Neuromorphic Computing Hardware with Silicon Photonics

B. J. Shastri[1], A. N. Tait[2], M. A. Nahmias[2], T. Ferreira de Lima[2], H.-T. Peng[2], and P. R. Prucnal[2]

[1]Queen's University, Kingston, Canada; [2]Princeton University, Princeton, USA
shastri@ieee.org

Artificial Intelligence (AI) is transforming our lives in the same way as the advent of the Internet and cellular phones has done. AI is revolutionizing the healthcare industry with complex medical data analysis, actualizing self-driving cars, and beating humans at strategy games such as Go. However, it takes thousands of CPUs and GPUs, and many weeks to train the neural networks in AI hardware. Over the last six years, this compute power has doubled every 3.5 months. Traditional CPUs, GPUs, and neuromorphic (i.e., brain-inspired) electronics such as the IBM TrueNorth and Google TPU will not be powerful enough to train the neural networks of the near future. There is a global race to solve this challenge.

Neuromorphic photonics [1] has experienced a recent surge of interest over the last few years, promising orders of magnitude improvements in both speed and energy efficiency over digital electronics using: artificial neural networks [2-5], spiking neural networks [6-8], and reservoir computing [9]. By combining the high bandwidth and efficiency of photonic devices with the adaptive, parallelism and complexity attained by methods similar to those seen in the brain, photonic processors have the potential to be at least ten thousand times faster than state-of-the-art electronic processors while consuming less energy/computation [10].

We will provide an overview of neuromorphic photonic systems and their application to optimization and machine learning problems. We will discuss the physical advantages of photonic processing systems and describe underlying device models that allow practical systems to be constructed. We also describe several real-world applications for control and deep learning inference. Lastly, we will discuss scalability in the context of designing a full-scale neuromorphic photonic processing system, considering aspects such as signal integrity, noise, and hardware fabrication platforms. This talk is intended for a wide audience and hopes to teach how theory, research, and device concepts from neuromorphic photonics could be applied in practical machine learning systems.

[1] P. R. Prucnal and B. J. Shastri. *Neuromorphic Photonics* (CRC Press, 2017).
[2] A. N. Tait et al. *arXiv preprint* arXiv:1812.11898 (2018).
[3] Y. Shen et al. *Nat. Photon.* **11**, 441 (2017).
[4] J. M. Shainline et al. *Phys. Rev. Appl.* **7**, 034013 (2017).
[5] A. N. Tait et al. *Sci. Rep.* **7**, 7430 (2017).
[6] H. T. Peng et al. *IEEE J. Sel. Top. Quantum Electron.* **24**, 6101715 (2018).
[7] B. J. Shastri et al. *Sci. Rep.* **5**, 19126 (2016).
[8] M. A. Nahmias et al. *IEEE J. Sel. Top. Quantum Electron.* **19**, 1800212 (2013).
[9] G. Van der Sande et al. *Nanophotonics*, **6**, 561 (2017).
[10] A. N. Tait et al. *J. Lightwave Technol.* **32**, 4029 (2014).