

NEWRON: A New Generalization of the Artificial Neuron to Enhance the Interpretability of Neural Networks

Federico Siciliano,¹ Maria Sofia Bucarelli,¹ Gabriele Tolomei,² Fabrizio Silvestri¹

¹ Department of Computer, Control and Management Engineering Antonio Ruberti, Sapienza University of Rome, Italy

² Department of Computer Science, Sapienza University of Rome, Italy

federico.siciliano@uniroma1.it, mariasofia.bucarelli@uniroma1.it, tolomei@di.uniroma1.it, fsilvestri@diag.uniroma1.it

Abstract

In this work, we formulate NEWRON: a generalization of the McCulloch-Pitts neuron structure. This new framework aims to explore additional desirable properties of artificial neurons. We show that some specializations of NEWRON allow the network to be interpretable with no change in their expressiveness. By just inspecting the models produced by our NEWRON-based networks, we can understand the rules governing the task. Extensive experiments show that the quality of the generated models is better than traditional interpretable models and in line or better than standard neural networks.

1 Introduction

Neural Networks (NNs) have now become the *de facto* standard in most Artificial Intelligence (AI) applications. The world of Machine Learning has moved towards Deep Learning, i.e., a class of NN models that exploit the use of multiple layers in the network to obtain the highest performance.

Research in this field has focused on methods to increase the performance of NNs, in particular on which activation functions (Apicella et al. 2021) or optimization method (Sun et al. 2019) would be best. Higher performances come at a price: (Arrieta et al. 2020) show that there is a trade-off between interpretability and accuracy of models. Explainable Artificial Intelligence (XAI) is a rapidly growing research area producing methods to interpret the output of AI models in order to improve their robustness and safety (see e.g. (Ghorbani, Abid, and Zou 2019) and (Bhatt, Ravikumar et al. 2019)). Deep Neural Networks (DNNs) offer the highest performance at the price of the lowest possible interpretability. It is an open challenge to attain such high performance without giving up on model interpretability.

The simplest solution would be to use a less complex model that is natively interpretable, e.g., decision trees or linear models, but those models are usually less effective than NNs. We ask the following question: can we design a novel neural network structure that makes the whole model interpretable without sacrificing effectiveness?

NNs are black-box models: we can only observe their input and output values with no clear understanding of how those two values are correlated according to the model's parameters. Although a single neuron in the NN performs a relatively simple linear combination of the inputs, there is no clear and straightforward link between the parameters

estimated during the training and the functioning of the network, mainly because of the stacking of multiple layers and non-linearities.

In this work, we propose a generalization of the standard neuron used in neural networks that can also represent new configurations of the artificial neuron. Thus, we discuss a specific example that allows us to interpret the functioning of the network itself.

We focus our efforts on tabular data since we investigate how NEWRON works only in the case of fully connected NNs. It is more straightforward to produce human-readable rules for this kind of data. We also remark that our goal is not to improve the performance of NNs, but rather to create interpretable versions of NNs that perform as well as other interpretable models (e.g., linear/logistic regression, decision trees, etc.) and similarly to standard NNs, when trained on the same data.

Motivating Example

Consider a simple dataset: MONK's¹. Each sample consists of 6 attributes, which take integer values between 1 and 4 and a class label determined by a decision rule based on the 6 attributes. For example, in MONK-2, the rule that defines the class for each sample is the following: “*exactly two*” out of the six attributes are equal to 1.

It is impossible to intuitively recover rules from the parameter setting from a traditional, fully connected NN.

We shall see in the following that our main idea is that of inverting the activation and aggregation. In NEWRON the nonlinearity directly operates on the input of the neuron. The nonlinearity acts as a thresholding function to the input, making it directly interpretable as a (fuzzy) logical rule by inspecting its parameters. Consider the following network, represented in Figure 1: 2 hidden layers, the first with 1 neuron, the second with 2 neurons, and 1 output neuron. The x_i 's are the inputs of the model, y is the output.

We present the form of a typical architecture composed by NEWRON in Figure 1. We show how we can interpret the parameters obtained from a trained network. The rectangles represent the plot of a function that divides the input domain into two intervals, separated by the number below the rectangle, taking values 1 and 0.

¹<https://archive.ics.uci.edu/ml/datasets/MONK%27s+Problems>




Figure 1: An example of a network for the MONK-2 dataset. x_i are the inputs, y is the output. The red and blue rectangles represent the plot of functions, with input range on the x -axis and output on the y -axis. The green rectangles contain the aggregation function. The numbers in bold represent the thresholds for the step functions.

The functions that process the input give output 1 only if the input is less than 1.1, given that inputs are integers and assume values only in $\{1, 2, 3, 4\}$, this means “if $x_i = 1$ ”. The sum of the output of all these functions, depicted in the green rectangle, then represents the degree of soundness of those rules are.

The second layer has two neurons: the first outputs 1 if it receives an input greater than 1.9, i.e. if at least 2 of the rules $x_i = 1$ are valid, while the second outputs 1 if it receives an input less than 2.1, i.e. if 2 or less of the rules $x_i = 1$ are valid. Notice that the two neurons are activated simultaneously only if $x_i = 1$ is true for exactly two attributes.

In the last layer, functions in the blue rectangles receive values in $\{0, 1\}$ and do not operate any transformation, keeping the activation rules unchanged. The sum of the outputs of these functions is then passed to the function in the red rectangle. This function outputs 1 only if the input is greater than 1.9. Since the sum is limited in $0, 1, 2$, this happens only when it receives 2 as input, which occurs only if the two central neurons are activated. As we have seen, this only applies if exactly 2 of the rules $x_i = 1$ are valid.

So we can conclude that the network gives output 1 just if “*exactly two*” of $\{x_1 = 1, x_2 = 1, x_3 = 1, x_4 = 1, x_5 = 1, x_6 = 1\}$ are true.

Contributions

The main contributions of this work are the following:

- We propose NEWRON, a generalization of the McCulloch-Pitt neuron allowing the definition of new artificial neurons. We show how special cases of NEWRON may pave the way towards interpretable, white-box neural networks.
- We prove the universal approximation theorem for three specializations of NEWRON, demonstrating that the new model does not lose any representation power in those cases.
- We experiment on several tabular datasets showing that NEWRON allows learning accurate Neural models, beating interpretable by design models such as Decision Trees and Logistic Regression.

2 Related Work

(Rosenblatt 1958) introduced the single artificial neuron: the

Perceptron. The Perceptron resembles the functioning of the human/biological neuron, where the signal passing through the neuron depends on the intensity of the received signal, the strength of the synapses, and the receiving neuron’s threshold. In the same way, the Perceptron makes a linear combination of the inputs received and is only activated if the result exceeds a certain threshold. Over the years, various improvements to neural networks have been proposed: Recurrent Units, Convolutional Layers, and Graph Neural Networks, but for Fully Connected NNs, research efforts have mainly focused on finding more efficient activation functions (Apicella et al. 2021). Two works that have focused on modifying the internal structure of the neuron are those of (Kulkarni and Venayagamoorthy 2009), and (Fan, Cong, and Wang 2018). In the former, a neuron is introduced that performs both a sum and a product of the inputs in parallel, applies a possibly different activation function for the two results, and then sums the two outcomes. Despite promising results, given the use of fewer parameters, better performance, and reduced training time compared to standard MLPs and RNNs, the proposed neuron, rather than being a generalization, is a kind of union between two standard neurons, one of which uses the product, instead of sum, as aggregation function. In the second paper, starting from the notion that the traditional neuron performs a first-order Taylor approximation, the authors propose a neuron using a second-order Taylor approximation. Although this improves the capacity of a single neuron, the authors do not demonstrate any gains in terms of training time or convergence. Indeed, this can be considered a particular case of the higher-order neural units (HONUs) (see, e.g., (Gupta et al. 2013)), i.e., a type of neurons that, by increasing the degree of the polynomial computed within them, try to capture the higher-order correlation between the input patterns. Recent works that focus on interpretation at neuron level ((Dalvi et al. 2019a), (Dalvi et al. 2019b), (Heo et al. 2019), (Nam et al. 2020)) often concentrate on extracting the most relevant neurons for a given task, but mostly deal with Recurrent or Convolutional neural networks. Although not designing an alternative version of the neuron, (Yang, Morillo, and Hospedales 2018) proposes an alternative neural network structure, based on a Binning Layer, which divides the single input features into several bins, and a Kronecker Product Layer, which takes

into account all the possible combinations between bins. The parameters estimated during training can be interpreted to translate the network into a decision tree through a clever design of the equations defining the network. Although in-terpretable, the main issue in this work is its scalability. The Kronecker Product Layer has an exponential complexity that makes training time unfeasible when the number of features grows.

3 The NEWRON Structure

A neuron, in the classical and more general case, is represented by the equation $y = f(b + \sum_{i=1}^n w_i x_i)$.




Figure 2: Structure of the standard artificial neuron. w_i and b are respectively weights and bias. f is the activation function. x_i 's are the inputs and y is the output.

b is called the bias, w_i are the weights, and x_i s are the inputs. f represents the activation function of the neuron. Usually, we use the sigmoid, hyperbolic tangent, or ReLU functions.

We first generalize the above equation, introducing NEWRON as follows:

$$y = f(G_{i=1}^n(h_i(x_i))) \quad (1)$$

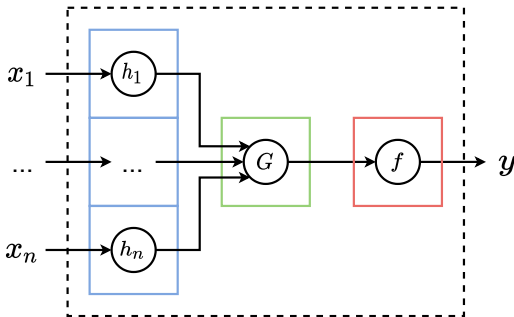


Figure 3: Structure of NEWRON, the generalized artificial neuron. The blue rectangles represent the processing function sections, the green rectangles contain the aggregation function, and the red rectangles represent the activation part. Same colors are also used in Figure 2

Each input is first passed through a function h_i , which we will call *processing function*, where the dependence on i indicates different parameters for each input. G , instead, represents a generic aggregation function.

Using NEWRON notation, the standard artificial neuron would consist of the following: $h_i(x_i) = w_i x_i$, $G = \sum_{i=1}^n$, and $f(z) = f^*(z + b)$.

G does not have any parameters, while b parametrizes the activation function.

Inverted Artificial Neuron (IAN)

We present 3 novel structures characterized by an inversion of the aggregation and activation functions. We name this architectural pattern: Inverted Artificial Neuron (IAN). In all the cases we consider the sum as the aggregation function and do not use any activation function: $G = \sum$, and $f(z) = z$.

Heaviside IAN The first case we consider uses a unit step function as activation. This function, also called the Heaviside function, is expressed by the following equation:

$$H(x) = \begin{cases} 1 & x \geq 0 \\ 0 & x < 0 \end{cases} \quad (2)$$

According to (1) we can define the processing function as follows:

$$h(x_i) = H(w_i(x_i - b_i)) = \begin{cases} H(w_i) & x_i \geq b_i \\ 1 - H(w_i) & x_i < b_i \end{cases} \quad (3)$$

where w_i and b_i are trainable parameters.

Sigmoid IAN We cannot train the Heaviside function using gradient descent, and it represents a decision rule that in some cases is too restrictive and not “fuzzy” enough to deal with constraints that are not clear-cut.

A natural evolution of the unit step function is therefore the sigmoid function $\sigma(x) = \frac{1}{1+e^{-x}}$. This function ranges in the interval $(0, 1)$, is constrained by a pair of horizontal asymptotes, is monotonic and has exactly one inflection point.

The sigmoid function can be used as a processing function with the following parameters: $h(x_i) = \sigma(w_i(x_i - b_i))$.

