

Photonic convolutional neural networks using integrated diffractive optics

Jun Rong Ong ¹, Thomas Yong Long Ang ², Chin Chun Ooi ², Soon Thor Lim ², and Ching Eng Png ²

¹Institute of High Performance Computing

²Affiliation not available

October 30, 2023

Abstract

With recent rapid advances in photonic integrated circuits, it has been demonstrated that programmable photonic chips can be used to implement artificial neural networks. Convolutional neural networks (CNN) are a class of deep learning methods that have been highly successful in applications such as image classification and speech processing. We present an architecture to implement a photonic CNN using the Fourier transform property of integrated star couplers. We show, in computer simulation, high accuracy image classification using the MNIST dataset. We also model component imperfections in photonic CNN and show that the performance degradation can be recovered in a programmable chip. Our proposed architecture provides a large reduction in physical footprint compared to current implementations as it utilizes the natural advantages of optics and hence offers a scalable pathway towards integrated photonic deep learning processors.

Photonic convolutional neural networks using integrated diffractive optics

Jun Rong Ong, Chin Chun Ooi, Thomas Y. L. Ang, Soon Thor Lim, Ching Eng Png

Abstract—With recent rapid advances in photonic integrated circuits, it has been demonstrated that programmable photonic chips can be used to implement artificial neural networks. Convolutional neural networks (CNN) are a class of deep learning methods that have been highly successful in applications such as image classification and speech processing. We present an architecture to implement a photonic CNN using the Fourier transform property of integrated star couplers. We show, in computer simulation, high accuracy image classification using the MNIST dataset. We also model component imperfections in photonic CNN and show that the performance degradation can be recovered in a programmable chip. Our proposed architecture provides a large reduction in physical footprint compared to current implementations as it utilizes the natural advantages of optics and hence offers a scalable pathway towards integrated photonic deep learning processors.

Index Terms—Artificial neural networks, Neuromorphics, Photonic integrated circuits, Silicon photonics

I. INTRODUCTION

Deep learning methods like CNNs have received a huge amount of interest from the research community as well as the general public after it was shown to approach human level performance in image recognition tasks [1]. This breakthrough was due in part to the availability of fast graphical processing units (GPUs) that greatly accelerated the implementation of deep neural networks [2]. Recent efforts in developing hardware machine learning accelerators include massively parallel high-throughput devices [3], as well as neuromorphic computing architectures in which aspects of the design mimic principles present in biological neural networks [4], [5]. The search for alternative computing paradigms is also fueled by the impending end of Moore’s law, which is a result of the fundamental limits of transistor scaling, as well as related bottlenecks in power dissipation and interconnect bandwidth [6].

On the other hand, optical computing platforms potentially offer a number of attractive advantages such as parallelism through wavelength and temporal multiplexing [7], [8] as well as ideally non-dissipative interconnects [9]. Such potential advantages could be harnessed in high-performance computing systems, such as dedicated hardware accelerators for machine learning. Artificial neural

networks are particularly suited for optical implementation as they mainly rely on computing matrix multiplications, which can be performed with high speed and throughput with optics [10], [11]. In fact, the inherent advantages of optical neural networks were studied in detail decades ago using bulk optics [12]. Recent progress in photonic integrated circuits and programmable photonics has enabled the demonstration of a chip-scale optical neural network (ONN) using a mesh of interferometers and phase shifters [13]. The authors used the ability of such meshes to implement general unitary transformations, and together with a singular value decomposition [14], showed a ONN equivalent to a fully-connected multi-layer perceptron (MLPs). CNNs, in contrast to such fully-connected networks, take advantage of hierarchical patterns in the underlying data by having shared weights between network nodes (i.e. a convolution operation) and hence have a reduced scale of complexity and connectedness. The CNN architecture takes its inspiration from the animal visual cortex [1] and is in that sense a neuromorphic computing system.

Coherent optical information processing systems, for example optical correlators, rely heavily on the natural Fourier transform property of optics [15]. Conventionally, the correlation filters are hand designed by human experts. More recently, diffractive optics systems have demonstrated a high degree of success in image classification tasks [16], [17]. Such systems are constructed using sequential amplitude and phase masks, with each individual pixel in the masks being a trainable parameter in a deep learning optimization algorithm. In this way, the filter masks are generated automatically using machine learning. With the lenses arranged in a “4f” system, the optics performs convolutions in a similar fashion as in the convolution layers in a CNN. However, most proposals of optical CNNs have focused on a hybrid optical-electronic system configuration, with an optical convolution front-end in combination with an electronic implementation of the fully-connected layers [18]–[20].

Programmable integrated photonics has made significant advances, potentially eliminating the need for such a hybrid system [21]. As compared to bulk optics, integrated photonics is a scalable solution in terms of alignment stability and total network size. Additionally, rapid advances in performance of integrated silicon photonics devices could enable such photonic CNNs to be faster and more energy efficient than electronic implementations [22]. In fact, a recent proposal implements a photonic CNN by

J.R. Ong, C.C. Ooi, T.Y.L. Ang, S.T. Lim and C.E. Png are with the Institute of High Performance Computing, Agency for Science, Technology and Research (A*STAR), 138632, Singapore. (e-mail: ongj@ihpc.a-star.edu.sg; pngce@ihpc.a-star.edu.sg)




Fig. 1: Schematic of $N \times M$ star coupler. R is the radius of the confocal circles that make up the free-space propagation region. θ_n is the angle of the n th input waveguide, θ'_m is the angle of the m th output waveguide. w is the waveguide mode width parameter.

employing the patching technique to vectorize the input and kernel matrices [23]. However, the proposed architecture requires integrating long on-chip delay lines which requires overcoming some severe engineering challenges. In this work, we propose to use a star coupler (i.e. integrated diffractive optics) [15], which implements the discrete Fourier transform (DFT), to perform the convolution operation. Combined with phase and amplitude masks, we construct an integrated photonics CNN. We present the performance of the photonic CNN on various datasets, study the performance degradation with imperfections and also provide discussion on scalability and possible future directions.

II. PHOTONIC CNN

A typical CNN architecture consists of convolution layers, pooling layers, activation layers, a fully-connected layer and finally an output layer. In this section, we first show the details of the star coupler DFT. Subsequently, we describe how to physically implement the convolution, pooling and activation layers using photonic components. Lastly, we present a variety of photonic CNN architectures together with prediction results on standard machine learning datasets.

