Polaritonic neuromorphic computing outperforms linear classifiers

Ballarini, Dario; Gianfrate, Antonio; Panico, Riccardo; Opala, Andrzej; Ghosh, Sanjib;
Dominici, Lorenzo; Ardizzone, Vincenzo; De Giorgi, Milena; Lerario, Giovanni; Gigli,
Giuseppe; Liew, Timothy Chi Hin; Matuszewski, Michal; Sanvitto, Daniele

2020

Ballarini, D., Gianfrate, A., Panico, R., Opala, A., Ghosh, S., Dominici, L., ... Sanvitto, D.
(2020). Polaritonic neuromorphic computing outperforms linear classifiers. Nano Letters,
20(5), 3506-3512. doi:10.1021/acs.nanolett.0c00435

<https://hdl.handle.net/10356/143433>

<https://doi.org/10.1021/acs.nanolett.0c00435>

This document is the Accepted Manuscript version of a Published Work that appeared in
final form in Nano Letters, copyright © American Chemical Society after peer review and
technical editing by the publisher. To access the final edited and published work see
<https://doi.org/10.1021/acs.nanolett.0c00435>

Downloaded on 26 Apr 2023 13:27:38 SGT

Polaritonic neuromorphic computing outperforms linear classifiers

D. Ballarini,¹ A. Gianfrate,¹ R. Panico,¹ A. Opala,²
 S. Ghosh,³ L. Dominici,¹ V. Ardizzone,¹ M. De Giorgi,¹ G. Lerario,¹ G. Gigli,¹
 T.C.H. Liew,³ M. Matuszewski,² D. Sanvitto,¹

¹CNR NANOTEC–Institute of Nanotechnology, Via Monteroni, 73100 Lecce, Italy

²Institute of Physics, Polish Academy of Sciences,
 Al. Lotnikw 32/46, PL-02-668 Warsaw, Poland

³School of Physical and Mathematical Science,
 Nanyang Technological University 637371, Singapore

Machine learning software applications are nowadays ubiquitous in many fields of science and society for their outstanding capability of solving computationally vast problems like the recognition of patterns and regularities in big datasets. One of the main goals of research is the realization of a physical neural network able to perform data processing in a much faster and energy-efficient way than the state-of-the-art technology. Here we show that lattices of exciton-polariton condensates accomplish neuromorphic computing using fast optical nonlinearities and with lower error rate than any previous hardware implementation. We demonstrate that our neural network significantly increases the recognition efficiency compared to the linear classification algorithms on one of the most widely used benchmarks, the MNIST problem, showing a concrete advantage from the integration of optical systems in reservoir computing architectures.

Introduction Artificial neural networks emerged as powerful information processing systems for the recognition and classification of patterns even with imperfect, noisy, or incomplete information. However, software implementations suffer from the so-called von Neumann bottleneck due to the physical separation of memory and processing units. On the other hand, neuromorphic designs attempt to adjust the structure of a physical system to the architecture of the neural network. This allows a native parallelism of tasks, increasing the efficiency of the computational process and avoiding the bottleneck of the von Neumann architecture.

Aside implementations using conventional binary logic to simulate neural networks, hardware implementations [1] have been based on memristors [2, 3, 4], spintronic [5], and optical systems [6]. Yet, the individual control of the weights of connections between nodes places a challenging requirement on hardware implementations of neural networks. Moreover, even being able to independently control many weight connections, one still has the task of choosing them. A promising solution appeared in the form of reservoir computing (RC) [7, 8], which operates with a random network of recurrent nodes, i.e., nodes with fixed bidirectional coupling, eliminating the need for their training and control. What is trained instead is a weight matrix that is applied to the output of the system. Such a design was implemented with photonic systems [9, 10], microwaves [11] and memristors [12], showing significant recognition rates on a standard set of hand-written digits (Mixed National Institute of Standards and Technology, MNIST). Despite these successfully demonstrations, linear classification algorithms applied directly on the MNIST dataset are still much more precise than any hardware realization of RC.

