeMScales).

11. Emerging hardware approaches for optimization

John Paul Strachan^{1,*} and Suman Datta²

¹Hewlett Packard Enterprise, United States of America ²University of Notre Dame, United States of America *Currently at Peter Grünberg Institute (PGI-14), Germany RWTH Aachen University, Germany E-mail: j.strachan@fz-juelich.de and sdatta@nd.edu

11.1. Status


This perspective outlines a roadmap of emerging hardware approaches that utilize neuromorphic and physicsinspired principles to solve combinatorial optimization problems faster and more efficiently than traditional CMOS in von Neumann architectures. Optimization problems are ubiquitous in modern society, needed in training ANN, building optimal schedules (e.g., airlines), allocating finite resources, drug discovery, path planning (VLSI and shipping), cryptography, and graph analytics problems (social networks, internet search). Such problems are often extremely challenging, requiring compute resources that scale exponentially with the problem size (i.e., NP-complete or NP-hard complexity). Mathematically, in a combinatorial optimization problem [223] one has a pre-defined cost function, c(x), that maps from a discrete domain \vec{X} (nodes, vectors, graph objects) to, the real number space, and the goal is to find the \vec{x}_{opt} that achieves the globally optimum cost value $c_{min}(\vec{x}_{opt})$.

While exact methods for solving optimization problems have been developed, these can be too timeconsuming for challenging or even modest-sized instances. Instead, there is steadily rising popularity for faster meta-heuristic approaches, such as simulated annealing [224] and evolutionary algorithms [225], computing models such as Boltzmann machines [226], Ising models [227, 228], and variations of Hopfield networks [229]. These take inspiration from physical and biological systems which solve optimization problems (figure 17) spontaneously. Many naturally-occurring phenomena, including the trajectories of baseballs and shapes taken by amoeba, are driven to extrema of objective functions by following simple principles (e.g., *least action* or *minimum power dissipation* [230]). In one example, proteins, which are long chains of amino acids, can contort into an exponentially large number of different shapes, yet they repeatably stabilize into a fixed shape on the time-scale of milliseconds. For a protein composed of only 100 peptide bonds, it is estimated that there are over 10^{300} different conformal shapes. Even exploring one every picosecond (10^{-12} s), takes more than the age of the Universe to explore them all. Instead, nature uses efficient dynamics to arrive at a solution in less than a second.

11.2. Current and future challenges

Neuromorphic and physics-based hardware complements the meta-heuristic algorithms and models described above. Emerging hardware approaches include quantum-based annealers [231, 232], optical or coherent Ising machines [233–235], CMOS-based digital annealers [236–238], analog resistive memory-based systems [239, 240], coupled oscillators [230, 241–243], and probabilistic bit logic [179, 244]. Some of these techniques are illustrated in figure 17 and table 2 highlights some of their respective features, strengths, and challenges.

We highlight and compare some of today's emerging approaches on the max-cut benchmark (partitioning a graph). D-wave's quantum annealer uses low-temperature superconducting devices comprising 5000 bits in a sparsely coupled network. Prior work shows solving a 200 node cubic graph in 11 ms, consuming around 25 kW power to operate at cryogenic temperature. The coherent Ising machine (CIM) based on optical parametric oscillators is another approach, where a fully connected 2000 node CIM uses a kilometer long fiber cavity to accommodate the degenerate optical parametric oscillator pulses and takes about 50 ms to solve maxcut problem on a 200 node cubic graph. CIM utilizes measurement-and-feedback schemes for coupling the spins, that is provided by traditional CMOS based field-programmable-gate-arrays (FPGAs). Several CMOSbased digital or mixed-signal hardware accelerators have also been developed. The Ising chip demonstrated by Hitachi uses CMOS static random access memory (SRAM) cells as spins while the coupling is realized using digital logic gates, used to implement the simulated annealing algorithm. Due to its non-von Neumann architecture, it exhibits a $50 \times$ lower energy-to solution over that of a CPU running a greedy algorithm to find max-cut of a 200 node random cubic graph. The energy dissipation is still orders of magnitude more than other Ising solvers recently implemented with emerging devices such as resistive RAMbased cross-bar arrays and insulator-to-metal phase transition (IMT)-based coupled oscillator arrays. Analog IMC using RRAM crossbar arrays has been utilized to demonstrate four orders of magnitude improvement in energy over CPUs [240]. IMT nano-oscillator-based Ising solvers may exhibit even lower energy consumption, primarily due to the ultra-low power dissipation of the IMT oscillators. The IMT oscillator network implements bi-directional



ferromagnetic and antiferromagnetic coupling using simple electrical elements such as resistance and capacitance, and can achieve highly parallelized all-to-all connectivity. The analog or continuous-time dynamics of these Ising solvers has an inherent advantage of parallelism which lowers the time to solution compared to CMOS annealers and CPUs operating in discrete time. The time-to-solution (or cycles-to-solutions) remains similar for both the IMT solver and RRAM-based hardware accelerator.

11.3. Advances in science and technology to meet challenges

Table 2 summarizes some of the current capabilities, strengths, and challenges to the highlighted emerging hardware approaches for optimization. In this section, we consider challenges shared by all approaches and the advances needed to address them.

At the lowest level, there is substantial room to improve both the devices used in the solution state (or neuron) representation and in the connections between them. Whether a magnetic, electronic, or optical medium is used, these provide the core computational elements. Desired properties include multi-bit levels (many problem beyond max-cut are no longer binary), high endurance and robustness, low variability, and rapid re-programmability. Material engineering targeting these properties is needed, and ideally kept compatible with today's CMOS technology for future integration and mass production.

To address important applications of the future, all hardware solutions will need to be able to scale to many thousands of variables and constraints. For example, railway companies operate tens of thousands of trains per week for which optimal crew schedules are desired [245]. Supporting this necessitates breaking large problems across many processing units that must communicate at low latency, low energy, and high bandwidth. While the efficiency of the processing sub-units may be very high, this communication quickly becomes the bottleneck and Amdahl's law limits the benefits of the core processing circuits if there is a high communication overhead. Related to this challenge is the inherent connectivity for the processing units. With limited connections between the nodes (e.g., only nearest or next-nearest neighbours), 'embeddings' are used to solve problems with higher density connections, and these come with exponential penalties [246]. Maintaining efficient connectivity for large problems that must span many processing units will challenge all of today's emerging approaches. Indeed, the massive fan-out capabilities in biological nervous systems averaging >10 000 connections between neurons shows how nature has addressed this problem and future neuromorphic systems tackling optimization may need to mimic these designs.

Table 2.	State-of-the-art and	l opportunities for	emerging hardware	approaches to	optimization.
----------	----------------------	---------------------	-------------------	---------------	---------------

	Digital CMOS annealer	Optical Ising machine	Analog resistive crossbar	Coupled oscillators	Probabilistic logic bits (P bit)	Quantum annealing	
Bit precision	Up to 64 bit	Binary (>5 bit with integ. photonics)	Analog, or >5 bit	Binary	Binary	Analog	
Problem re- programmability	Easy	Hard	Moderate	Hard	N/A	Moderate	
Dense problem sizes demonstrated	8192	2000	60	8	6	60	
Performant architecture developed for large-scale	Yes	No	No	No	No	No	
Room temperature operation	Yes	Yes	Yes	Yes	Yes	No	
Strengths	Mature technology, flexibility	Low latency, high speed, wavelength multiplexing capability	High density, continuous value and time capability	Low energy, continuous time capability	Inherent stochasticity	Leverage quantum resources, sophisticated control, software, and tools developed Embeddings for dense connectivity	
Main challenges short term	Noise injection	Power consumption of control system, photo-detectors, nonlinear operations	Device yield, repeatability, analog device drift	Frequency variability and drift	Integrating magnetic materials; multibit		
Main challenges long term	Interconnect and comm. Bottleneck, CMOS scaling, 3D integration	Component sizes, silicon integration, non-volatile storage of weights	Fully passive arrays, interconnect and parasitics, digital conversion	Continuous time stability, interconnect and parasitics	Only stochastic neurons, requires architectural integration with another approach	Cryogenic requirement, scaling of superconducting integrated circuits	

With increasing scale and fan-out there arises the inevitable challenge of significant device parasitics and variability. Non-idealities include interconnect/wire parasitics in terms of line-to-ground capacitance, line-to-line capacitance and frequency variability for the oscillator approaches. With increasing problem size and the concurrent increase in the size of the network, it will be increasingly difficult to find the globally optimal solution. The reduction in success probability can be mitigated by increasing the number of anneal cycles and/or executing larger trial batches, but only at the expense of time-to-solution. An alternate approach could be to exploit emerging monolithic three-dimensional integration technology that provides multiple tiers of interconnect that can be dynamically configured to provide an efficient, scalable and dense network on chip. This promising direction will provide new architectural opportunities for on-chip implementation of large dense networks with programmable connections that are beyond the capabilities of existing process and packaging technologies today.

We stress that optimization problems are highly diverse, and even within a problem category (e.g., scheduling) specific instances can have different traits and levels of difficulty, such as the characteristic scale of barriers between minima, the density of saddle points, or the relative closeness in value between local and global minima. Consequently, domain experts have developed techniques highly tailored to their problem class. This could entail parameter choices such as using different noise distributions or cooling schedules (simulated annealing), to algorithmic variations such as ensembles of models exchanging temperatures (parallel tempering) or populations exchanging and mutating characteristics (genetic algorithms). Thus, it is desired for any emerging hardware to support these rich variations as much as possible, exposing internal parameters to the user for control, as well as provisioning the architecture to efficiently realize the more promising algorithmic variations. Many optimization problems may also involve substantial pre- and post-processing computations. For example, transforming a practical airline crew scheduling problem into the prototypical NP-hard 'setcover' problem first involves constructing sub-sets from viable rotations. Such pre- and postprocessing, let alone midstream processing (replica exchange in parallel tempering), requires flexible and complex architectures that include traditional digital units in addition to neuromorphic and physics-based optimization solvers.

The above challenges highlight the need for hardware designs to be algorithm and 'software aware'. Equally important is the development of algorithms and tools that are strongly 'hardware aware'. These must be designed to exploit the strengths of the underlying processing units—such as cheap stochasticity or certain

types of parallelism [247]—while, simultaneously mitigating their respective weaknesses—such as reduced precision. Thus, constructing successful systems for optimization solving will require a deep co-design from materials, devices, and packaging, all the way up to algorithms and software tools.

11.4. Concluding remarks

This perspective has highlighted the promise in leveraging physics- and brain-inspired principles to tackle today's most intractable computational problems. There are future challenges spanning all levels of the computing stack. Highly diverse approaches are being explored leveraging electronic, magnetic, optical, or quantum systems. Ultimately, we expect the winning approach will be the one that gathers an excited community of users by building flexible, performant, and reliable optimization solvers, and thus begin the virtuous relationships between users, software, and hardware researchers.

Acknowledgements

The authors acknowledge helpful input for figure and table content from Thomas Van Vaerenbergh and Suhas Kumar.

12. Enabling technologies for the future heterogeneous neuromorphic accelerators

Elisa Vianello and Alexandre Valentian

Université Grenoble Alpes, France

12.1. Status

AI and in particular ANNs have demonstrated amazing results in a wide range of pattern recognition tasks including machine vision, natural language processing, and speech recognition. ANN hardware accelerators place significant demands on both storage and computation. Today's computing architectures cannot efficiently handle AI tasks: the energy costs of transferring data between memory and processor at the highest possible rates are unsustainably high. As a result, the development of radically different chip architectures and device technologies is fundamental to bring AI to power-constrained applications, such as data of things, combining edge analytics with IOT.


12.2. Current and future challenges

Most of the AI applications today are still running in the Cloud, i.e. in data centers which offer large storage capacity and processing power [5]. The learning phase, for which you need large datasets, is done in the Cloud, and inference tasks are performed on the same assets. Such a scheme is not sustainable in the long run: with the increasing demand for intelligent devices, the data centers will not be able to sustain the load. Part of it will have to be offloaded to the devices themselves. This is the current challenge that companies and research teams are taking on: to enable running inference tasks at the edge, thanks to dedicated hardware accelerators. The main pathway for the implementation of such constrained systems is the reduction of the power consumption. This also relates to longer battery life, and to heat dissipation, which cannot be afforded. Since the complexity of neural networks tends to grow over time, this problem will only become more acute. In data-centric application such as AI, the main source of power consumption is data movement [248]. It costs hundreds of times more energy to move data to/from an external memory than to compute on it. One of the solutions is to handle lightweight data, which is why quantization of weights and activations is a field of active research, down to the extreme of binary neural networks. Second, the integration of dense memories, as close as possible to the processing engines (PE) exploiting new memory technologies and 3D integration schemes is mandatory. Some promising solutions are presented in the next section. The ultimate evolution consists in removing the PEs altogether and computing directly inside the memory: this is the in-memory-compute paradigm.

Having accelerators dedicated to inference tasks at the edge is only the first step. The **future challenge** will be to perform the learning phase locally as well. Autonomous agents will need to learn from experience, in order to adapt to their environment, learn basic principles and rules, infer common sense. They will thus have to exhibit lifelong learning abilities. Such accelerators will have additional challenges, linked to the learning phase itself. For instance, the state-of-the-art method to train ANN is the back-propagation algorithm, which minimizes a given loss function based on gradient descent [249]. However, its implementation on GPUs is energy-consuming and it does not satisfy edge requirements.

12.3. Advances in science and technology to meet challenges

Artificial intelligent (AI) systems are data-hungry: first, large amounts of memory are required for storing network states and parameters; second, the energy cost associated to data transfer between the memory and the processor is the major source of energy dissipation. Resistive memory devices, also referred as memristors, can provide massive on-chip data storage with low voltage and lowlatency accessibility. Their basic working principle relies on the modification of the material at atomic level causing a change of resistance. These memories include resistive-switching random access memory (RRAM), PCM, magnetic random-access memory (MRAM), and the ferroelectric random access memory (FeRAM). They are currently implemented as a 1T1R structure, i.e. with one MOS transistor (1T) used for accessing one resistor (1R). The memory cell footprint is around $40F^2$ and is limited by the access transistor. Promising results have been recently demonstrated to increase the memory density: using one transistor to access multiple RRAMs (1T4R) [250], stacking multiple 1T1R thanks to monolithic 3D technology [251], replacing the MOS transistor by a stacked nanowire transistor [252] or by a backend selector [253]. This last technology option enables efficient crossbar arrays with low leakage currents, leading to the highest density, with a footprint of 4F². Another promising approach to achieve high density is to increase the number of resistance levels in a single cell, for storing multiple bits. RRAM arrays storing up to three bits were demonstrated [251, 254]. Those multiple level cells not only increase the density, they are also a great mean for efficiently implementing in-memory-compute functions: provided that the resistance levels are linearly allocated, MVMs and accumulation (MAC) of inputs can be done in an analog manner, simply exploiting Ohm's and Kirchhoff's laws [74, 255].



However, on-chip learning based on back-propagation remains a challenge on memristor-based hardware, because of non-ideal device properties: limited endurance, non-linear conductance modulation, as well as device variability [256]. Moreover, back-propagation does not allow continuous, incremental learning. This is tackled at technology level, but also at algorithm level. To achieve unsupervised learning, there have been multiple efforts implementing biologically plausible spike timing dependent plasticity (STDP)-variants and Hebbian learning using neuromorphic processors [257, 258]. However, they are local learning rules and they do not provide any guarantee that network performance will improve in multilayer or recurrent networks. Novel algorithms where both inference and learning could fully be achieved out of core physics have been recently proposed such as, threefactor spike-based learning rules [259, 260], direct-feedback alignment [261], equilibrium propagation (*Eq-prop*) [262]. Next, these algorithms have to be mapped on real hardware tacking into account device non-idealities. Some recent works explored this direction. For instance, it has been demonstrated that drift behavior of PCM devices (generally considered a device non-ideality) can be exploited to implement long lasting eligibility traces, a critical ingredient of three-factor learning rules [263]. A machine learning scheme that leverages the RRAM variability to implement Markov chain Monte Carlo (MCMC) sampling techniques to enable on-chip learning in Bayesian neural network has been demonstrated [264].

As can be imagined, many actors will provide accelerators for running AI workloads at the edge. They will use a combination of technologies, and different coding strategies: classical coding (FP32, BF16, INT8), spike coding (rate, temporal, time-to-first-spike), a combination of both. Since there is never one-size-fits- all solution, it is strongly believed that heterogeneous AI systems will play an important role for delivering products tailored to different application needs, as illustrated in figure 18. Such systems will leverage IP reuse thanks to 2.5D and 3D integration technologies, thus reducing nonrecurring engineering costs and time to market. Those technologies enable the integration of different chips, also known as chiplet [265], built with different manufacturing processes. Those chiplets are co-integrated on silicon or glass substrates, or an organic interposers. Communication with the outside world is ensured thanks to through-silicon-vias or through-mold-vias and redistribution layers.

For fulfilling that vision, issues such as interoperability and application mapping need to be addressed. A high-level AI system description language will be specified, for expressing the respective hardware resources

and their interactions. Deep learning frameworks will be complemented with export capabilities on those heterogeneous platforms.

12.4. Concluding remarks

AI and in particular machine learning have made tremendous progress in image, video, sound and speech recognition tasks. The next challenge is enable unsupervised and continuous learning at the edge. Innovations ranging from single memory devices to full-scale architectures exploiting heterogeneous integration are required to meet the computational needs of such applications.

Acknowledgements

We acknowledge funding support from the H2020 MeM-Scales project (871371), the ECSEL TEMPO project (826655) and the ECSEL ANDANTE project (876925).

13. Photonics

Johannes Feldmann¹, Xuan Li¹, Wolfram H P Pernice^{2,3} and Harish Bhaskaran¹

¹University of Oxford, United Kingdom ²University of Münster, Germany

³University of Münster, Germany

13.1. Status

The field of optical computing began with the development of the laser in 1960 and has since been followed by many inventions especially from the 1980 s demonstrating optical pattern recognition and optical Fouriertransform processing [266]. Although these optical processors never evolved to commercial products due to a limited application space and the high competition with emerging electronic computers, photonic computing again gained much interest in recent years to overcome the bottlenecks of electronic computing in the field of AI, where large datasets must be processed energy efficiently and at high speeds [267]. Optical computers are able to seriously challenge electronic implementations in these domains, particularly in throughput. Photonics further has allowed one to integrate optics on-chip enabling such optical neuromorphic processors to have several advantages compared to their electronic counterparts. One of them is based on the fact that photons are bosons and are able to occupy the same physical location (i.e. not subject to the Pauli exclusion principle). Thus, many can be transmitted through the same channel without mutual interference. This offers an intrinsically high degree of parallelization by wavelength and mode multiplexing techniques, enabling the use of the same physical processor to carry out multiple operations in parallel leading to high computing densities. Additionally, the data transport problem that is apparent in electronics at high signal speeds is easily addressed using photonic waveguides that serve as low power data links. Taken together with the fact that linear operations can be implemented in the optical domain with very high energy efficiency [268], photonics offers a promising platform for high speed and highly parallelised neuromorphic computing [269].

Many non-von Neumann photonic computing techniques have been demonstrated using integrated, fibrebased and free-space optics [268], showing a large variety of different approaches ranging from coherent neural networks [270], RC [271] and phase-change photonics [272–274] to hardware accelerators for the main computational bottlenecks (usually matrix multiplications) in conventional AI solutions [43, 275]. Most of these are analogous to IMC that has most prominently been developed by IBM [276, 277]. Further advances in photonic computing might first lead to optical co-processors that accelerate specific operations such as VMMs and are implemented together with conventional electronic processors. The next step could be photonic neuromorphic computers avoiding electro-optic conversions (figure 19).


13.2. Current and future challenges

The current main limitations of neuromorphic photonics lie in scalability, system stability and interfacing with electronic systems. As most approaches that rely on analogue computing, the precision of photonic approaches depends on the noise accumulated in the processor.

Whereas free-space implementations are often bulky and therefore difficult to scale, the challenges for integrated circuits are optical loss and the reliable fabrication of the individual components. Especially tuning the fabrication of resonant elements as ring resonators that are used for fast signal modulation and multiplexing require thermal tuning in order to operate properly. To gain from wavelength multiplexing and achieve highly parallel processing, the fabrication techniques need to provide sufficient reproducibility of the wavelength channels of the multiplexers.

In addition to fabrication challenges, many neuromorphic photonic processors rely on coherent light and require precise control of the optical phase. Because the phase is strongly temperature dependent, excellent thermal stability of the system is necessary, which in itself can become a significant component of the overall power consumption of the system. Another challenge for integrated photonic circuits is the integration of light sources on CMOS compatible silicon platforms. Whereas linear operations can be carried out in photonic circuits intrinsically very well, all-optical nonlinear elements working at sufficiently low power and high speeds are more challenging to develop. However, from a strictly neuromorphic computing standpoint this is crucial, as non-linear elements determine spiking functionality in neurosynaptic networks, at least at the device level.


From the overall system architecture standpoint for neuromorphic photonic computing, additional challenges regarding integration with conventional technologies exist. Given that most information processing is performed in the electronic domain, photonic neuromorphic systems usually have to be interfaced with electronic processors, e.g. to input the data or store the results of the computation. This imposes at least two electro-optic conversions, which decrease the overall energy efficiency of the system. As most neuromorphic processors are a type of analogue computing, additionally conversions between the digital and analogue



domain have to be performed. Especially analogue to digital converters (ADC) can make up a huge part of the power budget and scale badly in terms of energy with the number of bits and operation speed [278]. To be able to use the high modulation speeds accessible in modulating and detecting optical signals, significant improvements in digitizing the results of the computation have to be made.

13.3. Advances in science and technology to meet challenges

Like their electronic counterparts, photonic neuromorphic processors require precise fabrication. With increased interest in photonic information processing, several photonic foundry services are emerging and are continuing to enhance their capabilities. These will be crucial to progress this field. Although vastly improved



with good performance, the existing capabilities in photonic foundries are well behind those that exist in the more mature and established electronics. Improvements in fabrication techniques will give way to less variation in device specifications, e.g. in the wavelength specification of certain components like resonators or multiplexers and reduction of optical loss. This will be important to improve on the parameters that make photonic neuromorphic processors more advantageous, specifically the ability to wavelength multiplex. A useful tool in the fabrication process could be an additional tuning step after fabrication, to match the designed specifications such as measuring the resonance wavelength of a resonator and adjusting it to the desired wavelength as a post-processing correction. Advances in the standard components as modulators and detectors as well as the addition of new components to the libraries of photonic foundries and the development of new materials for non-volatile optical storage will enhance this field and bring these circuits closer to commercialization.

Yet another crucial component is efficient light sources that can be integrated on photonics alongside reliable many-channel on-chip multiplexers. Integrated optical frequency combs that provide a wide optical spectrum with a fixed channel spacing that can be exploited for computing as a coherent light source are a prime example of this [279]. Photonic neuromorphic circuits rely on electronic control and therefore improvements in high-speed electronic components such as digital-to analogue (DAC) and analogue-to digital converters (ADC) are also very important. Further research could also lead to alloptical DACs and ADCs circumventing the need for electro-optic conversions. In general, photonic neuromorphic processors that minimize conversions between the digital and analogue domains are preferable. A specific class of neural networks that could prove especially suitable for low power photonic processing are SNN, that reduce digital-to-analogue conversions by using binary spikes and their time dependence as information carriers.

As the non-linear optical coefficients for silicon are small, functional materials that allow for such nonlinearity or other added functionality are also important [280]. A promising class of materials are phasechange materials (PCMs) that switch their optical properties upon excitation and therefore effectively resemble a nonlinear element [281, 282]. Although PCMs can be switched with low optical powers, significant improvements have to be made in increasing the switching speed in order to keep up with high modulation speeds enabled by photonics. Another class of materials considered for low power optical non-linearities are epsilon-near-zero materials [283].

Operating with analogue signals results in a higher sensitivity to noise; recent advances in reducing the precision of neural networks to lower numbers of bits with low loss in prediction accuracy is one step to overcome this challenge [284] and further research in this area is also required.

As photonic integrated circuits become more and more complex, similar to electronics, a threedimensional implementation seems necessary to avoid crosstalk and loss when routing the signals and avoid waveguide crossings, which also requires investigation (figure 20).

13.4. Concluding remarks

Photonic computing has been a promising field of research over the last decade as photonics matures. New challenges in handling huge amounts of data in the fields of AI and machine learning that bring conventional electronic processors to their limits have resulted in a surge interest in photonic computing. This is because of their inherent advantage that enables high throughput, a high degree of parallelization together with the

ability to carry out linear operations at very low energies. This makes photonic neuromorphic processors a very promising route to tackle the upcoming challenges in AI applications.

In spite of the challenges photonic computing concepts that can overcome the limitations of electronic processors have been demonstrated in the recent years, and a roadmap to address their march into commercialization would be a huge benefit to society.

Acknowledgements

This research was supported by EPSRC via Grants EP/J018694/1, EP/M015173/1, and EP/M015130/1 in the United Kingdom and the Deutsche Forschungsgemeinschaft (DFG) Grant PE 1832/5-1 in Germany. WHPP gratefully acknowledges support by the European Research Council through Grant 724707. We further acknowledge funding for this work from the European Union's Horizon 2020 Research and Innovation Program (Fun-COMP project, #780848).

14. Large-scale neuromorphic computing platforms

Steve Furber

The University of Manchester, United Kingdom

14.1. Status

The last decade has seen the development of a number of large-scale neuromorphic computing platforms. Notable among these are the SpiNNaker [204] and BrainScaleS [285] systems, developed prior to, but supported under the auspices of, the EU Flagship Human Brain Project, and somewhat later the Intel Loihi [202] system. These systems all have largescale implementations and widespread user communities.

All three systems are based upon conventional CMOS technology but with different architectural approaches. SpiNNaker uses a large array of conventional small, embedded processors connected through a bespoke packet-switched fabric designed to support largescale SNN in biological real time and optimised for brain modelling applications. BrainScaleS uses above threshold analogue circuits to model neurons running 10 000 times faster than biology, implemented on a wafer-scale substrate, optimised for experiments involving accelerated learning. Loihi sits somewhere between these two, using a large array of asynchronous digital hardware engines for modelling and generally running somewhat faster than biological real time, with the primary purpose of accelerating research to enable the commercial adaptation of future neuromorphic technology.

In order to support their respective user communities these systems have extensive software stacks, allowing users to describe their models in a high-level neural modelling language such as PyNN [286] (used for both SpiNNaker and BrainScaleS) so that straightforward applications can be developed without a detailed understanding of the underlying hardware.

These large-scale systems have been up and running reliably for some time, supporting large user communities, and offer readily accessible platforms for experiments in neuromorphic computing. Access to neuromorphic technology is no longer a limiting factor for those who wish to explore its potential and capabilities, including using these existing platforms to model future technologies (figure 21).

14.2. Current and future challenges

The major challenge for neuromorphic computing technology, if it is to gain widespread adoption, is to offer a compelling demonstration of its commercial viability. It must demonstrate some significant advantage over competing technologies in terms of its capabilities, performance and/or energy-efficiency. The applicability of neuromorphic technologies to brain science is easier to justify, though even there it is competing with more conventional computing technologies such as GPUs [287] and HPC [288], and clear wins [289] are hard to come by. But brain science is a niche research application, and something with much broader applicability is required.

Although there have been a few demonstrations of neuromorphic superiority in application domains such as keyword recognition [290] these are still rather thin on the ground, but there is growing confidence that something will emerge in the near future. This confidence is underlined by the growing number of venture capital backed start-up companies in the neuromorphic arena. Most of these are focussed on small-scale AI applications in edge computing, though some such as SpiNNcloud Systems GmbH have an eye on data centre scale neuromorphic technology, in this case based on the second-generation SpiNNaker2 chip [203].

Large-scale systems of spiking neurons amplify many of the challenges faced by smallerscale systems in areas such as the design of homeostatic mechanisms that control the operating point of a neural population, adaptation mechanisms that compensate for local failures, unsupervised learning mechanisms that allow the system to focus resources around the statistical distribution of observed inputs, synaptogenesis, neurogenesis, on-line learning, and so on. But as the system scales up, the need for built-in automatic adaptation only increases as the time and cost of optimising by hand increases rapidly beyond reasonable bounds.


A challenge for the future will be to combine the scalability demonstrated by the current large-scale neuromorphic systems with the enhanced characteristics offered by the advanced device technologies described elsewhere in this review.

14.3. Advances in science and technology to meet challenges

Given the challenges that face neuromorphic technologies at all scales, where should we look for inspiration for the solutions to those challenges? There are two obvious places to seek such inspiration: biology, and mainstream AI.

Biological brains, from small insects up to human scale, demonstrate the flexibility, adaptability and energyefficiency that we aspire to for our engineered systems, if only we could understand how they work. This is





why neuromorphic engineers track advances in brain science, exploring the potential of advances such as new insights into dendritic computation to improve the capabilities of engineered systems.

Similarly, the explosion over the last decade of applications of AI based upon ANN offers insights into the effective organisation of neurons, whether spiking or not. There is a strong sense that the success of ANN must be telling us something about how brains work, despite the absence of evidence in biology for, for example, the error backpropagation learning mechanism that is so effective in ANN. Some form of gradient descent (the principle underlying backprop) must be at work in biological learning, and recent developments in algorithms such as e-prop [260] offer a glimpse of how that could work.

The prospect of the convergence of neuromorphic engineering with brain science and mainstream AI is tantalising for all three branches of science/engineering.

14.4. Concluding remarks

Large-scale neuromorphic computing platforms offer reliable, well-supported services that can be used to research neuromorphic technologies at the system level at minimal cost. They are accompanied by software stacks that abstract users away from low-level hardware details, allowing them to focus on higher-level issues such as network architectures, learning rules, brain modelling, etc. As such they play an important role in training the next generation in neuromorphic technology and in providing flexible development platforms for exploring neuromorphic system development without the constraints imposed by smaller research prototypes.

Acknowledgements

The design and construction of the SpiNNaker machine was supported by EPSRC (the UK Engineering and Physical Sciences Research Council) under Grants EP/D07908X/1 and EP/G015740/1. Ongoing development of the software is supported by the EU ICT Flagship Human Brain Project (FP7-604102, H2020-720270, H2020-785907, H2020-945539).

Section 3. Neuromorphic algorithms

15. Learning in spiking neural networks

Emre Neftci

Peter Grünberg Institute - Neuromorphic Software Ecosystems Forschungszentrum Jülich, Germany Faculty of Electrical Engineering and Information Technology RWTH Aachen, Germany University of California, Irvine

15.1. Status


The dynamical nature of SNN circuits and their spatiotemporal sparsity supported by asynchronous technologies makes them particularly promising for fast and efficient processing of dynamical signals (section 10). Here, we discuss learning in SNNs, which refers to the tuning of their states and parameters to learn new behaviors, achieve homeostasis and other basic computations. In biology, this is achieved via local plasticity mechanisms that operate at various spatial and temporal scales. While several neural and synaptic plasticity rules investigated in neurosciences have been implemented in neuromorphic hardware [307], recent work has shown that many of these rules can be captured through three factor rules (3F) of the type [259, 299, 307]

$$\delta W = F_{\rm pre} F_{\rm post} M_{\rm post}$$

where factors F_{pre} and F_{post} correspond to functions over presynaptic and postsynaptic states, respectively, and the factor M_{post} is a post-synaptic modulation term (see also section 18 for a specific example). The modulation is a task-dependent function, which can for example represent error in supervised learning task, surprise in an unsupervised learning task, or reward in reinforcement learning. Given the generality of the three-factor rule in representing existing learning rules and paradigms, this section focuses on the requirements for implementing 3F plasticity in neuromorphic hardware. By analogy to the brain, the most intuitive implementation of synaptic plasticity is on-chip, i.e. plasticity is achieved at or the synapse circuit or equivalent circuit near the SNN (figure 22, top). Neuromorphic engineers have extensively implemented learning dynamics derived from computational neurosciences, such as STDP variants [3, 291, 310] and more recently, 3F rules [296]. On-chip learning requires precious memory and routing resources [309], which hinders scalability. On digital technologies, this problem can be sidestepped by time-multiplexing a dedicated local plasticity processor [295, 301]. The timemultiplexing approach however suffers from the same caveats as a von Neumann computer due to the separation between the SNN and the associated plasticity processor. Other promising alternatives for true local plasticity are emerging devices (section 1) and related architectures (section 9), which allow storage and computation for plasticity to occur at the same place.

A more practical approach to learning in SNNs is off-chip (figure 22, bottom), which relies on a separate general purpose computer to train a model of the SNN, where memory and computational are potentially more abundant. In this approach, once the SNN is trained, the parameters are then mapped to the hardware. Provided a suitable model of the hardware substrate or a method to convert parameters from a conventional network to an SNN, off-chip learning generally achieves the best inference accuracy on practical tasks [298, 315]. Heterogenous approaches combining on-chip and off-chip approaches (also called chip-in-the-loop, figure 22, middle) have been successful at smaller scales [300], although scalability there remains hindered by the access to the local states necessary for plasticity in the chip. The suitability of onchip or off-chip learning is highly dependent on the task. The former is best for continual learning (section 18) and the latter is best when a large dataset is already available and the SNN model and parameter conversion are near-exact. If the model is not exact, hybrid learning is often the most suitable method. Onchip and hybrid learning also have the advantage that learning can occur online, i.e. during task performance.