A. DFT using star couplers

The unitary DFT operation \mathcal{F} and inverse \mathcal{F}^{-1} for a discrete signal $x[n]$ of length N can be defined,

$$X[m] = \frac{1}{\sqrt{N}} \sum_{n=n_0}^{n_0+N-1} x[n] e^{-i2\pi \frac{nm}{N}} \quad (1)$$

$$x[n] = \frac{1}{\sqrt{N}} \sum_{m=m_0}^{m_0+N-1} X[m] e^{+i2\pi \frac{nm}{N}} \quad (2)$$

where n_0 and m_0 equals $-\frac{N}{2}$ if N is even and $-\frac{N-1}{2}$ if N is odd. Here, we use a star coupler to implement the DFT, as in [24], [25].

A star coupler is a $N \times M$ device, with N single-mode input waveguides and M single-mode output waveguides connected by a “free-space” propagation region like a slab waveguide [26]. The input and output waveguides are arranged along the circumference of two confocal circles of radius R . Under scalar diffraction theory and using the paraxial approximation, the coupling between an input waveguide at angle θ_n to an output waveguide at an angle θ'_m is [27]

$$\kappa(\theta_n, \theta'_m) = U(\theta'_m) \int \Phi(\theta' - \theta'_m) e^{-i\tilde{k}R(\theta' - \theta'_m) \sin \theta_n} R d\theta' \quad (3)$$

where

$$U(\theta'_m) = \frac{e^{i\tilde{k}R}}{\sqrt{i\tilde{\lambda}R}} \int \Phi(\theta - \theta_n) e^{-i\tilde{k}R \sin \theta \sin \theta'_m} R \cos \theta d\theta \quad (4)$$

with $\tilde{k} = \frac{2\pi}{\tilde{\lambda}} n_s$ and $\tilde{\lambda} = \frac{\lambda}{n_s}$, n_s being the slab effective index. Φ is the waveguide mode field, which we take to be a power normalized Gaussian, $\Phi(\theta) = \sqrt{\frac{2}{\pi w^2}} e^{-(R\theta/w)^2}$ and w is the width parameter of the waveguide. Since $w \ll R$, we can approximate $\theta' = \theta'_m$ and $\theta = \theta_n$ for the phase terms in Eq. 3 and 4.

Then, apart from a constant phase term,

$$\kappa(\theta_n, \theta'_m) \propto e^{-i\frac{2\pi}{N} \frac{NR}{\tilde{\lambda}} \sin \theta_n \sin \theta'_m}. \quad (5)$$

Comparing with Eq. 1, assuming $N \geq M$, to get correspondence with DFT

$$\theta_n = \sin^{-1} \left(n \sqrt{\frac{\tilde{\lambda}}{NR}} \right) \quad (6)$$

$$\theta'_m = \sin^{-1} \left(m \sqrt{\frac{\tilde{\lambda}}{NR}} \right). \quad (7)$$

Hence, by choosing the angular locations of the n th input and m th output waveguides by Eq. 6 and 7, the star coupler can implement a DFT.

In Fig. 2, we plot the phase and amplitude response of a 21×21 star coupler DFT, \mathcal{F}_{sc} , calculated using Eq. 3. Note that we have subtracted the constant phase term in the plot. The star coupler physical parameters are detailed in the figure captions and are chosen to ensure that the paraxial approximation is satisfied. Comparing with the ideal DFT, the greatest deviation in amplitude and phase response occurs at the waveguides furthest from the center. Adopting the fidelity measure as a distance metric [28],

$$F = \left| \frac{\text{Tr}(|\mathcal{F}_{sc}\rangle\langle\mathcal{F}|)}{N} \right|^2 \quad (8)$$




Fig. 2: (Top) Phase and amplitude response of 21×21 star coupler DFT, calculated using Eq. 3, with physical parameters: $\lambda = 1550$ nm, $n_s = 2.85$, $w = 500$ nm, $R = 340.9$ μ m. (Bottom) Ideal DFT phase response and relative phase error in radians of star coupler DFT.

gives $F = 0.997$ for the star coupler described above, where $\|\cdot\|$ denotes division by the Frobenius norm. Another useful summary metric is the overall transmission

$$T = \frac{\text{Tr}(\mathcal{F}_{sc}^\dagger \mathcal{F}_{sc})}{N} \quad (9)$$

with $T = 0.162$ for the star coupler above. There is a trade-off between T and F and in a later section we will study this in more detail.

We can compare the star coupler DFT with several other existing designs that implement the DFT with integrated photonics. Since the DFT is a unitary matrix, we can implement it directly using a mesh of Mach-Zehnder interferometers (MZIs) and phase shifters [29], [30]. Such a direct implementation of $N \times N$ DFT requires $\frac{N(N-1)}{2}$ MZIs, each consisting of twice that number of phase shifters and beam splitters. A more efficient design using the Cooley-Tukey FFT algorithm reduces the number of MZIs needed to $\frac{N \log_2(N)}{2}$ [31], [32]. However, the waveguide crossings needed grows as a triangular number $T_{\frac{N}{2}-1}$, which can introduce a significant insertion loss. Alternatively, 3D integration would be required to circumvent the need for crossings [33]. In comparison, the $N \times N$ star coupler DFT needs only a single “free-space” propagation region, which is a major reduction in complexity.

B. Convolution, pooling and activation layers

The convolution and pooling layers are linear and can be straightforwardly implemented optically. The convolution layer can be implemented optically in a “4f” system, by using a cascade of two optical DFT operations with a phase and amplitude filter mask in between. Hence the convolution layer is defined $\mathcal{C}^N = \mathcal{F}_{sc}^{NN} \cdot A^N \cdot \mathcal{F}_{sc}^{NN}$, with

the superscripts denoting the size of the matrices. Note that the second star coupler performs an inverse transformation back from the Fourier domain, except that the data is flipped vertically. Programmable phase and amplitude modulation can be applied to each individual waveguide using fast phase shifting mechanisms like thermo-optic and electro-optic effects [34]–[36]. Alternatively, if a fast response is not essential, then reconfigurability using post-fabrication trimming or phase-change materials are attractive alternatives that do not need additional energy and control to maintain the state [37]–[40].

The pooling layer can be implemented as a low-pass filter [41], by passing only the $M < N$ low-frequency components of the optical DFT, which correspond to the waveguides near to the central waveguide. Hence, the pooling layer is defined $\mathcal{P} = \mathcal{F}_{sc}^{MN}$. We may also combine convolution and pooling functions by transforming back from Fourier domain like so: $\mathcal{C}^{NM} = \mathcal{F}_{sc}^{MM} \cdot A^M \cdot \mathcal{F}_{sc}^{MN}$.