Following a recent theoretical proposal [13], here we experimentally implement a nonlinear neural network made of a lattice of driven-dissipative condensates of exciton-polaritons to realize a reservoir computer. Exciton-polaritons (hereafter polaritons) are formed in semiconductors (in our case, GaAs quantum wells) placed inside a cavity when excitons hybridize with photons [14]. Thanks to the interactions inherited from the electronic component, polaritons

possess strong Kerr nonlinearities while keeping their lifetime in the picosecond timescale[15]. Many experimental milestones have demonstrated successful realisation of proof of principle polariton devices, such as logic gates [16], switches [17], routers [18] and transistors [19, 20] with femtojoule energy cost per operation and few tens of picosecond recovery time [21]. More recently, lattices of localized polariton condensates have been used to form topological protected states [22, 23] and solving many-body hamiltonians [24, 25]. Rather than using polaritons to reproduce binary logic systems, we demonstrate here an alternative neuromorphic information processing architecture and show how the polariton nonlinearity can be harnessed to perform very efficient pattern recognition.

We demonstrate the recognition rate at the level of 93% for the MNIST dataset, higher than in any other hardware implementation of RC to date [12, 26] and the first to go beyond linear classification for pattern recognition. To achieve highly efficient and ultrafast pattern recognition we exploit the exceptional properties of exciton-polaritons with a non-conventional approach to reservoir computing, where data representing a single digit is processed simultaneously instead of sequentially.

We implement the polaritonic reservoir by creating a lattice of 8×8 coupled nodes, which act as nonlinear artificial neurons. The transmission intensity of the interacting nodes is the output of the reservoir, which is fed to the final output layer implemented on a computer. Note that this last layer can be implemented all-optically as it requires only a single vector-matrix multiplication at the testing stage, which is a linear operation [27]. Although the size of our system is modest, we readily demonstrate a high recognition rate for a MNIST dataset which consists of 28×28 handwritten digits. Our result is the first experimental demonstration of MNIST pattern recognition in an optical RC.

Figure 1: Scheme of the experimental configuration. **a**, An index j is assigned to each pixel of a 28×28 input image, such that the input intensities are a_j . These inputs are multiplied by an $8^2 \times 28^2$ sparse random matrix giving the input as $b_i = W_{ij}a_j$. Here i indexes different pixels of an 8×8 image, which directly couples to the 8×8 sites of the reservoir. The same procedure is used for all input images. **b**, The resulting dataset is sent to the SLM to pattern the laser beam. The nonlinear transmission (**c**) of the polariton RC produces an output (**d**), which is recorded by a CCD camera and multiplied by the weight matrix (**e**) to obtain the digit classification (**f**). **g**, Scheme of the experimental setup corresponding to the steps described in **b**, **c**, **d**.

Results The schematic of the experimental procedure is shown in Fig. 1. To distribute the information over the whole reservoir, each digit is converted into a matrix with the same number of elements as the number of nodes in the polariton network. This is done by multiplying the input digit (Fig. 1a) by a random matrix, which is taken the same for all digits. The particular choice of the random matrix has a limited effect on the final recognition efficiency (see Supplemental Information for details on the dataset preparation). Each node in the network is optically created by shaping with a spatial light modulator (SLM) the pump laser beam (wavelength $\lambda = 836$ nm) into an array of spots with controllable intensity and phase. The laser light is injected into a semiconductor microcavity, in which exciton-polaritons are created by tuning the laser frequency at the polariton resonance (see SI). The information of the digit is encoded in the intensities of the spots (Fig. 1b), which excite different realizations of the polariton network, while the junction between neurons is created by imposing a phase difference of $\Delta\theta = \pi$ between adjacent nodes [28]. The nonlinearity is provided by the Kerr type interactions between polaritons: for resonant excitation, the curve of the transmission (output) intensity as a function of input power shows a nonlinear behavior (Fig. 1c), which depends on the frequency detuning of the laser beam with respect to the polariton resonance [29]. Thus each network node can be considered as an artificial neuron with a nonlinear response function, while the fast polariton propagation within the cavity plane provides the effective connectivity between nodes. The threshold power (of the laser outside the cavity) for each node is typically around few meV, while the fingerprint of the whole polariton RC is limited to only $150\text{ }\mu\text{m} \times 150\text{ }\mu\text{m}$.

