

Training a Neural Network with Exciton-Polariton Optical Nonlinearity

A. Opala^{1,*}, R. Panico^{2,3}, V. Ardizzone³, B. Piętka⁴, J. Szczytko⁴, D. Sanvitto^{3,5},
M. Matuszewski¹ and D. Ballarini³


¹*Institute of Physics, Polish Academy of Sciences, Al. Lotników 32/46, Warsaw PL-02-668, Poland*

²*Dipartimento di Matematica e Fisica E. De Giorgi, Università del Salento, Campus Ecotekne, via Monteroni, Lecce 73100, Italy*

³*CNR NANOTEC-Institute of Nanotechnology, Via Monteroni, Lecce 73100, Italy*

⁴*Institute of Experimental Physics, Faculty of Physics, University of Warsaw, ul. Pasteura 5, Warsaw PL-02-093, Poland*

⁵*INFN, Sez. Lecce, Lecce 73100, Italy*

 (Received 4 February 2022; revised 25 April 2022; accepted 14 June 2022; published 10 August 2022)

In contrast to software simulations of neural networks, hardware implementations have limited or no tunability. While such networks promise great improvements in terms of speed and energy efficiency, their performance is limited by the difficulty of applying efficient training. We propose and realize experimentally an optical system where highly efficient backpropagation training can be applied through an array of highly nonlinear, nontunable nodes. The system includes exciton-polariton nodes realizing nonlinear activation functions. We demonstrate high classification accuracy in the MNIST handwritten digit benchmark in a single-hidden-layer system.

DOI: [10.1103/PhysRevApplied.18.024028](https://doi.org/10.1103/PhysRevApplied.18.024028)

I. INTRODUCTION

Over the course of the last decade, artificial intelligence and artificial neural networks have become extremely valuable tools in industry, research, and everyday life. Deep neural networks excel in many tasks, including image and speech recognition, language processing, and autonomous driving [1]. A substantial part of the success can be attributed to the development of effective training algorithms, and the widely applied backpropagation method [2] in particular. While the effectiveness of existing solutions is unquestionable, it is commonly believed that further progress, in particular in edge computing applications, can only be sustained if software simulations of neural networks are replaced by systems where neural structure of the network is implemented in hardware [3–5]. This is dictated by the necessity of developing systems characterized by high speed and high energy efficiency, which is difficult to achieve in the von Neumann computer architecture, where huge amounts of data are transmitted back and forth between memory and computing units.

In recent years, hardware neural networks have been realized in many systems, including complementary metal-oxide-semiconductor electronics, memristors, and photonic systems [3,5–18]. In particular, recent realizations using exciton-polaritons in optical microcavities have

achieved state-of-the-art accuracy in the Modified National Institute of Standards and Technology (MNIST) handwritten digit recognition benchmark [19–21]. Exciton-polaritons are composite quasiparticles that result from strong quantum coupling of semiconductor excitons and cavity photons [22,23]. They are characterized by efficient transport via the photonic component and strong inter-particle interactions due to the matter component. These properties make them promising candidates for future applications in efficient nonlinear information processing [19,24–32]. Exciton-polariton systems can potentially provide orders-of-magnitude improvements in speed and energy efficiency as compared to electronics and other optical systems [33].

A precise, differentiable model of a physical system and its tunability is a prerequisite for the implementation of efficient training algorithms such as backpropagation. In the absence of such a model, one often designs the system according to the reservoir computing paradigm [34–36], where the majority of synaptic weights are unchanged during the training phase [15,37–41]. On the other hand, a large number of reservoir nodes are usually required to achieve high levels of accuracy, and even in the case of very large networks it cannot match the accuracy of backpropagation-based networks.

In recent years, backpropagation has been implemented in several optical systems [18,42–44]. However, experiments have mostly been limited to linear systems where

*opala@ifpan.edu.pl

light propagated in free space or in a linear medium. In fact, it is well known that propagation of light is well suited for linear operations such as vector-matrix multiplication [45–47]. On the other hand, nonlinearity is a key component of the majority of efficient neural networks. The implementation of nonlinear physical nodes requires a system characterized by a suitable nonlinear response, a precise and differentiable model, reproducibility, stability, and a low level of noise. All these requirements are difficult to fulfill and their absence strongly influences the accuracy of the predictions of the network. It was recently demonstrated that physics-aware training can lead to improved accuracy [43], but the optimal physical platform is still not known.

Here, we use the backpropagation algorithm to teach optical neural networks in which nonlinear hardware nodes are nontunable. We demonstrate a system that includes exciton-polariton nodes exhibiting a strong nonlinear input-output dependence that can be measured precisely. We show that such a precise characterization of the nodes can be used to perform efficient training. We physically separate the tunable linear weights from the nontunable nonlinear nodes, whose task is to apply a nonlinear activation function. The uncontrollable and static activation functions of each of the nodes are determined experimentally before the training phase, which allows the application of the backpropagation algorithm offline.

