



HAL
open science

The impact of feature representation on the accuracy of photonic neural networks

Mauricio Gomes de Queiroz, Paul Jimenez, Raphael Cardoso, Mateus Vidaletti Costa, Mohab Abdalla, Ian O'connor, Alberto Bosio, Fabio Pavanello

► To cite this version:

Mauricio Gomes de Queiroz, Paul Jimenez, Raphael Cardoso, Mateus Vidaletti Costa, Mohab Abdalla, et al.. The impact of feature representation on the accuracy of photonic neural networks. *APL Machine Learning*, 2024, 2 (3), 10.1063/5.0226172 . hal-04739864

HAL Id: hal-04739864

<https://hal.science/hal-04739864v1>

Submitted on 16 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

RESEARCH ARTICLE | SEPTEMBER 04 2024

The impact of feature representation on the accuracy of photonic neural networks

Mauricio Gomes de Queiroz ; Paul Jimenez ; Raphael Cardoso ; Mateus Vidaletti Costa ; Mohab Abdalla ; Ian O'Connor ; Alberto Bosio ; Fabio Pavanello 

 Check for updates

APL Mach. Learn. 2, 036109 (2024)

<https://doi.org/10.1063/5.0226172>



View Online



Export Citation

Articles You May Be Interested In

Complex-valued forecasting of the global solar irradiation

J. Renewable Sustainable Energy (August 2013)

Complex-valued neural network for short-term load forecasting of Yogyakarta hybrid renewable power plant

AIP Conf. Proc. (March 2024)

Global Mittag-Leffler stability and existence of the solution for fractional-order complex-valued NNs with asynchronous time delays

Chaos (November 2021)



Special Topics Open for Submissions

[Learn More](#)

The impact of feature representation on the accuracy of photonic neural networks

Cite as: APL Mach. Learn. 2, 036109 (2024); doi: 10.1063/5.0226172

Submitted: 29 June 2024 • Accepted: 31 July 2024 •

Published Online: 4 September 2024



View Online



Export Citation



CrossMark

Mauricio Gomes de Queiroz,^{1,a)}  Paul Jimenez,¹  Raphael Cardoso,¹  Mateus Vidaletti Costa,^{1,2} 
Mohab Abdalla,^{1,2}  Ian O'Connor,¹  Alberto Bosio,¹  and Fabio Pavanello³ 

AFFILIATIONS

¹Ecole Centrale de Lyon, INSA Lyon, CNRS, Universite Claude Bernard Lyon 1, CPE Lyon, INL, UMR5270, 69130 Ecully, France

²School of Engineering, RMIT University, Melbourne VIC 3000, Australia

³Univ. Grenoble Alpes, Univ. Savoie Mont Blanc, CNRS, Grenoble INP, CROMA, 38000 Grenoble, France

Note: This paper is part of the APL Machine Learning Special Topic on Neuromorphic Technologies for Novel Hardware AI.

^{a)}Author to whom correspondence should be addressed: mauricio.gomes@ec-lyon.fr

ABSTRACT

Photonic neural networks (PNNs) are gaining significant interest in the research community due to their potential for high parallelization, low latency, and energy efficiency. PNNs compute using light, which leads to several differences in implementation when compared to electronics, such as the need to represent input features in the photonic domain before feeding them into the network. In this encoding process, it is common to combine multiple features into a single input to reduce the number of inputs and associated devices, leading to smaller and more energy-efficient PNNs. Although this alters the network's handling of input data, its impact on PNNs remains understudied. This paper addresses this open question, investigating the effect of commonly used encoding strategies that combine features on the performance and learning capabilities of PNNs. Here, using the concept of feature importance, we develop a mathematical methodology for analyzing feature combination. Through this methodology, we demonstrate that encoding multiple features together in a single input determines their relative importance, thus limiting the network's ability to learn from the data. However, given some prior knowledge of the data, this can also be leveraged for higher accuracy. By selecting an optimal encoding method, we achieve up to a 12.3% improvement in the accuracy of PNNs trained on the Iris dataset compared to other encoding techniques, surpassing the performance of networks where features are not combined. These findings highlight the importance of carefully choosing the encoding to the accuracy and decision-making strategies of PNNs, particularly in size or power constrained applications.

© 2024 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/10.1063/5.0226172>

I. INTRODUCTION

Artificial Intelligence (AI) systems gained widespread relevance in recent years,¹ finding diverse applications ranging from image classification² to speech recognition.³ These systems have traditionally been implemented in electronic hardware, benefiting from the steady performance improvements driven by the miniaturization of electronic integrated circuits. However, with components now shrinking to the atomic scale, the limitations of this platform become apparent.⁴ At this size, for example, quantum effects may disrupt functionality,⁵ and the heat from densely packed devices becomes hard to dissipate.⁶ In response, new technologies are being explored to enable further improvements in AI. These emerging technologies are often not subject to the same constraints of their electronic

counterparts and, thus might offer more efficient alternatives for certain applications.⁷

Photonic neural networks (PNNs) are hardware implementations of AI systems that perform computations on optical signals, rather than on electronic ones. Using light, they are able to leverage several of its properties to potentially enable high parallelization, low latency, and reduced power consumption.⁸ For example, PNNs have been demonstrated to perform sub-nanosecond image classification⁹ and to achieve up to 10^{12} multiply-accumulate operations per second.¹⁰ However, transitioning from electronics to photonics remains challenging. Practical applications of medium- to large-scale systems are currently limited by the large physical footprint of photonic circuits,^{11,12} their loss accumulation, and the high power consumption of some of its electro-optic devices.¹³

One way of alleviating these issues is by optimizing circuits,^{14,15} or carefully designing PNNs to minimize circuit size. A common practice found in the literature involves taking advantage of the complex representation of light (using amplitude and phase) to represent multiple features in a single input, thus combining multiple real-valued features into fewer complex-valued inputs. By using fewer inputs, a circuit requires fewer components and a smaller network, which leads to a reduction in the overall footprint. Such technique aligns well with the capabilities of photonic circuits, which are able to process complex inputs through complex transformations.¹⁶

However, the way we represent features in neural networks (NNs) greatly influences the difficulty of the problems they solve. For example, in tasks with radial symmetry centered around the origin, opting for a polar coordinate system can emphasize the relevant feature relationships necessary for accurately solving the task, short-cutting the network’s need to learn it. This approach can significantly reduce the computational complexity required for achieving high accuracy. Moreover, the choice of feature representation also shapes the network’s approaches to solve tasks, as NNs tend to rely on the most straightforward cues available within the data.¹⁷ This highlights the need to understand which feature relationships are emphasized by the representation strategies used in PNNs. By doing so, we can ensure that these networks not only achieve high accuracy but also adopt desirable decision-making strategies.