Although Hebbian STDP variants have been instrumental for modeling in neuroscience, mathematically rigorous rules derived from task objectives such as 3F rules have a clear upper hand in terms of practical performance [304, 315, 322]. This is arguably because some forms of spatial credit assignment are necessary to learn in non-shallow networks [292]. Thus, we anticipate that mathematically motivated (top-down driven) rules grounded in neuroscience are likely to drive the majority of future research in SNN learning and their neuromorphic implementation. Already today, the success of top-down modeling of learning to efficiently train SNNs ushered in a new wave of inspiration from machine learning (ML) [313], and accelerated the quest to build neuromorphic learning machines. In the following, we focus on specific challenges of 3F learning approaches.



are mapped on to the neuromorphic device (blue). In the chip-inthe-loop approach, updates are computed partially off-chip, but using states recorded from the chip. In the on-chip implementation, the updates are computed locally to the neurons or synapses. While brainlike, efficient, continual learning can only be achieved using on-chip learning, off-chip approaches also play an important role in pre-training the model and prototyping new algorithms, circuits and devices, or when learning is not necessary (fast SNN inference).

15.2. Current and future challenges

While the technological advantages of spike-based sensing are now increasingly evident (sections 2.2 and 4), it is natural to question the advantages of SNNs and their related learning algorithms. In search of an answer, it is fruitful to consider some outstanding challenges of state-of-the-art AI technologies, which are largely based on deep learning: the mismatch between the training dataset and the real world (generalizability), and their high power requirements [248]. Learning online can overcome the former by learning continuously during task performance, so as to correct errors at the moment when they occur (provided issues related to sequential learning, e.g. catastrophic interference, are avoided [305]). On-chip learning on the other hand lends itself to distributing the plasticity processes across the hardware substrate, where it occurs, as is the case for memristor crossbar arrays [21]. Learning then becomes a local process, which can reduce power requirements, provided the necessary learning factors of the learning rule can be computed efficiently. This is possible for example by communicating errors using spikes [308, 318].

Because the challenges of on-chip learning form a superset of those faced by off-chip and hybrid learning strategies, and can solve key problems in state-of-the-art AI, we focus our challenge description on the on-chip case and the 3F rule. SNN learning requires the computation and storage of suitable learning factors, which is dictated by the requirements of space-time credit assignment, i.e. how blame or credit should be assigned given the performance on a given task. The form of the 3F rule described above is compatible with this requirement by virtue of the third factor M_{post} , which modulates the learning rule according to a (potentially global) learning signal. Phenomenological rules hypothesize the nature of these factors from neuroscientific evidence. However, by taking inspiration from ML, these terms can be derived from first principles, such as gradient descent on suitable loss functions. For example, in the gradient backpropagation (BP) rule, a form of gradient descent applied to differentiable networks, M_{post} can be identified with the backpropagated errors. However, as BP requires non-local signals, it is difficult to implement on a physical substrate and it is inconsistent with the brain [292]. As such, BP has significant scaling problems [302]. This could be circumvented if the errors can be estimated using local signals [306]. Thus, achieving local forms of BP is an active field of research that is now steadily advancing towards viable BP alternatives in SNNs [294, 304, 322]. Interestingly, the terms $F_{\rm pre}$ and F_{post} of the 3F rules correspond to neuronand synapse-specific traces of postsynaptic and presynaptic activity, respectively. An important challenge is the computation and storage of synaptic-specific traces locally, which are distinct from weight and parameter storage. Emerging volatile devices are central to this effort [296, 319]. Among the various nonidealities of emerging materials and devices, device-to-device variability remains a major unsolved problem for successful learning in SNNs [311]. As no device will be perfect, the community is now striving to co-design neuron, synapse, plasticity and algorithms that are tailored to the device and the application [307].

15.3. Advances in science and technology to meet challenges

State-of-the-art learning algorithms are specifically designed for conventional von Neumann computers. While relying on deep learning technologies for SNN learning via off-chip learning will help reap early results, real

breakthroughs will take place when software and algorithms are designed specifically for the constraints and dynamics of neuromorphic hardware. This involves moving beyond the concept of mapping conventional neural network architectures and operations to SNNs, and instead modeling and exploiting the computational properties of biological neurons and circuits. Beyond advances in the emerging devices themselves (section 1), one key enabler of such breakthroughs will be a differentiable programming library (e.g. tensorflow) operating at the level of spike-events and temporal dynamics that facilitates the scalable composition and tracing of operations [293]. Such a framework can in turn facilitate the computation of the necessary learning factors. While recent work demonstrated SNN learning with ML frameworks [314, 316], the mapping of the 3F computations on dedicated hardware does not yet exist. This is due to a lack of applications and the more stringent requirements for learning. Additionally, current technologies are not optimized to training large-scale SNNs, which remains today very slow and memory-intensive due to the high complexity of the underlying dynamics and gradients [320, 321]. However, provided that SNN models capture key features of the brain, namely that average spike rate of neurons in the brain is at most 2 Hz and that connectivity is locally dense but globally sparse [297], specialized computers capable of sparse matrix operations can greatly accelerate offline training compared to conventional computers. This is because a neuron that does not spike does not elicit any additional computations or learning at the afferent neurons. Spurred by the hardware efficiency of binarized neural networks [312], some ML hardware now support efficient sparse operations [303] which could be exploited in SNN computations. A community-wide effort in these directions (software and general-purpose hardware) are likely to boost several research areas, including the discovery of new (spatial) credit assignment solutions, the identification and control of the distinctive dynamics of SNNs (multiple compartments, dendrites, feedback dynamics, reward circuits etc), and the evaluation of new materials and devices, all in the light of community-accepted benchmarks. Undertaking such device evaluations prior to the design and fabrication cycle, for instance via a suitable surrogate model of the device, can save precious resources and dramatically accelerate the development of emerging devices.

The ability to cross-compile models in a software library can blur the line between hardware and software. This resonates well with the idea of on-chip and off-chip learning working in concert. That approach is attractive because the difficulties of online learning can be mitigated in hardware with multiple stages of training, for example by first training offline and then fine-tuning online [317]. Furthermore, fewer learning cycles entail fewer perturbations of the network, and thus mitigating the problems of sequential learning. At the same time, learning is achieved after a much smaller number of observations (e.g. few-shot learning), which is essential in continual learning tasks (section 3.4). The success of such *meta*-learning hinges on a good task set definition and is compute- and memory-intensive. Once again, general-purpose computers supporting sparse matrix operations, associated ML libraries and community-wide efforts are essential to achieve this at scale. Although ML is not the only approach for SNN learning, the tools developed to enable ML-style learning algorithms are central to other learning models and approaches. These include hyperdimensional computing, variational inference algorithms, and neural Monte Carlo sampling, all of which rely on well-controlled models and stochasticity that can be supported by such tools.

15.4. Concluding remarks

Beyond hardware developments, the hidden figure of deep learning's success has been software libraries. It enabled the composition and differentiation of hardware-optimized routines for inference and learning. Neuromorphic computing, and more specifically learning in neuromorphic hardware is arguably at a similar standpoint today. Neuromorphic computing entails a different set of constraints (distributed, sparsity, locality) that are determined by the physics of the underlying hardware. The next step in exploring this new set of constraints requires a software library tailored to SNN dynamics and sparsity capable processors, which in turn can greatly accelerate the discovery of new materials, architectures, and learning algorithms for SNNs.

Acknowledgements

This work was supported by the National Science Foundation under Grant 1652159 and 1823366.

16. Learning-to-learn for neuromorphic hardware

Franz Scherr and Wolfgang Maass

Institute of Theoretical Computer Science, TU Graz

16.1. Status

An important goal for neuromorphic hardware is to support fast on-chip learning in the hand of a user. Two problems need to be solved for that:

- (a) A sufficiently powerful learning method has to run on the chip, such as stochastic gradient descent.
- (b) This on-chip learning needs to converge fast, requiring ideally just a single example (oneshot learning).

Evolution has found methods that enable brains to learn a new class from a single or very few examples. For instance, we can recognize a new face in many orientations, scales, and lighting conditions after seeing it just once, or at least after seeing it a few times. But this fast learning is supported by a long series of prior optimization processes of the neural networks in the brain during evolution, development, and prior learning. In addition, insight from cognitive science suggests that the learning and generalization capability of our brains is supported by innate knowledge, e.g. about basic properties of objects, 3D space, and physics. Hence, in contrast to most prior on-chip learning experiments in neuromorphic engineering, neural networks in the brain do not start from a tabula rasa state when they learn something new.

Learning from few examples has already been addressed in modern machine learning and AI [323]. Of particular interest for neuromorphic applications are methods that enable recurrently connected neural networks RNNs to learn from single or few examples. RNNs are usually needed for online temporal processing—an application domain of particular interest for energy-efficient neuromorphic hardware. The gold standard for RNN-learning is backpropagation through time (BPTT). While BPTT is inherently an offline learning method that appears to be off-limit for online onchip learning, it has recently been shown that BPTT can typically be approximated quite well by computationally efficient online approximations. In particular, one can port the online broadcast alignment heuristic from feedforward to recurrent neural networks [324]. In addition, one can emulate the common LSTM (long short-term memory) units of RNNs in machine learning by neuromorphic hardware-friendly adapting spiking neurons. Finally, a computationally efficient online approximation of BPTT—called e-prop—exists that also works well for recurrent networks of spiking neurons (RSNNs) with such adapting neurons [260]. The resulting algorithm for on-chip training of the weights W_{ji} for neuron *i* to neuron *j* of an RSNN—for reducing some arbitrary but differentiable loss function E—takes there the form

$$\frac{\mathrm{d}E}{\mathrm{d}W_{ji}} = \sum_{t} L_{j}^{t} e_{ji}^{t}.$$

The so-called learning signal L_j^t at time, is some online available approximation to the derivative of the loss function *E* with regard to the spike output of neuron *j*, and the eligibility trace e_{ji}^t is an online and locally computable eligibility trace. While this would usually require even more training examples than BPTT, one can speed it up substantially by optimizing the learning signal L_j^t and the initial values of the synaptic weights W_{ji} to enable learning from few examples for a large—in general even infinitely large—family \mathcal{F} of on-chip learning tasks [325]. This can be achieved through learning-to-learn (L2L) [326]. A scheme for the application of L2L to enable fast on-chip learning is shown in figure 23.

16.2. Current and future challenges

The main choices that have to be made for such realization of fast on-chip learning are the choice of the family \mathcal{F} of tasks, the choice of the optimization method for offline priming through the definition of hyperparameters, and the choice of the hyperparameters. Options for the latter are for example:

- (a) Just the learning rate parameters of on-chip learning rules are hyperparameters.
- (b) Also the values of all synaptic weights of the RSNN are hyperparameters.
- (c) Only the initial values of the synaptic weights of the RSNN are hyperparameters.
- (d) In addition all parameters of an auxiliary NN—the learning signal generator—that generates online learning signals L^t for fast convergence of e-prop on the chip are hyperparameters.

Option 1 has been explored for RSNNs in [327, 328], and with an application to analog neuromorphic hardware in [329]. Option 2 is arguably the most commonly considered application of L2L in machine learning and computational neuroscience models [326, 330–334]. An attractive feature of this option for realizing fast on-chip learning is that it requires no synaptic plasticity for that. Rather, it uses hidden variables of the RNN




Figure 23. Scheme for the application of L2L for offline priming of a neuromorphic chip. Hyperparameters Θ of the RSNN on the chip are optimized for supporting fast learning of arbitrary tasks *C* from a family \mathcal{F} that captures learning challenges that may arise in the hands of a user. The resulting hyperparameters are then loaded onto the chip. Note that the desired generalization capability is here more demanding than usually: we want that the chip can also learn tasks *C* from the family \mathcal{F} very fast that did not occur during offline priming (but share structural properties with other tasks in the family \mathcal{F}).




Figure 24. (Left) Learning architecture for fast on-chip learning with e-prop. A learning signal generator produces online learning signals for fast on-chip learning. The weights of the learning signal generator as well as the initial weights of the learning network result from offline priming. (Right) Example application for fast learning. In this task C, the learning network has to learn the new command 'connect' from a single utterance, so that it recognizes it also from other speakers. The learning signal generator is activated when the new command is learnt (leftmost green segment).

for storing information from the few training examples that are needed for fast learning. In the case of machine learning, these hidden variables are the values of memory cells of LSTM units. In SNN these are the current values of firing thresholds of adapting neurons. An alternative is to choose only some synaptic weights to be hyperparameters, and to leave others open for fast on-chip learning [335]. Option 3 is used by the MAML approach of [336], where only very few updates of synaptic weights via BPTT are required in the inner loop of L2L. It also occurs in [325] in conjunction with option 4, see figure 24 for an illustration.

One common challenge that underlies the success of all mentioned options, is the efficacy of the training algorithm for the offline priming phase, the outer loop of L2L. While option 1 can often be carried out by gradient-free methods, the more demanding network optimizations of the other options tend to require BPTT for offline priming of the RNN.

16.3. Advances in science and technology to meet challenges

It is quite realistic to enable according to this L2L method fast on-chip learning on neuromorphic hardware. The most demanding aspect for the hardware is to be able to run the on-chip learning algorithm that is required. This can be implemented on most neuromorphic hardware if only simple local rules for synaptic plasticity are required in the inner loop of L2L, as in option 1. In the case of option 2 a spike-based neuromorphic hardware just needs to be able to emulate adapting spiking neurons. This can be done for example on SpiNNaker [337] and Intel's Loihi chip [202]. Using BPTT for on-chip learning appears to be currently infeasible, but on-chip learning with e-prop is supported by SpiNNaker and the next generation of Loihi. Then option 4 can be used for enabling more powerful fast on-chip learning. The only additional requirement for the hardware is that an offline primed learning signal generator can be downloaded onto the chip (once and for all), and that the chip supports communication of its learning signal for gating local synaptic plasticity rules according to eprop.

An illustration of a sample application which becomes then realistic is shown in figure 24: on-chip learning of a new spoken command from a single example in such a way that the same command can then also be recognized under different acoustic conditions and from different speakers.

Future advances need to address the challenge of training extended learning problems during the offline phase. Besides improved gradient-based algorithms, also gradient-free training methods such as evolution strategies [338] are attractive for that. In fact, since the latter paradigm allows to employ neuromorphic hard-ware directly for evaluating the learning performance, this approach can benefit from the speed and efficiency of fast neuromorphic devices, as in [329]. Particularly fast neuromorphic hardware such as Brainscales [339] might support then even more powerful offline priming with training algorithms that could not be carried out on GPU-based hardware, thereby providing the basis for superior hybrid systems.

16.4. Concluding remarks

Learning in neuromorphic hardware is likely to become split into two phases that each have different goals and require different learning methods: an extensive offline priming phase—either on the actual hardware or a software model for it—that optimizes selected hyperparameters but possibly also the network architecture for a large family of potential on-chip learning tasks in the hands of the user. Resulting hyperparameter values and network architectures will be downloaded onto the neuromorphic hardware before it gets into the hands of the user. The hardware is then primed so that remaining open parameters can be learnt on-chip from very few examples, possibly even just one example.

It is conceivable that this method can be expanded to provide another useful property for neuromorphic hardware in the hands of the user: that on-chip learning cannot bring the chip into an operating regime which is unsafe, or undesired for other reasons. It has already been verified that the outer loop of L2L can impose powerful priors for subsequent computing and learning of RSNNs [334].

Acknowledgements

This research/project was supported by the Human Brain Project (Grant Agreement Number 785907) of the European Union and a Grant from Intel.

17. Computational neuroscience

Srikanth Ramaswamy

Newcastle University

17.1. Status

Understanding the brain is probably the final frontier of modern science. Rising to this challenge can provide fundamental insights into what makes us human, develop new therapeutic treatments for brain disorders, and design revolutionary information and communication tools. Recent years have witnessed phenomenal strides in employing mathematical models, theoretical analyses and computer simulations to understand the multi-scale principles governing brain function and dysfunction—a field referred to as 'computational neuroscience' [340, 341]. Computational neuroscience aims at distilling the necessary properties and features of a biological system across multiple spatio-temporal scales—from membrane currents, firing properties, neuronal morphology, synaptic responses, structure and function of microcircuits and brain regions, to higher-order cognitive functions such as memory, learning and behavior. Computational models enable the formulation and testing of hypotheses, which can be validated by further experiments.

The multidisciplinary foundations of computational neuroscience can be broadly attributed to neurophysiology, and the interface of experimental psychology and computer science. The first school of thought, neurophysiology, is exemplified by the model of action potential initiation and propagation proposed by Hodgkin and Huxley [342] and theoretical models of neural population dynamics [229]. Whereas, the second school of thought, at the interface of experimental psychology and computer science focuses on information processing and learning, which could be traced back to models of ANN that were developed about half a century ago [343]. Computational neuroscience became its own nascent field about three decades ago and has rapidly evolved ever since [344].

In the early stages of its conception, computational neuroscience focused almost entirely on states of sensory processing, mainly due to the fact that studies of cognitive function were restricted to the domain of psychology, which was beyond what empirical neuroscience could offer. However, since then, rapid strides in tools and techniques have enabled tremendous advances in our knowledge of the neural mechanisms brain underlying cognitive states such learning and memory, reward and decision-making, arousal and attention [345–347]. Consequently, the dynamic field of neuroscience bestows many opportunities and challenges. A recent development is the symbiosis between computational neuroscience and deep learning [313]. Deep learning models have enabled efficient means to analyze vast amounts of data to catalyze computational modeling in brain research. However, the current framework of deep learning is mostly restricted to object recognition or language translation. Identifying the fundamental mechanisms responsible for the emergence of higher cognitive functions such as attention and decision making, appropriately recapitulated into algorithms by computational models, will influence the next generation of intelligent devices.

17.2. Current and future challenges

Thus far, computational neuroscience has been limited to modelling local circuits—through point neurons or multicompartmental models. However, the advent of 'big' spatial and temporal data, from single-cell transcriptomes, micro, meso and macro-scale connectomes, large-scale neurophysiology, and functional brain activity mapping, which is nothing short of an industrial revolution in neuroscience, are rapidly enabling computational models to incorporate greater biological realism and move beyond local circuits. The grand challenge, however, is to assimilate the exponentially increasing datasets into theories and models of global brain function that complement local circuit and large-scale network models. One possibility towards surmounting this challenge is through collaborative endeavors, which have pioneered a high-throughput, team science approach to neuronal modelling. These initiatives have brought together disparate disciplines such as experimental biology, computational neuroscience, high-performance computing and data science that attempt to model neural circuits with thousands of detailed multicompartmental neurons and millions of synapses and simulate these models on supercomputers. The first of these was probably the Blue Brain Project, followed by the Allen Institute for Brain Science, the US Brain Initiative, Japan Brain and Minds, the Human Brain Project, and the International Brain Laboratory among several others [348–355].

17.3. Advances in science and technology to meet challenges

The brain is undoubtedly the most sophisticated information processing device known to humankind. However, the underlying principles of its function and dysfunction seem to vastly differ from those understood through conventional computing hardware. Although we are far from fully understanding the principles of brain function, which operates by consuming a few watts of power, it still manages to solve complex problems by devising algorithms that appear to be intractable with current computing resources. The brain is robust, reliable and resilient in its operation despite the fact that its building blocks could fail. All of these are highly advantageous features to instruct the better design of the next generation of computing hardware [356].

In the future, computational models and simulations of brain function and dysfunction will be better informed through its unique capabilities to model and predict the outside environment, underlying design principles, their mechanisms and multi-scale organization and operation. The interface of experimental and computational neuroscience will shed new light on the unique biological architecture of the brain and help translate this knowledge into the development of brain-inspired technologies.

We are still at the tip of the iceberg in dealing with and solving diverse challenges. It is possible that 'neuromorphic' computing systems of the future will comprise billions of artificial neurons and the development, design, configuration and testing of radically different hardware systems will require new software compatible with the organizing principles of brain function. This will require a deep theoretical understanding of the way the brain implements its computational principles. Knowledge of the cognitive architectures underlying capabilities such as attention, visual and sensory perception can enable us to implement biological features that current computing systems lack.

Recent advances in 'connectomics' allow unprecedented reconstructions of biological brain networks. These connectomes display rich structural properties, which include heavy-tailed degree distributions, segregated ensembles and small world networks [357]. Despite these advances, how network structure determines function computation remains unknown. Going forward, approaches from computational modelling and neuroscience could better inform the design of neuromorphic systems to unravel how structure leads to function, how the same network configuration could result in a spectrum of cognitive tasks depending on the network state and how different network architectures could support the emergence of similar cognitive states.

17.4. Concluding remarks

The fledgling field of computational neuroscience has now transformed into a dynamic and vibrant global community. Computational models and simulations offer an unprecedented framework to understand brain function and dysfunction. But what could computational neuroscience possibly achieve in the next two or three decades? Can models of neural networks explain any higher-order brain function? For example, visual function involves progressive information processing with across brain regions with similar network architectures but different modalities of representation and computation [358]. Most neural network models to date are exclusively designed with functionally homogeneous populations of neurons, comprising simplified point neuron models or 'ball and stick' compartmental models that do not necessarily capture the rich features of brain organization and architecture.

Although computational neuroscience initially took off in the 'reductionist' direction that enabled researchers to abstract approaches and insights from disparate fields such as mathematics, physics and engineering it probably resulted in overly simplified assumptions about the unprecedented complexity of the building blocks of the brain. Future models of neural networks and theories would have to take cognizance of the astounding assortment of neuron types, impact of neuromodulatory systems, metabolic factors, co-release of neurotransmitters, neuro-glial-vasculature interactions adaptive mechanisms of neurons and synapses across diverse spatio-temporal scales among many other biological details.

In conclusion, computational neuroscientists that use their preferred level of abstraction and modeling traditions should come together as a community, which surprisingly, is rather difficult. As a famous neuroscientist and Nobel laureate, Sydney Brenner remarked 'we are drowning in a sea of data and starving for knowledge'. Computational neuroscientists should interact a lot more with experimentalists and learn from each other. More theory is the need of the hour to crack the brain.

Acknowledgements

S.R. acknowledges support from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 842492 and a Newcastle University Academic Track (NUAcT) Fellowship.

18. Stochastic computing

Jonathan Tapson

University of Technology, Sydney

18.1. Status

The human brain is extraordinarily efficient in computation, using at least five orders of magnitude less power than the best neuromorphic silicon circuits [359]. Nonetheless, it still consumes approximately 20%–25% of a human's available metabolic energy, and it is safe to assume that the evolutionary pressure to optimize for power efficiency in the brain was extremely severe [360]. It therefore comes as a suprise that the transmission of signals through the brain's synaptic junctions is apparently noisy and inefficient, with probabilities of 0.4–0.8 for transmission of an axonal spike being typical (see [361] for a detailed review). This raises the question: is this transmission variability a bug or a feature? Also, can any brain-inspired computational system which does not include synaptic stochasticity capture the essence of human thought? Perhaps it serves to regularize biological neural networks, in the same way that machine learning techniques such as dropout are used to make ANN more robust (figure 25).

This, and many other similar questions, drive the field of stochastic computation [362, 363]. The field covers a large number of techniques in which some kind of probabilistic function, filter or network is used to create a computational output which would not be possible with deterministic systems. For example, in neuromorphic neural networks, the use of nonlinear random projections has become a commonplace method for raising the dimensionality of an input space prior to a learned solution layer. Technologies as diverse as silicon device mismatch, memristor mismatch, and even random networks of conductive fibres have been proposed and tested for this purpose [364]. Generally, stochastic computation methodologies fall into a number of categories:


- (a) Systems where noise or randomness is used to add energy to a system, enabling it to traverse or establish states which were otherwise inaccessible. Energy in this sense means potential or kinetic energy (in mechanical, electrical or chemical form), rather than general system power consumption. The various phenomena of stochastic resonance and stochastic facilitation [362] are typical examples of these systems.
- (b) Systems where the data or input streams are intrinsically random or noisy, and rather than filter or otherwise reduce the uncertainty in the signals, a computational system is devised which processes the raw signal to produce a computationally optimal output. Recently, many of these systems apply Bayesian models [365], particularly when derived from biological principles.
- (c) Systems in which it is required to project the input space nonlinearly to a higher dimension, in order to facilitate linear or other numerical solutions to some regression or classification problem. This obviously includes conventional neural networks; however, there is an increasing body of research in both human neuroscience and machine learning in which random nonlinear projections are found to be optimal for some function.

18.2. Current and future challenges

Perhaps the most exciting theme for the future of stochastic computation is the modelling of random or noisy processes in the brain, and applying these to real world computational problems. For example, there is increasing evidence that random projection layers are used in the brain to improve the versatility and selectivity of neural circuits [366]. This has parallels in random projection ANN, but the utility of these systems has, until now, been overshadowed by the performance of application specific machine learning systems. Additionally, there is evidence that in SNN, noise appears to enable probabilistic inference [367]. As these systems move from academia into commercial fields of exploitation, the adapability and robustness of the random networks is likely to shine.

There are some significant issues in exploiting stochastic computation. For example, if randomness is created using intrinsically variable hardware means such as device mismatch, then the randomness is unique to each device, and therefore each device will need to be characterised before it can be used. This is reasonable in laboratory research, but infeasible as part of the workflow in producing highvolume consumer goods.

A further problem for stochastic computation is that many applications—for example, biomedical or aerospace systems—require deterministic and provably failsafe operation. Systems that compute bsed on probability and statistical likelihoods may struggle to meet industrial standards, notwithstanding that their performance may be superior to current methods.



18.3. Advances in science and technology to meet challenges

The foundations of stochastic computation are not mature. The field at present consists largely of proofs-ofconcept, with few widely-applicable principles that can be applied in a generic way. Noise and randomness are still seen as obstacles to be overcome in signal processing, so unless a practitioner has some prior experience of stochastic computing, there is little likelihood of applying these methods as part of a standard process. Nonetheless, the increasing interest in Bayesian methods, and their existence and application in biological neural systems, is generating growing interest in this approach. What the field requires for advancement are the following:

- A foundational theory of probabilistic computing, integrating Bayesian, SR and random projections into a more general framework which answers the question: 'I have noisy and nonlinear input data, or a noisy and nonlinear system. How do I take advantage of the energy and nonlinearity rather than simply filtering it out?'
- Current methods of implementing random connective weights and random numbers in silicon are not ideal. In particular, there are few methods which allow the deterministic generation of 'random' weights in ways that obviate the requirement to characterise each individual part prior to computation. Similarly, generation of random numbers 'on the fly' in computational platforms such as custom silicon or FPGAs requires a significant proportion of hardware and energy to be dedicated to it.

18.4. Concluding remarks

The real world and natural organisms are intrinsically noisy and variable, and yet the most efficient computer known—the human brain—computes very successfully using noisy and variable hardware on noisy and variable signals. As our understanding of this phenomenon grows, the desire to apply the principles in artificial computation will grow accordingly. In stochastic computation, we have the beginnings of a model for this process, and it seems likely that in the future we will regard computational systems which do not use probabilistic principles as being as limited as we now regard classification and regression using the purely analytic methods that were commonplace before machine learning.

19. Convolutional spiking neural networks

Priyadarshini Panda and Youngeun Kim

Yale University, United States of America

19.1. Current status

Recent machine learning literature show that convolutional neural networks have been adopted successfully in various vision applications [5, 368]. To apply SNNs to diverse vision tasks, convolutional SNNs (figure 26) should be explored as a priority. Consequently, there have been many attempts to train a convolutional SNN architecture. Conversion [373] proposes a weight balancing technique in order to convert pre-trained convolutional ANNs to convolutional SNNs. Note, we make a distinction between convolutional ANN and convolutional SNN based on the type of input processing and the neuronal activation as shown in figure 26. The overall architecture still uses convolution and pooling operation suitable for enhanced performance in vision tasks. The conversion method yields competitive accuracy, but requires careful weight scaling to emulate the floating point activation value of ANNs and therefore its application is restricted to simple static image recognition. Surrogate gradient learning circumvents the non-differentiability of a LIF neuron by defining an approximate gradient during backward propagation [374]. Training SNNs directly on spike trains can enable parsing various types of input representation beyond static images, and potentially lead to SNN implementation for diverse computer vision applications.


Further, the development of convolutional SNNs is in tandem with a growing interest in dynamic vision sensors (DVS). Different from a standard frame-based camera, a bio-plausible DVS camera emulates the human retina and generates asynchronous and binary spikes according to illumination changes. Therefore, this visual sensor has the advantage of a high frame rate, low energy consumption, and less blurring effect compared to a conventional frame-based camera [375]. A convolutional ANN fails to exploit the advantages of DVS camera inputs since ReLU neurons cannot capture temporal information. On the other hand, owing to temporal spiking neuronal dynamics, convolutional SNNs are a natural fit for processing DVS data. Also, combining DVS camera inputs with SNNs yields a fully-spiking and energy-efficient system. Despite the potential advantages, an algorithm for training convolutional SNNs on DVS data is still underdeveloped.

19.2. Challenges

The majority of previous works on convolutional SNNs classify static images. SNNs require a spike train to process in the temporal domain. Thus, various coding schemes have been proposed for converting static images to spikes [386]. It is important to select a proper coding strategy since the energy consumption on an asynchronous neuromorphic hardware is approximately proportional to the number of spikes. Currently, rate coding is the most widely-used coding scheme since it yields high application-level accuracy. But, rate coding generates spikes, where the number of spikes is proportional to the pixel intensity. This causes multiple (and sometimes redundant) spikes per neuron and therefore reduces the energy-efficiency of the overall system. In order to bring more energy advantages, a new coding scheme with fewer spikes should be explored.

Another challenge is directly training deep convolutional SNNs. From ANN literature, it is well known that network depth is a crucial factor for achieving high accuracy on vision tasks. ANN-toSNN conversion enables deep SNNs with competitive accuracy. But, emulating float activation with multiple binary spikes requires a large number of time-steps, which in turn increases overall energy and latency. Surrogate gradient learning allows short latency and can be used with flexible input representations, but it suffers convergence issues when we scale up the depth. Therefore, convolutional SNNs with surrogate learning are still restricted to shallow networks on trivial datasets. Overall, effective spike-based training techniques for deep convolutional SNNs is necessary to reap the full energy-efficiency advantages of SNNs.

Finally, there is a need to investigate SNNs beyond the perspective of accuracy and energyefficiency. Good-fellow *et al* [376] showed that unrecognizable noise can induce a significant accuracy drop in ANNs. This questions the reliability of ANNs since humans do not misclassify when presented with such perturbed adversarial inputs. In this light, there is a need to analyze the robustness of SNNs. Furthermore, the internal spike behavior of SNNs still remains to be a 'blackbox' as that of conventional ANNs. In the ANN domain, several interpretation tools have been proposed and provide cues for advanced computer vision applications such as visual-question answering. In a similar vein, an SNN interpretation tool should be explored because of its potential usage for real-world applications where interpretability in addition to high energy-efficiency is crucial.



19.3. Advances in algorithms and hardware to meet challenges

Beyond rate coding, temporal coding generates one spike per one neuron in which spike latency is inversely proportional to the pixel intensity. Mostafa *et al* [377] applied an exponential kernel to derive locally exact gradients of spiking neurons. The authors trains SNNs on XOR problem and MNIST dataset with much fewer number of spikes compared to standard rate coding. Han and Roy [378] tackled the memory access issue for tracking a synapse at every time-step until the synapse receives its first spike. They use two signed spikes that denote the start and finish time-step respectively, improving energy-efficiency. Besides temporal coding, phase coding has also been proposed based on biological observation [379], which encodes temporal information into spike patterns based on a global oscillator in an energy-efficient manner [380].

In order to address the shortcomings of ANN-to-SNN conversion and surrogate gradient learning, Rathi *et al* [381] proposed hybrid training-a conversion process followed by surrogate gradient learning. As a result, they achieve competitive performance with significantly lower latency. Wu *et al* [382] presented a spatio-temporal backpropagation technique for direct training convolutional SNNs. The authors evaluate their method on static MNIST as well as DVS-based N-MNIST. Recently, Kim and Panda [383] addressed the scalability problem of direct spike-based training by proposing a time-specific batch normalization technique, called batch normalization through time (BNTT). The authors successfully trained convolutional SNNs from scratch with BNTT spike-based training on complicated datasets such as CIFAR100, Tiny-ImageNet and DVS CIFAR10.

Few recent studies have shown that SNNs have more robustness compared to their ANN counterparts. For the first time, Sharmin *et al* [384] observed that SNNs are more robust with respect to adversarial images. They provided a comprehensive analysis on VGG9 networks trained on CIFAR10 dataset with various attack scenarios including black-box and white-box attacks. Also, the authors of [385] asserted that a Poisson coding generator and non-differentiable neuronal dynamics are the main reasons for SNN robustness. Recently, the authors of [370] presented a heatmap visualization tool that shows where SNN focuses on during classification. The results support that SNNs show more robust attention on discriminative regions with respect to adversarial samples. We believe that research on interpretability is essential to open up the possibility of an interpretable and reliable neuromorphic system.