In our proof-of-principle demonstration, we realize a single-hidden-layer feedforward neural network using optically excited exciton-polariton nodes, where both input and output weights are applied electronically. Despite the experimental imperfections, we achieve the MNIST inference accuracy of 96%, close to that of a software simulation carried out in the Tensorflow package [48]. This result is similar to that obtained with physics-aware training method using second harmonic optical nonlinearity in a five-layer network including more than 1000 physical nodes [43]. Our polariton network consists of only a few tens of nodes and training does not require hardware gradient estimation, which drastically reduces training time. We emphasize that linear weighting in our network could in principle be realized all-optically, as has been demonstrated previously [42,45–47,49–53]. Our work opens the way to more complex realizations, including deep and recurrent networks in systems with limited hardware tunability [3].

II. RESULTS

Implementation of machine learning typically consists of two separate stages. At the training stage, the system is taught to classify or predict using data from the training dataset. At this stage the synaptic connection weights are tuned, with the aim of increasing the accuracy of predictions. The second stage, called the inference (or testing)

stage, begins after all training samples have been processed. At this stage, synaptic weights are no longer tuned, and the system is not learning any more. The system is processing a testing dataset that consists of samples that it has not seen before. The accuracy of predictions at the inference stage is the most important benchmark of the network.

While training is usually a time-consuming and demanding process, once taught, the system is able to make predictions for an arbitrary number of samples at the inference stage. In many practical applications, it is the inference stage that requires the largest amount of resources, and may take an indefinitely long time. For example, language processing models can be used to process an arbitrarily large number of sentences without any need for readjustment after initial training. Therefore, from a practical point of view, it is valid to seek methods that improve the efficiency of inference, even when not increasing the efficiency of training at the same time [54].

In our approach we focused on the efficient hardware implementation of inference. To this end, we add an additional initial stage before training (the measurement of the response of each physical node), allowing construction of an accurate software model describing each of the hardware neurons. This allows us to perform training in software, while inference is realized in hardware. The entire process is schematically depicted in Fig. 1. A standard artificial neuron realizes two functions: (i) tunable synaptic weighting of inputs; and (ii) a nontunable activation function. We separate physically function (i) from function (ii), the latter being realized by a set of nontunable nonlinear nodes. In our case, these nodes are optically excited exciton-polariton modes of an optical microcavity. The knowledge of node response, measured at the initial stage as depicted in Fig. 1(a), allows us to perform the training stage entirely in software, using the backpropagation method as shown in Fig. 1(b). The resulting synaptic weights are implemented physically only at the inference stage, as in Fig. 1(c). We choose a simple single-hidden-layer feedforward neural network model to facilitate the backpropagation training procedure. Importantly, in contrast to reservoir computing method, we adjust *all* synaptic weights of the network, including those in the input layer. We will demonstrate that this allows us to significantly improve the accuracy of predictions.

Figure 1(a) presents the scheme of the experimental setup. A phase-only spatial light modulator (SLM) is used to modulate the intensity profile of a laser beam into an array of bright spots with individually tunable intensity. These spots, distributed in a 3×3 square lattice, are imaged on the surface of a semiconductor microcavity at resonance with the exciton-polariton energy. Due to the significant exciton-polariton nonlinearities, the light intensity transmitted through the microcavity follows a sigmoid-like behavior as a function of the input intensity