In this paper, we explore the role of feature representation in the accuracy and decision-making strategies of PNNs. We investigate the common practice of combining various features into a single input, using explainable AI (XAI) methods to compare relative importance of the combined features. To our knowledge, only one work investigated such practice as a means of improving accuracy in PNNs.¹⁸ However, the consequences of the feature combination itself are still unknown, and different feature representations were not explored. Our work tackles these open questions with a mathematical analysis of feature combination focused on photonic implementations, where networks and circuits are constrained

by size. We point out how different data representations and hardware implementations can be exploited for higher accuracy and lower complexity, as well as the shortcomings of the current solutions.

The rest of this paper is structured as follows: in Secs. II and III, we review the basics of photonic implementations of AI and feature importance metrics. In Sec. IV, we calculate the relative importance of features that share the same input. Sections V and VI discuss practical examples and simulations of artificial neural networks (ANNs) and PNNs. Finally, Sec. VII concludes the discussions brought up in this paper.

II. PHOTONIC NEURAL NETWORKS

In this section, we provide a review of ANNs and their photonic implementations. We also address common strategies of representing features in light, which will be used in our further discussions in Sec. IV.

A. Artificial neural networks

Artificial neural networks (ANNs), first proposed in the 1940s,¹⁹ are mathematical functions loosely inspired by how the human brain processes information. These functions are known to be universal approximators,²⁰ hence their ability to handle a wide variety of tasks. The network’s behavior, i.e., the way it processes inputs, is determined by their connection strengths (called “weights”) and non-linearities (referred to as “activation functions”).²¹ Typically, these parameters are obtained through training, approximating the ANN to a probability function associated with the given task. For instance, in classification tasks, ANNs are designed to assign a class to an input by approximating a function that calculates the likelihood of belonging to each class.²²

Consider a fully connected, feed forward (FF) NN, consisting of L layers and designed with N inputs and N outputs, as shown in Fig. 1(a). The process by which a given layer l transforms its inputs is described as

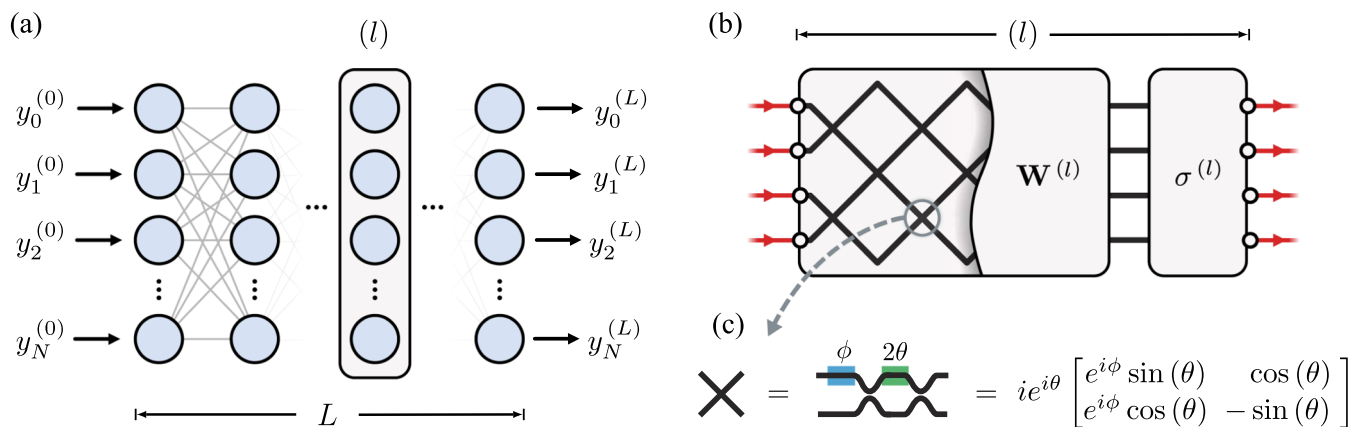


FIG. 1. (a) Representation of a generic neural network. An arbitrary layer l is highlighted. (b) Schematic of the photonic implementation of a neural network layer using meshes of Mach-Zehnder interferometers (MZIs). (c) An MZI and its two phase shifters ($\phi, 2\theta$) are illustrated alongside the transfer matrix representation of the transformation it performs over the field amplitude.

$$\bar{z}^{(l)} = \mathbf{W}^{(l)} \cdot \bar{y}^{(l-1)} + \bar{b}^{(l)}, \quad (1)$$

$$\bar{y}^{(l)} = \sigma^{(l)}(\bar{z}^{(l)}). \quad (2)$$

Initially, inputs are combined through weighted sums by a weight matrix \mathbf{W}^0 to obtain z^0 . Then, the element-wise application of an activation function $\sigma(\cdot)$ to z^0 introduces non-linearity and yields the output of the layer, where $y^{(0)}$ is the input of the network and $y^{(l)}$ the output. A bias \bar{b} might be added before the activation function to allow for the network to better adjust to the data. The entire network, from the first to last layer, can be seen as a sequence of such transformations, written as $\bar{y}^{(L)} = f(\bar{y}^{(0)})$.

Thus, ANNs implement input–output mappings that can be either real or complex. real-valued neural networks (RVNNs) are characterized by real parameters and inputs, with $f: \mathbb{R}^N \mapsto \mathbb{R}^N$. In these networks, each layer scales and combines inputs before non-linearly transforming them. Complex-valued neural networks (CVNNs), on the other hand, operate in the complex domain, meaning that both the input vector and the network's parameters are complex-valued and $f: \mathbb{C}^N \mapsto \mathbb{C}^N$.²³ In that case, each layer has the ability to not only scale and combine but also rotate inputs in the complex plane. This rotation, inherent to complex algebra, makes CVNNs more suitable for tasks where phase information is important, such as in audio processing²⁴ or optical communications.²⁵

B. Photonic implementations

Photonic computing is emerging as a promising approach to improve ANN implementations for specific applications by computing with light. This allows us to leverage its unique characteristics to potentially enable faster and more energy-efficient AI systems. For example, in the optical domain, linear transformations can be done passively²⁶ and information can be easily parallelized and processed at high speeds.¹⁰

PNNs are implementations of ANNs through photonic inputs, components, and transformations.²⁷ Although no single photonic component acts as an artificial neuron, a circuit can be designed to perform the mathematical operations of an ANN. This is achieved by using several components, such as waveguides, interferometers, and modulators, which guide and manipulate light signals. These circuits operate on complex signals and implement complex transformations, meaning that PNNs can act as RVNNs and CVNNs, depending on the task at hand.

Several PNN circuits were suggested and demonstrated experimentally. They can be broadly categorized by how different inputs are distinguished, whether through spatial, wavelength, or time domains.²⁸

In this study, we focus on PNNs that use spatial differentiation of inputs. These networks assign a separate input to each optical signal and implement weight matrix multiplications by making different inputs interfere with each other. Most notably, this is achieved by using meshes of Mach–Zehnder interferometers (MZIs).^{26,29} The interference, and hence the specific mathematical operation performed by the mesh, can be selected by adjusting the phase shifters found in these devices. Activation functions, on the other hand, can be implemented by using any of the devices and circuits that exhibit optical non-linearity.^{30,31} The schematics of an

ANN implementation and an MZI are shown in Figs. 1(b) and 1(c), respectively.

