

Research Article

HorNets: Learning from Discrete and Continuous Signals with Routing Neural Networks

Boshko Koloski^{1,2}, Nada Lavrač¹, Blaž Škrlj¹

1. Jožef Stefan Institute, Ljubljana, Slovenia; 2. Jožef Stefan Postgraduate School, Slovenia

Construction of neural network architectures suitable for learning from both continuous and discrete tabular data is a challenging research endeavor. Contemporary high-dimensional tabular data sets are often characterized by a relatively small instance count, requiring data-efficient learning. We propose HorNets (Horn Networks), a neural network architecture with state-of-the-art performance on synthetic and real-life data sets from scarce-data tabular domains. HorNets are based on a clipped polynomial-like activation function, extended by a custom discrete-continuous routing mechanism that decides which part of the neural network to optimize based on the input's cardinality. By explicitly modeling parts of the feature combination space or combining whole space in a linear attention-like manner, HorNets dynamically decide which mode of operation is the most suitable for a given piece of data with no explicit supervision. This architecture is one of the few approaches that reliably retrieves logical clauses (including noisy XNOR) and achieves state-of-the-art classification performance on 14 real-life biomedical high-dimensional data sets. HorNets are made freely available under a permissive license alongside a synthetic generator of categorical benchmarks.

Boshko Koloski and Blaž Škrlj equally contributed to this work.

Corresponding authors: Boshko Koloski, boshko.koloski@ijs.si; Blaž Škrlj, blaz.skrlj@ijs.si

1. Introduction

In recent years, neural network models have shown great promise when modeling complex real-life data sets such as images, texts, and networks^[1]. Data sets are most frequently represented in the form

of propositional tables. Bader et al.^[2] discussed the link between propositional logic and feedforward neural networks. They have successfully demonstrated that non-conventional activations taking values on the interval $[-1, 1]$ need to be adopted to model simple propositional expressions. They also discuss how such expressions can be composed into search trees and binary decision diagrams. Similar ideas were explored by Al Iqbal^[3], where decision trees were co-trained with neural networks to extract simple rules. The more recent work of Labaf et al.^[4] also explored how background knowledge can be incorporated into the learning of logical expressions by approximating the decision process of neural networks with knowledge-aware logical rule learners.

In parallel, this work is also inspired to a large extent by the recent progress in *neuro-symbolic machine learning*. For example, the works of Amizadeh et al.^[5] explored the link between logical induction, coupled to low-level (neural) pattern recognition. Similarly, works like DeepProbLog^[6] explore the language design required to bridge the two learning paradigms (i.e. extension of Prolog with a neural predicate).

Such insights are key motivators of this study, as we show that neural networks can be used to produce, and operate directly with *interpretable* logical expressions. The main contributions of this paper are the following:

- We propose HorNets (Horn Networks), a neural network architecture capable of operating with Horn clauses, that demonstrably achieves state-of-the-art performance on synthetic discrete/categorical and continuous data sets.
- HorNets are based on polyClip activation function – an activation for which we show a probabilistic link to zero-order (propositional) logic.
- To assess the HorNets capability to model logical gates, we conducted extensive benchmarks on synthetic data obtained by a generator implemented as part of this work. We demonstrate that HorNets neural network architecture achieves the best recall in terms of being able to model the most common logic gates (and with that, binary interactions).
- Competitive performance of HorNets is demonstrated on more than ten real-life data sets from the biomedical domain, demonstrating it outperforms strong approaches such as gradient-boosted trees and selected AutoML-based classifiers.

The rest of this paper is structured as follows. We begin by discussing the related work in Section 2, followed by the proposed methodology and its theoretical implications in Section 3. In Section 4, we present the considered experimental setting (on synthetic data in Section 4.1 and on real data in

Section 4.2) , followed by the results in Section 5. Finally, in Section 6, we discuss the achievements and drawbacks of this work and present future research directions.

2. Related work

This section overviews some of the related work that impacted the development of HorNets. We begin by discussing binarized neural networks, followed by the body of literature focusing on *neuro-symbolic learning*.

Exploration of how rule lists can be transformed into networks capable of classification was performed by^[7]. Their networks could account for AND, OR and NOT connectives, offering expressions in full propositional logic. The subfield of studying neural networks that also resonates with the proposed approach concerns binary neural networks. For example, the work of Rastegari et al. ^[8] demonstrates that a neural network capable of approximating the XNOR truth table offers sufficient performance on ImageNet, indicating high expressiveness of such networks – by being able to model XNOR, this was one of the first approaches that demonstrated capability to model general logical statements. Similar results were achieved on CIFAR-10 and MNIST data sets^[9]. Binary neural networks were also shown as a promising variant of models that can be computed on specialized hardware such as FPGAs^{[10][11]}.

Alternatively to considering neural networks that directly approximate logical expressions, such expressions can be generated before being fed into a neural network. This second branch of methods that inspired this work revolves around the notion of *neuro-symbolic learning*. This paradigm explores how neural network-based low-level recognition (e.g., vision) can be coupled with the notion of *reasoning*. Such ideas culminated into the concept known as *deep relational machines*^{[12][13]}, offering a state-of-the-art performance for e.g., molecule classification tasks. DeepProbLog^[14] combined the idea of probabilistic logic and deep learning, enabling probabilistic inference over deep latent structures, thus providing a tool for tasks such as distant supervision tasks. DeepProbLog^[14] integrates probabilistic logic with deep learning, enabling inference over deep latent structures. This framework is ideal for tasks that require background knowledge for validation and utilize deep learning for representation learning in distant supervision scenarios. Contrary to the work of^[2], where programs were extracted from feedforward neural networks, an inductive logic program was first used to construct features, which were used to train a deep neural network model. More

recently, the link between deep learning and propositionalization was explored by^[15], resulting in two fundamentally different ways to learn latent representations from a propositionalized relational database.^[16] ^[16] recently a deeper overview on inductive logic programming and its connections to neuro-symbolic methods.

The recent work on closing the loop between recognition and reasoning^[17] introduced a grammar model as a symbolic prior to bridge neural perception and symbolic reasoning alongside a top-down, human-like induction procedure. This work demonstrated that such a combined approach significantly outperforms the conventional reinforcement learning-based baselines. The Microsoft research division recently explored the interplay between visual recognition and reasoning^[5]. They introduced a framework to isolate and evaluate the reasoning aspect of visual question answering separately from its perception, followed by a calibration procedure that explores the relation between reasoning and perception. Further, a neuro-symbolic approach to logical deduction was proposed as *Neural Logic Machines*^[18]. This architecture was shown to have inductive logic learning capabilities, which were demonstrated on simple tasks such as sorting. Incorporating semantics-aware logical reasoning to understand better network substructure (communities) was recently explored^[19].

Using neural networks as a tool for learning rules has recently gained importance.^[20] explored the application of neural networks to learn rule sets for binary classification using the network structure. ^[21] proposed DR-Net – a two-layer neural network architecture for learning logical rule sets. The first layer maps directly interpretable if-then rules, while the output layer forms a disjunction of these rules. Regularisation based on sparsity is used to derive simpler but representative logical rules. On the other hand,^[22] proposed explainable neural rule learning in a differentiable way. The main idea of the approach is the initial construction of atomic propositions (learned as separate operators via semi-supervised learning) and their subsequent evolution to a binary tree topology to express multiple rules via a neural architecture search. Finally, an ensemble of such trees is created and later used for classification by voting. Similar to the work of^[20], it is limited to binary classification only. ^[23] recently proposed RuleNet – an approach that builds on the work of DR-Net, with the additional constraint that the rules must be ordered. In rigorous evaluation, they showcase the network's strong performance on synthetic and real-life classification tasks. However, one significant drawback of the proposed method is the expectation that the input is binary, which is unrealistic for most real-world applications.