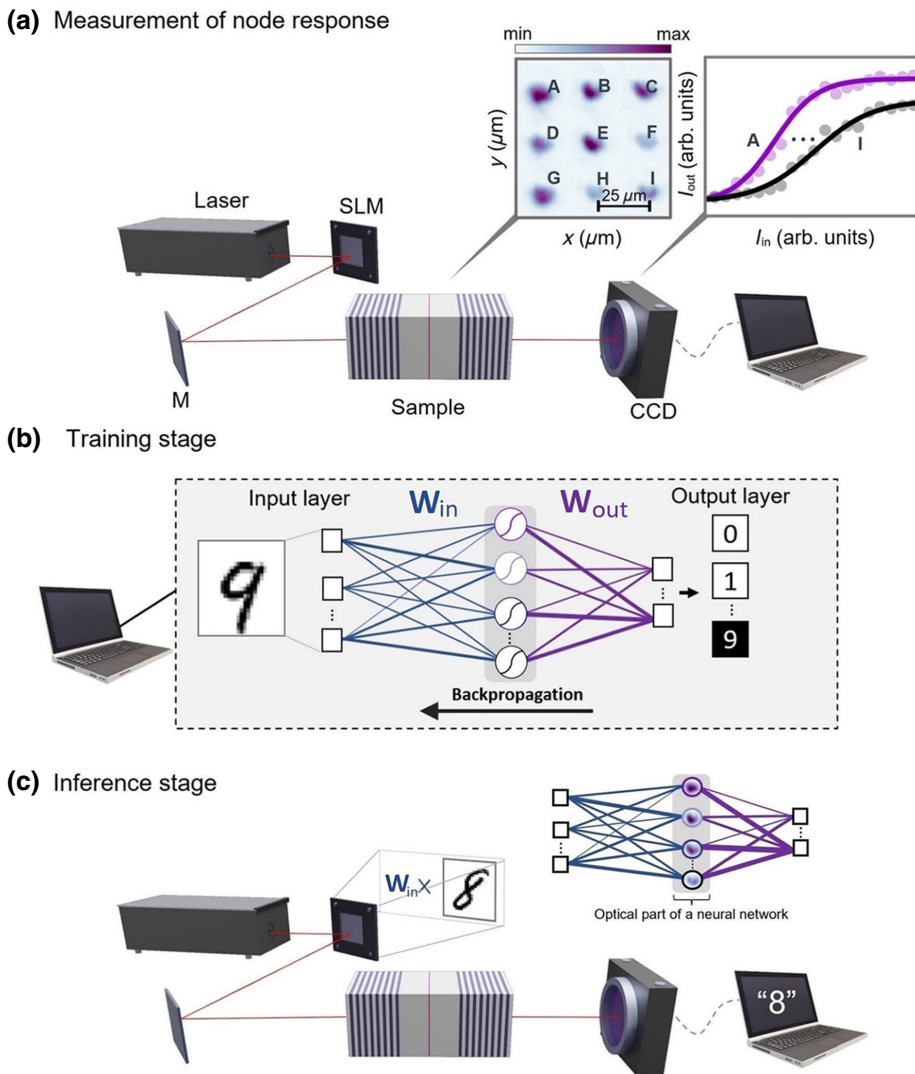


FIG. 1. Schematic of the experiment. (a) In the initial stage, we measure the input-output dependence of light transmission through the microcavity. The spatial pattern of input light creates a 3×3 array of uncoupled exciton-polariton nodes characterized by a sigmoid-like response. (b) The training stage is performed in software, where the previously measured node response functions are reproduced to perform backpropagation in a single-hidden-layer feed-forward network. Training results in optimal input weights W_{in} and output weights W_{out} . (c) At the inference stage, the obtained weights are used to classify handwritten digits using the nonlinear response of polariton nodes.

[20], as shown in the inset of Fig. 1(a). The light intensity of each polariton node, labeled (A, B, C, ..., I) in the figures, is measured by a charge-coupled device (CCD) camera and collected on a computer. Each node is spatially separated from the others to ensure that the output intensity of a polariton node depends only on its input intensity, independently of the input intensity of the other polariton nodes. This configuration is chosen for simplicity, as it facilitates the training procedure. However, our method is not limited to the case of isolated nodes. In the case of interconnected nodes, the system could be more accurate, even if requiring longer times and more complex methods to perform training.

The training stage, depicted in Fig. 1(b), is realized entirely in software. We use the Tensorflow package to simulate a single-layer feedforward neural network shown schematically in the figure. We consider the MNIST handwritten digit dataset, which contains 60 000 samples in the training set and 10 000 samples in the testing set, each sample being a 28×28 grayscale image of a digit between

0 and 9 and the corresponding label [55]. The number of nodes in the hidden layer of the network is equal to $9M$, where M is the number of experimental shots that will be used to process a single digit at the inference stage. The activation function of each hidden node is given by the node response measured in the initial stage. In the output layer, the softmax function is used to choose the class corresponding to one of the ten digits. We use the backpropagation algorithm to teach the network, which provides optimal values of synaptic weights in both in the hidden layer (W_{in}) and output layer (W_{out}) connections.

At the inference stage, shown in Fig. 1(c), we use the spatial light modulator to encode input data from the testing set. Additionally, data is multiplied by input synaptic weights, which in our case is realized in software. The light intensity incident on each microcavity node is the sum of inputs multiplied by the corresponding input synaptic weights. The role of the microcavity is to apply the nonlinear activation function via exciton-polariton interactions. The transmitted light corresponds therefore to the neuron

outputs. Further, the intensity measured on a CCD camera is collected, multiplied by output synaptic weights in software, and a softmax function is used to determine the predicted class.

In Fig. 2(a), we show the measured response of each polariton node, together with the corresponding analytical fits. The analytical functions are necessary to apply the backpropagation training method, which requires knowledge of the derivative of an activation function in order to calculate the update of input synaptic weights w_{ij}^{in} according to the formula

$$\Delta w_{ij}^{\text{in}} = -\eta \sum_{k=1}^M (y_k - d_k) \frac{dy_k}{dv_i} \frac{dv_i}{dw_{ij}^{\text{in}}} x_j, \quad (1)$$

where η is the learning rate, $\mathbf{x} = [x_1, x_2, x_3, \dots, x_N]^T$ and $\mathbf{y} = [y_1, y_2, y_3, \dots, y_M]^T$ are the input and output vectors respectively, and $\mathbf{d} = [d_1, d_2, d_3, \dots, d_M]^T$ is the target output vector. Activations of neurons in the hidden layer are given by the equation $v_i = \varphi_i(\sum_{j=0}^N w_{ij}^{\text{in}} x_j)$, where $\varphi_i(I_i^{\text{in}})$ is an activation function and $N = 28^2$ is the number of inputs. We approximate the polariton response by sigmoid analytical activation functions

$$\varphi_i = \frac{a_i}{1 + \exp(-(I_i^{\text{in}} - t_i)/c_i)} + d_i, \quad (2)$$

where a_i , t_i , c_i and d_i , $i = \{1, 2, 3, \dots, 9M\}$, are parameters obtained by fitting to the experimental data. Overall, we fit nine sigmoid functions to the response of nine polariton nodes in the (3×3) network. Figure 2(b) shows the results of the experimentally obtained accuracy at the inference stage, that is, predictions for 500 digits from the testing set. For comparison, in Fig. 2(c) we show the accuracy obtained in a Tensorflow software simulation of the same network and with the same dataset. The accuracy achieved in the experiment of 96.2% is comparable to the result obtained with a binarized network in [21] and higher than that obtained in a reservoir computing approach with a similar number of nodes [20]. In comparison to the binarized network [21], the setup presented here is much simpler, the number of nodes is much lower (tens instead of tens of thousands), and there is no need to create optical binary gates first. Consequently it has the potential to achieve faster data processing and is more scalable.