If PNNs use coherent light inputs, they can be represented in the complex domain. In these networks, the i th input is characterized by an amplitude A_i and phase ϕ_i . Thus, the input vector can be expressed as $\bar{y}^{(0)} = [A_1 e^{i\phi_1}, \dots, A_N e^{i\phi_N}]^T \in \mathbb{C}^N$. Given the two degrees of freedom available for each input, feature encoding can be achieved using various methodologies. We divide common approaches found in the literature into two distinct groups: real and complex encoding.

Real encoding simplifies the input representation by encoding data solely in the amplitude of the optical signals, maintaining a uniform initial phase across all inputs (in practice having $\phi_i = 0 \forall i$ and thus $\bar{y}^{(0)} \in \mathbb{R}^N$). Several researchers employ this encoding method for its compatibility with RVNNs used in electronic computers.^{8,32} It allows for an easy mapping of weights from electronically trained networks to photonic transformations. In these networks, while the nature of the transformations of individual MZIs is inherently complex, the overall behavior can effectively be real-valued. Since no phase information is used, only the amplitude of the outputs is of interest, which simplifies the detection scheme. However, it is important to ensure that different inputs experience the same phase before reaching the network to maintain phase consistency, which might not be simple to achieve experimentally.

In contrast, complex encoding uses both amplitude and phase at the same time, having inputs that lie in the complex plane, that is $\bar{y}^{(0)} \in \mathbb{C}^N$. The transformations in the PNN in this case are complex, and thus, detection of both intensity and phase in the outputs might be used, adding to the electronic complexity of the circuit. In image classification tasks, for example, real-valued input images can be transformed into Fourier space representation to obtain phase and amplitude information,^{33–35} or have different sections mapped to the real and imaginary parts of complex numbers,^{18,36} which reduces by half the number of inputs.

The encoding choice for PNNs influences the network's behavior, the type of information that is detected at the output, and the overall size of the circuit, as it may imply the use of additional peripheral devices. Beyond hardware specifications, this choice might also impact how features are processed within the network. When two features share the same input, the network may process them differently from the way they would be processed individually. Understanding these dynamics is crucial for optimizing the PNN performance.

III. FEATURE IMPORTANCE

In this section, we look to the field of XAI for methods of evaluating feature importance in ANNs, to later study the impact of combining features in PNNs. We focus on gradient-based techniques, particularly sensitivity analysis.

ANNs, especially those with several layers, are highly non-linear models that use numerous parameters. The network's complexity often leads them to be regarded as opaque or “black-box” systems, since their decision-making processes are difficult to grasp intuitively. That is, while we can mathematically describe how a given output is obtained, it is difficult to specify “why” with an intuitive explanation.

Nonetheless, being able to explain the decision-making strategies of a model has a number of practical applications. Clear explanations can, for instance, enhance our understanding of a problem or be used to demonstrate fair treatment. In photonics research, XAI is currently used to explain the inverse design of circuits³⁷ or to aid in the description of physical models.³⁸ The concept of “explainability” is still subject of an ongoing debate^{39,40} and, consequently, a variety of methods have been proposed to attain it.^{41,42} Highlighting which input features are considered as important to an ANN is a common way to explain its outputs. Several methods estimate such feature importance, of which we emphasize sensitivity analysis.

Sensitivity analysis quantifies feature importance by examining how sensitive the output of the model is to small variations in each feature.^{43–45} The underlying principle is that if small changes in an input lead to significant changes in the output, then that input is likely to be important for the network, i.e., it contributes to the prediction of this output. In such case, the importance of the i th input, $y_i^{(0)}$, to the c th output of the network, $y_c^{(L)}$, is denoted by $R_{i \rightarrow c}$,

$$R_{i \rightarrow c} = \left| \frac{\partial y_c^{(L)}}{\partial y_i^{(0)}} \right|. \quad (3)$$

Gradient-based explanations are frequently used in image classification tasks to generate saliency maps⁴⁶ and also show fair performance in matching feature importance in simulated data.⁴⁷ Over time, other methods built up on sensitivity analysis, addressing some of its drawbacks by suggesting additional forms of estimating feature importance.⁴⁸ For example, adding Gaussian noise to the input and averaging their resulting gradients helps in generating more consistent saliency maps.⁴⁹ These techniques are often easy to implement, given that the necessary partial derivatives can be computed through back-propagation.

IV. ANALYTICAL DERIVATION OF RELATIVE IMPORTANCE

Here, we use the concepts elaborated in previous sections to investigate how the importance of features is shaped in PNNs. Initially, we employ the sensitivity analysis shown in Sec. III to obtain the importance of an arbitrary feature encoded in one input. Then, we introduce the concept of *encoding functions* to describe the different feature encoding processes and representations in photonics, given in Sec. II.

Consider a set of features $\mathcal{X} = \{x_1, \dots, x_n\} \in \mathbb{R}$. Assume that we want all the elements in \mathcal{X} to be used by our model. However, due to either a prohibitively large quantity of features or size restrictions on our network, we also wish to use a number of inputs that is less than the number of elements in \mathcal{X} . To achieve both objectives, we combine features into complex inputs, as shown in Sec. II. In this context, we calculate the relative importance of such combined features, to understand what relationships are highlighted by our inputs.

Our first objective is to obtain the importance of an arbitrary feature, say $x_j \in \mathcal{X}$, with respect to an arbitrary output of the model. This feature is represented only in input $y_i^{(0)}$, and its importance is

assessed in relation to the c th output of the network. Applying the chain rule to Eq. (3), we write this importance $R_{j \rightarrow c}$ as

$$R_{j \rightarrow c} = \left| \frac{\partial y_c^{(L)}}{\partial x_j} \right| = \left| \frac{\partial y_c^{(L)}}{\partial y_i^{(0)}} \frac{\partial y_i^{(0)}}{\partial x_j} \right|. \quad (4)$$

The partial derivative $\partial y_i^{(0)}/\partial x_j$ of Eq. (4) relates to the way x_j is represented in the input. The process of creating an input from elements of \mathcal{X} is what we term feature encoding. An input $y_i^{(0)}$ obtained from the feature x_j is hence written as $y_i^{(0)} = g_i(x_j)$, where g_i is the encoding function for the i th input. Considering the encoding process as such, we can write

$$R_{j \rightarrow c} = \left| \frac{\partial y_c^{(L)}}{\partial y_i^{(0)}} \frac{\partial g_i(x_j)}{\partial x_j} \right|. \quad (5)$$

It should be noted how the importance depends on both the network, represented in the derivative from output to input, and the feature encoding process, given the presence of the encoding function. The modulus operation ensures that the importance is always positive and real-valued. Here, we assume the network to be derivable in the vicinity of the current input, which might not be the case for some CVNN architectures.

