Photonic Neural Cellular Automata for Self-Organized Image Classification

Gordon H.Y. Li[1], Christian R. Leefmans[1], James Williams[2], Robert M. Gray[2], Midya Parto[2], and Alireza Marandi[1,2,*]

[1]*Department of Applied Physics, California Institute of Technology, Pasadena, CA 91125, USA.*
[2]*Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125, USA.*
**marandi@caltech.edu*

Abstract: Neural networks based on Cellular Automata (CA) have recently yielded more robust, reliable, and parameter-efficient machine learning models. We experimentally demonstrate the first photonic implementation of CA which successfully performs image classification on the Fashion-MNIST dataset.

Photonic Neural Networks (PNNs) are emerging as faster and more energy-efficient hardware accelerators for deep learning computing workloads [1]. So far, much of the focus has been on developing the devices needed for crucial neural network operations such as linear matrix-vector multiplication [2], convolutions [3], and nonlinear activations [4]. These building blocks are now comparable or surpassing their electronic counterparts and can enable scalable end-to-end PNNs. However, finding system-level architectures that maximize the advantages of photonics whilst minimizing potential disadvantages remains an open problem. The conventional wisdom has been to apply the photonic building blocks to implement deep learning architectures resembling those that have been widely successful for digital electronics, such as multi-layer perceptrons (MLPs) and convolutional neural networks (CNNs) [5]. We argue that these architectures may not be well suited for photonics, which operates in a fundamentally different fashion compared to digital electronics. For instance, MLPs and CNNs are inherently brittle and susceptible to noise or adversarial attacks [6], which is exacerbated by the noisy and analog nature of photonic devices. Furthermore, these architectures neither generalize nor recognize out-of-distribution data [7], which raises artificial intelligence safety concerns and necessitates online learning that is infeasible for fixed-weight PNNs. Finally, current PNNs contain only a relatively small number of programmable parameters, typically $< 10^4$, which greatly limits their real-world usefulness for implementing state-of-the-art neural networks containing billions of parameters [8].

To address these limitations, we experimentally demonstrate a PNN based on the recently proposed architecture of neural cellular automata (NCA) [9] building on our first time-multiplexed photonic realization of Cellular Automata (CA) [10]. CA are a class of computational models consisting of a regular lattice of cells, which have states that update according to local interactions with neighboring cells [11]. Remarkably, repeated iteration of even very simple local update rules can yield complex and emergent phenomena including universal computation [12]. To harness this complexity to perform useful computation, NCA use modern deep learning techniques to learn the local update rule needed to perform a specific task [9]. These NCA have been shown to be more robust, reliable, and parameter-efficient compared to conventional MLPs and CNNs for image processing tasks [13]. Crucially, the local update rules of CA require only sparse connections and fixed weights, which favors photonics and is poorly optimized on current electronic hardware accelerators.

As a proof-of-concept, our photonic NCA uses standard optical fiber and tabletop components, with the experimental schematic shown in Fig. 1(a). Lattice sites are represented by short pulses of light with a fixed repetition rate produced by a mode-locked laser. An electro-optic modulator (EOM) is used to encode cell states onto the amplitudes of light pulses. The light pulses are then split into three paths with different delays and variable optical attenuator (VOA) in each delay line to induce interactions between different light pulses upon coherent recombination. Finally, the recombined light pulses pass through a reverse-proton-exchanged periodically-poled lithium niobate (PPLN) waveguide [14], which performs an all-optical nonlinear activation function based on pump-depleted and partially phase-matched second harmonic generation shown in Fig. 1(b). Therefore, we make use of time-multiplexing to map the lattice of cells onto a synthetic temporal dimension, which allows us to reuse a single nonlinear node. The resultant light pulses are photodetected, then stored on a field-programmable-gate-array (FPGA), which drives the EOM to input the measured cell states for the next iteration. This photonic NCA can be interpreted as having an update rule encoded by a single perceptron neuron composed of a linear weighted sum followed by a nonlinear activation. In this way, the cell neighborhood participating in the update rule, akin to the receptive field of a 1D convolution, is determined by the relative temporal delays in the delay lines, and the attenuations of the VOAs represent learnable weights.

To demonstrate the efficacy of this approach, we train (in silicon) our photonic NCA to perform binary image classification of 28×28 grayscale images of sneakers and trousers in the Fashion-MNIST dataset [15]. To perform image classification, we concatenate the columns of the image to form a 1D array of pixels compatible with our

