

RESEARCH ARTICLE | JUNE 11 2026

Experimental investigation of time series classification using a self-pulsing microring resonator network

Alessandro Foradori  ; Alessio Lugnan  ; Lorenzo Pavesi  ; Peter Bienstman 

APL Photonics 11, 066111 (2026)

<https://doi.org/10.1063/5.0329322>

Articles You May Be Interested In

"Möbius" microring resonator

Appl. Phys. Lett. (March 2019)

Accurate resonant frequency spacing of microring filters without postfabrication trimming

J. Vac. Sci. Technol. B (December 2006)

Enhancing noise robustness in SiN microring photonic reservoirs enabled by 2D material nonlinearities

 AIP
Publishing

APL Photonics

Special Topics Open for Submissions

[Learn More](#)

Experimental investigation of time series classification using a self-pulsing microring resonator network

Cite as: APL Photon. 11, 066111 (2026); doi: 10.1063/5.0329322

Submitted: 17 February 2026 • Accepted: 26 May 2026 •

Published Online: 11 June 2026



Alessandro Foradori,^{1,2,a)}  Alessio Lugnan,²  Lorenzo Pavesi,²  and Peter Bienstman¹ 

AFFILIATIONS

¹ Photonics Research Group, Department of Information Technology, Ghent University-imec, Technologiepark Zwijnaarde 15, Ghent 9052, Belgium

² Nanoscience Laboratory, Department of Physics, University of Trento, Via Sommarive 14, Trento 38123, Italy

^{a)} Author to whom correspondence should be addressed: alessandro.foradori@ugent.be

ABSTRACT

Photonic neuromorphic computing offers compelling advantages in power efficiency and parallel processing, but it often falls short in realizing scalable nonlinearity and long-term memory. These limitations can be overcome by silicon microring resonator (MRR) networks. These integrated photonic circuits enable compact, high-throughput neuromorphic computing by simultaneously exploiting spatial, temporal, and wavelength dimensions. This work provides an in-depth study of MRR networks for photonics-based machine learning. We investigate the system's effectiveness on two widely used image classification benchmarks, MNIST and Fashion-MNIST, by encoding images directly into time sequences. In particular, we enhance the computational performance of a linear readout classifier within the reservoir computing paradigm through the strategic use of multiple physical output ports, diverse laser wavelengths, and varied input power levels. Moreover, we explore a single-pixel classification setting, where inference does not require digital memory, thanks to the inherent memory and parallelism of our MRR network.

© 2026 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0329322>

I. INTRODUCTION

Within the evolving field of artificial intelligence (AI), neuromorphic computing systems are gaining attention as hardware platforms that emulate the distributed, nonlinear functions of the biological brain. Among these, photonics-based neuromorphic architectures emerge due to their inherent advantages. In fact, they offer high throughput, low latency, energy-efficient linear operations and, in general, high parallelism by exploiting multiplexing techniques across multiple optical degrees of freedom, such as the optical wavelength or polarization.^{1,2} These properties make neuromorphic photonic solutions not only interesting for accelerating large-scale artificial neural networks (ANNs) in data centers,³ but also attractive for compact, low-power edge computing applications, specifically when information is natively encoded in the optical domain, as in optical communications^{4,5} and in a broad variety of optical sensing implementations.^{6–8}

However, optical computing systems face challenges in providing energy-efficient scalable nonlinear nodes,² required to implement artificial neurons and, in general, to host powerful machine learning (ML) models. For example, hybrid nonlinear nodes for neuromorphic computing based on optical-to-electrical-to-optical (OEO) conversion have been introduced,⁹ at the cost of increased complexity and energy consumption, which hinder scalability to large numbers of artificial neurons. In addition, several edge computing applications require efficient and low-latency time-dependent signal processing, which would benefit from volatile or easily resettable optical memory effects.¹⁰ Although this is relatively easy to obtain for memory durations on the nanosecond scale or shorter, for example, via optical delay lines,¹¹ it has traditionally been difficult for longer timescales, such as in the millisecond to second range typical of most time-dependent sensing applications. All-optical non-volatile memory has been introduced in integrated photonics by employing phase-change materials (PCMs), and their

application for neuromorphic computing has been proposed in several works.^{12–14} However, PCM-based all-optical memory requires switching the PCM with optical pulses, introducing relatively high power losses and difficulties in resetting the memory states in large optical networks.¹⁵

Recently, we demonstrated that photonic integrated networks of silicon microring resonators (MRRs) can provide all-optical nonlinearity and volatile long-term memory,^{16,17} with high-throughput production of nonlinear dynamical representations of the input signal. These representations can be exploited for computationally efficient ML¹⁸ based on the *reservoir computing* (RC) approach,^{19,20} where a recurrent neural network (on a photonic circuit in our case) is randomly initialized and only a linear readout is trained to carry out a time-dependent ML task.

MRRs are versatile and widely used photonic devices, whose applications for neuromorphic computing have been proposed in several research works and diverse ways.^{15,17,21–25} Traditionally, MRRs function as wavelength-selective filters, exploiting their resonance properties to control the spectral components of optical signals. In the context of neuromorphic computing, MRRs have often been proposed as linear nodes implementing optically weighted connections.^{9,21,22} In this case, an MRR simultaneously extracts and weights a selected optical wavelength from a waveguide. Therefore, multiple MRRs can be applied to a single waveguide to implement the weighted multi-channel input of an optical artificial neuron with improved scalability. The weights assigned to MRRs can be controlled by tuning their resonance conditions. This can be achieved through several mechanisms, such as via a local heater exploiting the thermo-optic effect, using an integrated pn-junction for electro-optic modulation, or incorporating PCMs.

In addition, a silicon MRR can easily be driven in a nonlinear regime by raising the power of the resonant input light to the milliwatt range.^{25,26} Such a nonlinearity is due to the free carriers and heat accumulated in the silicon ring waveguide (mediated by the two-photon absorption effect), which change the ring refractive index, in turn producing a shift in the MRR resonant wavelength. Interestingly, the nonlinear effect based on free carriers blueshifts the MRR resonance and is faster (lifetime around 1–45 ns) but weaker than the thermo-optic nonlinearity (lifetime around 60–280 ns), which produces a resonance redshift instead. It is important to note that, because of the different lifetimes of these two competing effects, a nonlinear MRR exhibits volatile memory with two different timescales. Although these effects are present in any silicon waveguide, in a MRR they are amplified and coupled together by its optical resonance, resulting in a much more energy-efficient and dynamic nonlinearity and memory.^{21,27} Thus, MRRs can be used as efficient photonic neurons. For example, the interplay between a MRR's free carriers and temperature under suitable conditions gives rise to self-pulsing behavior, where a slowly varying optical input is translated into a pulsed output signal.²¹ Leveraging this property, several works have demonstrated that MRRs can be employed as artificial spiking neurons that are compact and simple to fabricate.^{16,25,27–30}

In this work, we expand the investigation of a MRR network for photonics-based ML by testing our system on two popular image classification benchmark tasks, leveraging the multiple responses of

the system across different wavelength and spatial ports. It should be stressed that, even if the MRR network is fixed after fabrication, it can produce a large number of different data representations by varying input wavelength, power, and output port. Therefore, the user can easily explore such a configuration space (given that RC systems allow for very fast and efficient training) and select those representations that are most suitable to learn a target ML task, which can be seen as a sort of *train-to-train* procedure. In addition, we study the case where the classification is performed from a single time sample, which eliminates the need for digital memory. Finally, our system and the related performances are compared with those of comparable digital ANNs to quantify the advantages presented by the proposed approach over conventional ML systems. Specifically, the original contributions of this work present the following differences with respect to our previous studies on MRR networks:^{17,18}

- Although it appears similar, the MRR network employed here is physically and functionally quite different with respect to the one employed in Ref. 18, as its long-term memory does not originate from lossy PCMs but, in contrast, comes from complex dynamics due to coupled self-pulsing in multiple MRRs (not present in Ref. 18).
- As opposed to Ref. 18, here we investigate how the ML accuracy changes by varying multiple key parameters, namely, the number of output network ports, the encoding modality, and the input sample rate. Moreover, we consider an additional ML task (Fashion-MNIST in addition to MNIST) and study the network's functional memory by employing a single time sample per network output for ML.
- The full ML datasets are here employed with minimal preprocessing, which is easily obtained without using digital processing but just by means of an optical modulator.
- We present a comparison with digital ANN counterparts reaching similar accuracies in terms of computational cost.
- In Ref. 17, the same MRR network considered here is employed to demonstrate a fundamentally different ML approach, that is, reservoir computing with non-volatile memory, aimed at exploiting the network's output much after the end of the input signal insertion. For this purpose, basic ad-hoc ML tasks were tackled, which are much simpler than the ones considered in this work, which is instead based on conventional reservoir computing.