Next, we examine the scenario where two arbitrary features, x_j and x_k , are represented using a single input $y_i^{(0)}$, comparing their relevance. We define the relative importance between x_j and x_k to the c th output, $R_{j,k \rightarrow c}$ as the following ratio:

$$R_{j,k \rightarrow c} = \frac{R_{j \rightarrow c}}{R_{k \rightarrow c}} = \left| \frac{\frac{\partial g_i(x_j, x_k)}{\partial x_j}}{\frac{\partial g_i(x_j, x_k)}{\partial x_k}} \right|. \quad (6)$$

We see that the component of Eq. (4) related to the network is canceled, leaving only the derivatives of the encoding function. Thus, $R_{j,k \rightarrow c}$ is solely determined by the way features are encoded into y_i , and hence, it is independent of the considered output. To simplify the notation, we drop the subscript indicating the output for the rest of this paper. One of the consequences of Eq. (6) is that the encoding function chosen to combine x_j and x_k defines how these features are perceived by the model relative to one another.

Although encoding functions are a method of pre-processing features, in the context of PNNs, they can also be implemented in hardware. The incorporation of encoding functions in the circuit is particularly interesting for low-latency applications, as the speed at which inputs are transformed and combined would be limited only by the reconfigurability of the driving electronics. We now explore two types of complex encoding functions to see how they dictate relative feature importance. We also point out how they could be implemented in hardware.

A. Exponential encoding

Since we are dealing with complex-valued inputs, one intuitive way to encode two features x_j and x_k into a single input would be to encode x_j in its amplitude and x_k in its phase. This encoding function can be written as

$$g(x_j, x_k) = x_j e^{ix_k}. \quad (7)$$

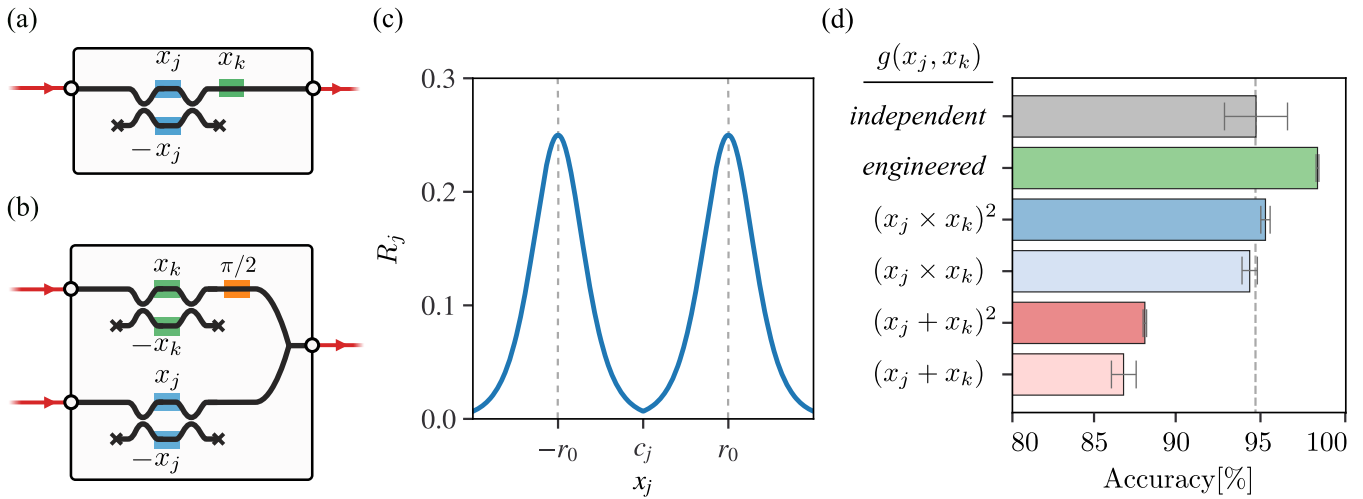


FIG. 2. Panels (a) and (b) show photonic circuits that implement encoding functions. By modulating the phase shifters to the indicated value, one can achieve functions similar to exponential encoding (a) or linear encoding (b). (c) Importance of a feature x_j when $r_p = |x_j|$, which is equal to moving P alongside the x_j axis. (d) Mean accuracy of 100 ANNs trained with several encoding functions. “Independent” refers to the situation where features are not combined and “engineered” to the implementation of Eq. (16).

The relative importance between these two features is calculated as

$$R_{j,k \rightarrow c} = \left| \frac{e^{ix_k}}{ix_j e^{ix_k}} \right| = \frac{1}{|x_j|}. \quad (8)$$

In this case, the relative importance between the two features is dynamic, establishing an amplitude-dependent relation between the importance of amplitude and phase.

A hardware version of an exponential encoding function is shown in Fig. 2(a), where a balanced MZI and a phase shifter are used to modulate the amplitude and phase of an input, respectively. The encodings and importance are not exactly the same since this amplitude modulation scheme is mediated by a sine function, it implements $g(x_j, x_k) = i \sin(x_j) \exp(x_k i)$. Here, Eq. (7) can be achieved short of a global phase shift by mapping x_j to $\arcsin(x_j)$.

B. Linear encoding

Another way to combine two features is by encoding one in the real part and the other in the imaginary part of a complex input. This can be represented by the function,

$$g(x_j, x_k) = x_j + ix_k. \quad (9)$$

Here, we find their relative importance to be

$$R_{j,k} = \frac{1}{|i|} = 1, \quad (10)$$

thus indicating that both features will be considered to have the same importance for the network. Since they are independent from the weights of the network, their relative importance cannot be unlearned, i.e., it cannot be modified by further training. This might pose problems when the chosen encoding leads to relative importance that do not match the data.

An encoding function similar to that of Eq. (9) implemented in hardware, is shown in Fig. 2(b). There, two MZIs are used as amplitude modulators, while one of their outputs has its phase shifted by $\pi/2$ to encode the respective input in the imaginary axis. In that case, $g(x_j, x_k) = i[\sin(x_j) + \sin(x_k)i]$. Equation (9) can be achieved short of a global phase shift by mapping x_j and x_k to $\arcsin(x_j)$ and $\arcsin(x_k)$.

V. ON THE IMPACT OF ENCODING FUNCTIONS TO ANNS

In this section, we address the practical implications of the discussions brought up in Sec. IV. Here, our objective is to demonstrate how a well-engineered encoding function can significantly improve the accuracy of an ANN on a test task. We begin by defining such a task and studying the relative feature importance found in a solution to it. Later, we create an encoding function that reproduces these importances on trained ANNs, finally comparing its use against others.