Product of tanh IAN Another option we consider as a processing function is the multiplication of hyperbolic tangent (\tanh). For simplicity, we will use the term “tanh-prod”.

The \tanh function $\tanh(x) = \frac{e^{2x}-1}{e^{2x}+1}$ is on its own very similar to the sigmoid. An interesting architecture is that using M \tanh simultaneously. Each \tanh applies its own weights, on each individual input.

While the sigmoid is monotonic with only one inflection point, roughly dividing the input space into two sections, the multiplication of \tanh , by being not monotonic, allows us to divide the input space into several intervals. The multiplication would remain in $(-1, 1)$, but can be easily rescaled to $(0, 1)$.

We can therefore write the processing function in the case of the \tanh multiplication as follows:

$$h(x_i) = \frac{\left(\prod_{m=1}^M \tanh(w_{im}(x_i - b_{im}))\right) + 1}{2} \quad (4)$$

Note how, in this case, the weights depend on both the input i and the m -th function. Such a neuron will therefore have M times more parameters than the Heaviside and sigmoid cases.

Output layer The output layer would produce values ranging in the interval $(0, N)$ ($\{0, 1, \dots, N\}$ for the Heaviside case), where N represents the number of neurons in the penultimate layer. This is because the last neuron makes the sum of N processing functions restricted in the interval $(0, 1)$ ($\{0, 1\}$ for the Heaviside case). To allow the last layer to have a wider output range and thus make our network able to reproduce a wider range of functions, we modify the last layer processing function h^* as follows: $h^*(x_i) = \alpha_i h(x_i)$, where α_i are trainable parameters.

In the same way, as for a traditional neural network, it is important, in the output layer, to choose an adequate activation function. We need, indeed, to match the range of the output of the network and the range of the target variable. In particular, in the case of output in $(0, 1)$, we use a sigmoid centered in b^* : $f^*(z) = \sigma(z - b^*)$

In the case of a classification problem with more than 2 classes, a softmax function ($s(z_j) = \frac{e^{z_j}}{\sum_i e^{z_i}}$) is used to output probabilities.

Note(s) The writing $w(x - b)$ is theoretically identical to that $w^*x + b^*$, where simply $w^* = w$ and $b^* = -bw$. This notation allows us to interpret the weights directly. From b , we already know the inflection point of the sigmoid; while looking at w , we immediately understand its direction.

4 Interpretability

(Arrieta et al. 2020) presented a well-structured overview of concepts and definitions in the context of Explainable Artificial Intelligence (XAI).

They make a distinction among the various terms that are mistakenly used as synonyms for interpretability. According to them:

- **Interpretability:** is seen as a passive feature of the model and represents the ability of a human to understand the underlying functioning of a decision model, focusing more on the cause-effect relationship between input and output.
- **Transparency:** very similar to interpretability, as it represents the ability of a model to have a certain degree of interpretability. There are three categories of transparency, representing the domains in which a model is interpretable. Simulatable models can be emulated even by a human. Decomposable models must be explainable in their individual parts. For algorithmically transparent models, the user can understand the entire process followed by an algorithm to generate the model parameters and how the model produces an output from the input.
- **Explainability:** can be seen as an active feature of a model, encompassing all actions that can detail the inner workings of a model. The explanation represents a kind of interface between a human and the model and must at the same time represent well the functioning of the model and be understandable by humans.

In this paper, we show decomposable models that, in some cases, are also algorithmically transparent.

Heaviside

The interpretability of an architecture composed of Heaviside IANs has to be analyzed by discussing its four main sections separately.

First layer - Processing function A single processing function $h(x) = H(w(x - b))$ divides the space of each variable x in two half-lines starting from b , one of which has a value of 1 and one of which has a value of 0, depending on the sign of w .

Aggregation Using sum as the aggregation function, the output takes values in $\{0, 1, \dots, n\}$; where 0 corresponds to a deactivation for each input, and n represents an activation for all inputs, and the intermediate integer values $\{1, 2, \dots, k, \dots, n - 1\}$ represent activation for k of inputs.

$$y = \sum_{i=1}^n h_i = \begin{cases} n & h_i = 1 \forall i \in \{1, \dots, n\} \\ k & h_i = 1 i \in S \subseteq \{1, \dots, n\}, |S| = k \\ 0 & h_i = 0 \forall i \in \{1, \dots, n\} \end{cases} \quad (5)$$

where we simplified the notation using $h_i = h(x_i)$.

2+ Layer - Processing function Let us define an M -of- N rule as true if at least M of the N rules of a given set are true.

The Heavisides of the layers after the first one receive values in $\{0, 1, \dots, n\}$, where n represents the number of inputs of the previous layer. In the case where $0 \leq b \leq n$ and $w > 0$, the Heaviside will output 1 only if the input received is greater than or equal to b , therefore only if at least $\lceil b \rceil$ of the rules R_i of the previous layer are true, which corresponds to a rule of the type $\lceil b \rceil - of - \{R_1, R_2, \dots, R_n\}$. In the opposite case, where $0 \leq b \leq n$ and $w < 0$, Heaviside will output 1 only if the input received is less than or equal to b , so only if no more than $\lfloor b \rfloor$ of the rules of the previous layer are true. This too can be translated to an M -of- N rule, inverting all rules R_j and setting M as $\lceil n - b_i \rceil$: $\lceil n - b_i \rceil - of - \{\neg R_1, \neg R_2, \dots, \neg R_n\}$.

Last layer - Aggregation In the last layer we have to account for the α factors used to weigh the contribution of each input:

$$y = \sum_{i=1}^n \alpha_i h_i(x_i) = \sum_{i=1}^n \alpha_i H(w_i(x_i - b_i)) \quad (6)$$

We have an activation rule for each of the n Heavisides forcing us to calculate all the 2^n possible cases. The contribution of each input is exactly α_i . So, the output corresponds to the sum of the α_i 's for each subset of inputs considered.

Sigmoid

In the case of sigmoid IAN, b_i represents the inflection point of the function, while the sign of w_i tells us in which direction the sigmoid is oriented; if positive, it is monotonically increasing from 0 to 1, while if negative, it is monotonically

decreasing from 1 to 0. The value of w_i indicates how fast it transitions from 0 to 1, and if it tends to infinity, the sigmoid tends to the unit step function.

Sigmoid Interpretation The sigmoid can be interpreted as a fuzzy rule of the type $x_i > b_i$ if $w_i > 0$ or $x_i < b_i$ if $w_i < 0$, where the absolute value of w_i indicates how sharp the rule is. The case $w_i = 0$ will always give value 0.5, so that the input does not have any influence on the output.

If w_i is very large, the sigmoid tends to the unit step function. If, on the other hand, w_i takes values for which the sigmoid in the domain of x_i resembles a linear function, what we can say is that there is a direct linear relationship (or inverse if $w_i < 0$) with the input.

The fuzzy rule can be approximated by its stricter version $x_i > b_i$, interpreting fall under the methodology seen for Heaviside. However, this would result in an approximation of the operation of the network.

It is more challenging to devise clear decision rules when we add more layers. Imagine, as an example, a second layer with this processing function:

$$h(y) = \sigma(w^*(y - b^*)) \quad (7)$$

where y is the aggregation performed in the previous layer of the outputs of its processing functions, its value roughly indicates how many of the inputs are active. In the second layer, consider as an example a value of $w^* > 0$. To have an activation, this means that we might need k inputs greater than or equal to b^*/k . Although this does not deterministically indicate how many inputs we need to be true, we know how the output changes when one of the inputs changes.

The last case to consider takes into account the maximum and minimum values that the sigmoid assumes in the domain of x . If they are close to each other, that happens when w is very small, the function is close to a constant bearing no connection with the input.

Product of tanh

The multiplication of tanh has more expressive power, being able to represent both what is represented with the sigmoid, as well as intervals and quadratic relations.

tanh-prod Interpretation In this case, it is not possible to devise as quickly as in the previous case decision rules. Indeed, it is still possible to observe the trend of the function and draw some conclusions. When the product of the two tanh resembles a sigmoid, we can follow the interpretation of the sigmoid case. In other cases, areas with quadratic relations can occur, i.e., bells whose peak indicates a more robust activation or deactivation for specific values.

Summary of Interpretation

The advantage of this method lies in the fact that it is possible to analyze each input separately in each neuron, thus easily graph each processing function. Then, based on the shape taken by the processing function, we can understand how the input affects the output of a neuron.

The Heaviside is the most interpretable of our models, allowing a direct generation of decision rules.

Sigmoid and tanh-prod cases depend on the parameter w . When it is close to 0, the activation is constant regardless of the input. When w is large enough, the processing function is approximately a piecewise constant function taking only values 0 and 1.

In all the other cases, the processing function approximates a linear or bell-shaped function. Even if we can not derive exact decision rules directly from the model, in these cases, we can infer a linear or quadratic relation between input and output.

Each layer aggregates the interpretations of the previous layers. For example, the processing function of a second layer neuron gives a precise activation when its input is greater than a certain threshold, i.e., the bias b of the processing function. The output of the neuron of the first layer must exceed this threshold, and this happens if its processing functions give in output values whose sum exceeds this threshold.

A separate case is the last layer, where the α parameters weigh each of the interpretations generated up to the last layer.

We can interpret a traditional individual neuron as a linear regressor. However, when we add more layers, they cannot be interpreted. Our structure, instead, remains interpretable even as the number of layers increases.

5 Universality

A fundamental property of neural networks is that of universal approximation. Under certain conditions, multilayer feed-forward neural networks can approximate any function in a given function space. In (Cybenko 1989) it is proved that a neural network with a hidden layer and using a continuous sigmoidal activation function is dense in $C(I_n)$, i.e., the space of continuous functions in the unit hypercube in \mathbb{R}^n . (Hornik, Stinchcombe, and White 1989) generalized to the larger class of all sigmoidal functions.

To make the statement of theorems clearer we recall that the structure of a two-layer network with IAN neurons and a generic processing function h is

$$\psi(x) = \sum_{j=1}^N \alpha_j h(w_j (\sum_{i=1}^n h(w_{ij}(x_i - b_{ij})) - b_j)) \quad (8)$$

where $w_j, w_{ij}, \alpha_j, b_j, b_{ij} \in \mathbb{R}$.

When the processing function is the Heaviside function we proved that the network can approximate any continuous function on I_n , Lebesgue measurable functions on I_n and functions in $L^p(\mathbb{R}^n, \mu)$ for $1 \leq p < \infty$, with μ being a Radon measure. More precisely, the following theorems hold; we detail the proofs of the theorems in the appendix.

Theorem 5.1. *When the processing function is the Heaviside function the finite sums of the form (8) are dense in $L^p(\mathbb{R}^n, \mu)$ for $1 \leq p < \infty$, with μ being a Radon measure on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ (\mathcal{B} denote the Borel σ -algebra).*

Theorem 5.2. *When the processing function is the Heaviside function the finite sum of the form (8) are m -dense in M^n . Where M^n is the set of Lebesgue measurable functions on the n -dimensional hypercube I_n .*

Theorem 5.3. Given $g \in C(I_n)$ and given $\epsilon > 0$ there is a sum $\psi(x)$ of the form (8) with Heaviside as processing function such that

$$|\psi(x) - g(x)| < \epsilon \quad \forall x \in I_n.$$

When the processing function is the sigmoid function or tanh-prod, we proved that the finite sums of the form (8) are dense in the space of continuous functions defined on the unit n -dimensional hypercube.

Theorem 5.4. When the processing function is a continuous sigmoidal function the finite sums of the form (8) are dense in $C(I_n)$.

Theorem 5.5. Let $\psi(x)$ be the family of networks defined by the equation (8) when the processing function is given by (4). This family of functions is dense in $C(I_n)$.

6 Experiments

Datasets

We selected a collection of datasets from the UCI Machine Learning Repository. We only consider classification models in our experiments. However, it is straightforward to apply NEWRON architectures to regression problems. The description of the datasets is available at the UCI Machine Learning Repository website or the Kaggle website.