Moreover, the scope of the mentioned previous studies was to demonstrate novel systems and ML approaches. This article instead, given the growing interest in nonlinear MRRs for neuromorphic photonics, aims to provide the community with a robust and in-depth analysis of the ML capabilities of these devices when coupled together into a network.

The rest of this paper is structured as follows: Sec. II, Principles of Operation, describes the physics of the optical system, specifically its response to a constant signal (Sec. II A), and explores how this behavior can be exploited as a reservoir for machine

learning tasks like classifying images encoded as time-dependent signals (Sec. II B). This section also discusses the dimensionality expansion available in space, wavelength, and power (Sec. II C), as well as the considered machine learning tasks (Sec. II D) and the specifics of input signal encoding, including sampling rate and power offset (Sec. II E). Section III, Results, presents our findings, covering classification using full time-series representations across multiple ports (Sec. III A), including a comparison with digital ANNs, and classification using single-pixel representations with optical memory (Sec. III B). Finally, the Conclusion, Sec. IV, summarizes our work, while the Methods, Sec. V, provides a detailed description of our experimental setup and procedures.

II. PRINCIPLES OF OPERATION

Reservoir Computing (RC) is a machine learning paradigm where a fixed, untrained recurrent neural network transforms an input signal into a high-dimensional representation, which is then fed to a trained linear regressor or classifier. It is particularly suitable for hardware implementation, since any complex enough, nonlinear, and dynamical physical system can be employed as a reservoir, confining digital computation and training to a computationally cheap readout linear combination.

To function effectively as a reservoir, a physical system must exhibit nonlinearity and memory, and it must have the ability to expand the dimensionality of its input.²⁰ In this context, the network of coupled microring resonators (MRRs) proposed here meets these requirements. Its nonlinear behavior, memory, multiple accessible output ports, and wavelength multiplexing capability make it a promising candidate for a compact integrated photonic reservoir with high throughput.^{17,24}

Our system consists of an 8×8 matrix of coupled MRRs (Fig. 1, details in Sec. V A). Each MRR is coupled to an upper and lower straight waveguide, which is terminated by grating couplers. These function as input or output ports, as they allow users to insert and extract light via optical fibers. In particular, the grating coupler on the left in the schematics of the circuit in Fig. 1 is used as input, while the ones on the right are used as output.

A single MRR acts as a filter around a resonant frequency, exhibiting a Lorentzian shape in its spectrum. Ideally, when coherent light is injected into the waveguide on the top left side of a ring (input port), part of it can pass directly through to the top right side (through port) or couple into the ring and exit at the bottom left side (drop port), depending on how closely the optical frequency matches the MRR resonant frequency. In particular, if the light is out of resonance, it will not couple with the MRR, and thus will continue unperturbed to the through port. On the other hand, the more the incoming light is resonant, the more it will be redirected backward through the opposite straight waveguide (part of it will be absorbed by the MRR) and the less it will proceed past the MRR (we refer the reader to Ref. 21 for more in-depth explanations).

The resonant frequency of an MRR is fundamentally determined by its material properties and geometric parameters, such as the effective refractive index and the optical path length. However, the resonance exhibits sensitivity to temperature variations because




FIG. 1. Schematics of the MRR network. The network consists of 64 silicon microring resonators, coupled by straight waveguides and linked to one optical input port (left) and multiple output ports (right) consisting of grating couplers. The output ports used for signal acquisition (ports 1, 2, 3, 5, 7, and 9, counted from top to bottom) are highlighted in green; the remaining gray ports were not used due to a low signal-to-noise ratio. The inset shows that light enters the microring resonator from the input port, passes either directly to the through port or couples into the ring and exits at the drop port, depending on whether the input optical frequency matches the resonant frequency.

of the thermo-optic effect, inducing changes in the refractive index. We control the temperature with a Peltier cell, which stabilizes the chip temperature within a range of 0.1 K. In addition, fabrication tolerances introduce slight differences in the optical path and thus in the resonant frequency of MRRs, even when they are designed to be nominally identical. Each ring resonance is thus fixed by construction but is randomly different from the others. To mitigate the impact of these variations, we designed MRRs with a relatively low quality factor, that is, with a broader resonance linewidth, to ensure spectral overlap, and thus coupling, between different rings [Fig. 2(a)].

It should be stressed that the scalability of this type of MRR network in terms of the number of coupled nodes (ensuring output readability) depends on the nonlinear dynamics originated by a certain optical input. Indeed, thanks to the nonlinear effects in MRRs, the ring resonances hindering a given optical path can be moved by a powerful enough input so that a readable output is obtained at an output port. For this reason, it is difficult to predict how large employable MRR networks can be. In addition to the effective network topology at a certain wavelength and input power, key factors are the Q factors of the MRRs (resonances of MRRs with a higher Q factor are more easily “moved” by an optical input), the sensitivity of the employed photodetectors, and the time characteristics of the input signal. Nevertheless, we expect that MRR networks quite larger (e.g., with double the size) than the one employed here could be used with similar MRR Q factors and input power levels, as is also evidenced in a previous study.¹⁸ Beyond this, higher MRR Q factors can enable the employment of much larger networks (or combinations thereof), as well as the usage of on-chip optical amplifiers.




FIG. 2. Response to a constant signal. (a) Spectrum of port 1 acquired at low input power, in the absence of nonlinear effects. The x-axis (laser frequency) aligns with that used in the self-pulsing map. (b) Color-map of the self-pulsing estimated period at the first output port of the network, as a function of input power and optical (laser) frequency. The pink region indicates a response without self-pulsing, while green tones correspond to different self-pulsing periods. Cells marked with * denote those responses for which we could not detect a periodicity (see Sec. V C for details). Two representative signals, with periods of 49 and 1.4 μ s, are displayed on the left. Each signal is colored according to its repeating pattern, which visually represents the estimated period (also indicated by a dashed sinusoid).

A. Response of the MRR network to a constant signal

A sufficiently high input optical power (above roughly 1 mW) temporarily and significantly affects the resonance frequency of an MRR via silicon nonlinear effects, due to changes in free carrier density and the temperature of the ring waveguide.²¹ In an MRR network, high enough power can excite these nonlinear states in multiple rings simultaneously.²⁶ Crucially, this phenomenon depends on the input power and on the detuning between the optical frequency of the input signal and the MRRs' resonance frequency. Moreover, different nonlinear responses can also be observed at different output ports, as the light follows distinct propagation paths and couples to different resonators along the network. Therefore, by employing different output ports and by varying the input laser power and optical frequency, our system can produce a large number of different nonlinear transformations of its input. Then, a suitable subset of representations can be selected to tackle a target ML task.

When a constant input is applied to the system, the output does not necessarily remain steady. Instead, the system can exhibit dynamic behaviors characterized by non-constant outputs. Depending on the specific input condition and system parameters,

the output can show periodic oscillations or more complex, non-periodic fluctuations.¹⁶ These dynamic responses are referred to as self-pulsing.

In Fig. 2(d), the estimated period of the self-pulsing oscillations at the first output port of the system is shown as a function of the constant input power and optical frequency, in those cases where some periodicity could be detected (see Sec. V C in the Methods for further details). The variety of detected periods, spanning over a relatively broad range (from around 1 to tens of microseconds), highlights the richness and diversity of the network's response.

