Neural Networks Based on Ultrafast Time-Delayed Effects in Exciton Polaritons

R. Mirek,[1] A. Opala,[2] M. Furman,[1] M. Król,[1] K. Tyszka,[1] B. Seredyński,[1] W. Pacuski,[1] J. Suffczyński,[1] J. Szczytko,[1] M. Matuszewski,[2] and B. Piętka[1,*]

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

[2] *The Institute of Physics, Polish Academy of Sciences, Aleja Lotników 32/46, Warsaw PL-02-668, Poland*

(Received 31 December 2021; accepted 2 May 2022; published 23 May 2022)

We demonstrate that time-delayed nonlinear effects in exciton polaritons can be used to construct neural networks where information is coded in optical pulses arriving consecutively on the sample. The highly nonlinear effects are induced by time-dependent interactions with the excitonic reservoir. These nonlinearities allow the creation of a nonlinear XOR logic gate that can perform operations on the picosecond timescale. An optoelectronic neural network based on the constructed logic gate performs the classification of spoken digits with a high accuracy rate.

DOI: 10.1103/PhysRevApplied.17.054037

I. INTRODUCTION

The concept of artificial neural networks (ANNs) is a widely used approach to machine learning, designed to deal with large amounts of data. ANNs are highly demanded for various applications such as image and speech recognition, pattern detection, optimization, prediction of stock exchange variations, dedicated advertising, autonomous cars, etc. [1]. However, the speed and efficiency of software ANN implementations is limited by the von Neumann architecture of traditional computers, where memory and computing units are physically separated. To circumvent this issue, neuromorphic computing implements neural networks in hardware where access to external memory is not necessary.

In order to achieve high speed and energy efficiency, photonic systems are natural candidates [2–11]. In electronic devices, communication requires charging of connection wire capacitance for every bit of information, which leads to significant energy dissipation. At very high data rates, this leads to losses, which are unacceptable from the practical point of view. This problem is absent in the case of optical links, where information travels at the speed of light and with almost no energy cost. For this reason, electronic wires are being replaced by optical links at smaller and smaller scales, from long-haul optical fibers, to interconnections in data centers, to direct on-chip optical connections [12]. It has also stimulated recent interest in optical computing, leading to remarkable achievements including efficient optical vector-matrix multiplication and neural networks [13,14]. One of the main practical obstacles in implementing optical data processing is the weakness of the nonlinear response of optical media, or equivalently photon-photon interactions, which is necessary to implement an activation function of a neuron or a transistor. From this point of view, semiconductor exciton polaritons, where photons and matter excitations (excitons) coexist in a quantum superposition state [15], are exceptionally promising candidates [16]. The excitonic component is providing the strong interactions required for low-threshold nonlinear operation [17,18]. Thanks to the photonic component, they are able to process data at very short timescales and transport data at the speed of light [19–21].

Neural networks based on exciton polaritons have been recently investigated both theoretically [16,22–24] and experimentally [25–27]. Results presented so far relied on reservoir computing [24] and binarized extreme learning machines [28] approaches where most of the network connections are static and only the output layer is trainable. Recently, it was demonstrated that the highly efficient backpropagation training algorithm can be implemented in an exciton-polariton network, leading to an increase of prediction accuracy [27].

In the experimental realizations presented in the literature, input data has been provided either by light pulses arriving at the same time or continuous light beams shaped by spatial light modulators. However, in many practical applications, such as signal processing or speech recognition, the input data has the form of a time-coded series. Moreover, dynamical systems are believed to provide solutions to the problems in miniaturization as they allow the number of spatial nodes to be decreased [29]. Therefore,

*barbara.pietka@fuw.edu.pl

the computing system needs to be able to process data arriving at different instants.

