

ENN: A Neural Network With DCT Adaptive Activation Functions

Marc Martínez-Gost¹, Graduate Student Member, IEEE, Ana Pérez-Neira², Fellow, IEEE,
and Miguel Ángel Lagunas³, Life Fellow, IEEE

Abstract—The expressiveness of neural networks highly depends on the nature of the activation function, although these are usually assumed predefined and fixed during the training stage. Under a signal processing perspective, in this article we present Expressive Neural Network (ENN), a novel model in which the non-linear activation functions are modeled using the Discrete Cosine Transform (DCT) and adapted using backpropagation during training. This parametrization keeps the number of trainable parameters low, is appropriate for gradient-based schemes, and adapts to different learning tasks. This is the first non-linear model for activation functions that relies on a signal processing perspective, providing high flexibility and expressiveness to the network. We contribute with insights in the explainability of the network at convergence by recovering the concept of bump, this is, the response of each activation function in the output space. Finally, through exhaustive experiments we show that the model can adapt to classification and regression tasks. The performance of ENN outperforms state of the art benchmarks, providing above a 40% gap in accuracy in some scenarios.

Index Terms—Neural networks, adaptive activation functions, discrete cosine transform, explainable machine learning.

I. INTRODUCTION

FUNCTION approximation is a fundamental problem across many domains, such as data analysis, control systems and communications. When the explicit function is not available but input-output data pairs are, the function can be revealed by minimizing a criterion loss in a supervised setting. The problem increases in complexity when the function is non-linear, for which many signal processing techniques have been developed.

Manuscript received 30 June 2023; revised 16 January 2024; accepted 19 January 2024. Date of publication 1 February 2024; date of current version 3 July 2024. This work was supported in part by IRENE, funded by MCIN/AEI/10.13039/501100011033, under Grant PID2020-115323RB-C31 and in part by Catalan government through the project SGR-Cat under Grant 2021-01207. The guest editor coordinating the review of this manuscript and approving it for publication was Prof. Maria Sabrina Greco. (*Corresponding author: Marc Martínez-Gost.*)

Marc Martínez-Gost is with the Centre Tecnològic de Telecomunicacions de Catalunya, 08860 Castelldefels, Spain, and also with the Department of Signal Theory, Communications, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain (e-mail: mmartinez@cttc.es).

Ana Pérez-Neira is with the Centre Tecnològic de Telecomunicacions de Catalunya, 08860 Castelldefels, Spain, also with the Department of Signal Theory, Communications, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain, and also with ICREA Acadèmia, 08010 Barcelona, Spain (e-mail: ana.isabel.perez@upc.edu).

Miguel Ángel Lagunas is with the Department of Signal Theory, Communications, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain (e-mail: m.a.lagunas@cttc.es).

Digital Object Identifier 10.1109/JSTSP.2024.3361154

For instance, least squares [1], orthogonal function approximation [2], [3], kernel methods [4] and neural networks [5], [6], [7], among others. The last decades have suffered an unprecedented growth in the development of artificial neural networks for function approximation due to its empirical success. The inception of neural networks as universal approximators boosted its development across many fields and applications, with different architectures built according to the task and data types to handle. Nevertheless, the expressiveness of the neural network is related to the non-linear activation function, which is usually assumed fixed. An overlooked field of research is in adaptive activation functions (AAF), where not only the weights of the neural networks are trained, but the non-linearities too [8].

In our previous work [9] we introduced the Discrete Cosine Transform (DCT) to approximate an univariate non-linear function in a joint communication and computing setting. Further, in [10] we show how a gradient-based algorithm can be used to tune the DCT coefficients to approximate a function in a supervised setting. However, extending the results to multivariate functions is not trivial: The number of required parameters to approximate the function increases exponentially with the number of input variables, and their corresponding indexes are unknown when the explicit function is not available. In other words, it is cumbersome how the top relevant coefficients can be learnt in a supervised fashion using labeled data.

In this work we propose to extend the capabilities of the DCT with a novel neural network model that integrates the DCT to represent and adapt the activation functions. We call this model Expressive Neural Network (ENN). We exploit the fact that a 2-layer neural network can theoretically represent any function and expand the representation capabilities of the network by adapting the activation functions at each neuron. The advantage of approximating an univariate function with the DCT is twofold: A small number of coefficients is required due to its high energy compaction, and the approximation error is easily controlled by the magnitude of the disregarded coefficients. In this work we also show that the DCT coefficients can be learnt using backpropagation in a supervised fashion and in the same pass as the standard network linear weights. In this way, the architecture is no different from a standard feed-forward neural network with fixed activation functions. From the learning perspective, using the DCT to model the activation functions brings the following benefits:

- **Network size:** The number of parameters in the network grows linear with the number of neurons, which is small due to energy compaction of DCT.
- **Backpropagation:** The tuning algorithm can be implemented because the DCT coefficients are real and ordered in decreasing magnitude. Besides, there exist analytical closed-form solutions for backpropagation.
- **Gradient behavior:** The basis functions (cosines) are real and bounded, which prevents exploding gradients. Likewise, since the Fourier representation creates a periodic function that does not saturate, it also prevents vanishing gradients.
- **Task adaptability:** The output non-linearity is automatically adapted depending on the task (e.g., classification or regression) without changing the loss function.

While we provide a general formulation for multivariate real functions, we constrain the analysis to bivariate functions. This allows to visualize the results and intuitively understand how the network is adapting the activation functions. In this respect, we recover the concept of *bump*, which is the non-linear enclosure that each activation function generates in the output space [11]. The global response of the network corresponds to a weighted sum of all bumps generated at the hidden layer. This concept allows to gain insights in how the network decides to exploit the periodic nature of the DCT model and create the boundaries for classification problems. In this work we focus on two general problems of function approximation, namely classification and regression. In the former there are two hypothesis associated to a function, this is, \mathcal{H}_0 when $f(x_1, x_2) < 0$ and \mathcal{H}_1 otherwise. In regression the goal is to approximate the function $f(x_1, x_2)$. Our primary goal is to show the interpretability and expressiveness of the DCT-based non-linearities in small networks, which is why we do not consider standard datasets and large models. We leave these considerations for future work, along with other learning aspects, such as overfitting. The source code of this study is openly available on <http://github.com/marcmartinezgost/enn>.