Finally, from a hardware implementation perspective, memristive crossbar array based neural network accelerators have been shown to be area and energy efficient by many orders of magnitude than fully digital CMOS implementations [386]. Thus, it is natural to explore the possibility of implementing convolutional SNNs on such crossbars. But, the convolutional architecture requires a careful weight mapping process. This is because the matrix-vector-multiplication operation of the crossbar arrays are suitable for a fully-connected layer. On the other hand, a convolutional layer uses spatially shared weight kernel where in, same weight kernel interact with different input patches to produce different outputs. Thus, weight sharing can lead to ineffective

crossbar utilization (see [386] for more details). To address the mapping problem due to the weight sharing, the authors in [371, 372] present a mapping protocol for convolutional ANNs. They also provide a simulation tool for crossbar implementation, which evaluates the energy and performance of a network during inference on crossbars. We believe that similar mapping protocols can be extended to convolutional SNNs.

19.4. Concluding remarks

Convolutional architecture plays a key role in AI. Therefore, this architecture is also important in the neuromorphic community because of its practical usage for expanding SNN applicability for diverse vision scenarios. Moreover, from a neuroscientific perspective, convolutional architecture is an interesting research topic itself. For example, Hubel and Wiesel [368] discovered a hierarchical model like convolutional architectures in the primary visual cortex (V1) of cats. After two decades, Fukushima and Miyake [369] suggested a more structured neural system, which consists of S-cells extracting visual features and C-cells working as a non-linear function (like ReLU). Today, convolutional ANNs have found use in neuroscience studies to interpret the activity recorded from different brain regions and explain about the functionality of the region. Since convolutional SNNs are more bio-plausible given the temporal spike processing, we believe convolutional SNNs alongside suitable neuromorphic hardware can be used as interpretation tools for neuroscience. Finally, in order to bring the convolutional architecture to a neuromorphic system, recent works have focused on energy-efficient coding schemes and training techniques. In addition to the development of SNNs with higher energy-efficiency and accuracy, robustness and interpretability of convolutional SNNs are significant research directions. Overall, many technical challenges, such as training, coding, and interpretability need to be addressed to build reliable and accurate convolutional SNNs compatible with neuromorphic hardware.

Acknowledgements

The research was funded in part by C-BRIC, one of six centres in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA, the National Science Foundation, the Technology Innovation Institute (Abu Dhabi) and the Amazon Research Award.

20. Reservoir computing

Gouhei Tanaka

The University of Tokyo

20.1. Status

Resevoir computing (RC) is a machine learning framework capable of fast learning, suited mainly for temporal/sequential information processing [387]. The general concept of RC is to transform sequential input data into a high-dimensional dynamical state using a 'reservoir' and then perform a pattern analysis for the reservoir state in a 'readout'. This concept was originally conceived with a special class of recurrent neural network (RNN) models (see figure 27(a)), such as echo state networks (ESNs) [388] and liquid state machines (LSMs) [389]. The main characteristic is that the reservoir is fixed and only the readout is adapted or optimized using a simple (mostly linear) learning algorithm, thereby enabling fast model training. Owing to the computational efficiency, software-based RC on general-purpose digital computers has been widely applied to pattern recognition, such as classification, prediction, system control, and anomaly detection, for various time series data. To improve computational performance, many variants of RC models have been actively studied [390].

On the other hand, hardware-based RC is an attracting option for realizing efficient machine learning devices. A reservoir can be constructed not only with RNNs but also with other nonlinear systems. In fact, a rich variety of physical reservoirs have been demonstrated using electrical, photonic, spintronic, mechanical, material, biological, and many other systems (see figure 27(b)) [391]. Such physical RC is promising for developing novel machine learning devices as well as for finding unconventional physical substrates available for computation. The system architectures of hardware-based reservoir can be mainly classified into several types, including network-type reservoirs consisting of nonlinear nodes, single-nonlinear-node reservoirs with time-delayed feedback [392], and continuous medium reservoirs [393]. Many efforts are currently underway to improve computational performance, enhance energy efficiency, reduce computational cost, and promote implementation efficiency of the physical reservoirs. They are often combined with a software-based readout or a readout device based on reconfigurable hardware capable of multiply-accumulate operation.


Further advances in physical RC would contribute to realizing novel AI chips, which are distinguished from AI chips for deep learning. One of their potential targets is edge computing [394]. High-speed machine learning computation for data stream obtained from sensors and terminal devices would lead to data traffic reduction and data security enhancement in the Internet of Things (IoT) society.

20.2. Current and future challenges

In exchange for the merit that speedy training is possible for RC systems with simple readouts, the fixed reservoir needs to be well designed for high computational performance. For ESNs and LSMs, there are plenty of theoretical and experimental results helpful in designing 'good' RNN-based reservoirs [395–398]. So far, the potential of RC framework has been shown for many benchmark tasks and practical applications (see table 3) [391]. Current challenges on software-based RC include enhancing the utility of RC models for industrial applications, comparing their cost-performance ratio with that of other machine learning methods, and developing extremely efficient RC models through various extensions.

For hardware-based RC, however, design principles have yet to be fully established. Therefore, a fundamental challenge is to make recipes for constructing 'good' physical reservoirs as well as identifying the main components governing the computational performance. The computational performance of RC systems relies on dynamical states of reservoirs when driven by input signals, which should have nonlinearity and memory (or history dependency) for temporal pattern recognition. The response characteristics of reservoirs can differ depending on the type of reservoir architecture, the physical property of the media/substrates in the reservoir, and the pre- and post-processing of signals to and from the reservoir [391, 399]. It is not straightforward to clarify how each of these factors influences the computational ability of the whole RC system, because they are not independent of each other in many cases [400]. In hardware implementation of physical reservoirs, it is needed to consider physical constraints, such as possible time resolution and measurement precision in signal processing, which could also affect the computational property [401]. Therefore, more numerical and experimental studies are necessary for elucidating the relationship between physical and computational properties of RC systems [402].

Toward practical applications of hardware-based RC systems, it is demanded to evaluate each physical reservoir in terms of different aspects such as computational performance, processing speed, power efficiency, scalability, miniaturizability, and robustness. After an accumulation of these evaluations for various pattern recognition tasks, it would be possible to make a comparison between different RC systems and determine






Table 3. Examples of subjects in RC applications. Table reproduced from (Tanaka et al 2019). CC BY 4.0.

Category	Examples	
Biomedical	EEG, fMRI, ECG, EMG, heart rates, biomarkers, BMI, eye movement, mammogram, lung images	
Visual	Images, videos	
Audio	Speech, sounds, music, bird calls	
Machinery	Vehicles, robots, sensors, motors, compressors, controllers, actuators	
Engineering	Power plants, power lines,	
0 0	renewable energy, engines, fuel cells,	
	batteries, gas flows, diesel oil,	
	coal mines, hydraulic excavators,	
	steam generators, roller mills,	
	footbridges, air conditioners	
Communication	Radio waves, telephone calls, internet traffic	
Environmental	tal Wind power and speed, ozone concentration, PM2.5, wastewater, rainfall, seismicity	
Security	Cryptography	
Financial	Stock price, stock index, exchange rate	
Social	Language, grammar, syntax, smart phone	

which physical RC system meets a specific purpose. It is also significant to promote an integration of RC-based machine learning devices with IoT devices.

20.3. Advances in science and technology to meet challenges

RC models can be more efficient by using additional techniques and extended architectures. For instance, a hybrid of an ESN and a knowledge-based model enables better prediction of spatiotemporal chaotic behaviour in high-dimensional dynamical systems [403]. RC models can also be combined with some other machine learning techniques found in this roadmap. Another extended RC model is the deep RC model consisting of multiple reservoirs [404], which is effective for diversifying temporal representations and generating rich dynamical behaviour.

Methodologies for designing and implementing hardware-based RC systems are currently under intensive investigation. For instance, substantial progress has been made for electronic RC [405] and photonic RC [406]. Compared with them, other physical RC systems are still in the initial stage of development [391, 399]. More studies are required to harness nonlinearity and memory of physical reservoirs for efficient computation.

On the one hand, it is significant to improve computational performance of physically implemented RC hardware such that it is competitive with other machine learning hardware. On the other hand, mathematical modelling and simulations of physical reservoirs are useful for proof-of-concept of novel hardware-based RC systems and analyses of their mechanisms. Both approaches are complementary for addressing the challenges mentioned above.

SNNs, which is one of the central topics in neuromorphic computing, have often been used for implementing a reservoir, mostly in the context of LSMs. Neuromorphic chips based on SNNs (e.g. BrainScaleS 2 prototype system) are available for emulating LSM models and thereby useful for exploring optimal setting of SNN-based reservoirs [407]. In the future, efficient SNNbased machine learning hardware for real-time computation could be realized based on the RC framework with biologically plausible mechanisms.

20.4. Concluding remarks

Toward realizing extremely efficient RC systems and hardware for temporal pattern recognition, both algorithmic and implementation efforts are necessary [408]. Moreover, both physical and computational viewpoints

are required for development of physical RC. Therefore, the progress of physical RC would be accelerated by interdisciplinary collaborations between experts in different research areas.

Acknowledgements

This work was partially based on results obtained from a project, JPNP16007, commissioned by the New Energy and Industrial Technology Development Organization (NEDO), and supported in part by JSPS KAK-ENHI Grant Number 20K11882, JST CREST Grant Number JPMJCR19K2, and JST-Mirai Program Grant Number JPMJMI19B1.

21. Computing with spikes

Simon Thorpe

CerCo-CNRS

21.1. Status

Deep learning architectures now dominate AI. But although they are superficially neurally-inspired, there are significant differences with biology. The 'neurons' in such systems typically send floating-point numbers, whereas real neurons send spikes. Attempting to model a system with the complexity of the human brain with floating-point numbers seems doomed to failure. The brain has 86 billion neurons, with around 7000 synapses each on average. Real-time simulation of such a system with a resolution of 1 millisecond would require $(8.6 \times 10) \times (7.0 \times 10^3) \times (1.0 \times 10^3)$ floatingpoint operations a second—over 600 PetaFLOPS, even without worrying about the details of individual neurons. This would saturate the most powerful supercomputer on the planet and require 30 Megawatts of power—over one million times the brain's remarkable 20 W budget. How does the brain achieve such a low energy budget? It seems very likely that spikes could be a key to this efficiency and a reason why, since the late 1990 s, SNN have attracted increasing interest [6, 409, 410].

A first critical advantage is that computation only occurs when there are spikes to process. The AER protocol (address event representation), first proposed in 1992 [411], communicates by sending lists of spikes. It is used in many neuromorphic systems, including DVS (see section 21) and the multi-million processor SpiN-Naker project [203, 412]. An early event-driven spiking neuron simulator was the original version of SpikeNet [413, 414]. At the time, the joke was that such a system could simulate the entire human brain in real-time—as long as none of the neurons spiked!

Second, spikes allow the development of far more efficient coding schemes. Researchers in both neuroscience and neural networks typically assume that neurons send information using a firing rate code. And yet, the very first recording of responses of the optic nerve by Lord Adrian in Cambridge in the 1920 s demonstrated that while increasing the luminosity of a flashed stimuli increased both the peak firing rate and maintained firing rate of fibres, there was also a striking reduction in latency [415]. Thus, even with a flashed stimulus, response latency is not fixed. Sensory neurons effectively act as intensity-to-delay convertors—a fact effectively ignored for over six decades. But in 1990, it was proposed that spike-arrival times across a population of neurons could be a highly efficient code [416], an idea confirmed experimentally for the retina in 2008 [417].

21.2. Coding with spiking neurons

If we accept that sensory neurons can act as intensity-to-latency convertors, it opens a whole range of interesting options. Figure 28 compares various spike coding schemes for transmitting the activation levels for 16 input channels labelled A-P with 4 bit resolution. The 'standard' method would be rate coding, which counts the number of spikes in a fixed period (e.g. 100 ms). But there are numerous drawbacks of such an approach. First, how do you determine the length of the observation window? Second, rate coding is exceptionally vulnerable to changes in global intensity or contrast of the input profile that completely disrupt firing rates.


Using a temporal coding scheme avoids these problems. One option uses the order of firing as a code—rank order coding [418]. Synaptic weights are set to values that depend on the input's rank. Then a simple feed-forward shunting inhibition mechanism makes neurons selective by progressively reducing the effectiveness of later spikes. Although feasible, rank order coding is perhaps too powerful. Even with just 16 inputs, there are factorial 16 possible firing orders—nearly 21 trillion. Each pattern could be detected by a neuron with the appropriate set of weights.

Another option, N-of-M coding [419], is much simpler. It uses a feedback inhibition circuit to count input spikes and block further firing once a given number of inputs have fired—an inexpensive way to implement a k-WTA (Winner-take-all circuit). If a target neuron has W connections, it can detect specific combinations of inputs. Effectively, each neuron gets a level of excitation that depends on how well the N inputs spikes match their set of weights. Crucially, this can be done with binary synapses, greatly simplifying the implementation.

Interestingly, this sort of coding had already been implemented in a commercial image processing package developed by SpikeNet Technology around two decades ago. The one published paper on the technology [420] explained that synapses between the eight orientation maps and neurons in the recognition maps were very sparse (less than 1% of possible connections). But it failed to mention that all the connections were binary—simply because the company wanted to keep its 'secret sauce' secret!

21.3. Learning with spiking neurons

Switching to temporal coding with spiking neurons also makes it possible to envisage much more efficient learning algorithms. Standard backpropagation can undoubtedly train systems to solve particular problems



window. **Bottom**: N of M coding which limits the number of spikes that a transmitted allows very rapid and efficient transmission of large amounts of data.

but requires huge numbers of training trials with labelled data. Fortunately, a human infant's brain does not need to be trained with millions of images of dogs and cats to categorize new images correctly! Instead, they can learn about new objects very rapidly. There is now good evidence that humans learn to detect virtually anything new that repeats, with no need for labelled data. If humans listen to meaningless Gaussian noise containing sections that repeat, they rapidly notice the repeating structure and form memories lasting for weeks [421]. And in experiments where random images from the ImageNet database are flashed at rates of up to 120 frames per second, humans notice images that repeat, even with only 2–5 presentations [422]. None of the existing floating-point (or rate-based) supervised learning schemes could explain such learning. In contrast, a simple spike-time dependent plasticity rule that reinforces synapses activated before the target neuron fires makes them develop selectivity to patterns of input spikes that repeat [423], and will even find the start of the pattern [424]. Similar methods have also been used to generate selectivity to repeating patterns in the output of a dynamic vision sensor corresponding to cars going by on a freeway—again in a totally unsupervised way [425].

These STDP based learning rules use continuously variable synaptic weights and typically require tens of repeats to find the repeating pattern, even when all parameters are optimized. But we have recently developed a new learning rule called JAST using binary weights [426] that can match our ability to spot repeating patterns

in as few as 2–5 presentations. The target neuron starts with a fixed number of binary connections. Then, instead of varying the strength of the synapses (as in conventional STDP learning), the algorithm effectively swaps the locations of the connections to match the repeating input pattern.

The algorithm was originally implemented on a low-cost Spartan-6 FPGA, already capable of implementing a network with 4096 inputs and 1024 output neurons, and calculating the activation level and updating all the outputs 100 000 times a second. The circuit also included the learning algorithm on-chip.

21.4. Concluding remarks

Although the deep learning revolution has transformed the landscape and allowed the development of artificial systems that rival or surpass human levels of performance, current systems lag behind their biological equivalents on several fronts. They are several orders of magnitude less energy efficient than biological systems, and the mode of learning is highly non-biological.

SNN could overcome many of these limitations in the next few years. Firstly, current systems are handicapped by very inefficient rate-based coding schemes or their floating-point equivalents. Spikes allow the use of temporal coding schemes such as rank order coding and N-of-M coding that are way more efficient. Second, spikes allow the use of spike-time dependent plasticity rules, some of which operate with binary connections and can learn to recognize repeating patterns very efficiently. Such coding schemes avoid the internal noise problems associated with Poisson rate coding, since they are effectively completely determinist. Of course, there will always be some degree of noise in the sensory data and in the mechanism leading to spike initiation in the sensors themselves.

Finally, sending information as spikes will allow the implementation of vast networks with billions of neurons and trillions of connections even with currently available technology. The critical question is the level of sparseness that can be achieved. It is quite possible that many cortical neurons are silent a lot of the time ('neocortical dark matter'), and that each neuron may only need a small number of (binary) connections. If so, the future impact of spike-based computation could be enormous.

Acknowledgements

The author thanks the CNRS for support.

Section 4. Applications

22. Robotics

Chiara Bartolozzi

Istituto Italiano di Tecnologia (IIT)

22.1. Status

Neuromorphic systems, being inspired by how the brain computes, are a key technology for the implementation of artificial systems that solve problems that the brain solves, under very similar constraints and challenges. As such, they hold the promise to efficiently implement autonomous systems capable of robustly understanding the external world in relation to themselves, and plan and execute appropriate actions.

The first neuromorphic robots were proof of concepts based on ad hoc hardware devices that emulated biological motion perception [427]. They relied on the know-how of chip designers, who had to manually turn knobs to tune the chip behaviour. That seed could grow into a mature field thanks to the availability of hardware that could be more easily tuned by nonexperts with standard software tools [428] and of quality DVS [429] and neuromorphic computing chips and systems [412, 430] featuring many instances of neurons and (learning) synapses that could be used as computational primitives for perception and decision making. Since then, neuromorphic robotics followed three main paths, with the development of visual perception for robots using event-driven (dynamic) vision sensors [375, 431], proof-of-concept systems linking sensing to control [432] and SNN for the control of motors [433, 434]. At the same time, the neurorobotics community started developing models of perception, cognition and behaviour based on SNN, with recent attempts to implement those on neuromorphic platforms [435–437]. Finally, the computational neuroscience community has developed learning theories to reconcile DNNs with biologically inspired spike-based learning and to directly develop spiking neural models for motor control that in the future could be implemented on neuromorphic hardware [438–440].

In this rich and lively scenario, the multiple contributing communities and research fields have the potential to lead to the next breakthrough, whereby neuromorphic sensing and computing support the development of smart, efficient and robust robots. This research is timely and necessary, as robots are moving from extremely controlled environments, to spaces where they collaborate with humans, where they must dynamically adapt, borrowing from neural computational principles (figure 29).

22.2. Current and future challenges

The complexity of robotics poses a number of challenges for the neuromorphic community, that needs to demonstrate scalable methods for sensing, perception, decision making and control, that can be deployed in complex applications.


The event-driven readout should be applied to the majority of sensors on the platforms, including those that convey information about the robot's state (inertial, F/T sensors, encoders, etc). Additionally, more sophisticated event-driven encoding could be implemented, by using local computation that emulates the neural encoding of specific sensory modalities.


Understanding the neural code relative to stimuli properties up to decision making [441] will guide the design of SNN to extract meaningful information and take appropriate decisions.

Controlling a robot requires translating information from the sensory to the joint domain, coordinating different joints to implement stable and smooth trajectories, and counteracting actuation errors. Also in this domain, it is crucial to understand principles of biological control and actuation that lead from the cortical motor commands, down to muscles' recruitment and the relative feedback signals.

At the same time, the big question is how much these principles are shaped on the properties of muscles, and how they have to change to be applied to any type of physical (artificial) actuation. Hidden in this research domain is the intrinsic link between the neural code for actuation and the sensory feedback code that interacts with control signals to effectively drive the system.

Given the number of components that need to be developed and orchestrated, a possible path is the progressive integration of specific neuromorphic modules within a 'traditional' robotic architecture. However, the overall system's architectures might be based on very different assumptions and this hybrid architecture might not fully exploit the advantages of the neuromorphic paradigm. A fully neuromorphic system should be the final goal, requiring an exponential growth of the community working on all the bits and pieces that compose a robotic system.





The signature of neuromorphic robots will be continuous learning and adaptation to different environments, different tasks, changes in the robot plant, different collaborators. This must be supported by hardware capable of handling plasticity at multiple temporal scales and a strong knowledge of how the brain implements such mechanisms.

At the technological level, it is paramount to develop neuromorphic devices that can be embedded on robots, increasing the neurons and synapses count and fan-in fan-out capabilities, while maintaining a low power budget. Ideally, those devices have standard interfaces that do not require the use of additional components to be connected to the software infrastructure of the robots. With growing task complexity, and the need of multiple hardware platforms to run different computational modules, the neuromorphic paradigm could take advantage of robotic middlewares, such as ROS, or YARP, that are currently seamlessly integrated with neuromorphic sensors and computing devices.

22.3. Advances in science and technology to meet challenges

The main components in which robotics is traditionally categorised—perception, decisionmaking and control—are strictly connected and influence each other. Studying their interplay and understanding how the neural code reflects sensory information, decision processes and how these are influenced by action and vice versa, is the future research challenge to bring neuromorphic agents to a level of complexity capable of enabling robots to effectively interact with humans. An advancement in the understanding of the role of different brain areas, their working principles and their interaction with other areas across different temporal and spatial scales shall guide the design of artificial architectures using spiking neuron models, synaptic plasticity mechanisms, connectivity structures to implement specific functionalities. It is crucial to find the right level of detail and abstraction of each neural computational primitive and develop a principled methodology to combine them. Starting from highly detailed models of brain areas, the community shall find reduced models that can implement their basic functionality and that can be implementable on neuromorphic hardware.

As the community is now developing SNN to extract information from a single sensory modality, the next step would be to take into account information from other sensory modalities, so that decisions depend on the state of the environment, of the robot and of the ongoing task. Among many others, a key area to take inspiration from is the cerebellum, that supports the acquisition of motor plans and their adaptation to the current (sensed) conditions [442]. The resulting computational frameworks shall therefore include dynamic and continuous learning and adaptation.

On the other hand, progress is necessary in the neuromorphic hardware supporting those new frameworks. New circuits for the emulation of additional computational primitives are needed, as well as the possibility to support dynamic, continuous learning and adaptation at multiple timescales.

Specific to the robotic domain, neuromorphic devices should be truly embeddable. To this aim, standardisation of communication protocols, programming tools, online integration with the robot's middleware must be developed. The necessary miniaturisation to pack more computational resources on a single system that can be mounted on a robot goes through the integration of silicon devices with memristive devices. On a longer term, nanotechnology and flexible electronics could represent a viable solution to further miniaturize, or distribute computational substrates that can de-localise computation to the periphery, or create folded structures similar to the cortex, that through folding increased the surface available for computation, achieving higher computational capabilities (figure 30).

22.4. Concluding remarks

Neuromorphic sensing, computing and actuation are based on design principles that will make them excel in applications for which animals have evolved. At the same time, animals' behaviour does not only depend on the brain but also on its embodiment, as the shape and distribution of sensors and of muscles, tendons, bones and their interaction influence how the sensory signals are acquired and must be interpreted, and how the control signal has to be delivered. The continuous interplay between brain, body and environment shapes neural computation and this must be taken into account when tailoring neuromorphic computation to robots.

The result of this research path will be crucial and timely for developing the next generation of robots that face the extremely hard challenge of collaborating with humans in humandesigned habitat.

The interdisciplinary nature of the neuromorphic approach results in a call for action towards different research communities and different souls within the neuromorphic community. Computational neuroscience and machine learning are called to be the theoretical backbone of neuromorphic computing, micro- and nano-electronics, engineering, physics and material science are called to develop the next generation physical substrate for neuromorphic sensing and computing.

Acknowledgements

The author would like to thank E Donati for fun and insightful discussions and brainstorming on the topic.

23. Self-driving cars

Jonathan Tapson

University of Technology, Sydney

23.1. Status

Self-driving cars have been a staple of science fiction for decades; more recently, they have seemed like an attainable goal in the near future. The machine-learning (ML) boom of the period 2015–2020 gave great cause for optimism, with experts such as the US Secretary of Transport Anthony Foxx declaring in 2016 [443] that 'By 2021, we will see autonomous vehicles in operation across the country in ways that we [only] imagine today...My daughter, who will be 16 in 2021, will not have her driver's license. She will be using a service'.

This optimism has faded away in the last three years, with the recognition that while it is straightforward to make cars autonomous in simple environments such as freeway driving, there are a multitude of situations where driving becomes too complex for current solutions to achieve autonomy. It is tempting to refer to these as 'corner cases' or 'edge cases'—in the sense of being a highly unlikely combination of circumstances, at a 'corner' or 'edge' of the feature space, which produces a situation where a machine learning algorithms fails to operate correctly—except that, in the real-world of driving, these situations appear to be far more common than was originally expected.

It may be helpful to use the industry terminology when discussing self-driving cars. Self-driving is more formally known as Advanced Driver Assistance Systems (ADAS) and the industry generally uses the Society for Automotive Engineering's (SAE) five-level ADAS model, illustrated below, when discussing autonomous driving capabilities (figure 31).

The more recent perception of ADAS progress can be summed up in a quote from Prof Mary Cummings, Director of Duke University's Humans and Autonomy Laboratory [444]: 'there are basically two camps. First are those who understand that full autonomy is not really achievable on any large scale, but are pretending they are still in the game to keep investors happy. Second are those who are in denial and really believe it is going to happen'.

Between the optimistic and pessimistic extremes, there is a consensus view amongst ADAS researchers that while full level 5 ADAS is unlikely to be available in the next five years, level 4 ADAS is both an attainable and useful target.


23.2. Current and future challenges

One core challenge in achieving high levels of ADAS is perception, and particularly visual perception. It has become apparent that the visual perception of human drivers is extraordinarily hard to reproduce artificially. The problem may be summed up with the following example: most, if not all, human drivers can infer the intent of a pedestrian viewed at a distance of 100 m. This is critically important for driving at moderate to high speeds (70–100 km h⁻¹) in non-freeway environments, and is as yet impossible to achieve with machine vision [445]. Consider that a pedestrian 1.5 m high when viewed at 100 m distance subtends an angle of less than 1° vertically and perhaps 0.1° horizontally. When imaged by a video imager through a moderately wide-angle lens, such as is necessary for the forward-facing cameras in an ADAS system, this corresponds to about 3×16 pixels in an HD system. It might be thought that the system is improved by increasing the image resolution (it would be 10×85 pixels in a 4 K system) but this also increases the space which must be searched for the pixels of interest.

A second core challenge in ADAS is the complexity of the world that must be modeled by the machine learning systems. Consider the example of a driver who sees a soccer ball bounce across the road in front of their moving car. A human driver is able to draw on an entire world experience concerning balls (and their relationship to perhaps children, or dogs) in reacting to this event. A machine learning system can (generally) only draw upon cases in their training set in determining a response. The effect is that a human is able to respond to events which have never previously occurred, in their or anyone's experience; machine learning systems are, as yet, not reliable in this regard. This is utterly different to the problem faced by, say, a Go-playing ML system—the rules (the world model) of Go can be written in a single page of text. It has taken some time for ML researchers to fully appreciate the difference between these two problems.

A third problem is the power cost of computation required for ADAS. It is obviously difficult to anticipate the power cost of something that hasn't been achieved yet, but the table below gives typical figures. For context, the 2019 Tesla full self driving (FSD) chip is reported to achieve 36 TOPS at 72 W power consumption (table 4).

These figures raise questions as to the extent to which ADAS systems will reduce the power and range available, particularly for electric vehicles. One model suggests a range reduction of the order 9%–20% for the






Table 4.	Compute requirements for various ADAS levels (source:
horizon i	robotics).

ADAS capability	Compute requirements
L2	2 TOPS
L3	24 TOPS
L4	320 TOPS
L5	4000 + TOPS

control systems for level 4 ADAS [446], which is not insignificant, particularly for electric vehicles where range is a critical issue.

23.3. Advances in science and technology to meet challenges

The perceptive reader will see that the problems facing the self-driving car industry are those which are addressed by neuromorphic engineering: modeling and reproducing human-quality perception; building cognitive models of the world; and reducing the power required for real-time perception cognition-action loop computation.

There has been some interest in using neuromorphic DVS [447] for ADAS vision, but this is hampered by the low spatial resolution of current DVS sensors (which has a complex relationship to the spatio-temporal event processing capacity of the interfaces and downstream processors used) and a perception, perhaps inaccurate, that ADAS will always require conventional (static, framed) sensors, and therefore the inclusion of DVS systems in the senor pack is a costly redundancy. This problem may be solved in due course by hybrid DVS/static imagers. In the bigger picture, the benefits offered by the high dynamic range and spatio-temporal resolution of DVS imagers has led to a general interest in event-based visual processing, which shows some potential to improve ADAS vision processing independently of the use of DVS.

Perhaps the most important feature of human vision which enables driving is the extraordinary resolution of the fovea, coupled with saccadic motion to apply this resolution to small regions of interest (such as pedestrians). There is significant (mostly unpublished) interest in using models of human salience [448] and attention to direct machine vision to areas of the visual (imager) field which require extra or high-resolution processing. Given the real-time nature of ADAS computation, and the necessity to process correlated streams of visual and 3D point-cloud data (from lidar systems), there is some expectation that event-based neuromorphic computation may be more suitable than current GPU-type computational hardware. At least one neuromorphic event-based hardware startup is focused on real-time vision processing for this purpose [449].

In terms of building cognitive models of the world, we are reaching a point where brute-force approaches to ML are producing diminishing returns. Language models such as GPT-3 [450] can produce impressive passages of text, but it becomes clear that there is no real insight being generated; and these models are trained on orders of magnitude more text than any human could assimilate in a lifetime, suggesting that there is an unfilled deficiency in the model. Neuromorphic approaches such as Legendre memory units [451] are offering equal performance to GPT-3 architectures with $10 \times$ lower training and memory requirement, suggesting that this may help to close this gap. Similarly, the use of neuromorphic hardware such as Intel's Loihi [202] and GML's GrAIOne chips [449], which are strictly event-based and intrinsically sparse, may provide a computational platform that enables these more biologically realistic machine learning methods.

23.4. Concluding remarks

Achieving truly autonomous vehicles is something of a holy grail for engineering in the current era. Recently, there has been a reckoning that it is unlikely to be solved by incremental improvements based on current technologies. Neuromorphic engineering offers exactly the technological leaps that are required in perception, computation and cognition in order to achieve this goal. In particular, building machine vision systems that acknowledge and model the unique features of human vision, and building computational systems that exploit the event-based sensory flow that is common to both human and ADAS systems, seem like clear areas for fruitful research.

24. Olfaction and chemosensation

Thomas A Cleland

Cornell University

24.1. Status

Artificial olfactory systems were early adopters of biologically-inspired design principles. Persaud and Dodd constructed an electronic nose in 1982 based explicitly on the principles of the mammalian olfactory system—specifically, the deployment of a diverse set of broadly-tuned chemosensors, with odorant selectivity arising from a convergent feature detection process based on the pattern of sensor responses to each odorant [452]. Such cross-sensor patterns, inclusive of sampling error and other sources of noise, invite machine learning strategies for classification. Gardner and colleagues subsequently trained ANN to recognize odorant-associated response patterns from chemosensor arrays [453], and constructed a portable, field-deployable system for this purpose [454].

The biomimetic principle of chemical sensing by arrays of partially selective chemosensors has remained the state of the art [455, 456]. Arrays obviate the need to develop highly selective sensors for analytes of interest, as high specificity can be readily achieved by the deployment of larger numbers of partially selective sensors [457]. Moreover, such systems are responsive to a wide range of chemical diversity (odorant quality space; figure 32), enabling the identification of multiple chemical species and diagnostic odorant mixtures and effectively representing their similarity relationships. The intrinsic redundancy of such chemosensor arrays also renders their responses more robust to contamination or interference, provided the analysis method is able to use the redundant information effectively. In contrast, strategies for post-sampling signal processing and analysis have varied. Typically, chemosensor array responses are conditioned by electronic preprocessors and then analyzed by one of a range of methods including linear discriminant analysis, principal components analysis, similaritybased cluster analyses, and support vector machines, along with a variety of artificial neural networkbased techniques [455, 456, 458, 459]. However, more directly brain-inspired techniques also have been applied to both the conditioning and analysis stages of processing. For example, the biological olfactory bulb (OB) network (figure 33) decorrelates similar inputs using contrast enhancement [460]. When applied as signal conditioning to artificial sensor data, this operation improved the performance of a naïve Bayes classifier [461]. Similarly, inhibitory circuit elements inspired by the analogous insect antennal lobe (AL) have been deployed to enhance the performance of support vector machines [459, 462]. Finally, fully neuromorphic circuits for analysis and classification have been developed that are based directly on OB/AL circuit architectures (first by [463]; reviewed in [464]; more recently [465, 466]). These approaches are discussed below.

24.2. Current and future challenges

In many ways, the development of functional/deployable neuromorphic systems is presently limited by algorithms. The energetic and efficiency benefits of neuromorphic engineering—arising from strictly local computation, the colocalization of memory and compute, and spike-mediated communication—are clear in principle, but these properties demand qualitatively new computational strategies to address real-world problems effectively. Developing and optimizing these algorithms for particular tasks of interest is a central challenge for contemporary neuromorphic design.