B. Classification of images encoded to time-dependent signals

This variety of responses can be exploited for practical purposes, as we show here. In particular, the system's nonlinearity and memory make it suitable for RC. For example, as shown in Fig. 3, these dynamics can be used to increase the dimensionality of time-dependent input sequences. A handwritten digit from the MNIST dataset (a popular ML benchmark dataset comprising ten classes of 28×28 images³¹) is encoded as a time sequence and injected into the




FIG. 3. Multiple representations of images encoded to time-dependent signals. Each MNIST digit is first converted into a time sequence and then processed optically through the reservoir. The resulting temporal responses vary depending on the optical frequency and output port. These distinct output signals are reconstructed into images, illustrating the diversity of representations generated by the system.

MRR network (more details can be found in Sec. V B in the Methods). The resulting output time sequences provide diverse nonlinear representations (with memory) of the input sequence, thus performing a dimensionality expansion. This transformation projects the input into a high-dimensional space where features could become more easily linearly separable. Such behavior aligns with the principle of reservoir computing, where the system dynamics create a fixed recurrent network and only a simple trainable readout layer is required for classification.

C. Dimensionality expansion

Leveraging the system's capabilities to produce different responses, the same input signal can be processed under different conditions. In particular, we process the input by acquiring outputs from different ports, at multiple optical frequencies and input power levels, in order to map it to a higher-dimensional space.

1. Physical ports

The output signals are measured at ports 1, 2, 3, 5, 7, and 9, counting from top to bottom on the right side with respect to Fig. 1, while the other output ports present a too low signal-to-noise ratio (SNR). In principle, measurements from all selected ports could be acquired simultaneously. However, due to practical limitations in the current experimental setup, these measurements are performed sequentially.

2. Optical frequencies

Different optical frequencies, resonant with some MRRs in the network, are employed for signal acquisition. In particular, we use ten equally spaced frequencies (unless otherwise stated) ranging from 192.68 to 192.86 THz. This frequency range corresponds to a region where the spectral content at port 1 is significantly rich, as shown in Fig. 2(a).

3. Input power levels

The system is also operated at different input power levels to modulate the network's nonlinearity. Because of the high complexity of the dynamics, the network's output signals can significantly change even for moderate variations of input power, as can be appreciated in Fig. 2(d). Average on-chip input powers are varied within the range of 0.7–7.0 mW for self-pulsing characterization, and within 1.1–5.5 mW for the machine learning task. For machine learning applications, 5 (unless otherwise stated) equidistant power levels are employed, allowing the system to explore a wide set of dynamic behaviors and enhancing its processing versatility.

4. Baseline at out-of-resonance frequency

An additional measurement on the first output port, with an optical frequency of 192.60 THz and maximum input power, is performed as a baseline reference. At this frequency, the light is not coupled to the MRRs and is directly conveyed to the first output port with negligible nonlinear effects. Therefore, this measurement allows us to check what the achieved ML performance is when the MRR network is not used, thus providing a direct estimate of the advantage brought by our photonic system. For this out-of-resonance measurement, the maximum input power is employed (in this case, the response is linear and thus does not depend on the power level) in order to maximize the output SNR.

D. ML tasks

In this work, we focus on the classification tasks provided by the MNIST and Fashion MNIST datasets, which are collections of small-size grayscale images (28×28 pixels) of handwritten digits and of clothes, respectively, each comprising ten classes.^{31,32} The training set includes 60 000 images, while the test set contains 10 000 images, following the standard train/test split conventionally adopted. The choice of these datasets is motivated by several complementary considerations.

First, they constitute widely recognized benchmarks in the machine learning community and in the emerging field of neuromorphic photonics. Moreover, the transformation operated by the physical system on the input data is directly visualizable. Indeed, since the raw temporal sequences generated by the reservoir are abstract and difficult to interpret, the spatial structure of the input data is leveraged to reconstruct the processed signals back into the image domain. This approach allows for a qualitative appreciation of the nonlinear dynamics and memory effects introduced by the MRR network.

Second, the considered tasks are not a spatial pattern recognition problem. Each image is serialized into a one-dimensional time series by concatenating pixel values sequentially. Thus, the classification target is not an image but a time series, fully compatible with the recurrent, time-dependent processing paradigm of reservoir computing.

Third, while tasks such as parity (XOR of the previous N bits) and delayed memory recall (recovery of the N th bit in the past) are standard benchmarks in the photonic reservoir computing literature, they probe a specific and relatively constrained form of memory. Image classification processed as a time series demands the integration of information across 784 sequential inputs,

exercising the memory of the reservoir in a richer and more complex way.

The MRR network operates as a photonic reservoir whose output is collected across different ports, frequencies, and input powers. These states are then fed to a linear classifier in software, which acts as a readout layer to perform the final classification.

E. Input signal and data acquisition

To process the MNIST and Fashion-MNIST datasets in the photonic reservoir system, we first define an appropriate encoding scheme. Both datasets consist of grayscale images with a resolution of 28×28 pixels and 8-bit intensity levels, ranging from 0 (typically corresponding to the background in the image) to 255 (representing the strokes of the handwritten digits or of the dress).

Each image is flattened into a one-dimensional sequence of 784 points by reading pixel values row by row. This transformation produces a temporal sequence where each time sample corresponds to the intensity of a single pixel, with the encoding strategy determining the relationship between grayscale intensity and optical output [Fig. 4(a)]. In particular, we employ the following encodings that are straightforward to achieve with a simple intensity modulator, without the need for further preprocessing.




FIG. 4. Encoding and processing of image data by the photonic network. (a) Overview of the data encoding and acquisition pipeline. A grayscale image from the MNIST or Fashion-MNIST dataset (28×28 pixels) is first flattened row by row into a 784-point sequence. This sequence is converted into a time series and used to drive an Arbitrary Waveform Generator (AWG), which modulates infrared laser light. The resulting optical signal is then injected into the MRR network, where it is optically processed. The output is detected by a photodetector and sampled by an oscilloscope (see Sec. V B in Methods for details). To facilitate intuitive insight into the optical transformation, the acquired temporal signal is reshaped back into an image for a better visual comparison. (b) Time-encoded transmission of each image is followed by a brief pause (on the order of $10 \mu\text{s}$, Table V in Methods), allowing the MRR network to return to its resting state. This ensures that the system's response to each input is independent and unaffected by residual dynamics from previous samples. (c) List of combinations of encoding modes (normal, inverse, and box), feature durations (i.e., the time step of the AWG), and datasets (MNIST and Fashion-MNIST).

1. *Normal*: A background pixel value of 0 corresponds to the minimum achievable optical power, while a value of 255 corresponds to the maximum achievable optical power. The background typically carries negligible optical power, and the nonlinear transformations only occur in the short pulses that correspond to the strokes of the image.
2. *Inverse*: This encoding inverts the image, such that a pixel value of 0 corresponds to the maximum optical power, and 255 corresponds to the minimum optical power. In this configuration, the strokes effectively act as interruptions of light insertion and, as a consequence, of light transmission through the MRR network.
3. *Box*: The minimum pixel value is mapped to 40% of the highest achievable optical power. Values between 0 and 255 are scaled accordingly within this restricted range. The background acts as a pump to the MRR network, possibly driving it into a nonlinear and self-pulsing state, which is then perturbed by the strokes.

Each feature (pixel) of the input sequence corresponds to a specific time step with a fixed duration. In this work, we consider time step durations of 2, 10, and 20 ns, allowing us to explore which timescale is most effective for exploiting the self-pulsing behavior of the microring resonators. These values are not chosen arbitrarily: the self-pulsing period of the device, operating in a non-chaotic regime, lies in the range of $\sim 1\text{--}50\ \mu\text{s}$, and the total duration of a single input pattern (784 samples per image) spans 2, 10, and $20\ \text{ns} \times 784 \approx 1.6, 8, \text{ and } 16\ \mu\text{s}$, which is commensurate with this period. Visual representations of the combinations of datasets, encodings, and feature durations used are provided in Fig. 4(c).

The complete input sequence is generated by modulating infrared laser light using an arbitrary waveform generator (AWG), which is programmed to continuously reproduce the entire dataset in a sequential loop. Each image is transmitted as a time-encoded sequence, followed by a brief pause to allow the microring resonator system to relax and return to its unexcited state before processing the next image [Fig. 4(b)]. This pause ensures that the system's response to each input remains independent and is not influenced by residual dynamics from previous samples.