We also used 4 synthetic datasets of our creation, composed of 1000 samples with 2 variables generated as random uniforms between -1 and 1 and an equation dividing the space into 2 classes. The 4 equations used are bisector, xor, parabola, and circle.

We give more details about the datasets in the appendix.

Methods

We run a hyperparameter search to optimize the IAN neural network structure, i.e., depth and number of neurons per layer, for each dataset. We tested IAN with all three different processing functions. In the tanh-prod case, we set $M = 2$.

Concerning the training of traditional neural networks, we tested the same structures used for NEWRON, i.e., the same number of layers and neurons. Finally, we also ran a hyperparameter search to find the best combinations in the case of Logistic Regression (LR), Decision Trees (DT), and Gradient Boosting Decision Trees (GBDT). We include all the technical details on the methods in the appendix.

Results

Table 1 presents on each row the datasets used while on the columns the various models. Each cell contains the 95% confidence interval for the accuracy of the model that obtains the best performance.

Results obtained with the new IAN neurons are better than those obtained by DTs and LR (interpretable) models. Moreover, IAN’s results are on par, sometimes better than, results of traditional NNs and GBDT classifiers. These last two methods, though, are not transparent.

Amongst the Heaviside, sigmoid, and tanh-prod cases, we can see that the first one obtains the worst results. The

reason may be that it is more challenging to train, despite being the most interpretable among the three cases. tanh-prod instead performs slightly better than sigmoid, being more flexible. Sigmoid, being more straightforward to interpret than tanh-prod, could be a good choice at the expense of a slight decrease in accuracy that remains, however, similar to that of a traditional neural network.

Circle dataset example

In order to first validate our ideas, we show what we obtained by applying a single neuron using multiplication of 2 tanh in the case of our custom dataset circle.

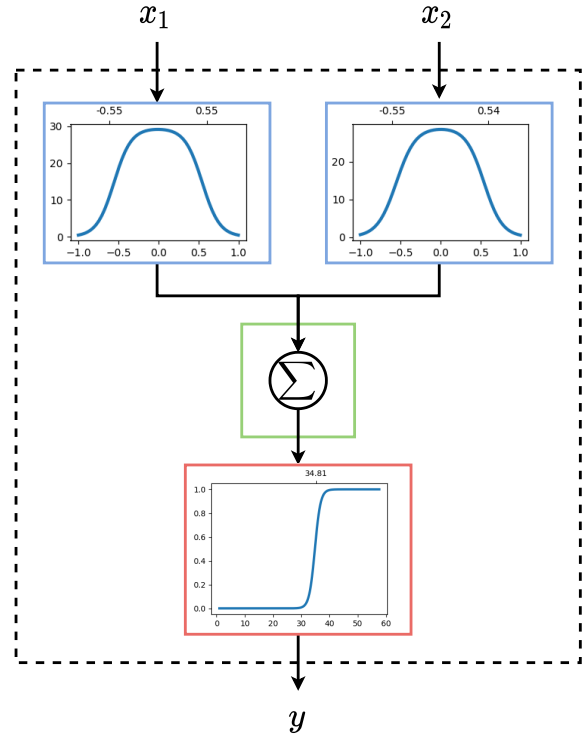


Figure 4: tanh-prod Neural Network trained on the circle dataset. The figure follows the color convention used for NEWRON in Figure 3. x_1 and x_2 are the inputs of the network and y is the output. The processing and activation functions are plotted with input on the x -axis and output on the y -axis. Coordinates of the inflection points are indicated above the plots.

In Figure 4 we can see how the multiplication of tanh has converged to two bells centred in 0, while α_1 and α_2 have gone to 30. According to the IANinterpretation method, values below 30 correspond to an activation function output of 0, while it is 1 for values above 38. In the middle range, the prediction is more uncertain. Combining this data with the previous prediction, we can conclude that we need the sum of the two values output by the two processing functions to be greater than 38 to have a prediction of class 1. Therefore, if one of the two inputs is 0 (output 30), it is enough for

Dataset	IAN models			Interpretable models		Non-interpretable models	
	Heaviside	sigmoid	tanh-prod	LR	DT	GBDT	NN
adult	80.2 (± 0.06)	82.6 (± 0.05)	82.3 (± 0.06)	76.2 (± 0.07)	81.5 (± 0.06)	<u>87.5 (± 0.05)</u>	83.1 (± 0.06)
australian	86.5 (± 0.51)	87.0 (± 0.5)	88.7 (± 0.4)	88.7 (± 0.4)	87.0 (± 0.41)	<u>90.2 (± 0.47)</u>	88.0 (± 0.4)
b-c-w	98.9 (± 0.16)	98.9 (± 0.16)	98.9 (± 0.16)	97.8 (± 0.23)	97.7 (± 0.23)	<u>98.3 (± 0.21)</u>	98.9 (± 0.17)
car	95.1 (± 0.2)	95.9 (± 0.21)	100.0 (± 0.0)	51.4 (± 0.45)	98.5 (± 0.11)	<u>100.0 (± 0.0)</u>	<u>99.8 (± 0.04)</u>
cleveland	65.6 (± 1.02)	60.1 (± 1.1)	62.9 (± 1.13)	60.8 (± 1.13)	53.6 (± 1.19)	<u>61.5 (± 1.01)</u>	65.6 (± 1.01)
crx	86.2 (± 0.51)	85.4 (± 0.58)	86.5 (± 0.5)	84.6 (± 0.45)	88.0 (± 0.42)	82.9 (± 0.58)	<u>87.7 (± 0.44)</u>
diabetes	73.3 (± 0.56)	72.7 (± 0.68)	76.1 (± 0.61)	75.6 (± 0.6)	74.1 (± 0.63)	75.1 (± 0.64)	74.2 (± 0.65)
german	78.2 (± 0.53)	77.0 (± 0.53)	75.5 (± 0.52)	75.1 (± 0.52)	68.3 (± 0.57)	76.6 (± 0.55)	76.7 (± 0.54)
glass	77.0 (± 1.17)	81.6 (± 1.04)	85.6 (± 1.02)	72.1 (± 1.08)	72.7 (± 1.19)	<u>87.3 (± 0.9)</u>	82.5 (± 0.91)
haberman	76.9 (± 0.94)	76.1 (± 0.92)	77.2 (± 0.88)	73.0 (± 1.05)	64.4 (± 1.08)	<u>72.5 (± 1.09)</u>	76.1 (± 0.92)
heart	88.7 (± 0.67)	86.3 (± 0.85)	82.7 (± 0.8)	82.4 (± 0.95)	81.4 (± 1.02)	81.7 (± 0.98)	82.9 (± 0.95)
hepatitis	84.7 (± 1.26)	85.1 (± 1.23)	82.5 (± 1.16)	79.1 (± 1.45)	79.1 (± 1.33)	81.7 (± 1.32)	82.4 (± 1.13)
image	93.0 (± 0.11)	94.0 (± 0.1)	94.4 (± 0.09)	90.4 (± 0.12)	90.6 (± 0.12)	<u>95.8 (± 0.08)</u>	92.6 (± 0.11)
ionsosphere	94.4 (± 0.48)	96.7 (± 0.34)	96.5 (± 0.37)	92.0 (± 0.51)	94.5 (± 0.45)	<u>95.4 (± 0.37)</u>	<u>96.7 (± 0.34)</u>
iris	100.0 (± 0.0)	100.0 (± 0.0)	100.0 (± 0.0)	100.0 (± 0.0)	97.3 (± 0.52)	97.3 (± 0.52)	<u>100.0 (± 0.0)</u>
monks-1	94.4 (± 0.21)	100.0 (± 0.0)	100.0 (± 0.0)	66.0 (± 0.46)	90.6 (± 0.27)	<u>100.0 (± 0.0)</u>	<u>100.0 (± 0.0)</u>
monks-2	100.0 (± 0.0)	100.0 (± 0.0)	100.0 (± 0.0)	54.5 (± 0.45)	82.7 (± 0.33)	<u>94.2 (± 0.21)</u>	<u>87.6 (± 0.27)</u>
monks-3	97.1 (± 0.15)	97.1 (± 0.15)	97.1 (± 0.15)	81.2 (± 0.31)	97.2 (± 0.16)	96.2 (± 0.16)	90.3 (± 0.25)
sonar	93.3 (± 0.74)	96.8 (± 0.48)	95.2 (± 0.53)	89.5 (± 0.75)	83.4 (± 0.98)	88.1 (± 0.9)	89.4 (± 0.87)
bisector	98.9 (± 0.13)	99.3 (± 0.09)	99.3 (± 0.09)	100.0 (± 0.0)	97.7 (± 0.18)	98.3 (± 0.16)	<u>100.0 (± 0.0)</u>
xor	100.0 (± 0.0)	100.0 (± 0.0)	99.2 (± 0.11)	53.2 (± 0.65)	99.2 (± 0.12)	<u>100.0 (± 0.0)</u>	<u>100.0 (± 0.0)</u>
parabola	98.8 (± 0.15)	100.0 (± 0.0)	99.6 (± 0.07)	77.8 (± 0.52)	97.6 (± 0.18)	<u>97.7 (± 0.17)</u>	<u>100.0 (± 0.0)</u>
circle	96.8 (± 0.22)	99.3 (± 0.1)	99.6 (± 0.07)	52.4 (± 0.67)	98.8 (± 0.13)	<u>97.6 (± 0.2)</u>	<u>99.2 (± 0.11)</u>

Table 1: Datasets accuracy ($\pm 95^{th}$ percentile standard error) results of the best performing model. In **bold** we indicate the best performing model amongst the interpretable ones. If GBDT or NN exceeds this accuracy, the corresponding result is underlined.

the other to be between -0.65 and 0.65 (output greater than 8). Otherwise, we may need an output of at least 19 from both outputs, corresponding to input values between -0.5 and 0.5 , i.e., the area covered by the circle. We show more examples in the appendix.

Current limitations

The extraction of proper rules from the network can be hard; in the Heaviside case, they might be too long in the sigmoid and tanh-prod cases because their simplicity depends on the final value parameters. Nevertheless, methods of regularization during training or additional Rule Extraction methods may help to simplify interpretability. We defer the study of regularization to future works.

Also, we have not compared NEWRON against state-of-the-art Deep Learning models for tabular data, as our main goal was to show that our formulation was more suitable than traditional neurons compared to “traditional” interpretable models. Comparisons with more advanced solutions for tabular data will be the subject of future work.

7 Conclusions and Future Work

We have introduced the concept of a generalized neuron and proposed three different specializations, along with the corresponding method to interpret the behavior of the network. Also, in cases where from the network we cannot devise exact rules (e.g., in the sigmoid and tanh-prod cases), the structure of the neuron and the parameters allow the visual-

ization of its behavior. Indeed, for every input, we apply the nonlinearity operation before the aggregation reducing it to a one-dimensional space allowing the analysis of each input separately. Through universal approximation theorems, we have proved that the new structure retains the same expressive power as a standard neural network. In future studies we will investigate more in detail the expressiveness of IAN based models with respect to the number of layers or neurons in arbitrarily deep but width-limited networks and arbitrarily wide but depth-limited networks. Experiments conducted on both real and synthetic datasets illustrate how our framework can outperform traditional interpretable models, Decision Trees, and Logistic Regression, and achieve similar or superior performance to standard neural networks. In the future, we will investigate the influence of hyperparameters (network depth, number of neurons, processing functions) and initialization on the model quality. Also, we will refine the analysis of the tanh-prod case as the number of tanh increases. In addition, we will investigate IAN with additional processing functions, such as ReLU and SeLU. Finally, we will extend this method to other neural models, such as Recurrent, Convolutional and Graph Neural Networks.