The main contributions of this article are described in the following:

- 1) We define ENN, a novel neural network model with non-linear AAF that are parameterized by the DCT. This allows to adapt DCT coefficients in a supervised fashion and learn specific non-linearities according to the task. This results in a highly flexible and expressive model.
- 2) We develop analytical closed-form expressions to adapt the non-linearities with backpropagation. While we choose the Least Mean Squares (LMS) algorithm to update the network parameters, the architecture remains a feed-forward neural network and any alternative algorithm can be used.
- 3) We provide insights in the field of explainable machine learning (XML). We recover the concept of bump, which allows to interpret how the non-linearities are adapted and what is the global response of the neural network. Furthermore, we show how the DCT model helps to dimension the network width, this is, the number of hidden neurons. Particularly, the network converges to duplicated activation functions with opposed linear weights in the output layer. This is, the network cancels out the information coming

from several neurons when the task does not require that many parameters.

- 4) We provide extensive experiments in both classification and regression setups for which ENN outperforms all the benchmarks. In classification tasks, ENN outperforms fixed activation functions up to 40% in accuracy. With that we show how the expressiveness of network highly depends on the activation function, without the need of increasing the size of the network.

The remainder of this article is organized as follows. In Section II, we present a literature review on neural networks for function approximation. Fourier models for non-linear representation are presented in Section III. We present ENN in Section IV and propose the learning procedure for supervised tasks in Section V. The simulation results are shown in Section VI and we conclude the paper in Section VII.

Notation: Lowercase and uppercase bold symbols correspond to vectors and matrices, respectively; $\mathbf{s}[m]$ corresponds to the m -th entry of vector \mathbf{s} ; \mathbb{R} stands for the set of real numbers and ∇ for the gradient.

II. LITERATURE REVIEW

The universal approximation theorem is a well-known result in mathematics stating that a 2-layer neural network can represent any continuous function with an arbitrary number of neurons [12], [13]. The theory behind neural networks has been further developed, providing bounds on the number of required neurons.

In a different line of research, the Kolmogorov-Arnold (KA) representation theorem shows how a multivariate function can be represented by functions of only univariate functions [14]:

$$f(x_1, \dots, x_n) = \sum_{i=1}^{2n+1} \Phi_i \left(\sum_{j=1}^n \phi_{ij}(x_j) \right), \quad (1)$$

where Φ_i and ϕ_j are termed the outer and inner functions, respectively. This result seems to be tightly connected to a 2-layer neural network, since the inner functions correspond to the hidden layer transformation and the outer functions to the output neuron. What is more, the inner functions do not depend on the function f to implement, which resembles the activation functions in neural network architectures. However, the formulation is not exactly identical: the KA representation requires n inner functions for each input variable, while a neural network implements a unique function for each linear combination of inputs. In this way, although there is an extensive literature motivating the development of neural networks with the KA theorem [15], [16], [17], there are still many gaps to be resolved. What is more, the theorem is not constructive and the functions in (1) are highly non-linear.

Despite its success in many applications, neural networks undergo several shortcomings that limit the interpretability of the results: optimization algorithms are easily trapped in local minima, convergence heavily depends on initialization and fail to converge when high non-linearities exist. Usually, each neuron in feedforward neural networks implements a linear combination and a non-linear mapping. The former is usually trained,

while the latter remains fixed. The non-linear activation function allows to generate non-linear mappings, which increases the expressiveness of the network. A common choice is the sigmoid function, although it exhibits well-known issues in its implementation: since the sigmoid saturates at large magnitudes the propagated gradients vanish, slowing down the learning process. The rectified linear unit (ReLU) replaced the sigmoid activation function because it does not suffer from vanishing gradients and it is computationally efficient, which results in faster convergence. Despite its popularity, ReLU also experiences several weaknesses that hinder the learning capacity of the model: the corresponding neuron can become dead when the output remains negative for a long time; likewise, the output is unbounded, which may produce the opposite effect of exploding gradient, making the network unstable.

Many variations and novel activation functions have been designed to enhance the performance of neural networks, although this highly depends on the application and the statistics of the data. In this respect, few authors have tackled the problem of AAF, in which the non-linear mapping is also trained to enhance the expressiveness of the network. In [8] the saturation level and slope of an S-shaped activation function can be tuned independently, which increases the expressiveness with respect to the sigmoid function. In a different vein, several authors proposed to approximate an arbitrary activation function using cubic splines [18], [19]. Later on, some authors introduced an activation function as a weighted sum of basis, such as sine, Gaussian or sigmoid [20], [21]. More recently, the authors in [22] propose an asymmetric S-shaped activation function, but only for regression tasks. The issues derived from these perspective is that they either constrain the geometry of the activation function, or the parametrization is too complex. In this work we empirically prove that having pre-specified activation functions result in suboptimal performance.

In the context of deep learning, several AAF have been designed as well. In [23] the authors propose a gradient adaptive piecewise linear activation function as a sum of hinge-shaped functions. Although they theoretically prove that any piecewise linear function can be expressed as so, this constrains the function to be linear at both extremes. In other words, it does not prevent the neural network from exploding gradients. In [24] an S-shaped rectified linear activation function is proposed. While this model can learn non-linear and non-convex representations, it is not smooth and still constrains the activation to be defined in a piece-wise fashion. The authors in [25] develop a piece-wise polynomial form which can approximate any continuous function. However, the exponential nature of the polynomials may increase the dynamic range of the parameters, which is a non-desirable property for gradient-based procedures. In [26], the ReLU is substituted by an S-shaped activation, which slightly outperforms the ReLU and its variants in different deep learning tasks. In [27], [28], [29], [30] different authors propose to model the activation as a sum of basis functions for deep learning tasks.

There have been also attempts to implement activation functions using Fourier models. In [31] the authors consider a binarized neural network, where tunable parameters only have 1 b. This makes the activation function to be the step function and,

therefore, the gradient is zero almost everywhere. To make the model trainable, the step function is approximated with a Fourier series, which prevents the gradient vanishing effect. Other works propose numerical solutions to specific mathematical models by applying Extreme Learning Machines (i.e., feed-forward neural networks) with Fourier basis functions. In [32], the Fourier series basis is used to solve the one-dimensional asset pricing model. In [33] the authors extend the work with a product of trigonometric functions to cope with the specific multi-dimensional case that solves the generalized Black–Scholes partial differential equation. In the same vein of ELM with trigonometric activation functions is [34], where numerical solution is found to integro-differential equations for risk theory. However, all these works, which do not resort to gradient-based backpropagation to work, are designed to solve specific numerical system of equations. Conversely, the ENN tries to solve the general problem of multi-dimensional function approximation.

In the following section we propose a Fourier-based parametrization for non-linear functions. This is the first approach to model AAF from a signal processing perspective. The proposed model does not constrain the shape of the activation function and keeps the number of trainable parameters low while circumventing the shortcomings of the previously proposed non-linear models.