The transmission pattern from the polariton lattice (Fig. 1d) is recorded by the camera and the image, sampled with a chosen resolution, is taken as the output of the reservoir. Due to the limited throughput of our SLM, a random choice of 5000 digits (80% for the training stage, 20% for the testing stage) is used with no preprocessing to evaluate the performance of the network. Training is realized by finding weights in the output layer (Fig. 1e) that minimize

the error rate of predictions. We implement logistic regression, an efficient algorithm for linear classification, on a computer to adjust the output weights in the readout matrix. In the testing phase, the trained network is used to predict the recognition rate on examples that the system has not seen before. Note that at this stage the computer could be replaced by an all-optical device performing a linear vector-matrix multiplication operation [27]. We always compare our results against the pure logistic regression on input data to evaluate the gain of the recognition rate due to the nonlinear transformation performed by the network. To this scope, the transfer function of the SLM is calibrated to be strictly linear (see SI), in order to measure solely the contribution of the polariton nonlinearity to the pattern recognition efficiency. This is important because, while the final efficiency can be increased by adding a nonlinear transfer function to the SLM, the polariton array can work at much higher operational speed than the electronically controlled SLM. Indeed, while the SLM is used here as a convenient way to generate an array of polariton nodes, polariton lattices of a few thousands nodes can be realized through lithographic patterning for on-chip integration and this would allow direct excitation with the optical input at repetition rates up to the THz range [30].

In Fig. 2, we test the recognition efficiency of the polariton RC on the MNIST dataset. To avoid any spurious effect, such as the possibility that a continuous drift in the temperature would artificially increment the recognition efficiency, the input frames in the dataset are randomly shuffled, confirming the proper operation of the polariton RC in the classification task (see SI). Given the limited reservoir size considered here, we reduce the resolution of the MNIST digits from 28x28 to 7x7 pixels (top row “M” in Fig. 2 left). Although this operation loses some information, the recognition rate of the polariton RC remains higher than that set by logistic regression, which is illustrated with the shaded grey region in Fig. 2 right for different training set sizes. The exact recognition rate depends on which examples in the MNIST dataset are taken for training and testing. Thus the different data points have error bars, calculated from statistical

Figure 2: Comparison of polariton RC vs linear classifier. On the left, samples of four digits (1, 3, 5 and 7) processed in the polariton RC. First row “M”: MNIST digits with 7x7 resolution. Second row “S”: SLM mask after multiplication of “M” by the random matrix. Third row “L”: laser output after the SLM. Bottom row “P”: polariton RC, where the quadratic deviation from the mean is shown to highlight the nonlinear redistribution of the information. The spatial scale of the polariton RC is $150\text{ }\mu\text{m} \times 150\text{ }\mu\text{m}$. On the right, the success rate obtained with polariton RC is compared to that of linear classifier (shaded region, full MNIST data set of 60000 training digits and 10000 testing digits with 28x28 resolution) and with the state-of-the-art electronic and spintronic RC and terahertz deep neural network (D²NN) [12, 31, 32]. Training and testing set sizes are reported as (training:testing). The “*” indicates that in Ref. [31] the recognition is measured on a preselection of digits, actually giving an absolute efficiency of $\approx 81\%$. The logistic regression recognition rates on the same training and testing sets used in the experiments are shown as black and green open squares for M, S (91%) and L (90.8%), respectively. Logistic regression recognition rate (grey points) are calculated as an average over 10 randomly chosen subsets of the training dataset, and their error bars correspond to the standard error of the mean. For the experiments, the error bar is calculated by partitioning the testing set of 1000 digits into 10 subsets of 100 digits each as the standard error of the mean recognition rate.