8 Acknowledgements

This research was supported by the Italian Ministry of Education, University and Research (MIUR) under the grant “Dipartimenti di eccellenza 2018–2022” of the Department

of Computer Science and the Department of Computer Engineering at Sapienza University of Rome. Partially supported by the ERC Advanced Grant 788893 AMDROMA “Algorithmic and Mechanism Design Research in Online Markets”, the EC H2020RIA project “SoBigData++” (871042), and the MIUR PRIN project ALGADIMAR “Algorithms, Games, and Digital Markets”. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

- Apicella, A.; Donnarumma, F.; Isgrò, F.; and Prevete, R. 2021. A survey on modern trainable activation functions. *Neural Networks*.
- Arrieta, A. B.; Díaz-Rodríguez, N.; Del Ser, J.; Bennetot, A.; Tabik, S.; Barbado, A.; García, S.; Gil-López, S.; Molina, D.; Benjamins, R.; et al. 2020. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58: 82–115.
- Bhatt, U.; Ravikumar, P.; et al. 2019. Building human-machine trust via interpretability. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 9919–9920.
- Cannarsa, P.; and D’Aprile, T. 2015. *Introduction to Measure Theory and Functional Analysis*. UNITEXT. Springer International Publishing. ISBN 9783319170183.
- Cybenko, G. 1989. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(2): 303–314.
- Dalvi, F.; Durrani, N.; Sajjad, H.; Belinkov, Y.; Bau, A.; and Glass, J. 2019a. What is one grain of sand in the desert? analyzing individual neurons in deep nlp models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 6309–6317.
- Dalvi, F.; Nortonsmith, A.; Bau, A.; Belinkov, Y.; Sajjad, H.; Durrani, N.; and Glass, J. 2019b. NeuroX: A toolkit for analyzing individual neurons in neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 9851–9852.
- Fan, F.; Cong, W.; and Wang, G. 2018. A new type of neurons for machine learning. *International journal for numerical methods in biomedical engineering*, 34(2): e2920.
- Ghorbani, A.; Abid, A.; and Zou, J. 2019. Interpretation of neural networks is fragile. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 3681–3688.
- Gupta, M. M.; Bukovsky, I.; Homma, N.; Solo, A. M.; and Hou, Z.-G. 2013. Fundamentals of higher order neural networks for modeling and simulation. In *Artificial Higher Order Neural Networks for Modeling and Simulation*, 103–133. IGI Global.
- Heo, B.; Lee, M.; Yun, S.; and Choi, J. Y. 2019. Knowledge transfer via distillation of activation boundaries formed by hidden neurons. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 3779–3787.
- Hornik, K.; Stinchcombe, M.; and White, H. 1989. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5): 359–366.
- Kulkarni, R. V.; and Venayagamoorthy, G. K. 2009. Generalized neuron: Feedforward and recurrent architectures. *Neural networks*, 22(7): 1011–1017.
- Nam, W.-J.; Gur, S.; Choi, J.; Wolf, L.; and Lee, S.-W. 2020. Relative attributing propagation: Interpreting the comparative contributions of individual units in deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 2501–2508.
- Rosenblatt, F. 1958. The Perceptron: A Probabilistic Model for Information Storage and Organization in The Brain. *Psychological Review*, 65–386.
- Rudin, W. 1976. *Principles of mathematical analysis*. International series in pure and applied mathematics. New York: McGraw-Hill.
- Rudin, W. 1987. *Real and Complex Analysis*. New York: McGraw-Hill.
- Stein, E. M.; and Shakarchi, R. 2005. *Real analysis: measure theory, integration, and Hilbert spaces*. Princeton lectures in analysis. Princeton, NJ: Princeton Univ. Press.
- Sun, S.; Cao, Z.; Zhu, H.; and Zhao, J. 2019. A survey of optimization methods from a machine learning perspective. *IEEE transactions on cybernetics*, 50(8): 3668–3681.
- Wheeden, L., Richard; and Zygmund, A. 2015. *Measure and integral: An introduction to real analysis*. CRC Press. Second edition.
- Yang, Y.; Morillo, I. G.; and Hospedales, T. M. 2018. Deep neural decision trees. *arXiv preprint arXiv:1806.06988*.

Supplementary Materials

A Universality Theorems

This is the appendix to the Universality section in the main article. In this section, we shall prove the mathematical results concerning the universal approximation properties of our IAN model. In particular, we restrict ourselves to some specific cases. We consider the cases where the processing function is the Heaviside function, a continuous sigmoidal function, or the rescaled product of hyperbolic tangents.

Heaviside IAN

Theorem 5.1. *The finite sums of the form*

$$\psi(x) = \sum_{j=1}^N \alpha_j H(w_j \sum_{i=1}^n H(w_{ij}(x_i - b_{ij})) - b_j) \quad (9)$$

with $N \in \mathbb{N}$ and $w_{ij}, w_j, \alpha_j, b_{ij}, b_j \in \mathbb{R}$ are dense in $L^p(\mathbb{R}^n, \mu)$ for $1 \leq p < \infty$, with μ a Radon measure on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ (\mathcal{B} denote the Borel σ -algebra).

In other words given, $g \in L^p(\mathbb{R}^n, \mu)$ and $\epsilon > 0$ there is a sum $\psi(x)$ of the above form for which

$$\|\psi - g\|_p^p = \int_{\mathbb{R}^n} |\psi(x) - g(x)|^p d\mu(x) < \epsilon.$$

To prove that a neural network defined as in equation (9) is a universal approximator in L^p , for $1 \leq p < \infty$ we exploit that step functions are dense in L^p and that our network can generate step functions.

Proposition 1. Let \mathcal{R} be the set of the rectangles in \mathbb{R}^n of the form

$$R = \prod_{k=1}^n [a_k, b_k) \quad a_k, b_k \in \mathbb{R}, \quad a_k < b_k$$

We denote by \mathcal{F} the vector space on \mathbb{R} generated by $\mathbb{1}_R, R \in \mathcal{R}$ i.e.

$$\mathcal{F} = \left\{ \sum_{i=1}^m \alpha_i \mathbb{1}_{R_i} \mid m \in \mathbb{N}, \alpha_i \in \mathbb{R}, R_i \in \mathcal{R} \right\} \quad (10)$$

\mathcal{F} is dense in $L^p(\mathbb{R}^n, \mu)$ for $1 \leq p < \infty$, with μ a Radon measure on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$.

Proof. See chapter 3, L^p Spaces, in (Cannarsa and D'Aprile 2015). \square

Lemma 2. Given $\rho(x) \in \mathcal{F}$, with \mathcal{F} defined as in equation (10), there exists a finite sum $\psi(x)$ of the form (9) such that $\rho(x) = \psi(x) \forall x \in \mathbb{R}^n$.

Proof. To prove that a neural network described as in equation (9) can generate step functions we proceed in two steps. First, we show how we can obtain the indicator functions of orthants from the first layer of the network. Then we show how, starting from these, we can obtain the step functions.

An *orthant* is the analogue in n -dimensional Euclidean space of a quadrant in \mathbb{R}^2 or an octant in \mathbb{R}^3 . We denote by *translated orthant* an orthant with origin in a point different

from the origin of the Euclidean space O . Let A be a point in the n -dimensional Euclidean space, and let us consider the intersection of n mutually orthogonal half-spaces intersecting in A . By independent selections of half-space signs with respect to A (i.e. to the right or left of A) 2^n orthants are formed.

Now we shall see how to obtain translated orthant with origin in a point A of coordinates (a_1, a_2, \dots, a_n) from the first layer of the network i.e. $\sum_{i=1}^n H(w_i(x_i - b_i))$.

For this purpose we can take $w_i = 1 \quad \forall i \in \{1, \dots, n\}$.

The output of $\sum_{i=1}^n H(x_i - b_i) \in \{0, \dots, n\}$ and depends on how many of the n Heaviside functions are activated. We obtain the translated orthant with origin in A by choosing $b_i = a_i \quad \forall i \in \{1, \dots, n\}$. In fact,

$$H(x_i - a_i) = \begin{cases} 0 & \text{if } x_i < a_i \\ 1 & \text{if } x_i \geq a_i. \end{cases}$$

The i -th Heaviside is active in the half-space $x_i \geq a_i$ delimited by the hyperplane $x_i = a_i$ that is orthogonal to the i -th axis. Therefore, the Euclidian space \mathbb{R}^n is divided in 2^n regions according to which value the function $\sum_{i=1}^n H(x_i - a_i)$ takes in each region. See Figure 5 for an example in \mathbb{R}^2 .




Figure 5: Partition of \mathbb{R}^2 , according to output of the function $H(x_1 - a_1) + H(x_2 - a_2)$. A is a point of coordinates (a_1, a_2) .

There is only one region in which the output of the sum is n , which corresponds to the orthant in which the condition $x_i \geq a_i \forall i = 1, \dots, n$ holds. We denote it as *positive orthant* (the red colored orthant in the example shown in Figure 5).

Going back to equation (9), let us now consider the Heaviside function applied after the sum. As before, we can choose $w_j = 1$. If we take $b_j > n - 1$, the value of the output is 0 for each of the 2^n orthants except for the positive orthant. This way, we get the indicator function of the positive orthant.

The indicator function of a rectangle in \mathcal{R} can be obtained as a linear combination of the indicator function of the positive orthants centered in the vertices of the rectangle. See Figure 6 for an example of the procedure in \mathbb{R}^2 .

In general, the procedure involves considering a linear combination of indicator functions of positive orthants centered in the vertices of the rectangle in such a way that op-

posite values are assigned to the orthants corresponding to adjacent vertices.

For example, suppose we want to obtain the indicator function of the right-closed left-open square $[0, 1]^2$ in \mathbb{R}^2 (see the illustration in Figure 6). Denoting by $\mathbb{1}_{(x_P, y_P)_\leftarrow}$ the indicator function of the positive orthant centered in the point of coordinates (x_P, y_P) , we can write:

$$\mathbb{1}_{[0,1]^2} = \mathbb{1}_{(0,0)_\leftarrow} - \mathbb{1}_{(1,0)_\leftarrow} - \mathbb{1}_{(0,1)_\leftarrow} + \mathbb{1}_{(1,1)_\leftarrow}.$$




Figure 6: How to obtain the indicator function on the square $[0, 1]^2$ from the linear combination of four indicator functions of positive orthants centered in the vertices of $[0, 1]^2$. $\mathbb{1}_{[0,1]^2} = \mathbb{1}_{(0,0)_\leftarrow} - \mathbb{1}_{(1,0)_\leftarrow} - \mathbb{1}_{(0,1)_\leftarrow} + \mathbb{1}_{(1,1)_\leftarrow}$. The numbers in the orthants shows the sum of the indicator functions that are active in that orthant. For instance if $x = (x_1, x_2)$ belongs to the blue part of the plane, i.e. it is true that $0 < x_1 < 1$ and $x_2 > 1$, we have that $\mathbb{1}_{(0,0)_\leftarrow}(x) - \mathbb{1}_{(1,0)_\leftarrow}(x) - \mathbb{1}_{(0,1)_\leftarrow}(x) + \mathbb{1}_{(1,1)_\leftarrow}(x) = 1 - 0 - 1 + 0 = 1 - 1$.

