Silicon Photonics for Neuromorphic Computing and Artificial Intelligence: Applications and Roadmap

B. J. Shastri[1,2,3], C. Huang[2], A. N. Tait[2], T. Ferreira de Lima[2], and P. R. Prucnal[2]

[1]Department of Physics, Engineering Physics & Astronomy
Queen's University, Kingston, ON K7L 3N6, Canada
[2]Department of Electrical Engineering, Princeton University
Princeton, NJ 08544, USA
[3]Vector Institute for Artificial Intelligence, Toronto, ON M5G 1M1, Canada

Abstract— Artificial intelligence and neuromorphic computing driven by neural networks has enabled many applications. Software implementations of neural networks on electronic platforms are limited in speed and energy efficiency. Neuromorphic photonics aims to build processors in which optical hardware mimic neural networks in the brain. These processors promise orders of magnitude improvements in both speed and energy efficiency over purely digital electronic approaches. However, integrated optical neural networks are much smaller (hundreds of neurons) than electronic implementations (tens of millions of neurons). This raises a question: what are the applications where sub-nanosecond latencies and energy efficiency trump the sheer size of processor? We provide an overview of neuromorphic photonic systems and their real-world applications to machine learning and neuromorphic computing.

1. INTRODUCTION

Neuromorphic (i.e., neuron-isomorphic) computing aims to bridge the gap between the energy efficiency of von Neumann computers and the human brain [1, 2]. The rise of neuromorphic computing can be attributed the widening gap between current computing capabilities and current computing needs [3, 4]. Consequently, this has spawned research into novel brain-inspired algorithms and applications uniquely suited to neuromorphic processors. These algorithms attempt to solve artificial intelligence (AI) tasks in real-time while using less energy. We posit that we can make use of the

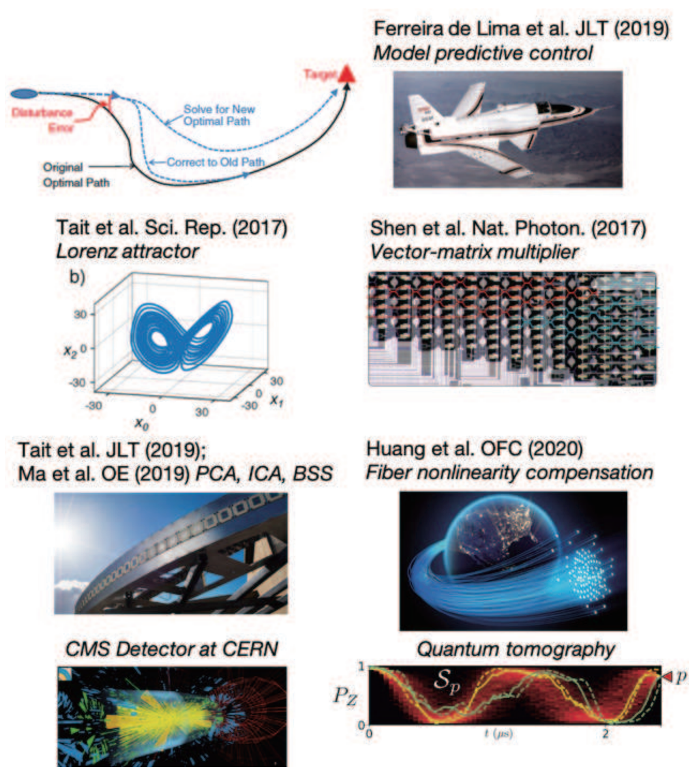

Figure 1: Examples of recently demonstrated photonic neural network architectures.

2022 Photonics & Electromagnetics Research Symposium (PIERS) | 978-1-6654-6023-1/22/$31.00 ©2022 IEEE | DOI: 10.1109/PIERS55526.2022.9792850

high parallelism and speed of photonics to bring the same neuromorphic algorithms to applications requiring multiple channels of multi-gigahertz analog signals, which digital processing struggles to process in real-time.

By combining the high bandwidth and parallelism of photonic devices with the adaptability and complexity attained by methods similar to those seen in the brain, photonic neural networks (PNNs) have the potential to be orders of magnitude faster than state-of-the-art electronic processors while consuming less energy per computation [5]. The goal of neuromorphic photonic processors is not to replace conventional computers, but to enable applications that are unreachable at present by conventional computing technology — those requiring low latency, high bandwidth and low energies [6, 7]. As shown in Figure 1, examples of applications for ultrafast neural networks include: 1) Enabling fundamental physics breakthroughs: qubit read-out classification, high-energy-particle collision classification, fusion reactor plasma control; 2) Nonlinear programming: solving nonlinear optimization problems (robotics, autonomous vehicles, predictive control) and partial differential equations; 3) Machine learning acceleration: vector–matrix multiplications, deep learning inference, ultrafast or online learning; 4) Intelligent signal processing: wideband radio-frequency signal processing, fibre-optic communication.