One strategy for developing effective neuromorphic algorithms is to adapt circuit motifs from the biological brain. Armed with well-developed computational circuit models of olfactory brain structures, neuromorphic olfaction has embraced this strategy, incorporating features such as layered columnar organization, recurrent lateral inhibition, temporal dynamics, spike phase coding, distinct classes of excitatory and inhibitory neurons, custom synaptic learning rules, and adaptive network expansion [460]. Notably, even within these biomimetic constraints, algorithms can differ substantially. For example, lateral inhibition is often applied to decorrelate inputs, but also has served in different implementations to implement attractor dynamics [458, 463, 466], unsupervised clustering [465], and attractor-based denoising [466] in service to different functional goals.

The present challenge is to adapt these biomimetic algorithms to practical applications. Neuromorphic technology is well suited for low-power edge devices, putting a premium on properties like rapid, online learning and adaptation to local conditions that may be partially or wholly unpredictable. Notably, both properties are weaknesses of deep networks [466], which require extensive training with examples of all planned targets including reasonable estimates of predicted future variance. Fundamentally, online learning in neuromorphic olfaction arises from *selective plasticity*, in which the network architecture directs the potentially disruptive effects of new learning to specific circuit elements, coupled with *adaptive network expansion*, which expands the physical network to encompass this new learning and render it accessible. Systems with sufficient





Figure 32. Illustration of the capacities of chemosensor arrays to distinguish small changes in odorant quality. (A) Axes denote a 2D quality (Q-) space of physicochemical similarity, ellipses depict the selectivities of three different chemosensors sampling that space. Sensors with broader receptive fields cover a given sensory space more effectively. Discrimination capacity (denoted by *hot colors*) is maximized where the dropoff of sensitivity is steepest, and where the chemoreceptive fields of multiple sensors overlap. (B) Example two-dimensional Q-space with 30 sensors (ellipse pairs, distinguishing selectivity from sensitivity) and 40 chemical ligands (points) deployed. (C) Mean discrimination capacity depends on the number of sensors deployed into a Q-space, shown as a function of the number of competing ligands deployed into the Q-space illustrated in (B). Deploying additional chemosensors reliably improves system performance. Adapted from [457].



capacity for expansion can exhibit *lifelong learning* capabilities. We have referred to this collection of properties as *learning in the wild* [460, 467], and focused on the capacity of such olfaction-inspired algorithms to learn targets from one- or few-shot learning and identify known targets amidst unpredictable interference [466], function in statistically unpredictable environments [468], and mitigate the effects of sensor drift and decay [467]. Notably, working neuromorphic olfaction algorithms have been deployed on diverse edgecompatible hardware platforms including Intel Loihi, IBM TrueNorth, fieldprogrammable gate arrays, and custom neuromorphic devices [458, 463, 466, 469, 470].

24.3. Advances in science and technology to meet challenges

The theoretical performance capacities of machine olfaction are not clearly constrained by anything short of the fundamental signal to noise limits of their deployed environment. The basic capacity to distinguish among similar odorants can be steadily increased by deploying arrays with larger numbers of different chemosensors, provided those chemosensors are responsive to a range of relevant ligands and distinguish between some of them (figure 32; [457]). The commensurate expansion in network size is tractable on neuromorphic devices in terms of both execution time and energy expenditure (illustrated on Intel Loihi by [466]). The redundancy derived from sensors' overlapping chemoreceptive fields also offers improved resistance to interference from

other odorant sources. The actual capacity for such signal restoration under noise depends on the development of circuits that leverage this capacity, and while early efforts are promising [466], there is substantial room for algorithm improvement, such as integrating the pattern completion and clustering capabilities of piriform cortex circuitry, developing cognitive computing methods such as hierarchical category learning to optimize speed-precision tradeoff decisions, and improving performance in the wild. Making these capacities robust requires the development of larger-scale chemosensor arrays, including compact, high-density arrays that can be deployed in the field. Different sensor technologies, optimized both for different sample phases (gas, liquid) and different chemoreceptive ranges of sample quality (e.g., food odors for quality control, toxic gases for safety), will be required. Large libraries of candidate sensors can be screened [455, 471], reducing the need for predictive models of sensors' chemoreceptive fields in this process. However, molecular imprinting technology, developed to produce highly specific chemosensors, now provides this capacity [472], and in principle could be adapted to produce broader receptive fields by imprinting analyte mixtures.

Neuromorphic circuits are not readily adaptable to arbitrary tasks; the domain-specific architectures that underlie their efficient operation also delimit the range of their applications. Olfaction-inspired networks are not limited to chemosensory applications [467], but they are not likely to be effective when tasks do not match their structural priors. However, the characterization and analysis of such fully functional neuromorphic circuits enables the identification and extraction of computational motifs, yielding toolkits that can be intelligently applied to new functional circuits. Moreover, new techniques for spike-based gradient descent learning have successfully demonstrated few-shot learning in neuromorphic circuits in the future is likely to depend on combinations of these strategies, with qualitative circuit elements drawn from theory and generalized domains established therein via emerging optimization strategies.

24.4. Concluding remarks

Machine olfaction has been an early and successful adopter of neuromorphic strategies for sampling and computation, in part because of the detailed elucidation of olfactory system networks by experimental and computational neuroscience. Larger-scale sensor arrays and the continued development of postsampling neuromorphic circuitry for signal conditioning and contextually-aware category learning in the wild are critical emphases for near-term progress. More broadly, the benchmarking and analysis of the various computational motifs contained in these functional biomimetic architectures may address broader theoretical and task-dependent questions such as where and when to gate plasticity in selectively plastic architectures, and the transformative capacities of local learning rules deployed within particular network contexts.

Acknowledgements

This work was supported by NIH/NIDCD Grant R01 DC014701.

25. Event vision sensors

Christoph Posch

Prophesee

25.1. Status

Neuromorphic event-based (EB) vision sensors take inspiration from the functioning of the human retina, trying to recreate its visual information acquisition and processing operations on VLSI silicon chips. The first device of this kind out of C Mead's group at Caltech, named the 'Silicon Retina', made it on the cover of Scientific American in 1991 [473]. In contrast to early more biologically faithful models, often modelling many different cell types and signalling pathways, in turn leading to very complex designs with limited practical usability, in recent years more focus has been put on the creation of practical sensor designs, usable in real-world artificial vision applications. In [474], a comprehensive history and state-of-the-art review of neuromorphic vision sensors is presented.

Today, the majority of EB sensor devices are based on the 'temporal contrast' or 'change detection' (CD) type of operation, loosely mimicking the transient magno-cellular pathway of the human visual system (figure 34). In contrast to conventional image sensors, CD sensors do not use one common sampling rate (=frame rate) for all pixels, but each pixel defines the timing of its own sampling points in response to its visual input by reacting to changes of the amount of incident light [429, 477, 478]. Consequently, the entire sampling process is no longer governed by an artificial timing source but by the signal to be sampled itself, or more precisely by the variations over time of the signal. The output generated by such a sensor is not a sequence of images but a quasi-time-continuous stream of pixelindividual contrast events, generated and transmitted conditionally, based on the dynamics happening in the scene. Acquired information is encoded and transmitted in the form of data packets containing the originating pixel's *X*, *Y* coordinate, time stamp, and often contrast polarity. Other families of EB devices complement the pure asynchronous temporal contrast function with the additional acquisition of sustained intensity information, either pixel individually [475] or in the form of frames like in conventional image sensors [476].

Due to the high temporal precision of acquired visual dynamics, inherent data sparsity, and robust high dynamic range operation, EB sensors gain increasing prevalence as visual transducer for artificial vision systems in applications where the need for high-speed or low-latency operation, uncontrolled lighting conditions and limited resources in terms of power budget, post-processing capabilities or transmission bandwidth, coincide, e.g. in various automotive, IoT, surveillance, mobile or industrial use cases [375].

25.2. Current and future challenges

Despite their undeniably beneficial characteristics, EB sensors face challenges regarding practical usability and competitiveness in an industrial product. Given the non-deterministic asynchronous nature of the data generation process, leading to e.g. non-constant data rates out of the sensor, the topic of integration of an EB sensor with post-processing into a vision system raises different questions around on-chip event data preparation and interfacing.

Competing technologies for EB vision systems range from conventional computer vision to radar, lidar, ultrasound, passive infrared (PIR) and more. The challenges to make EB a successful technology are related to (1) electro-optical performance of the sensor itself, (2) system integration issues, and (3) challenges around the topic of post-processing for extracting the relevant information from the event data stream.

Pixel size is one of the most important properties for competitiveness with respect to conventional image sensor-based systems. Silicon area, optical format, compatibility with standard optics, camera module dimensions, and form factors dominate cost and applicability of the vision system. Pixel area limits the sensor resolution but also influences the dynamic range of the sensor, in particular its sensitivity at low light levels, which are important factors in key applications such as automotive, mobile or IoT. Thanks to semiconductor technology advances including 3D integration, back-side illumination and improving photodiode efficiency, continuous progress is being made on these issues. With increasing pixel array sizes, the throughput of the sensor readout system and data interface are becoming increasingly relevant to retain the pixel data temporal precision in high-speed/low-latency applications. Power consumption is a key aspect in many artificial vision applications, particularly at the edge. EB vision systems promise to deliver low-power operation thanks to the data sparsity and related efficiency in sensing and processing.

Processing of EB sensor data can be coarsely divided into (1) algorithms that prepare and optimize the raw pixel event data for more efficient transmission and processing on an (external) computation platform, and (2) algorithms that extract higher-level application-specific information out of the event data, to solve a





Figure 34. (a) Simplified three-layer retina model and (b) corresponding CD pixel circuitry; in (c) typical signal waveforms of the pixel circuit are shown. The upper trace represents an arbitrary voltage waveform at the node V_p tracking the photocurrent through the photoreceptor. The bipolar cell circuit responds with spike events of different polarity to positive and negative gradients of the photocurrent, while being monitored by the ganglion cell circuit that also transports the spikes to the next processing stage; the rate of change is encoded in inter-event intervals; (d) shows the response of an array of CD pixels to a natural scene (person moving in the fieldof-view of the sensor). Spikes, also called 'events', have been collected for some tens of milliseconds and are displayed as an image with ON (going brighter) and OFF (going darker) events drawn as white and black dots.

vision task such as e.g. object detection, classification, tracking, optic flow, etc. The first group is preferably implemented close to where the raw data are generated, i.e. near-sensor or in-sensor. Typically implemented in an on-chip HW data pipeline, algorithms pre-process the raw pixel data for more efficient transmission and post-processing, also with respect to memory access and processing algorithms requirements. This data conditioning pipeline can include functions such as recoding, formatting, rearranging, compressing, thinning, filtering, binning, histogramming, framing etc. The latter group includes all application-specific vision processing using computer-vision and/or ML-based algorithms and compute models, typically running on some form of application processor. The question of the optimal compute fabric and architecture to be used with EB sensors is unresolved today, and the optimal choice application dependent. However, as discussed widely in other parts of this review, emerging non-von-Neumann architectures, in particular neuromorphic approaches such as SNN, are better suited to realize an efficient EB system than e.g. general purpose CPUs. Much progress is being made in this area, however challenges remain around the absence of wellestablished deep learning architectures, including training techniques, for event data, or the lack of largescale datasets.

25.3. Advances in science and technology to meet challenges

EB sensors benefit from advances in semiconductor technology. Significant progress was made since early designs with array dimensions of about thousand pixels and pixel sizes of $1000 \ \mu\text{m}^2$, to today, about 15 years later, where the most advanced EB sensors have one-megapixel arrays ($1000 \times \text{larger}$) with 20 $\ \mu\text{m}^2$ pixels ($50 \times \text{smaller}$) [478] (figure 35). In contrast to conventional image sensors, typically having one photodiode and 4 or 5 transistors in each pixel, pixels of EB sensors are small analog computers, typically using 50+ transistors for the required signal processing and communication functions. Early EB sensor designs employed standard CMOS processes and later front-side illuminated (FSI) CIS processes. All transistors and other devices needed to be placed next to the photodiode, leading to large pixel sizes and low fill factors. The introduction of back-side illuminated CIS processes relaxed the situation and today, 3D wafer stacking combined with small pitch metal bonding allows to place most of the pixel circuitry underneath the photodiode. Latest generation EB sensors today reach pixel pitches below 5 $\ \mu\text{m}$ and fill factors above 75% [478].



Following the CMOS technology and integration roadmaps will yield EB devices with increasing industrial applicability. Further advances in production and packaging technologies like triple wafer stacking, diestacking system-in-package and wafer-level optics will support the trend to autonomous ultra-low power/small form-factor edge perception devices and AI systems where the sensor is highly integrated and tightly packaged with pre-processing and application processing, thereby significantly reducing power consumption and transmission bandwidth requirements of an artificial vision system, e.g. in IoT, mobile or perception networks applications.

A big impact on the usability and competitiveness of EB systems is expected to come from future advances in neuromorphic computing and event-based processing techniques. SNN are a natural fit for post-processing to the data generated by EB sensors [202, 479]. But the sparse data output of EB sensors is also a good match to future hardware accelerators for conventional DNN that exploit activation and network sparsity [480].

Recently, new kinds of neuromorphic vision devices beyond CMOS have been demonstrated, exploiting different electro-optical material properties and fabrication techniques to further advance the tight integration of sensing and processing, often combining photon transduction and analog neural network (ANN) functions into a single fabric [129, 130, 481, 482]. Even though these devices are in their early proof-of-concept phase, interesting and promising results have already been demonstrated.

25.4. Concluding remarks

Integrating event-based vision sensing and processing with neuromorphic computation techniques is expected to yield solutions that will be able to penetrate the artificial vision market and gain considerable market share in the coming years [483]. This new class of extremely low-power and lowlatency AI systems could, in a world where power-hungry deep learning techniques are becoming a commodity, and at the same time, environmental concerns are increasingly pressuring our way of life, become an essential component of a sustainable society.

26. Neuromorphic audition

ShihChii Liu

University of Zurich and ETH Zurich

26.1. Status

Neuromorphic audition technology is inspired by the amazing capability of human hearing. Humans understand speech even in difficult auditory scenarios and using a tiny fraction of the brain's entire 10 W. Matching the capability of human hearing is an important goal of the development of algorithms, hardware technology and applications for artificial hearing devices.

Brief history: human hearing starts with the biological cochlea which uses a space-to-rate encoding. The incoming sound is encoded as asynchronous output pulses generated by a set of broadly frequency-selective channels [484]. For frequencies below 3 kHz, these pulses are phase locked to the frequency [485]. This encoding scheme leads to sparser sampling of frequency information from active frequency channels instead of the maximal sampling rate used on a single audio input. The first silicon cochlea designs starting with the work of Lyon and Mead (electronic cochlea) [486], model the basilar membrane (BM) of the cochlea by a set of coupled filter stages. Subsequent designs include those with better matching properties for the filter stages and using coupled filter architectures ranging from the originally proposed cascaded type modeling the phenomenological output of the cochlea [486], to a resistively-coupled bank of bandpass filters that models the role of the BM and the cochlear fluid more explicitly [487, 488].

Later designs include models of the inner-hair cells on the BM, that transduce the BM and fluid vibrations into an electrical signal. They are frequently modelled as half-wave rectifiers in silicon designs. Some designs include the automatic gain control mechanism of outer hair cells that are useful for dealing with larger sound volume ranges from 60–120 dB. Cochlea designs starting from the early 2000 s include circuits that generate asynchronous binary outputs (or spikes) encoded using the address-event representation. Details and historical evolution of these VLSI designs are described in [487, 488]. Recent spiking cochlea designs in more advanced technologies such as 65 nm and 180 nm CMOS demonstrate better power-efficiency (e.g., <1 uW/channel in [489]). These new designs show competitive power efficiency compared to other audio front end designs that compute spectrogram features from regular samples of a single audio source.


Importance of field: in the early 2000s, cochlea circuits were developed for audio bionic applications [490] and models of biological auditory localization circuits [491]. With increasing prevalence of voicecontrolled devices in everyday life, neuromorphic and bio-inspired solutions can potentially be interesting because of the need for low-latency and energy-efficient design solutions in audio edge application domains.

26.2. Current and future challenges

Big research issues: neuromorphic bio-inspired features such as the sparse sampling (e.g., the nonsampling during silent pauses), the event-driven form of brain computing, the natural temporal encoding carried by the asynchronous events (pulses, or sets of pulses, within a specified time window), can enable more energy-efficient solutions for hardware-friendly models that solve an auditory task. The timing information from the asynchronous events of the spiking silicon cochleas is ideal for extracting interaural time differences, which is useful for spatial audition at lower latencies [492, 493]. Speech recognition examples that use spiking cochlea inputs include the reservoir networks of spiking neurons applied to digit recognition [494] (figure 36).

With the advent of deep learning, DNNs have been used successfully for many audio benchmarks. Even though training of these networks requires global supervised methods and large datasets, they can be applied on features from spiking cochlea events for always-on low-level audio tasks such as voice activity detection (VAD) and key word spotting (KWS). These modules can then activate the more energy-expensive audio tasks such as speech recognition which require running larger networks on the cloud. The energy savings of these neuromorphic solutions come from the event-driven way of processing, i.e., the processing is triggered by asynchronous information carried by the frequency-selective events rather than on regular time-stepped spectrogram frames; compute savings come naturally with silent pauses and non-changing inputs. Hardware versions of this combination include a complete system using an FPGA recurrent network processing audio features [496]. ASIC implementations that include both sensor and network show extremely lowpower consumption, e.g., the <1 uW ASIC VAD chip [497].

Mapping spiking neural network (SNN) algorithms trained for a task to SNN hardware platforms require additional considerations such as the variability of the network parameters from transistor mismatch and the timing jitter noise of the transmitted spikes from the transmission of asynchronous spikes in real-time on a specified SNN hardware platform. While more work is needed on local spikebased learning rules [498]



to configure SNNs to reach similar accuracy compared to continuous-valued artificial neural network (ANN) solutions on a specific task, conversion techniques that map trained ANNs to SNNs [499] and global supervised training methods for SNNs have been effective for this goal [500].

26.3. Advances in science and technology to meet challenges

For edge applications, besides the high energy-efficiency needed of the hardware platform, maintaining high audio task accuracy is also important. More work is needed to show that neuromorphic audition systems can be competitive for these domains, e.g., through benefits of spike events for efficient models or local learning.

Audio spectrogram features are dominant features used in conventional audio applications. Studies such as [494, 501] show that the accuracy on a simple speech recognition task using spiking cochlea features while showing lower numbers than spectrogram features for clean conditions, can maintain lower accuracy loss over decreasing signal-to-noise conditions. These ideas need further investigation for larger datasets and tasks. Features extracted from the spikes, e.g., exponential features [502], could help to determine the utility of dynamic information. The real-time extracted azimuth information allows source separation from a mixture of two speakers, e.g., by using the streamed source spikes for speaker identification and keyword spotting using deep network algorithms [502]. Future research is needed for models that operate robustly in wide dynamic range (>60 dB) auditory scenes.

More extensive work is needed to develop training or spike-based learning methods to configure for highaccuracy networks on specified audio tasks. Training directly on spiking networks instead of using conversion ANN to SNN methods would capitalize better on implicit SNN encoding capabilities. The approximation of rate codes to analog activations of ANNs is not ideal. Temporal coding methods would be more energy efficient but need extension to more complex datasets. Combination of the auditory systems together with other modalities for a robotic platform would also be interesting for future developments.

Technology: technology advances include the continual support of analog integrated circuit designs as technology nodes scale down for digital transistors. Of importance for low-power devices are transistors with low leakage current (for low standby power) and fabrication processes that still allow low transistor mismatch without requiring large transistor dimensions. The standby power consumption of the ASIC is an important

IOP Publishing

metric for always-on audio devices. Other challenges include algorithms that are hardware aware, e.g., to the variability of the network parameters after the ASIC fabrication, and approaches to reduce memory access or to create predictable memory access patterns to reduce energy loss from the unpredictable memory accesses of SNNs. The emerging largescale availability of high-density local memory is also an interesting component of future research for the ASIC development.

26.4. Concluding remarks

Progress in neuromorphic audition critically depends on advances in both silicon technology and algorithmic development. Recent ASIC prototypes that combine binary networks and spectral features report competitive low-power numbers (<1 μ W) for KWS [503]. Evaluation of these solutions with ASICs that use a spiking cochlea front-end will help determine the latency-accuracy-power tradeoff between these approaches for audio edge applications. Other considerations are the system-level power of a complete audio system including microphones (current commercial audio assistants are ~1 W), that continue to benefit from the ultra-low power numbers of the audio ASICs. Also needed are event-driven algorithms to address more complex tasks beyond KWS and VAD while maintaining the energy efficiency and low latency benefits, and how we can capitalize on the new emerging highdensity memory technologies to reduce the power dissipation from unpredictable memory accesses of events/spikes. The combination of the DAS with other spiking sensors, e.g., the dynamic vision sensor, will be interesting for robotic platforms. Hybrid solutions that bring knowledge from auditory science, signal processing, and machine learning will enrich future artificial neuromorphic technology solutions for audio application domains.

Acknowledgements

We acknowledge the Sensors Group members and colleagues who have worked on the Dynamic Audio Sensor design and audio systems. Partial funding provided by the Swiss National Science Foundation, HEAR-EAR, 200021172553.

27. Biohybrid systems for brain repair

Gabriella Panuccio and Mufti Mahmud

Istituto Italiano di Tecnologia, Italy Nottingham Trent University, United Kingdom

27.1. Status

Biohybrid systems are established by biological and artificial components interacting in a unidirectional or bidirectional fashion. In this section, we specifically refer to neurons or brain tissue as the biological component of biohybrid systems for brain repair.

The first demonstration of a biohybrid dialogue was achieved *in vitro* at the beginning of the 1990 s by Renaud-LeMasson and colleagues, who established a communication between a biological neuronal network and a computational model neuron [504]. Soon after, Chapin and colleagues brought the biohybrid paradigm to the *in vivo* setting by providing the first proof of concept of interfacing the brain with a robotic end-effector [505], a paradigm that has recently become a reality in the clinical research [506].

Biohybrid systems are now a widespread approach to address brain dysfunction and devise novel treatments for it [507]. Representative examples are electronic devices coupled to biological neurons *in vitro* [219] or to the brain *in vivo* [508] and establishing a bidirectional communication through a closedloop architecture. A key feature of such systems is the real-time processing and decoding of neural signals to drive an actuator for brain function modulation or replacement. To this end, enhancement of biohybrid systems with AI is the emerging strategy to achieve an adaptive interaction between the biological and artificial counterparts. Neuromorphic engineering represents the latest frontier for enhancing biohybrid systems with hardware intelligence [509] and distributed computing [510], offering unprecedented brain-inspired computational capability, dynamic learning of and adaptation to ongoing brain activity, power-efficiency, and miniaturization to the micro-scale. In particular, the intrinsic learning and adaptive properties of neuromorphic devices present the key to bypass the typical trial-and-error programming along with the stiff pre-programmed behaviour of current brain implantable devices, such as those used for deep-brain stimulation. In turn, such a unique potential enables surpassing the drawbacks of current mechanistic approaches with a phenomenological (evidence-based) operating mode. Overall, these features serve as an asset to attain a physiologically-plausible interaction between the biological and artificial counterparts.


The latest avenue for biomedical applications is neuromorphic-based functional biohybrids for brain regeneration. These are hybridized brain tissue grafts (figure 37), wherein the neuromorphic counterpart(s) emulate and integrate brain function, aiming at guiding the integration of the biological graft into the host brain. This crucial aspect cannot be attained by a purely biological regenerative approach. Further advances in neuromorphic biohybrids are thus expected to bring unparalleled strategies in regenerative medicine for the brain: by providing symbiotic artificial counterparts capable of autonomous and safe operation for controlled brain regeneration, they herald a paradigm shift in biomedical interventions for brain repair, from interaction to integration.

27.2. Current and future challenges

As biohybrid grafts for brain repair are fully integrated autonomous constructs, their achievement entails the full physical integration of the neuromorphic counterpart(s) within the bioengineered brain tissue to be grafted. However, while neuromorphic neuroprostheses coupled to bioengineered brain tissue may constitute in themselves a biohybrid system, they do not make up a biohybrid graft in the strict sense, being neuroprostheses conceived as an exogenous body (figure 38).

In functional biohybrids, the integrated neuromorphic counterpart(s) must not perturb the mechanical equilibrium of the biological neurons; hence, depending on the intended application (i.e., extracellular or intracellular biosensing from neurons and their networks), these must be scaled down to meet the size of a small neuronal ensemble (<150 μ m [511]) or even to achieve intracellular residency (<3 μ m [512, 513]). In both cases, the primary challenge, and the pre-requisite, is the aggressive miniaturization of the integrated neuromorphic devices(s) without compromising their computational capability. Thus, functional biohybrids entail the realization of neuromorphic dust. Further challenges arise due to the physical confinement of the neuromorphic counterparts within the biohybrid graft, stemming in the requirement of:

(a) Power-autonomy: device powering cannot rely on a wired power supply unit, such as commonly used subcutaneous batteries: while continuous device operation without the need of battery replacement is of utmost importance in brain regeneration, wiring neural dust is unfeasible (the devices are ultrasmall




and physically inaccessible). Further, the operation of an autonomous system, by definition, should not depend on external components.

- (b) Wireless operation: this is required to follow the graft's evolving function during the regeneration process, to enable wireless device re-programming and hardware failure monitoring.
- (c) On-chip learning, supported by application-specific integrated circuits for advanced signal processing, to follow the evolving temporal dynamics of the graft during its integration within the host brain, without the aid of an external controller.
- (d) Bioresorbable property: in aiming at healing brain damage, the neuromorphic counterparts should be regarded as a temporary aid in the process. Thus, they should be removable upon completion of brain repair. While non-invasive micro-surgery techniques, such as high-intensity focused ultrasound, may permit removal of mm-sized devices, this is not technically feasible in the case of ultrasmall (and, even more so, intracellular) devices. Thus, particularly relevant to functional biohybrids is that the neuromorphic counterparts should be bioresorbable.

27.3. Advances in science and technology to meet challenges

Miniaturization. The device size will reflect the complexity of the required architecture. Such a complex system requires integrated interfacing circuits, for which the roadmap toward 1 nm-gate length of CMOS technology [514], along with avant-garde fabrication and integration technologies hold promise to meet this outstanding challenge. Further, novel materials for improved performance are expected to emerge in the near future to enable higher computational capability within smaller sizes. Organic materials hold great promise to merge advanced computation with extreme miniaturization [515].

Power autonomy. While radiofrequency [516] and ultrasonic waves [517] are successfully deployed for powering deep biomedical implants, they require an external transmitter. Energy harvesting from the human body may represent the strategy of the future [518], wherein ATP seems to be the most promising solution to a long-standing issue [519]. However, caution must be made to avoid negative physical effects on the patient, as the amount of harvested energy must be much less than what available. Thus, the device components must be ultra-low-power and energy efficient.



Wireless operation. While autonomous operation is a key feature of the neuromorphic dust, the need of patient monitoring, device fine-tuning, and hardware failure checks should not be underestimated for guaranteeing the patient's safety. To this end, wireless access to the device is fundamental. Thus, dedicated integrated circuits are required, which must be ultrasmall so not to introduce bottlenecks in device miniaturization. As stated above, advanced CMOS technology holds promise to enable these wireless features in neuromorphic dust. Further, protocols tailored to energy efficient wireless communication are needed.

On-chip learning. Understanding spatiotemporal patterns is a key feature to address the evolving dynamics of neuronal networks and reverse-engineer brain dysfunction. So far, these features have been achieved by pre-programming and the use of a microcontroller [520]. Further advances must be made in order to achieve the same level of performance through on-chip learning. This would enable to address the inter-individual variability of the human brain while overcoming the drawbacks of trial-and-error (re)programming and of the need of a wired controller.

Bioresorbable materials. The device materials must be fully biocompatible so not to release cytotoxic compounds in the patient's brain. While outstanding advances have been made in the biosensors field [521], a major efforts must be put in the field of neuromorphic engineering, where the performance of the device strongly depends on materials. In this regard, organic materials may present the key to beat this challenge.

27.4. Concluding remarks

Much progress has been made to incorporate artificial devices within neural spheroids [522], or inside single living cells [512, 513]. However, these are biosensors of simple architecture, far from the complexity of a neuromorphic system. The achievement of functional biohybrids thus requires the parallel effort from diverse disciplines to address the numerous challenges linked to them. The exponential progress in fabrication and miniaturization strategies, energy harvesting, learning algorithms, wireless technology, and bioresorbable bioelectronics heralds the feasibility of cutting-edge biohybrid neurotechnology for controlled and safe brain regeneration.

Acknowledgements

This work was funded by the European Union under the Horizon 2020 framework programme through the FET-PROACTIVE project HERMES - Hybrid Enhanced Regenerative Medicine Systems, Grant Agreement No. 824164.

28. Embedded devices for neuromorphic time-series assessment

Arnab Neelim Mazumder, Morteza Hosseini and Tinoosh Mohsenin


University of Maryland, Baltimore County

28.1. Status

Neuromorphic computing aims to mimic the brain to create energy-efficient devices capable of handling complicated tasks. In this regard, analysis of multivariate time-series signals has led to advancements in different application areas ranging from speech recognition and human activity classification to electronic health evaluation. Exploration of this domain has led to unique bio-inspired commercial off-the-shelf device implementations in the form of fitness monitoring devices, sleep tracking gadgets, and EEG-based brain trauma marker identifying devices. Even with this deluge of work over the years, the necessity of evolving the research direction with day-to-day needs relating to this sphere is still pivotal. The key idea behind the wealth of research in these domains comes from the fact that it is very difficult to generalize human abilities and activities, and it is even more difficult to create devices that can operate at a level as accurate as human-level perception. This is where contemporary machine learning and the more modern deep learning frameworks shine. The current scenario of using automated devices for a variety of health-related applications requires that these devices become more sensitive, specific, user-friendly, and lastly accurate for their intended tasks. This relates to further advancements in the region of algorithm construction and constraint-based design of implementable hardware architectures. The current crop of research in this area investigates DNNs architectures for the purpose of feature extraction, object detection, classification, etc DNN models utilize the capacity of CNNs, recurrent neural networks (RNNs), and even to some extent fully connected layers to extract spatial features for time-series assessment which was previously exhaustively calculated via different hand-engineered feature extraction techniques coupled with simple classification algorithms. Along with this, RNNs and their advanced equivalents in the form of long short term memory networks (LSTMs) and gated recurrent unit has also been integrated into the deep learning architectures to handle timeseries signals. The idea behind this integration stems from the fact that RNNs and LSTMs are modeled in such a way that they can keep track of previous instances of the input data in order to make a prediction, which makes these architectures very effective for pattern and dependency detection within the time-series data. The other aspect of developing these diverse DNN models is to make them readily implementable in terms of hardware accelerators and therein lies the issue of hardware constrained efficient designs. As a consequence, the computation and model size specifications of different hardware-oriented approaches will result in the advancement of application-oriented software designs which will, in turn, increase the reliability and efficiency of these embedded devices.

28.2. Current and future challenges


There are several challenges associated with managing time-series signals for classification or recognition tasks. One of the foremost issues of time-series classification is to make these signals interpretable by the DNNs as these signals contain multiple variables relaying information about concurrent actions and it is difficult to process these signals in their raw form. Authors in [523] proposed a solution to this problem by transforming these multimodal signals into windowed images based on their sampling frequencies. Another obstacle that is related to time-series analysis pertains to skewed or imbalanced information belonging to multimodal variables as the data collection procedure with different sensors might not always be the same. As a way around, a common practice is to use weighted sampling of the input features during the training of the DNN models so as to balance the impact of all features. Pruning outliers in the dataset by eliminating unnecessary sensor data can alleviate this problem as demonstrated by the authors in [524], however, it is not always feasible to delete multimodal information as the sensor data for multiple variables might be correlated. In addition to this, many of the software frameworks dedicated to time-series classification do not consider the large computation overhead of the DNNs. This has a significant impact when these frameworks are replicated on to resource-limited and low-power embedded platforms where the use of off-the-chip-memories becomes essential. As a result, the performance remains limited by the memory bandwidth while the power consumption stays high due to the rapid accessing of off-the chip memories. The extent of these complications has introduced shallow networks [525], approaches to quantizing model parameters [526] along with ternary [527] and binary [528] models that focus on reducing the memory overhead for efficient resource-constrained hardware accelerator implementation. Authors in [529] provide an example of a fixed-point CNN classifier involving 4 bit fixed point arithmetic that suggests negligible accuracy degradation and authors in [530] present fast BNN inference accelerators to meet the FPGA on-chip memory requirements. Reducing memory footprints in hardware accelerators is also tied up to the cost-effective designing of memory units. On top of this, managing and limiting frequent accesses of these memory units also contribute to latency, power, and energy efficiency as a whole.




Thus, a critical challenge in terms of hardware design is to maintain high frequency and energy efficiency with low energy consumption.