The MRR network's output signal is received by a photodetector and converted into an electrical signal acquired by an oscilloscope. The sampling time of the oscilloscope is shorter than that of the AWG, thus increasing the number of features per image by a factor of ~ 1.5 or 2.5 , depending on the ratio between the AWG and oscilloscope sampling rates (see Sec. V B 1 in the Methods for further details).

The reservoir dynamics associated with the best classification performance can differ across time steps, reflecting the fact that each time step selects a distinct operating point along the input modulation timescale. The time step acts as a hyperparameter of the reservoir system: while it can in principle be tuned, the optimal value is task-dependent and no universal optimum exists.

III. RESULTS

In the first ML approach we consider in this work, we apply a linear readout classifier, as detailed in Sec. III A, to different

representations of the input time series (i.e., a flattened image) from different output ports, with fixed input power and frequency. Then, we present some comparisons with fully digital ANNs to highlight the decrease in computational cost allowed by our photonic implementation.

In the second ML approach (Sec. III B), we select a single time sample (corresponding approximately to a pixel of the original image). We then combine all output representations of this single sample, gathered across various optical frequencies and output ports but fixing the input power level, and we feed them to a linear readout classifier. This strategy leverages the intrinsic optical memory of the MRR network, thus removing the requirement for an external digital memory component.

In both cases, we evaluate the advantage brought by our photonic system via a comparison with suitable baseline results.

A. Linear readout on full time-series representations across multiple ports

To evaluate the network's computational capabilities, we encode each full flattened image as a time-domain signal and feed it into our MRR network. The system produces multiple parallel output signals, each of which represents a distinct feature projection of the input. These parallel representations can either be employed individually, by considering the signal from a single output port, or jointly, by combining signals across multiple ports. We assess classification performance under both scenarios, for a fixed optical frequency and input power. The aim of this analysis is to demonstrate how performance improvements are linked to both the internal dynamics of the photonic hardware and the combination of its representations. After measuring the responses of all ports, at first, each port's representation is independently used to train and test a linear classifier on the selected dataset. Next, we combine the representations from multiple ports to assess the benefits of exploiting different optical paths as reservoir nodes. Specifically, in this case, we use 2–6 ports, and we explore different combinations of available representations (the first port is always included to avoid that this exploration takes too much time). It should be noted that not all the acquired output signals are included in this analysis due to a too low SNR. Moreover, it should be stressed that for this first ML approach, each classification is made considering output signals that can be easily acquired in parallel just by employing a single input laser frequency and using a photodetector for each output port.

1. Classification of hand-written digits: MNIST

In Fig. 5, we present the best classification accuracies obtained for the MNIST dataset using a single linear classifier, while a complete set of final results for the different encodings is presented in Table I. Each result is compared to a baseline, which is calculated by applying the same linear classifier directly to the original dataset, without any photonic processing (see Sec. V D 2 in the Methods for more details). The accuracy of this baseline model is 92.03%. For each combination of optical frequency and input power, we report in Fig. 5(b) the highest accuracy achieved across all the output port combinations in the inverse configuration with a 20 ns timestep. In




FIG. 5. Classification performances on the MNIST dataset using representations from multiple ports. (a) Schematic example of the multiport classification strategy. Representations acquired from different output ports, under fixed optical frequency and input power, are stacked together and fed into a single linear classifier. (b) Heatmap showing the best test classification accuracy obtained for each combination of laser frequency and input power, using the MNIST dataset with inverse encoding and a 20 ns AWG time step. For each pair, the number of ports used can range from one up to six (all available ports). If multiple ports are used, the first port is always included. The colormap indicates the test accuracy improvement over a baseline linear classifier trained on the raw (non-photonic) data, whose accuracy is 92.03%. (c) Boxplots showing test accuracy as a function of the number of combined ports, evaluated for different system configurations. For each case, the optical frequency and input power providing the highest accuracy are chosen. A dashed line connects the highest accuracy (colored point) achieved at each port count, highlighting both the improvements given by combining multiple ports and the overfit that occurs with an excessive number of ports. Outliers are shown as empty points.

TABLE I. Classification accuracy on the MNIST dataset using a single linear classifier applied to different combinations of output ports. The table presents the best test accuracy (fourth column) achieved for a given configuration of encoding (first column) and timestep duration (second column). The number of ports for which the maximum is obtained is shown in the third column, while the best accuracy obtained using a single port is in the fifth column and the one in the out-of-resonance condition is in the sixth column. The last row reports the test accuracy of the baseline linear classifier applied to the original (non-photonic) data.

Encoding	Duration (ns)	No. of ports	Accuracy (%)	Single port (%)	Out-of-resonance (%)
Normal	2	5	94.50	92.29	91.54
Normal	10	4	94.82	91.77	91.57
Normal	20	4	94.58	92.26	91.78
Box	10	3	95.77	93.30	91.77
Box	20	4	95.94	93.32	91.80
Inverse	10	5	94.93	92.99	91.93
Inverse	20	5	96.49	94.31	92.12
Software LC			92.03		

most cases, the photonic processing leads to a clear improvement in performance over the baseline. The out-of-resonance configuration shows that a purely linear transformation by the employed photonic circuit (even though non-idealities of the employed modulator, optical amplifier, and photodetector could introduce some simple nonlinearity and memory effects), that is, without engaging the nonlinear and memory effects of the MRR network, does not provide a significant performance advantage (up to 92.17% accuracy).

The obtained accuracy improvements are therefore attributed to the MRR network's dynamics, which perform a nonlinear mixing of the input features. This effect is particularly evident for the configuration (2.2 mW, 192.76 THz), where, by using only one output port, we achieve an accuracy of up to 94.32%. It should be stressed that, from the perspective of the data encoding, there is no dimensionality expansion provided by a single output port of the MRR network, since the number of temporal points (or pixels) is approximately the same for the input and for the output signals (for details, see Sec. III A 3). However, the dynamical transformation performed by our photonic system spreads the information about the handwritten digit all over the background of the image, which would otherwise have a constant pixel value and thus carry no information in its intensity values (see examples in Fig. 3). This effect corresponds to an effective dimensionality expansion of the information about the handwritten digit, thus resulting in a significant performance improvement.

In contrast, the small performance improvement (up to 92.26%) observed in the normal configuration, where the nonlinear mixing is confined to the short duration of the digit strokes, highlights the importance of the dynamics occurring in the background. This is evident when comparing it to the performance of the box (93.32%) and inverse configurations, where the background effectively carries information.

Moreover, the additional dimensionality expansion provided by combining multiple ports contributes to further performance gains. For instance, we achieve the best accuracy of 96.49% in the 20 ns inverse configuration with

(1.1 mW, 192.74 THz) when combining signals from five ports.

The boxplots in Fig. 5(c) further highlight that these accuracy improvements are consistently observed across different combinations of ports, confirming the robustness of the multi-port strategy. Although similar improvement trends are observed when using the other encoding schemes or sampling times, the overall performance improvements in these cases are generally less pronounced.

Finally, we note that the decrease in the best accuracy observed in the plots of Fig. 5(c) when using a higher number of ports is ascribed to overfitting.

2. Classification of clothing items: Fashion-MNIST

The Fashion-MNIST dataset is another widely used benchmark in machine learning, offering a more challenging classification task compared to the original MNIST dataset. Like MNIST, it consists of 28×28 grayscale images, but instead of handwritten digits, it contains images of various clothing items, such as shirts, sneakers, and coats, each assigned to one of ten classes.

Due to the increased complexity, Fashion-MNIST is significantly harder for simple models. For example, a linear classifier typically achieves a test accuracy of around 85.17% on Fashion-MNIST, compared to 92.03% on MNIST.

The approach relying on a single linear classifier and representations obtained from multiple ports yields only a modest improvement in accuracy (85.81%, Table II) over the baseline (85.17%) of the fully software-based linear classifier. This limited gain is ascribed to the noise introduced during the photonic transformation and the discretization constraints imposed by the oscilloscope's finite voltage resolution. These effects are further supported by the lower performance obtained using out-of-resonance representations, which reach a maximum of 83.84% in accuracy. Moreover, unlike MNIST, which consists mainly of high-contrast features (sharp edges between a white background and dark strokes), Fashion-MNIST includes features distributed on a grayscale with smaller variation between each other.