III. FOURIER MODELS FOR NON-LINEAR FUNCTIONS

Consider a scalar univariate function $f(x)$ to be approximated over the input variable range $x \in [-1, 1]$. One of the most popular approximations for non-linear systems is the Volterra model. Given the function representation by the inverse Fourier transform,

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} F(w) e^{jwx} dw, \quad (2)$$

where w is the frequency variable and $F(w)$ are the Fourier coefficients, the Volterra model is obtained by using the Taylor series expansion of the exponential family:

$$\begin{aligned} f(x) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} F(w) \left(\sum_{n=0}^{\infty} \frac{(jw)^n}{n!} x^n \right) dw \\ &= \sum_{n=0}^{\infty} \underbrace{\left(\frac{1}{2\pi} \frac{1}{n!} \int_{-\infty}^{\infty} F(w) (jw)^n dw \right)}_{c_n} x^n \end{aligned} \quad (3)$$

The building blocks of (3) are power functions of the input variable, which generate a polynomial approximation. Notice that the coefficients c_n correspond to the n -th derivative at the origin divided by the factorial of the index. Volterra has not been widely used in practice because the exponential nature of the basis functions heavily increase the dynamic range, even when the input is bounded. Moreover, outside the dynamic range the approximation is unbounded and cannot be controlled. This happens because Volterra is a Taylor approximation and it presents high accuracy near the origin only. Moreover, the coefficients change with respect to how many are preserved, as the kernels

are not orthogonal. All these concerns make this approximation not suitable for gradient-based learning algorithms [10].

One possible approach to solve these issues is the representation using a finite number of terms in the Fourier approximation in (2), which corresponds to the Discrete Fourier Transform (DFT). Outside the dynamic range of x the function is periodically extended, generating discontinuities at the border. As a result, the number of coefficients required to approximate the function is highly sensitive to these discontinuities. To prevent this phenomenon the function can be extended with even symmetry and, then, periodically extended. This smooths out the edges and its corresponding derivatives, reducing the number of required coefficients with respect to the DFT. This is known as the Discrete Cosine Transform (DCT), which has the following expression:

$$f(x) \approx \sum_{q=1}^Q g_q F_q \cos\left(\frac{\pi(q-1)(2z+1)}{2N}\right), \quad (4)$$

with $z = \frac{N}{2}(1+x)$, $g_1 = 1/\sqrt{N}$ and $g_q = \sqrt{2/N}$ otherwise. The $F_q \in \mathbb{R}$ are termed the DCT coefficients. Notice that for functions with odd symmetry, only the odd coefficients are retained. In general, the quality of approximation is more than sufficient for $Q = 12$ (i.e., 6 coefficients in odd functions). For the sake of brevity, the following definition will be used when needed:

$$\cos_i(x) = \cos\left(\frac{\pi}{2N}(2i-1)(N(x+1)+1)\right) \quad (5)$$

In [10] we propose an adaptive design in which the DCT coefficients of (4) are tuned using the LMS algorithm in a supervised setting to approximate an univariate function. There are plenty of advantages in using the DCT representation: The required coefficients to provide the same quality (i.e., approximation error) are fewer and these are real and ordered in decreasing magnitude. Since the basis functions are orthogonal, the approximation error can be easily controlled by the magnitude of the disregarded coefficients, and it simplifies the convergence of the learning procedure. Furthermore, see that the coefficient index appears in the phase of (4), so the approximation is real and bounded, even when the input exceeds the dynamic range. All these features make the DCT an appropriate function approximation, whose coefficients can be learnt by a gradient-based rule.

A. Extension of the DCT to Multiple Input Variables

The extension of the DCT representation to multiple input variables is not trivial. Consider a bivariate function, which could be parameterized by a 2-dimensional (2D) DCT as

$$\begin{aligned} f(x_1, x_2) &= \sum_{n=1}^N \sum_{m=1}^N F_{nm} \cos_n(x_1) \cos_m(x_2) \\ &\approx \sum_{n=1}^Q \sum_{m=1}^Q F_{nm} \cos_n(x_1) \cos_m(x_2) \end{aligned} \quad (6)$$

Note that this model has the following drawbacks: In general the number of coefficients grows exponentially with the number of inputs, M_0 , as Q^{M_0} , as well as the complexity of the DCT;

while the coefficients in the DCT are ordered in decreasing magnitude, the structure of the indexes is broken in the 2D-DCT. This is, the location (n, m) of the relevant coefficients changes with the function of interest; this implies that, unless the function is known a priori, the indexes (n, m) are unknown and a supervised learning procedure is hard to implement. This happens for instance in classification problems, where the function is unknown and to be discovered.

In the following we will present ENN, a neural network model integrating the DCT in a single dimension and whose coefficients can be trained in a supervised fashion.

IV. DCT-ADAPTIVE ACTIVATION FUNCTIONS

A perceptron (or neuron) is a non-linear processor involving a weighted sum and a non-linear function. It is described by the following expressions:

$$z = a_0 + \sum_{i=1}^{M_0} a_i x_i = \mathbf{a}^T [1 \ \mathbf{x}^T]^T \quad (7)$$

$$\bar{z} = \frac{N}{2}(z+1) \quad (8)$$

$$\begin{aligned} \hat{y} = \sigma(\bar{z}) &\approx \sum_{q=1}^{Q/2} F_q \cos\left(\frac{\pi(2q-1)(2\bar{z}+1)}{2N}\right) \\ &= \sum_{q=1}^{Q/2} F_q \cos_q(z) \end{aligned} \quad (9)$$

In (7), \mathbf{x} contains M_0 inputs and a 1 is appended for the bias term a_0 . The normalization in (8) is needed to map the input to $[0, N]$, assuming the input is confined to $[-1, 1]$. Nevertheless, as it will be seen later, the first linear transformation may map \bar{z} to a different range, providing expressiveness to the network. Finally, the non-linear activation function $\sigma(\cdot)$ is approximated by the DCT with $Q/2$ coefficients. Without loss of generality, g_q is assumed to be integrated in the DCT coefficient F_q . In (9) we assume the non-linearity to have odd symmetry, so that only the odd coefficients are retained. As it will be explained later on in Section IV-B and shown empirically in Section VI, this does not prevent the network from learning only odd activation functions. A total of $M_0 + Q/2 + 1$ parameters per perceptron are to be trained, which only represents an increment of $Q/2$ coefficients with respect to a standard perceptron with a fixed activation functions. As mentioned in the previous section, Q is small due to the energy compaction property of the DCT (e.g., around $Q/2 = 6$ coefficients).