2. NEUROMORPHIC PHOTONICS APPROACHES

Neuromorphic photonic [6, 8] approaches can be divided into two main categories (Figure 2): coherent (single wavelength) and incoherent (multiwavelength) approaches. Neuromorphic systems based on reservoir computing [9–11] and Mach-Zehnder interferometers [12, 13] are example of coherent approaches. In reservoir computing the predefined random weights of their hidden layers cannot be modified. An alternative approach uses silicon photonics to design fully programmable neural networks [14, 15], with a so-called broadcast-and-weight protocol [16]. In this architecture, photonic neurons output optical signals with unique wavelengths. These are multiplexed into a single waveguide and broadcast to all others, weighted, and photodetected. Each connection between a pair of neurons is configured independently by one microring resonator (MRR) weight, and the wavelength division multiplexed (WDM) carriers do not mutually interfere when detected by a single photodetector. Consequently, the physics governing the neural computation is fully analog and does not require any logic operation or sampling, which would involve serialization and sampling. Thus, they exhibit distinct, favorable trends in terms of energy dissipation, latency, crosstalk, and bandwidth when compared to electronic neuromorphic circuits [5]. The advantage of this approach over the aforementioned approaches is that it has already demonstrated fan-in, inhibition, time-resolved processing, and autaptic cascadability [15].

However, the same physics also introduce new challenges, especially reconfigurability, integration, and scalability. Information carried by photons is harder to manipulate compared to electronic signals, especially nonlinear operations and memory storage [6]. Photonic neurons described

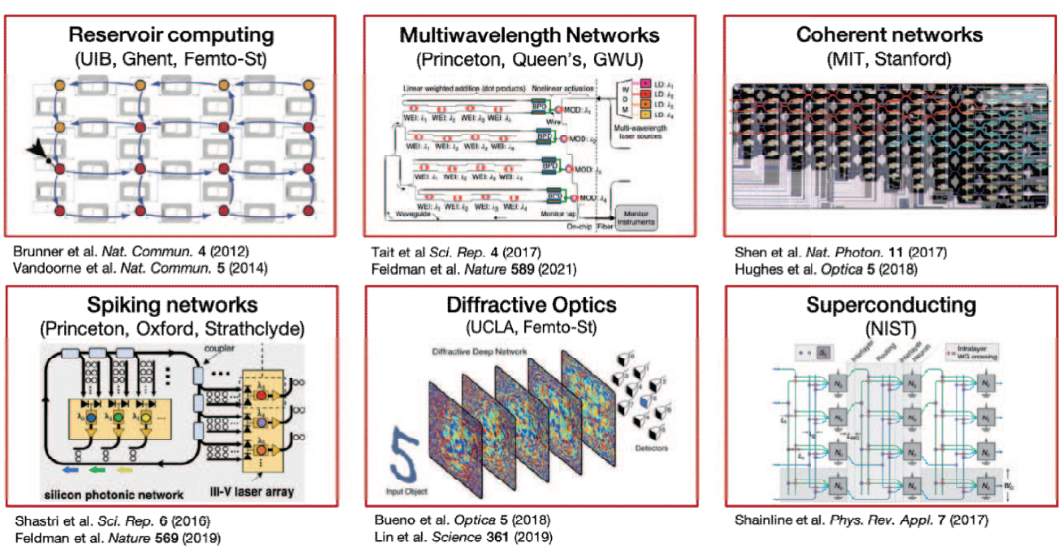

Figure 2: Examples of recently demonstrated photonic neural network architectures.

here solve that problem by using optoelectronic components (O/E/O), which can be mated with standard electronics providing reconfigurability. However, neuromorphic photonic circuits are challenging to scale up because they do not benefit from digital information, memory units and a serial processor, and therefore requires a physical unit for each element in a neural network, increasing size, area and power consumption. Here, integration costs must also be considered, since the advantages of using analog photonics (high parallelism and high bandwidth) must outweigh the costs of interfacing it with digital electronics (requiring both O/E and analog/digital conversion).