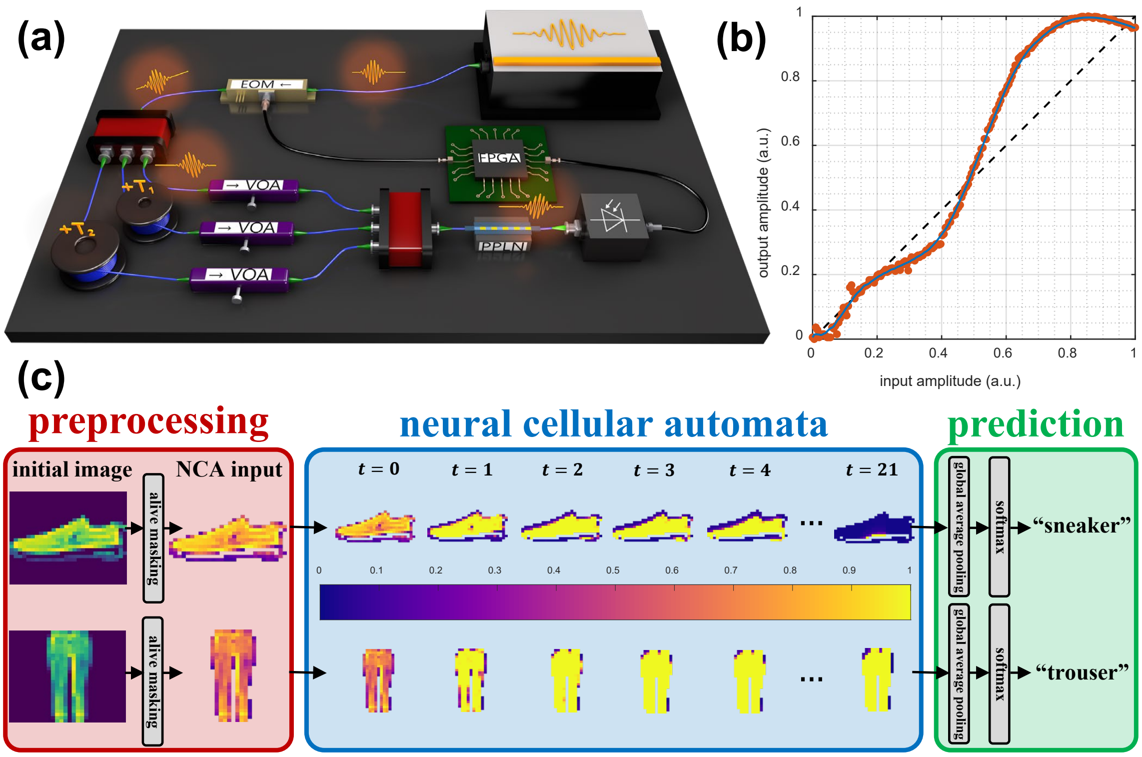

Fig. 1: (a) Schematic of experimental setup for neural CA using time multiplexing in photonics. (b) All-optical nonlinear activation function of PPLN. (c) Experimentally obtained time evolution of example images for sneaker and trouser classes.

temporal lattice in the experiment and set the appropriate time delays to achieve the desired 3-cell neighborhood in the 2D image. The initial cell states for the NCA are given by the corresponding pixel values. First, the image is pre-processed using *alive masking*, which designates any cell with initial pixel value < 0.1 as *dead* with a constant state of 0 that does not update but can still affect the state of neighboring cells. Next, the initial condition given by the *alive* cells pass through the NCA for 21 iterations until they converge to near-steady state. Finally, the classification is performed by taking the global average. A global average closer to 0 indicates the predicted class is *sneaker*, and conversely a global average closer to 1 indicates the predicted class is *trouser*. Therefore, the cells self-organize to form a collective agreement on the class through only local message-passing, inspired by the process of morphogenesis in biology [9], which is fundamentally different to how conventional MLPs and CNNs output classification labels.

Experimentally obtained examples of the time evolution for both sneaker and trouser in the NCA are shown in Fig. 1(c). Using this approach, we achieved a simulated test accuracy on 2000 images of 99.3% on the binary image classification task using only 3 learnable parameters. This is orders-of-magnitude fewer parameters required compared to MLPs and CNNs, whilst only suffering a slight decrease in accuracy. This NCA approach is general and can be applied to more complicated multi-class classification and generative tasks [9, 13] by using more complicated rules, e.g. encoded by a small, shallow, and multi-channel neural network instead of a single perceptron. Therefore, instead of storing an entire MLP or CNN with dense fully connected layers, the photonic hardware need only encode the local connections and relatively fewer number of parameters needed to encode the NCA update rule. In summary, in this work we experimentally demonstrate a scalable end-to-end PNN based on NCA with hardware much simpler than previous demonstrations of PNNS without sacrificing performance.

References

1. X. Guo *et al.*, Adv. Photonics Res. **2**, 2000212 (2021).
2. Y. Shen *et al.*, Nat. Photonics **11**, 441–446 (2017).
3. J. Feldmann *et al.*, Nature **589**, 52–58 (2021).
4. G. H. Li *et al.*, Nanophotonics (2022).
5. Y. LeCun *et al.*, Nature **521**, 436–444 (2015).
6. C. Szegedy *et al.*, arXiv preprint arXiv:1312.6199 (2013).
7. C. Guo *et al.*, in *International Conference on Machine Learning,* (PMLR, 2017), pp. 1321–1330.
8. T. Brown *et al.*, Adv. Neural Inf. Process. Syst. **33**, 1877–1901 (2020).
9. A. Mordvintsev *et al.*, Distill **5**, e23 (2020).
10. G. H. Li *et al.*, arXiv preprint arXiv:2211.01811 (2022).
11. S. Wolfram, Rev. Mod. Phys. **55**, 601 (1983).
12. M. Cook, Complex systems **15**, 1–40 (2004).
13. E. Randazzo *et al.*, Distill **5**, e00027–002 (2020).
14. C. Langrock *et al.*, Opt. Lett. **32**, 2263–2265 (2007).
15. H. Xiao *et al.*, arXiv preprint arXiv:1708.07747 (2017).