28.3. Advances in science and technology to meet challenges

Deep learning frameworks have been widely successful for classifying time-series signals. However, the challenges mentioned in the previous section make this task ever more difficult. To further boost the performance of deep learning methods for time-series data, some form of digital signal processing is commonly required. To this extent, a common practice is to convert these raw waveforms into windowed overlapping time-series frames. A sliding window of some specific size along with a stepping size is passed through all variables, creating a set of images of shape as desired by the user. Since most time-series signals contain label information at precise time intervals, it is fairly easy to determine the label of the images. Another facet of dealing with time-series signals requires feature extraction relevant to the application that is being targeted. With classical machine learning algorithms, this was achieved using several mathematical and analytical processes to determine the correlation between variables. In contrast, one of the strengths of using CNNs or RNNs in deep learning ensures that the relevant features are being extracted in image or time-space. A general practice in time-series classification is to deploy CNN or RNN layers in conjunction with pooling layers as illustrated in figure 39. The pooling layers reduce the feature map size so that the cost of computation for the following fully connected layers is minimized. Additionally, the feasibility of hardware deployment of these deep learning algorithms depends on the computational complexity and size of these architectures. It is imperative that these frameworks are reduced in size via quantization, pruning, or by making the networks shallow in the first place so that they fit on embedded devices with small memories. Hence, there comes a point where the





designer has to find the sweet spot between the accuracy of the model and the practicality of its size being suitable for low power embedded platforms while also ensuring that the energy efficiency of the target device is also satisfactory. Figure 40 shows a comparison among different models with a variety of applications for their model size, classification/detection accuracy, and energy efficiency which establishes that depending on the application, deep learning models can fit on low-power embedded devices with standard performance. Also, a modification to these frameworks can take in additional information in the form of vectors from a separate model to enhance the overall accuracy of the model as demonstrated in figure 39.

28.4. Concluding remarks

Human-related time-series data analysis encompasses a wide range of tasks including speech recognition, keyword spotting, health monitoring, and human activity recognition to name a few. This also allows the dedicated development of embedded devices suited for accelerating such tasks. Challenges in processing such time-variant data for device implementation range from pre-processing the raw signals and removing noise and outliers to interpreting long and short dependencies that exist within the nature of the data. Windowing the continuous stream of data into overlapping frames to be processed using a simple DNN or CNN is a common practice for real-world applications in which the long dependencies in data are negligible. On the other hand, novel approaches such as RNNs and LSTMs can improve the overall confidence of analysis for time-series data with long dependencies. When implementing all these methods on resource-bound hardware in which power, energy, memory footprint, and application latency are all limited, it is of utmost importance to design deep learning algorithms with small model sizes and low computation that meet all the application requirements and hardware limitations. In conclusion, there must be a trade-off between performance and implementation feasibility to justify the use of low-power embedded devices to replicate deep learning applications of time-series assessment.

29. Electromyography processing using wearable neuromorphic technologies

Elisa Donati

University of Zurich and ETH Zurich, Switzerland

29.1. Status

Electromyography (EMG) is a neurophysiological technique for recording muscle movements. It is based on the principle that whenever a muscle contracts, a burst of electric activity is propagated through the close tissue. The source of the electrical signal in EMG is the summation of action potentials of motor units (MUs) [533]. An MU is composed of muscle fibers innervated by axonal branches of a motorneuron, that is intermingled with fibers of other MUs. The recorded electric activity is linearly correlated to the strength of the contraction and the number of recruited MUs. EMG signals can be acquired both invasively, using needle electrodes, and superficially, by placing electrodes on the skin—called surface EMG (sEMG).

EMG signals have been and are relevant in several clinical and biomedical applications. In particular, they are extensively employed in myoelectric prosthetics control for classifying muscle movements. Wearable solutions for this application already exist, but they have a large margin for improvement, from increasing the granularity of movement classification to reducing computational resources needed and consequently power consumption.

Like any other signal, EMG is susceptible to various types of noises and interferences, such as signal acquisition noise, and electrode displacement. Hence, a pre-processing phase is the first step to perform proper signal analysis, which involves filtering, amplification, compression, and feature extraction both in time and frequency domains [534]. The mainstream approach for movement classification is machine learning (ML), which delivers algorithms with very high accuracy [535], although the high variability in test conditions and their high computational load limit their deployment to controlled environments. These drawbacks can be partially solved by using deep learning techniques that allow for better generalization to unseen conditions but remain computationally expensive, requiring bulky power-hungry hardware, that hinder wearable solutions [536].

Neuromorphic technologies offer a solution to this problem by processing data with low latency and lowpower consumption mimicking the key computational principles of the brain [3]. Compared to state-of the-art ML approaches, neuromorphic EMG processing shows a reduction of up to three orders of magnitude in terms of power consumption and latency [537–539], with limited loss in accuracy (5%–7%) [540, 541].


New approaches have been proposed that directly extract the motorneurons activity from EMG signals as spike trains [543]. They represent a more natural and intuitive interface with muscles but currently limit themselves by processing spikes with traditional ML techniques and do not consider the possibility of using more appropriate frameworks such as SNNs.

29.2. Current and future challenges

Although the performance of myoelectric prosthetics increased conspicuously in the last decade [544], they still cannot be used in daily life. The fine-grained control is in fact limited by the number of electrodes. This issue can be overcome by using high-density EMG (HD-EMG), which typically uses hundreds of electrodes, allowing to monitor larger areas and effectively increasing the precision of the measurements [545]. However, HD-EMG uses more computational resources, in terms of power and time required to classify movements and to generate motor commands. Current technologies are not able to process such an amount of data *in situ* and with low latency simultaneously. For this reason, the EMG signals are transmitted, for example via bluetooth, to a remote system that is quite bulky and heavy, making a wearable solution impractical.

Neuromorphic technologies represent a solution to all the described limitations by processing data in parallel, with low latency, and taking advantage of the low-power nature of analog computing and spiking communication, as the biological system they are inspired from. Although recent results show promising advances, the current challenge of neuromorphic technology is to fill the gap with state-of-the-art ML approaches, in terms of accuracy. One of the main reasons behind this gap is the different amount of resources invested in the respective research fields. In addition, current research that focuses on adopting ML methods and implementing them in neuromorphic hardware faces challenges governed by the unsuitability of such substrates which are primarily targeted for SNNs [542].

To get the most from neuromorphic computing we need a change of paradigm, where the neuromorphic technology can directly interface with motorneurons' spiking activity, instead of continuous sEMG signals. This represents a matching condition between inputs and outputs that optimize the information transfer between the muscle activity and the processing and control unit. The spike trains of motorneurons can be extracted from sEMG signals by means of decomposition algorithms. In particular, the spatial distribution



of MUs action potentials can be assessed with activation maps obtained from HD-EMG signals [543]. Nevertheless, current implementations are still computationally expensive, and only recently it was possible for their deployment in real-time. After the decomposition, the spike trains are translated and processed using ML methods instead of better-suited SNNs [546].

Designing neuromorphic systems able to extract and process motorneurons activity from EMG signals will pave the way to a new class of wearable devices that can be miniaturized and directly interface with the electrodes.

29.3. Advances in science and technology to meet challenges

A concrete roadmap towards wearable neuromorphic EMG processing, see figure 41, could be constructed with short and long-term objectives. In the short term, we should advance neuromorphic computation to bridge the gap with ML methods for EMG classification, and optimize decomposition algorithms to make them run real-time on embedded systems. In the long-term, the decomposition algorithm should be ported into a neuromorphic chip to implement a fully spiking pipeline while the technological breakthroughs in surface smart electrodes could potentially be able to record directly motorneurons' spike trains.

Bridge-the-gap. The first step is to understand the requirements to improve the accuracy of EMG movements classification. The front-end, which includes pre-processing and spike conversion, has the largest margin for improvement. Signal-to-spike conversion produces spike trains required by neuromorphic devices. The most common signal-to-spike converter is the delta-sigma [538] which is widely applied in biomedical applications, thanks to its lower circuit complexity compared to multi-bit ADCs. However, the delta-modulator generates a high sampling rate and larger data size that can easily push the neurons' firing rate into saturation, making them insensitive to further input variations. Furthermore, SNNs for EMG classification should be optimized and learning algorithms could make them adaptable to different patients. Embedded decomposition sEMG decomposition into spike trains is generally based on shape-based algorithms, also called template matching [547] or blind source separation algorithms [548]. The decomposition of the complex sEMG is a computationally expensive procedure in a multidimensional constraint space. To run these algorithms on embedded platforms and in real-time it is imperative to (i) reduce the complexity and (ii) optimize it for the selected digital embedded architecture (e.g. PULP platform [549]) and exploit its hardware capabilities. The extracted spike trains are then sent to a neuromorphic chip, creating a hybrid digital-analog framework for spike encoding low-power computation.

Spike-based EMG decomposition to build a fully spiking pipeline that can be integrated into a single neuromorphic chip, the MUs identification algorithm needs to be translated into a spiking version. Embedding the entire process into a single chip that can be miniaturized and connected directly to the electrodes will allow online processing, which is optimal for real-time closed-loop applications and less vulnerable to interferences either caused by humans or the environment.

Smart electrodes another long-term game-changer would be the technological breakthroughs that will allow the single electrode to be able to record directly the activity of a single MU, removing the need for decomposition algorithms.

29.4. Concluding remarks

The need of improving myoelectric prosthetic control to increase the life quality of the patient poses new challenges for implementing real-time, compact, and low-power EMG processing systems. A wearable device based on neuromorphic technology can enable *in situ* EMG signal processing and decomposition, without information transfer and external computation. In particular, mixed-signal SNNs implemented on neuromorphic processors can be integrated directly with the sensors to extract temporal data streams in real-time with lowpower consumption.

This roadmap presents the specific case of prosthetic control, nevertheless, the development of this technology could reveal useful to more applications where continuous monitoring is required. In clinical settings, continuous monitoring of EMG signals can be utilized to detect degenerative diseases of motorneurons [550] even for very large time spans such as weeks or months. In rehabilitation, EMG can be used as feedback to adapt the patient training accordingly to its muscular status, after a stroke or neurological impairments [551].

With the current rate of technological and computational improvements the proposed objectives could be realistically achieved within a decade. If successfully executed, this roadmap will bring technology that will improve the quality of life for amputees and patients with motorneuron diseases.

30. Collaborative autonomous systems

Silvia Tolu and Roberto Galeazzi

Technical University of Denmark (DTU)

30.1. Status

Collaborative autonomous systems (CAS) (see figure 42) are entities that can cooperate among themselves and with humans, with variable levels of human intervention (depending on the level of autonomy) in performing complex tasks in unknown environments. Their behaviour is driven by the availability of perception, communication, cognitive and motor skills and improved computational capabilities (on/off-board systems). The high level of autonomy enables the execution of dependable actions under changing internal or external conditions. Therefore, CAS are expected to be able to: (1) perceive and understand their own condition and the environment they operate in; (2) dependably interact with the physical world despite of sudden changes; (3) intelligently evolve through learning and adaptation to unforeseen operational conditions; (4) self-decide their actions based on their understanding of the environment.


Currently, CAS (e.g., collaborative robots–cobots) show limited performances when accomplishing physical interaction tasks in complex scenarios [552]. Recent studies have demonstrated that autonomous robots can outperform the task they are programmed for, but they are limited in the ability to adapt to unexpected situations [553] and to different levels of human-robot cooperation [552]. These limitations are mainly due to the lack of generalization capabilities, i.e., cobots cannot transfer knowledge across multiple situations (environments, tasks, and interactions). One of the most viable pathways to solve this issue is to build intelligent autonomous cobots by incorporating AI-based methods into the control systems [554]. These bio-inspired controllers [555] allow taking a different perspective from the classical control approaches, which require a deeper understanding of the mechanics of the interactions and of the intrinsic limitations of the systems beforehand. Main current research directions [556] are focused on the understanding of the biological working principles of the central nervous system in order to build innovative neuromorphic computing algorithms and hardware that will bring significant advances in this field; in particular, they will provide computational efficiency and powerful control strategies for robust and adaptive behaviours.

In the next decades, there will be significant developments in CAS related to self-capabilities such as selfinspection, -configuration, -adaptation, -healing, -optimization, -protection, and -assembly. This will be a great enabler of systems acting in real-world unstructured scenarios, such as in remote applications (deep sea or space), in hazard situations (disasters), in healthcare interventions (assistive, rehabilitation, or diagnosis), and in proximity to people.

30.2. Current and future challenges

Several fundamental challenges demand to be addressed to enable the deployment of heterogenous autonomous systems able to collaborate towards the achievement of common mission objectives. These challenges span across different research topics including online mission planning and execution for multi-agent systems under uncertainty. Future mission planners [557] should integrate several factors to determine the optimal allocation of agents to the fulfilment of the mission tasks. These factors, among many, include energy availability and depletion rates, physical capabilities of the agents, probability of failures, and amount of collaboration needed. The mission execution demands the development of a revolutionary control paradigm that enables true collaboration among CAS with different functionalities. Cooperative control [558] has been so far limited to consensus and synchronization to enable the coordinated dynamic evolution of mostly homogeneous multi-agent systems to perform the same type of actions. The execution of tasks in uncertain environments calls for robust learning/adaptation methods to enable baseline control systems to ensure robust cooperation and coordination of heterogeneous multi-agent robots in various real applications [558]. Another big question is how to endow CAS with a high-level of **fault tolerance** capabilities in order to ensure dependability under a wide variety of operational conditions [559]. Despite the large research effort pursued by the community over the past four decades, condition monitoring and fault tolerant control are lacking efficiency due to the ever-increasing complexity of the systems.

Future work will aim to provide insights about how a CAS will show **robust, compliant, and intelligent physical interactions** with the environment, human beings, or other systems. In this regard, real-time, energy-efficient computing is required to advance the type of primitive collaborations that are achievable so far. With this aim, systems should be equipped with small processors able to ensure low energy consumption, and, at the same time, increase the memory bandwidth. Current alternatives (e.g., multicore central processing units—CPUs [560], new graphical processors—GPUs [561], parallel processing core—SpiNNaker [562]) still



suffer from an extremely high energy demand that is not sustainable and they cannot be easily scaled. Additionally, a limited number of processes can run simultaneously, and the speed of the response is still low. Consequently, new neuromorphic architectures are the most promising alternatives to address the increasing demand to create CAS able of a seamless interaction with human beings.

30.3. Advances in science and technology to meet challenges

In this section, we discuss the foremost advances in science and technology that will address the main aforementioned challenges.

Mission planning. Novel AI-based heuristic methods will be developed to equip mission planners with key functionalities that will increase the value for the human operators. These include: the close-loop decomposition of missions to achieve an adaptive task allocation by leveraging information gathered at mission execution; automated survivability prediction to assess the likelihood of vehicle loss based on faults and failures occurred in past missions; automated reliability assessment to forecast the probability of mission failure based on past missions' information; automated learning from previous missions' performance to tune the future missions' parameters; inclusion of services to extend the mission endurance [563, 564].

Fault-tolerant and cooperative control. Paradigms based on cooperation will be created to fulfil the advances in multi-agent systems [559]. Cooperation among agents offers the possibility of achieving fault-tolerance towards sensors and actuators faults through the design of diagnostic solutions that leverage shared proprioceptive and exteroceptive information. Prescribe-time fault tolerant cooperative control solutions for safety critical cyber-physical systems will be achieved; these will provide the basis for efficient fault-tolerant algorithms able to trade-off between fast convergence and acceptable fault-tolerance performance.

Robust, compliant and intelligent physical interactions. New physical mechanisms will be designed to provide passive properties to the system, to increase the physical interaction performances, and include advanced control aspects for achieving simultaneous robustness and compliance. The advances in neuro-robotics and neuromorphic computing will influence the development of the next generation of intelligent agents [565]. Current neuromorphic computing systems already exploit learning and adaptive skills in systems compared to conventional von Neumann machines thanks to non-volatile memories and power efficiency performance [566]. However, new types of sensors and actuators will be introduced to enhance the cognitive and learning functionalities of the systems and deal with safety and robustness concerns. Advanced bio-inspired platforms, e.g., brain-on-the-chip devices, will be designed for processing complex brain-inspired computing techniques that will support autonomy, more connectivity, increased decentralization, and high-performance computing. Indeed, neuromorphic technologies will be able to process complex unstructured data and learn to self-respond to external unknown stimuli enabling their use in critical edge applications, for example in autonomous navigation, human-machine interactions and smart healthcare markets.

Finally, innovative applications could be generated through the development of self-reconfigurable modular CAS, systems able to adapt their morphology and functionality to varied environments including unforeseen conditions [567]. This will require self-learning capabilities to develop new knowledge and to decide upon the previous accumulated experience.

30.4. Concluding remarks

This paper has presented the future perspectives of CAS and the main challenges and research issues that need to be addressed toward their realization. Further to these scientific and technological challenges, there are ethical, social, and legal issues when realising CAS, though these are beyond the scope of this article.

CAS working alongside humans have already been deployed and they support humans' work ensuring high productivity, speed, and accuracy [568]; they also relieve us of many heavy and time-consuming tasks and reduce the overall risk of collisions. CAS provide an economically viable entry-point to automation of processes, i.e., accelerated testing scenarios on products, environmental impacts. Fusion of fundamental and applied research in both technical and natural sciences will facilitate the development of new theoretical frame-works for the design of intelligent CAS. Multiple disciplines will be merged to pursue a systematic innovation within cyber-physical systems with variable level of autonomy and cooperation; the use of AI and Internet of Everything technologies future proofs the system to address changing market demands and expectations in several technological areas. Applications will be many and varied including, and not limited to, manufacturing, health care, inspection and maintenance, precision farming, autonomous marine operations, and education.

Section 5. Ethics

31. The ethics of developing neuromorphic technology

Martin Ejsing Christensen¹ and Sune Holm²

¹The Danish Council on Ethics ²University of Copenhagen E-mail: mech@dketik.dk and suneh@ifro.ku.dk

Like the development of other forms of AI, the development of neuromorphic technology may raise a number of ethical questions [569, 570] (figure 43).

One issue concerns privacy and surveillance. The development of most forms of AI depends upon access to data, and as far as these data can be seen as private or personally identifiable, it raises a question about when it is (ethically) defensible to use such data. On the one hand, some argue that persons have a right to be let alone and exercise full control over information about themselves, so that any use of such data presupposes fully informed consent. On the other hand, others recognize the importance of privacy but argue that it may sometimes be outweighed by the fact that reliable applications for the good of everyone presuppose access to high quality representative data [571].


Another issue concerns opacity. Many forms of AI support decision making based on complex patterns extracted from huge data sets. Often, however, it will be impossible not only for the person who makes the final decision but also for the developer to know what the system's recommendations are based on and it is in this sense that it is said to be opaque. For some such opacity does not matter as long as there are independent ways of verifying that the system delivers an accurate result, but others argue that it is important that the system is explainable [572]. In this way, a tension is often created between accuracy and transparency, and what the right trade-off is may often depend upon the concrete context.

Opacity is closely connected with the question of bias since opacity may hide certain biases. There are different forms of bias but in general, bias arises when automated AI decision support systems are based on data that is not representative of all the individuals that the system supports decisions in relation to [573]. There are different opinions as to when the existence of bias in automated decision support systems poses a serious problem. Some argue that 'traditional' unsupported human decision-making is biased, too, and that the existence of bias in automated AI decision support systems only pose a serious problem if the bias is more significant than the pre-existing human bias. Others argue that features such as opacity or the lack of suitable institutional checks and balances may tend to make the existence of bias in automated decision support systems more problematic than 'ordinary' human bias [574]. A separate problem is created by the fact that it sometimes will be easier to identify and quantify bias in AI systems than in humans, making a direct comparison more difficult.

The development of forms of AI based on neuromorphic technology also raises questions about manipulation of human behavior, online as well as offline. One context in which such questions arise is advertising and political campaigning, where AI generated deep knowledge about individuals' preferences and beliefs, which may be used to influence them in a way that escapes the individuals' own awareness. Similar issues may also arise in connection with other forms of AI such as chatbots and care or sex robots that simulate certain forms of human behavior without being 'the real deal'. Even if persons develop some form of emotional attachment to such systems, some argue that there is something deeply problematic and deceptive about such systems [575], while others point out that there is nothing intrinsically wrong with such systems as long as they help satisfy human desires [576]. If, as described in section 4.1, neuromorphic technologies will make it possible for robots to move from extremely controlled environments to spaces where they collaborate with humans and exhibit continuous learning and adaptation, it may make such questions more pressing.

A distinct set of issues are raised by the possibility of developing AI systems that do not just support human decision making but operate in a more or less autonomous way such as 'self-driving' cars and autonomous weapons. One question that such systems raise concerns the way in which they should be programmed in order to make sure that they make ethically justifiable decisions (in most foreseeable situations). Another question concerns how responsibility and risk should be distributed in the complex social system they are a part of. If, as described in section 4.2, neuromorphic engineering offers the kind of technological leaps required for achieving truly autonomous vehicles, the development of neuromorphic technologies may make such questions more pressing than at present.

A distinct issue relates to sustainability. As pointed out in the introduction, 5%–15% of the world's energy is spent in some form of data manipulation (transmission or processing), and as long as a substantial amount of that energy comes from sources that contribute to climate change through the emission of greenhouse gases,



it raises a question as to whether all that data manipulation is really necessary or could be done in a more energy efficient way. And in so far as neuromorphic technologies, as e.g. pointed out in section 28, shows a reduction of up to three orders of magnitude in terms of power consumption compared to state-of the-art ML approaches, it seems to provide robust ethical support for the development of neuromorphic technologies.

As mentioned in the beginning of this section, the ethical questions raised by the development of neuromorphic technology is not unique to this technology but related to the development of AI as such. The successful development of neuromorphic technology may make some of the issues more pressing, and a central task for future work on the ethics of neuromorphic technology will, accordingly, be to inquire into the exact way in which the issues are raised by the development of neuromorphic technology. But the existing forms of AI already raise many of the questions described so far. Besides these questions, however, the development of neuromorphic technology (as well as other forms of AI) may also raise a number of questions that are more speculative either because it is unclear whether the development will take place, when it will happen or what the precise consequences will be.

One such issue has to do with automation and unemployment. AI systems have already replaced humans in certain job functions (e.g., customer service), but it has been suggested that most job functions will be affected by the development of AI at one point [577]. Because such a development has the potential to disrupt the social order (e.g., through mass unemployment) it raises an important ethical (and political) question as to how AI systems should be introduced into society [578].

Another more speculative issue relates to artificial moral agents and so-called robot rights. If the development of neuromorphic (and other) forms of AI leads to the creation of systems that possess some or all the traits that make us ascribe rights and responsibilities to humans, it may thus raise a question about whether such rights and responsibilities should be ascribed to artificially intelligent systems [579, 580].

Thirdly, some have also pointed out that the development of neuromorphic (and other) forms of AI may create issues related to the so-called singularity. The idea is that the technological development may lead to the creation of general forms of AI that surpass the human level of intelligence and then begin to control the further development of AI in ways that may not be in the interests of the human species and perhaps even threaten its very existence. Whether such a scenario is likely has been questioned [581], but some argue that even a slight risk should be taken serious given the potentially devastating consequences [582].

No matter what one thinks is the right answer to the ethical questions raised by the development of neuromorphic technology, it is, finally, worth noticing that it still leaves an important practical question: how best to make sure that the actual development and implementation of neuromorphic technology will take place in an ethically defensible way. For some questions, governmental regulation may be the best means. For others, the best solution may be to trust the community of developers to make the right, valuebased decisions when designing systems, while some questions, perhaps, should be left to the enlightened citizenry. In the end, however, it will probably be up to an inquiry into the concrete situation to decide when one or the other approach—or combination of approaches—provides the best means of securing an ethically defensible development of neuromorphic technology.

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

ORCID iDs

Dennis V Christensen D https://orcid.org/0000-0003-0048-7595 Abu Sebastian b https://orcid.org/0000-0001-5603-5243 Manuel Le Gallo D https://orcid.org/0000-0003-1600-6151 Stefan Slesazeck D https://orcid.org/0000-0002-0414-0321 Thomas Mikolajick b https://orcid.org/0000-0003-3814-0378 Sabina Spiga D https://orcid.org/0000-0001-7293-7503 Ilia Valov D https://orcid.org/0000-0002-0728-7214 Gianluca Milano D https://orcid.org/0000-0002-1983-6516 Carlo Ricciardi D https://orcid.org/0000-0002-4703-7949 Shi-Jun Liang b https://orcid.org/0000-0002-3466-8063 Feng Miao D https://orcid.org/0000-0002-0962-5424 Mario Lanza D https://orcid.org/0000-0003-4756-8632 Scott T Keene b https://orcid.org/0000-0002-6635-670X Alice Mizrahi D https://orcid.org/0000-0003-2043-049X J Joshua Yang D https://orcid.org/0000-0001-8242-7531 Giacomo Indiveri D https://orcid.org/0000-0002-7109-1689 John Paul Strachan D https://orcid.org/0000-0002-5718-7924 Elisa Vianello D https://orcid.org/0000-0002-8868-9951 Srikanth Ramaswamy b https://orcid.org/0000-0001-6642-7136 Gabriella Panuccio D https://orcid.org/0000-0002-4059-7193 Mufti Mahmud b https://orcid.org/0000-0002-2037-8348 Elisa Donati D https://orcid.org/0000-0002-8091-1298 Silvia Tolu D https://orcid.org/0000-0003-1825-8440 Daniele Ielmini D https://orcid.org/0000-0002-1853-1614 N Pryds D https://orcid.org/0000-0002-5718-7924

References

- [1] Vidal J 2017 'Tsunami of Data' Could Consume One Fifth of Global Electricity by 2025 (Climate Home News)
- [2] Mead C 1990 Neuromorphic electronic systems Proc. IEEE 78 1629-36
- Chicca E, Stefanini F, Bartolozzi C and Indiveri G 2014 Neuromorphic electronic circuits for building autonomous cognitive systems Proc. IEEE 102 1367–88
- [4] Chicca E and Indiveri G 2020 A recipe for creating ideal hybrid memristive-CMOS neuromorphic processing systems Appl. Phys. Lett. 116 120501
- [5] LeCun Y, Bengio Y and Hinton G 2015 Deep learning Nature 521 436-44
- [6] Maass W 1997 Networks of spiking neurons: the third generation of neural network models Neural Netw. 10 1659-71
- [7] Yole 2021 Neuromorphic computing and sensing 2021 Yole Reports www.yole.fr
- [8] Zidan M A, Strachan J P and Lu W D 2018 The future of electronics based on memristive systems Nat. Electron. 1 22-9
- [9] Chua L 1971 Memristor-the missing circuit element *IEEE Trans. Circuit Theory* 18 507–19
- [10] Strukov D B, Snider G S, Stewart D R and Williams R S 2008 The missing memristor found Nature 453 80-3
- [11] Yang J J, Strukov D B and Stewart D R 2013 Memristive devices for computing Nat. Nanotechnol. 8 13–24
- [12] Dittmann R and Strachan J P 2019 Redox-based memristive devices for new computing paradigm APL Mater. 7 110903
- [13] Ielmini D and Wong H-S P 2018 In-memory computing with resistive switching devices Nat. Electron. 1 333-43
- [14] Li C et al 2018 Analogue signal and image processing with large memristor crossbars Nat. Electron. 1 52-9
- [15] Ielmini D, Wang Z and Liu Y 2021 Brain-inspired computing via memory device physics *APL Mater.* **9** 050702
- [16] Milano G, Pedretti G, Fretto M, Boarino L, Benfenati F, Ielmini D, Valov I and Ricciardi C 2020 Brain-inspired structural plasticity through reweighting and rewiring in multi-terminal self-organizing memristive nanowire networks Adv. Intell. Syst. 2 2000096
- [17] Shastri B J, Tait A N, Ferreira de Lima T, Pernice W H P, Bhaskaran H, Wright C D and Prucnal P R 2021 Photonics for artificial intelligence and neuromorphic computing *Nat. Photon.* 15 102–14
- [18] Marković D and Grollier J 2020 Quantum neuromorphic computing Appl. Phys. Lett. 117 150501
- [19] Thorpe S, Fize D and Marlot C 1996 Speed of processing in the human visual system Nature 381 520-2
- [20] Le Gallo M and Sebastian A 2020 An overview of phase-change memory device physics J. Phys. D: Appl. Phys. 53 213002

- [21] Sebastian A, Le Gallo M, Khaddam-Aljameh R and Eleftheriou E 2020 Memory devices and applications for in-memory computing *Nat. Nanotechnol.* **15** 529–44
- [22] Giannopoulos I, Singh A, Le Gallo M, Jonnalagadda V P, Hamdioui S and Sebastian A 2020 Adv. Intell. Syst. 2 2000141
- [23] Karunaratne G, Le Gallo M, Cherubini G, Benini L, Rahimi A and Sebastian A 2020 In-memory hyperdimensional computing Nat. Electron. 3 327–37
- [24] Joshi V et al 2020 Accurate deep neural network inference using computational phase-change memory Nat. Commun. 11 2473
- [25] Tsai H, Ambrogio S, Narayanan P, Shelby R M and Burr G W 2018 Recent progress in analog memory-based accelerators for deep learning J. Phys. D: Appl. Phys. 51 283001
- [26] Nandakumar S R *et al* 2020 Mixed-precision deep learning based on computational memory *Front. Neurosci.* 14 406
- [27] Sebastian A, Le Gallo M, Burr G W, Kim S, BrightSky M and Eleftheriou E 2018 Tutorial: brain-inspired computing using phase-change memory devices *J. Appl. Phys.* **124** 111101
- [28] Ambrogio S, Ciocchini N, Laudato M, Milo V, Pirovano A, Fantini P and Ielmini D 2016 Unsupervised learning by spike timing dependent plasticity in phase change memory (PCM) synapses Front. Neurosci. 10 56
- [29] Tuma T, Pantazi A, Le Gallo M, Sebastian A and Eleftheriou E 2016 Stochastic phase-change neurons Nat. Nanotechnol. 11 693
- [30] Zuliani P, Conte A and Cappelletti P 2019 The PCM way for embedded non volatile memories applications 2019 Symp. VLSI Technology T192–3
- [31] Boniardi M, Ielmini D, Lavizzari S, Lacaita A L, Redael A and Pirovano A 2010 Statistics of resistance drift due to structural relaxation in phase-change memory arrays *IEEE Trans. Electron Devices* 57 2690–6
- [32] Crespi L et al 2015 Modeling of atomic migration phenomena in phase change memory devices 2015 IEEE Int. Memory Workshop (IMW) 1–4
- [33] BrightSky M et al 2015 Crystalline-as-deposited ALD phase change material confined PCM cell for high density storage class memory 2015 IEEE Int. Electron Devices Meeting (IEDM) 13.6.-4
- [34] Kau D et al 2009 A stackable cross point phase change memory 2009 IEEE Int. Electron Devices Meeting (IEDM) 127.1.-4
- [35] Arnaud F et al 2020 High density embedded PCM cell in 28 nm FDSOI technology for automotive micro-controller applications 2020 IEEE Int. Electron Devices Meeting (IEDM) 124.2.–4
- [36] Redaelli A, Pellizer F and Pirovano A 2012 Phase change memory device for multibit storage EP Patent 2034536
- [37] Koelmans W W, Sebastian A, Jonnalagadda V P, Krebs D, Dellmann L and Eleftheriou E 2015 Projected phase-change memory devices Nat. Commun. 6 8181
- [38] Giannopoulos I et al 2018 8 bit precision in-memory multiplication with projected phase-change memory 2018 IEEE Int. Electron Devices Meeting (IEDM) 127.7.–4
- [39] Ding K et al 2019 Phase-change heterostructure enables ultralow noise and drift for memory operation Science 366 210-5
- [40] Salinga M et al 2018 Monatomic phase change memory Nat. Mater. 17 681-5
- [41] Boybat I et al 2018 Neuromorphic computing with multimemristive synapses Nat. Commun. 9 2514
- [42] Rios C, Youngblood N, Cheng Z, Le Gallo M, Pernice W H, Wright C D, Sebastian A and Bhaskaran H 2019 In-memory computing on a photonic platform *Sci. Adv.* 5 eaau5759
- [43] Feldmann J et al 2021 Parallel convolutional processing using an integrated photonic tensor core Nature 589 52-8
- [44] Cheng Z, Milne T, Salter P, Kim J S, Humphrey S, Booth M and Bhaskaran H 2021 Antimony thin films demonstrate programmable optical nonlinearity Sci. Adv. 7 eabd7097
- [45] Valasek J 1921 Piezo-electric and allied phenomena in Rochelle salt Phys. Rev. 17 475
- [46] Buck D A 1952 Ferroelectrics for digital information storage and switching MIT Digital Computer Laboratory Report
- [47] Bondurant D 1990 Ferroelectronic RAM memory family for critical data storage Ferroelectrics 112 273-82
- [48] Mikolajick T, Schroeder U and Slesazeck S 2020 The past, the present, and the future of ferroelectric memories *IEEE Trans. Electron Devices* **67** 1434–43
- [49] Ross I 1957 Semiconductive translating device USA Patent 2791760A
- [50] Zhang X, Takahashi M, Takeuchi K and Sakai S 2012 64 kbit ferroelectric-gate-transistor-integrated NAND flash memory with 7.5 V program and long data retention Japan. J. Appl. Phys. 51 04DD01
- [51] Esaki L, Laibowitz R and Stiles P 1971 Polar Switch (IBM Technical Disclosure Bulletin) vol 13 p 2161
- [52] Garcia V, Fusil S, Bouzehouane K, Enouz-Vedrenne S, Mathur N, Barthelemy A and Bibes M 2009 Giant tunnel electroresistance for non-destructive readout of ferroelectric states *Nature* 460 81–4
- [53] Böscke T, Müller J, Bräuhaus D, Schröder U and Böttger U 2011 Ferroelectricity in hafnium oxide thin films Appl. Phys. Lett. 99 102903
- [54] Max B, Hoffmann M, Mulaosmanovic H, Slesazeck S and Mikolajick T 2020 Hafnia-based double-layer ferroelectric tunnel junctions as artificial synapses for neuromorphic computing ACS Appl. Electron. Mater. 2 4023–33
- [55] Mulaosmanovic H, Chicca E, Bertele M, Mikolajick T and Slesazeck S 2018 Mimicking biological neurons with a nanoscale ferroelectric transistor Nanoscale 10 21755–63
- [56] Beyer S et al 2020 FeFET: a versatile CMOS compatible device with game-changing potential IEEE Int. Memory Workshop (IMW) 1–4
- [57] Sally A 2004 Reflections on the memory wall Proc. Conf. Comput. Front. p 162
- [58] Okuno J et al 2020 SoC compatible 1T1C FeRAM memory array based on ferroelectric Hf_{0.5}Zr_{0.5}O₂ Symp. VLSI Technology 1–4
- [59] Pešić M et al 2016 Physical mechanisms behind the field-cycling behavior of HfO₂ based ferroelectric capacitors Adv. Funct. Mater. 26 4601–12
- [60] Slesazeck S and Mikolajick T 2019 Nanoscale resistive switching memory devices: a review Nanotechnology 30 352003
- [61] Deng S, Jiang Z, Dutta S, Ye H, Chakraborty W, Kurinec S, Datta S and Ni K 2020 Examination of the interplay between polarization switching and charge trapping in ferroelectric FET *Int. Electron Device Meeting (IEDM)* (San Francisco)
- [62] Wei Y *et al* 2018 A rhombohedral ferroelectric phase in epitaxially strained $Hf_{0.5}Zr_{0.5}O_2$ thin films *Nat. Mater.* **17** 1095–100 [63] Pesic M, Schroeder U, Slesazeck S and Mikolajick T 2018 Comparative study of reliability of ferroelectric and anti-ferroelectric
- memories IEEE Trans. Device Mater. Relib. 18 154–62
 [64] Eichtnar S. Wolff N. Locht F. Kianla Land Warner P. 2010 AlSo N. e. H. V. somiconductor based formalactric L. Appl. Phys. 125
- [64] Fichtner S, Wolff N, Lofink F, Kienle L and Wagner B 2019 AlScN: a III–V semiconductor based ferroelectric J. Appl. Phys. 125 114103
- [65] Spiga S, Sebastian A, Querlioz D and Rajendran B 2020 Memristive Device for Brain-Inspired Computing: From Materials, Devices, and Circuits to Applications-Computational Memory, Deep Learning and Spiking Neural Networks (Oxford: Woodhead Publishing)
- [66] Waser R, Dittmann R, Staikov G and Szot K 2009 Redox-based resistive switching memories— nanoionic mechanisms, prospects, and challenges Adv. Mater. 21 2632–63