3. VISION OF A NEUROMORPHIC PROCESSOR

Recently, in our tutorial, Ref. [17], we proposed a vision for a neuromorphic processor. We discussed how such a neuromorphic chip could potentially be interfaced with a general-purpose computer (Figure 3), i.e., a CPU, as a coprocessor to target specific applications. In general, there are two levels of complexity associated with co-integrating a general-purpose electronic processor with an application-specific optical processor. Firstly, a CPU processes a series of computation instructions in an undecided amount of time and is not guaranteed to be completed. Neural networks, on the other hand, can process data in parallel and in a deterministic amount of time. CPUs have a concept of a 'fixed' instruction set on top of which computer software can be developed. However, a neuromorphic processor would require a hardware description language (HDL) because it describes the intended behavior of a hardware in real-time. Secondly, seamlessly interfacing a photonic integrated circuit with an electronic integrated circuit will take several advances in science and technology including on-chip lasers and amplifiers, co-integration of CMOS with silicon photonics, system packaging, high-bandwidth digital-to-analog converters (DAC) and analog-to-digital converters (ADCs).

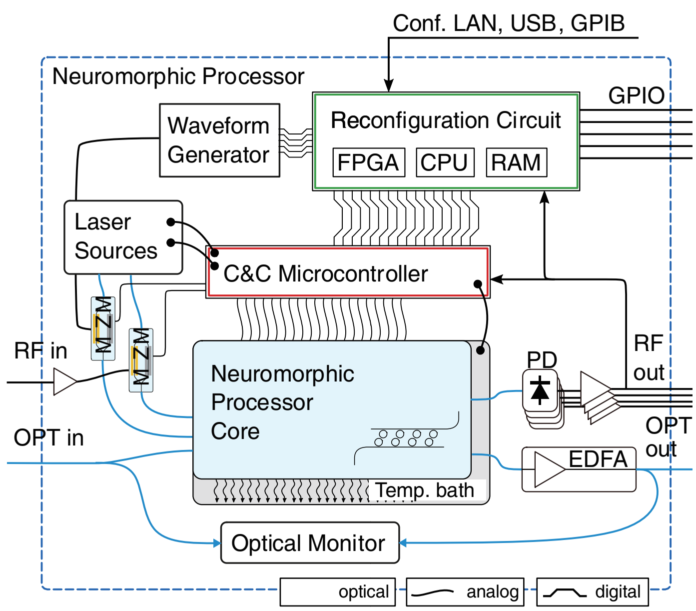

Figure 3: Simplified schematics of a neuromorphic processor. Thanks to integrated laser sources and photodetectors, it can input and output RF signals directly as an option to optically modulated signals. The waveform generator allows for programming arbitrary stimulus that can be used as part of a machine learning task. Reproduced from [17].

4. MACHINE LEARNING APPLICATION: FIBER NONLINEARITY IMPAIRMENT COMPENSATION

The world is witnessing an explosion of internet traffic. The global internet traffic has reached 5.3 exabytes per day in 2020 and will continue doubling approximately every 18 months. Innovations in fiber communication technologies are required to sustain the long-term exponential growth of data traffic [18]. Increasing data rate and system crosstalk has imposed significant challenges on the digital signal processing (DSP) chips in terms of analog-to-digital converters (ADCs) performances, circuit complexity, and power consumption. A key to advancing the deployment of DSP relies on the consistent improvement in CMOS technology [19]. However, the exponential hardware scaling of ASIC based DSP chips, which is embodied in Moore's law as other digital electronic hardware, is fundamentally unsustainable. In parallel, many efforts are focused on developing new DSP algorithms to minimize computational complexity, but usually at the expense of reducing transmission link performances [20].

Instead of embracing such a complexity-performance trade-off, an alternative approach is to explore new hardware platforms that intrinsically offer high bandwidth, speed, and low power consumption [6, 21, 22]. Machine learning algorithms, especially neural networks, have been found effective in performing many functions in optical networks, including dispersion and nonlinearity impairments compensation, channel equalization, optical performance monitoring, traffic prediction, etc [23].

PNNs are well suited for optical communications because the optical signals are processed directly in the optical domain. This innovation avoids prohibitive energy consumption overhead and speed reduction in ADCs, especially in data center applications. In parallel, many PNN approaches are inspired by optical communication systems, making PNNs naturally suitable for processing optical communication signals. For example, we proposed synaptic weights and neuron networking architecture based on the concept of WDM to enable fan-in and weighted addition [14]. This architecture can provide a seamless interface between PNNs and WDM systems, which can be applied as a front-end processor to address inter-wavelength or inter-mode crosstalks problems that DSP usually lacks the bandwidth or computing power to process (e.g., fiber nonlinearity compensation in WDM systems). Moreover, PNNs combine high-quality waveguides and photonic devices that have been initially developed for telecommunications. Therefore, PNNs, by default, can support fiber optic communication rates and enable real-time processing. For example, the scalable silicon PNN proposed by the authors is composed of microring resonator (MRR) banks for synaptic weighting and O/E/O neurons to produce standard machine learning activation functions. The MRR weight bank is inspired by WDM filters, and the O/E/O neurons use typical silicon photodetector and modulator. Therefore, the optimization of associated devices in PNNs can utilize the fruits of the entire silicon photonic ecosystem that is driven by telecommunications and data center applications.