Example 1 (Linear discriminant): Assume $M_0 = 2$, we want to discriminate when $x_1 > x_2$. This can be framed as a classification problem, for which we can take $\mathbf{a} = a[0 \ 1 \ -1]^T$, where a is a positive constant. Then, any odd monotone increasing function will discriminate the two hypothesis. For the sake of simplicity, take $a = 1$ and the non-linearity approximated with only one coefficient (i.e., $F_1 = -1$). This results in

$$z = x_1 - x_2 \quad (10)$$

$$\bar{z} = \frac{N}{2}(x_1 - x_2 + 1) \quad (11)$$

$$\hat{y} = -\cos\left(\frac{\pi(2\bar{z} + 1)}{2N}\right) \approx \sin\left(\frac{\pi}{2}(x_1 - x_2)\right) \quad (12)$$

The solution in (12) provides a soft-decision, whereas a hard-decision would take the sign, i.e., $\text{sign}(\hat{y})$. While there are infinite solutions for this example, the proposed scheme only increases the complexity by one extra coefficient (i.e., F_1). Implementing a linear $\sigma(\cdot)$ would require many coefficients and provide no further benefit in terms of performance.

In Example 1 there are infinite minimums, all optimal. Nevertheless, in more complex problems (e.g., high-order discriminants) the expressiveness of a single perceptron is not enough. In light of the universal approximation theorem, in the following we will increase the capabilities of the network by including a hidden layer of several neurons. This model, termed ENN, will be fully-adaptive as both the linear weights and the activation functions will be trained in a supervised fashion.

A. Expressive Neural Network (ENN)

As a single layer perceptron limits the capabilities of the network, we will include another processing layer to increase its expressiveness. For a general multi-layer perceptron of L layers, the following expressions show the perceptron signals at the l -th layer:

$$\mathbf{z}_l = \mathbf{A}_l^T [1 \ \mathbf{s}_{l-1}^T]^T \quad (13)$$

$$\bar{\mathbf{z}}_l = \frac{N}{2}(\mathbf{z}_l + 1) \quad (14)$$

$$\mathbf{s}_l = \sigma_{l,m}(\bar{\mathbf{z}}_l) \quad (15)$$

for $l = 1, \dots, L$. The first input and last output correspond to $\mathbf{s}_0 = \mathbf{x}$ and $\hat{y} = \mathbf{s}_L$, respectively. The matrix of linear weights is $\mathbf{A}_l = [\mathbf{a}_l^{(1)} \ \dots \ \mathbf{a}_l^{(M_l)}]$ and M_l stands for the number of perceptrons at layer l , with M_0 being the number of inputs. The operations in (14) and (15) are performed element-wise. The AAF in the ENN, this is, the m -th element of (15) is computed as

$$\mathbf{s}_l[m] = \sum_{q=1}^{Q/2} F_{q,l}^{(m)} \cos_q(\mathbf{z}_l[m]), \quad (16)$$

where $F_{q,l}^{(m)}$ corresponds to the q -th coefficient of the m -th perceptron at the l -th layer. As explicitly shown in (15) and (16) the activation function is not necessarily the same at each neuron, although the number of DCT coefficients Q is kept constant for the whole network. The last activation function could be substituted by an step function in the classification setup and a linear for regression. However, this will be maintained so that the network can self-adapt depending on the nature of the problem.




Fig. 1. A 2-layer perceptron with M_1 neurons in the hidden layer.




Fig. 2. Sigmoid function (dashed) and its DCT representation in the $[-1, 1]$ range (solid).

As opposed to the M_0 -dimensional DCT in Section III-A, the proposed model does not require implementing the DCT in M_0 dimensions and can be trained in a supervised fashion. Furthermore, the number of coefficients grows linearly as M_0Q . Throughout the rest of the paper we assume a two input variable, namely $M_0 = 2$, and $L = 2$ layers. These layers are sequentially termed hidden and output layer. Fig. 1 shows the architecture of the two-layer perceptron with M_1 neurons in the hidden layer and expression (17) shown at the bottom of this page, shows the input-output relationship.

B. Expressiveness of Periodic Activation Functions

As it will be shown in the experimental results, the expressiveness of the DCT comes from its periodic nature. Fig. 2 shows the sigmoid function and the corresponding DCT representation in the $[-1, 1]$ range using $Q/2 = 6$ coefficients. While the sigmoid function saturates outside the range, the DCT approximation offers a periodic infinite non-linearity. This offers more capacity to the activation function, as the network may

$$\hat{y} = \sum_{m=1}^{Q/2} F_{m,2}^{(1)} \cos_m \left(a_2^{(1)}[0] + \sum_{k=1}^{M_1} a_2^{(1)}[k] \sum_{q=1}^{Q/2} F_{q,1}^{(k)} \cos_q \left(a_1^{(k)}[0] + a_1^{(k)}[1]x_1 + a_1^{(k)}[2]x_2 \right) \right) \quad (17)$$




Fig. 3. Bump generation. Mapping of a linear combination of two inputs over the response of a non-linear function.

choose to work at an increasing range (e.g., $[-1, 0]$), decreasing range (e.g., $[-3, -1]$), bump range (e.g., $[-1, 3]$), valley range (e.g., $[-3, 1]$), or even with several periods simultaneously (e.g., $[-5, 3]$). Imposing odd symmetry in the DCT representation does not constrain the resulting activation function to be odd. This Fourier model for non-linearities is the first one to provide such flexibility to activation functions, which is not possible with a fixed non-linearity not implemented with the DCT.

To understand how the non-linear activation function maps the input to the output space, we use the concept of *bump*. Consider a single neuron with $M_0 = 2$. Notice that all input pairs satisfying the following equality are mapped to the same non-linear output $\sigma(c)$:

$$a_0 + a_1x_1 + a_2x_2 = c, \quad (18)$$

where c is a constant value. Fig. 3(a) shows how expression (18) corresponds to a line in the input space. Fig. 3(b) shows the mapping $\sigma(c)$ along with the limits of the function in dashed lines, that correspond to extreme values of c . Fig. 3(c) shows how all input pairs with $\sigma(c)$ are mapped in the output space. The bump is generated by evaluating all the input data pairs through the linear transformation and the non-linear mapping.