Further, a link between combinatorial optimization solvers and deep learning was proposed^[24], offering novel insights into solving hard, e.g., graph-based combinatorial problems. In recent works concerning *program synthesis*, the auto-encoding logical programs^[25] as well as the DreamCoder approach^[26] both demonstrated that symbolic representations can be learned (albeit differently), applicable to solving e.g., physics-related equation discovery (DreamCoder) and learning compact relational representations of a given relational structure (auto-encoding logical programs).

Finally, the proposed work bases some of the ideas from the recent advancements in *understanding* of black-box models. Many deep neural network models are inherently black-box, offering little to no insight into the most representative patterns being learned. As such, the recent trend of model *explainability* in terms of *post-hoc* approximations of model decisions is actively studied. Widely used tools for such approximations are e.g., LIME^[27] and SHAP^[28]. However, both methods are prone to adversarial attacks, making their ubiquitous use questionable^[29]. Alternatively to *post-hoc* explanation of predictions, recent attempts such as TabNet^[30] and propositional Self Attention Networks (SANs)^[31] indicate that highly relevant features can be extracted via the attention and similar mechanisms directly.^[32] proposed KAN, a network built on the concept that weights should be placed on the edges, with node activations represented as interactions between functions of the edges, inspired by the Kolmogorov-Arnold representation theorem. This approach provides for interpretable networks, although it could be more scalable.

3. HorNets Architecture Overview

We begin the section with an overview of the HorNets architecture, followed by a theoretical analysis of its expressiveness and computational complexity.

3.1. General overview and intuition

We begin by discussing a high-level overview of the HorNets architecture, followed by a more detailed inspection in the following sections. The architecture consists of three main components: the *routing operator*, *categorical*, and *continuous blocks*. A schematic overview is shown in Figure 1.

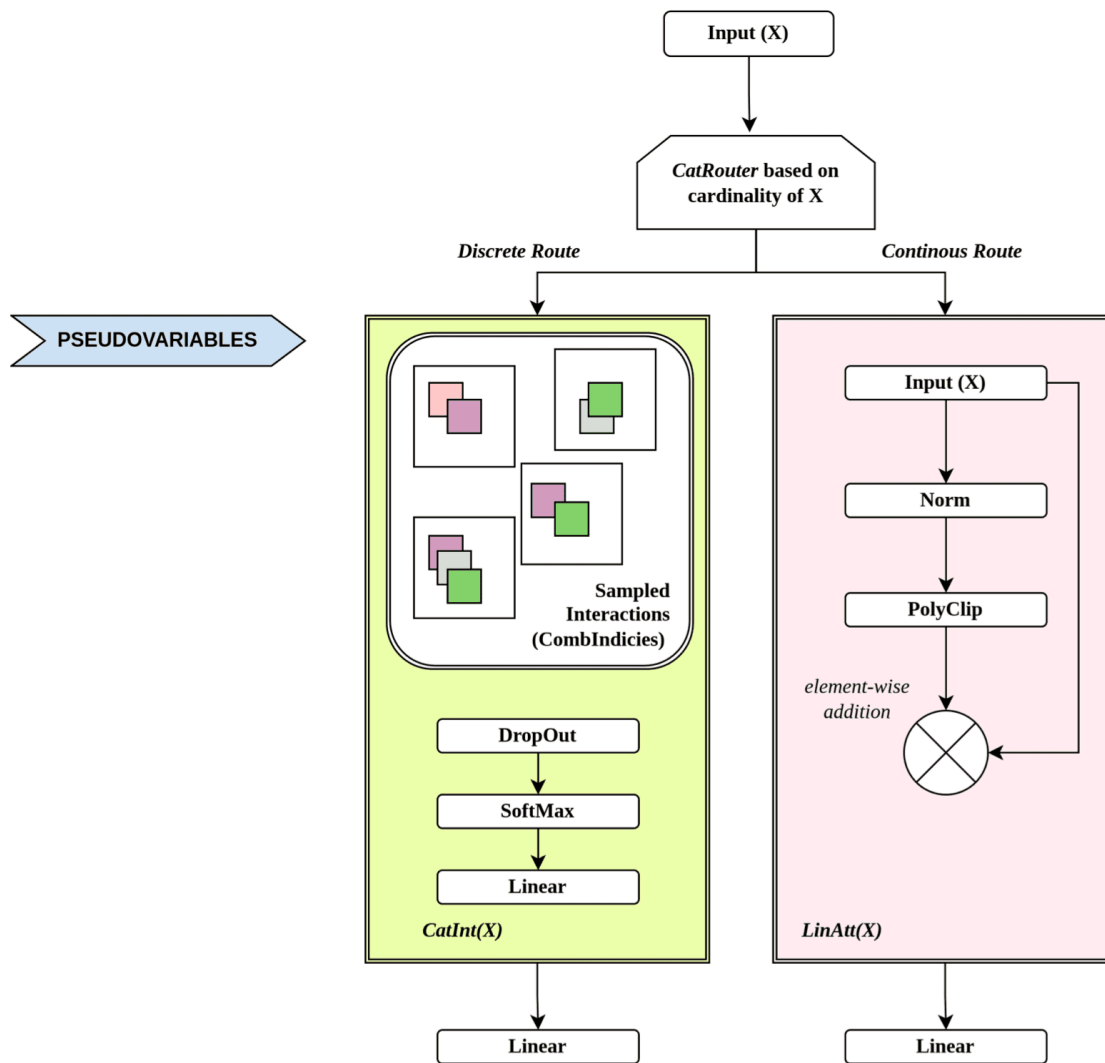


Figure 1. Overview of HorNets architecture. The *CatRouter* (Category Router) component decides whether to treat input as discrete or continuous. If discrete, *CatInt* layer is invoked, conducting factorization-based estimation of a subspace of interactions. If *LinAtt* is invoked, an efficient, element-wise POLYCLIP-based operation is considered. Both routes end up with a linear layer.

Inputs are batches of tabular data. For each batch, the number of unique values is computed and used to decide which of the subsequent layers to activate. This routing logic is inspired by recent work in the mixture-of-experts-based large language model architectures^[33]. In this work, all experiments consider categorical layers only if cardinality is lower than 3 (binary matrices). If cardinality exceeds this bound, the continuous part of the architecture takes over. This design choice stems from empirical observation that even if the categorical part of the architecture is substantially scaled, it

converges much slower compared to the continuous one – one of the key purposes of this work is to demonstrate that HorNets is capable of handling (automatically) both scenarios in its current form.