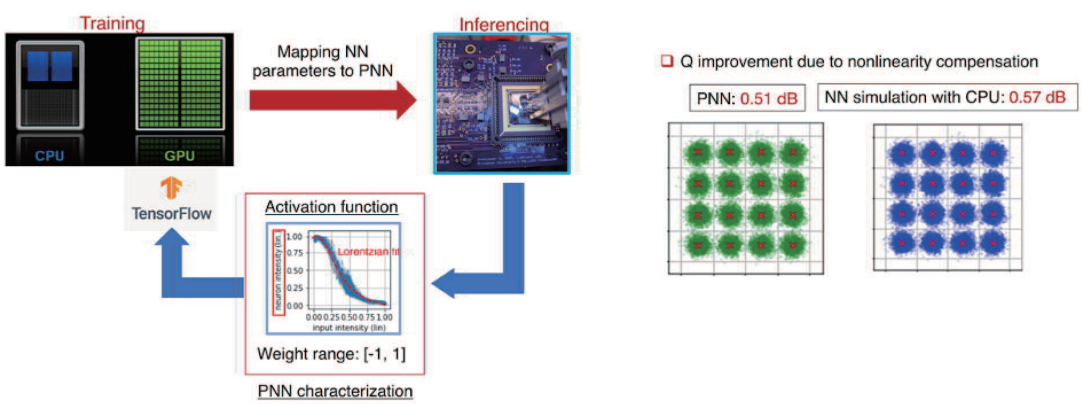

(a) (b)

Figure 4: (a) Concept of training and implementing photonic neural networks. Inset shows a transfer function of the photonic neural network measured with real-time signals. (b) Constellations of X-polarization of a 32 Gbaud PM-16QAM, with the ANN-NLC gain of 0.57 dB in Q-factor and with the PNN-NLC gain of 0.51 dB in Q-factor [22].

To truly demonstrate photonics can excel over DSP, careful considerations are required to identify different application scenarios (i.e., long-haul, short reach) and system requirements (i.e., performances, energy). Continuous research is needed to improve photonic hardware and to develop hardware-compatible algorithms. Here, we discuss several approaches to train and apply PNNs for optical communications.

Long-haul communication systems prioritize high performances in terms of distance reach and spectral efficiency. This requirement allows the use of coherent technology, along with dense wavelength multiplexing and polarization multiplexing schemes, to maximize the fiber capacity. In long-haul fiber optic transmission systems, fiber nonlinearity remains a challenge to the achievable

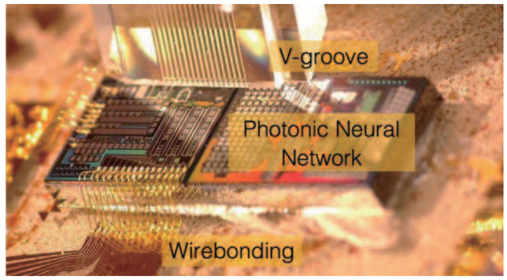

Figure 5: Micrograph image of a wirebonded photonic neural network chip from Princeton University's Lightwave Lab.

capacity and transmission distance. One reason is that the nonlinear interplay between signal, noises, and optical fibers negates the accuracy of conventional nonlinear compensation algorithms based on digital backpropagation. Another reason is, the implementation of most nonlinear compensation algorithms in DSP chips demands excessive resources. In contrast, the neural network approach can learn and approximate the nonlinear perturbation from the abundant training data, rather than solely relying on the physical fiber model (known as stochastic nonlinear Schrodinger equation). Based on the perturbation methods, the derived neural network algorithm has enabled compensating the nonlinear distortion in a 10800 km fiber transmission link with 32 Gbaud signals [24]. In Ref. [15], we developed a PNN platform based on the so-called "neuromorphic" approach, aiming to map physical models of optoelectronic systems to abstract models of neural networks (which differs from the reservoir approaches). By doing so, the PNN system can leverage existing machine learning algorithms (i.e., backpropagation) and map training results from simulations to heterogeneous photonic hardware. The concept is shown in Figure 2. A proof-of-concept experiment demonstrates the real-time implementation of a trained neural network model using an integrated silicon PNN chip [22]. In this work, the authors experimentally demonstrated that the silicon PNN can produce a similar Q factor improvement compared to the simulated neural network for nonlinear compensation as shown in Figure 4, but it promises to process the communication data in real-time and with high bandwidth and low latency.