In this work, we use the dynamics of exciton polaritons to construct a binarized neural network [30,31] where hidden-layer neurons are XOR logic gates processing time-coded data. We have recently demonstrated that spatial-coded polariton binarized neural networks can be realized [26] reaching state-of-the-art accuracy at the Modified National Institute of Standards and Technology pattern recognition task. Here, we show that the dynamical properties of polaritons can be exploited to use the time domain, instead of the spatial domain, to process the input signal. This is enabled by time-delayed interactions successfully realized in many photonic systems [3,32,33] and also recently in polaritons arranged in spatial lattices [34]. In our realization based on polaritons, the exceptionally strong time-delayed nonlinear effects are provided by the interplay between the timescales of polariton condensate lifetime and the decay of the so-called reservoir of uncondensed excitons. As a proof-of-principle demonstration, we show that our system is able to make highly accurate classification predictions in a spoken digit recognition task.

II. SAMPLE AND EXPERIMENTAL DETAILS

The investigated sample is a semiconductor heterostructure based on CdTe microcavity with two distributed Bragg reflectors (DBRs) made of quaternary (Cd, Zn, Mg)Te materials. The microcavity contains three pairs of semimagnetic quantum wells of 20 nm thickness with small concentration (about 0.5%) of manganese ions. The structure is grown on a (100)-oriented GaAs substrate by molecular beam epitaxy. The sample is placed in an optical cryostat at liquid helium temperature and excited nonresonantly with pulses of 3 ps width at frequency of 76 MHz. The excitation energy is matching the first minimum of the DBRs stopband on the high-energy side. The emission from the sample is collected on the CCD camera or streak camera.

Our experiment relies directly on the properties of exciton polaritons. These quasiparticles result from the strong coupling of excitons in quantum wells to photon modes in a microcavity [15]. At high enough densities, they may undergo a transition to a nonequilibrium Bose-Einstein condensate and form a quantum fluid of light [35]. We create a polariton population using two independent laser beams with an adjustable time delay between the pulses, as illustrated in Fig. 1. The first pulse creates a long-lived reservoir (uncondensed particles) allowing the formation of exciton polaritons through energy relaxation and scattering. The reservoir remains populated when the second pulse excites the sample. This leads to the dramatic increase of the density of polaritons and more efficient relaxation to the ground state due to condensation by stimulated scattering. Collective polariton population is

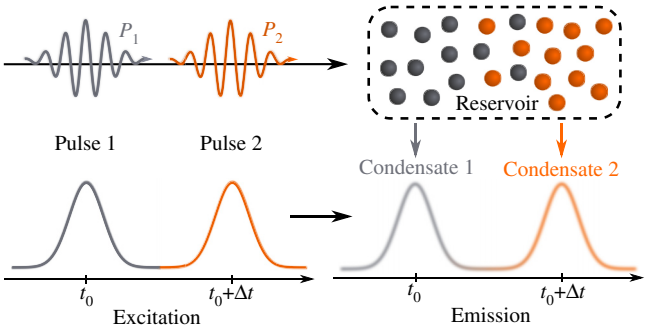

FIG. 1. Scheme of the time-delayed interactions. Two different pulses P_1 and P_2 delayed in time create a long-living excitonic reservoir. Relaxation from the reservoir leads to a nonlinear polariton population delayed in time.

therefore coupled through a single reservoir, providing time-delayed interactions.

In the realization of our polaritonic logic gate, the pulses correspond to two inputs where logic states are encoded via the excitation power and delay between the pulses. Emission from the condensate serves as the output signal.

III. TIME-CODED HIGHLY NONLINEAR XOR GATE

Our method provides large nonlinearities originating from time-delayed interactions of polaritons with a long-lived excitonic reservoir and from the buildup of circular polarization in the condensate. It should be noted that the nonlinear effects are negligible below the condensation threshold. Figure 2 illustrates the emission signal from exciton polaritons excited with two pulses delayed in time (brown line). The first peak originates from the weak emission of polaritons below the condensation threshold. The first pulse, however, creates a long-lived reservoir that remains active and the second pulse increases the population of excitons that provide gain to reach the condensation threshold. This results in an observation of a second peak with much higher intensity than the first one. To demonstrate the nonlinear nature of this phenomenon, we excite the sample separately with the corresponding laser pulses. The gray and orange lines indicate emission originating from the excitation with the first and the second pulse separately. Both peaks are multiplied 5 times for better visibility. We observe that while the first emission peak has almost the same intensity in both schemes, the second peak induced by a single pulse has significantly lower intensity than the peak resulting from both pulses. The magnitude of nonlinear effects depends on the time delay between the pulses. The inset in Fig. 2 shows the intensity of the second peak measured for different delays between the pulses. The observed dependence has a strongly nonlinear character and a shape similar to the sigmoid function (marked with a red line).