analysis of different random choices of the training and testing. More training samples results in better performance, but it is remarkable that the polariton RC performs well even with a limited number of training samples. One can also compare to the recognition rate obtained from performing logistic regression on the 7x7 input digit set (point M in Fig. 2 right), performing

logistic regression on the 7×7 image SLM mask (point S in Fig. 2 right), and performing logistic regression on the laser beam image after the SLM mask (point L in Fig. 2 right). It can also be seen that the recognition rates obtained in previous hardware implementations of neuromorphic computing were lower.

In Fig. 3, we show that the polariton RC works extremely well also when the resolution of the input image is further reduced to only 4×4 pixels. The logistic regression, performed on the nominal 4×4 input dataset, gives a recognition rate of 81.6%. Once the signal has been processed by the polariton array, the efficiency increases up to 86.3%, showing that polariton nonlinearities allow an increment of 4.7% over the limit set by logistic regression. This is possible by repeating the same testing and learning phases with different random matrices, as shown in Fig. 3. The combination of different realizations of the nonlinear array effectively increases the size of the reservoir, allowing to decrease the error rate. In Fig. 3, an improvement of $\approx 3\%$ is measured after combining six low-resolution outputs. An equivalent result is obtained on a single repetition by adopting a higher output resolution (see SI).

Finally, we highlight the importance of using hybrid light-matter particles to simultaneously obtain high nonlinearity and efficient connectivity between nodes. Indeed, the nonlinear behaviour of polariton nodes can be simulated by a nonlinear calibration of the SLM transfer function and redirecting its output to the camera. However, even though the success rate is slightly increased over the linear classifier, in this case the result is still lower than that obtained by processing the information in the polariton RC. This is due to the lack of efficient connections between nodes in the SLM array, which are instead present in the polariton network. In Fig. 4, we show that the hopping of polaritons between nodes provides the required connectivity. In Fig. 4b, all the nodes except the central one, marked with A, are pumped using an almost homogeneous SLM pattern. When the intensity of the surrounding nodes is increased, the switched-off node A is activated as well, reflecting the intensity redistribution in the array

Figure 3: Polariton RC on low resolution images. Linear transformation into a 8×8 matrix is used to distribute the input digit (4×4 resolution) on the whole array by multiplication of the 16 input intensities by a 16×64 matrix. The success rate obtained from logistic regression, both on the 4×4 input digits (samples of 1, 3, 5, 7 input digits are shown in the top row on the left) and on the 8×8 array (second row on the left), is 81.6%. The signal after the nonlinear transformation performed by the polariton RC (bottom row on the left, where the quadratic deviation from the mean is shown to highlight the nonlinear redistribution of the information) allows an increase of the recognition rate up to 83.6% already with 64 nodes. The spatial scale of the polariton RC in the bottom row P is $150 \mu\text{m} \times 150 \mu\text{m}$. On the right, the combination of the information obtained from 6 different random masks allows to increase the effective reservoir size and further enhance the recognition rate to 86.3% (dashed line is a guide to the eye), which is better than logistic regression by a 4.7%.

(blue line in Fig. 4a). Remarkably, connections are not limited to first neighbours but extend farther in the polariton network. In Fig. 4c, also first neighbour nodes A' are switched off and node A can be activated only by the second neighbour nodes (indicated as A'' in Fig. 4). The red line in Fig. 4a shows that nodes A'' contribute to A with a connection strength of about one half as compared to nodes A' .