TABLE I

Photonic CNN physical implementation

Network Layer	Operation	Optics
Convolution	DFT	$N \times N$ star coupler
	Filter	Phase/Amp. mod.
	DFT	$N \times N$ star coupler
Pooling	DFT, Low-pass	$N \times M$ star coupler
Activation	modReLU	O-E-O

Activation functions are nonlinear functions that allow the neural network to learn complex mappings between inputs and outputs [42]. Nonlinear activation functions are regarded as one of the key reasons for the power of deep learning compared to classical machine learning methods. Adding nonlinearity into a photonic network substantially changes the functionality compared with previously demonstrated linear photonic circuits [21], [43], [44]. Several different kinds of optical nonlinearities have been proposed for implementation in optical neural networks, such as saturable absorption, optical bistability and two-photon absorption, to name a few [45]–[47]. However, all-optical nonlinearities are generally very weak and hence require high signal powers. As such, enhancing optical nonlinearities remains an area of active research. Recent advances in photonic integration have enabled demonstration of low-energy and high speed optical-electrical-optical or O-E-O devices, which act as pseudo-optical nonlinear devices [48]. Such O-E-O devices are reconfigurable to show a variety of output responses, including an approximation of the widely used ReLU function [49], [50].

Figure 3 shows a comparison of a typical CNN architecture schematic and an equivalent photonic CNN implementation. We implemented the photonic CNN using TensorFlow [51] with the data encoded as the amplitudes of the complex field u_0 at the input waveguides. The filter mask A is a diagonal matrix with complex-valued entries $a_n e^{i\phi_n}$, which are the trainable parameters for the convolution layers. This can be physically implemented as phase shifters and attenuators at each of the n waveguides. For the pointwise activation function G , we formulate it as

$\text{modReLU}(z_n, \{b_n, \varphi_n\} = 0) = \text{abs}(z_n)$, i.e. discarding the phase. See the Appendix for more details on modReLU . Hence, the complex field at the $(k+1)$ th convolution layer is related to the k th layer as

$$u_{k+1}^M = G(\mathcal{C}_k^{MN} \cdot u_k^N) = G(\mathcal{F}_k^{MM} \cdot A_k^M \cdot \mathcal{F}_k^{MN} \cdot u_k^N) \quad (10)$$

for a generic $N \times M$ convolution-pooling layer. Finally, we implemented the fully-connected layers in the same way as in Ref. [13]. Optically, it will be

$$u_{k+1}^M = G(W_k^{MN} \cdot u_k^N) = G(U_k^{MM} \cdot \Sigma_k^M \cdot (V_k^{NN})^\dagger \cdot u_k^N) \quad (11)$$

with W as the weight matrix decomposed as $W = U\Sigma V^\dagger$ using singular value decomposition (SVD). More details of the implementation and network training procedures are found in the Appendix.

C. Results on MNIST dataset

In Table II, we consider the performance of various different photonic ONNs on a standard machine learning task of digit recognition on handwritten images (MNIST) [52]. The 784 real-valued image pixels are fed into each of the 784 input waveguides using amplitude-only encoding. The PCNN-784 architecture is as follows: $\mathcal{C}^{784} \rightarrow \mathcal{C}^{392} \rightarrow \mathcal{C}^{196} \rightarrow W^{56} \rightarrow W^{10}$, where we omit the input size for clarity. We compare four different variations of PCNN-784 with modifications in *only* the convolution layers: amplitude and phase modulation (i.e. $A_k^M = \text{diag}(a_m e^{i\phi_m})$), amplitude-only modulation (i.e. $A_k^M = \text{diag}(a_m)$), phase-only modulation (i.e. $A_k^M = \text{diag}(e^{i\phi_m})$), and finally phase-only modulation with a linear activation function (i.e. $\text{modReLU}(z_m, b_m = 0)$). We find the best performance is obtained with amplitude and phase modulation, but only narrowly better than phase-only modulation. Amplitude-only modulation has the poorest performance. Below, we choose to focus on phase-only modulation since it achieves very good accuracy with less parameters, and it is also well adopted in existing literature [17], [18], [53].

The training (test) accuracy obtained is 99.2%(97.9%), which is below the state-of-the-art test accuracy of $\approx 99.6\%$. However, we expect the result to improve with advanced techniques like data augmentation, learning rate scheduling etc. We also tested the performance of the phase-only PCNN-784 on a more challenging dataset (F-MNIST) [54] and obtained an accuracy of 91.0%(88.6%) which is comparable to generic CNNs.

To test the potential advantage of the PCNN architecture as compared to existing ONN architectures, we also simulate the performance of two other kinds of networks: MLP-784 is a fully-connected network ($W^{16} \rightarrow W^{10}$) of the kind found in Ref. [13]; while D2NN-16 is a stack of 17 star couplers (\mathcal{F}_{sc}^{784}) sandwiching 16 activation and phase layers, similar to Ref. [17]. We intentionally chose the network architectures to have roughly the same number of trainable parameters for comparison. The results (Table II) show that the PCNN, which has a combination of

TABLE II

Comparison of different ONN architectures

ONN	Param.	Data	Acc. (%)
PCNN-784 (Amp. & Phase)	14280	MNIST	99.6 (98.2)
" (Amp.)	12908	MNIST	97.4 (95.3)
" (Phase)	12908	MNIST	99.2 (97.9)
" (Phase, Linear)	"	MNIST	99.0 (97.6)
" (Phase)	"	F-MNIST	91.0 (88.6)
MLP-784 [13]	12704	MNIST	95.9 (93.8)
D2NN-16 [17]	12544	MNIST	96.6 (95.6)
PCNN-256-32	4800	MNIST	97.7 (96.6)
PCNN-256-16	2592	MNIST	96.1 (95.6)
PCNN-112-32	2280	MNIST	93.9 (93.6)
PCNN-112-16	1224	MNIST	91.2 (91.1)

convolution and fully-connected layers, indeed has an additional benefit in this classification task.