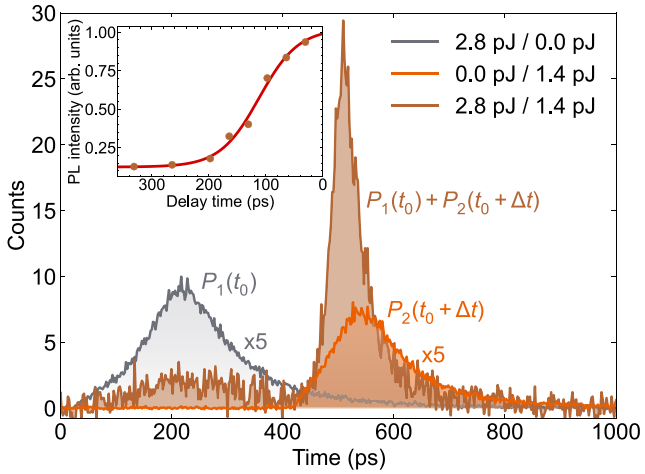

FIG. 2. Evidence of nonlinearity. Time-resolved photoluminescence from the exciton polaritons pumped with the first pulse $P_1(t_0)$, the second pulse $P_2(t_0+\Delta t)$, and two time-delayed pulses $P_1(t_0)+P_2(t_0+\Delta t)$. The delay between the input pulses is 390.7 ps. Emission originating from polaritons excited with a single pulse is increased 5 times for better visibility. The legend shows the pulse energy of the first and the second pulse. The inset illustrates intensity of the second emission peak as a function of delay between the pulses.

The nonlinear response of the system is used to construct a XOR gate, which can be considered the simplest form of a nonlinear machine-learning problem that cannot be implemented by a simple thresholding of a linear combination of inputs, as is the case for OR and AND gates. The XOR classification task is designed to distinguish patterns in two classes, consisting of input combinations (00, 11) in the first class and (01, 10) in the second class. Moreover, it was demonstrated that XOR gates can be used as building blocks of binarized neural networks, which are able to perform complex tasks such as image recognition very efficiently [30,31]. In this case, XOR gates are used in the role of neurons. The truth table of the XOR problem is shown in Fig. 3(a). The XOR logic gate is an example of a problem that is not linearly separable, as schematically represented in Fig. 3(b). The binary inputs correspond to the x and y axes in the figure. The linear classification task can be performed at most with 75% accuracy as shown with the dashed line. To overcome this problem, we use the nonlinear transformation of the inputs by the polariton condensate. A linear combination of gate inputs x_1, x_2 and a sufficiently nonlinear output from the condensate y reduce this problem into a linearly separable one in the three-dimensional input-output space [26]. In this case, the result of the XOR operation is given by the sign of the expression $z = w_1x_1 + w_2x_2 + w_3y(x_1,x_2) + b$, where w_1 and w_2 are the input weights, w_3 is the weight of the nonlinear output feature and b is the bias. The task of finding the proper weights and thresholding are

(a)

x_1	x_2	$x_1 \oplus x_2$
0	0	0
0	1	1
1	0	1
1	1	0

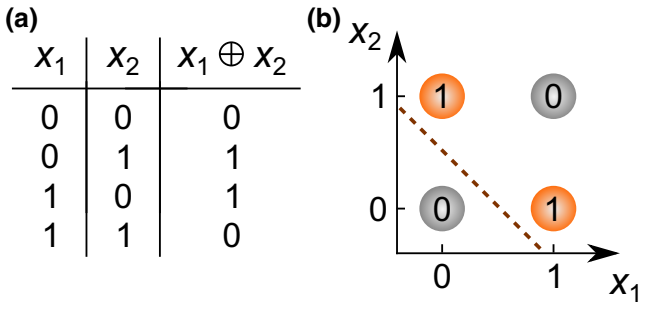

FIG. 3. XOR logic gate as a problem not solvable by linear separation. (a) A truth table. (b) The input parameter space and linear separation represented by dashed line.

both implemented in software. The weights are determined using a linear-regression algorithm.