Moreover, the redistribution dynamics within the network takes place on a timescale comparable to the polariton lifetime. As shown by simulations (see SI), the steady state is reached

Figure 4: Connectivity between nodes. **a**, Comparison between the output intensities of node A activated by the contributions from first and second neighbour nodes A' and A'' (blue line), and from second neighbour nodes A'' only (red line). **b**, Region of the network with pumping power at node A intentionally set to $P=0$. **c**, Same as in **b**, but with power set to $P=0$ as well at the first neighbour nodes A' . Node A is activated also in this case with a transmitted intensity of about $1/3$ with respect to the configuration shown in **b**.

in this case after about 50 ps. Even though the present experiments use a continuous wave pump, we anticipate that the recognition rate can be further improved by following the transient dynamics of the system (see the section on time multiplexing in SI).

In conclusion, neuromorphic computing in a physical network of multiple connected, non-linear nodes is demonstrated. Up to now, all optical implementations of artificial neural networks have been limited by weak nonlinearities, mostly without overcoming the efficiency of the bare linear classification applied to the raw input data. This lack of fast, energy efficient non-linearity is one of the most challenging obstacles to be overcome before photon-based neural networks can outperform standard computers. We show that exciton-polaritons are a suited system thanks to their mixed light-matter components: polariton-polariton interactions brings the

desired nonlinearities while the photonic component assures the connectivity between nodes and high operational speeds. We measure the recognition efficiency on the MNIST dataset, finding that the error rate is lower than in any other physical implementation used so far. We note that a further increase of the recognition rate can be obtained in polariton networks by exploiting the spin degree of freedom and the phase difference between nodes [33, 34]. With larger polariton arrays, success rates comparable to software implementations of reservoir models can be achieved with present technologies, moving towards a realistic integration of artificial intelligence based on optical systems.

References

- [1] Misra, J. & Saha, I. Artificial neural networks in hardware: A survey of two decades of progress. *Neurocomputing* **74**, 239 – 255 (2010).
- [2] Serb, A. *et al.* Unsupervised learning in probabilistic neural networks with multi-state metal-oxide memristive synapses. *Nature Communications* **7**, 12611 (2016).
- [3] Hu, S. *et al.* Associative memory realized by a reconfigurable memristive hopfield neural network. *Nature Communications* **6**, 7522 (2015).
- [4] Prezioso, M. *et al.* Training and operation of an integrated neuromorphic network based on metal-oxide memristors. *Nature* **521**, 61–64 (2015).
- [5] Sengupta, A., Banerjee, A. & Roy, K. Hybrid spintronic-cmos spiking neural network with on-chip learning: Devices, circuits, and systems. *Phys. Rev. Applied* **6**, 064003 (2016).
- [6] Shen, Y. *et al.* Deep learning with coherent nanophotonic circuits. *Nature Photonics* **11**, 441–446 (2017).

- [7] Maass, W., Natschlger, T. & Markram, H. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation* **14**, 2531–2560 (2002).
- [8] Jaeger, H. & Haas, H. Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science* **304**, 78–80 (2004).
- [9] Vandoorne, K. *et al.* Experimental demonstration of reservoir computing on a silicon photonics chip. *Nature Communications* **5**, 3541 (2014).
- [10] Brunner, D., Soriano, M. C., Mirasso, C. R. & Fischer, I. Parallel photonic information processing at gigabyte per second data rates using transient states. *Nature Communications* **4**, 1364 (2013).
- [11] Torrejon, J. *et al.* Neuromorphic computing with nanoscale spintronic oscillators. *Nature* **547**, 428–431 (2017).
- [12] Du, C. *et al.* Reservoir computing using dynamic memristors for temporal information processing. *Nature Communications* **8**, 2204 (2017).
- [13] Opala, A., Ghosh, S., Liew, T. C. & Matuszewski, M. Neuromorphic computing in ginzburg-landau polariton-lattice systems. *Phys. Rev. Applied* **11**, 064029 (2019).
- [14] Carusotto, I. & Ciuti, C. Quantum fluids of light. *Rev. Mod. Phys.* **85**, 299–366 (2013).
- [15] Walker, P. M. *et al.* Ultra-low-power hybrid light-matter solitons. *Nature Communications* **6**, 8317 (2015).
- [16] Leyder, C. *et al.* Interference of coherent polariton beams in microcavities: Polarization-controlled optical gates. *Phys. Rev. Lett.* **99**, 196402 (2007).