We emphasize that this excellent result has been obtained with a relatively small network, including only 90 nodes in the hidden layer ($M = 10$). This results from the use of backpropagation, which adjusts synaptic weights in all the network layers. Achieving a similar accuracy with a reservoir computing network requires a much larger number of nodes in the hidden layer [19]. To investigate the advantage of our approach in more detail, in Fig. 3(a) we compare the accuracy of our network, simulated in software, to the accuracy of an “extreme learning machine”

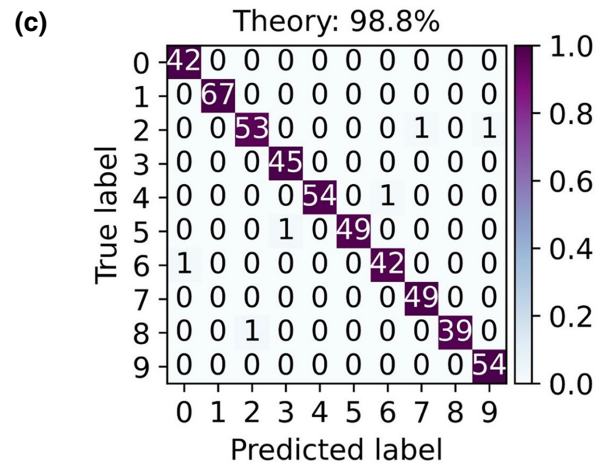




FIG. 2. (a) The measured response of polariton nodes (circles) together with the corresponding analytical fits (solid lines). (b) Confusion matrix resulting from the experimental inference in the case of the testing set. Numbers and color scale indicate the number of samples for a given pair of true and predicted labels. The diagonal corresponds to correct predictions. (c) Confusion matrix for an identical network and identical dataset simulated in the Tensorflow package.

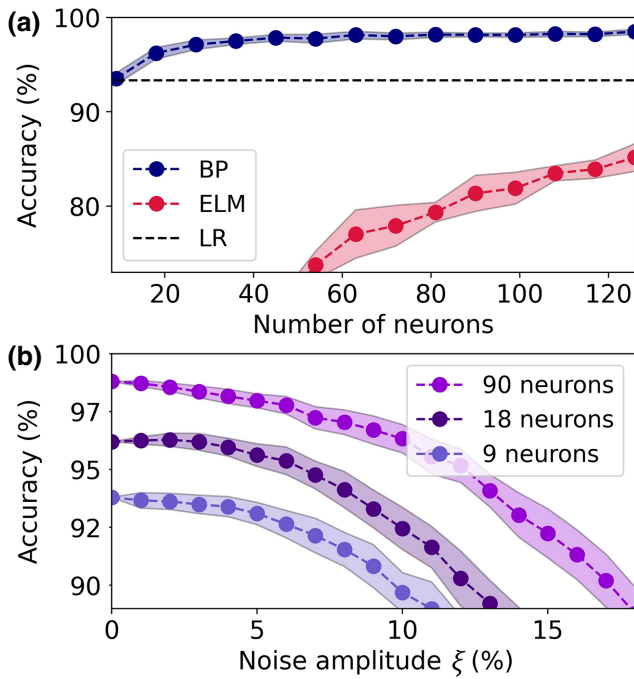


FIG. 3. (a) Numerically estimated inference accuracy of back-propagation (BP) networks, compared to “extreme learning machine” (ELM) networks, where only the output weights are adjusted, and the linear classification using logistic regression (LR). The horizontal axis corresponds to the number of neurons in the hidden layer. (b) Accuracy of inference in function of the amplitude of noise.

(ELM) network [56,57]. The latter network has an identical architecture to ours, but the input synaptic weights W_{in} are random and not adjusted in the training phase. Since only the synaptic weights in the output layer are adjusted, it can be considered a simplified feedforward analog of a reservoir computing network. Additionally, we compare these results with the accuracy level obtained using a linear classification method (taught using logistic regression), where nonlinear transformation of data is absent. It is clear that the use of backpropagation to adjust the weights results in a significant improvement of accuracy as compared to the ELM network, and in contrast to reservoir computing, it surpasses the linear classification method even for a network with a very small number of nodes.

We also consider the effect of noise on the performance of the network. In Fig. 3(b) we show accuracy as a function of the amplitude of noise. We add to the activation functions φ_i random variables $\delta\varphi_i$, $\tilde{\varphi}_i = \varphi_i + \delta\varphi_i$, that have Gaussian distributions with zero mean and standard deviations equal to

$$\sigma(\delta\varphi_i) = \xi \Delta_i, \quad (3)$$

where $\Delta_i = \max(I_i^{\text{out}}) - \min(I_i^{\text{out}})$ is the amplitude of the i th polariton neuron response and ξ is the relative

amplitude of noise. The accuracy typically remains high if the noise is weak, even in the case of small neural networks. Only large-amplitude noise can significantly deteriorate the network performance, while larger networks appear to be more resistant to strong noise.