Since the bulk of the trainable parameters in the PCNN come from the fully-connected layers, whereas the parameters in the convolution layers only grow linearly with the size of the input, we studied the trade-off between performance and complexity by being more aggressive with the pooling layers. The PCNN-x-y architecture is: $\mathcal{C}^x \rightarrow \mathcal{C}^{x/2} \rightarrow W^y \rightarrow W^{10}$, with phase-only modulation and linear activation in the convolution layers. We have introduced a pooling operation in the first layer to extract the low frequency components. A similar strategy of retaining only the low frequency features was adopted in Ref. [50], [55], in which the authors reported achieving a high accuracy of 98.9%(97.8%). There is some degradation of accuracy but potentially a great reduction in size and complexity of the network.

III. COMPONENT IMPERFECTIONS

Here, we study the effect on performance of two kinds of imperfections in the photonic CNN: first, the inherent imperfection of the star coupler DFT; second, imperfection in implementation of pre-trained ideal components parameters.

A. Imperfect DFT implementation using star couplers

The radius R of a $N \times N$ DFT star coupler is given by Eq. 6 as,

$$\sqrt{R} = \frac{1}{|\sin \theta_{n_0}|} \sqrt{\frac{|n_0|^2}{N}} \tilde{\lambda}. \quad (12)$$

Since N is determined by the size of the input, we have freedom to choose the angle of the waveguide furthest from the central line, θ_{n_0} , as long as the paraxial approximation is satisfied. In Fig. 4, we plot the fidelity F and transmission T of 784×784 star couplers with $5^\circ < \theta_{n_0} < 15^\circ$, i.e. within paraxial limits. We defined a N independent normalized radius $R_{norm} = \frac{\tilde{\lambda}}{|\sin \theta_{n_0}|^2}$, with $R \approx \frac{N}{4} R_{norm}$. As can be seen, a very good fidelity is obtained with small θ_{n_0} , but the transmission suffers and the star coupler radius is very large. By choosing a bigger θ_{n_0} , there is much better transmission and nearly order of magnitude reduction in radius, with some sacrifice of fidelity. The

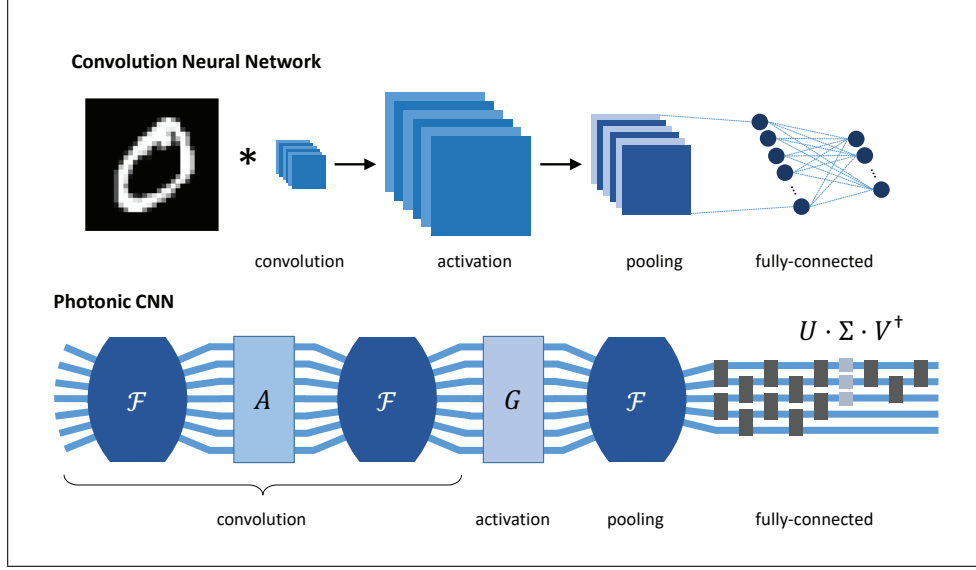


Fig. 3: Comparison between generic CNN architecture and corresponding photonic CNN implementation. The convolution is performed using a first star coupler DFT, then applying the filter mask in the Fourier domain and finally a second star coupler. Pooling is performed by passing only the low-frequency components through the output of the star coupler. Fully-connected layers are implemented as meshes of MZIs (dark gray boxes) and amplifiers/attenuators (light gray boxes).




Fig. 4: (Left) Trade-off between fidelity and transmission for star coupler DFT when varying normalized radius R_{norm} . (Right) Change in prediction accuracy versus input star coupler radius R for different PCNN architectures, using the MNIST and F-MNIST data.

reason for this trend is the following: choosing a small θ_{n_0} (large radius) concentrates the receiving waveguides near the axis and thus some of the optical power further away from the axis is not captured. On the other hand, choosing a large θ_{n_0} (small radius) angularly spreads out the receiving waveguides which are better able to cover the Gaussian shaped far-field envelope of the emitting waveguides.

To study the effect of the reduced fidelity, we simulated different PCNN with varying θ_{n_0} . In Fig. 4, we plot the

accuracy of these PCNN when trained on the MNIST and F-MNIST image recognition tasks. We observed almost no degradation of accuracy with increasing θ_{n_0} in the MNIST task and a small reduction of accuracy of about 0.6% in the F-MNIST task. This indicates that the reduced fidelity can be compensated by the network training and is an advantageous trade-off for the gains in transmission and reduction in star coupler footprint.

B. Non-idealities and fabrication imperfections

The current standard procedure to implement a optical neural network begins with training performed using a software simulation model of the system, followed by translation of the trained parameters to the optical device parameters and finally fabrication of the device. This method relies heavily on the accuracy of the mapping process from trained parameters to physical parameters. This includes, for example, non-idealities introduced by thermal cross-talk and the finite precision of electronic control circuits [13]. In addition, there will inevitably be fabrication imperfections of the photonic components (star couplers, waveguides, phase shifters, beam-splitters etc.) which break the correspondence between the trained software model and the hardware implementation [28], [56]. Such additional uncertainty introduced by physical implementation becomes non-negligible especially when scaling to a large number of components, hence it is important to evaluate its effects on the photonic CNN performance.