Consider a simple classification problem with a known solution in the real domain: determining whether points lie inside or outside an n -sphere. An n -sphere is the generalization of a circle to $n + 1$ dimensions, similar to how hyperplanes generalize planes. It is defined by a set of points $S^{(n)}$ that are equidistant from a central point $c_0 = (c_1, \dots, c_{n+1})$ by a radius r_0 . The distance of a point $P = (x_1, \dots, x_{n+1})$ to c_0 is

$$r_p(x_1, \dots, x_{n+1}) = \sqrt{\sum_{i=1}^{n+1} (x_i - c_i)^2}. \quad (11)$$

Naturally, P is considered outside of the n -sphere if r_p exceeds r_0 and inside otherwise. In this context, a mathematical model that outputs a probability of P being outside of $S^{(n)}$ can be constructed

using a logistic function. The logistic function $\sigma(x) = 1/(1 + e^{-x})$ is bounded between 0 and 1 with a smooth sigmoid transition and is typically used in binary classification problems. Given the coordinates of P , this model can be expressed as

$$y = f(x_1, \dots, x_{n+1}) = \sigma(r_p - r_0). \quad (12)$$

Here, y represents the probability that $r_p > r_0$ given the coordinates of P . When $y = 0.5$, Eq. (12) delineates the boundary defined by $S^{(n)}$, allowing for accurate classification of points based on this threshold.

Since Eq. (12) can be used to accurately classify any point P , we conjecture that its relative feature importances are desirable to other models that wish to do it as well. Thus, we examine the sensitivity of y to an arbitrary feature x_j , which can be calculated according to Eq. (4) as

$$R_j = \left| \frac{\partial y}{\partial x_j} \right| = \left| \sigma(r_p - r_0)(1 - \sigma(r_p - r_0)) \frac{x_j - c_j}{r_p} \right| \forall r_p \neq 0. \quad (13)$$

As shown in Fig. 2(c), the importance of a feature peaks when it is at the boundary of the n -sphere. Exactly at that point, small variations in x_j cause the largest deviations of the probability of P being outside of $S^{(n)}$. The relative feature importance between two features x_j and x_k is

$$R_{j,k} = \left| \frac{x_j - c_j}{x_k - c_k} \right|. \quad (14)$$

Now assume that, as in Sec. IV, we are constrained by size and thus wish to combine different features to reduce the number of inputs. However, for this particular example (and by design), we have prior knowledge of the relative importance of features before combining them. Taking advantage of this, we propose an encoding function $g(x_j, x_k)$ that achieves the desired reduction in dimensionality while preserving the relationships given by Eq. (14). Given that after combining features, their relative importance should follow Eq. (6), we can obtain one such $g(x_j, x_k)$ by solving the following system of partial derivatives:

$$\begin{cases} \left| \frac{\partial g(x_j, x_k)}{\partial x_j} \right| = |x_j - c_j|, \\ \left| \frac{\partial g(x_j, x_k)}{\partial x_k} \right| = |x_k - c_k|. \end{cases} \quad (15)$$

This system can lead to one of many solutions of the form

$$g(x_j, x_k) = \frac{1}{2} (x_j^2 + x_k^2) - c_j x_j - c_k x_k + C, \quad (16)$$

where C is a constant.

In the same manner, we can obtain other functions that express different relative importance. A constant $R_{j,k} = 1$, for instance, is achieved by using $g(x_j, x_k) = (x_j + x_k)^n \forall n \in \mathbb{R}$. Alternatively, $g(x_j, x_k) = (x_j \times x_k)^n \forall n \in \mathbb{R}$ leads to $R_{j,k} = |x_k/x_j|$, which is the inverse of Eq. (14) when c_j and c_k are zero. To compare

the use of these encoding functions, we trained several ANNs, benchmarking them on networks that do not combine inputs (called “independent” here). We are particularly interested in the performance of Eq. (16), which we call “engineered” encoding function. The description of the training procedures is in the following.

A dataset of 1000 points in four dimensions was created, where each coordinate value was randomly chosen between -2 and 2 . The points were labeled as either inside or outside of a three-sphere $S^{(3)}$ of radius 1, centered at the origin $c_0 = (0, 0, 0, 0)$, according to their position. In order to obtain a balanced dataset, we generated the same amount of points inside and outside the sphere. The networks trained to solve this task were composed of an input layer containing either two or four neurons (depending on the combination of features or not), a hidden layer of six neurons, and an output layer with a single neuron. A logistic activation function σ was used for every layer. Each encoding function was used to train 100 different networks, thus accounting for the random initialization of weights and random shuffling of the dataset prior to training. The networks were trained on 70% of the available data for 100 epochs with a learning rate of 0.001 and tested on the remaining data.

The results of these experiments are shown in Fig. 2(d). We notice that some representations can render the task harder to solve, while others maintain, to some extent, the accuracy achieved by the use of independent inputs. The engineered encoding function in Eq. (16) outperformed all others. With this example, we show that the way we combine features plays a role in the accuracy of ANNs. Given prior knowledge on how features relate to one another, which may come from domain-specific knowledge or from inspecting the data (noticing symmetries or class distributions, for example), we could estimate relative feature importance and obtain an encoding function that aligns with them. Combining features with said encoding function could improve the network performance.

VI. APPLICATION OF ENCODING FUNCTIONS IN PNNs

In this section, we retake the subject of this study and explore the use of different encoding functions in PNNs trained on the Iris dataset,⁵⁰ a standard benchmark for classification algorithms. Our goal is to show how carefully chosen encoding functions might lead to higher accuracies in PNNs. To this end, we compare the performance of several encoding functions by means of simulations of PNNs, which differ significantly from the ANNs of Sec. V in terms of their complex-valued inputs and transformations.

The Iris flower classification task involves categorizing three different Iris species (Setosa, Versicolour, and Virginica) based on four features: the lengths and widths of sepals and petals. The dataset, consisting of 150 labeled data points, has considerable class overlaps, such that no single feature alone can distinguish all the species, making this an ideal candidate for our experiments. Visualizations of feature distributions and class overlaps are shown in Fig. 3(a).

Our experimental design involves training PNNs by combining features in pairs, as shown in Fig. 3(b). We assess their performance by averaging the accuracy of 100 trained PNNs, benchmarking them against a PNN that does not combine features. This sample size was chosen to allow for convergence in the average values obtained for accuracy, given the variability in the training process.

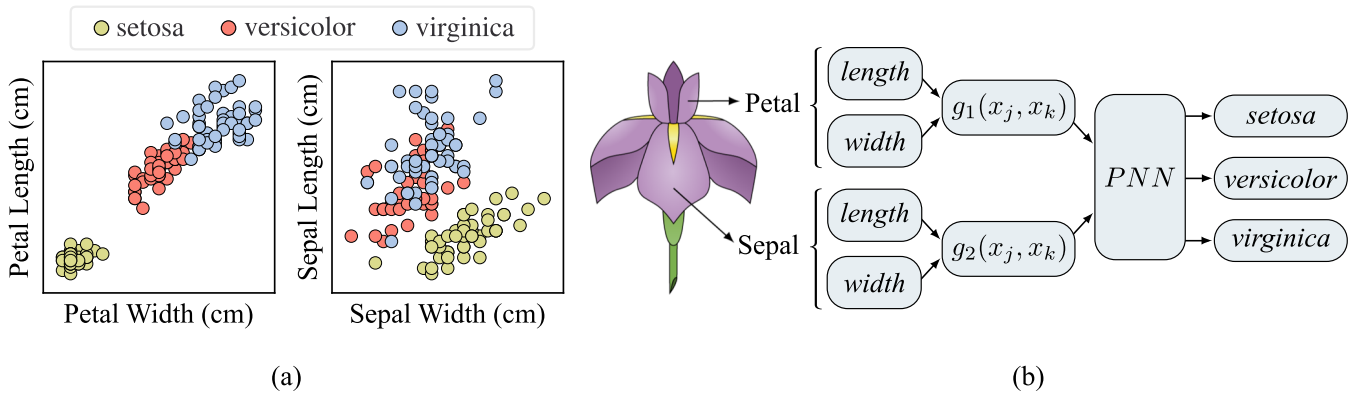


FIG. 3. (a) Class distribution of normalized features of the Iris dataset. (b) Schematic of the process of combining features utilized in the training of PNNs.