III. DISCUSSION

Exciton-polaritons enable high speed and energy efficiency of all-optical computation [33]. Here, the light intensity per neuron measured before the sample was approximately 1 mW (or $9 \mu\text{W}/\mu\text{m}^2$ with $6 \mu\text{m}$ node radius). For comparison, the intensity of the control beam in Zuo *et al.* [50] was up to $50 \mu\text{W}$ per node. However, the most important figure of merit is the energy cost, which depends also on the data rate of the device. In polariton systems, the natural time-scale is of the order of picoseconds, which makes it possible to achieve optical nonlinearities at the cost of femtojoules. For comparison, the energy cost of nonlinear activation based on phase change materials was a few hundred picojoules [13]. One could also compare the energy cost to electronics, where the highest energy cost arises from linear operations, and is of the order of 100 pJ per multiply and accumulate (MAC) in graphics processing unit systems and 10 pJ per MAC in tensor processing unit systems in the case of fully connected networks [58]. On the other hand, the above estimates for optical systems do not take into account optical signal generation, detection, and additional electronic elements of the system. The energy cost of these needs to be reduced to achieve high overall system efficiency.

We emphasize that while in our experimental implementation the tunable weights are implemented in software, there are known all-optical methods for vector-matrix multiplication. These methods have been realized in many experiments, both in the case of coherent and incoherent light [45,46,49–53]. In particular, multiplication of a vector by a matrix containing 3000 elements has recently been implemented [47]. Thus, the inference stage could be realized in an all-optical system, without use of any electronic elements. Such an optical network could be completely passive, as it would not require any external power supply except for the laser source. Room-temperature polaritons can be used to avoid the cost of sample cooling [59–61]. Moreover, replacing the spatial light modulator with ultrafast modulators working at the gigahertz data rate, one could take advantage of the very short, picosecond time-scales of the optical system. Another possibility is to use an on-chip integrated version of the system, where transmission is tuned by optoelectronic modulators using the Stark effect acting on the exciton component [62]. Physics-aware backpropagation training [43] could be also used to perform backpropagation, possibly leading to considerable improvement of the accuracy in a multilayer system.

ACKNOWLEDGMENTS

MM acknowledges support from National Science Center, Poland grant 2020/37/B/ST3/01657. AO acknowledges support from National Science Center, Poland grant 2020/36/T/ST3/00417. BP acknowledges support from National Science Center, Poland grant 2020/37/B/ST3/01657. RP, VA, DB and DS acknowledge the Italian Ministry of University (MIUR) for funding through the PRIN project INPhoPOL (grant 2017P9FJBS) and the support of the joint project CNR-RFBR (CNR Project ID 14123, RFBR No. 20-52-7816).

APPENDIX A: MODEL

Our implementation is a dense feedforward neural network that contains a single hidden layer between the input and output layers. The hidden layer contains $i = 9M$ neurons, activations of which were obtained by fitting functions (2) to the measured polariton node responses. The particular activation function φ_i depends on the physical node (A, B, C, . . . , I) that it corresponds to. The neural network transforms the input vector \mathbf{x} into the output $y(\mathbf{x})$ according to

$$y(\mathbf{x}) = f(\mathbf{W}_{\text{out}}\phi(\mathbf{W}_{\text{in}}\mathbf{x} + \mathbf{b}_{\text{in}}) + \mathbf{b}_{\text{out}}) \quad (\text{A1})$$

where $f(\cdot)$ is the softmax function, \mathbf{W}_{out} is the weight matrix of connections between the output and the hidden layer. The function ϕ applies neuron activation functions $\phi(\mathbf{I}^{\text{in}}) = [\varphi_1(I_1^{\text{in}}), \varphi_2(I_2^{\text{in}}), \varphi_3(I_3^{\text{in}}), \dots, \varphi_{9M}(I_{9M}^{\text{in}})]$. The matrix \mathbf{W}_{in} contains the weight between the input and the hidden layer, while \mathbf{b}_{in} and \mathbf{b}_{out} are the input and output bias vectors, respectively. We optimize the weight matrices and bias vectors using the ADAM optimizer. We use the supervised learning method and the backpropagation algorithm. The application of the backpropagation method is possible thanks to the simulation of neuron activations by sigmoid functions, which are continuous and differentiable.

APPENDIX B: METHODS

The semiconductor microcavity used in these experiments is a planar cavity with three 8 nm $\text{In}_{0.04}\text{Ga}_{0.96}\text{As}$ quantum wells embedded between two AlAs/GaAs distributed Bragg reflectors, and kept at a temperature of about 7 K. The high quality factor ($Q \sim 10^4$) of this sample results in a polariton lifetime of 10 ps. In particular, the region of the sample we used exhibits an exciton-cavity detuning of $\delta \simeq 1.7$ meV and polaritons are pumped by a continuous-wave laser tuned at $\lambda = 836.64$ nm.

To shape the profile of the laser beam we used a spatial light modulator (SLM), a liquid crystal display with a surface area of approximately 2 cm² and a 1920 × 1080 resolution. Applying a voltage to the cells changes the

orientation of the liquid crystals and in turn the effective refractive index seen by the incident light. The control of the birefringence of each pixel allows us to spatially design the amplitude and phase of the reflected wave.