We also proposed a photonic architecture enabling all-to-all continuous-time recurrent neural networks (RNN) [15]. Recurrent neural networks can resemble optical fiber transmission systems: the linear neuron-to-neuron connections with internal feedback is analog to linear multiple-input multiple-output (MIMO) fiber channel with dispersive memory. With neuron nonlinearity, RNNs can be ideally used to approximate all types of linear and nonlinear effects in a fiber transmission system and compensate for different transmission impairments. RNNs, consisting of many feedback connections, are computationally expensive for digital hardware and require at least milliseconds to conduct a single inference. Contrarily, in photonic RNN, the feedback operations are simply done by busing the signals on photonic waveguides, allowing photonic hardware capable of converging to the solution within microseconds. This architecture also adopts the neuromorphic approach and thus allows to train PNNs externally using standard machine learning algorithms.

5. NEUROMORPHIC COMPUTING APPLICATION: NONLINEAR PROGRAMMING — SOLVING OPTIMIZATION PROBLEMS AND ORDINARY DIFFERENTIAL EQUATIONS

Solving mathematical optimization problems lies at the heart of various applications present in modern technology such as machine learning, resource optimization in wireless networks, and drug discovery. Many optimization problems can be written as a quadratic program. For example, the least squares regression method can be mathematically mapped to a relatively easy quadratic program (with a positive definite quadratic matrix). Quadratic programming (QP) refers to algorithms related to solving the optimization problem of finding the extremes of a quadratic objective function subject to linear constraints. The general formulation of a quadratic program is usually solved iteratively, often requiring many time steps to reach the desired solution. The difficulty of QP grows exponentially with the dimension of the problem. Algorithms that can deal with large dimensions involve more computationally intensive techniques such as genetic algorithms or particle swarm optimization. As a result, conventional digital computers must either be limited to solving quadratic programs of very few variables, or to applications where the computation time is non-critical. Therefore, traditional computers are not appropriate to implement algorithms depending on QP for high-speed applications such as signal processing and control systems. In machine learning, many algorithms, such as SVM, require offline training because of the computational complexity of QP, but would be much more effective were they trained online.

A number of high-speed control problems, e.g., controlling plasma in aircraft actuations, fusion power plants, guiding of drones etc., are currently bottlenecked by the speed and latency of the control algorithms. Model predictive control (MPC) is an advanced technique to control complex systems, outperforming traditional PID control methods because it can predict control violations, rather than react to it. However, its control loop involves solving a quadratic problem at every control step, and therefore it is not computationally tractable for systems requiring speeds higher than kHz. Photonic neural networks help overcome this tradeoff by using techniques such as wavelength division multiplexing, which enables hundreds of high bandwidth signals (20 GHz) to be guided through a single optical waveguide.

Neural networks were demonstrated to solve general-purpose quadratic optimization problems by Hopfield and Tank [25, 26]. Until now, Hopfield networks have not been commonly implemented in hardware due to its all-to-all connectivity, which creates an undesirable tradeoff between neuron speed and neural network size — in an electronic circuit, as the number of connections increases, the bandwidth at which the system can operate decreases [14]. This means a photonic Hopfield network implementation can simultaneously tackle quadratic programs with large dimensions and converge in nanoseconds [7]. Implemented in photonic hardware, model predictive control can be employed in systems operating in the MHz regime.

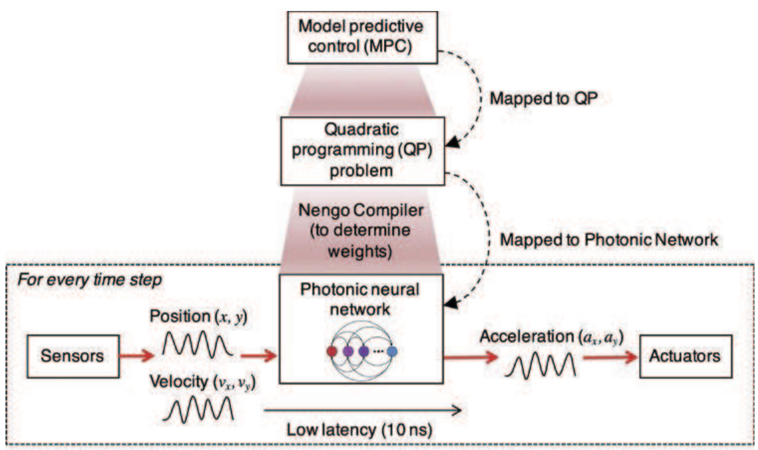

Figure 6: Schematic figure of the procedure to implement the MPC algorithm on a neuromorphic photonic processor. Firstly, map the MPC problem to QP. Then, construct a QP solver with continuous-time recurrent neural networks (CT-RNN) (Cichocki et al., 1993). Finally, build a neuromorphic photonic processor to implement the CT-RNN. The details of how to map MPC to QP, and how to construct a QP solver with CT-RNN are given in De Lima et al. (2019). Adapted from Ferreira De Lima et al. (2019) [7].