TABLE II. Classification accuracy on the Fashion-MNIST dataset using a single linear classifier applied to different combinations of output ports. The table presents the best test accuracy (fourth column) achieved for a given configuration of encoding (first column) and timestep duration (second column). The number of ports for which the maximum is obtained is shown in the third column, while the best accuracy obtained using a single port is in the fifth column and the one in the out-of-resonance condition is in the sixth column. The last row reports the test accuracy of the baseline linear classifier trained on the raw (non-photonic) data.

Encoding	Duration (ns)	No. of ports	Accuracy (%)	Single (%)	Out-of-resonance (%)
Normal	10	4	85.76	84.15	83.32
Normal	20	4	85.81	84.55	83.84
Software LC			85.17		

3. Comparison with digital ANNs

To contextualize the effectiveness of our neuromorphic photonic system, we compare the classification results presented in Secs. III A 1 and III A 2 with those obtained using purely digital ANNs. This comparison serves not to establish competitive accuracy benchmarks, but rather to situate our photonic implementation within the broader landscape of neural network models of comparable architectural complexity. We stress that competing with digital neural networks on standard classification benchmarks is explicitly outside the scope of this work. The goal is instead to demonstrate that a physical nonlinear photonic substrate is capable of meaningful computation under realistic hardware constraints, with a number of digital linear and nonlinear operations lower than that required by conventional digital ANNs. This comparison highlights the advantage offered by our MRR network in the trade-off between computational efficiency and accuracy, namely fewer linear and nonlinear operations compared to digital ANNs. While convolutional neural networks are better suited for image classification, for a fair comparison, we consider digital ANNs that, like our photonic network, are not specifically structured to carry out the considered ML tasks (namely, multilayer perceptrons).

For the digital ANNs, we consider multilayer perceptrons (MLPs), which, like our photonic network, are not specifically

structured to carry out the considered ML tasks. Although convolutional neural networks (CNNs) generally achieve superior performance on image classification tasks when implemented in software, a direct comparison with CNNs would be misleading for two reasons. First, in this work, the image data are presented to the photonic reservoir as temporally serialized sequences, thereby recasting what is nominally a spatial classification problem as a time-series classification task. Second, a one-dimensional convolutional layer operating over a span of N time steps is, in essence, a finite-memory function over the input features. In this sense, 1D-CNNs and reservoir computing systems share a common functional principle: both exploit a form of memory to encode temporal context into a fixed-dimensional representation. The distinction lies in whether that memory is implemented through learned convolutional filters or through the dynamics of a physical nonlinear system. Given this equivalence, the MLP represents a more appropriate and architecturally neutral reference for benchmarking purposes, as it imposes no implicit temporal structure on the input and thus reflects more faithfully the computational overhead required in the absence of task-specific inductive biases.

We evaluate the classification performance on the MNIST dataset. In Table III, we show the results obtained using different software machine learning architectures. The first one is a linear classifier (logistic regression, which is often used as a baseline test

TABLE III. Best classification accuracy on the MNIST dataset, comparing photonic and digital network architectures. The first three rows show results from fully digital networks trained on the original data, while the last two rows present results from a photonic readout trained using an inverse representation with a 20 ns timestep. The first column specifies the network type. The second column provides a detailed description of the network's architecture by listing the number of nodes for each layer (input – [hidden –] output). Note that the MRR network output is oversampled (1225 time samples) with respect to the original number of pixels (784). The remaining columns quantify the number of trainable parameters, the number of nonlinear operations, and the best test accuracy (eventually expressed in concise uncertainty notation, with the error given as the standard deviation over ten repeated measurements; see Sec. V D 2 in Methods) achieved for each model.

System	Layers	Trainable parameters	Nonlinear operations	Accuracy (%)
Linear classifier (LC)	784 – 10	7850	10	92.03
Multilayer ANN	784 – 77 – 10	61 225	87	96.07(10)
Multilayer ANN	784 – 120 – 10	95 410	130	96.53(7)
MRR net (1 output) + LC	1225 – 10	12 260	10	94.32
MRR net + LC	(1225 × 5) – 10	61 260	10	96.49

accuracy) trained with Stochastic Gradient Descent (SGD).³³ The second one is a feedforward neural network composed of an input layer (784 features), a hidden layer with a variable number of neurons with a rectified linear unit (ReLU) activation function, and an output layer with ten nodes.³³ The weights are trained using SGD (see Sec. V D 2 in Methods). For each configuration, we report the test accuracy along with the number of trainable parameters and the total number of software-based nonlinear operations required. These metrics serve as a reference point for comparison with our photonic reservoir computing approach based on the MRR network.

We first establish a baseline using a simple linear classifier (logistic regression). This model, when applied to the original MNIST data, achieves an accuracy of 92.03%. By processing the data with our MRR network using only a single port and then applying the linear classifier, we achieve a higher accuracy of 94.32%. Interestingly, this improvement was accomplished with only a slight increase in trainable parameters, due to the fact that we oversample the waveforms in the photonics case.

Then, we compare our system to conventional multilayer ANNs with one hidden layer. First, we configure a digital ANN to have a similar number of trainable parameters as our multipoint photonic system ($\sim 60 \times 10^3$). This ANN, with 77 neurons in its hidden layer, reaches an accuracy of 96.07(10)%. In contrast, our system achieved a higher accuracy of 96.49% with fewer nonlinear operations.

Next, we trained a more complex digital ANN to match our system's accuracy of $\sim 96.5\%$. To achieve this, the ANN required a hidden layer with 120 neurons, increasing the number of trainable parameters to 95×10^3 and its nonlinear operations to 130 (compared to just ten for our system's softmax readout layer). This highlights an advantage of our photonic approach: it delivers the same accuracy with a lower number of trainable parameters and nonlinear software operations.

B. Classification using single pixel representations with optical memory

In Sec. III A, we have shown that the nonlinear dynamics and memory provided by our MRR network can generate multiple representations of an input image, which are useful to improve the classification accuracy of a linear classifier. In this case, however, digital memory is still required for ML inference, since the linear classifier is applied to the recorded signals produced by the photonic network. In this section, we will instead study a case where no digital memory (of the signals to be processed) is required for ML inference, that is, when the linear classifier is directly applied to the MRR network's outputs across various optical frequencies and spatial ports at a specific time, corresponding to a single pixel of the processed image [Fig. 6(a)]. It should be stressed that this demonstration is not about solving image classification with high accuracy,




FIG. 6. Classification using a single time sample (or pixel). (a) A linear classifier is applied to a specific time sample (or image pixel) of the nonlinear representations generated by the MRR network. (b) Bar plot presenting the corresponding maximum test accuracies (photonic network, green bars) obtained for the considered classification tasks (x-axis, see Sec. II E). They are compared with their baseline counterparts (out-of-resonance, orange bars), obtained without processing from our MRR network (and completely digitally in the case of the blue bars). (c) Pixel-by-pixel accuracy in the case with the highest maximum classification accuracy, that is, the MNIST classification with *inverse* encoding and 10 ns time sample duration. (d) Corresponding classification accuracy, for the best-performing pixel (with index 264), as a function of the number of employed representations.

TABLE IV. Classification accuracy using a single time sample (or pixel). The table presents the best test accuracy (fourth column) achieved for a given dataset (first column) and a given configuration of encoding (second column) and timestep duration (third column). The accuracy obtained in the out-of-resonance condition is shown in the fifth column.

Task	Encoding	Duration (ns)	Accuracy (%)	Out-of-resonance (%)
MNIST	Normal	10	49.30	24.58
MNIST	Normal	20	46.43	23.64
MNIST	Box	10	64.13	24.35
MNIST	Box	20	62.05	24.90
MNIST	Inverse	10	68.57	23.67
MNIST	Inverse	20	64.09	24.29
Fashion-MNIST	Normal	10	53.13	25.16
Fashion-MNIST	Normal	20	50.01	27.36

but it is meant to provide the reader with an idea of how much useful information can be conveyed by the MRR network output at a single time. This setting can be particularly useful when tackling simpler ML tasks requiring memory on an edge computing system with stringent restrictions on computational complexity, energy consumption, or latency.