To create the node lattice on the sample we reconstructed the real space image of the SLM, reduced in size by a factor of 50. On top of the displayed lattice pattern we used a blazed grating pattern that serves two purposes: on the one hand, it allows us to block the zeroth-order reflection from the SLM that brings the nonmodulated part of the laser; on the other hand, by changing the diffraction efficiency of the grating we are able to tune the intensity of the individual nodes. Furthermore a phase difference of π is applied between each node, for a better separation.

The pattern is focused onto the microcavity sample by a camera objective with a focal length of 5 cm. The pump beam frequency is slightly blue-detuned with respect to the polariton resonance: increasing the pump power results in an energy shift of the polariton resonance due to the exciton-mediated interactions. The resulting transmission intensity is a highly nonlinear function of the pump power, realizing a physical implementation of a sigmoidal response function. The emission is collected with a 2 cm aspheric lens and recorded on a charge-coupled device.

-
- [1] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, *Nature* **521**, 436 (2015).
 - [2] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, *Learning internal representations by error propagation*, Tech. Rep. (California Univ San Diego La Jolla Inst for Cognitive Science, 1985).
 - [3] J. Zhu, T. Zhang, Y. Yang, and R. Huang, A comprehensive review on emerging artificial neuromorphic devices, *Appl. Phys. Rev.* **7**, 011312 (2020).
 - [4] J. Misra and I. Saha, Artificial neural networks in hardware: A survey of two decades of progress, *Neurocomputing* **74**, 239 (2010).
 - [5] D. Marković, A. Mizrahi, D. Querlioz, and J. Grollier, Physics for neuromorphic computing, *Nat. Rev. Phys.* **2**, 499 (2020).
 - [6] P. A. Merolla, J. V. Arthur, R. Alvarez-Icaza, A. S. Cassidy, J. Sawada, F. Akopyan, B. L. Jackson, N. Imam, C. Guo, Y. Nakamura, B. Brezzo, I. Vo, S. K. Esser, R. Appuswamy, B. Taba, A. Amir, M. D. Flickner, W. P. Risk, R. Manohar, and D. S. Modha, A million spiking-neuron integrated circuit with a scalable communication network and interface, *Science* **345**, 668 (2014).
 - [7] M. Davies *et al.*, Loihi: A neuromorphic manycore processor with on-chip learning, *IEEE Micro* **38**, 82 (2018).
 - [8] S. B. Furber, F. Galluppi, S. Temple, and L. A. Plana, The SpiNNaker project, *Proc. IEEE* **102**, 652 (2014).
 - [9] B. V. Benjamin, P. Gao, E. McQuinn, S. Choudhary, A. R. Chandrasekaran, J.-M. Bussat, R. Alvarez-Icaza, J. V. Arthur, P. A. Merolla, and K. Boahen, Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations, *Proc. IEEE* **102**, 699 (2014).