The architecture of the PNNs consists of a single hidden layer with six neurons. Depending on the configuration, the number of input neurons varies between 3 (when combining features) and 5 (when using independent inputs), where one input acts as a bias for both configurations. The output layer has three neurons, matching the number of classes. All the configurations use the same underlying circuit, where NN layers are implemented using meshes of MZIs with trainable phase shifters, as shown in Fig. 1(b). Every layer is followed by a *softplus* activation function, which can be implemented in integrated photonic circuits.⁵¹ Although its hardware implementation would change both the modulus and the phase of the signals, we model it by applying $\text{softplus}(x) = \log(1 + \exp(x))$ solely to the modulus of the complex numbers.³³ This approach allows us to simulate the gain and activation behavior while simplifying the model by avoiding additional phase changes. These phase changes can make the simulation and training more challenging and are less

critical to the primary function of the *softplus* activation in this context.

The circuits were simulated using the Photontorch Python package.⁵² The simulations were performed under ideal conditions, excluding noise and component imperfections. They were trained for 300 epochs on 70% of the dataset, reserving the remaining 30% for testing. The dataset was divided into five shuffled batches per epoch to enhance training stability. A *softmax* function was used to convert the output light intensity values into class probabilities.²² Weight updates were performed using a cross-entropy loss function combined with stochastic gradient descent. The initial learning rate was set at 0.01 and adjusted at learning plateaus. While higher accuracies may be achieved by further optimizing the training process to each specific case, we opted for a constant training procedure across all circuits to isolate the effects of different encoding functions.

16 October 2024 12:35:35

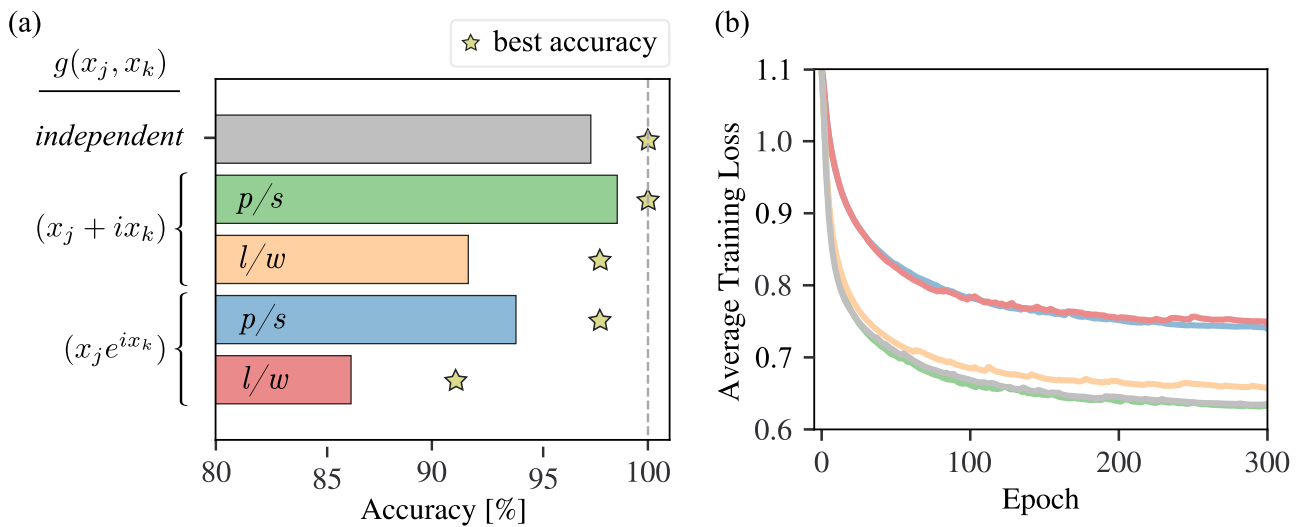


FIG. 4. (a) Accuracy of PNNs for different encoding functions. “Independent” represents the accuracy of a PNN trained without combining features. (b) Average training loss of the trained PNNs per epoch. The color scheme is the same as the one defined in panel (a).

Here, we investigate the use of the encoding functions detailed in Sec. II: linear and exponential encoding. Given the anisotropic nature of the Iris classification task, unlike the n -sphere problem, we also consider which features to combine. To explore the impacts of this choice on the obtained accuracy, we use two combination strategies for features: grouping by the lengths and widths (l/w) or by petal and sepal information (p/s). We benchmarked the performance of PNNs using different encoding functions and groupings of features to the independent case, where features were not combined.

The results of these experiments, shown in Fig. 4, are summarized as follows: exponential encoding exhibited the lowest performance, falling up to 11% in mean accuracy compared to the independent benchmark. In contrast, linear encoding, commonly used in the photonics community,^{18,33–36} was able to match the performance of the independent case. The difference between the best and worst performing encoding functions was 12.3%. These results highlight that both the manner in which features are combined and the combination of features itself play significant roles in the final accuracy of PNNs. When comparing different feature groupings, we found that (l/w) consistently performed worse than (p/s), demonstrating that the choice of which features to combine can also impact accuracy for some tasks.

These findings are supported by heuristics found in the data. A closer inspection of Fig. 3(a) reveals that petal length and petal width together are highly discriminative of the different classes. These features separate different species in a similar fashion, as evidenced by the distribution of classes along the diagonal of the plot, suggesting that they may have similar importance. Thus, the combination (p/s) with linear encoding would combine petal length and width with an equal relative importance, expressing such relationships.

VII. CONCLUSIONS

Combining features into single inputs in PNNs can lead to reduced number of inputs and associated devices as well as enable the use of smaller and more energy efficient NNs. These benefits would help render some circuits more feasible to be simulated, fabricated, tested, or deployed. However, this method of feature combination imposes predefined relationships among the features that may not necessarily reflect the nature of data or task at hand. Nonetheless, selecting or designing encoding functions based on an understanding of the dataset or from domain-specific knowledge can lead to improved accuracy. We have illustrated this first on an ideal simple example and then for simulated PNNs.