In particular, we consider all the available output representations for a given input optical power, and we downsample them by a factor of 3 to reduce experimental noise and computation time (see Sec. V D 3 in Methods for details). At each moment in time, we use the different outputs of the network as features for the linear classifier. For example, if we have 60 representations (ten laser frequencies per output port) at a given input power, our classifier will be applied to 60 values, each taken at the same time from the 60 parallel representations. We repeat this for all the available output pixels (i.e., the value corresponding to a single time sample of the downsampled output time series), and we select the one with the highest validation accuracy (the dataset is split into 55 000 training samples, 5000 validation samples, and 10 000 test samples). The test accuracies are subsequently computed using this selected index and are shown in Fig. 6(b) and reported in Table IV for the different classification cases considered (i.e., different encodings of the MNIST and the Fashion-MNIST datasets, explained in Sec. II E). An experimental baseline accuracy is computed for each case, considering a single out-of-resonance linear representation (discussed in Sec. II C), corresponding to the case where no processing by the MRR network occurs. Moreover, a fully digital baseline accuracy is computed for the MNIST and Fashion-MNIST tasks using the original datasets. It should be stressed that, due to setup limitations, we performed sequential measurements of network responses at different input wavelengths. We consider the corresponding obtained ML accuracies as approximations of the accuracies we would obtain in the fully parallel case, where different wavelengths are injected concurrently and might influence each other in the MRR network.

While the baseline accuracies are usually below 25% (except for the Fashion-MNIST case, where it reaches 27.4%), by employing our photonic network we reach 68.6% for the MNIST (with *inverse* encoding and a pixel duration of 10 ns) and 53.1% for the Fashion-MNIST (pixel duration of 10 ns). Such an improvement is ascribed to the diverse memory effects provided by the different representations generated by the MRR network, where the value of a single

pixel retains information about the previous image pixels. This conclusion is supported by the observed trend in classification accuracy as a function of the considered pixel index (in chronological order), shown in Fig. 6(c). The single-pixel accuracy notably increases with the pixel index (in addition to fluctuations due to the image spatial structure), as the network's memory builds up until a single pixel conveys enough information about a substantial part of the full image. As expected, this is clearly not the case for the experimental baseline accuracy because of the absence of optical memory.

Moreover, we investigate how the number of parallel nonlinear representations benefits single-pixel accuracy for the best-performing case (i.e., MNIST classification with *inverse* encoding and a 10 ns pixel duration). Specifically, we calculated the classification accuracy while increasing the number of employed representations (corresponding to the number of single-pixel features fed into the classifier), starting from the lowest laser frequency at the first output port and gradually adding more output ports and laser frequencies [Fig. 6(d)]. The corresponding plot presents a continuously increasing trend, thus showing that most of the different nonlinear representations generated by our MRR network are useful to significantly improve the linear classifier's performance.

IV. CONCLUSION

We studied the highly complex and diverse nonlinear dynamics hosted by a microring resonator (MRR) network (namely, a silicon photonic integrated circuit comprising 64 coupled MRRs) and its multiwavelength parallelism for the classification of images encoded as a time-dependent optical signal. In particular, we investigated the capability of the MRR network to expand its input dimensionality through the use of multiple physical output ports, laser wavelengths, and input power levels, to enhance the computational performance of a linear readout classifier (as per the *reservoir computing* scheme). To do so, we considered two popular machine learning benchmark tasks on image classification, respectively from the MNIST (handwritten digits) and Fashion-MNIST (clothing items) datasets. By stacking together the signals from different output ports, we achieved classification accuracies of 96.49% and 85.81%, respectively, starting from corresponding baselines (accuracy without the photonic network) of 92.12% and 83.83%.

This photonic approach provides an efficient alternative to traditional digital neural networks. For instance, when using $\sim 60 \times 10^3$ trainable parameters, our system achieved a 96.49% accuracy on the MNIST dataset. However, a standard digital ANN with the same number of trainable parameters only reached 96.07(10)%. To match our system's performance, a digital ANN would require a hidden layer with 120 neurons instead of 77, which would increase the number of trainable parameters to 95×10^3 and the number of nonlinearities to 130. In contrast, our system's readout linear classifier only requires ten softmax operations to achieve the same accuracy.

Moreover, we studied the case where the readout classifier is applied to a single time sample (or pixel) rather than to the whole image, which in principle enables machine learning inference without digital memory. We found that the optical memory and dimensionality expansion performed by the MRR network allowed classification accuracies of 68.57% and 53.13% for the MNIST and Fashion-MNIST tasks, respectively, starting from baselines of 23.67% and 25.16%. Although not directly relevant to the goal of accurate image classification, this result shows that the MRR network can greatly improve the performance of fast and computationally cheap digital models when tackling simple tasks that require working memory. This is significant, e.g., for edge computing applications with stringent restrictions on computational cost, energy consumption, or latency.

Further advantages of our MRR network are compactness (0.15 mm^2), low energy consumption (a few milliwatts of input optical power), and a technologically mature and CMOS-compatible platform. Furthermore, our system can be easily scaled up by employing more physical ports, MRRs, and wavelength channels. Exciting future directions are direct application to photonic sensors producing time-dependent signals and making the network internally trainable by applying optical modulators to the MRRs.

V. METHODS

A. Photonic integrated circuit

We employ a photonic integrated circuit (PIC) fabricated by imec (Leuven, Belgium) on a silicon-on-insulator (SOI) platform. The silicon waveguides have a cross-sectional dimension of $450 \times 220 \text{ nm}^2$ and are embedded in a silica cladding.

The circuit comprises a network of 64 microring resonators coupled with bus waveguides, as shown in Fig. 7(a). Adjacent MRRs in the same row are spaced by $22.7 \mu\text{m}$. MRRs are horizontally offset by $11.35 \mu\text{m}$ compared to the neighboring rows. Each MRR features a racetrack geometry, with a bend radius of $7 \mu\text{m}$ and straight coupling sections of $0.71 \mu\text{m}$. The coupling gap between the bus waveguides and the resonators is $0.2 \mu\text{m}$.

The circuit includes one input port and nine output ports at the end of the bus waveguides, each equipped with a grating coupler with a nominal insertion loss of $\sim 3 \text{ dB}$.

B. Experimental setup

The optical setup [Fig. 7(b)] used for working with integrated photonics consists of a transmission stage, in which the information is converted from an electric signal into an optical one, and a receiver stage, in which it is converted back to an electric signal. The light travels inside the photonic integrated circuit (PIC) or inside polarization-maintaining optical fibers, which link all the optic-related instruments required for the correct functioning of the system.

A continuous-wave tunable laser (TLS, Pure Photonics) provides optical input at a fixed power and selected wavelength, typically centered around 1550 nm . This wavelength corresponds to an optical frequency in the C-band and an energy slightly above half the bandgap of silicon. As such, it is sufficient to induce nonlinear effects, such as two-photon absorption, that are essential to the dynamical behavior of the microring resonator network.

The optical signal is produced by modulating the laser light with an electro-optic modulator (EOM, iXblue model MXAN-LN-10), driven by an arbitrary waveform generator (AWG, Spectrum model DN2.663-02).

Then, it is amplified with an erbium-doped fiber amplifier (EDFA, Thorlabs) and adjusted to the desired power level using an electronic variable optical attenuator (VOA).

Light, in addition to its intensity, has other degrees of freedom. One of the most important ones is polarization, which describes the orientation of the electric field vector of the wave. Photonic components are typically sensitive to polarization and are optimized for




FIG. 7. Micrograph of the MRR network and employed experimental setup. (b) The coherent light (wavelength around $1550 \mu\text{m}$) of a tunable laser source (TLS) is modulated by an electro-optical modulator (EOM) controlled by an arbitrary waveform generator (AWG). The resulting optical signal is then amplified by an erbium-doped fiber amplifier (EDFA) and passes through a polarization controller (PC) and a variable optical attenuator (VOA). Then, the optical beam is split by a 99%–1% beam splitter (BS), where the smaller fraction of light power is read by a slow photodetector (PD) as an input power reference. The rest of the optical power is coupled via a cleaved fiber into the input directional coupler of the integrated MRR network. Similarly, the output signal is picked up by a second cleaved fiber and fed into a variable optical attenuator (VOA) before being acquired by a fast PD. Both PDs' electric outputs are sampled and stored by an oscilloscope (OSC).

a specific polarization state. To ensure efficient coupling and consistent behavior of the microring resonator network, a polarization control (PC) stage is used to adjust the input light's polarization.