- [10] P. R. Prucnal and B. J. Shastri, *Neuromorphic Photonics* (CRC Press, Boca Raton, 2017).
- [11] G. Wetzstein, A. Ozcan, S. Gigan, S. Fan, D. Englund, M. Soljačić, C. Denz, D. A. Miller, and D. Psaltis, Inference in artificial intelligence with deep optics and photonics, *Nature* **588**, 39 (2020).
- [12] B. J. Shastri, A. N. Tait, T. F. de Lima, W. H. Pernice, H. Bhaskaran, C. D. Wright, and P. R. Prucnal, Photonics for artificial intelligence and neuromorphic computing, *Nat. Photonics* **15**, 102 (2021).
- [13] J. Feldmann, N. Youngblood, C. D. Wright, H. Bhaskaran, and W. H. P. Pernice, All-optical spiking neurosynaptic networks with self-learning capabilities, *Nature* **569**, 208 (2019).
- [14] 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.* **7**, 1 (2017).
- [15] K. Vandoorne, P. Mechet, T. Van Vaerenbergh, M. Fiers, G. Morthier, D. Verstraeten, B. Schrauwen, J. Dambre, and P. Bienstman, Experimental demonstration of reservoir computing on a silicon photonics chip, *Nat. Commun.* **5**, 3541 (2014).
- [16] X. Lin, Y. Rivenson, N. T. Yardimci, M. Veli, Y. Luo, M. Jarrahi, and A. Ozcan, All-optical machine learning using diffractive deep neural networks, *Science* **361**, 1004 (2018).
- [17] P. Antonik, N. Marsal, and D. Rontani, Large-scale spatiotemporal photonic reservoir computer for image classification, *IEEE J. Sel. Top. Quantum Electron.* **26**, 1 (2019).
- [18] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Laroche, D. Englund, and M. Soljacic, Deep learning with coherent nanophotonic circuits, *Nat. Photonics* **11**, 441 (2017).
- [19] A. Opala, S. Ghosh, T. C. H. Liew, and M. Matuszewski, Neuromorphic computing in Ginzburg-Landau polariton-lattice systems, *Phys. Rev. Appl.* **11**, 064029 (2019).
- [20] D. Ballarini, A. Gianfrate, R. Panico, A. Opala, S. Ghosh, L. Dominici, V. Ardizzone, M. De Giorgi, G. Lerario, G. Gigli, T. C. H. Liew, M. Matuszewski, and D. Sanvitto, Polaritonic neuromorphic computing outperforms linear classifiers, *Nano Lett.* **20**, 3506 (2020).
- [21] R. Mirek, A. Opala, P. Comaron, M. Furman, M. Król, K. Tyszka, B. Seredyński, D. Ballarini, D. Sanvitto, T. C. H. Liew, W. Pacuski, J. Suffczyński, J. Szczytko, M. Matuszewski, and B. Piętka, Neuromorphic binarized polariton networks, *Nano Lett.* **21**, 3715 (2021).
- [22] A. Kavokin, J. J. Baumberg, G. Malpuech, and F. P. Laussy, *Microcavities* (Oxford University Press, Oxford, 2017).
- [23] I. Carusotto and C. Ciuti, Quantum fluids of light, *Rev. Mod. Phys.* **85**, 299 (2013).
- [24] A. Dreismann, H. Ohadi, Y. del Valle-Inclan Redondo, R. Balili, Y. G. Rubo, S. I. Tsintzos, G. Deligeorgis, Z. Hatzopoulos, P. G. Savvidis, and J. J. Baumberg, A sub-femtojoule electrical spin-switch based on optically trapped polariton condensates, *Nat. Mater.* **15**, 1074 (2016).
- [25] A. Amo, T. C. H. Liew, C. Adrados, R. Houdré, E. Giacobino, A. V. Kavokin, and A. Bramati, Exciton-polariton spin switches, *Nat. Photonics* **4**, 361 (2010).
- [26] M. De Giorgi, D. Ballarini, E. Cancellieri, F. M. Marchetti, M. H. Szymanska, C. Tejedor, R. Cingolani, E. Giacobino, A. Bramati, G. Gigli, and D. Sanvitto, Control and ultrafast dynamics of a two-fluid polariton switch, *Phys. Rev. Lett.* **109**, 266407 (2012).
- [27] D. Ballarini, M. D. Giorgi, E. Cancellieri, R. Houdré, E. Giacobino, R. Cingolani, A. Bramati, G. Gigli, and D. Sanvitto, All-optical polariton transistor, *Nat. Commun.* **4**, 1778 (2013).
- [28] T. Gao, P. S. Eldridge, T. C. H. Liew, S. I. Tsintzos, G. Stavrinidis, G. Deligeorgis, Z. Hatzopoulos, and P. G. Savvidis, Polariton condensate transistor switch, *Phys. Rev. B* **85**, 235102 (2012).
- [29] A. V. Zasedatelev, A. V. Baranikov, D. Urbonas, F. Scaffirimoto, U. Scherf, T. Stöferle, R. F. Mahrt, and P. G. Lagoudakis, A room-temperature organic polariton transistor, *Nat. Photonics* **13**, 378 (2019).
- [30] A. V. Baranikov, A. V. Zasedatelev, D. Urbonas, F. Scaffirimoto, U. Scherf, T. Stöferle, R. F. Mahrt, and P. G. Lagoudakis, All-optical cascaded universal logic gate with sub-picosecond operation, arXiv preprint [arXiv:2005.04802](https://arxiv.org/abs/2005.04802) (2020).
- [31] T. C. H. Liew, A. V. Kavokin, and I. A. Shelykh, Optical circuits based on polariton neurons in semiconductor microcavities, *Phys. Rev. Lett.* **101**, 016402 (2008).
- [32] T. Espinosa-Ortega and T. C. H. Liew, Perceptrons with Hebbian learning based on wave ensembles in spatially patterned potentials, *Phys. Rev. Lett.* **114**, 118101 (2015).
- [33] M. Matuszewski, A. Opala, R. Mirek, M. Furman, M. Król, K. Tyszka, T. C. H. Liew, D. Ballarini, D. Sanvitto, J. Szczytko, and B. Piętka, Energy-efficient neural network inference with microcavity exciton polaritons, *Phys. Rev. Appl.* **16**, 024045 (2021).
- [34] H. Jaeger and H. Haas, Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication, *Science* **304**, 78 (2004).
- [35] W. Maass, T. Natschlag, and H. Markram, Real-time computing without stable states: A new framework for neural computation based on perturbations, *Neural Comput.* **14**, 2531 (2002).
- [36] M. Lukoševičius and H. Jaeger, Reservoir computing approaches to recurrent neural network training, *Comput. Sci. Rev.* **3**, 127 (2009).
- [37] 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.* **4**, 1364 (2013).
- [38] J. Torrejon, M. Riou, F. A. Araujo, S. Tsunegi, G. Khalsa, D. Querlioz, P. Bortolotti, V. Cros, K. Yakushiji, A. Fukushima, H. Kubota, S. Yuasa, M. D. Stiles, and J. Grollier, Neuromorphic computing with nanoscale spintronic oscillators, *Nature* **547**, 428 (2017).
- [39] C. Du, F. Cai, M. A. Zidan, W. Ma, S. H. Lee, and W. D. Lu, Reservoir computing using dynamic memristors for temporal information processing, *Nat. Commun.* **8**, 2204 (2017).