The input-output characteristic of time-coded nonlinear effects for the polariton XOR gate is illustrated in Fig. 4. Panels (a)–(d) show real-space emission from a spot pumped with σ^+-polarized pulses and detected in the opposite (σ^-) circular polarization. For this polarization configuration, the obtained nonlinearities are the largest. The condensate is excited with two pulsed laser beams with different delay between them (long and short delay encoded 0 and 1) and two different excitation powers (low and high power encoded 0 and 1, respectively). Emission is collected on a CCD camera with the acquisition time of 60 ms. In this case, increasing the excitation power of laser pulses and decreasing the delay time between them results in a buildup of a condensate. Due to the right circular polarization of input pulses, condensation of semimagnetic

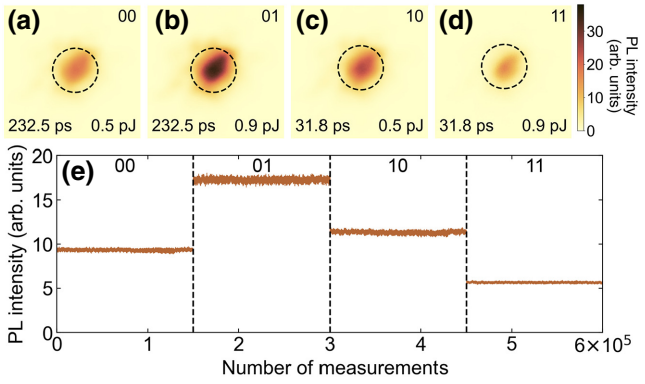

FIG. 4. Demonstration of time-coded nonlinearities used for XOR operation. (a)–(d) Real-space emission of a polariton condensate with different delay between excitation pulses (indicated at bottom left) and different power of the pulses (bottom right). Corresponding input configurations x_1x_2 are marked at the top right corner. Black dashed circles indicate the integration regions of 3 μm diameter used for further analysis. The color scale located on the right side is the same for all panels. (e) Output intensity for all four input combinations.

exciton polaritons results in a spontaneous buildup of σ^+ circular polarization. This leads to a vanishing σ^- component as observed in photoluminescence. This agrees with the emission observed in Figs. 4(a)–4(d) where the lowest signal intensity is measured for short delay and large power of the pulses. This method allows us to obtain nonlinearities exceeding significantly the noise existing in the system. It is worth noting that in this case the system dynamics is orders of magnitude faster than the speed of the CCD camera, what is discussed in Ref. IV.

We repeat the experiment to obtain training and testing datasets for the XOR problem. We perform linear classification on a data nonlinearly transformed by polariton condensate. As the nonlinear output signal, we use mean intensities in regions marked by black dashed circles in Figs. 4(a)–4(d). In Fig. 4(e) we show the intensities measured in the testing series. Each state is realized 150 000 times in the training phase and 30 000 times in the testing phase. It is worthwhile to note that the output intensities obtained for all input configurations do not have to resemble the XOR gate truth table. At this stage, only the high enough nonlinear output is crucial to further perform linear regression and obtain XOR logic gate with perfect accuracy.

The accuracy of the XOR gate depends on the ratio of the obtained nonlinearity to the noise present in the system. We calculate the degree of nonlinearity for a given set of data as [26]

$$\eta = \frac{\langle I_{00}\rangle + \langle I_{11}\rangle - \langle I_{01}\rangle - \langle I_{10}\rangle}{\sqrt{V_{00} + V_{01} + V_{10} + V_{11}}}, \tag{1}$$

where $\langle I_{ij}\rangle$ is the average output intensity for the ij input logic state and $V_{ij} = \langle \left(I_{ij} - \langle I_{ij}\rangle\right)^2\rangle$. It was demonstrated previously [26] that almost perfect XOR gate operation is achieved for nonlinearity degree $\eta > 5$. Here during the training phase, $\eta \approx 62$. The same procedure is implemented in the testing phase where $\eta \approx 57$. The differences between the values in the training and testing sets are originating mostly from the experimental noise. Despite the noise present in the system, it is clear that the proposed system is highly nonlinear. Furthermore, the nonlinearities are large enough to obtain 100% accuracy rate of the XOR gate in the whole testing dataset.