3.2. Modeling logical operators via activations

We first discuss one of the main contributions of this paper, POLYCLIP activation, its formulation and connections to zero-order (propositional) logic. Conventionally used activations surveyed in^[34], such as ReLU and Sigmoid, offer truncation of weight values to positive real values and were shown to offer state-of-the-art performance on many tasks. On the other hand, activations such as ELU and LReLU offer the inclusion of negative weights; however, they are not symmetric and thus offer little insight into the effect of positive and negative weights. Activations such as *tanH* indeed offer symmetry around zero, and incorporate negative values, however, their infinite bounds hinder interpretability/capability to truncate them to logical clauses directly.

In this work, we introduce the POLYCLIP family of activation functions. For a given parameter $k \in \mathbb{N}_0$, the corresponding POLYCLIP activation function¹ is defined as follows:

$$\text{POLYCLIP}(x, k) = \text{clip}(x^{2 \cdot k + 1}, -1, 1) = \frac{x^{2k+1}}{\max\{|x^{2k+1}|, 1\}} = \begin{cases} -1; x \leq -1 \\ 1; x \geq 1 \\ x^{2k+1}; \text{else.} \end{cases} \quad (1)$$

Here, the x represents real-valued input and k a parameter, determining the order of the considered odd polynomial.

Note that for the rest of this paper, if A is a real-valued matrix, we denote with $\text{polyClip}_k(A)$ the matrix in which the POLYCLIP function is applied to each element. Apart from being efficient in computing, POLYCLIP activation offers a probabilistic means of expressing statements in formal logic, which we will discuss next.

The activation projects a weight's value into one of the three main regions, each of which can be associated with the direct interpretation of the POLYCLIP's effect on the input's value. The three regions of relevance to this work are summarized in Figure 2. Given an activated weight's value, it can be interpreted as (partial) *negation* (leftmost, negative part of the possible values) and conjunction (rightmost, positive part of the values). Further, values close to 0 imply that a given feature is potentially irrelevant and can be neglected. Multiplication with zero-activated weights in subsequent

layers is expected to impact the overall architecture’s output less. To maintain the domain and with it the (implicitly) discrete operation on a given input, the main difference between POLYCLIP and the previously mentioned activation functions is *the order* in which it is applied to a given weight and multiplied with a given input².




Figure 2. Overview of the POLYCLIP activation and its implications/interpretation when fully discretized $(-1, 0, 1)$ range is considered. The curves were obtained by varying the k parameter.

We next discuss the theoretical insights, related both to the computational complexity of the main steps in HorNets, as well as its expressiveness and relation to formal logic. We first discuss the behavior of POLYCLIP-activated weights with respect to different types of inputs. Let o represent the output of multiplying an activated and *discretized* weight with a given input; thus, $o \in \{-1, 0, 1\}$ – we assume discretization by *rounding*. POLYCLIP-activated weights can express negation and conjunction (NOT and AND), and thus model these two types of logical gates. As an example consider a matrix D , where each entry is a *discretized* weight from the initial weight matrix, defined as $D_{ij} = \text{round}(w_{ij})$. We consider a single hidden-layer neural network-based architecture first. Here, the layer performs an inner product of a column of weights and a row of feature values, i.e. for the k -th input vector r_k ,

the output is calculated as $\sum_j r_{kj} \cdot D_{ji}$. In this example, the output will be large (and thus considered during the prediction phase) only if all the summands contribute large values, i.e. if the conjunction of the truth values, to which the individual values $r_{kj} \cdot D_{ji}$ are mapped, is true, thus showing that POLYCLIP-activated weights can model conjunctions. Note that the absence (no contribution) of a given feature can also be expressed (iff $D_{ij} = 0$). To show that POLYCLIP-activated weights also express negation, we need to consider the case when $D_{ij} = -1$. In this case, the contributing summand will be large if and only if the j -th feature value is negative, i.e. if its 'negation' is true. Note that the mentioned relations of activations' values and their impact on output hold true for a single-layer architecture, and can be true (yet it is not guaranteed) for a multi-layer architecture – interchanging signs can have unpredictable effects in this case. Further, the direct modeling of gates holds true if no other activations are considered simultaneously. Next, by being able to consider both AND and NOT connectives, we can see that POLYCLIP-activated and discretized weights can model XNOR gate. Let $r_{kj} \in \{-1, 1\}$ and $D_{ji} \in \{-1, 1\}$, i.e. each input contributes either positive or negative partial weight value. Next, we establish a mapping where 1 represents True and -1 False. By considering the outputs (o) when considering cartesian product of the two input spaces ($\{-1, 1\} \times \{-1, 1\}$), the following results of the multiplication are possible. The $o = 1$ if both inputs are 1 or -1. If one of the inputs is -1 and the other is 1 (or vice versa), $o = -1$ (without considering the absence of features)

The demonstrated capability illustrates that by considering symmetric positive-negative activation values, clipped to the same range as the input space (-1 and 1 in this case), multiplications between particular values manifest in known gate definitions. Having shown that logical operations can be modeled via discretized activation regime, we can extend the claim to the following. POLYCLIP-activated weights can express statements in propositional logic.

Following these definitions, let $D_{ji} = -1$. The inner product $r_{kj} \cdot D_{ji}$ directly translates to $\bigwedge_j \neg(F_j)$. By employing De Morgan's law, $\bigwedge_j \neg F_j \iff \neg(\bigvee_j F_j)$, the conjunctions can be directly expressed. Being able to model conjunction, negation and disjunction, POLYCLIP-activated weights can express statements in disjunctive normal form and hence in propositional logic.

3.3. HorNets Architecture formulation

We continue the discussion with a more detailed formulation of HorNets architecture. Let $\mathbf{X} \in \mathbb{R}^{b \times d}$ represent an input batch of data (b being batch size and d the dimension). Let X represent the set of values in this batch. The routing layer, named `catRouter`, denotes a category routing function that distinguishes between two categories: discrete (`linAtt` – Linear Attention) and continuous (`catInt` – Categorical Interactions). The `catRouter` mechanism can be in its general form written as

$$\text{CATROUTER}(X, y) = \begin{cases} \text{LINATT}(X, y); & |X| > 2; \\ \text{CATINT}(X, y); & \text{else.} \end{cases}$$

The `CATROUTER(X)` is thus an efficient operation with negligible practical cost³. The subsequent possible operations, `LINATT(x)` and `CATINT(x)` are discussed next. The `LINATT(x)` is defined as (recursive formulation aligned with code)

$$\begin{aligned} x_0 &= \frac{x}{\max(\|x\|_p, \varepsilon)} \\ x_1 &= \text{POLYCLIP}(x_0) \otimes x \\ x_2 &= \mathbf{x}_1 \cdot \mathbf{w} + \mathbf{b}. \end{aligned}$$

This part of the architecture consists of three main steps; normalization p norm (ε is a small constant), activation and elementwise product (\otimes), followed by a linear layer.