V. SUPERVISED LEARNING

A. Backpropagation Rules

Consider a dataset \mathcal{D} , consisting of $(\mathbf{x}_i, y_i) \in \mathcal{D}$, where the former is a M_0 -sized vector training sample and the latter is the associated reference (i.e., label in classification, or function value in neuromorphic computing). Given y_i and \hat{y}_i , the error can be computed and propagated throughout the network to adjust the learnable parameters. We assume the loss to be the mean squared error (MSE),

$$\varepsilon^2 = (y_i - \hat{y}_i(\mathbf{x}_i))^2, \quad (19)$$

for both classification and regression problems. Notice in (19) we explicitly write the output as a function of the input data. Given a learnable parameter w , the chosen algorithm to update it is LMS. The update rule at a given iteration corresponds to

$$w \leftarrow w - \mu \nabla \varepsilon = w - \mu \varepsilon \left(-\frac{\partial \hat{y}}{\partial w} \right), \quad (20)$$

this is, the parameter is updated by the product of the error by the instantaneous gradient of the output with respect to the parameter. The hyperparameter μ is the step size, which controls the convergence speed of LMS. We choose the simplest algorithm to learn the parameters in the neural network because the focus

of this work is on the model, not on the training algorithm. In this respect, the experimental results show that even using the instantaneous gradient, LMS converges to a minimum and outperforms state-of-the-art models. Thus, the LMS may be replaced by any other algorithm to speed up convergence, which is out of the scope of this work. Due to the same reason, we do not include the iteration index in the parameter updates.

For the 2-layer architecture in Fig. 1, there are 2 set of parameters per layer that need to be updated, namely, the DCT coefficients and the weights from the linear transformation. Starting at the output, the DCT coefficients of the layer $l = 2$ are updated as

$$F_{m,2} \leftarrow F_{m,2} + \mu \varepsilon \frac{\partial \tilde{y}}{\partial F_{m,2}} = F_{m,2} + \frac{4\alpha_1}{Q} \varepsilon \cos_m(z_2), \quad (21)$$

for $m = 1, \dots, Q/2$. The superscript in the parameter is omitted because there is only one perceptron at the output. In general, the step size is modeled as twice the mismatch α_1 divided by the power of the corresponding input. Since the $Q/2$ cosines used to approximate the non-linearity are orthonormal, the total power is $Q/2$. This exhibits another advantage of this parametrization as the power remains always constant.

In order to update the linear weights from the output layer, the backpropagation procedure allows to propagate the error across the network with respect to previously computed derivatives, which makes it a very efficient algorithm. Therefore, these are updates as

$$\begin{aligned} a_2[k] &\leftarrow a_2[k] + \mu \varepsilon \frac{\partial \tilde{y}}{\partial a_2[k]} = a_2[k] + \mu \varepsilon \frac{\partial \tilde{y}}{\partial z_2} \frac{\partial z_2}{\partial a_2[k]} \\ &= a_2[k] - \frac{2\alpha_3}{P_1} \varepsilon \frac{\pi}{2} s_1[k] \sum_{m=1}^{Q/2} F_{m,2} (2m-1) \cos_m(z_2), \end{aligned} \quad (22)$$

for $k = 0, \dots, M_1$, and $P_1 = \mathbf{s}_1^T \mathbf{s}_1$. The superscript has also been suppressed in this case. Accordingly, the DCT coefficients in the first layer are updated as

$$\begin{aligned} F_{q,1}^{(k)} &\leftarrow F_{q,1}^{(k)} + \mu \varepsilon \frac{\partial \tilde{y}}{\partial F_{q,1}^{(k)}} = F_{q,1}^{(k)} + \mu \varepsilon \frac{\partial \tilde{y}}{\partial z_2} \frac{\partial z_2}{\partial s_1[k]} \frac{\partial s_1[k]}{\partial F_{q,1}^{(k)}} \\ &= F_{q,1}^{(k)} \\ &\quad - \frac{4\alpha}{Q} \varepsilon \frac{\pi}{2} a_2[k] \cos_q(z_1[k]) \sum_{m=1}^{Q/2} F_{m,2} (2m-1) \sin_m(z_2), \end{aligned} \quad (23)$$

for $q = 1, \dots, Q/2$ and $k = 1, \dots, M_1$. And the linear combination of the first layer is updates as

$$\begin{aligned} a_1^{(k)}[m] &\leftarrow a_1^{(k)}[m] + \mu \varepsilon \frac{\partial \tilde{y}}{\partial a_1^{(k)}[m]} \\ &= a_1^{(k)}[m] + \mu \varepsilon \frac{\partial \tilde{y}}{\partial z_2} \frac{\partial z_2}{\partial s_1[k]} \frac{\partial s_1[k]}{\partial z_1[k]} \frac{\partial z_1[k]}{\partial a_1^{(k)}[m]} \\ &= a_1^{(k)}[m] + \end{aligned}$$

$$\frac{2\alpha_4}{P_0} \varepsilon \frac{\pi^2}{4} a_2[k] s_0[m] \sum_{p=1}^{Q/2} F_{p,2}(2p-1) \sin_p(z_2)$$

$$\sum_{q=1}^{Q/2} F_{q,1}^{(k)}(2q-1) \sin_q(z_1[k]), \quad (24)$$

for $m = 1, \dots, M_0$, $k = 1, \dots, M_1$, and $P_0 = \mathbf{s}_0^T \mathbf{s}_0$. Without loss of generality, $s_0[0] = s_1[0] = 1$, which correspond to the bias terms in the linear combination at each layer. Notice that in all the update rules, the derivative exists and inherits the periodic nature of the cosine transform.

To enhance the stability of the convergence, the power of each input, namely P_0 and P_1 , can be computed with a damping effect. For instance,

$$P_0 \leftarrow \beta P_0 + (1 - \beta) \mathbf{s}_0^T \mathbf{s}_0 \quad (25)$$

at each iteration. The α parameter is set smaller in the output layer, which intuitively reduces the propagation of artifacts to previous layers during training.

B. Initialization and Training Procedure

The ENN trains all parameters at every backpropagation step, this is, both AAF and linear weights. The AAF are initialized as linear functions, which has shown to provide the best performance. In the hidden layer, the linear weights are initialized to generate bump diversity. As seen in Fig. 3, linear weights can be interpreted as the orientation of the activation functions (or bump) in the output space. For instance, with $M_1 = 4$ the linear weights are initialized to $[0, 0, 1]$, $[0, 1, 0]$, $[0, 1, -1]$ and $[0, -1, 1]$. This helps in practice to start the learning process with bumps in different orientations. Finally, the output layer is initialized with linear weights in a uniform distribution in $[-0.5, 0.5]$. This allows to constrain the dynamics and shows that the convergence does not depend on the initialization.

Regarding the LMS setup, for the non-linearities we set $\alpha_1 = 10^{-3}$ in the hidden layer and $\alpha_2 = 10^{-4}$ in the output layer. For the linear weights we set $\alpha_3 = 5 \cdot 10^{-3}$ in the hidden layer and $\alpha_4 = 5 \cdot 10^{-5}$ in the output layer. We keep the step-size constant for all neurons in the same layer. The damping parameter is set to $\beta = 0.999$ for both layers. Under the assumption that the error between parameters is uncorrelated, this LMS configuration guarantees that

$$6\alpha_1 + 6\alpha_2 + \alpha_3 + \alpha_4 \approx 0.01 < 0.1, \quad (26)$$

where the right hand side is an upper bound on the error at convergence. Notice that the convergence speed at the output layer is lower than at the hidden layer, which grants an adequate gradient propagation.