- [17] Amo, A. *et al.* Excitonpolariton spin switches. *Nature Photonics* **4**, 361–366 (2010).
- [18] Marsault, F. *et al.* Realization of an all optical exciton-polariton router. *Applied Physics Letters* **107**, 201115 (2015).
- [19] Ballarini, D. *et al.* All-optical polariton transistor. *Nature Communications* **4**, 1778 (2013).
- [20] Zasedatelev, A. V. *et al.* A room-temperature organic polariton transistor. *Nature Photonics* **13**, 378–383 (2019).
- [21] Dreismann, A. *et al.* A sub-femtojoule electrical spin-switch based on optically trapped polariton condensates. *Nature Materials* **16**, 1074 – 1078 (2016).
- [22] St-Jean, P. *et al.* Exciton-polariton topological insulator. *Nature Photonics* **11**, 651656 (2017).
- [23] Klemmt, S. *et al.* Exciton-polariton topological insulator. *Nature* **562**, 552556 (2018).
- [24] Ohadi, H. *et al.* Spin order and phase transitions in chains of polariton condensates. *Phys. Rev. Lett.* **119**, 067401 (2017).
- [25] Berloff, N. G. *et al.* Realizing the classical xy hamiltonian in polariton simulators. *Nature Materials* **16**, 1120–1126 (2017).
- [26] Wang, Q., Li, Y., Shao, B., Dey, S. & Li, P. Energy efficient parallel neuromorphic architectures with approximate arithmetic on fpga. *Neurocomputing* **221**, 146 – 158 (2017).
- [27] Tamir, D. E., Shaked, N. T., Wilson, P. J. & Dolev, S. High-speed and low-power electro-optical dsp coprocessor. *J. Opt. Soc. Am. A* **26**, A11–A20 (2009).
- [28] Caputo, D. *et al.* Josephson vortices induced by phase twisting a polariton superfluid. *Nature Photonics* **13**, 488–493 (2019).

- [29] Baas, A., Karr, J. P., Eleuch, H. & Giacobino, E. Optical bistability in semiconductor microcavities. *Phys. Rev. A* **69**, 023809 (2004).
- [30] Milićević, M. *et al.* Edge states in polariton honeycomb lattices. *2D Materials* **2**, 034012 (2015).
- [31] Lin, X. *et al.* All-optical machine learning using diffractive deep neural networks. *Science* **361**, 1004–1008 (2018).
- [32] Jiang, W. *et al.* Physical reservoir computing built by spintronic devices for temporal information processing. *arXiv:1901.07879* (2019).
- [33] Colas, D. *et al.* Polarization shaping of poincar beams by polariton oscillations. *Light: Science & Applications* **4**, e350 (2015).
- [34] Ohadi, H. *et al.* Nontrivial phase coupling in polariton multiplets. *Phys. Rev. X* **6**, 031032 (2016).

Acknowledgments

The ERC ”ElecOpteR” grant N. 780757 is acknowledged. S.G. and T. L. were supported by the Ministry of Education (Singapore), grant numbers MOE2017-T2-1-001 and MOE2018-T2-02-068. A.O. acknowledges support from the National Science Center, Poland via Grant No. 2016/22/E/ST3/00045. M.M. acknowledges support from the National Science Center, Poland via Grant No. 2017/25/Z/ST3/03032, under the QuantERA program. We thank R. Houdré for growing the sample.

Data availability: The raw experimental and numerical data used in this study are available from the corresponding author upon reasonable request.