The second part of the architecture `catRouter(X)` can be formulated as follows. Let $\mathbf{M} \in \mathbb{R}^{o \times |C|}$ represent the *interaction factorization* matrix, where o represents the order of interaction and C the space of input combinations to be considered. This matrix serves as the basis for representing different feature interactions separately, enabling *fine-grained control at interaction level*. The forward pass for this part of the architecture, however, is computationally substantially more expensive and can be formulated as

$$\begin{aligned} \mathbf{F} &= \text{COMBACTOP}(\mathbf{M}, \mathbf{x}, \text{CombIndices}) \\ x_0 &= \text{DROPOUT}(\mathbf{F}) \\ x_1 &= \text{SOFTMAX}(x_0) \\ x_2 &= x_1 \cdot \mathbf{w} + \mathbf{b}. \end{aligned}$$

The key component is COMBACTOP. This operator enables the traversal of considered feature combination space with incremental factorization (and activation). It is defined for a particular feature combination (and its indices) as

$$\text{COMBACTOP}(\mathbf{M}, \mathbf{x}, \text{CombIndices}) = \mathbf{F}[:, \text{combIndex}] = \text{polyClip}(x[:, \text{CombIndices}] \cdot \mathbf{M}[:, \text{CombIndices}]),$$

where CombIndices represent indices mapping to features in the factorization matrix, and CombIndex represents the index of the interaction's representation.

The main motivation for explicit factorization of interactions in this component is based on the observation that this type of intermediary factorization substantially simplifies accurate modeling when binary discrete spaces are considered as input. Consider a learning problem of the form

$$\{0, 1\}^{a \times b} \rightarrow \{0, 1\}^a; \quad b > 2.$$

Here b features are used to learn the association between the two spaces. Let us assume target space is defined as XOR(b_1, b_2), i.e. an XOR between the first two features. By considering all interactions of order two, i.e., Comb($b, 2$), one of the interactions will also require explicit factorization of representations for b_1 and b_2 . Assuming a neural network is capable of modeling XOR, by being explicitly forced to consider this interaction, it is highly probable it will identify the two features as key to modeling the relation – linking output to any other feature pair would imply the network learned a wrong interaction that does not generalize. In the context of HorNets, because the considered categorical input space is binary⁴, interaction factorization will take place as opposed to a linear counterpart that is more suitable for continuous inputs. In this work, we considered cross-entropy loss for end-to-end optimization of the architecture. The loss is defined as (for a pair of inputs x and y).

$$\ell(x, y) = -w_{y_n} \log \frac{\exp(x_{n,y_n})}{\sum_{c=1}^C \exp(x_{n,c})} \cdot \mathbb{1}_{\{y_n \neq \text{ignore_index}\}}.$$

Here, C is the set of all classes. Mean aggregation is considered (per batch).

3.4. Pseudovariables and the curse of higher order interactions

While developing the architecture, we observed that unless specific interactions were retrieved, performance on simple tasks such as modeling of logic gates could have been more consistent. By

considering higher-order interactions way beyond what is normally considered in matrix factorization (e.g., order 16 and more)⁵, we can force HorNets to sample from larger feature subspaces. However, if e.g., a subspace of 16 interactions contains the two features that govern the output signal, yet the remaining 14 represent pure noise, optimization (backpropagation) was observed to have issues with convergence. To remedy this shortcoming, we introduced the notion of *pseudovariabes* – artificially added features that, in the context of HorNets, serve as mask variables. These variables are initialized as constant space of ones, and don't change the results of intermediary computation. Each HorNets evaluation automatically introduced k pseudovariabes, where k is the order of interactions being modeled. The rationale for this design choice is that by being able to mask most of a higher-order interaction, lower-order interactions can be reliably retrieved, while at the same time, fewer samples of the feature space are required, thus reducing overall computational complexity – this observation is also one of the reasons HorNets are approachable on commodity hardware in its current form⁶. An interesting side-effect of pseudovariabes is increased interpretability. As the categorical part of the architecture performs softmax per-interaction, highest-scored interactions can be directly inspected, and serve as possible candidates for underlying rules.

If, for example, a rule of form $f_1 \wedge f_2 \wedge p_3 \wedge p_4$ is obtained, where p_x are pseudovariabes, we can reliably claim that only f_1 and f_2 have the impact on learning (as by definition p_3 and p_4 contribute nothing).

A note on computational complexity. Having discussed the architecture, it is worth noting that the two main components between which the router layer decides are subject to notably different computational complexity. In particular the LINATT(x, y) operation is linear with regards to the number of input dimensions, while CATINT(x, y) can be subject to a combinatorial explosion (many features, higher-order interactions), and is thus *exponential in the limit*⁷. In practice, however, we observed that a relatively small number of combinations of higher order is sufficient for modeling categorical data well. Many reasons for this are based on empirical evaluation *pseudovariabes*. Further, we implemented a Monte Carlo-based sampling of interactions within the network to mitigate for the explosion, effectively changing which interactions are considered per epoch. Hence, sufficiently long training cycles in the limit sample all possible combinations, even though in practice HorNets

converge faster. The architecture was built to be practical – all experiments discussed in the following sections can be replicated on a moderate off-the-shelf laptop with no GPU.

4. Evaluation

In this section, we set up the evaluation scenario. We focus on assessing the predictive power of our method – first on synthetic data, where we want to analyze the logic modeling capabilities and second on real-life downstream classification tasks.

4.1. Empirical evaluation on synthetic data

The first round of experiments revolves around benchmarking of HorNets' logic modeling capabilities. To systematically evaluate this trait, we designed a synthetic data generator that enables the construction of arbitrary-dimensional binary training datasets, where target space is defined as a logical combination of a given feature subspace. We considered the main logic operators, namely AND, OR, NOT, XOR and XNOR. We refer to these gates as OP. The dimension parameter d parametrizes the generator, and the instance count parameter c . It first generates a binary matrix $\mathbf{B} \in \{0, 1\}^{c \times d}$, where $\forall \mathbf{k} \in \mathbf{B} | \mathbf{k} \sim \text{BERNOULLI}(0.5)$. Note that OP are binary operators – they comprise second-order interactions. For each operator, generator computes the target vector $\mathbf{t} \in \{0, 1\}^c$ as $t_i = \text{OP}(k_{i,j_0}, k_{i,j_1})$. The remainder of the \mathbf{B} remains random (noise). The higher the d , the harder it is estimated that a given classifier can retrieve the explicitly encoded interaction (two features). Note that XNOR and XOR are of particular interest because both inputs must be considered simultaneously (non-myopic model). We generated 30 data sets for each gate and d of interest. Further, c was set to 128 in all synthetic experiments. To evaluate the relative performance of our method, we compare it with Logistic Regression (LR), Random Forest (RF) [35], TabNet [30] and an MLP classifier. We consider two variants of our method: one with the ReLU activation function and another with the proposed POLYCLIP variant.

4.2. Empirical evaluation on real-life data

We proceed by an overview of the HorNets performance on real-life biomedical data. The motivation for this is twofold: First, biomedical data features an abundant number of variables but a scarce number of examples, which leads to problematic behavior when MLPs are applied directly. Second, neural network-based approaches such as TabNET and MLPs are hardly interpretable, posing

challenges for the real-life application of machine learning models in medicine. The data sets [36] considered represent different classification problems that aim to associate gene expression signals with the target output (e.g. tumor presence). A summary of data sets is shown in Table 1.