While neural networks are often used for their learning properties, they can also be programmed directly. Direct programming of analog systems has a major pitfall in that the components are unreliable and subject to parameter variation. One approach is to represent variables as population states that are robust to parameter variations. The neural engineering framework (NEF) [27] provides an algorithm to program ensembles of imperfect analog devices to perform operations on population coded variables.

Tait et al. [15] demonstrated a programmable network of two photonic neurons. The NEF algorithm was fed the responses of photonic devices, resulting in a weight matrix that allows the network to approximate variables, operations, and differential equations. It was shown in simulation how 24 neurons could emulate the Lorenz attractor. The approximation improves with more neurons.

The example of the Lorenz attractor is an example of a task-based benchmark for a 24-neuron network. Task-based benchmarks play an important role in validating experimental systems as they begin to incorporate more neurons and parameters. The example also demonstrated the compatibility between PNNs and the NEF, including all the NEF's key principles and consequent

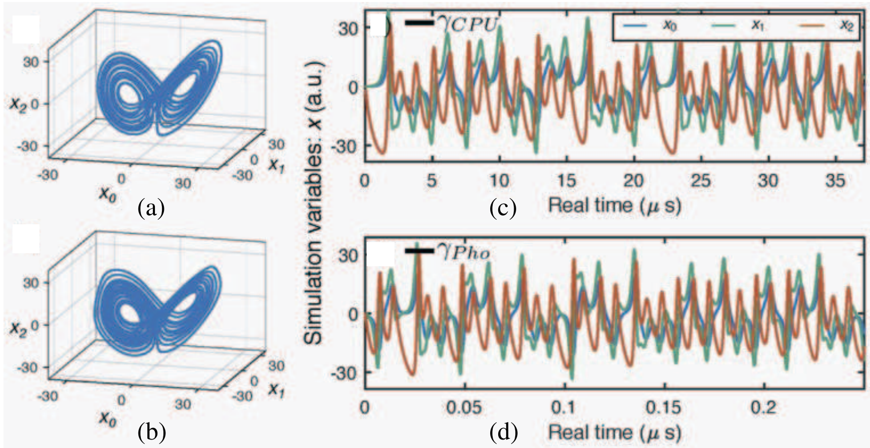

Figure 7: Photonic neural network benchmarking against a CPU. (a), (b) Phase diagrams of the Lorenz attractor simulated by (a) a conventional CPU and (b) a photonic neural network. (c), (d) Time traces of simulation variables for (c) a conventional CPU and (d) a photonic CTRNN. The horizontal axes are labeled in physical real time, and cover equal intervals of virtual simulation time, as benchmarked by γCPU and γPho. The ratio of real-time values of γ's indicates a 294-fold acceleration. From [15].

functionality. Using this neural compiler provides a route to a variety of known applications and other benchmarks [28].

6. ROADMAP SUMMARY: ADVANCES IN SCIENCE AND TECHNOLOGY TO MEET CHALLENGES

The renaissance of field of neuromorphic photonics is enabled by the confluence of three area (Figure 8): technological advances in integrated photonics due to silicon photonics, the algorithmic advances in machine learning algorithms, and advances in analogue photonic signal processing. In the recent roadmap article [29], we outlined some scientific and technological advances necessary to meet the challenges to envision a practical neuromorphic processor outlined in the previous section.

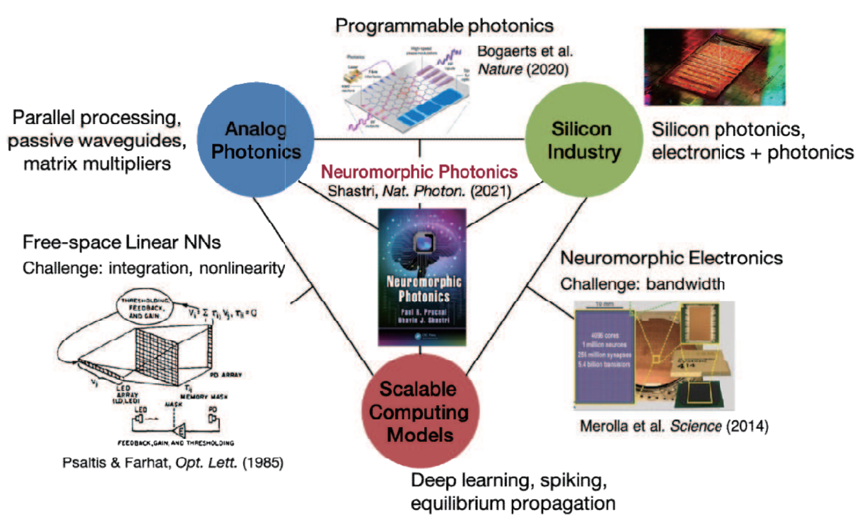

Figure 8: The advent of neuromorphic photonics is due to the convergence of recent advances in photonic integration technology, resurgence of scalable computing models (e.g., spiking, deep neural networks), and a large-scale silicon industrial ecosystem.