- [67] Brivio S, Ly D R B, Vianello E and Spiga S 2021 Non-linear memristive synaptic dynamics for efficient unsupervised learning in spiking neural networks *Front. Neurosci.* **15** 580909
- [68] Payvand M, Demirag Y, Dalgaty T, Vianello E and Indiveri G 2020 Analog weight updates with compliance current modulation of binary ReRAMs for on-chip learning pp 1–5
- [69] Covi E, George R, Frascaroli J, Brivio S, Mayr C, Mostafa H, Indiveri G and Spiga S 2018 Spike-driven threshold-based learning with memristive synapses and neuromorphic silicon neurons J. Phys. D: Appl. Phys. 51 344003
- [70] Zhang W, Gao B, Tang J, Li X, Wu W, Qian H and Wu H 2019 Analog-type resistive switching devices for neuromorphic computing *Phys. Status Solidi RRL* 13 1900204
- [71] Covi E, Brivio S, Serb A, Prodromakis T, Fanciulli M and Spiga S 2016 Analog memristive synapse in spiking networks implementing unsupervised learning *Front. Neurosci.* 10 6–13
- [72] Mochida R et al 2018 A 4M synapses integrated analog ReRAM based 66.5 TOPS/W neuralnetwork processor with cell current controlled writing and flexible network architecture 2018 IEEE Symp. VLSI Technology (Honolulu, HI, USA 18–22 June 2018)
- [73] Valentian A, Rummens F, Vianello E, Mesquida T, de Boissac C L M, Bichler O and Reita C 2019 Fully integrated spiking neural network with analog neurons and RRAM synapses 2019 IEEE Int. Electron Devices Meeting (IEDM)
- [74] Yao P, Wu H, Gao B, Tang J, Zhang Q, Zhang W, Yang J J and Qian H 2020 Fully hardware-implemented memristor convolutional neural network *Nature* 577 641–6
- [75] You T et al 2015 Engineering interface-type resistive switching in BiFeO₃ thin film switches by Ti implantation of bottom electrodes Sci. Rep. 5 18623
- [76] Moon K, Fumarola A, Sidler S, Jang J, Narayanan P, Shelby R M, Burr G W and Hwang H 2018 Bidirectional non-filamentary RRAM as an analog neuromorphic synapse: I. Al/Mo/Pr_{0.7}Ca_{0.3}MnO₃ material improvements and device measurements *IEEE J*. *Electron Devices Soc.* 6 146–55
- [77] Gao L, Wang I-T, Chen P-Y, Vrudhula S, Seo J-S, Cao Y, Hou T-H and Yu S 2015 Fully parallel write/read in resistive synaptic array for accelerating on-chip learning *Nanotechnology* **26** 455204
- [78] Govoreanu B *et al* 2015 a-VMCO: a novel forming-free, selfrectifying, analog memory cell *IEEE 2015 Symp. VLSI Technology* T132–3
- [79] Kim S et al 2019 Metal-oxide based, CMOS compatible ECRAM for deep learning accelerator 2019 IEEE Int. Electron Device Meeting (San Francisco, USA 07 November 2019)
- [80] Li Y et al 2020 Filament-free bulk resistive memory enables deterministic analogue switching Adv. Mater. 32 2003984
- [81] Bengel C, Siemon A, Cüppers F, Hoffmann-Eifert S, Hardtdegen A, von Witzleben M, Helllmich L, Waser R and Menzel S 2020 Variability-aware modeling of filamentary oxide based bipolar resistive switching cells using SPICE level compact models TCAS vol 67 p 46184630
- [82] Puglisi F M, Larcher L, Padovani A and Pavan P 2016 Bipolar resistive RAM based on HfO₂: physics, compact modeling, and variability control *IEEE J. Emerg. Sel. Top. Circuits Syst.* 6 171–84
- [83] Zhao M, Gao B, Tang J, Qian H and Wu H 2020 Reliability of analog resistive switching memory for neuromorphic computing Appl. Phys. Rev. 7 011301
- [84] Cueppers F, Menzel S, Bengel C, Hardtdegen A, von Witzleben M, Boettger U, Waser R and Hoffmann-Eifert S 2019 Exploiting the switching dynamics of HfO₂-based ReRAM devices for reliable analog memristive behavior APL Mater. 7 091105
- [85] Menzel S, Böttger U, Wimmer M and Salinga M 2015 Physics of the switching kinetics in resistive memories Adv. Funct. Mater. 25 6306–25
- [86] Kozicki M N and West W C 1998 Programmable metallization cell structure and method of making same US Patent 5761115A
- [87] Swaroop B, West W C, Martinez G, Kozicki M N and Akers L A 1998 Programmable current mode Hebbian learning neural network using programmable metallization cell *IEEE Int. Symp. Circuits and Systems (ISCAS)* 33–6
- [88] Hasegawa T, Terabe K, Tsuruoka T and Aono M 2012 Atomic switch: atom/ion movement controlled devices for beyond von Neumann computers Adv. Mater. 24 252–67
- [89] Valov I, Linn E, Tappertzhofen S, Schmelzer S, van den Hurk J, Lentz F and Waser R 2013 Nanobatteries in redox-based resistive switches require extension of memristor theory *Nat. Commun.* 4 1771
- [90] Kozicki M N, Mitkova M and Valov I 2016 Electrochemical Metallization Memories (New York: Wiley) pp 483-514
- [91] Raeis-Hosseini N and Lee J-S 2017 Resistive switching memory using biomaterials J. Electroceram. 39 223–38
- [92] Gao S, Yi X, Shang J, Liu G and Li R-W 2019 Organic and hybrid resistive switching materials and devices Chem. Soc. Rev. 48 1531–65
- [93] Midya R et al 2017 Anatomy of Ag/Hafnia-based selectors with 10¹⁰ nonlinearity Adv. Mater. 29 1604457
- [94] Gonzalez-Velo Y, Barnaby H J and Kozicki M N 2017 Review of radiation effects on ReRAM devices and technology *Semicond*. *Sci. Technol.* **32** 083002
- [95] Chen W, Chamele N, Gonzalez-Velo Y, Barnaby H J and Kozicki M N 2017 Low-temperature characterization of Cu–Cu: silica-based programmable metallization cell *IEEE Electron Device Lett.* 38 1244–7
- [96] Wang Z *et al* 2016 Memristors with diffusive dynamics as synaptic emulators for neuromorphic computing *Nat. Mater.* 16 101–8
 [97] Ohno T, Hasegawa T, Tsuruoka T, Terabe K, Gimzewski J K and Aono M 2011 Short-term plasticity and long-term potentiation
- mimicked in single inorganic synapses *Nat. Mater.* **10** 591–5 [98] Cha J-H, Yang S Y, Oh J, Choi S, Park S, Jang B C, Ahn W and Choi S-Y 2020 Conductive-bridging random-access memories for
- emerging neuromorphic computing *Nanoscale* 12 14339–68
 [99] Valov I, Waser R, Jameson J R and Kozicki M N 2011 Electrochemical metallization memories-fundamentals, applications,
- [99] Valov I, waser K, Jameson J K and Kozicki M N 2011 Electrochemical metallization memories-rundamentals, applications, prospects Nanotechnology 22 254003
- [100] Valov I 2014 Redox-based resistive switching memories (ReRAMs): electrochemical systems at the atomic scale ChemElectroChem 1 26–36
- [101] Valov I 2017 Interfacial interactions and their impact on redox-based resistive switching memories (ReRAMs) Semicond. Sci. Technol. 32 093006
- [102] Lübben M, Cüppers F, Mohr J, von Witzleben M, Breuer U, Waser R, Neumann C and Valov I 2020 Design of defect-chemical properties and device performance in memristive systems Sci. Adv. 6 eaaz9079
- Belmonte A, Celano U, Degraeve R, Fantini A, Redolfi A, Vandervorst W, Houssa M, Jurczak M and Goux L 2015
 Operating-current dependence of the Cu-mobility requirements in oxide-based conductive-bridge RAM *IEEE Electron Device* Lett. 36 775–7
- [104] Yeon H et al 2020 Alloying conducting channels for reliable neuromorphic computing Nat. Nanotechnol. 15 574–9

- [105] Valov I and Tsuruoka T 2018 Effects of moisture and redox reactions in VCM and ECM resistive switching memories J. Phys. D: Appl. Phys. 51 413001
- [106] Kandel E R, Schwartz J H, Jessell T M, Siegelbaum S A and Hudspeth A J 2013 Principles of Neural Science 5th edn (New York: McGraw-Hill)
- [107] Xia Q and Yang J J 2019 Memristive crossbar arrays for brain-inspired computing Nat. Mater. 18 309–23
- [108] Stieg A Z, Avizienis A V, Sillin H O, Martin-Olmos C, Aono M and Gimzewski J K 2012 Emergent criticality in complex turing B-type atomic switch networks Adv. Mater. 24 286–93
- [109] Stieg A Z, Avizienis A V, Sillin H O, Martin-Olmos C, Lam M-L, Aono M and Gimzewski J K 2014 Self-organized atomic switch networks Japan. J. Appl. Phys. 53 01AA02
- [110] Diaz-Alvarez A, Higuchi R, Sanz-Leon P, Marcus I, Shingaya Y, Stieg A Z, Gimzewski J K, Kuncic Z and Nakayama T 2019 Emergent dynamics of neuromorphic nanowire networks Sci. Rep. 9 14920
- [111] Milano G, Porro S, Valov I and Ricciardi C 2019 Recent developments and perspectives for memristive devices based on metal oxide nanowires Adv. Electron. Mater. 5 1800909
- [112] Sillin H O, Aguilera R, Shieh H-H, Avizienis A V, Aono M, Stieg A Z and Gimzewski J K 2013 A theoretical and experimental study of neuromorphic atomic switch networks for reservoir computing *Nanotechnology* 24 384004
- [113] Aono M and Ariga K 2016 The way to nanoarchitectonics and the way of nanoarchitectonics Adv. Mater. 28 989–92
- [114] Loeffler A et al 2020 Topological properties of neuromorphic nanowire networks Front. Neurosci. 14 184
- [115] Manning H G et al 2018 Emergence of winner-takes-all connectivity paths in random nanowire networks Nat. Commun. 9 3219
- [116] Li Q *et al* 2020 Dynamic electrical pathway tuning in neuromorphic nanowire networks *Adv. Funct. Mater.* 30 2003679
 [117] Li Q, Diaz-Alvarez A, Tang D, Higuchi R, Shingaya Y and Nakayama T 2020 Sleep-dependent memory consolidation in a
- neuromorphic nanowire network ACS Appl. Mater. Interfaces 12 50573-80
- [118] Mallinson J B, Shirai S, Acharya S K, Bose S K, Galli E and Brown S A 2019 Avalanches and criticality in self-organized nanoscale networks Sci. Adv. 5 eaaw8438
- [119] Pike M D et al 2020 Atomic scale dynamics drive brain-like avalanches in percolating nanostructured networks Nano Lett. 20 3935–42
- [120] Fu K *et al* 2020 Reservoir computing with neuromemristive nanowire networks 2020 Int. Joint Conf. Neural Networks (IJCNN) 1–8
- [121] Milano G et al 2021 In materia reservoir computing with a fully memristive architecture based on self-organizing nanowire networks Nat. Mater. 21 195–202
- [122] Usami Y et al 2021 In-materio reservoir computing in a sulfonated polyaniline network Adv. Mater. 33 2102688
- [123] Lilak S et al 2021 Spoken digit classification by in-materio reservoir computing with neuromorphic atomic switch networks Front. Nanotechnol. 3 1–11
- [124] Nirmalraj P N et al 2012 Manipulating connectivity and electrical conductivity in metallic nanowire networks Nano Lett. 12 5966–71
- [125] Sannicolo T et al 2018 Electrical mapping of silver nanowire networks: a versatile tool for imaging network homogeneity and degradation dynamics during failure ACS Nano 12 4648–59
- [126] Milano G et al 2020 Mapping time-dependent conductivity of metallic nanowire networks by electrical resistance tomography toward transparent conductive materials ACS Appl. Nano Mater. 3 11987–97
- [127] Diederichsen K M, Brow R R and Stoykovich M P 2015 Percolating transport and the conductive scaling relationship in lamellar block copolymers under confinement ACS Nano 9 2465–76
- [128] Zhou F and Chai Y 2020 Near-sensor and in-sensor computing Nat. Electron. 3 664-71
- [129] Mennel L, Symonowicz J, Wachter S, Polyushkin D K, Molina-Mendoza A J and Mueller T 2020 Ultrafast machine vision with 2D material neural network image sensors *Nature* 579 62–6
- [130] Wang C-Y et al 2020 Gate-tunable van der Waals heterostructure for reconfigurable neural network vision sensor Sci. Adv. 6 eaba6173
- [131] Sun L et al 2019 Self-selective van der Waals heterostructures for large scale memory array Nat. Commun. 10 3161
- [132] Pan C et al 2020 Reconfigurable logic and neuromorphic circuits based on electrically tunable two-dimensional homojunctions Nat. Electron. 3 383–90
- [133] Liu C, Yan X, Song X, Ding S, Zhang D W and Zhou P 2018 A semi-floating gate memory based on van der Waals heterostructures for quasi-non-volatile applications *Nat. Nanotechnol.* 13 404–10
- [134] Tian H et al 2017 Emulating bilingual synaptic response using a junction-based artificial synaptic device ACS Nano 11 71567163
- [135] Sangwan V K, Jariwala D, Kim I S, Chen K-S, Marks T J, Lauhon L J and Hersam M C 2015 Gate-tunable memristive phenomena mediated by grain boundaries in single-layer MoS₂ Nat. Nanotechnol. 10 403
- [136] Zhu J et al 2018 Ion gated synaptic transistors based on 2D van der Waals crystals with tunable diffusive dynamics Adv. Mater. 30 1800195
- [137] Zhu X, Li D, Liang X and Lu W D 2019 Ionic modulation and ionic coupling effects in MoS₂ devices for neuromorphic computing Nat. Mater. 18 141–8
- [138] Zhang F *et al* 2019 Electric-field induced structural transition in vertical MoTe₂- and Mo_{1-x} W_x Te₂-based resistive memories *Nat. Mater.* **18** 55–61
- [139] Wang C Y, Wang C, Meng F, Wang P, Wang S, Liang S J and Miao F 2020 2D layered materials for memristive and neuromorphic applications Adv. Electron. Mater. 6 1901107
- [140] Wang M et al 2018 Robust memristors based on layered two-dimensional materials Nat. Electron. 1 130
- [141] Shi Y et al 2018 Electronic synapses made of layered two-dimensional materials *Nat. Electron.* 1 458
- [142] Chen S et al 2020 Wafer-scale integration of two-dimensional materials in high-density memristive crossbar arrays for artificial neural networks Nat. Electron. 3 638–45
- [143] Ge R, Wu X, Kim M, Shi J, Sonde S, Tao L, Zhang Y, Lee J C and Akinwande D 2018 Atomristor: nonvolatile resistance switching in atomic sheets of transition metal dichalcogenides Nano Lett. 18 434–41
- [144] Hus S M et al 2021 Observation of single-defect memristor in an MoS₂ atomic sheet Nat. Nanotechnol. 16 58-62
- [145] Liu C, Wang L, Qi J and Liu K 2020 Designed growth of large-size 2D single crystals Adv. Mater. 32 2000046
- [146] Jang H, Liu C, Hinton H, Lee M H, Kim H, Seol M, Shin H J, Park S and Ham D 2020 An atomically thin optoelectronic machine vision processor Adv. Mater. 32 2002431
- [147] Wang S et al 2021 Networking retinomorphic sensor with memristive crossbar for brain-inspired visual perception Natl Sci. Rev. 8 nwaa172

- [148] Zeng F, Li S, Yang J, Pan F and Guo D 2014 Learning processes modulated by the interface effects in a Ti/conducting polymer/Ti resistive switching cell *RSC Adv.* 4 14822–8
- [149] Ouyang J, Chu C-W, Szmanda C R, Ma L and Yang Y 2004 Programmable polymer thin film and non-volatile memory device Nat. Mater. 3 918
- [150] Van De Burgt Y et al 2017 A non-volatile organic electrochemical device as a low-voltage artificial synapse for neuromorphic computing Nat. Mater. 16 414–8
- [151] Giovannitti A et al 2016 Controlling the mode of operation of organic transistors through side-chain engineering Proc. Natl Acad. Sci. USA 113 12017–22
- [152] Giovannitti A et al 2020 Energetic control of redox-active polymers toward safe organic bioelectronic materials Adv. Mater. 32 1908047
- [153] Go G-T, Lee Y, Seo D-G, Pei M, Lee W, Yang H and Lee T-W 2020 Achieving microstructure-controlled synaptic plasticity and long-term retention in ion-gel-gated organic synaptic transistors Adv. Intell. Syst. 2 2000012
- [154] Gkoupidenis P, Koutsouras D A and Malliaras G G 2017 Neuromorphic device architectures with global connectivity through electrolyte gating *Nat. Commun.* 8 15448
- [155] Kim Y et al 2018 A bioinspired flexible organic artificial afferent nerve Science 360 998-1003
- [156] Fuller E J et al 2019 Parallel programming of an ionic floating-gate memory array for scalable neuromorphic computing Science 364 570–4
- [157] Defranco J A, Schmidt B S, Lipson M and Malliaras G G 2006 Photolithographic patterning of organic electronic materials Org. Electron. 7 22–8
- [158] Zakhidov A A, Lee J-K, DeFranco J A, Fong H H, Taylor P G, Chatzichristidi M, Ober C K and Malliaras G G 2011 Orthogonal processing: a new strategy for organic electronics *Chem. Sci.* 2 1178–82
- [159] Keene S T, Melianas A, van de Burgt Y and Salleo A 2018 Mechanisms for enhanced state retention and stability in redox-gated organic neuromorphic devices Adv. Electron. Mater. 5 1800686
- [160] Melianas A et al 2020 Temperature-resilient solid-state organic artificial synapses for neuromorphic computing Sci. Adv. 6 eabb2958
- [161] Choi Y, Oh S, Qian C, Park J H and Cho J H 2020 Vertical organic synapse expandable to 3D crossbar array Nat. Commun. 11 4595
- [162] Spyropoulos G D, Gelinas J N and Khodagholy D 2019 Internal ion-gated organic electrochemical transistor: a building block for integrated bioelectronics Sci. Adv. 5 eaau7378
- [163] Bischak C G et al 2020 A reversible structural phase transition by electrochemically-driven ion injection into a conjugated polymer J. Am. Chem. Soc. 142 7434–42
- [164] Lenz J, del Giudice F, Geisenhof F R, Winterer F and Weitz R T 2019 Vertical, electrolyte-gated organic transistors show continuous operation in the MA cm⁻² regime and artificial synaptic behaviour *Nat. Nanotechnol.* 14 579–85
- [165] Shulaker M M et al 2015 Monolithic 3D integration of logic and memory: carbon nanotube FETs, resistive RAM, and silicon FETs Technical Digest—Int. Electron Devices Meeting IEDM 27.4.pp 1–4
- [166] Gumyusenge A et al 2018 Semiconducting polymer blends that exhibit stable charge transport at high temperatures Science 362 1131–4
- [167] Keene S T et al 2020 A biohybrid synapse with neurotransmitter-mediated plasticity Nat. Mater. 19 969–73
- [168] Grollier J, Querlioz D, Camsari K Y, Everschor-Sitte K, Fukami S and Stiles M D 2020 Neuromorphic spintronics Nat. Electron. 3 360–70
- [169] Ma Y, Miura S, Honjo H, Ikeda S, Hanyu T, Ohno H and Endoh T 2016 A 600 µW ultra-low-power associative processor for image pattern recognition employing magnetic tunnel junction-based nonvolatile memories with autonomic intelligent power-gating scheme Japan. J. Appl. Phys. 55 04EF15
- [170] Vincent A F et al 2015 Spin-transfer torque magnetic memory as a stochastic memristive synapse for neuromorphic systems IEEE Trans. Biomed. Circuits Syst. 9 166–74
- [171] Lequeux S, Sampaio J, Cros V, Yakushiji K, Fukushima A, Matsumoto R, Kubota H, Yuasa S and Grollier J 2016 A magnetic synapse: multilevel spin-torque memristor with perpendicular anisotropy Sci. Rep. 6 31510
- [172] Mansueto M *et al* 2019 Realizing an isotropically coercive magnetic layer for memristive applications by analogy to dry friction *Phys. Rev. Appl.* **12** 044029
- [173] Fukami S, Zhang C, DuttaGupta S, Kurenkov A and Ohno H 2016 Magnetization switching by spin-orbit torque in an antiferromagnet-ferromagnet bilayer system Nat. Mater. 15 535-41
- [174] Borders W A, Akima H, Fukami S, Moriya S, Kurihara S, Horio Y, Sato S and Ohno H 2016 Analogue spin–orbit torque device for artificial-neural-network-based associative memory operation *Appl. Phys. Express* 10 013007
- [175] Torrejon J et al 2017 Neuromorphic computing with nanoscale spintronic oscillators Nature 547 428–31
- [176] Romera M et al 2018 Vowel recognition with four coupled spin-torque nano-oscillators Nature 563 230
- [177] Mizrahi A, Hirtzlin T, Fukushima A, Kubota H, Yuasa S, Grollier J and Querlioz D 2018 Neural-like computing with populations of superparamagnetic basis functions Nat. Commun. 9 1533
- [178] Daniels M W, Madhavan A, Talatchian P, Mizrahi A and Stiles M D 2020 Energy-efficient stochastic computing with superparamagnetic tunnel junctions Phys. Rev. Appl. 13 034016
- [179] Borders W A, Pervaiz A Z, Fukami S, Camsari K Y, Ohno H and Datta S 2019 Integer factorization using stochastic magnetic tunnel junctions *Nature* 573 390–3
- [180] Yuasa S, Nagahama T, Fukushima A, Suzuki Y and Ando K 2004 Giant room-temperature magnetoresistance in single-crystal Fe/MgO/Fe magnetic tunnel junctions Nat. Mater. 3 868–71
- [181] Parkin S S P, Kaiser C, Panchula A, Rice P M, Hughes B, Samant M and Yang S-H 2004 Giant tunnelling magnetoresistance at room temperature with MgO (100) tunnel barriers *Nat. Mater.* 3 862–7
- [182] Zahedinejad M, Fulara H, Khymyn R, Houshang A, Fukami S, Kanai S, Ohno H and Åkerman J 2020 Memristive control of mutual SHNO synchronization for neuromorphic computing (arXiv:2009.06594)
- [183] Pinna D, Bourianoff G and Everschor-Sitte K 2020 Reservoir computing with random skyrmion textures Phys. Rev. Appl. 14 054020
- [184] Fernández-Pacheco A, Streubel R, Fruchart O, Hertel R, Fischer P and Cowburn R 2017 Three-dimensional nanomagnetism *Nat. Commun.* **8** 15756
- [185] Papp A, Porod W and Csaba G 2020 Nanoscale neural network using non-linear spin-wave interference (arXiv:2012.04594)

- [186] Khymyn R, Lisenkov I, Voorheis J, Sulymenko O, Prokopenko O, Tiberkevich V, Akerman J and Slavin A 2018 Ultra-fast artificial neuron: generation of picosecond-duration spikes in a current-driven antiferromagnetic auto-oscillator *Sci. Rep.* **8** 15727
- [187] Zázvorka J et al 2019 Thermal skyrmion diffusion used in a reshuffler device Nat. Nanotechnol. 14 658–61
- [188] Wang Z, Wu H, Burr G W, Hwang C S, Wang K L, Xia Q and Yang J J 2020 Resistive switching materials for information processing Nat. Rev. Mater. 5 173–95
- [189] Pi S, Li C, Jiang H, Xia W, Xin H, Yang J J and Xia Q 2019 Memristor crossbar arrays with 6 nm half-pitch and 2 nm critical dimension Nat. Nanotechnol. 14 35–9
- [190] Lin P et al 2020 Three-dimensional memristor circuits as complex neural networks Nat. Electron. 3 225-32
- [191] Prezioso M, Merrikh-Bayat F, Hoskins B D, Adam G C, Likharev K K and Strukov D B 2015 Training and operation of an integrated neuromorphic network based on metal-oxide memristors *Nature* 521 61–4
- [192] Yao P et al 2017 Face classification using electronic synapses Nat. Commun. 8 15199
- [193] Li C et al 2018 Efficient and self-adaptive in situ learning in multilayer memristor neural networks Nat. Commun. 9 2385
- [194] Li C et al 2019 Long short-term memory networks in memristor crossbar arrays Nat. Mach. Intell. 1 49-57
- [195] Liu Q et al 2020 33.2 A fully integrated analog ReRAM based 78.4 TOPS/W compute-in-memory chip with fully parallel MAC computing 2020 IEEE Int. Solid-State Circuits Conf. (ISSCC) 500–2
- [196] Wan W et al 2020 33.1 a 74 TMACS/W CMOS-RRAM neurosynaptic core with dynamically reconfigurable dataflow and in situ transposable weights for probabilistic graphical models 2020 IEEE Int. Solid-State Circuits Conf. (ISSCC) 498–500
- [197] Xue C-X et al 2021 A CMOS-integrated compute-in-memory macro based on resistive random-access memory for AI edge devices Nat. Electron. 4 81–90
- [198] Wang W, Song W, Yao P, Li Y, Van Nostrand J, Qiu Q, Ielmini D and Yang J J 2020 Integration and co-design of memristive devices and algorithms for artificial intelligence *Iscience* 23 101809
- [199] Li W, Xu P, Zhao Y, Li H, Xie Y and Lin Y 2020 TIMELY: pushing data movements and interfaces in PIM accelerators towards local and in time domain 2020 ACM/IEEE 47th Annual Int. Symp. Computer Architecture (ISCA) 832–45
- [200] Sangheon O, Shi Y, de Valle J, Salev P, Lu Y, Huang Z, Kalcheim Y, Schuller I V and Kuzum D 2021 Energy-efficient Mott activation neuron for full-hardware implementation of neural networks Nat. Nanotechnol. 16 680–7
- [201] Zhang X et al 2019 Experimental demonstration of conversion-based SNNs with 1T1R Mott neurons for neuromorphic inference 2019 IEEE Int. Electron Devices Meeting (IEDM) pp 6–7
- [202] Davies M et al 2018 Loihi: a neuromorphic manycore processor with on-chip learning IEEE Micro 38 82-99
- [203] Mayr C, Hoeppner S and Furber S 2019 Spinnaker 2: a 10 million core processor system for brain simulation and machine learning (arXiv:1911.02385)
- [204] Furber S and Bogdan P (ed) 2020 SpiNNaker: A Spiking Neural Network Architecture (Boston-Delft: Now Publishers)
- [205] Rasche C, Douglas R and Mahowald M 1997 Characterization of a pyramidal silicon neuron *Neuromorphic Systems: Engineering Silicon from Neurobiology* ed L S Smith and A Hamilton (Singapore: World Scientific)
- [206] van Schaik A 2001 Building blocks for electronic spiking neural networks Neural Netw. 14 617–28
- [207] Maldonado Huayaney F L, Nease S and Chicca E 2016 Learning in silicon beyond STDP: a neuromorphic implementation of multi-factor synaptic plasticity with calcium-based dynamics IEEE Trans. Circuits Syst. I 63 2189–99
- [208] Levi T, Nanami T, Tange A, Aihara K and Kohno T 2018 Development and applications of biomimetic neuronal networks toward brainmorphic artificial intelligence *IEEE Trans. Circuits Syst. II* 65 577–81
- [209] Abu-Hassan K, Taylor J D, Morris G, Donati E, Bortolotto Z A, Indiveri G, Paton J F R and Nogaret A 2019 Optimal solid state neurons Nat. Commun. 10 5309
- [210] Rubino A, Livanelioglu C, Qiao N, Payvand M and Indiveri G 2020 Ultra-low-power FDSOI neural circuits for extreme-edge neuromorphic intelligence IEEE Trans. Circuits Syst. I 68 45–56
- [211] Mead C 1989 Analog VLSI and Neural Systems (Reading, MA: Addison-Wesley)
- [212] Douglas R, Mahowald M and Mead C 1995 Neuromorphic analogue VLSI Annu. Rev. Neurosci. 18 255-81
- [213] Payvand M, Nair M V, Müller L K and Indiveri G 2019 A neuromorphic systems approach to in-memory computing with non-ideal memristive devices: from mitigation to exploitation *Faraday Discuss*. 213 487–510
- [214] Lillicrap T P and Santoro A 2019 Backpropagation through time and the brain Curr. Opin. Neurobiol. 55 82–9
- [215] Backus J 1978 Can programming be liberated from the von Neumann style? Commun. ACM 21 613-41
- [216] Indiveri G and Liu S-C 2015 Memory and information processing in neuromorphic systems Proc. IEEE 103 1379–97
- [217] Ganguli S, Huh D and Sompolinsky H 2008 Memory traces in dynamical systems Proc. Natl Acad. Sci. 105 18970-5
- [218] Park S, Chu M, Kim J, Noh J, Jeon M, Lee B H, Hwang H, Lee B and Lee B-g 2015 Electronic system with memristive synapses for pattern recognition Sci. Rep. 5 10123
- [219] Buccelli S et al 2019 A neuromorphic prosthesis to restore communication in neuronal networks iScience 19 402-14
- [220] Bauer F C, Muir D R and Indiveri G 2019 Real-time ultra-low power ECG anomaly detection using an event-driven neuromorphic processor IEEE Trans. Biomed. Circuits Syst. 13 1575–82
- [221] Donati E, Payvand M, Risi N, Krause R and Indiveri G 2019 Discrimination of EMG signals using a neuromorphic implementation of a spiking neural network *IEEE Trans. Biomed. Circuits Syst.* 13 795–803
- [222] Burelo K, Sharifshazileh M, Krayenbühl N, Ramantani G, Indiveri G and Sarnthein J 2021 A spiking neural network (SNN) for detecting high frequency oscillations (HFOs) in the intraoperative ECoG Sci. Rep. 11 6719
- [223] Papadimitriou C H and Steiglitz K 1998 Combinatorial Optimization: Algorithms and Complexity (Courier Corporation)
- [224] Kirkpatrick S, Gelatt C D and Vecchi M P 1983 Optimization by simulated annealing *Science* 220 671–80
- [225] Fogel D B 2006 Evolutionary Computation: Toward a New Philosophy of Machine Intelligence vol 1 (New York: Wiley)
- [226] Ackley D H, Hinton G E and Sejnowski T J 1985 A learning algorithm for Boltzmann machines Cogn. Sci. 9 147–69
- [227] Lucas A 2014 Ising formulations of many NP problems Front. Phys. 2 5
- [228] Barahona F 1982 On the computational complexity of Ising spin glass models J. Phys. A: Math. Gen. 15 3241
- [229] Hopfield J J 1982 Neural networks and physical systems with emergent collective computational abilities *Proc. Natl Acad. Sci.* 79 2554–8
- [230] Vadlamani S K, Xiao T P and Yablonovitch E 2020 Physics successfully implements Lagrange multiplier optimization Proc. Natl Acad. Sci. USA 117 26639–50
- [231] Kadowaki T and Nishimori H 1998 Quantum annealing in the transverse Ising model Phys. Rev. E 58 5355
- [232] Johnson M W et al 2011 Quantum annealing with manufactured spins Nature 473 194-8
- [233] Inagaki T et al 2016 A coherent Ising machine for 2000-node optimization problems Science 354 603-6
- [234] McMahon P L et al 2016 A fully programmable 100-spin coherent Ising machine with all-to-all connections Science 354 614-7