Dataset	Instances	Features	Classes	C1	C2	C-remainder
Multi-B ^[37]	32	5565	4	34.38	28.12	37.50
Breast-B ^[37]	49	1213	4	38.78	24.49	36.73
DLBCL-C ^[37]	58	3795	4	29.31	27.59	43.10
Breast-A ^[37]	98	1213	3	52.04	36.73	11.22
Prostate-GE ^[38]	102	5966	2	50.98	49.02	-
Multi-A ^[37]	103	5565	4	27.18	25.24	47.58
CLL ^[39]	111	11340	3	45.95	44.14	9.91
DLBCL-D ^[37]	129	3795	4	37.98	28.68	33.33
DLBCL-A ^[37]	141	661	3	35.46	34.75	29.79
TOX ^[40]	171	5748	4	26.32	26.32	47.36
DLBCL-B ^[37]	180	661	3	48.33	28.33	23.33
SMK ^[41]	187	19993	2	51.87	48.13	-
Lung ^[42]	203	3312	5	68.47	10.34	21.18
TCGA ^[43]	801	20531	5	37.45	18.23	44.32

Table 1. Dataset summary. For each dataset, we report the number of instances, features, classes, and class distribution (Class 1, Class 2, and combined distribution of remaining classes).

Motivated by the work of^[44], we aim to compare our method against various classifiers, including linear, ensemble, neural networks, and AutoML classifiers. Specifically, we have selected the following classifiers:

- **Linear:** Decision Tree (DT), Logistic Regression (Lasso (L1) and Ridge (L2)), and Support Vector Machine (SVM)^[45];
- **Ensemble:** Random Forest (RF)^[46] and XGBoost (XGB)^[47];
- **Neural:** Multi-Layer Perceptron (MLP), TabNET^[39], and LassoNet^[48];
- **AutoML:** TPOT^[35] – an Automated Machine Learning tool that optimizes machine learning pipelines using genetic programming.

All models, except the MLP model, were utilized from their official libraries (scikit-learn for linear and ensemble classifiers, PyTorch-TabNet for TabNET, and TPOT for TPOT). The MLP classifier was implemented using the PyTorch suite, consisting of an input layer followed by a dense layer of 32 neurons, preceding the final classification layer, optimized with the Adam optimizer for up to 100 epochs with early stopping after 10 epochs. We use the macro F1-score to compare classifiers, conducting 5-fold stratified cross-validation across five different seeds, predicting 25 times for a given dataset to assess the variability of results. We search through the following hyperparameters, reporting the best-performing per dataset: Activation: ReLU and our proposed PolyClip; Order: 4, 8, 16, 32, 64, 128, 512, 1024, 2048, 4096; Number of Rules: 4, 8, 16, 32, 64, 128, 512, 1024, 2048, 4096; Learning Rate: 0.001, 0.01, 0.1; Batch Size: fixed at 15.

5. Results

This section presents the experimental results of the synthetic and the real biomedical experiments, followed by the analysis of the execution time and impact of hyperparameters on the method.

5.1. Synthetic results

Table 2 presents the performance of our method on synthetic logic modeling datasets. Using the POLYCLIP variant, our method achieves a perfect score, demonstrating its ability to model arbitrary Horn clauses. Strong competitors such as Logistic Regression, Random Forest, and MLP could model logical operations like AND, NOT, OR, and XNOR when the number of variables was less than 8. All models, with the exclusion of the Random Forest, failed to model the XOR operator. Interestingly, the attention-based TabNET method failed to successfully model any interactions regardless of the number of variables across all problems. High-dimensional settings posed challenges for most methods, except for the proposed HorNet. These results highlight the ability of HorNets with the POLYCLIP variant to model arbitrary logical clauses effectively.

Problem	TabNet	Random Forest	MLP	Logistic Regression	HorNets- ReLU	HorNets- POLYCLIP
and(dim=3)	0.296 ± 0.06	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
and(dim=4)	0.204 ± 0.136	1.0 ± 0.0	0.905 ± 0.134	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
and(dim=16)	0.512 ± 0.137	0.981 ± 0.06	0.973 ± 0.063	0.981 ± 0.06	0.935 ± 0.178	1.0 ± 0.0
and(dim=32)	0.331 ± 0.147	0.901 ± 0.122	0.899 ± 0.131	0.912 ± 0.107	0.821 ± 0.21	1.0 ± 0.0
and(dim=8)	0.465 ± 0.089	0.978 ± 0.069	0.981 ± 0.06	1.0 ± 0.0	0.941 ± 0.187	1.0 ± 0.0
and(dim=64)	0.584 ± 0.151	0.783 ± 0.205	0.779 ± 0.21	0.805 ± 0.191	0.95 ± 0.083	1.0 ± 0.0
and(dim=128)	0.463 ± 0.085	0.484 ± 0.093	0.655 ± 0.178	0.627 ± 0.185	0.811 ± 0.209	1.0 ± 0.0
not(dim=3)	0.446 ± 0.023	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
not(dim=4)	0.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
not(dim=8)	0.629 ± 0.256	0.948 ± 0.164	0.948 ± 0.164	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
not(dim=16)	0.418 ± 0.036	1.0 ± 0.0	0.896 ± 0.219	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
not(dim=32)	0.146 ± 0.094	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
not(dim=64)	0.435 ± 0.041	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
not(dim=128)	0.521 ± 0.169	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
or(dim=3)	0.255 ± 0.135	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
or(dim=4)	0.574 ± 0.211	1.0 ± 0.0	0.961 ± 0.125	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
or(dim=8)	0.265 ± 0.116	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0