Photonic processors have light sources, passive and active devices. Currently, there is no single commercial fabrication platform that can simultaneously offer devices for light generation, wavelength multiplexing, photodetection, and transistors on a single die; state-of-the-art devices in each of these categories use different photonic materials (SiN, Ge, InP, GaAs, 2D materials, etc) with incongruous fabrication processes (silicon-on-insulator, CMOS, FinFETs). Silicon photonics is becoming an ideal platform for integrating these devices while offering a combination of foundry compatibility, device compactness, and cost that enables the creation of scalable photonic systems on chip.

Materials: Energy efficient and fast switching optical and electro-optical materials are needed for non-volatile photonic storage and weighting, as well as high-speed optical switching and routing, with low power consumption. Neural non-linearities are already possible on mainstream platforms using electrooptic transfer functions [16], but new materials promise significant performance opportunities. Phase change materials (PCMs), and graphene and ITO-based modulators can also be utilized for implementing non-linearities. Plasmonic PCMs can bridge the optical and electrical signals, through the dual operation modes [30]. A general material design method is in urgent need to develop appropriate photonic materials for different photonic components [31].

Lasers and amplifiers: On-chip optical gain and power will require co-integration with active InP lasers and semiconductor optical amplifiers. Current approaches involve either III–V to silicon wafer bonding (heterogeneous integration) or co-packaging with precise assembly (hybrid approach) [32]. Quantum dot lasers are another promising approach as they can be grown directly onto silicon, but fabrication reliability does not currently reach commercial standards [33].

Electrical control: Co-integrating CMOS controller chips with silicon photonics to provide electrical tuning control/stabilization will be critical. Candidates include wire-bonding, flip-chip bonding, 2.5D integration (interposers), 3D stacking (through-silicon-vias), and monolithic integration. Each has performance and design tradeoffs [34].

System packaging: A photonic processor must be interfaced with a computer. It would need to be self-contained, robust to temperature fluctuations, and with electrical inputs/outputs [6]. Currently, manufacturers do not assemble electrical/thermal elements and chip-to-fiber interconnects.

Algorithms: Significant advances will be required to map abstract neural algorithms to photonic processor to usher these platforms into the commercial space. So far, only individual devices and small control circuits are described in the literature. The goal is to enable neural network programming tools (TensorFlow) to directly reconfigure a neuromorphic photonic processor [7].

7. CONCLUSION

Neuromorphic photonics has reached an inflection point, benefiting from great opportunities as the world looks for alternative processor architectures. The physical limits of Dennard scaling is galvanizing the community to put forward candidates for next generation computing, from bio- to quantum computers. Photonics and in particular neuromorphic photonics are a formidable candidate for analog reconfigurable processing. We expect the development of this field to accelerate as neuroscience makes further leaps towards our understanding of the nature of cognition and artificial intelligence demands more computational resources for machine learning. As photonics technology matures and becomes more accessible to academic groups and small companies, we expect this acceleration to continue.

REFERENCES

1. Schuman, C. D., et al., "A survey of neuromorphic computing and neural networks in hardware," arXiv preprint arXiv:1705.06963, 2017.
2. Hasler, J. and H. B. Marr, "Finding a roadmap to achieve large neuromorphic hardware systems," *Front. Neurosci.*, Vol. 7, 118, 2013.
3. Merolla, P. A., et al., "A million spiking-neuron integrated circuit with a scalable communication network and interface."*Science*, Vol. 345, No. 6197, 668–673, 2014.
4. Davies, M., et al., "Loihi: A neuromorphic manycore processor with on-chip learning," *IEEE Micro*, Vol. 38, No. 1, 82–99, 2018.
5. Ferreira De Lima, T., et al., "Progress in neuromorphic photonics," *Nanophotonics*, Vol. 6, No. 3, 577–599, 2017.
6. Shastri, B. J., et al., "Photonics for artificial intelligence and neuromorphic computing," *Nat. Photon.*, Vol. 15, No. 2, 102–114, 2021.
7. Ferreira De Lima, T., et al., "Machine learning with neuromorphic photonics," *J. Lightwave Technol.*, Vol. 37, No. 5, 1515–1534, 2019.
8. Prucnal, P. R. and B. J. Shastri, *Neuromorphic Photonics*, CRC Press, 2017.