We monitor the average input power by directing about 1% of the signal to a slow photodetector (PD, 30 kHz bandwidth, New Focus model M2033), splitting the light with a 99:1 fiber splitter. The remaining part is injected into the PIC via a fiber. The terminal side of the optical fiber is mounted on a three-axis linear piezoelectric stage and inclined at a specific angle (about 10° with respect to the normal of the PIC). The light that exits the fiber is coupled to the input grating coupler. Analogously, the transmitted light is sent back to another optical fiber from an output grating coupler. Finally, the signal is collected by a fast photodetector (PD, 600 MHz, Menlo-system model FPD610-FC-NIR), which is preceded for safety by a second VOA. This second VOA safeguards the fast photodetector from excessively high optical power.

The electrical signals from both photodetectors are acquired by an oscilloscope (OSC, Picoscope model 6000). The chip temperature is actively stabilized by a thermostat system controlled by a PID controller connected to a Peltier cell and a 10 k Ω thermistor.

1. Data acquisition

Each dataset is encoded as an optical time sequence, as previously described in Sec. II E. Specifically, the sequence consists of flattened images interleaved with short breaks that reset the system to its initial state. A longer pause is inserted at the end to mark the sequence boundary. This entire representation is repeatedly generated by the AWG. The available AWG time steps are constrained by the discrete set of sampling rates supported: the AWG operates at a maximum rate of 500 MSa/s, with only a finite number of lower rates available. Note that the AWG and the oscilloscope operate at different sampling rates and are not synchronized, as detailed in Table V. To ensure adequate temporal resolution, we opted for a slightly denser sampling for the output signal compared to the input one, so as to avoid excessive oversampling, which would rapidly fill the acquisition memory and limit the recordable sequence length.

The following procedure is carried out for each available output port.

The data acquisition begins with an initial sequence captured at an out-of-resonance laser wavelength and maximum input power, aiming to maximize the SNR. At this stage, the system operates within a linear regime, ensuring that the measured signal is only subject to linear transformations.

TABLE V. Sampling time of the employed instrumentation used to generate and acquire the time sequences. The sampling times of the Arbitrary Waveform Generator (AWG) and the oscilloscope (OSC) are shown in the first two columns. The ratio between these two timesteps is provided in the third column, which also corresponds to the upsampling ratio from the original 784 features to the final number of features (fourth column). The duration of a single image and the brief pause between two consecutive images are displayed in the fifth and sixth columns, respectively.

AWG (ns)	OSC	Ratio	No. of features	Duration	Pause
2	0.8 ns	2.5	1960	1568 ns	2744 ns
10	6.4 ns	1.5625	1225	7.84 μ s	13.72 μ s
20	12.8 μ s	1.5625	1225	15.68 μ s	15.68 μ s

Subsequently, the system's response is recorded across a range of resonant laser frequencies and input power levels. The outer acquisition loop varies the laser frequency, while for each selected frequency, the input power is adjusted using the VOA.

During this process, the oscilloscope records a signal that lasts slightly more than twice the length of the full sequence. This ensures that at least one full, uninterrupted representation of the dataset is captured.

Finalizing the process, a second out-of-resonance sequence is recorded under the same conditions as the initial one. This final measurement verifies the stability of the system and ensures that no significant changes occurred during the acquisition, thus validating the reliability of the collected data.

C. Estimation of self-pulsing regime

The discrimination between self-pulsing and constant output is done by computing the sum of the absolute value of the Fourier transform of the acquired signals. This scalar feature separates the two regimes, forming two distinct clusters that do not overlap, making them linearly separable. The self-pulsing period is then computed as the lag greater than zero corresponding to the highest peak in the autocorrelation. If there is no peak, or if the peaks are lower than 0.5, then the series is considered non-periodic. The autocorrelation is computed via the Wiener-Khinchin theorem, since the Fourier transform is already available. We believe it likely that the self-pulsing states without a period estimate [the ones with an asterisk in the colormap in Fig. 2(d)] are indeed chaotic states, although a more in-depth analysis (beyond the scope of this work) is required to confirm this. Other possible but unlikely reasons for a missing period estimate are that our measured waveform was not long enough to detect periodicity or that the SNR ratio was too high.

D. Machine learning aspects

1. Linear classifier (LC)

For baseline comparison, we utilize a linear classifier (LC) implemented in *Tensorflow* and consisting of a single *Dense* layer.

The classifier's input size corresponds to the number of features. The ten classes were represented using a one-hot encoding scheme.

The model is trained using the *CategoricalCrossentropy* loss function, with the parameter *from_logits = True*. This setting enables the loss function to perform the final softmax operation internally, which enhances computational efficiency and numerical stability.

Model parameters are updated using Stochastic Gradient Descent (SGD). A batch size of 64 is used for training.

Input data are normalized using the training set to achieve a zero mean and a unit standard deviation.

Due to the convex nature of its loss landscape, a properly regularized linear classifier consistently converges to the same accuracy, which is why a single training run is sufficient for its evaluation.

a. Digital. The LC classifier is applied to the normalized original data.

b. Single port. The output signal from the MRR network serves as the input for the LC. A *Dropout* layer is placed after the input representation and before the *Dense* layer, which reduces the feature space to 784 during training. The chosen *Dropout* rate acts

as a regularization mechanism, and its value was chosen through hyperparameter optimization, which tunes it over a validation set. This approach also mimics the downsampling performed on an oscilloscope, where the signal is oversampled to mitigate noise and subsequently averaged to extract a single, representative value per time step.

c. Multiple ports. Similar to the single-port approach, each representation is reduced to 784 features during training using a *Dropout* layer. These features are then concatenated to form a single input for the LC. Therefore, the total number of input features for the LC is the product of the number of features per port and the number of ports.

2. Multilayer ANN in software

In contrast to the linear classifier, the training of a multilayer neural network can lead to different final solutions depending on the initial conditions. To mitigate this variability and obtain a more robust performance estimate, we trained the multilayer ANN ten times with different random weight initializations. The results presented are the average of the accuracies obtained from these ten training runs, along with their standard deviation.

Our multilayer ANN is relatively simple, consisting of an input layer followed by a single *Dense* hidden layer with a variable number of neurons and a ReLU activation function. A *Dropout* layer with a rate of 0.5 is applied after this hidden layer for regularization. The network concludes with a final *Dense* output layer for classification, configured identically to the one used in the linear classifier, with *CategoricalCrossentropy* as the loss function.

3. Classification using a single pixel

The output signals of the MRR network, corresponding to the flattened output images, are downsampled from 1225 time samples (with relative experimental indices) to 409 output pixels, in order to speed up computations and reduce the effect of experimental noise and low vertical resolution of the oscilloscope (averaging smoothens the signal). This is done using the Python function *skimage.measure.block_reduce* from *scikit-image*, based on the average of time samples and with a downsample factor of 3. It should be stressed that this operation approximates the use of a slower (and cheaper) photodetector, which would not introduce any additional computational cost. Moreover, to speed up the computations producing the ML accuracies reported in Fig. 6(b) and Table IV, we exclude the first 80 time samples from the investigation (with the exception of the digital baseline results), as they convey little information about the image class and consistently provide low classification accuracy. For this part of our work, we employ a multinomial logistic regression as a readout linear classifier, using the Python function *sklearn.linear_model.LogisticRegression* from the *scikit-learn* library, with the solver option *lbfgs*.

ACKNOWLEDGMENTS

We acknowledge fruitful discussions with Dr. Stefano Biasi. A.F. acknowledges funding by the European Union under Grant No. 101070238-NEUROPULS. A.L. acknowledges funding by the European Union under Grant No. 101064322-ARIADNE. Views and opinions expressed are, however, those of the author(s) only and do

not necessarily reflect those of the European Union or The European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

A.F. and A.L. contributed equally to this work.

A.F. and A.L. conceived the experiments, prepared the setup and performed the experimental measurements. A.L. designed the integrated photonic circuit. A.F. performed the data analysis for Sec. III A, A.L. for Sec. III B. A.F. and A.L. wrote the manuscript. P.B. and L.P. supervised the work. All authors contributed to the revision of the manuscript.