In the scenarios shown here, as it is seen in the literature, features are combined into a single input. As an alternative, we could distribute features across many inputs, circumventing the discussions brought up here and making it possible to learn other relative feature importance. For instance, principal component analysis (PCA) can be used for dimensionality reduction, distributing features across many inputs simultaneously.³² Expanding on this concept, a learnable encoding function that uses every feature available would be a fully connected layer of NN,³⁵ which is more complex and less efficient than what is explored in our work. In addition, the approach used here could be applied directly at a hardware level, using integrated photonics and CMOS-compatible platforms for volume production.

Here, the discussions highlight that there is no neutral way of using this feature combination strategy in PNNs. Combining features in this manner will necessarily emphasize certain feature relationships. Sometimes, a PNN might achieve good performance metrics despite combinations that are not ideal. However, even if high accuracy is achieved, these combinations can also introduce or amplify biases in the model outputs, depending on the specific features and their encoded interactions. Rather than leaving this to chance, we suggest to carefully assess how to encode features given the nature of the problem and data.

ACKNOWLEDGMENTS

The authors thank **Peter Bienstman** and **Thomas Van Vaerenbergh** for their valuable advice and discussions during the early stages of this work, as well as for reviewing the paper before submission.

This project received funding from École Centrale de Lyon and ANR (Grant No. ANR-20-THIA-0007-01). **Paul Jimenez** and **Fabio Pavanello** acknowledge ANR's support (Grant No. ANR-20-CE39-0004), and **Fabio Pavanello** acknowledges support from European Union's Horizon Europe research and innovation programme (Grant No. 101070238).

Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union. 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

Mauricio Gomes de Queiroz: Conceptualization (lead); Data curation (lead); Formal analysis (lead); Investigation (lead); Methodology (lead); Software (lead); Visualization (lead); Writing – original draft (lead); Writing – review & editing (equal). **Paul Jimenez:** Formal analysis (supporting); Methodology (supporting); Writing – review & editing (equal). **Raphael Cardoso:** Methodology (supporting); Writing – review & editing (equal). **Mateus Vidaletti Costa:** Methodology (supporting); Writing – review & editing (equal). **Mohab Abdalla:** Methodology (supporting); Writing – review & editing (equal). **Ian O'Connor:** Supervision (supporting); Writing – review & editing (equal). **Alberto Bosio:** Funding acquisition (equal); Supervision (supporting); Writing – review & editing (equal). **Fabio Pavanello:** Conceptualization (supporting); Formal analysis (supporting); Funding acquisition (equal); Methodology (supporting); Supervision (lead); Writing – review & editing (equal).

DATA AVAILABILITY

The code that reproduces the experiments and the data that support the findings of this study are available at github.com/mgomesq/feature-representation-pnns.