IV. ULTRAFAST XOR OPERATION

In order to prove the picosecond timescale of the XOR gate operations, we perform a time-resolved experiment with a streak camera detection. We study time-resolved emission in four x_1x_2 input configurations coded by different excitation power of the pulses and different delay time between the pulses. Figure 5 demonstrates the signal observed in σ^+ circularly polarized emission. In Fig. 5(a), low excitation power and long delay (432 ps) between the pulses (00) results in the formation of low density exciton

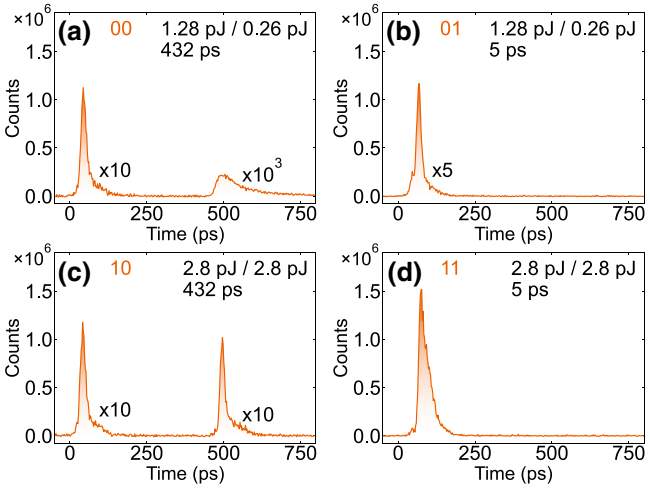

FIG. 5. Ultrafast XOR operation. (a)–(d) Time-resolved photoluminescence in the σ^+ circular polarization. Excitation power of the two excitation pulses and the delay between them are marked at the top right. Top left annotation indicates the realized x_1x_2 input configuration. Corresponding peaks are multiplied by marked factors for better visibility.

polaritons, which decay over a few hundred picoseconds. For two other input configurations (01) and (10), where the delay time is set to 5 ps or power of the pulses is high, we observe emission with similar intensities and decay times. On the other hand, the (11) configuration, with a short delay between high intensity pulses, results in a clearly stronger emission. To calculate the nonlinearities we integrate the signal in the first peak in each configuration, which leads to $\eta = 70$.

The same experimental conditions are applied during the measurements performed on a spectrometer with a CCD camera. Output is taken as the total intensity observed in the emission. The performed experiments resulted in a degree of nonlinearity above 50, which confirms that the nonlinearities remain high even for the signal averaged over many realizations.

Our results confirm that the time-coded XOR gate operates in a picosecond scale, but its output can be collected by a much slower detector, while still preserving high enough nonlinearities. In fact, there is only one factor playing a crucial role in a XOR gate operation speed. It is the lifetime of an excitonic reservoir, which imposes the minimum delay between the laser pulses and for this realization the value is about 250 ps. In particular, it is possible to create a XOR gate using constant delay between the pulses, defining the input parameter space with power of the first and second pulse.

V. SPEECH RECOGNITION

Previously, a binarized exciton-polariton neural network allowed pattern recognition to be performed with a high

accuracy in comparison to other neuromorphic hardware systems [26]. This work applies the same type of network architecture, realized using time-multiplexed XOR gates, for speech recognition. To determine the efficiency of our system, we consider a part of the TI 46-Word data set (available from the Linguistic Data Consortium) [36]. The database mentioned above is a commonly used speech recognition benchmark for artificial neural networks.

In our analysis, we use 500 training and 1275 testing samples in the form of audio waveform files. Each item is a recorded spoken word that corresponds to one of the ten digits (0 to 9). The training set includes words spoken by five different female voices, each word uttered 10 times. The testing set includes recordings of eight other females, repeating the digits approximately 16 times.