Now suppose we want the linear combination of the indicator functions of K rectangles with coefficients $\alpha_1, \dots, \alpha_K$. With suitably chosen coefficients the indicator function of a rectangle can be written as

$$\sum_{l=1}^{2^n} (-1)^l H(w_{jl} \sum_{i=1}^n H(w_{ij}(x_i - b_{ij})) - b_{jl})$$

that replacing $H(w_{jl} \sum_{i=1}^n H(w_{ij}(x_i - b_{ij})) - b_{jl})$ by H_l , to abbreviate the notation becomes

$$\sum_{l=1}^{2^n} (-1)^l H_l.$$

The linear combination of the indicator functions of K rectangles with coefficients $\alpha_1, \dots, \alpha_K$ can be derived as

$$\sum_{k=1}^K \alpha_k \sum_{l=1}^{2^n} (-1)^l H_{lk}. \quad (11)$$

The summation (11) can be written as a single sum, defining a sequence $\beta_j = (-1)^j \alpha_m$ with $m = \lceil \frac{j}{2^n} \rceil$ for $j = 1, \dots, 2^n K$. Thus (11) becomes

$$\sum_{j=1}^{N=2^n K} \beta_j H_j$$

that is an equation of form (9). We have therefore shown that for every step function ρ in \mathcal{F} there are $N \in \mathbb{N}$ and $\alpha_j, w_{ij}, b_{ij}, b_j, w_j \in \mathbb{R}$ such that the sum in equation (9) is equal to ρ . \square

Proof of Theorem 5.1. The theorem follows immediately from Lemma 2 and Proposition 1. \square

Remark 1. In Lemma 2 we proved that a network defined as in equation (9) can represent functions belonging to set \mathcal{F} defined as in equation (10). Note that if the input is bounded, we can obtain indicator functions of other kinds of sets. For example, suppose $x \in [0, 1]^n$. If we choose $w_{ij} = 1$ and $b_{ij} < 0 \forall i, j$ and if we choose the weights of the second layer so that they don't operate any transformation, we can obtain the indicator function of $[0, 1]^n$. By a suitable choice of parameters, (9) may also become the indicator functions of any hyperplane $x_i = 0$ or $x_i = 1$ for $i \in \{1, \dots, n\}$. Furthermore we can obtain any rectangle of dimension $n-1$ that belongs to an hyperplane of the form $x_i = 1$ or $x_i = 0$.

We have proven in Lemma 2 that a network formulated as in equation (9) can represent step functions. By this property and by Proposition 3 we shall show that it can approximate Lebesgue measurable functions on any finite space, for example the unit n -dimensional cube $[0, 1]^n$.

We denote by I_n the closed n -dimensional cube $[0, 1]^n$. We denote by M^n the set of measurable functions with respect to Lebesgue measure m , on I_n , with the metric d_m defined as follows. Let be $f, g \in M^n$,

$$d_m(f, g) = \inf\{\epsilon > 0 : m\{x : |f(x) - g(x)| > \epsilon\} < \epsilon\}$$

We remark that d_m -convergence is equivalent to convergence in measure (see Lemma 2.1 in (Hornik, Stinchcombe, and White 1989)).

Theorem 5.2. The finite sums of the form (9) with $N \in \mathbb{N}$ and $w_{ij}, w_j, \alpha_j, b_{ij}, b_j \in \mathbb{R}$ are d_m -dense in M^n . M^n is the set of Lebesgue measurable functions on I_n .

This means that, given g measurable with respect to the Lebesgue measure m on I_n , and given an $\epsilon > 0$, there is a sum ψ of the form (9) such that $d_m(\psi, g) < \epsilon$.

Proposition 3. Suppose f is measurable on \mathbb{R}^n . Then there exists a sequence of step functions $\{\rho_k\}_{k=1}^\infty$ that converges pointwise to $f(x)$ for almost every x .

Proof. See Theorem 4.3 p. 32 in (Stein and Shakarchi 2005). \square

Proof of Theorem 5.2. Given any measurable function, by Proposition 3 there exists a sequence of step functions that converge to it pointwise. By Lemma 2 we have that equation (9) can generate step functions. Now $m(I_n) = 1$ and for a finite measure space pointwise convergence implies convergence in measure, this concludes the prof. \square

Remark 2. Notice that for Theorem 5.2 we don't need the fact that I_n is a closed set. For this theorem in fact it is sufficient that it is a bounded set (so that its Lebesgue measure is finite). The compactness of I_n will be necessary for the next theorem.

Theorem 5.3. Given $g \in C(I_n)$ and given $\epsilon > 0$ there is a sum $\psi(x)$ of the form (9) such that

$$|\psi(x) - g(x)| < \epsilon \quad \forall x \in I_n.$$

Proof. Let g be a continuous function from I_n to \mathbb{R} , by the compactness of I_n follows that g is also uniformly continuous (see Theorem 4.19 p. 91 in (Rudin 1976)). In other words, for any $\epsilon > 0$, there exists $\delta > 0$ such that for every $x, x' \in [0, 1]^n$ such that $\|x - x'\|_\infty < \delta$ it is true that $|g(x) - g(x')| < \epsilon$. To prove the statement of Theorem 5.3, let $\epsilon > 0$ be given, and let $\delta > 0$ be chosen according to the definition of uniform continuity.

As we have already seen in Lemma 2 the neural network described in (9) can generate step functions with support on right-open left-closed n -dimensional rectangles and on $(n-1)$ -dimensional rectangles that belongs to an hyperplane of equation $x_i = 0$ or $x_i = 1$ for some $i \in \{1, \dots, n\}$ as seen in Remark 1. There exists a partition of $[0, 1]^n$, (R_1, \dots, R_N) , consisting of right-open left-closed n -dimensional rectangles and of $(n-1)$ -dimensional rectangles that belongs to an hyperplane of equation $x_i = 0$ or $x_i = 1$ for some $i \in \{1, \dots, n\}$, such that all side lengths are no greater than δ . An example of a set of rectangles with this property is the set of right-open left-closed cubes of side length $\frac{1}{\tilde{m}}$, $\tilde{m} > \lceil \frac{1}{\delta} \rceil$ with the $(n-1)$ -dimensional rectangles with the same side length which we need to cover all the boundary of $[0, 1]^n$ not covered by the right-open left-closed rectangles.

Suppose that for all $j \in \{1, \dots, N\}$ we choose $x_j \in R_j$, and we set $\alpha_j = g(x_j)$. If $x \in [0, 1]^n$ there is j so that $x \in R_j$, hence x satisfies $\|x - x_j\|_\infty \leq \delta$, and consequently $|g(x) - g(x_j)| \leq \epsilon$. Therefore the step function $h = \sum_{j=1}^N \alpha_j \mathbf{1}_{R_j}$ satisfies

$$\begin{aligned} & \sup_{x \in I_n} |h(x) - g(x)| = \\ &= \sup_{j \in \{1, \dots, N\}} \sup_{x \in R_j} |h(x) - g(x)| = \\ &= \sup_{j \in \{1, \dots, N\}} \sup_{x \in R_j} |\alpha_j - g(x)| \leq \epsilon \end{aligned}$$

□

Sigmoid IAN

Definition 4. A function $\sigma : \mathbb{R} \rightarrow [0, 1]$ is called sigmoidal if

$$\lim_{x \rightarrow -\infty} \sigma(x) = 0, \quad \lim_{x \rightarrow +\infty} \sigma(x) = 1$$

Theorem 5.4. Let σ be a continuous sigmoidal function. Then the finite sums of the form:

$$\psi(x) = \sum_{j=1}^N \alpha_j \sigma(w_j (\sum_{i=1}^n \sigma(w_{ij}(x_i - b_{ij})) - b_j)) \quad (12)$$

with $w_{ij}, \alpha_j, b_{ij}, b_j, w_j \in \mathbb{R}$ and $N \in \mathbb{N}$ are dense in $C(I_n)$.

In other words, given a $g \in C(I_n)$ and given $\epsilon > 0$ there is a sum $\psi(x)$ of the form (12) such that

$$|\psi(x) - g(x)| < \epsilon \quad \forall x \in I_n.$$

Proof. Since σ is a continuous function, it follows that the set U of functions of the form (12) with $\alpha_j, w_{ij}, b_{ij}, w_j, b_j \in \mathbb{R}$ and $N \in \mathbb{N}$ is a linear subspace of $C(I_n)$. We claim that the closure of U is all of $C(I_n)$.

Assume that U is not dense in $C(I_n)$, let S be the closure of U , $S \neq C(I_n)$. By the Hahn-Banach theorem (see p. 104 of (Rudin 1987)) there is a bounded linear functional on $C(I_n)$, call it L , with the property that $L \neq 0$ but $L(S) = L(U) = 0$.

By the Riesz Representation Theorem (see p. 40 of (Rudin 1987)), the bounded linear functional L , is of the form

$$L(f) = \int_{I_n} f(x) d\mu$$

for some signed regular Borel measures μ such that $\mu(K) < \infty$ for every compact set $K \subset I_n$ (i.e. μ is a Radon measure). Hence,

$$\int_{I_n} h(x) d\mu = 0, \quad \forall h \in U. \quad (13)$$

We shall prove that (13) implies $\mu = 0$, which contradicts the hypothesis $L \neq 0$.

Using the definition of U , equation (13) can also be written as

$$\sum_{j=1}^N \alpha_j \int_{I_n} \sigma(w_j (\sum_{i=1}^n \sigma(w_{ij}(x_i - b_{ij})) - b_j)) d\mu = 0,$$

for any choice of $\alpha_j, w_{ij}, w_j, b_{ij}, b_j \in \mathbb{R}$ and $N \in \mathbb{N}$.

Note that for any $w, x, b \in \mathbb{R}$ we have that the continuous functions

$$\sigma_\lambda(w(x - b)) = \sigma(\lambda w(x - b) + \phi)$$

converge pointwise to the unit step function as λ goes to infinity, i.e.

$$\lim_{\lambda \rightarrow \infty} \sigma_\lambda(w(x - b)) = \gamma(w(x - b))$$

with

$$\gamma(y) = \begin{cases} 1 & \text{if } y > 0 \\ \sigma(\phi) & \text{if } y = 0 \\ 0 & \text{if } y < 0 \end{cases}$$

By hypothesis is true that for all λ_1, λ_2 in \mathbb{R}

$$\int_{I_n} \sigma_{\lambda_2}(w_j (\sum_{i=1}^n \sigma_{\lambda_1}(w_{ij}(x_i - b_{ij})) - b_j)) d\mu = 0.$$

It follows that for all λ_2 :

$$\lim_{\lambda_1 \rightarrow \infty} \int_{I_n} \sigma_{\lambda_2}(w_j (\sum_{i=1}^n \sigma_{\lambda_1}(w_{ij}(x_i - b_{ij})) - b_j)) d\mu = 0.$$

Now applying the Dominated Convergence Theorem (see Theorem 11.32 p 321 of (Rudin 1976)) and the fact that σ is continuous:

$$\int_{I_n} \lim_{\lambda_1 \rightarrow \infty} \sigma_{\lambda_2} \left(w_j \left(\sum_{i=1}^n \sigma_{\lambda_1} (w_{ij} (x_i - b_{ij})) - b_j \right) \right) d\mu = \int_{I_n} \sigma_{\lambda_2} \left(w_j \left(\sum_{i=1}^n \gamma (w_{ij} (x_i - b_{ij})) - b_j \right) \right) d\mu.$$

Again, by Dominated Convergence Theorem we have:

$$\lim_{\lambda_2 \rightarrow \infty} \int_{I_n} \sigma_{\lambda_2} \left(w_j \left(\sum_{i=1}^n \gamma (w_{ij} (x_i - b_{ij})) - b_j \right) \right) d\mu = \int_{I_n} \gamma \left(w_j \left(\sum_{i=1}^n \gamma (w_{ij} (x_i - b_{ij})) - b_j \right) \right) d\mu.$$

Hence we have obtained that $\forall \alpha_j, w_{ij}, b_{ij}, w_j, b_j \in \mathbb{R}$ and $\forall N \in \mathbb{N}$

$$\int_{I_n} \sum_{j=1}^N \alpha_j \gamma \left(w_j \left(\sum_{i=1}^n \gamma (w_{ij} (x_i - b_{ij})) - b_j \right) \right) d\mu = 0.$$

The function γ is very similar to the Heaviside function H , the only difference is that $H(0) = 1$ while $\gamma(0) = \sigma(\phi)$. Let R_i denote an open rectangle, $\partial_a R_i$ its left boundary (i.e. the boundary of a left-closed right-open rectangle) and $\partial_b R_i$ its right boundary (i.e. the boundary of a right-closed left-open rectangle). Repeating the construction seen in Lemma 2 to obtain rectangles, with the difference that here γ takes value $\sigma(\phi)$ on the boundaries, we get that

$$\sigma(\phi) \mu(\partial_a R_i) + (1 - \sigma(\phi)) \mu(\partial_b R_i) + \mu(R_i) = 0$$

for every open rectangle R_i . Taking $\phi \rightarrow \infty$, implies

$$\mu(\partial_a R_i) + \mu(R_i) = 0 \quad \forall \text{ open rectangle } R_i.$$

Every open subset A of I_n , can be written as a countable union of disjoint partly open cubes (see Theorem 1.11 p.8 of (Wheeden and Zygmund 2015)). Thus, from the fact that the measure is σ -additive we have that for every open subset A of I_n , $\mu(A) = 0$. Furthermore $\mu(I_n) = 0$. To obtain I_n from

$$\sum_{j=1}^N \alpha_j \gamma \left(w_j \left(\sum_{i=1}^n \gamma (w_{ij} (x_i - b_{ij})) - b_j \right) \right)$$

it is sufficient to choose the parameters so that $w_{ij} (x_i - b_{ij}) > 0 \forall x_i \in [0, 1]$ and so that w_j, b_j maintains the condition on the input.