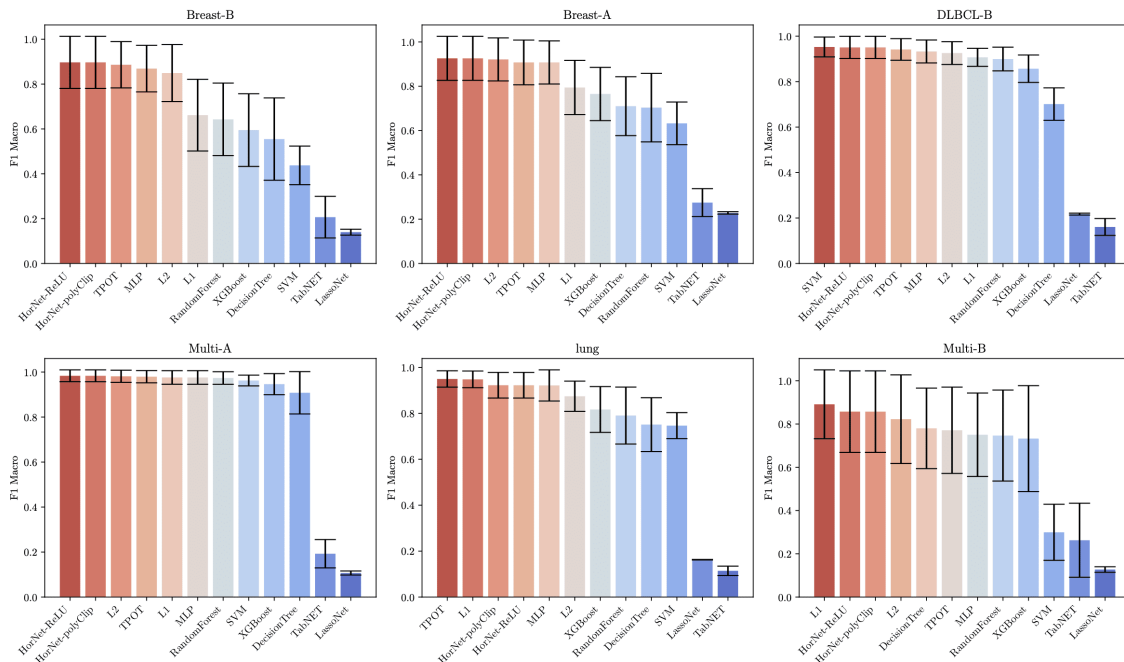
Problem	TabNet	Random Forest	MLP	Logistic Regression	HorNets- ReLU	HorNets- POLYCLIP
or(dim=16)	0.399 ± 0.101	0.988 ± 0.039	0.99 ± 0.031	0.988 ± 0.039	1.0 ± 0.0	1.0 ± 0.0
or(dim=32)	0.48 ± 0.175	0.91 ± 0.123	0.87 ± 0.168	0.92 ± 0.078	1.0 ± 0.0	1.0 ± 0.0
or(dim=64)	0.34 ± 0.081	0.741 ± 0.201	0.73 ± 0.197	0.748 ± 0.215	1.0 ± 0.0	1.0 ± 0.0
or(dim=128)	0.285 ± 0.073	0.558 ± 0.244	0.764 ± 0.199	0.699 ± 0.227	1.0 ± 0.0	1.0 ± 0.0
xnor(dim=3)	0.547 ± 0.24	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
xnor(dim=4)	0.093 ± 0.067	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
xnor(dim=8)	0.729 ± 0.286	1.0 ± 0.0	0.948 ± 0.164	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
xnor(dim=16)	0.4 ± 0.047	1.0 ± 0.0	0.948 ± 0.164	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
xnor(dim=32)	0.19 ± 0.086	1.0 ± 0.0	0.948 ± 0.164	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
xnor(dim=64)	0.543 ± 0.244	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
xnor(dim=128)	0.623 ± 0.261	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
xor(dim=3)	0.483 ± 0.126	1.0 ± 0.0	0.688 ± 0.141	0.535 ± 0.211	1.0 ± 0.0	1.0 ± 0.0
xor(dim=4)	0.511 ± 0.18	1.0 ± 0.0	0.747 ± 0.15	0.497 ± 0.189	1.0 ± 0.0	1.0 ± 0.0
xor(dim=8)	0.326 ± 0.12	0.975 ± 0.041	0.992 ± 0.025	0.459 ± 0.106	0.984 ± 0.033	1.0 ± 0.0
xor(dim=16)	0.42 ± 0.104	0.842 ± 0.099	0.913 ± 0.077	0.465 ± 0.169	0.975 ± 0.081	1.0 ± 0.0
xor(dim=32)	0.436 ± 0.125	0.678 ± 0.158	0.647 ± 0.237	0.554 ± 0.188	0.911 ± 0.172	1.0 ± 0.0
xor(dim=64)	0.471 ± 0.09	0.451 ± 0.081	0.522 ± 0.162	0.476 ± 0.146	0.908 ± 0.184	1.0 ± 0.0

Problem	TabNet	Random Forest	MLP	Logistic Regression	HorNets- ReLU	HorNets- POLYCLIP
xor(dim=128)	0.45 ± 0.097	0.473 ± 0.146	0.505 ± 0.143	0.584 ± 0.107	0.966 ± 0.082	1.0 ± 0.0

Table 2. Overview of F1 performance of different algorithms on a collection of synthetic data sets (30 repetitions, mean and deviation reported).

5.2. Real data sets

We show the results on the real biomedical experiments runs in Figure 3.



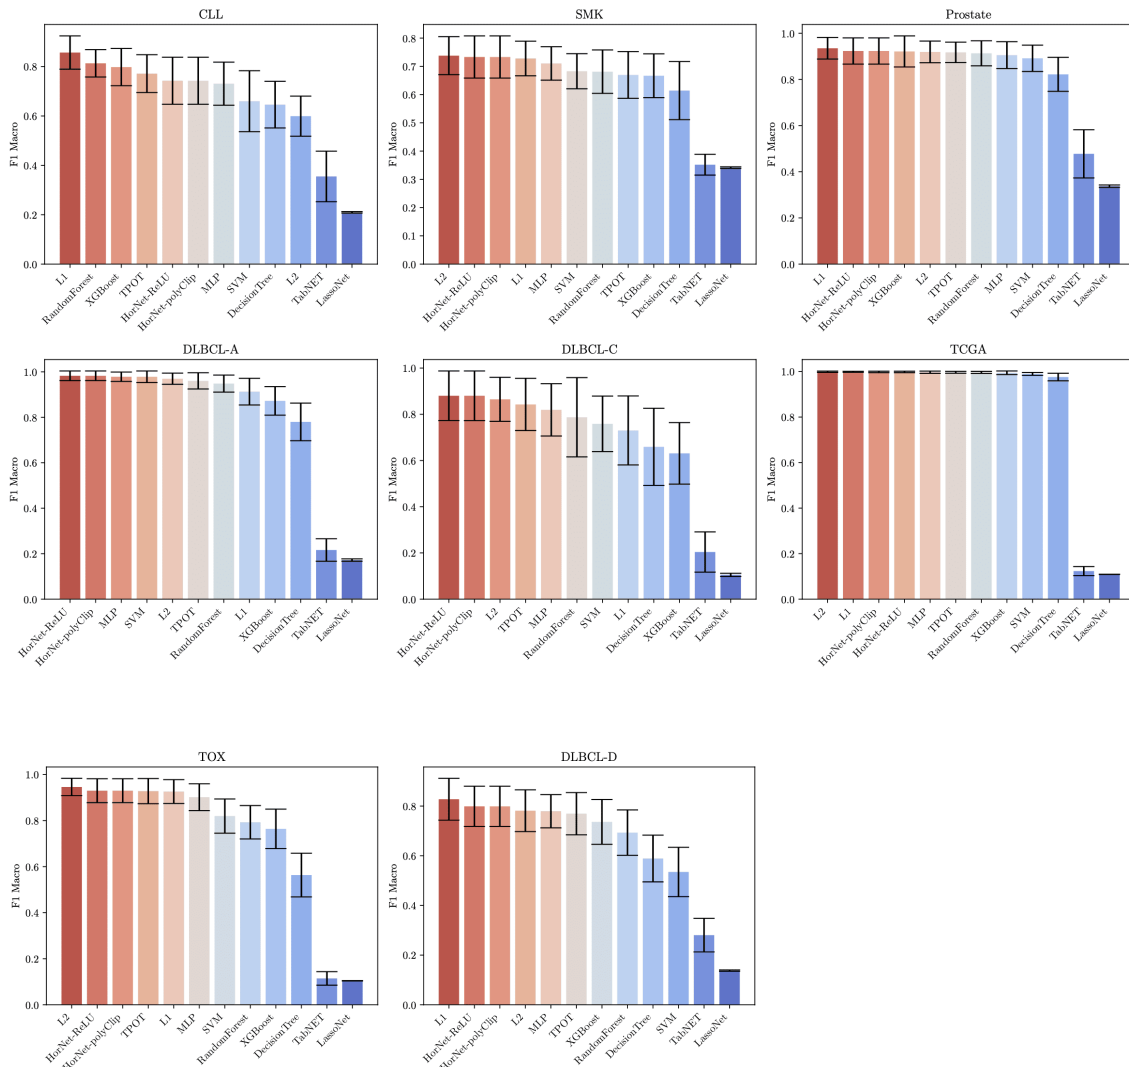


Figure 3. Overview of algorithm performance (14 data sets). It can be observed that HorNets consistently rank among the top three performers, indicating the capability of this algorithm to operate on different problems consistently. Its main competitor in terms of tabular neural network learning, TabNet, failed to achieve high performance – it is hypothesized this is due to high dimensionality/multiple classes of data considered.