Using raw audio files as an input to a neural network is an ineffective method of speech recognition. To obtain the optimal classification results, the recorded signals are mapped onto the time-frequency domain using an acoustic transformation [37,38]. This method can be considered as a feature extraction technique, commonly used in automatic speech recognition. In our work, we employ two methods of audio waveform transformation: the Fourier transform (FT) and the Lyon cochlear-ear model (LCM). The Lyon cochlear-ear model is a nonlinear transformation used to decompose acoustic signals, which mimics the operation of a biological cochlear. It was confirmed that the application of LCM to acoustic signals greatly increases the recognition rate [37].

The conceptual scheme of the proposed neural network is shown in Fig. 6. Each audio sample from a data set is preprocessed and transformed into the time-frequency domain using N_f frequency channels. Next, all maps are normalized and resized to the size $N_f \times N_f$. Time-frequency patterns are converted into eight black and white bitmaps denoted by c_i, where $i = \{1, 2, \ldots, 8\}$, and c_i is given by the formula $c_i = \left| \lfloor I/2^{(i-1)} \rfloor \right|_2$, where $|a|_b$ is the modulo operation returning the remainder of division of a/b. The binary time-frequency maps are used as neural-network inputs. From each map, we select N pairs of pixels denoted by $p_1, p_2, p_3, \ldots, p_N$ which allow us to reveal nontrivial correlations between pixels in the bitmaps. The positions of the selected pixels are chosen randomly and did not change during the entire learning process. The N pairs of pixels are used as binary inputs of N gates, implemented by time multiplexing the XOR operation, realized experimentally. Next, we use the N outputs of the XOR gates as inputs for the last layer implemented in the software and perform a linear classification, denoted in the figure as LC, to predict the spoken word. To avoid overfitting, we use a stochastic gradient-descent method with regularization. As we train only the weights in the output layer used for linear classification, the network is an instance of an extreme learning machine [28].

In this experiment, we use the time-delayed interactions in the system only to a limited degree, to implement the XOR operation using the dynamics of a single node. However, one can envisage that in a more sophisticated setup, time coding of entire samples could be used to take the full advantage of the dynamical nature of the system and reduce the need for time multiplexing as demonstrated in Ref. [38]. On the other hand, the use of multiple spatial nodes of the polariton microcavity instead of time multiplexing could substantially increase the processing speed, with many operations taking place in parallel.

We examine the efficiency of the network using both spectrograms and cochleagrams as input data. We consider time-frequency maps of size (16×16), (32×32), and (64×64). Data samples delivered to the network are combinations of eight binary maps merged into vectors $[c_1, c_2, c_3, \ldots, c_8]$. The obtained results are presented in Table I. To measure the extent to which the hidden layer of XOR gates improves the accuracy, we compare the results with the case when the hidden layer is absent, and the data is processed directly by the linear classifier. We find that the hidden layer of XOR gates results in the improvement of accuracy up to a few percent. The highest increase of accuracy, obtained from the formula $\Delta = A_{\mathrm{XOR}} - A_{\mathrm{LC}}$, where A_{XOR} is the accuracy of the network including the

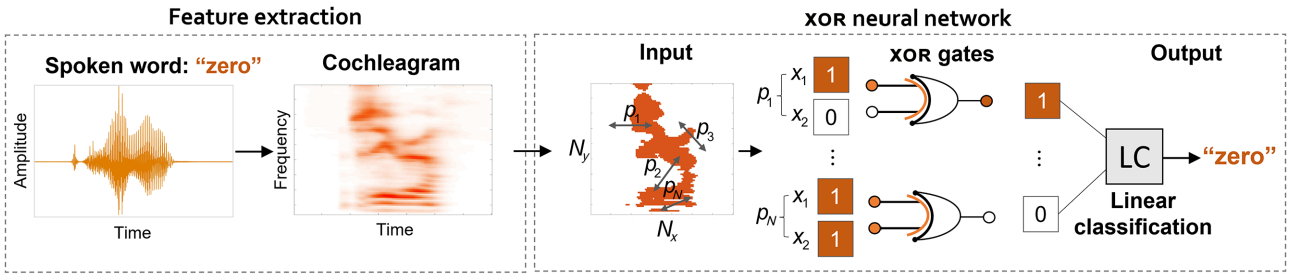

FIG. 6. Scheme of the binarized neural network used for the speech recognition task. The recorded audio signal transformed into a cochleagram is presented as an orange and white bitmap. Then the N randomly selected pixel pairs from the bitmap are used as inputs to a set of XOR logic gates, acting as neurons in the hidden layer. The layer of XOR gates is implemented experimentally by time multiplexing. The result of gate operations is used for linear classification in the output layer.