To study the effects of such imperfections, we evaluated the degradation of the MNIST classification accuracy of a pre-trained PCNN-784 (phase, linear) by introducing both




Fig. 5: Degradation in prediction accuracy with increasing noise. Additional phase noise has standard deviation of $2\pi\sigma$, while amplitude noise is $a_\delta = 1 - |\delta|$ with δ having standard deviation of σ . “Final” and “Full” indicates restoration of accuracy after re-training the weights of the output layer and the full network when phase noise is added. (Inset) Zoom-out of plot, showing the complete randomization of the network when noise is large.

amplitude and phase gaussian noise to the star coupler matrices \mathcal{F}_k and the complex-valued filter masks A_k . We avoided adding noise to the fully-connected layers as this has been studied in previous literature along with several methods to ameliorate its effects being suggested [28], [56]. The added phase noise $\Delta\phi$ is zero mean normally distributed with width $2\pi\sigma$, whereas the amplitude noise is modeled as an additional loss i.e. $a_\delta = 1 - |\delta|$, where δ is zero mean normally distributed with width σ . Hence, each element of the matrices \mathcal{F}_k and A_k is multiplied by a random complex factor $a_\delta \cdot e^{i\Delta\phi}$.

Previous studies have considered phase noise of width up to 0.02 rad, which is justified for low index contrast platforms [28], [56]. However, for high index contrast platforms like silicon-on-insulator, the phase errors resulting from imperfections could be up to 2 orders of magnitude greater [57] and hence we considered much larger phase errors.

Figure 5 shows a comparison of the resulting degradation when adding purely phase noise, purely amplitude noise and complex-valued noise. We plot the mean and standard deviation of the accuracies, with each point in the plot consisting of 20 random instances. We can see from the results that the performance begins to degrade when $\sigma > 0.02$. Additional phase noise is especially detrimental to the PCNN accuracy, as expected for such a coherent optical system, and it is imperative that mitigation strategies are in place for high index contrast platform like silicon photonics [57], [58]. Several such proposals and demonstrations of post-fabrication reconfiguration and optimization of programmable coherent optical meshes can

be found in the literature [13], [55], [59], [60]. As hoped, re-training PCNN-784 with added phase noise restores the network accuracy (see Fig. 5), except in the most noisy configurations (see Appendix for details on training). Drawing inspiration from substantial previous work on randomly weighted networks in the literature [61]–[64], we attempted to restore the performance of noisy PCNN-784 by training only the weights of the output layer. Although the restored accuracy is reduced from the ideal case, we see a substantial improvement from the noisy state, which suggests that full reconfigurability may not be necessary for photonic CNNs to function despite the presence of a large amount of noise.

IV. PHYSICAL FOOTPRINT

Recently, progress has been made in demonstrations of silicon photonic integrated circuits with a large number (~ 1000 s) of reconfigurable components [13], [21], [43], [59], [65]–[67] and it should be possible to implement photonic neural networks with a comparable number of trainable parameters. As the complexity of the ONN increases, the physical footprint of photonics components will be an important limiting factor on scalability. As mentioned previously, the star coupler DFT potentially provides substantial reduction in footprint required. For concreteness, let us consider a 256×256 DFT. A typical MZI will have physical size of $\sim 60 \mu\text{m} \times 100 \mu\text{m}$ [50]. Implementing the DFT using the Cooley-Tukey FFT algorithm requires 1024 MZIs, which will have a footprint of 6.14 mm^2 . In comparison, for a star coupler DFT of $R_{\text{norm}} = 10 \mu\text{m}$, the footprint would be $\sim 0.6 \text{ mm} \times 0.3 \text{ mm}$, which is 34 times smaller. This considerable reduction in footprint makes deep CNNs feasible to be implemented using photonic integrated circuits.

V. CONCLUSION

In conclusion, we have proposed and simulated a scalable architecture for photonic convolutional neural networks using the Fourier transform property of star couplers. We described in detail methods to use photonic components to implement various layers of a generic CNN, including convolution and pooling. We compared our proposed architecture to existing designs and found a boost in performance as well as a significant reduction in complexity and footprint. We also considered effects of component imperfections and noted that the photonic CNN is robust to small amounts of noise and in the case of very noisy networks, required minimal re-training to restore network accuracy.

Real implementations of photonic CNNs still requires important engineering work in practical areas such as latency, energy consumption etc. [50]. Fortunately, interest in the successes of deep learning has spurred significant efforts towards the realization of photonic neural networks, as evidenced by the numerous publications in recent years. We expect our proposed architecture to enable the implementation of scalable deep CNN on integrated photonic




Fig. 6: Loss and accuracy history over 80 epochs training of PCNN-784 (phase-only, linear) for MNIST training and test set.

platforms and further the efforts towards the goal of fully optical neuromorphic computing platforms.

APPENDIX

A. Network training

We implemented the photonic convolutional neural network (PCNN) using TensorFlow [51]. The PCNN consists of convolutional layers and fully-connected layers. Classification is done by taking the location of the maximum power at the output layer.

We used the stochastic gradient descent algorithm, Adam [68], for training. The loss function used for training is the cross-entropy with softmax function applied to the power at the output layer. The training batch size and the number of epochs was set to be 8 and 80, respectively. Figure 6 shows an example of loss and accuracy training history. The MNIST and Fashion-MNIST datasets consist of 28×28 images split into a training set of 60,000 examples and a test set of 10,000 examples.

B. Layer details

The PCNN convolutional layers are made up of star couplers and filter masks. The k -th layer $N \times M$ star coupler coupling matrices \mathcal{F}_k^{MN} are calculated using Eq. 3 and are fixed. Since \mathcal{F}_k^{MN} has no trainable parameters, it is unchanged during network training.

The k -th layer filter masks A_k^M have complex-valued coefficients $a_m e^{i\phi_m}$. We define the filter mask coefficients a_m and ϕ_m as [19]

$$a_m = \frac{|\alpha_m|}{\max_{1 \leq m \leq M} |\alpha_m|} \quad (13)$$

$$\phi_m = 2\pi\theta_m \quad (14)$$

with α_m and θ_m being the trainable parameters. The normalization of a_m ensures that it remains in the interval $[0,1]$, whereas ϕ_m does not need normalization due to the periodicity of the phase $e^{i\phi_m}$. For phase-only modulation, we fix all $a_m = 1$.

For the nonlinear activation functions, we considered the modReLU function acting on the complex number $z_m = a_m e^{i\varphi_m}$ [69]. As an example, for the complex field vector \mathbf{z} of the neural network layer, the activation on the m -th element

$$\text{modReLU}(z_m, b_m) = \text{ReLU}(a_m + b_m) \cdot e^{i\varphi_m} \quad (15)$$

where the bias vector \mathbf{b} is a trainable parameter. Table IIIa shows the accuracy results on the MNIST data for the phase-only PCNN-784 network when using different variations of modReLU. With b_0 , we assume a single shared trained value for the bias vector \mathbf{b} . When $\{\mathbf{b}, \boldsymbol{\varphi}\} = 0$, modReLU is equivalent to taking the $\text{abs}(z_m)$ as in the main text. Also, when $\mathbf{b} = 0$, modReLU is just a linear activation. Table IIIb shows the effect of zero mean normally distributed bias noise of width $\Delta\mathbf{b}$.