Hence, $\mu(A^c) = \mu(I_n) - \mu(A) = 0$. It follows that for all compact set K of I_n , $\mu(K) = 0$.

From the regularity of the measure, it follows that μ is the null measure. \square

tanh-prod IAN

Theorem 5.5. *The finite sums of the form*

$$\psi(x) = \sum_{j=1}^N \frac{\alpha_j}{2} \left[\prod_{l=1}^{M_j} \tanh(w_{jl}(z_j(x) - b_{jl})) + 1 \right] \quad (14)$$

$$z_j(x) = \sum_{i=1}^n \frac{1}{2} \left[\prod_{k=1}^{m_i} \tanh(w_{ik}(x_i - b_{ik})) + 1 \right]$$

with $w_{jl}, w_{ijk}, \alpha_j, b_{jl}, b_{ijk} \in \mathbb{R}$ and $M_j, N, m_i \in \mathbb{N}$, are dense in $C(I_n)$.

In other words given $g \in C(I_n)$ and given $\epsilon > 0$ there is a sum $\psi(x)$ as above such that

$$|\psi(x) - g(x)| < \epsilon \quad \forall x \in I_n.$$

Since \tanh is a continuous function, it follows that the family of functions defined by equation (14) is a linear subspace of $C(I_n)$. To prove that it is dense in $C(I_n)$ we will use the same argument we used for the continuous sigmoidal functions.

This is, called U the set of functions of the form (14), we assume that U is not dense in $C(I_n)$. Thus, by the Hahn-Banach theorem there exists a not null bounded linear functional on $C(I_n)$ with the property that it is zero on the closure of U . By the Riesz Representation Theorem, the bounded linear functional can be represented by a Radon measures. Then using the definition of U we will see that this measure must be the zero measure, hence the functional associated with it is null contradicting the hypothesis.

We define

$$h_\lambda(x) = \frac{1}{2} \left[\prod_{k=1}^m \tanh(\lambda(w_k(x - b_k)) + \phi) + 1 \right]. \quad (15)$$

To proceed with the proof as in the case of the proof for continuous sigmoidal functions, we need only to understand to what converges the function

$$\psi_{\lambda_2, \lambda_1}(x) = \sum_{j=1}^N \frac{\alpha_j}{2} h_{j\lambda_2} \left(\sum_{i=1}^n h_{i\lambda_1}(x) \right) \quad (16)$$

when λ_1 and λ_2 tend to infinity, and $h_{i\lambda}$ indicates the processing function related to input i .

Once we have shown that for some choice of the parameters they converge pointwise to step function we can use the same argument we used in the proof of Theorem 5.4.

The first step is therefore to study the limit of equation (16). Let us focus of the multiplication of \tanh in the first layer, given by equation (15).

The pointwise limit of $h_\lambda(x)$ for $\lambda \rightarrow \infty$ depends on the sign of the limit of the product of \tanh , that in turn depends on the sign of $w_k(x - b_k)$ for $k \in \{1, \dots, m\}$.

Remark 3. *We remark that for $x \in [0, 1]$, from the limit of equation (15) we can obtain the indicator functions of set of the form $x > b$ or $x < b$ for any $b \in \mathbb{R}$. We just have to choose the parameters in such a way that only one of the \tanh in the multiplication is relevant. Let us define $Z = \{k \in \{1, \dots, m\} : w_k(x - b_k) > 0 \quad \forall x \in [0, 1]\}$. If $|Z| = m - 1$, i.e. there is only one $i \in \{1, \dots, m\}$ so that its weight are significant it holds that*

$$\lim_{\lambda \rightarrow \infty} h_\lambda(x) = v(x) = \begin{cases} 1 & \text{if } w_i(x - b_i) > 0 \\ \sigma(2\phi) & \text{if } w_i(x - b_i) = 0 \\ 0 & \text{if } w_i(x - b_i) < 0 \end{cases}$$

taking into account that $\sigma(2\phi) = \frac{1}{2} (\tanh(\phi) + 1)$.

Proof of Theorem 5.5. Considering Remark 3, the proof of Theorem 5.5 is analogous to that of Theorem 5.4. \square

B Experimental settings

All code was written in Python Programming Language. In particular, the following libraries were used for the algorithms: tensorflow for neural networks, scikit-learn for Logistic Regression, Decision Trees and Gradient Boosting Decision Trees.

A small exploration was made to determine the best structure of the neural network for each dataset. We used a breadth-first search algorithm defined as follows. We started with a network with just one neuron, we trained it and evaluated its performance. At each step, we can double the number of neurons in each layer except the output one or increase the depth of the network by adding a layer with one neuron. For each new configuration, we build a new structure based on it, initialize it and train it. If the difference between the accuracy achieved by the new structure and that of the previous step is lower than 1%, then a patience parameter is reduced by 1. The patience parameter is initialized as 5 and is passed down from a parent node to its spawned children, so that each node has its own instance of it. When patience reach 0, that configuration will not spawn new ones.

Before the neural network initialization, a random seed was set in order to reproduce the same results. As for the initialization of IAN, the weights w are initialised using the glort uniform. For the biases b of the first layer a uniform between the minimum and the maximum of each feature was used, while for the following layers a uniform between the minimum and the maximum possible output from the neurons of the previous layer was used.

For the network training, Adam with a learning rate equal to 0.1 was used as optimization algorithm. The loss used is the binary or categorical crossentropy, depending on the number of classes in the dataset. In the calculation of the loss, the weight of each class is also taken into account, which is inversely proportional to the number of samples of that class in the training set. The maximum number of epochs for training has been fixed at 10000. To stop the training, an early stopping method was used based on the loss calculated on the train. The patience of early stopping is 250 epochs, with the variation that in these epochs the loss must decrease by at least 0.01. Not using a validation dataset may have led to overfitting of some structures, so in future work we may evaluate the performance when using early stopping based on a validation loss. The batch size was fixed at 128 and the training was performed on CPU or GPU depending on which was faster considering the amount of data. The Heaviside was trained as if its derivative was the same as the sigmoid.

For Decision Trees (DT) and Gradient Boosting Decision Trees (GBDT), an optimisation of the hyperparameters was carried out, in particular for `min_samples_split` (between 2 and 40) and `min_samples_leaf` (between 1 and 20). For GBDT, 1000 estimators were used, while for DT the `class_weight` parameter was set. For the rest of the parameters, we kept the default values of the python sklearn library.

C Datasets

19 out of 23 datasets are publicly available, either on the UCI Machine Learning Repository website or on the Kaggle website. Here we present a full list of the datasets used, correlated with their shortened and full-length name, and the corresponding webpage where the description and data can be found.

The 4 synthetic datasets of our own creation are composed of 1000 samples with 2 variables generated as random uniforms between -1 and 1 and an equation dividing the space into 2 classes. The 4 equations used are:

- bisector: $x_1 > x_2$
- xor: $x_1 > 0 \wedge x_2 > 0$
- parabola: $x_2 < 2x_1^2 - \frac{1}{2}$
- circle $x_1^2 + x_2^2 < \frac{1}{2}$

These datasets are also represented in Figure 7.




Figure 7: The synthetically generated datasets we used to assess the soundness of our methodology.

D Examples

Heart dataset - Heaviside IAN

The Statlog Heart dataset is composed of 270 samples and 13 variables of medical relevance. The dependent variable is whether or not the patient suffers from heart disease. In Figure 8 you can find the network based on Heaviside IAN trained on the heart dataset. Only the inputs with a relevant contribution to the output are shown. From now on, we will indicate with $R_{k,j,i}$ the rule related to the processing function corresponding to the i -th input, of the j -th neuron, of the k -th layer. From the first neuron of the first layer we can easily retrieve the following rules: $R_{1,1,1} = x_1 \leq 54.29$, $R_{1,1,3} = x_3 \leq 3.44$, $R_{1,1,4} = x_4 \leq 123.99$, $R_{1,1,5} = x_5 \geq 369,01$, $R_{1,1,9} = x_9 \leq 0.48$, $R_{1,1,10} = x_{10} \leq 1.22$, $R_{1,1,11} = x_{11} \leq 1.44$, $R_{1,1,12} = x_{12} \leq 0.52$, $R_{1,1,13} = x_{13} \leq 6.26$. The second neuron of the first layer is not shown for

Short name	Full-length name	Webpage
adult	Adult	<UCI_MLR_URL>/adult
australian	Statlog (Australian Credit Approval)	<UCI_MLR_URL>/statlog+(australian+credit+approval)
b-c-w	Breast Cancer Wisconsin	<UCI_MLR_URL>/Breast+Cancer+Wisconsin+(Diagnostic)
car	Car Evaluation	<UCI_MLR_URL>/car+evaluation
cleveland	Heart Disease	<UCI_MLR_URL>/heart+disease
crx	Credit Approval	<UCI_MLR_URL>/credit+approval
diabetes	Diabetes	https://www.kaggle.com/uciml/pima-indians-diabetes-database
german	Statlog (German Credit Data)	<UCI_MLR_URL>/statlog+(german+credit+data)
glass	Glass Identification	<UCI_MLR_URL>/glass+identification
haberman	Haberman’s Survival	<UCI_MLR_URL>/haberman%27s+survival
heart	Statlog (Heart)	<UCI_MLR_URL>/statlog+(heart)
hepatitis	Hepatitis	<UCI_MLR_URL>/hepatitis
image	Statlog (Image Segmentation)	<UCI_MLR_URL>/Statlog+(Image+Segmentation)
ionsphere	Ionosphere	<UCI_MLR_URL>/ionsphere
iris	Iris	<UCI_MLR_URL>/iris
monks-1	MONK’s Problems	<UCI_MLR_URL>/MONK%27s+Problems
monks-2	MONK’s Problems	<UCI_MLR_URL>/MONK%27s+Problems
monks-3	MONK’s Problems	<UCI_MLR_URL>/MONK%27s+Problems
sonar	Connectionist Bench	<UCI_MLR_URL>/Connectionist+Bench+(Sonar,+Mines+vs.+Rocks)

Table 2: Publicly available datasets, with the short name used in in our work, their full-length name and the webpage where data and description can be found. The UCI_MLR_URL is the following: <https://archive.ics.uci.edu/ml/datasets/>

lack of space, but its obtained rules are $R_{1,2,2} = x_2 \geq 0.79$, $R_{1,2,3} = x_3 \geq 3.59$, $R_{1,2,4} = x_4 \geq 99.95$, $R_{1,2,5} = x_5 \geq 253.97$, $R_{1,2,8} = x_8 \leq 97.48$, $R_{1,2,9} = x_9 \leq 0.04$, $R_{1,2,10} = x_{10} \geq 2.56$, $R_{1,2,11} = x_{11} \geq 1.53$, $R_{1,2,12} = x_{12} \geq 0.52$, $R_{1,2,13} = x_{13} \geq 5.47$. Moreover, input x_7 gives always 1, so this must be taken into consideration in the next layer.