TABLE I. Classification accuracy determined for spectrograms and cochleagrams using a network with a linear classifier only (LC) and with a network that additionally includes a hidden layer of 5000 XOR gates (XOR). Parameter Δ is the improvement of the classification accuracy achieved by transforming the binarized data by the XOR layer. Each value is an average of five different realizations of neural-network teaching. Each realization corresponds to a different selection of randomly selected pixel pairs.

Size of time-frequency map			(16 × 16)	(32 × 32)	(64 × 64)
Methods	XOR	FT (%)	76.08	87.16	94.06
		LCM (%)	90.96	95.50	96.44
	LC	FT (%)	75.50	86.12	92.72
		LCM (%)	86.12	94.40	95.50
Δ(%)		Δ_{FT} (%)	0.58	1.04	1.34
		Δ_{LCM} (%)	4.84	1.10	0.94

XOR gate layer, and A_{LC} is the linear classification accuracy of binarized data, is equal to 4.84%, and is achieved for input in the form of a cochleagram containing 16 frequency channels. We observe modest improvement for other configurations. This indicates that the XOR gate layer is not very well suited to process data obtained directly from the Fourier transform, and on the other hand, that larger cochleagrams (32 × 32 and 64 × 64) already include large amount of nonlinear data and the use of XOR gate layer cannot improve the result significantly. The efficiency of the neural network presented in Table I describes the system with "ideal" XOR gates. These gates can be realized when the nonlinearity degree is $\eta > 5$.

VI. SUMMARY

We develop a method for the realization of exciton-polariton-based binarized neurons using time-delayed effects. We use time-coded addressing of a single exciton-polariton condensate with laser pulses delayed in time. The system is used to implement two machine-learning tasks, the nonlinear XOR problem, and spoken digit recognition. Our proof-of-principle experiment demonstrates the capability of exciton-polariton systems to process time-dependent data on the picosecond timescale.

ACKNOWLEDGMENTS

This work is supported by National Science Center, Poland: R.M. acknowledges 2019/33/N/ST3/02019, B.P. acknowledges 2017/27/B/ST3/00271, M.K. acknowledges 2015/18/E/ST3/00558, M.F. acknowledges 2020/37/B/ST3/01657, K.T. acknowledges 2020/04/X/ST7/01379, A.O. acknowledges 2019/35/N/ST3/01379, M.M. acknowledges 2020/37/B/ST3/01657.

[1] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, Nature **521**, 436 (2015).
[2] D. Brunner and D. Psaltis, Competitive photonic neural networks, Nat. Photonics **15**, 323 (2021).
[3] 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).
[4] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund, and M. Soljacic, Deep learning with coherent nanophotonic circuits, Nat. Photonics **11**, 441 (2017).
[5] 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).
[6] A. N. Tait, M. A. Nahmias, B. J. Shastri, and P. R. Prucnal, Broadcast and weight: An integrated network for scalable photonic spike processing, J. Lightwave Technol. **32**, 3427 (2014).
[7] P. Antonik, N. Marsal, D. Brunner, and D. Rontani, Human action recognition with a large-scale brain-inspired photonic computer, Nat. Mach. Intell. **1**, 530 (2019).
[8] G. Wetzstein, A. Ozcan, S. Gigan, S. Fan, D. Englund, M. Soljačić, C. Denz, D. A. B. Miller, and D. Psaltis, Inference in artificial intelligence with deep optics and photonics, Nature **588**, 39 (2020).
[9] B. J. Shastri, A. N. Tait, T. Ferreira de Lima, W. H. P. Pernice, H. Bhaskaran, C. D. Wright, and P. R. Prucnal, Photonics for artificial intelligence and neuromorphic computing, Nat. Photonics **15**, 102 (2021).
[10] X. Xu, M. Tan, B. Corcoran, J. Wu, A. Boes, T. G. Nguyen, S. T. Chu, B. E. Little, D. G. Hicks, R. Morandotti, A. Mitchell, and D. J. Moss, 11 TOPS photonic convolutional accelerator for optical neural networks, Nature **589**, 44 (2021).
[11] 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).
[12] C. Sun, M. T. Wade, Y. Lee, J. S. Orcutt, L. Alloatti, M. S. Georgas, A. S. Waterman, J. M. Shainline, R. R. Avizienis, and S. Lin, *et al.*, Single-chip microprocessor that communicates directly using light, Nature **528**, 534 (2015).
[13] 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).
[14] 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).
[15] A. Kavokin, J. J. Baumberg, G. Malpuech, and F. P. Laussy, *Microcavities* (Oxford University Press, Oxford, New York, 2017).
[16] 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).
[17] A. Delteil, T. Fink, A. Schade, S. Höfling, C. Schneider, and A. Imamoglu, Towards polariton blockade of confined exciton-polaritons, Nat. Mater. **18**, 219 (2019).