VI. EXPERIMENTAL RESULTS

In the following we will test the ENN for both binary classification and regression problems. As benchmarks, we will test also the following models, in which they all have the same architecture and differ in the activation functions:

- *ReLU*: It uses a fixed non-trainable ReLU, this is, $\sigma(z) = \max\{0, z\}$.
- *Sigm*: It uses a fixed non-trainable sigmoid activation function. Notice that the output saturates for an input of large magnitude.
- *F-DCT*: It uses a fixed non-trainable sigmoid activation function, but modeled with the DCT. Thus, the function does not saturate and it is periodic.

Although not an odd function, we include the ReLU because it is the standard activation function in current neural networks. The Sigm model is also included to compare it with F-DCT and assess the gains of a periodic activation function. We acknowledge that the benchmarks have significantly less parameters than the ENN and they will exhibit lower learning capabilities. However, the scope of this work is to show the expressiveness through the adaptability of non-linear functions instead of increasing the width of the architecture. We constrain the number of neurons because we work on interpretability and want to provide a signal processing perspective to the problem.

All models are built with $M_1 = 6$ neurons in the hidden layer, and ENN with $Q = 12$ parameters in all the non-linearities and $N = 512$ samples. However, recall that only $Q/2$ coefficients are different from zero because we impose odd symmetry. These are initialized to approximate the identity function. Notice that the output activation function in ReLU and Sigm has to be specifically selected to be either sigmoid in classification or linear (i.e., identity) in regression. Conversely, the ENN will automatically adapt it to approach the required function, which is an advantage of this model with respect to the benchmarks.

For both classification and regression problems, a synthetic dataset is generated. All samples come from independent uniform distributions in the $[-1, 1]$ range for each input variable. The train and test sets contain 800.000 and 50.000 samples, respectively.

A. Classification

Table I shows the different classification problems that have been evaluated. It displays the ideal decision map along with the order of the discriminant. The metric chosen to compare the different models is the test accuracy, this is, the percentage of correctly classified samples in the test set.

As expected, the complexity of the problem increases with the order of the discriminant. As seen in Example 1, a single neuron is sufficient to implement a linear discriminant, which results in an excellent performance for all models in (P1). With no surprise, for quadratic (P2) and cubic (P3) discriminants all models are capable of finding the appropriate boundary. However, ENN attains almost an excellent performance. This may be claimed due to the larger number of parameters in the proposed model with respect to the benchmarks. When moving to high-order classification problems is when the non-adaptive models are not capable of providing a satisfactory solution. Conversely, the ENN model manages to find a local minimum with high accuracy and with very consistent results for a wide variety of problems.

Fig. 4 shows the AAF for (P7). Notice that only one of the activation functions has been adapted, namely Fig. 4(b), which

TABLE I
ACCURACY FOR DIFFERENT BINARY CLASSIFICATION PROBLEMS

N°	Map	Order	ReLU	Sigm	F-DCT	ENN
(P1)		1	99.97%	99.90%	99.84%	99.85%
(P2)		2	97.62%	97.82%	97.43%	99.85%
(P3)		3	96.70%	96.74%	96.17%	99.90%
(P4)		> 3	95.86%	96.00%	87.36%	99.16%
(P5)			91.76%	89.68%	89.87%	99.48%
(P6)			75.72%	75.20%	75.20%	99.19%
(P7)			51.93%	54.56%	54.22%	99.62%
(P8)			75.68%	80.16%	79.32%	98.22%

For each case, the ideal map and the order of the discriminant are shown.




Fig. 4. AAF in the hidden layer for (P7).

exploits the periodic nature of the DCT. The corresponding bumps are shown in Fig. 5, where we clearly confirm that the other AAF have a dynamic range around $[-0.5, 0.5]$, which results in zero bumps and, thus, in non-active neurons. The global response of the ENN is shown in Fig. 6(a). This is obtained by adding the bumps weighted by the corresponding linear weights of the output layer and, finally, applying the output AAF. As expected, the output AAF for classification problems is an approximation of the step function and it is shown in Fig. 10(a). Since there is only one neuron active, the global response corresponds to the step (or sign) function applied to the bump in Fig. 4(b). Notice, however, that there is a flip between the bump and the global response, which happens because the corresponding linear weight of the output layer is negative. The fact of constraining the model to $Q = 12$ coefficients makes it infeasible to generate an ideal step function in the output layer. Nevertheless, this solution is steeper than a standard sigmoid function, which reduces the error at data close

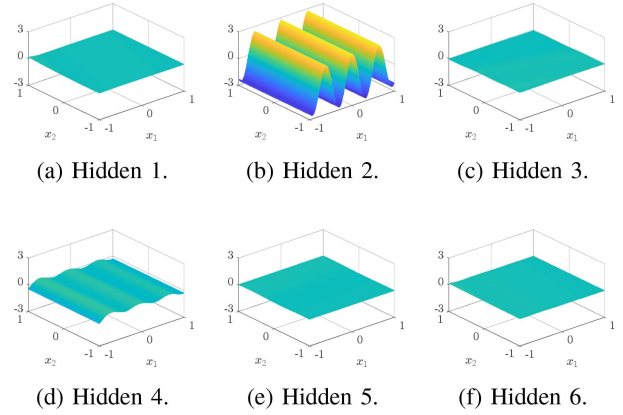


Fig. 5. Bumps corresponding to the AAF in Fig. 4.

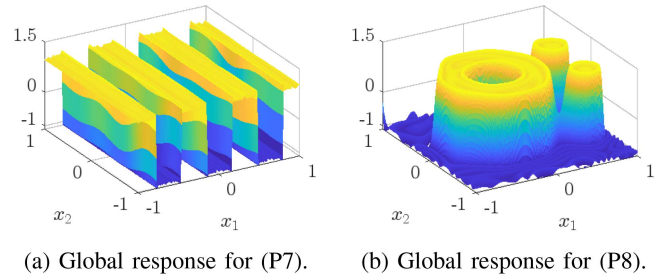


Fig. 6. Global response of the ENN for two classification problems.




Fig. 7. AAF in the hidden layer for (P8).

to the boundary and increases the accuracy. From this experiment we can also conclude that another advantage of ENN is that it allows to model the network width, this is, the required number of hidden neurons. In other experiments we found out that the ENN generates pairs of identical bumps but with opposite sign in the output linear weights. This means that the network neglects the information that comes from these two branches. See that this is impossible for the non-adaptive models, which have a very small accuracy even with 6 neurons.