We notice that both variants of the HorNet – the one based on the POLYCLIP activation and the one based on the ReLU activation surpass the performance on 5 out of 14 datasets, surpassing the performance of the state-of-the-art models like the TPOT and the L1 and L2 models.

We further investigate the statistical difference between these models by conducting Friedman-Nemenyi test^[49] in Figure 4. The results show that both variants perform similarly, outperforming

the remaining models and performing within critical distance to the best-ranked ones – L1, L2, TPOT and MLP models.




Figure 4. Friedman-Nemenyi test demonstrating the competitive performance of our method’s F1-macro metric.

The Bayesian Hierarchy test^[50] between our model and the TPOT model shown in Figure 5 also shows that the model is significantly better with a probability of 65%, while the TPOT model is better with a probability of 35% (within a region-of-practical-equivalence of 0.01).

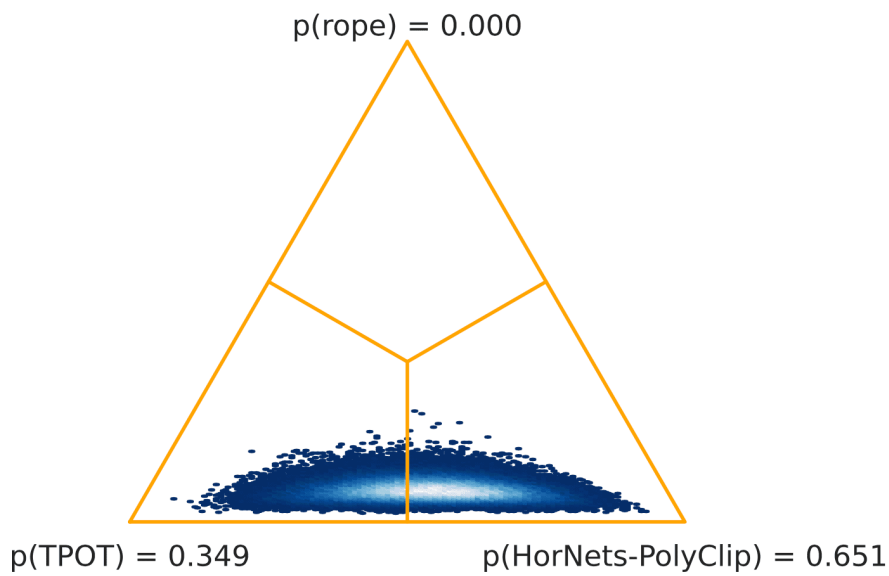


Figure 5. Bayesian hierarchical t-test assessing differences between the HorNets-POLYCLIP variant and the TPOT AutoML model. The test indicates insignificant differences between HorNets-POLYCLIP and the state-of-the-art AutoML approach. The proposed approach is even marginally better.

5.3. Execution time

Next, we analyze the execution time of the methods shown in Figure 6. Our results show that both the POLYCLIP and ReLU variants exhibit indistinguishable performance. The proposed method significantly outperforms advanced neural baselines such as LassoNet and TPOT (which exhaustively search for solutions). It also outperforms ensemble methods such as RandomFrost and XGBoost. Furthermore, we find that the execution time of all methods increases with the size of the dataset, both in terms of the number of features and the number of samples.

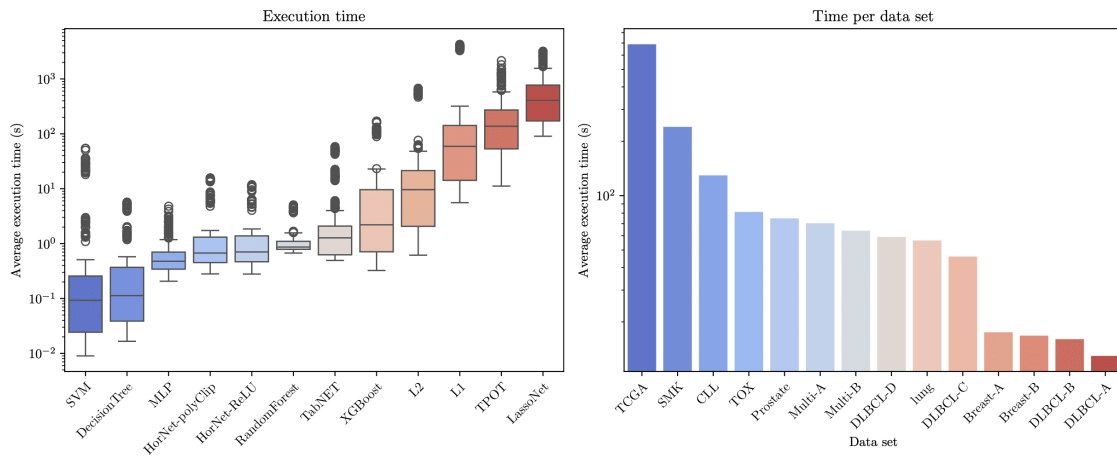


Figure 6. Execution time comparison of our method against the baselines.

In Figure 7 we analyze the effects of the learning rates and the different number of parameters on the execution time. The empirical results confirm the theoretical bounds on the execution time, with the order parameter increasing exponentially and the number of rules increasing linearly.