TABLE III

(a) Comparison of different nonlinear activations

Convolution	Fully-connected	Acc. (%)
modReLU(\mathbf{b})	modReLU(\mathbf{b})	98.8 (96.7)
modReLU($\mathbf{b} = 0.01$)	modReLU($\mathbf{b} = 0.01$)	97.9 (96.6)
modReLU($\mathbf{b} = b_0$)	modReLU($\mathbf{b} = b_0$)	98.7 (96.6)
modReLU($\mathbf{b} = b_0$)	modReLU($\{\mathbf{b}, \boldsymbol{\varphi}\} = 0$)	99.3 (97.7)
modReLU($\mathbf{b} = 0$)	modReLU($\mathbf{b} = 0$)	92.6 (92.4)

(b) Effect of bias noise $\Delta\mathbf{b}$ on accuracy of modReLU(\mathbf{b})

$\Delta\mathbf{b}$	Train Acc. (%)	Test Acc. (%)
0.005	98.3 \pm 0.1	96.6 \pm 0.1
0.01	97.1 \pm 0.2	95.7 \pm 0.2
0.02	92.7 \pm 0.9	91.9 \pm 0.9
0.05	63.7 \pm 3.4	63.8 \pm 3.5
0.1	36.6 \pm 3.1	37.2 \pm 3.0
0.2	23.0 \pm 3.6	23.5 \pm 3.6
0.5	16.9 \pm 3.7	17.3 \pm 4.0

The k -th fully-connected layers of the PCNN are described by a real-valued weight matrix W_k^{MN} , which can be decomposed by SVD into a product of two unitary matrices and a non-negative real diagonal matrix, i.e. $W = U\Sigma V^\dagger$. For the purposes of training the network, we take the tunable parameters as the real-valued elements of W , with the knowledge that it can be implemented optically through the decomposition [13]. In general, attenuators can be used to implement a scaled matrix $\Sigma' = \frac{\Sigma}{\beta}$, such that the singular values are ≤ 1 . In that case, a global optical amplification β is needed and the weight matrix is $W = \beta \cdot U\Sigma'V^\dagger$ [28]. In the main text, the choice of activation function is multiplicative, i.e. $\text{abs}(\beta z) = \beta \cdot \text{abs}(z)$. Hence, using a scaled weight matrix $W' = U\Sigma'V^\dagger$ would give the same prediction result as using W .

C. Re-training noisy networks

For the re-training of noisy networks, the training batch size and the number of epochs was set to be 8 and 10, respectively. Gaussian noise was added to the star coupler matrices \mathcal{F}_k and the complex-valued filter masks A_k . During re-training, only the A_k and W_k matrices are trainable, while the \mathcal{F}_k are fixed in their noisy state.

REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, p. 436, 2015.
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [3] N. Jouppi, C. Young, N. Patil, and D. Patterson, "Motivation for and evaluation of the first tensor processing unit," *IEEE Micro*, vol. 38, no. 3, pp. 10–19, May 2018.
- [4] S. B. Furber, F. Galluppi, S. Temple, and L. A. Plana, "The spinnaker project," *Proc. IEEE*, vol. 102, no. 5, pp. 652–665, 2014.
- [5] S. K. Essera, P. A. Merollaa, J. V. Arthura, A. S. Cassidya, R. Appuswamy, A. Andreopoulou, D. J. Berga, J. L. McKinstyria, T. Melanoa, D. R. Barcha *et al.*, "Convolutional networks for fast energy-efficient neuromorphic computing," *Proc. Nat. Acad. Sci. USA*, vol. 113, no. 41, pp. 11 441–11 446, 2016.
- [6] R. K. Cavin, P. Lugli, and V. V. Zhirnov, "Science and engineering beyond moore's law," *Proc. IEEE*, vol. 100, no. Special Centennial Issue, pp. 1720–1749, 2012.
- [7] A. N. Tait, T. F. De Lima, E. Zhou, A. X. Wu, M. A. Nahmias, B. J. Shastri, and P. R. Prucnal, "Neuromorphic photonic networks using silicon photonic weight banks," *Sci. Rep.*, vol. 7, no. 1, p. 7430, 2017.
- [8] D. Brunner, M. C. Soriano, C. R. Mirasso, and I. Fischer, "Parallel photonic information processing at gigabyte per second data rates using transient states," *Nat. Commun.*, vol. 4, p. 1364, 2013.
- [9] R. G. Beausoleil, J. Ahn, N. Binkert, A. Davis, D. Fattal, M. Fiorentino, N. P. Jouppi, M. McLaren, C. Santori, R. S. Schreiber *et al.*, "A nanophotonic interconnect for high-performance many-core computation," in *2008 16th IEEE Symposium on High Performance Interconnects*. IEEE, 2008, pp. 182–189.
- [10] J. W. Goodman, A. Dias, and L. Woody, "Fully parallel, high-speed incoherent optical method for performing discrete fourier transforms," *Opt. Lett.*, vol. 2, no. 1, pp. 1–3, 1978.
- [11] L. Yang, R. Ji, L. Zhang, J. Ding, and Q. Xu, "On-chip cmos-compatible optical signal processor," *Opt. Express*, vol. 20, no. 12, pp. 13 560–13 565, 2012.
- [12] H. J. Caulfield, J. Kinser, and S. K. Rogers, "Optical neural networks," *Proc. IEEE*, vol. 77, no. 10, pp. 1573–1583, 1989.
- [13] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund *et al.*, "Deep learning with coherent nanophotonic circuits," *Nat. Photonics*, vol. 11, no. 7, p. 441, 2017.
- [14] D. A. Miller, "Self-configuring universal linear optical component," *Photonics Research*, vol. 1, no. 1, pp. 1–15, 2013.
- [15] J. W. Goodman, *Introduction to Fourier optics*. Roberts and Company Publishers, 2005.
- [16] X. Lin, Y. Rivenson, N. T. Yardimci, M. Vel, Y. Luo, M. Jarrahi, and A. Ozcan, "All-optical machine learning using diffractive deep neural networks," *Science*, vol. 361, no. 6406, pp. 1004–1008, 2018.
- [17] T. Yan, J. Wu, T. Zhou, H. Xie, F. Xu, J. Fan, L. Fang, X. Lin, and Q. Dai, "Fourier-space diffractive deep neural network," *Phys. Rev. Lett.*, vol. 123, no. 2, p. 023901, 2019.
- [18] J. Chang, V. Sitzmann, X. Dun, W. Heidrich, and G. Wetzstein, "Hybrid optical-electronic convolutional neural networks with optimized diffractive optics for image classification," *Sci. Rep.*, vol. 8, no. 1, p. 12324, 2018.
- [19] D. Mengü, Y. Luo, Y. Rivenson, and A. Ozcan, "Analysis of diffractive optical neural networks and their integration with electronic neural networks," *IEEE J. Sel. Topics Quantum Electron.*, 2019.
- [20] S. Colburn, Y. Chu, E. Shilzerman, and A. Majumdar, "Optical frontend for a convolutional neural network," *Appl. Opt.*, vol. 58, no. 12, pp. 3179–3186, 2019.
- [21] N. C. Harris, J. Carolan, D. Bunandar, M. Prabhu, M. Hochberg, T. Baehr-Jones, M. L. Fanto, A. M. Smith, C. C. Tison, P. M. Alsing, and D. Englund, "Linear programmable nanophotonic processors," *Optica*, vol. 5, no. 12, pp. 1623–1631, Dec 2018.
- [22] A. H. Atabaki, S. Moazeni, F. Pavanella, H. Gevorgyan, J. Notaros, L. Alloatti, M. T. Wade, C. Sun, S. A. Kruger, H. Meng *et al.*, "Integrating photonics with silicon nanoelectronics for the next generation of systems on a chip," *Nature*, vol. 556, no. 7701, p. 349, 2018.
- [23] H. Bagherian, S. Skirlo, Y. Shen, H. Meng, V. Ceperic, and M. Soljacic, "On-chip optical convolutional neural networks," *arXiv preprint arXiv:1808.03303*, 2018.
- [24] K. Takiguchi, T. Kitoh, A. Mori, M. Oguma, and H. Takahashi, "Optical orthogonal frequency division multiplexing demultiplexer using slab star coupler-based optical discrete fourier transform circuit," *Opt. Lett.*, vol. 36, no. 7, pp. 1140–1142, 2011.
- [25] G. Cincotti, "What else can an awg do?" *Opt. Express*, vol. 20, no. 26, pp. B288–B298, 2012.
- [26] C. Dragone, "Efficient $n \times n$ star couplers using fourier optics," *Journal of Lightwave Technology*, vol. 7, no. 3, pp. 479–489, 1989.
- [27] A. Klekamp and R. Munzner, "Calculation of imaging errors of awg," *Journal of lightwave technology*, vol. 21, no. 9, pp. 1978–1986, 2003.
- [28] M. Y.-S. Fang, S. Manipatruni, C. Wierzynski, A. Khosrowshahi, and M. R. DeWeese, "Design of optical neural networks with component imprecisions," *Opt. Express*, vol. 27, no. 10, pp. 14 009–14 029, 2019.
- [29] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani, "Experimental realization of any discrete unitary operator," *Phys. Rev. Lett.*, vol. 73, no. 1, p. 58, 1994.
- [30] W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walmsley, "Optimal design for universal multiport interferometers," *Optica*, vol. 3, no. 12, pp. 1460–1465, 2016.
- [31] R. Barak and Y. Ben-Aryeh, "Quantum fast fourier transform and quantum computation by linear optics," *JOSA B*, vol. 24, no. 2, pp. 231–240, 2007.
- [32] J. K. George, H. Nejadriahi, and V. J. Sorger, "Towards on-chip optical fts for convolutional neural networks," in *2017 IEEE International Conference on Rebooting Computing (ICRC)*. IEEE, 2017, pp. 1–4.
- [33] A. Crespi, R. Osellame, R. Ramponi, M. Bentivegna, F. Flamini, N. Spagnolo, N. Viggianiello, L. Innocenti, P. Mataloni, and F. Sciarrino, "Suppression law of quantum states in a 3d photonic fast fourier transform chip," *Nat. Commun.*, vol. 7, p. 10469, 2016.
- [34] P. Orlandi, F. Morichetti, M. J. Strain, M. Sorel, A. Melloni, and P. Bassi, "Tunable silicon photonics directional coupler driven by a transverse temperature gradient," *Opt. Lett.*, vol. 38, no. 6, pp. 863–865, 2013.
- [35] M. R. Watts, J. Sun, C. DeRose, D. C. Trotter, R. W. Young, and G. N. Nielson, "Adiabatic thermo-optic mach-zehnder switch," *Opt. Lett.*, vol. 38, no. 5, pp. 733–735, 2013.
- [36] E. Timurdogan, C. V. Poulton, M. Byrd, and M. Watts, "Electric field-induced second-order nonlinear optical effects in silicon waveguides," *Nat. Photonics*, vol. 11, no. 3, p. 200, 2017.
- [37] Y. Shen, I. B. Divliansky, D. N. Basov, and S. Mookherjee, "Electric-field-driven nano-oxidation trimming of silicon microrings and interferometers," *Opt. Lett.*, vol. 36, no. 14, pp. 2668–2670, 2011.
- [38] R. Topley, L. O'Faolain, D. Thomson, F. Gardes, G. Mashanovich, and G. Reed, "Planar surface implanted diffractive grating couplers in soi," *Opt. Express*, vol. 22, no. 1, pp. 1077–1084, 2014.
- [39] K. J. Miller, R. F. Haglund, and S. M. Weiss, "Optical phase change materials in integrated silicon photonic devices," *Optical Materials Express*, vol. 8, no. 8, pp. 2415–2429, 2018.
- [40] J. Feldmann, N. Youngblood, C. Wright, H. Bhaskaran, and W. Pernice, "All-optical spiking neurosynaptic networks with self-learning capabilities," *Nature*, vol. 569, no. 7755, p. 208, 2019.
- [41] O. Rippel, J. Snoek, and R. P. Adams, "Spectral representations for convolutional neural networks," in *Advances in neural information processing systems*, 2015, pp. 2449–2457.
- [42] K. Hornik, M. Stinchcombe, and H. White, "Multilayer feedforward networks are universal approximators," *Neural networks*, vol. 2, no. 5, pp. 359–366, 1989.
- [43] J. Carolan, C. Harrold, C. Sparrow, E. Martín-López, N. J. Russell, J. W. Silverstone, P. J. Shadbolt, N. Matsuda, M. Oguma, M. Itoh *et al.*, "Universal linear optics," *Science*, vol. 349, no. 6249, pp. 711–716, 2015.