- [40] G. Tanaka, T. Yamane, J. B. Héroux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, and A. Hirose, Recent advances in physical reservoir computing: A review, *Neural Netw.* **115**, 100 (2019).
- [41] L. Larger, A. Baylón-Fuentes, R. Martinenghi, V. S. Udaltsov, Y. K. Chembo, and M. Jacquot, High-speed photonic reservoir computing using a time-delay-based architecture: Million words per second classification, *Phys. Rev. X* **7**, 011015 (2017).
- [42] T. Zhou, X. Lin, J. Wu, Y. Chen, H. Xie, Y. Li, J. Fan, H. Wu, L. Fang, and Q. Dai, Large-scale neuromorphic optoelectronic computing with a reconfigurable diffractive processing unit, *Nat. Photonics* **15**, 367 (2021).
- [43] L. G. Wright, T. Onodera, M. M. Stein, T. Wang, D. T. Schachter, Z. Hu, and P. L. McMahon, Deep physical neural networks trained with backpropagation, *Nature* **601**, 549 (2022).
- [44] T. Wang, S.-Y. Ma, L. G. Wright, T. Onodera, B. C. Richard, and P. L. McMahon, An optical neural network using less than 1 photon per multiplication, *Nat. Commun.* **13**, 123 (2022).
- [45] J. W. Goodman, A. Dias, and L. Woody, Fully parallel, high-speed incoherent optical method for performing discrete Fourier transforms, *Opt. Lett.* **2**, 1 (1978).
- [46] 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.* **8**, 1 (2018).
- [47] J. Spall, X. Guo, T. D. Barrett, and A. Lvovsky, Fully reconfigurable coherent optical vector–matrix multiplication, *Opt. Lett.* **45**, 5752 (2020).
- [48] M. Abadi *et al.*, Tensorflow: Large-scale machine learning on heterogeneous distributed systems (2015).
- [49] J. Bueno, S. Maktoobi, L. Froehly, I. Fischer, M. Jacquot, L. Larger, and D. Brunner, Reinforcement learning in a large-scale photonic recurrent neural network, *Optica* **5**, 756 (2018).
- [50] Y. Zuo, B. Li, Y. Zhao, Y. Jiang, Y.-C. Chen, P. Chen, G.-B. Jo, J. Liu, and S. Du, All-optical neural network with nonlinear activation functions, *Optica* **6**, 1132 (2019).
- [51] N. H. Farhat, D. Psaltis, A. Prata, and E. Paek, Optical implementation of the Hopfield model, *Appl. Opt.* **24**, 1469 (1985).
- [52] M. Gruber, J. Jahns, and S. Sinzinger, Planar-integrated optical vector-matrix multiplier, *Appl. Opt.* **39**, 5367 (2000).
- [53] T. Lu, S. Wu, X. Xu, and T. Francis, Two-dimensional programmable optical neural network, *Appl. Opt.* **28**, 4908 (1989).
- [54] A. Reuther, P. Michaleas, M. Jones, V. Gadepally, S. Samsi, and J. Kepner, in *2020 IEEE High Performance Extreme Computing Conference (HPEC)* (IEEE, New York, 2020), p. 1.
- [55] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* **86**, 2278 (1998).
- [56] G.-B. Huang, Q.-Y. Zhu, and C.-K. Siew, Extreme learning machine: Theory and applications, *Neurocomputing* **70**, 489 (2006).
- [57] D. Pierangeli, G. Marcucci, and C. Conti, Photonic extreme learning machine by free-space optical propagation, *Photon. Res.* **9**, 1446 (2021).
- [58] N. P. Jouppi *et al.*, in *Proceedings of the 44th annual International Symposium on Computer Architecture* (ISCA, Toronto, 2017), p. 1.
- [59] S. Christopoulos, G. B. H. von Högersthal, A. J. D. Grundy, P. G. Lagoudakis, A. V. Kavokin, J. J. Baumberg, G. Christmann, R. Butté, E. Feltin, J.-F. Carlin, and N. Grandjean, Room-temperature polariton lasing in semiconductor microcavities, *Phys. Rev. Lett.* **98**, 126405 (2007).
- [60] A. Fieramosca, L. Polimeno, V. Ardizzone, L. De Marco, M. Pugliese, V. Maiorano, M. De Giorgi, L. Dominici, G. Gigli, D. Gerace, D. Ballarini, and D. Sanvitto, Two-dimensional hybrid perovskites sustaining strong polariton interactions at room temperature, *Sci. Adv.* **5**, eaav9967 (2019).
- [61] F. Li, L. Orosz, O. Kamoun, S. Bouchoule, C. Brimont, P. Disseix, T. Guillet, X. Lafosse, M. Leroux, J. Leymarie, M. Mexis, M. Mihailovic, G. Patriarche, F. Réveret, D. Solnyshkov, J. Zuniga-Perez, and G. Malpuech, From excitonic to photonic polariton condensate in a zno-based microcavity, *Phys. Rev. Lett.* **110**, 196406 (2013).
- [62] D. G. Suárez-Forero, F. Riminucci, V. Ardizzone, M. D. Giorgi, L. Dominici, F. Todisco, G. Lerario, L. N. Pfeiffer, G. Gigli, D. Ballarini, and D. Sanvitto, Electrically controlled waveguide polariton laser, *Optica* **7**, 1579 (2020).