Fig. 7 shows the AAF for (P8), along with the corresponding bumps in Fig. 8. As expected, the ENN exploits the periodic nature of the DCT and different orientations to generate diversity. Notice that there is almost no drop in accuracy between (P5)

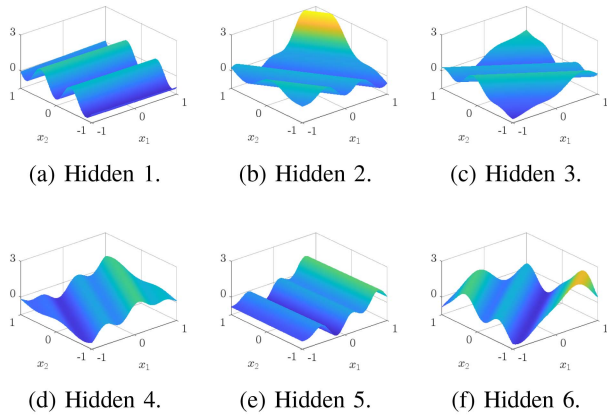


Fig. 8. Bumps corresponding to the AAF in Fig. 7.




Fig. 9. Output decision map of (P8) for each of the models.




Fig. 10. Output AAF for classification and regression.

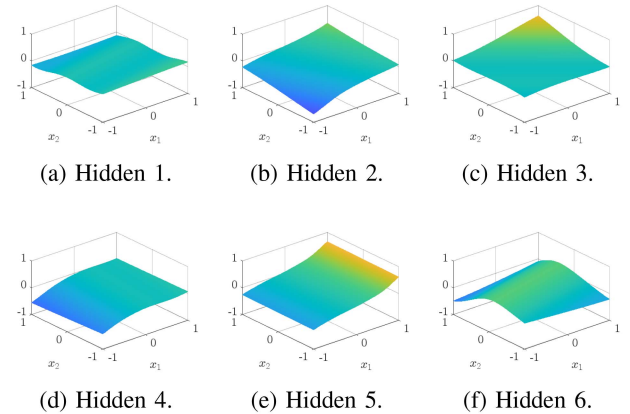
and (P8), although the latter is more complex. This is due to the expressiveness of the DCT, which allows to generate several peaks per bump. In other words, by allowing the activation to be non-monotone it reuses non-linearities to approximate the ring and both circles still with 6 neurons. Fig. 6(b) shows the global response of the network and Fig. 9 shows the different decision boundaries for the different models. Despite the accuracy of the non-adaptive models in Table I being around 80%, the quality is very poor and only the ENN boundary recovers the ideal map. The output activation function converges to the same configuration for all classification problems and it is shown in Fig. 10(a).

B. Regression

The network can also be trained for regression problems, this is, to approximate a function. Particularly, we compute bivariate scalar functions. Table II shows the MSE achieved by the four different models for three different experiments. While it may be impractical to implement a sum or product with a neural network as the model itself contains many sums and products,

TABLE II
MSE IN REGRESSION TASKS

Func.	ReLU	Sigm	F-DCT	ENN
$\frac{x_1 + x_2}{2}$	$2.9 \cdot 10^{-3}$	$2.1 \cdot 10^{-2}$	$3.5 \cdot 10^{-3}$	$2.3 \cdot 10^{-6}$
$\frac{x_1^2 + x_2^2}{2}$	$2.2 \cdot 10^{-1}$	$2.2 \cdot 10^{-1}$	$2.2 \cdot 10^{-1}$	$9.5 \cdot 10^{-6}$
$x_1 x_2$	$5.1 \cdot 10^{-2}$	$3.6 \cdot 10^{-2}$	$1.6 \cdot 10^{-1}$	$7.6 \cdot 10^{-5}$

Fig. 11. Bumps for the regression problem $(x_1^2 + x_2^2)/2$.

these experiments allow to test the learning capabilities of the ENN with respect to the benchmarks.

As expected, the gap in performance with respect to the benchmarks is larger in the second and third problem, since these are non-linear problems. Fig. 11 shows the different bumps for the regression problem $(x_1^2 + x_2^2)/2$. Since the function is symmetric around the center, the ENN chooses to work with non-periodic ranges of the non-linearities. In fact, all bumps correspond to flat surfaces that curve toward the edges. Recall that the curvature of the bump is undetermined as it also depends on the linear weights of the output layer. In fact, the weights of neurons 2, 4 and 6 are negative, meaning that all bumps are convex functions, as they need to be to generate a convex function.

Regarding the output activation, it is adapted to the learning task and different from classification. As seen in Fig. 10(b) the activation function approaches a linear function for all regression problems. The function is not perfectly linear because we constrain the system to $Q = 12$ coefficients. We underline the idea that in a non-adaptive setting one needs to choose the last non-linearity according to the task. This represents an improvement with respect to non-adaptive models as it is general enough to encompass different tasks.

C. Effect of the Number of Coefficients

Increasing the number of coefficients Q in the DCT representation theoretically provides more expressiveness to the network. However, in practice, increasing the number of trainable parameters hinders the convergence of the algorithm because there are more local minima. Fig. 12 shows the output AAF for $Q = \{12, 20\}$. Increasing the number of coefficients allows to create a steeper transition, reducing the number of errors around




Fig. 12. Output AAF in classification for different Q values.

the boundary. Doing so increases the Gibbs effect, although it produces no harm to the learning process because those samples are far from the boundary and always correctly classified as either 1 or -1 .

VII. CONCLUSION

In this article we have presented ENN, a novel neural network model with adaptive activation functions. Under a signal processing motivation we use the DCT to design the non-linear functions, whose coefficients can be trained using backpropagation. Specifically, the ENN is able to adapt the activation functions for classification and regression problems. We provide insights in the interpretability of the network by recovering the concept of bump, this is, the response of each neuron in the output space. Through extensive experiments we determine that the expressiveness of a neural network highly depends on the activation function. Particularly, the key strength of ENN is the periodic nature of the DCT, providing high accuracy for a wide range of non-linear classification problems. In some cases outperforming state-of-the-art non-adaptive activation functions up to 40% in accuracy.