[18] 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).

[19] D. Ballarini, M. De 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).

[20] 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).

[21] A. V. Zasedatelev, A. V. Baranikov, D. Urbonas, F. Scafirimuto, U. Scherf, T. Stöferle, R. F. Mahrt, and P. G. Lagoudakis, A room-temperature organic polariton transistor, Nat. Photonics **13**, 378 (2019).

[22] 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).

[23] H. Xu, S. Ghosh, M. Matuszewski, and T. C. H. Liew, Universal Self-Correcting Computing with Disordered Exciton-Polariton Neural Networks, Phys. Rev. Appl. **13**, 064074 (2020).

[24] 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).

[25] 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).

[26] 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).

[27] A. Opala, R. Panico, V. Ardizzone, B. Piętka, J. Szczytko, D. Sanvitto, M. Matuszewski, and D. Ballarini, Teaching a neural network with non-tunable exciton-polariton nodes, (2021), ArXiv:2107.11156.

[28] G.-B. Huang, Q.-Y Zhu, and C.-K. Siew, Extreme learning machine: Theory and applications, Neurocomputing **70**, 489 (2006).

[29] L. Appeltant, M. C. Soriano, G. Van der Sande, J. Danckaert, S. Massar, J. Dambre, B. Schrauwen, C. R. Mirasso, and I. Fischer, Information processing using a single dynamical node as complex system, Nat. Commun. **2**, 468 (2011).

[30] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio, in *Advances in Neural Information Processing Systems 29*, edited by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., 2016), p. 4107.

[31] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi, XNOR-Net: in *Computer Vision - ECCV 2016*, edited by Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Springer International Publishing, Cham, 2016), p. 525.

[32] 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).

[33] J. Robertson, M. Hejda, J. Bueno, and A. Hurtado, Ultrafast optical integration and pattern classification for neuromorphic photonics based on spiking VCSEL neurons, Sci. Rep. **10**, 6098 (2020).

[34] J. D. Töpfer, H. Sigurdsson, L. Pickup, and P. G. Lagoudakis, Time-delay polaritonics, Commun. Phys. **3**, 2 (2020).

[35] I. Carusotto and C. Ciuti, Quantum fluids of light, Rev. Mod. Phys. **85**, 299 (2013).

[36] M. Liberman, R. Amsler, K. Church, E. Fox, C. Hafner, J. Klavans, M. Marcus, B. Mercer, J. Pedersen, P. Roossin, D. Walker, S. Warwick, and A. Zampolli, "TI 46-Word," https://catalog.ldc.upenn.edu/LDC93S9 (1993).

[37] F. A. Araujo, M. Riou, J. Torrejon, S. Tsunegi, D. Querlioz, K. Yakushiji, A. Fukushima, H. Kubota, S. Yuasa, M. D. Stiles, and J. Grollier, Role of non-linear data processing on speech recognition task in the framework of reservoir computing, Sci. Rep. **10**, 328 (2020).

[38] 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).