- [235] Hamerly R *et al* 2019 Experimental investigation of performance differences between coherent Ising machines and a quantum annealer *Sci. Adv.* 5 eaau0823
- [236] Aramon M, Rosenberg G, Valiante E, Miyazawa T, Tamura H and Katzgraber H G 2019 Physics-inspired optimization for quadratic unconstrained problems using a digital annealer *Front. Phys.* 7 48
- [237] 2018 Fujitsu quantum-inspired digital annealer cloud service to rapidly resolve combinatorial optimization problems-fujitsu global https://fujitsu.com/global/about/resources/news/press-releases/2018/0515-01.html (accessed 05 February 2021)
- [238] Takemoto T, Hayashi M, Yoshimura C and Yamaoka M 2020 A 2 × 30 k-spin multi-chip scalable CMOS annealing processor based on a processing-in-memory approach for solving large-scale combinatorial optimization problems *IEEE J. Solid-State Circuits* 55 145–56
- [239] Mahmoodi M R, Prezioso M and Strukov D B 2019 Versatile stochastic dot product circuits based on nonvolatile memories for high performance neurocomputing and neurooptimization Nat. Commun. 10 5113
- [240] Cai F et al 2020 Power-efficient combinatorial optimization using intrinsic noise in memristor Hopfield neural networks Nat. Electron. 3 409–18
- [241] Chou J, Bramhavar S and Ghosh S (W H-S Reports, and Undefined) 2019 Analog coupled oscillator based weighted Ising machine nature.com available: https://nature.com/articles/s41598-019-49699-5 (accessed: 06 February 2021)
- [242] Dutta S, Khanna A, Gomez J, Ni K, Toroczkai Z and Datta S 2019 Experimental demonstration of phase transition nano-oscillator based ising machine 2019 IEEE Int. Electron Devices Meeting (IEDM) pp 37–8
- [243] Xiao T 2019 Optoelectronics for Refrigeration and Analog Circuits for Combinatorial Optimization (UC Berkeley)
- [244] Camsari K Y, Sutton B M and Datta S 2019 p-bits for probabilistic spin logic Appl. Phys. Rev. 6 011305
- [245] Abbink E J W, Albino L, Dollevoet T, Huisman D, Roussado J and Saldanha R L 2011 Solving Large Scale Crew Scheduling Problems in Practice vol 3 (Berlin: Springer) pp 149–64
- [246] Hamerly R *et al* 2018 Scaling advantages of all-to-all connectivity in physical annealers: the coherent Ising machine vs D-wave 2000Q
- [247] Kalehbasti R, Ushijima-Mwesigwa H, Mandal A and Ghosh I 2020 Ising-based louvain method: clustering large graphs with specialized hardware (arXiv:2012.11391) (accessed: 05 February 2021)
- [248] Strubell E, Ganesh A and Mccallum A 2019 Energy and policy considerations for deep learning in NLP Proc. 57th Annual Meeting of the Association for Computational Linguistics (ACL) 3645–50
- [249] Rumelhart D E, Hinton G E and Williams R J 1986 Learning representations by back-propagating errors Nature 323 533-6
- [250] Hsieh E R et al 2019 High-density multiple bits-per-cell 1T4R RRAM array with gradual SET/RESET and its effectiveness for deep learning Proc. IEEE Int. Electron Devices Meeting (IEDM)
- [251] Esmanhotto E *et al* 2020 High-density 3D monolithically integrated multiple 1T1R multi-level-cell for neural networks *Proc. IEEE Int. Electron Devices Meeting (IEDM)*
- [252] Barraud S *et al* 2020 3D RRAMs with gate-all-around stacked nanosheet transistors for in-memory-computing *Proc. IEEE Int. Electron Devices Meeting (IEDM)*
- [253] Alfaro Robayo D *et al* 2019 Integration of OTS based back-end selector with HfO₂ OxRAM for crossbar arrays *Proc. IEEE Int. Electron Devices Meeting (IEDM)*
- [254] Le B Q, Grossi A, Vianello E, Wu T, Lama G, Beigne E, Wong H-S P and Mitra S 2019 Resistive RAM with multiple bits per cell: array-level demonstration of 3 bits per cell *IEEE Trans. Electron Devices* 66 641–6
- [255] Valentian A *et al* 2019 Fully integrated spiking neural network with analog neurons and RRAM synapses *Proc. IEEE Int. Electron* Devices Meeting (IEDM)
- [256] Ambrogio S et al 2018 Equivalent-accuracy accelerated neural-network training using analogue memory Nature 558 60-7
- [257] Payvand M and Indiveri G 2019 Spike-based plasticity circuits for alwayson on-line learning in neuromorphic systems *IEEE Int.* Symp. Circuits and Systems (ISCAS)
- [258] Ly D R B et al 2018 Role of synaptic variability in resistive memory-based spiking neural networks with unsupervised learning J. Phys. D: Appl. Phys. 51 444002
- [259] Gerstner W, Lehmann M, Liakoni V, Corneil D and Brea J 2018 Eligibility traces and plasticity on behavioral time scales: experimental support of neohebbian three-factor learning rules *Front. Neural Circuits* 12 53
- [260] Bellec G, Scherr F, Subramoney A, Hajek E, Salaj D, Legenstein R and Maass W 2020 A solution to the learning dilemma for recurrent networks of spiking neurons Nat. Commun. 11 3625
- [261] Crafton B, Parihar A, Gebhardt E and Raychowdhury A 2019 Direct feedback alignment with sparse connections for local learning Front. Neurosci. 13 525
- [262] Ernoult M et al 2019 Updates of equilibrium prop match gradients of backprop through time in an RNN with static input NeurIPS 2019 Proc. (arXiv:1905.13633)
- [263] Demirag Y *et al* 2021 PCM-trace: scalable synaptic eligibility traces with resistivity drift of phase-change materials *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS)* (to be published)
- [264] Dalgaty T, Castellani N, Turck C, Harabi K-E, Querlioz D and Vianello E 2021 In situ learning using intrinsic memristor variability via Markov chain Monte Carlo sampling Nat. Electron. 4 151–61
- [265] Vivet P et al 2020 A 220GOPS 96-core processor with 6 chiplets 3D-stacked on an active interposer offering 0.6 ns mm⁻¹ latency, 3-Tb/s/mm² inter-chiplet interconnects and 156 mW mm⁻² @ 82%-peakdfficiency DC–DC converters Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC) 46–8
- [266] Ambs P 2010 Optical computing: a 60 year adventure Adv. Opt. Technol. 2010 372652
- [267] Nahmias M A, De Lima T F, Tait A N, Peng H-T, Shastri B J and Prucnal P R 2020 Photonic multiply-accumulate operations for neural networks *IEEE J. Sel. Top. Quantum Electron.* 26 1–18
- [268] Lin X, Rivenson Y, Yardimci N T, Veli M, Luo Y, Jarrahi M and Ozcan A 2018 All-optical machine learning using diffractive deep neural networks Science 361 1004–8
- [269] Shastri B J et al 2020 Photonics for artificial intelligence and neuromorphic computing Nat. Photon. 15 102–14
- [270] Shen Y *et al* 2017 Deep learning with coherent nanophotonic circuits *Nat. Photon.* **11** 441–6
- [271] Tanaka G et al 2019 Recent advances in physical reservoir computing: a review Neural Netw. 115 100–23
- [272] Feldmann J, Youngblood N, Wright C D, Bhaskaran H and Pernice W H P 2019 All-optical spiking neurosynaptic networks with self-learning capabilities *Nature* **569** 208–14
- [273] Cheng Z, Ríos C, Pernice W H, Wright C D and Bhaskaran H 2017 On-chip photonic synapse Sci. Adv. 3 e1700160
- [274] Wu C, Yu H, Lee S, Peng R, Takeuchi I and Li M 2021 Programmable phase-change metasurfaces on waveguides for multimode photonic convolutional neural network Nat. Commun. 12 96

- [275] Xu X et al 2021 11 TOPS photonic convolutional accelerator for optical neural networks Nature 589 44-51
- [276] Sebastian A et al 2017 Temporal correlation detection using computational phase-change memory Nat. Commun. 8 1115
- [277] Ríos C et al 2019 In-memory computing on a photonic platform Sci. Adv. 5 eaau5759
- [278] Murmann B 2015 The race for the extra decibel: a brief review of current ADC performance trajectories IEEE Solid-State Circuits Mag. 7 58–66
- [279] Marin-Palomo P *et al* 2017 Microresonator-based solitons for massively parallel coherent optical communications *Nature* 546 274–9
- [280] Huang C, de Lima T F, Jha A, Abbaslou S, Shastri B J and Prucnal R 2019 Giant enhancement in signal contrast using integrated all-optical nonlinear thresholder *Optics InfoBase Conf. Papers*
- [281] Wuttig M, Bhaskaran H and Taubner T 2017 Phase-change materials for non-volatile photonic applications Nat. Photon. 11 465–76
- [282] Pernice W H and Bhaskaran H 2012 Photonic non-volatile memories using phase change materials Appl. Phys. Lett. 101 171101
- [283] Reshef O, De Leon I, Alam M Z and Boyd R W 2019 Nonlinear optical effects in epsilon-near-zero media Nat. Rev. Mater. 4 535–51
- [284] Gupta S, Agrawal A, Gopalakrishnan K and Narayanan P 2015 Deep learning with limited numerical precision 32nd Int. Conf. Machine Learning (ICML 2015)
- [285] Schemmel J, Brüderle D, Grübl A, Hock M, Meier K and Millner S 2021 A wafer-scale neuromorphic hardware system for large-scale neural modelling Proc. Int. Symp. Circuits System 1947–50
- [286] Davison A P, Brüderle D, Eppler J M, Kremkow J, Muller E, Pecevski D A, Perrinet L and Yger P 2009 PyNN: a common interface for neuronal network simulators *Front. Neuroinform.* 2 11
- [287] Knight J C and Nowotny T 2018 GPUs outperform current HPC and neuromorphic solutions in terms of speed and energy when simulating a highly-connected cortical model Front. Neurosci. 12 941
- [288] van Albada S J, Rowley A G, Senk J, Hopkins M, Schmidt M, Stokes A B, Lester D R, Diesmann M and Furber S B 2018 Performance comparison of the digital neuromorphic hardware SpiNNaker and the neural network simulation software NEST for a full-scale cortical microcircuit model *Front. Neurosci.* 12 291
- [289] Rhodes O, Peres L, Rowley A G D, Gait A, Plana L A, Brenninkmeijer C and Furber S B 2020 Realtime cortical simulation on neuromorphic hardware *Phil. Trans. R. Soc.* A 378 20190160
- [290] Blouw P, Choo X, Hunsberger E and Eliasmith C 2019 Benchmarking keyword spotting efficiency on neuromorphic hardware *Proc. NICE'19*
- [291] Arthur J and Boahen K 2006 Learning in silicon: timing is everything Advances in Neural Information Processing Systems 18 ed Y Weiss, B Schölkopf and J Platt (Cambridge, MA: MIT Press)
- [292] Baldi P, Sadowski P and Lu Z 2017 Learning in the machine: the symmetries of the deep learning channel Neural Netw. 95 110–33
- [293] Baydin A G, Pearlmutter B A, Radul A A and Siskind J M 2017 Automatic differentiation in machine learning: a survey J. Mach. Learn. Res. 18 5595–637
- [294] Bellec G, Scherr F, Hajek E, Salaj D, Legenstein R and Maass W 2019 Biologically inspired alternatives to backpropagation through time for learning in recurrent neural nets (arXiv:1901.09049)
- [295] Davies M, Srinivasa N, Lin T H, Chinya G, Joshi P, Lines A, Wild A and Wang H 2018 Loihi: a neuromorphic manycore processor with on-chip learning *IEEE Micro* 38 82–99
- [296] Yigit D, Moro F, Dalgaty T, Navarro G, Frenkel C, Indiveri G, Vianello E and Payvand M 2021 PCM-trace: scalable synaptic eligibility traces with resistivity drift of phase-change materials (arXiv:2102.07260)
- [297] Ercsey-Ravasz M, Markov N T, Lamy C, Van Essen D C, Knoblauch K, Toroczkai Z and Kennedy H 2013 A predictive network model of cerebral cortical connectivity based on a distance rule *Neuron* 80 184–97
- [298] Esser S K *et al* 2016 Convolutional networks for fast, energy-efficient neuromorphic computing *Proc. Natl Acad. Sci. USA* 113 11441–6
- [299] Fouda M E, Kurdahi F, Eltawil A and Neftci E 2019 Spiking Neural Networks for Inference and Learning: A Memristor-Based Design Perspective (Elsevier) ch 19 (accepted)
- [300] Simon F, Schemmel J, Grubl A, Hartel A, Hock M and Meier K 2017 Demonstrating hybrid learning in a flexible neuromorphic hardware system *IEEE Trans. Biomed. Circuits Syst.* 11 128–42
- [301] Galluppi F, Lagorce X, Stromatias E, Pfeiffer M, Plana L A, Furber S B and Benosman R B 2014 A framework for plasticity implementation on the spinnaker neural architecture *Front. Neurosci.* 8 429
- [302] Jaderberg M, Czarnecki W M, Osindero S, Vinyals O, Graves A and Kavukcuoglu K 2016 Decoupled neural interfaces using synthetic gradients (arXiv:1608.05343)
- [303] Jia Z, Tillman B, Maggioni M and Scarpazza D P 2019 Dissecting the graphcore IPU architecture via microbenchmarking (arXiv:1912.03413)
- [304] Kaiser J, Mostafa H and Neftci E 2019 Synaptic plasticity for deep continuous local learning 14 424
- [305] Kumaran D, Hassabis D and McClelland J L 2016 What learning systems do intelligent agents need? Complementary learning systems theory updated *Trends Cognit. Sci.* 20 512–34
- [306] Lillicrap T P, Santoro A, Marris L, Akerman C J and Hinton G 2020 Backpropagation and the brain Nat. Rev. Neurosci. 21 335-46
- [307] Neftci E O 2018 Data and power efficient intelligence with neuromorphic learning machines iScience 5 52-68
- [308] Payvand M, Fouda M E, Kurdahi F, Eltawil A and Neftci E O 2020 Error-triggered threefactor learning dynamics for crossbar arrays 2020 2nd IEEE Int. Conf. Artificial Intelligence Circuits and Systems (AICAS) 218–22
- [309] Pedroni B U, Joshi S, Deiss S R, Sheik S, Detorakis G, Paul S, Augustine C, Neftci E O and Cauwenberghs G 2019 Memory-efficient synaptic connectivity for spike-timing- dependent plasticity *Front. Neurosci.* 13 357
- [310] Pfeil T, Potjans T C, Schrader S, Potjans W, Schemmel J, Diesmann M and Meier K 2012 Is a 4 bit synaptic weight resolution enough? Constraints on enabling spike-timing dependent plasticity in neuromorphic hardware *Front. Neurosci.* 6 90
- [311] Prezioso M, Mahmoodi M R, Merrikh Bayat F, Nili H, Kim H, Vincent A and Strukov D B 2018 Spike-timing dependent plasticity learning of coincidence detection with passively integrated memristive circuits *Nat. Commun.* **9** 5311
- [312] Rastegari M, Ordonez V, Redmon J and Farhadi A 2016 Xnor-net: imagenet classification using binary convolutional neural networks *European Conf. Computer Vision* (Berlin: Springer) 525–42
- [313] Richards B A et al 2019 A deep learning framework for neuroscience Nat. Neurosci. 22 1761–70
- [314] Rueckauer B, Bybee C, Goettsche R, Singh Y, Mishra J and Wild A 2021 NxTF: an API and compiler for deep spiking neural networks on intel loihi (arXiv:2101.04261)

[315] Shrestha S B and Orchard G 2018 SLAYER: spike layer error reassignment in time Advances in Neural Information Processing Systems 1412-21 [316] Spilger P et al 2020 hxtorch: PyTorch for brainscales-2 IoT Streams for Data-Driven Predictive Maintenance and IoT, Edge, and Mobile for Embedded Machine Learning (Berlin: Springer) pp 189-200 Stewart K, Orchard G, Shrestha S B and Neftci E 2020 Online few-shot gesture learning on a neuromorphic processor IEEE J. [317] Emerg. Sel. Top. Circuits Syst. 10 512-21 [318] Thiele J C, Bichler O and Dupret A 2019 SpikeGrad: an ANN-equivalent computation model for implementing backpropagation with spikes (arXiv:1906.00851) [319] Wang Z et al 2018 Fully memristive neural networks for pattern classification with unsupervised learning Nat. Electron. 1 137 [320] Zenke F and Gerstner W 2014 Limits to high-speed simulations of spiking neural networks using generalpurpose computers Front. Neuroinf. 8 76 [321] Zenke F and Neftci E O 2021 Brain-inspired learning on neuromorphic substrates Proc. IEEE 109 935-50 [322] Zenke F and Ganguli S 2017 Superspike: supervised learning in multi-layer spiking neural networks (arXiv:1705.11146) [323] Lake B M, Ullman T D, Tenenbaum J B and Gershman S J 2017 Building machines that learn and think like people Behav. Brain Sci. 40 e253 [324] Murray J M 2019 Local online learning in recurrent networks with random feedback eLife 8 e43299 [325] Scherr F, Stöckl C and Maass W 2020 One-shot learning with spiking neural networks (bioRxiv) [326] Hochreiter S, Younger A S and Conwell R 2020 Learning to learn using gradient descent Int. Conf. Artificial Neural Networks (Berlin: Springer) pp 87-94 [327] Confavreux B, Zenke F, Agnes E J, Lillicrap T and Vogels T 2020 A meta-learning approach to (re) discover plasticity rules that carve a desired function into a neural network (bioRxiv) [328] Jordan J, Schmidt M, Senn W and Petrovici M A 2020 Evolving to learn: discovering interpretable plasticity rules for spiking networks (arXiv:2005.14149) [329] Bohnstingl T, Scherr F, Pehle C, Meier K and Maass W 2019 Neuromorphic hardware learns to learn Front. Neurosci. 13 483 [330] Wang J X, Kurth-Nelson Z, Tirumala D, Soyer H, Leibo J Z, Munos R and Botvinick M 2016 Learning to reinforcement learn (arXiv:1611.05763) [331] Duan Y, Schulman J, Chen X, Bartlett L, Sutskever I and Abbeel P 2016 Rl²: fast reinforcement learning via slow reinforcement learning (arXiv:1611.02779) [332] Wang J X, Kurth-Nelson Z, Kumaran D, Tirumala D, Soyer H, Leibo J Z, Hassabis D and Botvinick M 2018 Prefrontal cortex as a meta-reinforcement learning system Nat. Neurosci. 21 860-8 [333] Bellec G, Salaj D, Subramoney A, Legenstein R and Maass W 2018 Long short-term memory and learning-to-learn in networks of spiking neurons Advances in Neural Information Processing Systems 787-97 [334] Subramoney A, Bellec G, Scherr F, Legenstein R and Maass W 2020 Revisiting the role of synaptic plasticity and network dynamics for fast learning in spiking neural networks (bioRxiv) [335] Subramoney A, Scherr F and Maass W 2019 Reservoirs learn to learn (arXiv:1909.07486) [336] Finn C, Abbeel P and Levine S 2017 Model-agnostic meta-learning for fast adaptation of deep networks Proc. 34th Int. Conf. Machine Learning (PMLR) vol 70 pp 1126-35 [337] Furber S B, Lester D R, Plana L A, Garside J D, Painkras E, Temple S and Brown A D 2012 Overview of the spinnaker system architecture IEEE Trans. Comput. 62 2454-67 [338] Salimans T, Ho J, Chen X, Sidor S and Sutskever I 2017 Evolution strategies as a scalable alternative to reinforcement learning (arXiv:1703.03864) [339] Grübl A, Billaudelle S, Cramer B, Karasenko V and Schemmel J 2020 Verification and design methods for the brainscaleS neuromorphic hardware system J. Signal Process. Syst. 92 1277-92 [340] Sejnowski T J, Koch C and Churchland P S 1988 Computational neuroscience Science 241 1299-306 [341] Abbott L F 2008 Theoretical neuroscience rising Neuron 60 489-95 [342] Hodgkin A L and Huxley A F 1952 A quantitative description of membrane current and its application to conduction and excitation in nerve J. Physiol. 117 500-44 [343] Hopfield J J 1988 Artificial neural networks IEEE Circuits Dev. Mag. 4 3-10 [344] Schwartz E L 1993 Computational Neuroscience (Cambridge, MA: MIT Press) Thiele A and Bellgrove M A 2018 Neuromodulation of attention Neuron 97 769-85 [345] [346] Baddeley A 1992 Working memory Science 255 556-9 [347] Hollerman J R and Schultz W 1998 Dopamine neurons report an error in the temporal prediction of reward during learning Nat. Neurosci. 1 304-9 [348] Markram H et al 2015 Reconstruction and simulation of neocortical microcircuitry Cell 163 456-92 [349] Koch C and Reid R C 2012 Observatories of the mind Nature 483 397-8 [350] Koroshetz W et al 2018 The state of the NIH BRAIN initiative J. Neurosci. 38 6427-38 Amunts K, Ebell C, Muller J, Telefont M, Knoll A and Lippert T 2016 The human brain project: creating a European Research [351] infrastructure to decode the human brain Neuron 92 574-81 [352] Adams A et al 2020 International brain initiative: an innovative framework for coordinated global brain Research efforts Neuron **105** 212-6 [353] Okano H, Miyawaki A and Kasai K 2015 Brain/MINDS: brain-mapping project in Japan Phil. Trans. R. Soc. B 370 20140310 [354] Davison A P 2012 Collaborative modelling: the future of computational neuroscience? Netw. Comput. Neural Syst. 23 157-66 [355] Ramaswamy S et al 2015 The neocortical microcircuit collaboration portal: a resource for rat somatosensory cortex Front. Neural Circuits 9 44 [356] Markram H 2013 Seven challenges for neuroscience Funct. Neurol. 28 145-51 [357] Suárez L E, Richards B A, Lajoie G and Misic B 2021 Learning function from structure in neuromorphic networks Nat. Mach. Intell. 3 771-86 [358] Tsodyks M 2008 Computational neuroscience grand challenges—a humble attempt at future forecast Front. Neurosci. 2 21 [359] Hasler J and Marr B 2013 Finding a roadmap to achieve large neuromorphic hardware systems Front. Neurosci. 7 118 [360] Niven J E and Laughlin S B 2008 Energy limitation as a selective pressure on the evolution of sensory systems J. Exp. Biol. 211 1792-804 [361] Branco T and Staras K 2009 The probability of neurotransmitter release: variability and feedback control at single synapses Nat. Rev. Neurosci. 10 373-83

- [362] Hamilton T J, Afshar S, van Schaik A and Tapson J C 2014 Stochastic electronics: a neuro-inspired design paradigm for integrated circuits Proc. IEEE 102 843–59
- [363] Alawad M and Lin M 2019 Survey of stochastic-based computation paradigms IEEE Trans. Emerg. Top. Comput. 798–114
- [364] Pantone R D, Kendall J D and Nino J C 2018 Memristive nanowires exhibit small-world connectivity *Neural Netw*. 106 144–51 [365] Pigotti M, Parek O, Wordon M P, Wang X L, Day N D, Miller F K and Euci S 2013 The importance of mixed coloritiity in
- [365] Rigotti M, Barak O, Warden M R, Wang X-J, Daw N D, Miller E K and Fusi S 2013 The importance of mixed selectivity in complex cognitive tasks *Nature* 497 585–90
- [366] Thakur C S, Afshar S, Wang R M, Hamilton T J, Tapson J and van Schaik A 2016 Bayesian estimation and inference using stochastic electronics Front. Neurosci. 10 104
- [367] Maass W 2014 Noise as a resource for computation and learning in networks of spiking neurons *Proc. IEEE* **102** 860–80
- [368] Hubel D H and Wiesel T N 1962 Receptive fields, binocular interaction and functional architecture in the cat's visual cortex J. Physiol. 160 106–54
- [369] Fukushima K and Miyake S 1982 Neocognitron: a self-organizing neural network model for a mechanism of visual pattern recognition Competition and Cooperation in Neural Nets (Berlin: Springer) pp 267–85
- [370] Kim Y and Panda P 2021 Visual explanations from spiking neural networks using interspike intervals (arXiv:2103.14441)
- [371] Jain S, Sengupta A, Roy K and Raghunathan A 2020 RxNN: a framework for evaluating deep neural networks on resistive crossbars IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst. 40 326–38
- [372] Chen P-Y, Peng X and Yu S 2018 NeuroSim: a circuit-level macro model for benchmarking neuro-inspired architectures in online learning IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst. 37 3067–80
- [373] Diehl P U, Neil D, Binas J, Cook M, Liu S C and Pfeiffer M 2015 Fast-classifying, highaccuracy spiking deep networks through weight and threshold balancing 2015 Int. Joint Conf. Neural Networks (IJCNN) (Piscataway, NJ: IEEE) pp 1–8
- [374] Lee J H, Delbruck T and Pfeiffer M 2016 Training deep spiking neural networks using backpropagation Front. Neurosci. 10 508
- [375] Gallego G et al 2020 Event-based vision: a survey IEEE Trans. Pattern Anal. Mach. Intell. 44 154-80
- [376] Goodfellow I J, Shlens J and Szegedy C 2014 Explaining and harnessing adversarial examples (arXiv:1412.6572)
- [377] Mostafa H 2017 Supervised learning based on temporal coding in spiking neural networks *IEEE Trans. Neural Netw. Learn. Syst.* 29 3227–35
- [378] Han B and Roy K 2020 Deep spiking neural network: energy efficiency through time based coding European Conf. Computer Vision p 388
- [379] Montemurro M A, Rasch M J, Murayama Y, Logothetis N K and Panzeri S 2008 Phase-of-firing coding of natural visual stimuli in primary visual cortex Curr. Biol. 18 375–80
- [380] Kim J, Kim H, Huh S, Lee J and Choi K 2018 Deep neural networks with weighted spikes Neurocomputing 311 373-86
- [381] Rathi N, Srinivasan G, Panda P and Roy K 2020 Enabling deep spiking neural networks with hybrid conversion and spike timing dependent backpropagation (arXiv:2005.01807)
- [382] Wu Y, Deng L, Li G, Zhu J and Shi L 2018 Spatio-temporal backpropagation for training high-performance spiking neural networks *Front. Neurosci.* **12** 331
- [383] Kim Y and Panda P 2020 Revisiting batch normalization for training low-latency deep spiking neural networks from scratch (arXiv:2010.01729)
- [384] Sharmin S, Panda P, Sarwar S S, Lee C, Ponghiran W and Roy K 2019 A comprehensive analysis on adversarial robustness of spiking neural networks 2019 Int. Joint Conf. Neural Networks (IJCNN) (Piscataway, NJ: IEEE) pp 1–8
- [385] Sharmin S, Rathi N, Panda P and Roy K 2020 Inherent adversarial robustness of deep spiking neural networks: effects of discrete input encoding and non-linear activations *European Conf. Computer Vision* (Berlin: Springer) pp 399–414
- [386] Roy K, Jaiswal A and Panda P 2019 Towards spike-based machine intelligence with neuromorphic computing Nature 575 607-17
- [387] Lukoševičius M and Jaeger H 2009 Reservoir computing approaches to recurrent neural network training *Comput. Sci. Rev.* 3 127–49
- [388] Jaeger H 2001 The 'echo state' approach to analysing and training recurrent neural networkswith an erratum note *GMD Technical Report 148.34* (Bonn: German National Research Center for Information Technology) p 13
- [389] Maass W, Natschläger T and Markram H 2002 Real-time computing without stable states: a new framework for neural computation based on perturbations *Neural Comput.* 14 2531–60
- [390] Jaeger H 2007 Echo state network Scholarpedia 2 2330
- [391] Tanaka G, Yamane T, Héroux J B, Nakane R, Kanazawa N, Takeda S, Numata H, Nakano D and Hirose A 2019 Recent advances in physical reservoir computing: a review *Neural Netw.* 115 100–23
- [392] Appeltant L et al 2011 Information processing using a single dynamical node as complex system Nat. Commun. 2 468
- [393] Fernando C and Sojakka S 2003 Pattern recognition in a bucket *European Conf. Artificial Life* (Berlin: Springer)
- [394] Shi W, Cao J, Zhang Q, Li Y and Xu L 2016 Edge computing: vision and challenges *IEEE Internet Things J*. **3** 637–46
- [395] Jaeger H 2002 Tutorial on Training Recurrent Neural Networks, Covering BPPT, RTRL, EKF and the 'Echo State Network' Approach vol 5 (Bonn: GMD-Forschungszentrum Informationstechnik)
- [396] Bertschinger N and Natschläger T 2004 Real-time computation at the edge of chaos in recurrent neural networks Neural Comput. 16 1413–36
- [397] Lukoševičius M 2012 A practical guide to applying echo state networks Neural Networks: Tricks of the Trade (Berlin: Springer) pp 659–86
- [398] Grigoryeva L and Ortega J-P 2018 Echo state networks are universal Neural Netw. 108 495–508
- [399] Nakajima K and Fischer I (ed) 2021 Reservoir Computing: Theory, Physical Implementations, and Applications (Berlin: Springer)
- [400] Ma C et al 2021 Addressing limited weight resolution in a fully optical neuromorphic reservoir computing readout Sci. Rep. 11 3102
- [401] Araujo F A et al 2020 Role of non-linear data processing on speech recognition task in the framework of reservoir computing Sci. Rep. 10 328
- [402] Nakane R *et al* Spin waves propagating through a stripe magnetic domain structure and their applications to reservoir computing *Phys. Rev. Res.* (accepted)
- [403] Pathak J, Lu Z, Hunt B R, Girvan M and Ott E 2017 Using machine learning to replicate chaotic attractors and calculate Lyapunov exponents from data Chaos 27 121102
- [404] Gallicchio C, Micheli A and Pedrelli L 2017 Deep reservoir computing: a critical experimental analysis *Neurocomputing* 268 87–99
- [405] Hadaeghi F 2021 Neuromorphic Electronic Systems for Reservoir Computing (Reservoir Computing Natural Computing Series) ed K Nakajima and I Fischer (Berlin: Springer)