Moving on to the second layer, we can see in the first neuron that the second input is irrelevant, since the Heaviside is constant. The first processing function activates if it receives an input that is greater or equal to 2.99. Given that the input can only be an integer, we need at least 3 of the rules obtained for the first neuron of the first layer to be true: $R_{2,1,1} = 3 - of - \{R_{1,1,i}\}$. Following the same line of reasoning, in the second neuron of the second layer we see that we get $R_{2,2,1} = 5 - of - \{\neg R_{1,1,i}\}$ and $R_{2,2,2} = 5 - of - \{R_{1,2,i}\}$ (5 and not 6 because of x_7 processing function).

In the last layer, the first processing function has an activation of around 2.5 if it receives an input that’s less than 1.17. This can happen only if $R_{2,1,1}$ does not activate, so we can say: $R_{3,1,1} = \neg R_{2,1,1} = 7 - of - \{\neg R_{1,1,i}\}$. The second processing function gives a value of around -2.5 only if it gets an input less than 0.99, so only if the second neuron of the second layer does not activate. This means that $R_{2,2,1}$ and $R_{2,2,2}$ must be both false at the same time, so we get $R_{3,1,2} = \neg R_{2,2,1} \wedge \neg R_{2,2,2} = 5 - of - \{R_{1,1,i}\} \wedge 6 - of - \{\neg R_{1,2,i}\}$. Now there are 4 cases for the sum, i.e. the combinations of the 2 activations: $\{0 + 0, 2.5 + 0, 0 - 2.5, 2.5 - 2.5\} = \{-2.5, 0, 2.5\}$. Given that both have around the same value for the α parameter, the set is reduced to two cases. Looking at the processing function, we can see that is increasing with respect to the input, so since α_1 is positive, we can say that rule $R_{3,1,1}$ is

correlated to class 1, while rule $R_{3,1,2}$, having a negative α_2 , has an opposite correlation. Looking at its values, we can see that for both 0 and 2.5 inputs, the activation function gives an output greater than 0.5. If we consider this as a threshold, we can say that only for an input of -2.5 we get class 0 as prediction. This happens only if $R_{3,1,2}$ is true and $R_{3,1,1}$ is false. Summarizing we get $R_0 = R_{3,1,2} \wedge \neg R_{3,1,1} = 5 - of - \{R_{1,1,i}\} \wedge 6 - of - \{\neg R_{1,2,i}\} \wedge 3 - of - \{R_{1,1,i}\} = 5 - of - \{R_{1,1,i}\} \wedge 6 - of - \{\neg R_{1,2,i}\}$, so that we can say “if R_0 then predicted class is 0, otherwise is 1”.

Although we are not competent to analyse the above results from a medical perspective, it is interesting to note for example that the variables x_1 and x_4 , representing age and resting blood pressure respectively, are positively correlated with the presence of a heart problem.

Xor - sigmoid IAN

Our custom xor dataset divides the 2D plane in quadrants, with the opposites having the same label.

The network based on sigmoid IAN trained on xor dataset is represented in Figure 9. As we can see, all the processing functions of the first layer converged to nearly the same shape: a steep inverted sigmoid centered in 0. Therefore, we can say the rules obtained are $R_{1,1,1} = R_{1,2,1} = x_1 \leq 0$ and $R_{1,1,2} = R_{1,2,2} = x_2 \leq 0$. In the last layer, the first processing function has a value of about -15 for inputs in $[0, 1]$, then it starts growing slowly to reach almost 0 for an input of 2. This tells us that it doesn’t have an activation if both rules of the first neuron are true, so if $x_1 \leq 0 \wedge x_2 \leq 0$. On the other hand, the second processing function has no activation if its input greater than 1, that happens for example if we have a clear activation from at least one of the inputs in the second neuron of the first layer. So looking at it the opposite way, we need both those rules to be false

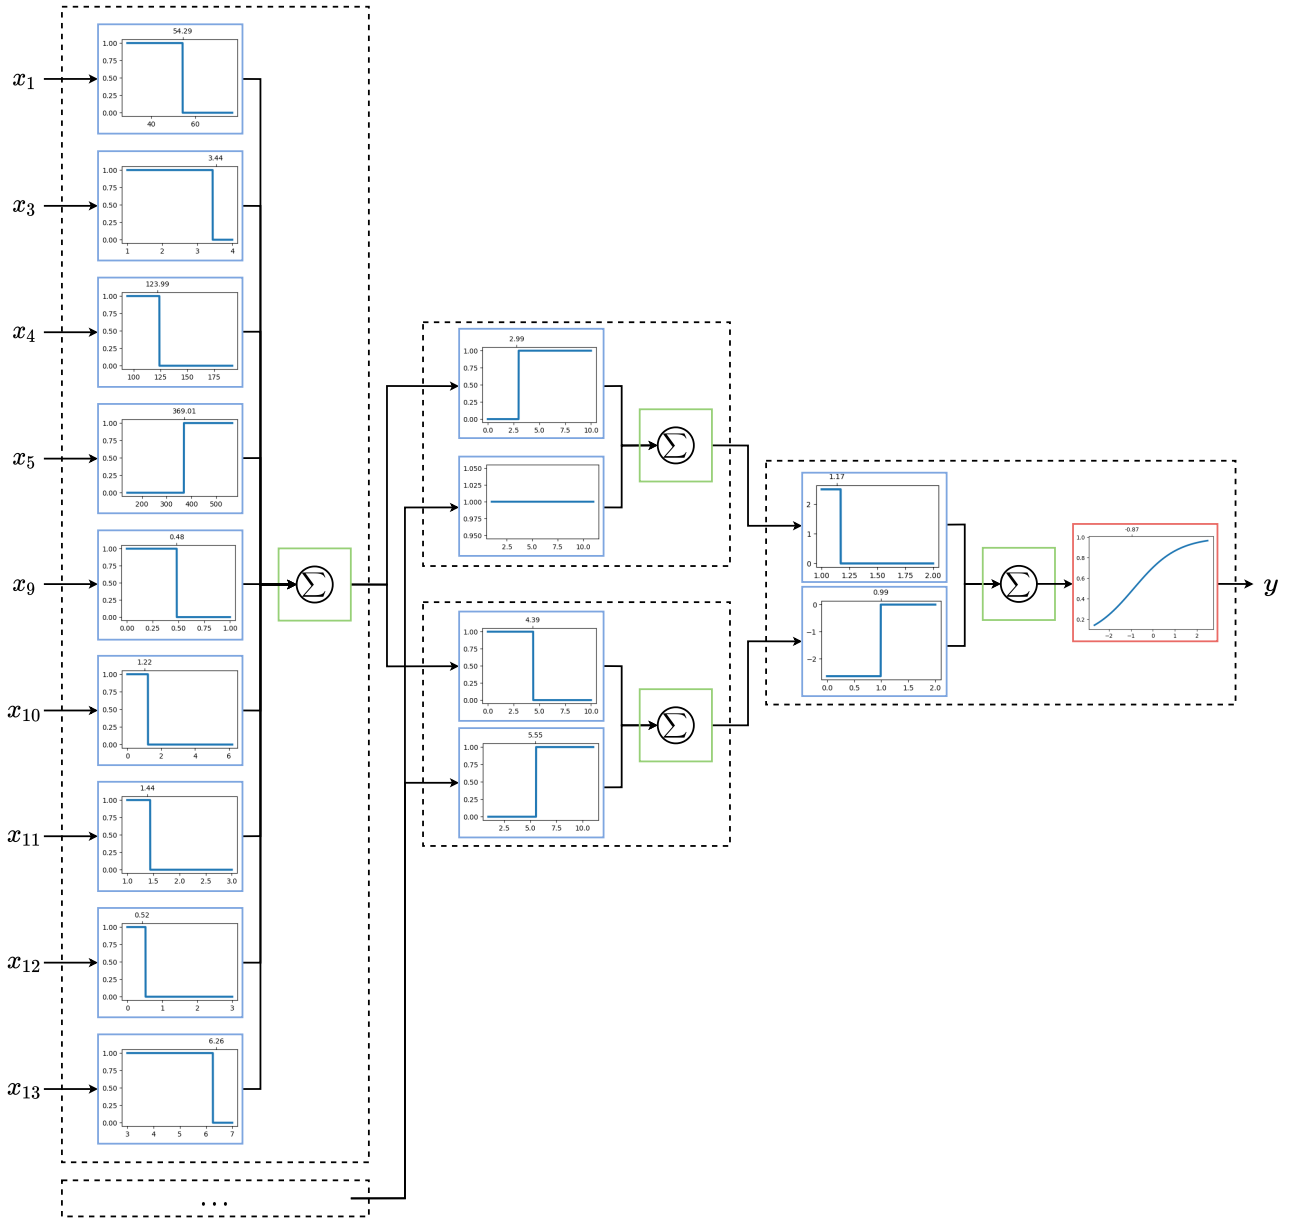


Figure 8: The Heaviside IAN Network trained on the heart dataset. The Figure follows the color convention used for NEWRON.

$(x_1 > 0 \wedge x_2 > 0)$ to have an activation of 12.5. The activation function is increasing with respect to the input, and to get a clear class 1 prediction, we need the input to be at least -5 . Considering if the processing functions could give only $\{-15, 0\}$ and $\{12.5, 0\}$ values, just in the case we got -15 from the first one and 0 from the second one it would give us a clear class 0 prediction. This happens only if $\neg(x_1 \leq 0 \wedge x_2 \leq 0) = x_1 > 0 \vee x_2 > 0$ and $\neg(x_1 > 0 \wedge x_2 > 0) = x_1 \leq 0 \vee x_2 \leq 0$, that can be summarised $(x_1 > 0 \vee x_2 > 0) \wedge (x_1 \leq 0 \vee x_2 \leq 0) = (x_1 > 0 \wedge x_2 \leq 0) \wedge (x_1 \leq 0 \wedge x_2 > 0)$. Since this rule describes the opposite to xor, for class 1 we get the exclusive or logical operation.

Iris dataset - tanh-prod IAN

A dataset widely used as a benchmark in the field of machine learning is the Iris dataset. This contains 150 samples, divided into 3 classes (setosa, versicolor and virginica) each representing a type of plant, while the 4 attributes represent in order sepal length and width and petal length and width.

In Figure 10 you can see the final composition of the network generated with the tanh-prod2 IAN neuron.