REFERENCES

- [1] S. Chen, C. F. N. Cowan, and P. M. Grant, "Orthogonal least squares learning algorithm for radial basis function networks," *IEEE Trans. Neural Netw.*, vol. 2, no. 2, pp. 302–309, Mar. 1991.
- [2] S.-S. Yang and C.-S. Tseng, "An orthogonal neural network for function approximation," *IEEE Trans. Syst., Man, Cybern., Part B.*, vol. 26, no. 5, pp. 779–785, Oct. 1996.
- [3] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. IEEE 27th Asilomar Conf. Signals, Syst. Comput.*, 1993, pp. 40–44.
- [4] O. L. Mangasarian and E. W. Wild, "Nonlinear knowledge in kernel approximation," *IEEE Trans. Neural Netw.*, vol. 18, no. 1, pp. 300–306, Jan. 2007.
- [5] S. Liang and R. Srikant, "Why deep neural networks for function approximation?," in *Proc. 5th Int. Conf. Learn. Representations*, 2017.
- [6] S. Yang, T. Ting, K. Man, and S.-U. Guan, "Investigation of neural networks for function approximation," *Procedia Comput. Sci.*, vol. 17, pp. 586–594, 2013.
- [7] S. Ferrari and R. Stengel, "Smooth function approximation using neural networks," *IEEE Trans. Neural Netw.*, vol. 16, no. 1, pp. 24–38, Jan. 2005.
- [8] C.-T. Chen and W.-D. Chang, "A feedforward neural network with function shape autotuning," *Neural Netw.*, vol. 9, no. 4, pp. 627–641, 1996.
- [9] M. M. Gost, A. Pérez-Neira, and M. A. Lagunas, "DCT-based air interface design for function computation," *IEEE Open J. Signal Process.*, vol. 4, pp. 44–51, 2023.
- [10] A. Pérez-Neira, M. Martinez-Gost, and M. A. Lagunas, "Adaptive function approximation based on the DCT," in *Proc. IEEE Int. Conf. Circuits, Syst., Commun. Comput.*, 2023, pp. 224–231.
- [11] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. New York, USA: Wiley, 2001.
- [12] K. Hornik, "Approximation capabilities of multilayer feedforward networks," *Neural Netw.*, vol. 4, no. 2, pp. 251–257, 1991.
- [13] M. Leshno, V. Y. Lin, A. Pinkus, and S. Schocken, "Multilayer feedforward networks with a nonpolynomial activation function can approximate any function," *Neural Netw.*, vol. 6, no. 6, pp. 861–867, 1993.
- [14] A. Kolmogorov, "The representation of continuous functions of many variables by superposition of continuous functions of one variable and addition," *Doklady Akademii Nauk SSSR*, vol. 114, pp. 953–956, 1957.
- [15] R. Hecht-Nielsen, "Kolmogorov's mapping neural network existence theorem," in *Proc. IEEE Int. Conf. Neural Netw.*, 1987, pp. 11–14.
- [16] J. Schmidt-Hieber, "The Kolmogorov-arnold representation theorem revisited," *Neural Netw.*, vol. 137, pp. 119–126, 2021.
- [17] M. Nakamura, R. Mines, and V. Kreinovich, "Guaranteed intervals for Kolmogorov's theorem (and their possible relation to neural networks)," *Interval Comput.*, vol. 3, pp. 183–199, 2004.
- [18] P. Campolucci, F. Capperelli, S. Guarnieri, F. Piazza, and A. Uncini, "Neural networks with adaptive spline activation function," in *Proc. IEEE 8th Mediterranean Electrotechnical Conf. Ind. Appl. Power Syst., Comput. Sci. Telecommun.*, 1996, pp. 1442–1445.
- [19] L. Vecchi, F. Piazza, and A. Uncini, "Learning and approximation capabilities of adaptive spline activation function neural networks," *Neural Netw.*, vol. 11, no. 2, pp. 259–270, Mar. 1998.
- [20] S. Xu and M. Zhang, "Justification of a neuron-adaptive activation function," in *Proc. IEEE-INNS-ENNS Int. Joint Conf. Neural Netw. IJCNN, Neural Comput.: New Challenges Perspectives New Millennium*, 2000, pp. 465–470.
- [21] S. Xu and M. Zhang, "A novel adaptive activation function," in *Proc. IEEE Int. Joint Conf. Neural Netw.*, 2001, pp. 2779–2782.
- [22] G. S. da S. Gomes, T. B. Ludermir, and L. M. Almeida, "Neural networks with asymmetric activation function for function approximation," in *Proc. IEEE Int. Joint Conf. Neural Netw.*, 2009, pp. 980–987.
- [23] F. Agostinelli, M. Hoffman, P. Sadowski, and P. Baldi, "Learning activation functions to improve deep neural networks," in *Proc. 3rd Int. Conf. Learn. Representations*, 2014.
- [24] X. Jin, C. Xu, J. Feng, Y. Wei, J. Xiong, and S. Yan, "Deep learning with s-shaped rectified linear activation units," in *Proc. 30th AAAI Conf. Artif. Intell.*, 2016, pp. 1737–1743.
- [25] L. Hou, D. Samaras, T. Kurc, Y. Gao, and J. Saltz, "ConvNets with smooth adaptive activation functions for regression," in *Proc. 20th Int. Conf. Artif. Intell. Statist.*, vol. 54, 2017, pp. 430–439.
- [26] D. Li and Y. Zhou, "Soft-root-sign: A new bounded neural activation function," in *Proc. Pattern Recognit. Comput. Vis.*, 2020, pp. 310–319.
- [27] P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for activation functions," in *Proc. 6th Int. Conf. Learn. Representations*, 2018.
- [28] F. Manessi and A. Rozza, "Learning combinations of activation functions," in *Proc. IEEE 24th Int. Conf. Pattern Recognit.*, 2018, pp. 61–66.
- [29] S. Qian, H. Liu, C. Liu, S. Wu, and H. S. Wong, "Adaptive activation functions in convolutional neural networks," *Neurocomputing*, vol. 272, pp. 204–212, 2018.
- [30] M. Varshney and P. Singh, "Optimizing nonlinear activation function for convolutional neural networks," *Signal, Image Video Process.*, vol. 15, no. 6, pp. 1323–1330, Sep. 2021.
- [31] S. Lee, H.-Y. Kwak, and J.-S. No, "Effect of the period of the fourier series approximation for binarized neural network," in *Proc. IEEE Int. Conf. Artif. Intell. Inf. Commun.*, 2022, pp. 262–265.
- [32] J. Yang and M. Ma, "The numerical solution of one-dimensional discrete asset pricing model based on the improved trigonometric extreme learning machine," *J. Comput. Sci.*, vol. 63, 2022, Art. no. 101809.
- [33] M. Ma, J. Yang, and R. Liu, "A novel structure automatic-determined Fourier extreme learning machine for generalized Black-Scholes partial differential equation," *Knowl.-Based Syst.*, vol. 238, 2022, Art. no. 107904.
- [34] T. Zhou, X. Liu, M. Hou, and C. Liu, "Numerical solution for ruin probability of continuous time model based on neural network algorithm," *Neurocomputing*, vol. 331, pp. 67–76, 2019.