- [406] Van der Sande G, Brunner D and Soriano M C 2017 Advances in photonic reservoir computing Nanophotonics 6 561-76
- [407] Cramer B et al 2020 Control of criticality and computation in spiking neuromorphic networks with plasticity Nat. Commun. 11 2853
- [408] Tanaka G et al 2020 Guest editorial: special issue on new frontiers in extremely efficient reservoir computing IEEE Trans. Neural Netw. Learn. Syst. (unpublished)
- [409] Gerstner W and Kistler W M 2002 Spiking Neuron Models: Single Neurons, Populations, Plasticity (Cambridge: Cambridge University Press)
- [410] Paugam-Moisy H and Bohte S M 2012 Computing with spiking neuron networks Handbook of Natural Computing vol 1 pp 1-47
- [411] Mahowald M 1992 VLSI Analogs of Neuronal Visual Processing: A Synthesis of Form and Function
- [412] Furber S B, Galluppi F, Temple S and Plana L A 2014 The spinnaker project Proc. IEEE 102 652-65
- [413] Delorme A, Gautrais J, Van Rullen R and Thorpe S 1999 SpikeNET: a simulator for modeling large networks of integrate and fire neurons Neurocomputing 26–27 989–96
- [414] Delorme A and Thorpe S J 2003 SpikeNET: an event-driven simulation package for modelling large networks of spiking neurons Netw. Comput. Neural Syst. 14 613–27
- [415] Adrian E D and Matthews R 1927 The action of light on the eye J. Physiol. 63 378-414
- [416] Thorpe S J 1990 Spike arrival times: a highly efficient coding scheme for neural networks Parallel Processing in Neural Systems and Computers ed R Eckmiller, G Hartmann and G Hauske (Amsterdam: North-Holland) pp 91–4
- [417] Gollisch T and Meister M 2008 Rapid neural coding in the retina with relative spike latencies Science 319 1108–11
- [418] Thorpe S and Gautrais J 1998 Rank order coding Computational Neuroscience: Trends in Research ed J Bower (New York: Plenum) pp 113–8
- [419] Furber S B, John Bainbridge W, Mike Cumpstey J and Temple S 2004 Sparse distributed memory using N-of-M codes Neural Netw. 17 1437–51
- [420] Thorpe S J, Guyonneau R, Guilbaud N, Allegraud J-M and Vanrullen R 2004 SpikeNet: real-time visual processing with one spike per neuron Neurocomputing 58–60 857–64
- [421] Agus T R, Thorpe S J and Pressnitzer D 2010 Rapid formation of robust auditory memories: insights from noise Neuron 66 610-8
- [422] Thunell E and Thorpe S J 2019 Memory for repeated images in rapid-serial-visual-presentation streams of thousands of images *Psychol. Sci.* **30** 989–1000
- [423] Masquelier T and Thorpe S J 2007 Unsupervised learning of visual features through spike timing dependent plasticity PLoS Comput. Biol. 3 e31
- [424] Masquelier T, Guyonneau R and Thorpe S J 2008 Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains *PLoS One* 3 e1377
- [425] Bichler O, Querlioz D, Thorpe S J, Bourgoin J-P and Gamrat C 2012 Extraction of temporally correlated features from dynamic vision sensors with spike-timing-dependent plasticity *Neural Netw.* 32 339–48
- [426] Thorpe S, Yousefzadeh A, Martin J and Masquelier T 2017 Unsupervised learning of repeating patterns using a novel STDP based algorithm J.Vis. 17 1079
- [427] Indiveri G 1999 Neuromorphic analog VLSI sensor for visual tracking: circuits and application examples IEEE Trans. Circuits Syst. II 46 1337–47
- [428] http://jaerproject.org
- [429] Lichtsteiner P, Posch C and Delbruck T 2008 A 128 × 128 120 dB 15 μs latency asynchronous temporal contrast vision sensor IEEE J. Solid-State Circuits 43 566–76
- [430] Qiao N et al 2015 A re-configurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128 K synapses Front. Neurosci. 9 141
- [431] Monforte M, Arriandiaga A, Glover A and Bartolozzi C 2020 Where and when: event-based spatiotemporal trajectory prediction from the iCub's point-of-view 2020 IEEE Int. Conf. Robotics and Automation (ICRA) (Piscataway, NJ: IEEE) pp 9521–7
- [432] Kreiser R, Renner A, Leite V R, Serhan B, Bartolozzi C, Glover A and Sandamirskaya Y 2020 An on-chip spiking neural network for estimation of the head pose of the iCub robot *Front. Neurosci.* **14** 551
- [433] Gutierrez-Galan D, Dominguez-Morales J P, Perez-Peña F, Jimenez-Fernandez A and Linares-Barranco A 2020 NeuroPod: a real-time neuromorphic spiking CPG applied to robotics *Neurocomputing* 381 10–9
- [434] Zhao J, Risi N, Monforte M, Bartolozzi C, Indiveri G and Donati E 2020 Closed-loop spiking control on a neuromorphic processor implemented on the iCub IEEE J. Emerg. Sel. Top. Circuits Syst. 10 546–56
- [435] Naveros F, Luque N R, Ros E and Arleo A 2020 VOR adaptation on a humanoid iCub robot using a spiking cerebellar model IEEE Trans. Cybern. 50 47444757
- [436] García D H, Adams S, Rast A, Wennekers T, Furber S and Cangelosi A 2018 Visual attention and object naming in humanoid robots using a bio-inspired spiking neural network *Robot. Auton. Syst.* 104 56–71
- [437] Rapp H and Nawrot M P 2020 Proc. Natl Acad. Sci. USA 117 28412-21
- [438] Kaiser J, Mostafa H and Neftci E 2020 Synaptic plasticity dynamics for deep continuous local learning (DECOLLE) Front. Neurosci. 14 424
- [439] Illing B, Gerstner W and Brea J 2019 Biologically plausible deep learning—but how far can we go with shallow networks? *Neural Netw.* 118 90–101
- [440] Klos C, Kossio Y F K, Goedeke S, Gilra A and Memmesheimer R M 2020 Dynamical learning of dynamics Phys. Rev. Lett. 125 088103
- [441] Panzeri S, Harvey C D, Piasini E, Latham E and Fellin T 2017 Cracking the neural code for sensory perception by combining statistics, intervention, and behavior Neuron 93 491507
- [442] Yang S, Wang J, Zhang N, Deng B, Pang Y and Azghadi M R 2021 CerebelluMorphic: large-scale neuromorphic model and architecture for supervised motor learning *IEEE Trans. Neural Netw. Learn. Syst.* 1–15
- [443] 2016 Video interview with anthony foxx available online at https://theverge.com/a/verge2021/secretary-anthony-foxx
- [444] Adams E 2020 The revolution will not be self-driven: the real future of autonomous cars available online at https://robbreport.com/motors/cars/self-driving-cars-explainer-2901586/
- [445] Jason Eichenholz, Luminar Technologies Inc 2019 personal communication
- [446] Gawron J H, Keoleian G A, De Kleine R D, Wallington T J and Kim H C 2018 Life cycle assessment of connected and automated vehicles: sensing and computing subsystem and vehicle level effects *Environ. Sci. Technol.* 52 3249–56
- [447] Lichtsteiner P and Delbruck T 2005 64 × 64 event-driven logarithmic temporal derivative silicon retina 2005 IEEE Workshop on Charge-Coupled Devices and Advanced Image Sensors (Nagano, Japan) pp 157–60

- [448] Itti L, Koch C and Niebur E 1998 A model of saliency-based visual attention for rapid scene analysis IEEE Trans. Pattern Anal. Machine Intell. 20 1254–9
- [449] Moreira O et al 2020 NeuronFlow: a hybrid neuromorphic—dataflow processor architecture for AI workloads 2nd IEEE Int. Conf. Artificial Intelligence Circuits (AICAS 2020)
- [450] Brown T et al 2020 Language models are few-shot learners (arXiv:2005.14165v4)
- [451] Voelker A, Rasmussen D and Eliasmith C 2020 A spike in performance (arXiv:2002.03553)
- [452] Persaud K and Dodd G 1982 Analysis of discrimination mechanisms in the mammalian olfactory system using a model nose Nature 299 352–5
- [453] Gardner J W, Hines E L and Wilkinson M 1990 Application of artificial neural networks to an electronic olfactory system Meas. Sci. Technol. 1 446–51
- [454] Hines E L and Gardner J W 1994 An artificial neural emulator for an odour sensor array Sensors Actuators B 19 661-4
- [455] Persaud K C, Marco S and Gutierrez-Galvez A 2013 *Neuromorphic Olfaction (Frontiers in Neuroengineering)* (Boca Raton, FL: CRC Press)
- [456] Pearce T C, Schiffman S S, Nagle H T and Gardner J W 2003 Handbook of Machine Olfaction: Electronic Nose Technology (New York: Wiley)
- [457] Gronowitz M E, Liu A, Qiu Q, Yu C R and Cleland T A A physicochemical model of odor sampling (bioRxiv)
- [458] Pearce T C, Karout S, Rácz Z, Capurro A, Gardner J W and Cole M 2013 Rapid processing of chemosensor transients in a neuromorphic implementation of the insect macroglomerular complex *Front. Neurosci.* 7 119
- [459] Vergara A, Fonollosa J, Mahiques J, Trincavelli M, Rulkov N and Huerta R 2013 On the performance of gas sensor arrays in open sampling systems using inhibitory support vector machines Sensors Actuators B 185 462–77
- [460] Cleland T A and Borthakur A 2020 A systematic framework for olfactory bulb signal transformations Front. Comput. Neurosci. 14 579143
- [461] Schmuker M and Schneider G 2007 Processing and classification of chemical data inspired by insect olfaction Proc. Natl. Acad. Sci. USA 104 20285–9
- [462] Huerta R, Vembu S, Amigó J M, Nowotny T and Elkan C 2012 Inhibition in multiclass classification Neural Comput. 24 2473-507
- [463] Koickal T J, Hamilton A, Tan S L, Covington J A, Gardner J W and Pearce T C 2007 Analog VLSI circuit implementation of an adaptive neuromorphic olfaction chip IEEE Trans. Circuits Syst. 1 54 60–73
- [464] Vanarse A, Osseiran A and Rassau A 2017 An investigation into spike-based neuromorphic approaches for artificial olfactory systems Sensors 17 2591
- [465] Diamond A, Schmuker M and Nowotny T 2019 An unsupervised neuromorphic clustering algorithm Biol. Cybern. 113 423-37
- [466] Imam N and Cleland T A 2020 Rapid online learning and robust recall in a neuromorphic olfactory circuit *Nat. Mach. Intell.* 2 181–91
- [467] Borthakur A and Cleland T A 2019 Signal conditioning for learning in the wild Presented at the Neuro-inspired Computational Elements Workshop (Albany, NY, USA) available: https://doi.org/10.1145/3320288.3320293
- [468] Borthakur A and Cleland T A 2017 A neuromorphic transfer learning algorithm for orthogonalizing highly overlapping sensor array responses *Presented at the ISOCS/IEEE Int. Symp. Olfaction and Electronic Nose (ISOEN)* (Montreal, Canada)
- [469] Guerrero-Rivera R and Pearce T C 2007 Attractor-based pattern classification in a spiking FPGA implementation of the olfactory bulb 2007 3rd Int. IEEE/EMBS Conf. Neural Engineering (Piscataway, NJ: IEEE) pp 593–9
- [470] Imam N, Cleland T A, Manohar R, Merolla P A, Arthur J V, Akopyan F and Modha D S 2012 Implementation of olfactory bulb glomerular-layer computations in a digital neurosynaptic core *Front. Neurosci.* 6 83
- [471] Marco S et al 2014 A biomimetic approach to machine olfaction, featuring a very large-scale chemical sensor array and embedded neuro-bio-inspired computation *Microsyst. Technol.* 20 729–42
- [472] BelBruno J J 2019 Molecularly imprinted polymers Chem. Rev. 119 94-119
- [473] Mahowald M and Mead C 1991 The silicon retina Sci. Am. 264 76-82
- [474] Posch C, Serrano-Gotarredona T, Linares-Barranco B and Delbruck T 2014 Retinomorphic event-based vision sensors: bioinspired cameras with spiking output Proc. IEEE 102 1470–84
- [475] Posch C, Matolin D and Wohlgenannt R 2011 A QVGA 143 dB dynamic range frame-free PWM image sensor with lossless pixel-level video compression and time-domain CDS IEEE J. Solid-State Circuits 46 259–75
- [476] Brandli C, Berner R, Yang M, Liu S-C and Delbruck T 2014 A 240 × 180 130 dB 3 μs latency global shutter spatiotemporal vision sensor *IEEE J. Solid-State Circuits* 49 2333–41
- [477] Son B et al 2017 A 640 × 480 dynamic vision sensor with a 9 μm pixel and 300 Meps addressevent representation 2017 IEEE Int. Solid-State Circuits Conf. (ISSCC) (San Francisco, CA) pp 66–7
- [478] Finateu T *et al* 2020 A 1280 × 720 back-illuminated stacked temporal contrast event-based vision sensor with 4.86 μm pixels, 1.066 GEPS readout, programmable event-rate controller and compressive data-formatting pipeline 2020 IEEE Int. Solid-State Circuits Conf. (ISSCC) (San Francisco, CA, USA) pp 112–4
- [479] Moradi S, Qiao N, Stefanini F and Indiveri G 2018 A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs) *IEEE Trans. Biomed. Circuits Syst.* 12 106122
- [480] Moreira O et al 2020 NeuronFlow: a neuromorphic processor architecture for live AI applications
- [481] Zhou F et al 2019 Optoelectronic resistive random access memory for neuromorphic vision sensors Nat. Nanotechnol. 14 776
- [482] Tian H, Wang X, Wu F, Yang Y and Ren T-L 2018 High performance 2D perovskite/graphene optical synapses as artificial eyes 2018 IEEE Int. Electron Devices Meeting (IEDM)
- [483] 2021 Neuromorphic Sensing & Computing Report (Yole Développement) www.yole.fr
- [484] Shamma S A 1985 Speech processing in the auditory system: I. The representation of speech sounds in the responses of the auditory nerve J. Acoust. Soc. Am. 78 1612–21
- [485] Palmer A R and Russell I J 1986 Phase-locking in the cochlear nerve of the Guinea-pig and its relation to the receptor potential of inner hair-cells Hear. Res. 24 1–15
- [486] Lyon R F and Mead C 1988 An analog electronic cochlea IEEE Trans. Acoust. Speech Signal Process. 36 1119–34
- [487] Lyon R F, Katsiamis A G and Drakakis E M 2010 History and future of auditory filter models *Proc. 2010 IEEE Int. Symp. Circuits* and Systems pp 3809–12
- [488] Liu S-C, Delbruck T, Indiveri G, Whatley A and Douglas R 2015 Event-Based Neuromorphic Systems (New York: Wiley)
- [489] Yang M, Chien C-H, Delbruck T and Liu S-C 2016 A 0.5 V 55 μ W 64 \times 2 channel binaural silicon cochlea for event-driven stereo-audio sensing *IEEE J. Solid-State Circuits* 51 2554–69

- [490] Sarpeshkar R, Salthouse C, Sit J-J, Baker M W, Zhak S M, Lu T K-T, Turicchia L and Balster S 2005 An ultra-low-power programmable analog bionic ear processor *IEEE Trans. Biomed. Eng.* 52 711–27
- [491] Horiuchi T K 2009 A spike-latency model for sonar-based navigation in obstacle fields IEEE Trans. Circuits Syst. I 56 2393-401
- [492] Chan V, Jin C and van Schaik A 2010 Adaptive sound localisation with a silicon cochlea pair Front. Neurosci. 4 196
- [493] Liu S-C, Rueckauer B, Ceolini E, Huber A and Delbruck T 2019 Event-driven sensing for efficient perception: vision and audition algorithms IEEE Signal Process. Mag. 36 29–37
- [494] Uysal I, Sathyendra H and Harris J 2008 Towards spike-based speech processing: a biologically plausible approach to simple acoustic classification Int. J. Appl. Math. Comput. Sci. 18 129–37
- [495] Gao C, Braun S, Kiselev I, Anumula J, Delbruck T and Liu S 2019 Real-time speech recognition for IoT purpose using a delta recurrent neural network accelerator 2019 IEEE Int. Symp. Circuits and Systems (ISCAS) pp 1–5
- [496] Tsai W Y et al 2016 Always-on speech recognition using truenorth, a reconfigurable, neurosynaptic processor IEEE Trans. Comput. 66 996–1007
- [497] Yang M, Yeh C-H, Zhou Y, Cerqueira J P, Lazar A A and Seok M 2019 Design of an always-on deep neural network-based 1 µW voice activity detector aided with a customized software model for analog feature extraction IEEE J. Solid-State Circuits 54 1764–77
- [498] Neftci E, Mostafa H and Zenke F 2019 Surrogate gradient learning in spiking neural networks IEEE Signal Process. Mag. 36 61-3
- [499] Rueckauer B, Lungu I, Hu Y, Pfeiffer M and Liu S-C 2017 Conversion of continuous-valued deep networks to efficient
- event-driven networks for image classification *Front. Neurosci* 11 682
 [500] Wu J, Yılmaz E, Zhang M, Li H and Tan K C 2020 Deep spiking neural networks for large vocabulary automatic speech recognition *Front. Neurosci.* 14 199
- [501] Zai A, Bhargava S, Mesgarani N and Liu S-C 2015 Reconstruction of audio waveforms from spike trains of artificial cochlea models Front. Neurosci. 9 347
- [502] Ceolini E, Anumula J, Braun S and Liu S 2019 Event-driven pipeline for low-latency low-compute keyword spotting and speaker verification system 2019 IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP) 7953–7
- [503] Shan W, Yang M-H, Xu J, Lu Y, Zhang S, Wang T, Yang J, Shi L and Seok M 2020 14.1 A 510 nW 0.41 V low-memory lowcomputation keyword-spotting chip using serial FFT-based MFCC and binarized depthwise separable convolutional neural network in 28 nm CMOS 2020 IEEE Int. Solid-State Circuits Conf. (ISSCC) 230–2
- [504] Renaud-LeMasson S, LeMasson G, Marder E and Abbot L F 1993 Hybrid circuits of interacting computer model and biological neurons *Presented at the Advances in Neural Information Processing Systems 5 (NIPS Conf.)*
- [505] Chapin J K, Moxon K A, Markowitz R S and Nicolelis M A L 1999 Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex Nat. Neurosci. 2 664–70
- [506] Hochberg L R et al 2012 Reach and grasp by people with tetraplegia using a neurally controlled robotic arm Nature 485 372-5
- [507] Panuccio G, Semprini M and Chiappalone M 2016 Intelligent biohybrid systems for functional brain repair New Horiz. Transl. Med. 3 162–74
- [508] Guggenmos D J, Azin M, Barbay S, Mahnken J D, Dunham C, Mohseni P and Nudo R J 2013 Restoration of function after brain damage using a neural prosthesis Proc. Natl Acad. Sci. 110 21177
- [509] Vassanelli S and Mahmud M 2016 Trends and challenges in neuroengineering: toward 'intelligent' neuroprostheses through brain-'brain inspired systems' communication Front. Neurosci. 10 438
- [510] Serb A et al 2020 Memristive synapses connect brain and silicon spiking neurons Sci. Rep. 10 2590
- [511] Pautot S, Wyart C and Isacoff E Y 2008 Colloid-guided assembly of oriented 3D neuronal networks Nat. Methods 5 735-40
- [512] Gómez-Martínez R, Hernández-Pinto A M, Duch M, Vázquez P, Zinoviev K, de la Rosa E J, Esteve J, Suárez T and Plaza J A 2013 Silicon chips detect intracellular pressure changes in living cells Nat. Nanotechnol. 8 517–21
- [513] Chen L Y, Parizi K B, Kosuge H, Milaninia K M, McConnell M V, Wong H-S P and Poon A S Y 2013 Mass fabrication and delivery of 3D multilayer μtags into living cells Sci. Rep. 3 2295
- [514] Desai S B *et al* 2016 MoS₂ transistors with one-nanometer gate lengths *Science* **354** 6308
- [515] van de Burgt Y, Melianas A, Keene S T, Malliaras G and Salleo A 2018 Organic electronics for neuromorphic computing Nat. Electron. 1 386–97
- [516] Tran L-G, Cha H-K and Park W-T 2017 RF power harvesting: a review on designing methodologies and applications Micro Nano Syst. Lett. 5 14
- [517] Basaeri H, Christensen D B and Roundy S 2016 A review of acoustic power transfer for bio-medical implants *Smart Mater. Struct.* 25 123001
- [518] Rebel G, Estevez F, Gloesekoetter P and Castillo-Secilla J M 2015 Energy harvesting on human bodies Smart Health: Open Problems and Future Challenges ed A Holzinger, C Röcker and M Ziefle (Berlin: Springer) pp 125–59
- [519] Roseman J M, Lin J, Ramakrishnan S, Rosenstein J K and Shepard K L 2015 Hybrid integrated biological-solid-state system powered with adenosine triphosphate Nat. Commun. 6 10070
- [520] Wang W *et al* 2018 Learning of spatiotemporal patterns in a spiking neural network with resistive switching synapses *Sci. Adv.* **4** eaat4752
- [521] Huang X 2018 Materials and applications of bioresorbable electronics J. Semicond. 39 011003
- [522] Lecomte A, Giantomasi L, Rancati S, Boi F, Angotzi G N and Berdondini L 2020 Surface-functionalized self-standing microdevices exhibit predictive localization and seamless integration in 3D neural spheroids *Adv. Biosys.* **4** 2000114
- [523] Jafari A, Ganesan A, Thalisetty C S K, Sivasubramanian V, Oates T and Mohsenin T 2019 SensorNet: a scalable and low-power deep convolutional neural network for multimodal data classification *IEEE Trans. Circuits Syst.* I 66 274–87
- [524] Hosseini M, Ren H, Rashid H-A, Mazumder A N, Prakash B and Mohsenin T 2020 Neural networks for pulmonary disease diagnosis using auditory and demographic information epiDAMIK 2020: 3rd epiDAMIK ACM SIGKDD Int. Workshop on Epidemiology Meets Data Mining and Knowledge Discovery
- [525] Dauphin Y N and Bengio Y 2013 Big neural networks waste capacity (arXiv:1301.3583)
- [526] Gong Y, Liu L, Yang M and Bourdev L 2014 Compressing deep convolutional networks using vector quantization (arXiv:1412.6115)
- [527] Alemdar H, Leroy V, Prost A B and Pétrot F 2017 Ternary neural networks for resource-efficient AI applications Int. Joint Conf. Neural Networks (IJCNN) (Anchorage, AK, USA)
- [528] Courbariaux M, Bengio Y and David J 2015 Binaryconnect: training deep neural networks with binary weights during propagations Advances in Neural Information Processing Systems vol 28 3123–31

- [529] Lo C Y, Lau F C M and Sham C W 2018 Fixed-point implementation of convolutional neural networks for image classification Int. Conf. Advanced Technologies for Communications (ATC) (Ho Chi Minh City, Vietnam)
- [530] Umuroglu Y, Fraser N J, Gambardella G, Blott M, Leong P, Jahre M and Vissers K 2017 FINN: a framework for fast, scalable binarized neural network inference Proc. 2017 ACM/SIGDA Int. Symp. Field-Programmable Gate Arrays (New York: ACM) pp 65–74
- [531] Ren H, Mazumder A N, Rashid H-A, Chandrareddy V, Shiri A and Mohsenin T 2020 End-to-end scalable and low power multi-modal CNN for respiratory-related symptoms detection *IEEE 33rd Int. System-on-Chip Conf. (SOCC)*
- [532] Khatwani M, Rashid H-A, Paneliya H, Horton M, Waytowich N, Hairston W D and Mohsenin T 2020 A flexible multichannel EEG artifact identification processor using depthwise-separable convolutional neural networks ACM J. Emerg. Technol. Comput. Syst. 17 1–23
- [533] Enoka R M 1995 Morphological features and activation patterns of motor units J. Clin. Neurophysiol. 12 538-59
- [534] Rodriguez-Tapia B, Soto I, Martinez D M and Arballo N C 2020 Myoelectric interfaces and related applications: current state of EMG signal processing—a systematic review IEEE Access 8 7792–805
- [535] Arteaga M V, Castiblanco J C, Mondragon I F, Colorado J D and Alvarado-Rojas C 2020 EMG-driven hand model based on the classification of individual finger movements *Biomed. Signal Process. Control* 58 101834
- [536] Park K H and Lee S W 2016 Movement intention decoding based on deep learning for multiuser myoelectric interfaces 2016 4th Int. Winter Conf. Brain-Computer Interface (BCI) (Piscataway, NJ: IEEE) pp 1–2
- [537] Donati E, Payvand M, Risi N, Krause R, Burelo K, Indiveri G, Dalgaty T and Vianello E 2018 Processing EMG signals using reservoir computing on an event-based neuromorphic system 2018 IEEE Biomedical Circuits and Systems Conf. (BioCAS) (Piscataway, NJ: IEEE) pp 1–4
- [538] Donati E, Payvand M, Risi N, Krause R and Indiveri G 2019 Discrimination of EMG signals using a neuromorphic implementation of a spiking neural network *IEEE Trans. Biomed. Circuits Syst.* 13 795–803
- [539] Ceolini E, Frenkel C, Shrestha S B, Taverni G, Khacef L, Payvand M and Donati E 2020 Handgesture recognition based on EMG and event-based camera sensor fusion: a benchmark in neuromorphic computing *Front. Neurosci.* 14 637
- [540] Behrenbeck J et al 2019 Classification and regression of spatio-temporal signals using NeuCube and its realization on SpiNNaker neuromorphic hardware J. Neural Eng. 16 026014
- [541] Ma Y, Chen B, Ren P, Zheng N, Indiveri G and Donati E 2020 EMG-based gestures classification using a mixed-signal neuromorphic processing system IEEE J. Emerg. Sel. Top. Circuits Syst. 10 578–87
- [542] Azghadi M R, Lammie C, Eshraghian J K, Payvand M, Donati E, Linares-Barranco B and Indiveri G 2020 Hardware implementation of deep network accelerators towards healthcare and biomedical applications *IEEE Trans. Biomed. Circuits Syst.* 14 1138–59
- [543] Del Vecchio A, Germer C M, Elias L A, Fu Q, Fine J, Santello M and Farina D 2019 The human central nervous system transmits common synaptic inputs to distinct motor neuron pools during non-synergistic digit actions J. Physiol. 597 5935–48
- [544] Valle G et al 2018 Biomimetic intraneural sensory feedback enhances sensation naturalness, tactile sensitivity, and manual dexterity in a bidirectional prosthesis Neuron 100 37–45
- [545] Moin A et al 2020 A wearable biosensing system with in-sensor adaptive machine learning for hand gesture recognition Nat. Electron. 4 54–63
- [546] Urh F, Strnad D, Clarke A, Farina D and Holobar A 2020 On the selection of neural network architecture for supervised motor unit identification from high-density surface EMG 2020 42nd Annual Int. Conf. IEEE Engineering in Medicine & Biology Society (EMBC) (Piscataway, NJ: IEEE) pp 736–9
- [547] Kleine B U, van Dijk J P, Lapatki B G, Zwarts M J and Stegeman D F 2007 Using two-dimensional spatial information in decomposition of surface EMG signals J. Electromyogr. Kinesiol. 17 535–48
- [548] Holobar A and Zazula D 2007 Multichannel blind source separation using convolution kernel compensation IEEE Trans. Signal Process. 55 4487–96
- [549] Rossi D et al 2015 PULP: a parallel ultra low power platform for next generation IoT applications 2015 IEEE Hot Chips 27 Symp. (HCS) (Piscataway, NJ: IEEE) pp 1–39
- [550] Chatterjee S, Roy S S, Bose R and Pratiher S 2020 Feature extraction from multifractal spectrum of electromyograms for diagnosis of neuromuscular disorders IET Sci. Meas. Technol. 14 817–24
- [551] Trigili E, Grazi L, Crea S, Accogli A, Carpaneto J, Micera S, Vitiello N and Panarese A 2019 Detection of movement onset using EMG signals for upper-limb exoskeletons in reaching tasks J. Neuroeng. Rehabil. 16 45
- [552] Ajoudani A, Zanchettin A M, Ivaldi S, Albu-Schäffer A, Kosuge K and Khatib O 2018 Progress and prospects of the human-robot collaboration Auton. Robots 42 957–75
- [553] Ingrand F and Ghallab M 2017 Deliberation for autonomous robots: a survey Artif. Intell. 247 10-44
- [554] Kunze L, Hawes N, Duckett T, Hanheide M and Krajnik T 2018 Artificial intelligence for long-term robot autonomy: a survey IEEE Robot. Autom. Lett. 3 4023–30
- [555] Capolei M C, Angelidis E, Falotico E, Lund H H and Tolu S 2019 A biomimetic control method increases the adaptability of a humanoid robot acting in a dynamic environment *Front. Neurorobot.* **13** 70
- [556] Indiveri G and Sandamirskaya Y 2019 The importance of space and time for signal processing in neuromorphic agents: the challenge of developing low-power, autonomous agents that interact with the environment *IEEE Signal Process. Mag.* 36 16–28
- [557] Thompson F and Galeazzi R 2020 Robust mission planning for autonomous marine vehicle fleets Robot. Auton. Syst. 124 103404
- [558] Zool H I and Nohaidda S 2019 A survey and analysis of cooperative multi-agent robot systems: challenges and directions Applications of Mobile Robots vol 1 (IntechOpen)
- [559] Yang H, Han Q-L, Ge X, Ding L, Xu Y, Jiang B and Zhou D 2020 Fault-tolerant cooperative control of multiagent systems: a survey of trends and methodologies *IEEE Trans. Ind. Inf.* 16 4–17
- [560] Fardet T and Levina A 2020 Simple models including energy and spike constraints reproduce complex activity patterns and metabolic disruptions PLoS Comput. Biol. 16 e1008503
- [561] Naveros F, Garrido J A, Carrillo R R, Ros E and Luque N R 2017 Event- and time-driven techniques using parallel CPU-GPU co-processing for spiking neural networks *Front. Neuroinform.* 11 7
- [562] Furber S 2016 Large-scale neuromorphic computing systems J. Neural Eng. 13 051001
- [563] Thompson F and Guihen D 2019 Review of mission planning for autonomous marine vehicle fleets *J Field Robot.* 36 333–54
 [564] Atyabi A, MahmoudZadeh S and Nefti-Meziani S 2018 Current advancements on autonomous mission planning and
 - management systems: an AUV and UAV perspective Annu. Rev. Control 46 196-215

- [565] Li J, Li Z, Chen F, Bicchi A, Sun Y and Fukuda T 2019 Combined sensing, cognition, learning, and control for developing future neuro-robotics systems: a survey IEEE Trans. Cognit. Dev. Syst. 11 148–61
- [566] Woo J, Kim J H, Im J-P and Moon S E 2020 Recent advancements in emerging neuromorphic device technologies Adv. Intell. Syst. 2 2070101
- [567] Hauser S, Mutlu M, Léziart P-A, Khodr H, Bernardino A and Ijspeert A J 2020 Roombots extended: challenges in the next generation of self-reconfigurable modular robots and their application in adaptive and assistive furniture *Robot. Auton. Syst.* 127 103467
- [568] Galin R, Meshcheryakov R, Kamesheva S and Samoshina A 2020 Cobots and the benefits of their implementation in intelligent manufacturing IOP Conf. Ser.: Mater. Sci. Eng. 862 032075
- [569] Müller V C 2020 Ethics of artificial intelligence and robotics *The Stanford Encyclopedia of Philosophy (Winter 2020 Edition)* ed E N Zalta (https://plato.stanford.edu/archives/win2020/entries/ethics-ai/)
- [570] Coeckelbergh M 2020 AI Ethics (Cambridge, MA: MIT Press)
- [571] Topol E 2019 Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again (New York: Basic Books)
- [572] Cohen I 2018 Is there a duty to share healthcare data? *Big Data, Health Law, and Bioethics* ed I Cohen, H Lynch, E Vayena and U Gasser (Cambridge: Cambridge University Press) pp 209–22
- [573] Binns R 2018 Fairness in machine learning: lessons from political philosophy Proc. 1st Conf. Fairness, Accountability and Transparency, Proc. Machine Learning Research vol 81 149–59
- [574] Ploug T and Holm S 2020 The right to refuse diagnostics and treatment planning by artificial intelligence Med. Health Care Philos. 23 107–14
- [575] Nyholm S and Frank L 2017 From sex robots to love robots: is mutual love with a robot possible? *Robot Sex: Social and Ethical Implications* vol 219–243 (Cambridge, MA: MIT Press)
- [576] Danaher J 2019 The philosophical case for robot friendship J. Posthuman Stud. 3 5-24
- [577] Baldwin R 2019 The Globotics Upheaval: Globalisation, Robotics and the Future of Work (New York: Oxford University Press)
- [578] Goos M 2018 The impact of technological progress on labour markets: policy challenges Oxford Rev. Econ. Pol. 34 362-75
- [579] Turner J 2019 Robot Rules: Regulating Artificial Intelligence (Berlin: Springer)
- [580] Coeckelbergh M 2016 Care robots and the future of ICT-mediated elderly care: a response to doom scenarios *AI Soc.* 31 455–62
 [581] Bostrom N 2014 Superintelligence: Paths, Dangers, Strategies (Oxford: Oxford University Press)
- [581] Bostrom N 2014 Superintelligence: Paths, Dangers, Strategies (Oxford: Oxford University Press)
 [582] Floridi L 2016 Should we be afraid of AI? Machines seem to be getting smarter and smarter and much better at human jobs, yet true AI is utterly implausible. Why? (Aeon) https://aeon.co/essays/true-ai-is-both-logically-possible-and-utterly-implausible