9. Brunner, D., et al., "Parallel photonic information processing at gigabyte per second data rates using transient states," *Nat. Commun.*, Vol. 4, No. 1, 1–7, 2013.
10. Vandoorne, K., et al., "Experimental demonstration of reservoir computing on a silicon photonics chip," *Nat. Commun.*, Vol. 5, No. 1, 1–6, 2014.
11. Larger, L., et al., "Photonic information processing beyond Turing: an optoelectronic implementation of reservoir computing," *Opt. Express*, Vol. 20, No. 3, 3241–3249, 2012.
12. Shen, Y., et al., "Deep learning with coherent nanophotonic circuits," *Nat. Photon.*, Vol. 11, No. 7, 441–446, 2017.
13. Hughes, T. W., et al., "Training of photonic neural networks through in situ backpropagation and gradient measurement," *Optica*, Vol. 5, No. 7, 864–871, 2018.
14. Tait, A. N., et al., "Neuromorphic photonic networks using silicon photonic weight banks," *Sci. Rep.*, Vol. 7, No. 1, 1–10, 2017.
15. Tait, A. N., et al., "Silicon photonic modulator neuron," *Phys. Rev. Appl.*, Vol. 11, No. 6, 064043, 2019.
16. Tait, A. N., et al., "Broadcast and weight: an integrated network for scalable photonic spike processing," *J. Lightwave Technol.*, Vol. 32, No. 21, 4029–4041, 2014.
17. Ferreira De Lima, T., et al., "Primer on silicon neuromorphic photonic processors: architecture and compiler," *Nanophotonics*, Vol. 9, No. 13, 4055–4073, 2020.
18. Cisco. (n.d.), "Cisco annual internet report — cisco annual internet report (2018–2023) white paper," https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html.
19. Pillai, B. S. G., et al., "End-to-end energy modeling and analysis of long-haul coherent transmission systems," *J. Lightwave Technol.*, Vol. 32, No. 18, 3093–3111, 2014.
20. Agrell, Erik, et al., "Roadmap of optical communications," *J. Opt.*, Vol. 18, No. 6, 063002, 2016.
21. Argyris, A., J. Bueno, and I. Fischer, "Photonic machine learning implementation for signal recovery in optical communications," *Sci. Rep.*, Vol. 8, No. 1, 1–13, 2018.
22. Huang, C., et al., "Demonstration of photonic neural network for fiber nonlinearity compensation in long-haul transmission systems," *2020 Optical Fiber Communications Conference and Exhibition, OFC). IEEE*, 2020.
23. Khan, F. N., et al., "An optical communication's perspective on machine learning and its applications," *J. Lightwave Technol.*, Vol. 37, No. 2, 493–516, 2019.
24. Zhang, S., et al., "Field and lab experimental demonstration of nonlinear impairment compensation using neural networks," *Nat. Commun.*, Vol. 10, No. 1, 1–8, 2019.
25. Cochocki, A. and R. Unbehauen, *Neural Networks for Optimization and Signal Processing*, John Wiley &, Vol. Sons, Inc., 1993.
26. Hopfield, J. J. and D. W. Tank, "Computing with neural circuits: A model," *Science*, Vol. 233, No. 4764, 625–633, 1986.
27. Stewart, T. C. and C. Eliasmith, "Large-scale synthesis of functional spiking neural circuits," *Proc. IEEE*, Vol. 102, No. 5, 881–898, 2014.
28. Stewart, T. C., et al., "Closed-loop neuromorphic benchmarks," *Front. Neurosci.*, Vol. 9, 464, 2015.
29. Berggren, K., et al., "Roadmap on emerging hardware and technology for machine learning," *Nanotechnology*, Vol. 32, No. 1, 012002, 2020.
30. Farmakidis, N., et al., "Plasmonic nanogap enhanced phase-change devices with dual electrical-optical functionality," *Sci. Adv.*, Vol. 5, No. 11, eaaw2687, 2019.
31. Zhang, W., et al., "Designing crystallization in phase-change materials for universal memory and neuro-inspired computing," *Nat. Rev. Mater.*, Vol. 4, No. 3, 150–168, 2019.
32. Liang, D. and J. E. Bowers, "Recent progress in lasers on silicon," *Nat. Photon.*, Vol. 4, No. 8, 511–517, 2010.
33. Chen, S., et al., "Electrically pumped continuous-wave III – V quantum dot lasers on silicon," *Nat. Photon.*, Vol. 10, No. 5, 307–311, 2016.
34. Atabaki, A. H., et al., "Integrating photonics with silicon nanoelectronics for the next generation of systems on a chip," *Nature*, Vol. 556, No. 7701, 349–354, 2018.