- [44] A. Ribeiro, A. Ruocco, L. Vanacker, and W. Bogaerts, "Demonstration of a 4×4 -port universal linear circuit," *Optica*, vol. 3, no. 12, pp. 1348–1357, 2016.
- [45] K. Nozaki, T. Tanabe, A. Shinya, S. Matsuo, T. Sato, H. Taniyama, and M. Notomi, "Sub-femtojoule all-optical switching using a photonic-crystal nanocavity," *Nat. Photonics*, vol. 4, no. 7, p. 477, 2010.
- [46] Z. Cheng, H. K. Tsang, X. Wang, K. Xu, and J.-B. Xu, "In-plane optical absorption and free carrier absorption in graphene-on-silicon waveguides," *IEEE J. Sel. Topics Quantum Electron.*, vol. 20, no. 1, pp. 43–48, 2013.
- [47] Y. Jiang, P. T. DeVore, and B. Jalali, "Analog optical computing primitives in silicon photonics," *Opt. Lett.*, vol. 41, no. 6, pp. 1273–1276, 2016.
- [48] K. Nozaki, S. Matsuo, T. Fujii, K. Takeda, A. Shinya, E. Kuramochi, and M. Notomi, "Femtofarad optoelectronic integration demonstrating energy-saving signal conversion and nonlinear functions," *Nat. Photonics*, vol. 13, pp. 454–459, 2019.
- [49] A. N. Tait, T. F. de Lima, M. A. Nahmias, H. B. Miller, H.-T. Peng, B. J. Shastri, and P. R. Prucnal, "Silicon photonic modulator neuron," *Phys. Rev. Appl.*, vol. 11, no. 6, p. 064043, 2019.
- [50] I. A. D. Williamson, T. W. Hughes, M. Minkov, B. Bartlett, S. Pai, and S. Fan, "Reprogrammable electro-optic nonlinear activation functions for optical neural networks," *IEEE J. Sel. Topics Quantum Electron.*, vol. 26, no. 1, pp. 1–12, 2019.
- [51] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin *et al.*, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," *arXiv preprint arXiv:1603.04467*, 2016.
- [52] "The mnist database of handwritten digits," <http://yann.lecun.com/exdb/mnist/>, accessed: 3-September-2019.
- [53] J. L. Horner and P. D. Gianino, "Phase-only matched filtering," *Appl. Opt.*, vol. 23, no. 6, pp. 812–816, 1984.
- [54] H. Xiao, K. Rasul, and R. Vollgraf. (2017) Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms.
- [55] S. Pai, I. A. Williamson, T. W. Hughes, M. Minkov, O. Solgaard, S. Fan, and D. A. Miller, "Parallel fault-tolerant programming of an arbitrary feedforward photonic network," *arXiv preprint arXiv:1909.06179*, 2019.
- [56] F. Flamini, N. Spagnolo, N. Viggianiello, A. Crespi, R. Osellame, and F. Sciarrino, "Benchmarking integrated linear-optical architectures for quantum information processing," *Sci. Rep.*, vol. 7, no. 1, p. 15133, 2017.
- [57] M. Gehl, D. Trotter, A. Starbuck, A. Pomerene, A. L. Lentine, and C. DeRose, "Active phase correction of high resolution silicon photonic arrayed waveguide gratings," *Opt. Express*, vol. 25, no. 6, pp. 6320–6334, 2017.
- [58] Y. Yang, Y. Ma, H. Guan, Y. Liu, S. Danziger, S. Ocheltree, K. Bergman, T. Baehr-Jones, and M. Hochberg, "Phase coherence length in silicon photonic platform," *Opt. Express*, vol. 23, no. 13, pp. 16890–16902, 2015.
- [59] A. Annoni, E. Guglielmi, M. Carminati, G. Ferrari, M. Sampietro, D. A. Miller, A. Melloni, and F. Morichetti, "Unscrambling light—automatically undoing strong mixing between modes," *Light: Science Applications*, vol. 6, p. e17110, 2017.
- [60] T. W. Hughes, M. Minkov, Y. Shi, and S. Fan, "Training of photonic neural networks through in situ backpropagation and gradient measurement," *Optica*, vol. 5, no. 7, pp. 864–871, 2018.
- [61] W. F. Schmidt, M. A. Kraaijveld, R. P. Duin *et al.*, "Feed forward neural networks with random weights," in *International Conference on Pattern Recognition*. IEEE COMPUTER SOCIETY PRESS, 1992, pp. 1–1.
- [62] A. M. Saxe, P. W. Koh, Z. Chen, M. Bhand, B. Suresh, and A. Y. Ng, "On random weights and unsupervised feature learning," in *ICML*, vol. 2, no. 3, 2011, p. 6.
- [63] G.-B. Huang, "What are extreme learning machines? filling the gap between frank rosenblatt's dream and john von neumann's puzzle," *Cognitive Computation*, vol. 7, no. 3, pp. 263–278, 2015.
- [64] A. Saade, F. Caltagirone, I. Carron, L. Daudet, A. Drémeau, S. Gigan, and F. Krzakala, "Random projections through multiple optical scattering: Approximating kernels at the speed of light," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 6215–6219.
- [65] J. Sun, E. Timurdogan, A. Yaacobi, Z. Su, E. S. Hosseini, D. B. Cole, and M. R. Watts, "Large-scale silicon photonic circuits for optical phased arrays," *IEEE J. Sel. Topics Quantum Electron.*, vol. 20, no. 4, pp. 264–278, 2013.
- [66] S. Han, T. J. Seok, N. Quack, B.-W. Yoo, and M. C. Wu, "Large-scale silicon photonic switches with movable directional couplers," *Optica*, vol. 2, no. 4, pp. 370–375, 2015.
- [67] D. Pérez, I. Gasulla, L. Crudgington, D. J. Thomson, A. Z. Khokhar, K. Li, W. Cao, G. Z. Mashanovich, and J. Capmany, "Multipurpose silicon photonics signal processor core," *Nat. Commun.*, vol. 8, no. 1, p. 636, 2017.
- [68] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [69] M. Arjovsky, A. Shah, and Y. Bengio, "Unitary evolution recurrent neural networks," in *International Conference on Machine Learning*, 2016, pp. 1120–1128.