Considering the first neuron of the first layer, we see that it generates the following fuzzy rules: $R_{1,1,2} = x_2 > 3.08$ (sepal width), $R_{1,1,3} = x_3 < 5.14$ (petal length) and $R_{1,1,4} = x_4 < 1.74$ (petal width). For the first attribute (sepal length) it does not generate a clear rule, but forms a

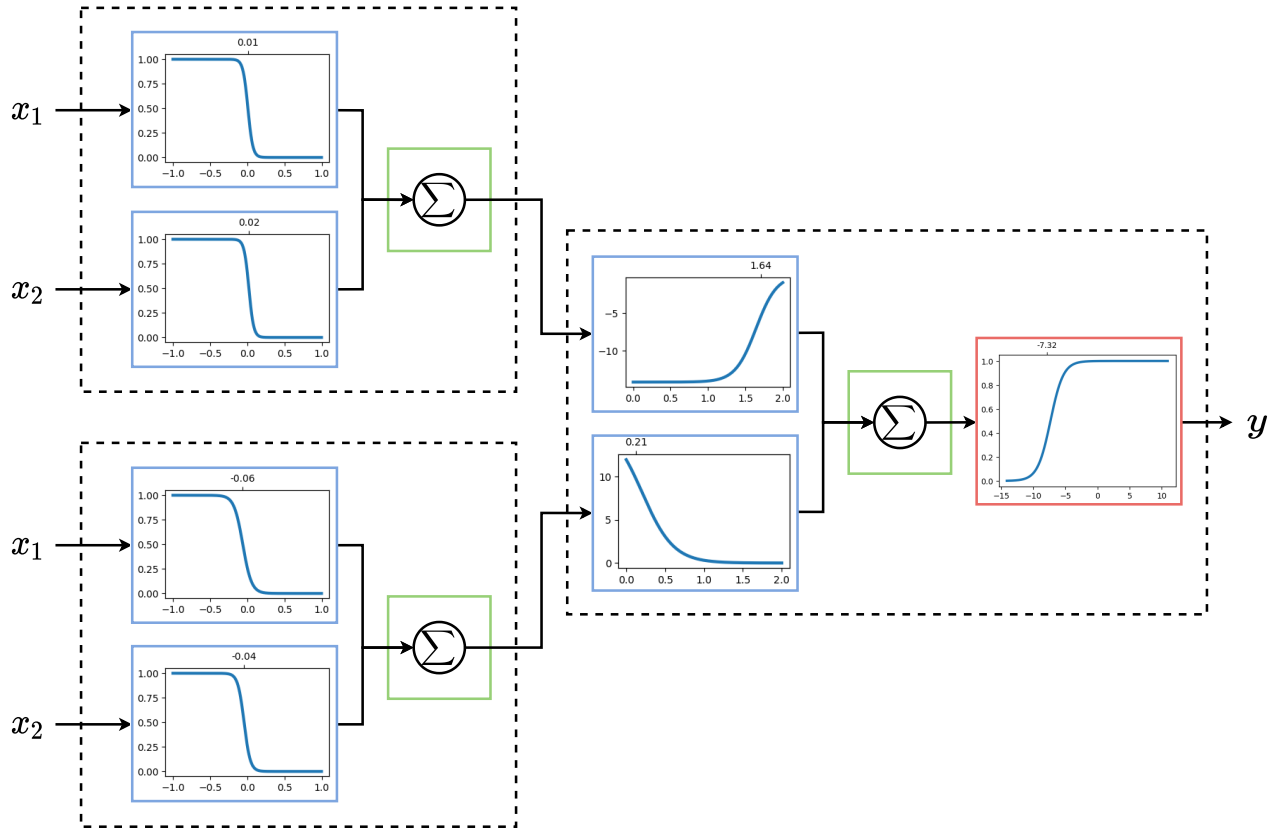


Figure 9: The sigmoid IAN Network trained on the xor dataset. The Figure follows the color convention used for NEWRON.

bell shape, reaching a maximum of 0.5. This tells us that x_1 is less relevant than the other attributes, since, unlike the other processing functions, it does not reach 1.

The second neuron has an inverse linear activation for the first attribute, starting at 0.7 and reaching almost 0. The second attribute also has a peculiar activation, with an inverse bell around 2.8 and a minimum value of 0.4. The third and fourth attributes have clearer activations, such as $R_{1,2,3} = x_3 < 2.51$ and $R_{1,2,4} = x_4 < 1.45$.

The fact that petal length and width are the ones with the clearest activations and with those specific thresholds are in line with what has previously been identified on the Iris dataset by other algorithms.

We denote by $y_{k,j}$ the output of the j -th neuron of the k -th layer. Moving on to the second layer, the first neuron generates the rules “if $y_{1,1} < 1.83$ ” and “if $y_{1,2} < 2.66$ ”, while the second one generates “if $y_{2,1} > 2.08$ ” and “if $y_{2,2} > 2.22$ ”. Combined with what we know about the previous layer, we can deduce the following: $y_{1,1}$ is less than 1.83 only if the sum of the input activation functions is less than 1.83, which only happens if no more than one of the last three rules is activated ($0 + 1 + 0 < 1.83$), while the first one, even taking its maximum value, is discriminative only when the input of one of the other rules is close to the decision threshold ($0.5 + 1 + 0 + 0 < 1.83$, while $0.5 + 1 + 0.5 + 0 > 1.83$). For $y_{1,2} < 2.66$, there are more cases. We can divide the second processing function of the second neuron of the first layer in

two intervals: one for which $x_2 < 3.2$ and the other when $x_2 \geq 3.2$. In the first interval, the processing function gives a value that is less than 0.66, greater in the second one. With this, we can say that $y_{1,2} < 2.66$ even if $R_{1,2,3}$ and $R_{1,2,4}$ activates, if $x_2 < 3.2$ and x_1 is near its maximum.

In the second neuron of the second layer, the first processing function is nearly the exact opposite to that of the other neuron; we need at least two of $R_{1,1,2}$, $R_{1,1,3}$ or $R_{1,1,4}$ to be true, while $R_{1,1,1}$ still doesn’t have much effect. The second processing function gives us $y_{1,2} > 2.22$. Considering that the minimum for the processing function related to x_2 is 0.4, we may need both rules $R_{1,2,3}$ and $R_{1,2,4}$ to be true to exceed the threshold, or just one of them active and x_1 to take on a low value and x_2 to be a high value.

For the last layer, remember that in this case since there are more than 2 classes, a softmax function is used to calculate the output probability, hence the arrows in the figure that join the layers of the last layer.

For the first output neuron, in order to obtain a clear activation, we need the first input to be less than 0.46 and the second greater than 1.42. This is because the α_i are 3 and -8 , and the output activation function starts to have an activation for values greater than -2 . This means that the first neuron of the second layer should hardly activate at all, while the other should activate almost completely. Considering the thresholds for $y_{1,1}$ and $y_{1,2}$, we need the first to be greater than 2.08 and the other to be greater than 2.66. So

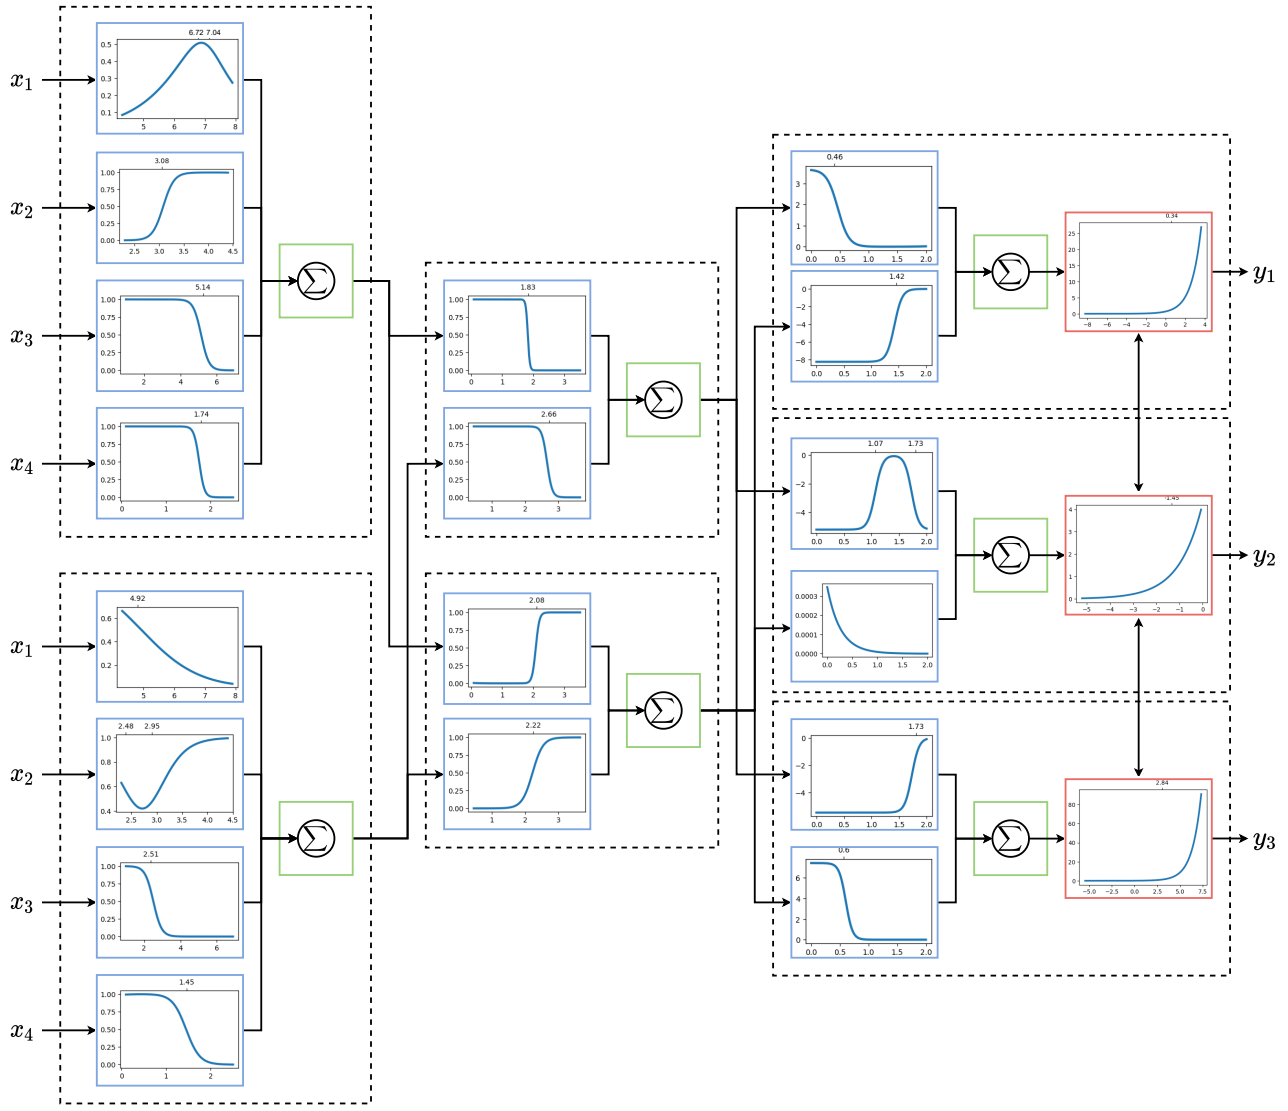


Figure 10: The tanh-prod IAN Network trained on the iris dataset. The Figure follows the color convention used for NEWRON.

$R_{3,1,1} = 2 - of - \{x_2 > 3.08, x_3 < 5.14, x_4 < 1.74\}$. For $R_{3,1,2}$ is more tricky to get a clear decision rule, but we can say that we may need both $R_{1,2,3}$ and $R_{1,2,4}$ to be true and $x_2 \geq 3.2$. If $x_2 < 3.2$, we need x_1 to not be near its maximum value. If just one of those two rules is true, we need $x_2 < 3.2$ and x_1 near 4, or $x_2 > 3.2$ but with a (nearly) direct correlation with x_1 , such that the more x_1 increases, the same does x_2 .

In the second output neuron, the second processing function is negligible, while the first one forms a bell shape between 1 and 2. This means that it basically captures when $y_{2,1}$ has a value of approximately 1.5, so when the decision is not clear. This is what gives this neuron maximum activation.

In the third and last output layer, since the first processing function has a negative α parameter and the activation function is increasing with respect to the input, we want it to

output 0, and this requires maximum activation for the first neuron of the second layer. Regarding the second processing function, we want it to output 8, so we need nearly no activation from the second neuron of the second layer. So we need the first neuron of the first layer to output a value lower than 1.83 and the second neuron to output a value lower than 2.22. This means that no more than one rule $R_{1,1,i}$ needs to be active and at most two rules of $R_{1,2,i}$ need to be true.

We can conclude by saying that both neurons of the first layer are positively correlated with class 1, while they are negatively correlated with class 3. This means that low values of x_3 and x_4 , or high values of x_2 increase the probability of a sample to belong to class 1, while x_1 has almost no effect. For class 2, what we can say is that it correlates with a non-maximum activation of both neurons of the first layer, meaning that it captures those cases in which the prediction of one of the other classes is uncertain.