REFERENCES

- ¹S. Dong, P. Wang, and K. Abbas, "A survey on deep learning and its applications," *Comput. Sci. Rev.* **40**, 100379 (2021).
- ²K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) (2014).
- ³H. Purwins, B. Li, T. Virtanen, J. Schlüter, S.-Y. Chang, and T. Sainath, "Deep learning for audio signal processing," *IEEE J. Sel. Top. Signal Process.* **13**, 206–219 (2019).
- ⁴T. N. Theis and H.-S. P. Wong, "The end of Moore's law: A new beginning for information technology," *Comput. Sci. Eng.* **19**, 41–50 (2017).
- ⁵J. R. Powell, "The quantum limit to Moore's law," *Proc. IEEE* **96**, 1247–1248 (2008).
- ⁶C. E. Leiserson, N. C. Thompson, J. S. Emer, B. C. Kuszmaul, B. W. Lampson, D. Sanchez, and T. B. Scharld, "There's plenty of room at the top: What will drive computer performance after Moore's law?," *Science* **368**, eaam9744 (2020).
- ⁷M. M. Waldrop, "The chips are down for Moore's law," *Nature* **530**, 144 (2016).
- ⁸Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund, and M. Soljačić, "Deep learning with coherent nanophotonic circuits," *Nat. Photonics* **11**, 441–446 (2017).
- ⁹F. Ashtiani, A. J. Geers, and F. Aflatouni, "An on-chip photonic deep neural network for image classification," *Nature* **606**, 501–506 (2022).
- ¹⁰X. Xu, M. Tan, B. Corcoran, J. Wu, A. Boes, T. G. Nguyen, S. T. Chu, B. E. Little, D. G. Hicks, R. Morandotti *et al.*, "11 TOPS photonic convolutional accelerator for optical neural networks," *Nature* **589**, 44–51 (2021).
- ¹¹T. Shibata, S. Kamei, T. Kitoh, T. Tanaka, and M. Kohtoku, "Compact and low insertion loss (~10 dB) Mach-Zehnder interferometer-synchronized arrayed-waveguide grating multiplexer with flat-top frequency response," *Opt. Express* **16**, 16546–16551 (2008).
- ¹²X. Xiao, M. B. On, T. Van Vaerenbergh, D. Liang, R. G. Beausoleil, and S. Yoo, "Large-scale and energy-efficient tensorized optical neural networks on III–V-on-silicon MOSCAP platform," *APL Photonics* **6**, 126107 (2021).
- ¹³A. N. Tait, "Quantifying power in silicon photonic neural networks," *Phys. Rev. Appl.* **17**, 054029 (2022).
- ¹⁴G. Mourgas-Alexandris, M. Moralis-Pegios, A. Tsakyridis, S. Simos, G. Dabos, A. Totovic, N. Passalis, M. Kirtas, T. Rutirawat, F. Gardes *et al.*, "Noise-resilient and high-speed deep learning with coherent silicon photonics," *Nat. Commun.* **13**, 5572 (2022).
- ¹⁵M. G. de Queiroz, R. Cardoso, P. Jimenez, M. Abdalla, I. O'Connor, A. Bosio, and F. Pavanello, "Power reduction in photonic meshes by MZI optimization," in *Frontiers in Optics* (Optica Publishing Group, 2023), pp. JW4A–7.
- ¹⁶H. Zhang, M. Gu, X. Jiang, J. Thompson, H. Cai, S. Paesani, R. Santagati, A. Laing, Y. Zhang, M. Yung *et al.*, "An optical neural chip for implementing complex-valued neural networks," *Nat. Commun.* **12**, 457 (2021).
- ¹⁷R. Geirhos, J.-H. Jacobsen, C. Michaelis, R. Zemel, W. Brendel, M. Bethge, and F. A. Wichmann, "Shortcut learning in deep neural networks," *Nat. Mach. Intell.* **2**, 665–673 (2020).
- ¹⁸R. Qiu, A. Eldebiky, L. Zhang, X. Yin, C. Zhuo, U. Schlichtmann, and B. Li, "OplixNet: Towards area-efficient optical split-complex networks with real-to-complex data assignment and knowledge distillation," in *Design, Automation and Test in Europe (DATE)*, 2024.
- ¹⁹W. S. McCulloch and W. Pitts, "A logical calculus of the ideas immanent in nervous activity," *Bull. Math. Biophys.* **5**, 115–133 (1943).
- ²⁰K. Hornik, M. Stinchcombe, and H. White, "Multilayer feedforward networks are universal approximators," *Neural Networks* **2**, 359–366 (1989).
- ²¹V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proc. IEEE* **105**, 2295–2329 (2017).
- ²²I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016).
- ²³J. Basse, L. Qian, and X. Li, "A survey of complex-valued neural networks," [arXiv:2101.12249](https://arxiv.org/abs/2101.12249) (2021).
- ²⁴G. Kim, D. K. Han, and H. Ko, "Sound source localization using complex-valued deep neural networks," in *2024 IEEE International Conference on Consumer Electronics (ICCE)* (IEEE, 2024), pp. 1–4.
- ²⁵S. Masaad, E. Gooskens, S. Sackesyn, J. Dambre, and P. Bienstman, "Photonic reservoir computing for nonlinear equalization of 64-QAM signals with a Kramers–Kronig receiver," *Nanophotonics* **12**, 925–935 (2023).
- ²⁶M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani, "Experimental realization of any discrete unitary operator," *Phys. Rev. Lett.* **73**, 58 (1994).
- ²⁷B. J. Shastri, A. N. Tait, T. Ferreira de Lima, W. H. Pernice, H. Bhaskaran, C. D. Wright, and P. R. Prucnal, "Photonics for artificial intelligence and neuromorphic computing," *Nat. Photonics* **15**, 102–114 (2021).
- ²⁸Y. Bai, X. Xu, M. Tan, Y. Sun, Y. Li, J. Wu, R. Morandotti, A. Mitchell, K. Xu, and D. J. Moss, "Photonic multiplexing techniques for neuromorphic computing," *Nanophotonics* **12**, 795–817 (2023).
- ²⁹W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walsmley, "Optimal design for universal multiport interferometers," *Optica* **3**, 1460 (2016).
- ³⁰I. A. D. Williamson, T. W. Hughes, M. Minkov, B. Bartlett, S. Pai, and S. Fan, "Reprogrammable electro-optic nonlinear activation functions for optical neural networks," *IEEE J. Sel. Top. Quantum Electron.* **26**, 7700412 (2020).
- ³¹A. Jha, C. Huang, and P. R. Prucnal, "Reconfigurable all-optical nonlinear activation functions for neuromorphic photonics," *Opt. Lett.* **45**, 4819–4822 (2020).
- ³²K. H. R. Mojaver, B. Zhao, E. Leung, S. M. R. Safaei, and O. Liboiron-Ladouceur, "Addressing the programming challenges of practical interferometric mesh based optical processors," *Opt. Express* **31**, 23851–23866 (2023).
- ³³S. Banerjee, M. Nikdast, and K. Chakrabarty, "Characterizing coherent integrated photonic neural networks under imperfections," *J. Lightwave Technol.* **41**, 1464–1479 (2023).
- ³⁴R. Hamerly, S. Bandyopadhyay, and D. Englund, "Asymptotically fault-tolerant programmable photonics," *Nat. Commun.* **13**, 6831 (2022).
- ³⁵R. Wang, P. Wang, C. Lyu, G. Luo, H. Yu, X. Zhou, Y. Zhang, and J. Pan, "Multicore photonic complex-valued neural network with transformation layer," *Photonics* **9**, 384 (2022).
- ³⁶M. Y.-S. Fang, S. Manipatruni, C. Wierzynski, A. Khosrowshahi, and M. R. DeWeese, "Design of optical neural networks with component imprecisions," *Opt. Express* **27**, 14009 (2019).
- ³⁷Z. Jia, W. Qarony, J. Park, S. Hooten, D. Wen, Y. Zhiyenbayev, M. Secli, W. Redjem, S. Dhuey, A. Schwartzberg, E. Yablonovitch, and B. Kanté, "Interpretable inverse-designed cavity for on-chip nonlinear photon pair generation," *Optica* **10**, 1529–1534 (2023).
- ³⁸C. Yeung, J.-M. Tsai, B. King, Y. Kawagoe, D. Ho, M. W. Knight, and A. P. Raman, "Elucidating the behavior of nanophotonic structures through explainable machine learning algorithms," *ACS Photonics* **7**, 2309–2318 (2020).
- ³⁹Z. C. Lipton, "The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery," *Queue* **16**, 31–57 (2018).
- ⁴⁰T. Miller, "Explanation in artificial intelligence: Insights from the social sciences," *Artif. Intell.* **267**, 1–38 (2019).
- ⁴¹W. Samek, G. Montavon, S. Lapuschkin, C. J. Anders, and K.-R. Müller, "Explaining deep neural networks and beyond: A review of methods and applications," *Proc. IEEE* **109**, 247–278 (2021).
- ⁴²G. Montavon, W. Samek, and K.-R. Müller, "Methods for interpreting and understanding deep neural networks," *Digital Signal Process.* **73**, 1–15 (2018).
- ⁴³A. H. Sung, "Ranking importance of input parameters of neural networks," *Expert Syst. Appl.* **15**, 405–411 (1998).
- ⁴⁴L. Fu and T. Chen, "Sensitivity analysis for input vector in multilayer feed-forward neural networks," in *IEEE International Conference on Neural Networks* (IEEE, 1993), pp. 215–218.
- ⁴⁵I. Dimopoulos, J. Chronopoulos, A. Chronopoulou-Sereli, and S. Lek, "Neural network models to study relationships between lead concentration in grasses and permanent urban descriptors in Athens city (Greece)," *Ecol. Modell.* **120**, 157–165 (1999).
- ⁴⁶K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," in *Workshop at International Conference on Learning Representations*, 2014.
- ⁴⁷J. D. Olden, M. K. Joy, and R. G. Death, "An accurate comparison of methods for quantifying variable importance in artificial neural networks using simulated data," *Ecol. Modell.* **178**, 389–397 (2004).

⁴⁸M. Ancona, E. Ceolini, C. Öztireli, and M. Gross, “Towards better understanding of gradient-based attribution methods for deep neural networks,” [arXiv:1711.06104](#) (2017).

⁴⁹D. Smilkov, N. Thorat, B. Kim, F. Viégas, and M. Wattenberg, “SmoothGrad: Removing noise by adding noise,” [arXiv:1706.03825](#) (2017).

⁵⁰R. A. Fisher, “The use of multiple measurements in taxonomic problems,” [Ann. Eugen.](#) **7**, 179–188 (1936).

⁵¹J. R. Rausell Campo and D. Pérez-López, “Reconfigurable activation functions in integrated optical neural networks,” [IEEE J. Sel. Top. Quantum Electron.](#) **28**, 8300513 (2022).

⁵²F. Laporte, J. Dambre, and P. Bienstman, “Highly parallel simulation and optimization of photonic circuits in time and frequency domain based on the deep-learning framework PyTorch,” [Sci. Rep.](#) **9**, 5918 (2019).