Alessandro Foradori: Conceptualization (equal); Data curation (equal); Investigation (equal); Methodology (equal); Software (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal). **Alessio Lugnan:** Conceptualization (equal); Data curation (equal); Investigation (equal); Methodology (equal); Software (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal). **Lorenzo Pavesi:** Funding acquisition (equal); Project administration (equal); Resources (equal); Supervision (equal); Writing – review & editing (equal). **Peter Bienstman:** Funding acquisition (equal); Project administration (equal); Resources (equal); Supervision (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request. Other data and the code supporting the results of this study are openly available at <https://doi.org/10.5281/zenodo.17105607>.

REFERENCES

- ¹R. Li, Y. Gong, H. Huang, Y. Zhou, S. Mao, Z. Wei, and Z. Zhang, “Photonics for neuromorphic computing: Fundamentals, devices, and opportunities,” *Adv. Mater.* **37**(2), 2312825 (2025).
- ²D. Brunner *et al.*, “Roadmap on neuromorphic photonics,” *arXiv:2501.07917* [cs.ET] (2025).
- ³N. Farmakidis, B. Dong, and H. Bhaskaran, “Integrated photonic neuromorphic computing: Opportunities and challenges,” *Nat. Rev. Electr. Eng.* **1**(6), 358–373 (2024).
- ⁴S. Masaad, S. Sackesyn, S. Sygletos, and P. Bienstman, “Experimental demonstration of 4-port photonic reservoir computing for equalization of 4 and 16 QAM signals,” *J. Lightwave Technol.* **42**(24), 8555–8563 (2024).
- ⁵A. Argyris, “Photonic neuromorphic technologies in optical communications,” *Nanophotonics* **11**(5), 897–916 (2022).
- ⁶J. Qin, S. Jiang, Z. Wang, X. Cheng, B. Li, Y. Shi, D. P. Tsai, A. Q. Liu, W. Huang, and W. Zhu, “Metasurface micro/nano-optical sensors: Principles and applications,” *ACS Nano* **16**(8), 11598–11618 (2022).

- ⁷A. Venkateswaran, N. Lalam, J. Wuenschell, P. R. Ohodnicki, Jr., M. Badar, K. P. Chen, P. Lu, Y. Duan, B. Chorpene, and M. Buric, "Recent advances in machine learning for fiber optic sensor applications," *Adv. Intell. Syst.* **4**(1), 2100067 (2022).
- ⁸U. M. N. Jayawickrema, H. M. C. M. Herath, N. K. Hettiarachchi, H. P. Sooriyaarachchi, and J. A. Epaarachchi, "Fibre-optic sensor and deep learning-based structural health monitoring systems for civil structures: A review," *Measurement* **199**, 111543 (2022).
- ⁹B. 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**(2), 102–114 (2021).
- ¹⁰J. Liu, G. Feng, W. Li, S. Hao, S. Han, Q. Zhu, B. Tian, C. Duan, and J. Chu, "Physical reservoir computing for Edge AI applications," *Innovation Mater.* **3**(2), 100127 (2025).
- ¹¹K. 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**(1), 3541 (2014).
- ¹²W. Zhou, X. Shen, X. Yang, J. Wang, and W. Zhang, "Fabrication and integration of photonic devices for phase-change memory and neuromorphic computing," *Int. J. Extreme Manuf.* **6**(2), 022001 (2024).
- ¹³W. H. P. Pernice and H. Bhaskaran, "Photonic non-volatile memories using phase change materials," *Appl. Phys. Lett.* **101**(17), 171101 (2012).
- ¹⁴C. Ríos, M. Stegmaier, P. Hosseini, D. Wang, T. Scherer, C. D. Wright, H. Bhaskaran, and W. H. P. Pernice, "Integrated all-photonic non-volatile multi-level memory," *Nat. Photonics* **9**(11), 725–732 (2015).
- ¹⁵J. Feldmann, N. Youngblood, C. D. Wright, H. Bhaskaran, and W. H. P. Pernice, "All-optical spiking neurosynaptic networks with self-learning capabilities," *Nature* **569**(7755), 208–214 (2019).
- ¹⁶S. Biasi, A. Lugnan, D. Micheli, and L. Pavesi, "Exploring the potential of self-pulsing optical microresonators for spiking neural networks and sensing," *Commun. Phys.* **7**(1), 380 (2024).
- ¹⁷A. Lugnan, S. Biasi, A. Foradori, P. Bienstman, and L. Pavesi, "Reservoir computing with all-optical non-fading memory in a self-pulsing microresonator network," *Adv. Opt. Mater.* **13**(11), 2403133 (2025).
- ¹⁸A. Lugnan, S. Aggarwal, F. Brückhoff-Plückelmann, C. D. Wright, W. H. P. Pernice, H. Bhaskaran, and P. Bienstman, "Emergent self-adaptation in an integrated photonic neural network for backpropagation-free learning," *Adv. Sci.* **12**(2), 2404920 (2025).
- ¹⁹H. Jaeger and H. Haas, "Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication," *Science* **304**(5667), 78–80 (2004).
- ²⁰G. 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 Networks* **115**, 100–123 (2019).
- ²¹S. Biasi, G. Donati, A. Lugnan, M. Mancinelli, E. Staffoli, and L. Pavesi, "Photonic neural networks based on integrated silicon microresonators," *Intell. Comput.* **3**(0067), 0067 (2024).
- ²²A. N. Tait, A. X. Wu, T. F. de Lima, E. Zhou, B. J. Shastri, M. A. Nahmias, and P. R. Prucnal, "Microring weight banks," *IEEE J. Sel. Top. Quantum Electron.* **22**(6), 312–325 (2016).
- ²³D. Bazzanella, S. Biasi, M. Mancinelli, and L. Pavesi, "A microring as a reservoir computing node: Memory/nonlinear tasks and effect of input non-ideality," *J. Lightwave Technol.* **40**(17), 5917–5926 (2022).
- ²⁴C. Mesaritakis, V. Papataxiarhis, and D. Syvridis, "Micro ring resonators as building blocks for an all-optical high-speed reservoir-computing bit-pattern-recognition system," *J. Opt. Soc. Am. B* **30**(11), 3048–3055 (2013).
- ²⁵T. Van Vaerenbergh, M. Fiers, P. Mechet, T. Spuesens, R. Kumar, G. Morthier, B. Schrauwen, J. Dambre, and P. Bienstman, "Cascadable excitability in microrings," *Opt. Express* **20**(18), 20292–20308 (2012).
- ²⁶M. Mancinelli, M. Borghi, F. Ramiro-Manzano, J. M. Fedeli, and L. Pavesi, "Chaotic dynamics in coupled resonator sequences," *Opt. Express* **22**(12), 14505–14516 (2014).
- ²⁷A. Lugnan, S. García-Cuevas Carrillo, C. D. Wright, and P. Bienstman, "Rigorous dynamic model of a silicon ring resonator with phase change material for a neuromorphic node," *Opt. Express* **30**(14), 25177–25194 (2022).
- ²⁸J. Xiang, Y. Zhang, Y. Zhao, X. Guo, and Y. Su, "All-optical silicon microring spiking neuron," *Photonics Res.* **10**(4), 939–946 (2022).
- ²⁹G. Donati, S. Biasi, L. Pavesi, and A. Hurtado, "All-optical spiking processing and reservoir computing with a passive silicon microring and wavelength-time division multiplexing," *Photonics Res.* **13**, 2641 (2025).
- ³⁰J. Xiang, Y. Zhao, A. He, J. Xiao, Y. Su, and X. Guo, "Photonic neuromorphic processing with on-chip electrically-driven microring spiking neuron," *Laser Photonics Rev.* **19**(3), 2400604 (2025).
- ³¹Y. LeCun, C. Cortes, and C. J. Burges, MNIST handwritten digit database, ATT Labs, available at: <http://yann.lecun.com/exdb/mnist>, 2010.
- ³²H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms," *arXiv:1708.07747* [cs.LG] (2017).
- ³³E. Alpaydin, *Introduction to Machine Learning* (MIT Press, 2020).