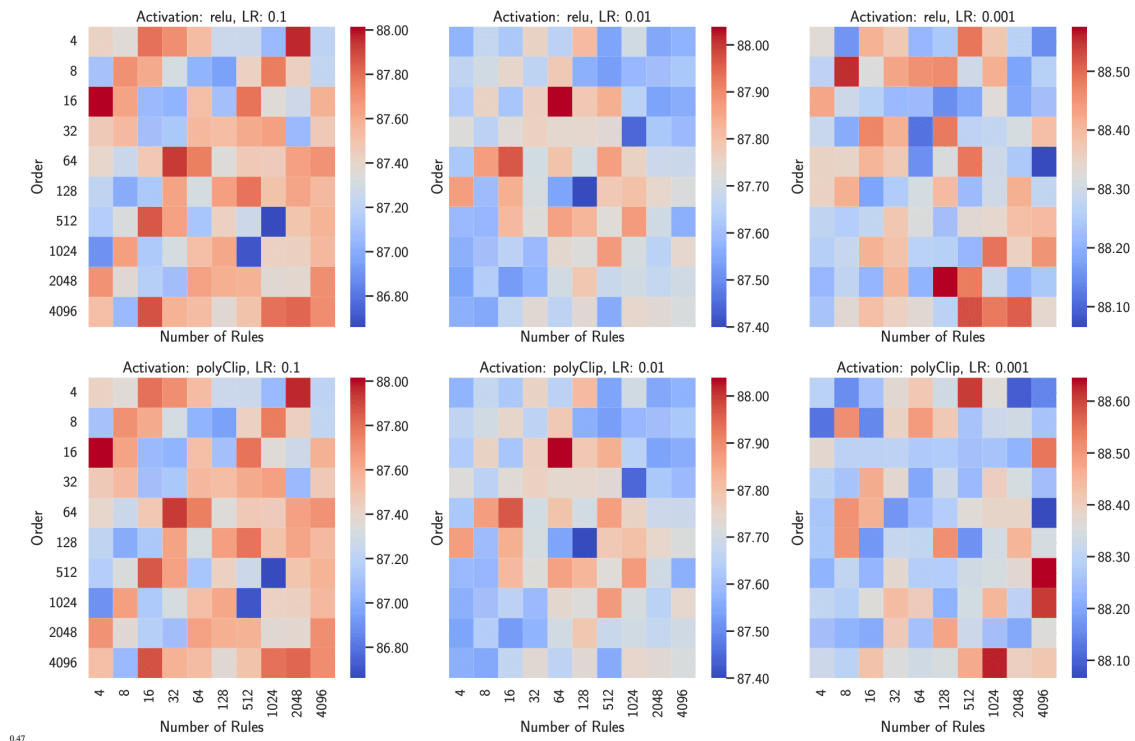


Figure 7. Effect of different hyper-parameters on execution time (seconds) of our method.

5.4. Impact of different hyperparameters

Next, we analyze the impact of different hyperparameters. Recall, we use the following grid to search and evaluate for hyperparameters Activation: ReLU and our proposed POLYCLIP; Order: 4, 8, 16, 32, 64, 128, 512, 1024, 2048, 4096; Number of Rules: 4, 8, 16, 32, 64, 128, 512, 1024, 2048, 4096; Learning Rate: 0.001, 0.01, 0.1; Batch Size: fixed at 15. For each tuple of parameters, we train on all datasets in the described 5 run 5-fold cross-validation setting. Figure 8 shows the optimization landscape across different hyperparameter settings, showing the interaction of different activation functions, the learning rates with the order and number of rules. The results show that no clear best setting can be found and that a theoretical parameter sweep is required when trying to optimize for the best results. Note that, however, the difference between the best-performing (88.60%) and worst-performing parameter setting (86.10%) on average is 1.98 percentage points, which we hypothesize.



0.47

Figure 8. Interaction of different order and number of rules at different learning rates. Results indicate learning rate is the key parameter that has to be tuned – if an appropriate value is selected, relatively low order/rule count are sufficient for good performance.

However, we find that there is no significant variance in the results between the different datasets. In Table 3 we present the top-5 performing parameter settings. The results indicate a negligible difference in performance across these top-5 configurations, despite variations in the combination of order and rules. As expected, a smaller learning rate consistently yielded the highest F1-scores. Table 4 showcases that a similar pattern is observed in the analysis of the top-5 worst-performing parameter settings. Here too, the differences between these combinations are minimal. Notably, the higher learning rates were associated with poorer results, reinforcing the notion that a lower learning rate is preferable for achieving better performance. Notably, the variability within datasets is consistently about 12.5 percentage points regardless of the parameter, pointing at the stability of the method.

Activation	Learning Rate	Order	Rules	F1-score
POLYCLIP	0.001	512	4096	88.65 ± 12.40
POLYCLIP	0.001	4096	1024	88.63 ± 12.31
POLYCLIP	0.001	1024	4096	88.61 ± 12.36
POLYCLIP	0.001	4	512	88.61 ± 12.40
ReLU	0.001	2048	128	88.58 ± 12.44

Table 3. Top 5 Best - Performing Parameters.

Activation	Learning Rate	Order	Rules	F1-score
POLYCLIP	0.10	1024	4	86.90 ± 12.72
ReLU	0.10	1024	512	86.70 ± 13.07
POLYCLIP	0.10	1024	512	86.70 ± 13.07
POLYCLIP	0.10	512	1024	86.66 ± 13.19
ReLU	0.10	512	1024	86.66 ± 13.19

Table 4. Bottom 5 Worst - Performing Parameters.

Finally, we present the sensitivity analysis, aggregated by the effects of “Order” and “Rules” across different datasets, as shown in Figure 9. Despite significant differences in the scales of “Order” and “Rules” (ranging from 4 to 4096), the variations in performance are minor. This suggests that the method exhibits a high degree of stability within the datasets.




Figure 9. Interaction of different order and number of rules at different learning rates.

6. Discussion and conclusions

In this section, we discuss the potential impact of the proposed HorNets based on their performance in synthetic and real-world problems. First, we review the results of the synthetic datasets. This evaluation aimed to empirically test the functional completeness of the proposed HorNets POLYCLIP capability. The results in Table 2 show that the network can learn any interaction between any number of variables for four neural logic operators: AND, NOT, OR, XNOR, and XOR. In contrast, the newer variants of neural networks – TabNET and MLP – as well as the naive rule learning approaches – the decision trees and random forests – were unable to learn. Since we can learn these combinations of features, we believe that the routing functionality means that we can learn arbitrary combinations of features that interact with each other. This, in turn, shows the potential for using the HorNets applications in many neurosymbolic systems that Garcez et al.^[51] see as the next generation of AI. On the other hand,^[52] emphasizes the problem of linking neurosymbolic learners and considers the production of rules about false relations as linking inferences. We show that the HorNets achieve maximum log-likelihood on both training and test data and thus successfully reconstruct the concept of ground truth – in the case of logical operations.

Using the results on the real biomedical datasets, we establish the effectiveness of the learned rules for subsequent application to real classification problems. The statistical tests confirm that the method can keep up with the strong AutoML baseline (TPOT), while generating interpretable rules

that experts can use as a basis for further analyses. This experiment shows that our model can a) train neural networks on small, large tabular datasets,^[44] and b) achieve interpretable results, something that related work highlights as necessary^{[53][54][55]} Furthermore, we show that the proposed HorNets are both data- and resource-intensive, making them ideal for applications related to computing on the edge. We presented a novel approach for learning rules using neural networks. Our method learns from both continuous and discrete signals, as evidenced by rigorous evaluations on both synthetic and real biomedical datasets. The architecture excels at obtaining arbitrary Horn clauses by combining a custom polynomial-like activation function with an attention mechanism that allows it to model complex interactions. This method outperforms advanced AutoML-based learners and is well suited for numerical and categorical data with large table density, while requiring significantly less computational resources than comparable architectures. Our results show that deep neural networks augmented with attention mechanisms and user-defined activation functions are very effective in learning rules. We plan to improve our framework with uncertainty-based principles to enable the extraction of probabilistic rules. We also plan to extend our approach to explore spaces with mixed inputs. To further improve the interpretability of the extracted rules, we will perform in-depth analyses with domain experts.

Appendix: Proof of the Exponentiality of the exact CatInt interactions Using Stirling's Approximation

Stirling's approximation states that for large n :

$$n! \approx \sqrt{2\pi n} \left(\frac{n}{e}\right)^n$$

Given the complexity expression for the combinatorial part of the CatRouter $\mathcal{O}(\text{Comb}(|I|, \text{order}))$, let us consider the combination $\text{Comb}(|I|, \text{order})$ which represents the binomial coefficient:

$$\text{Comb}(|I|, \text{order}) = \binom{|I|}{\text{order}} = \frac{|I|!}{\text{order}! \cdot (|I| - \text{order})!}$$

Using Sterling's approximation for factorials, we have:

$$|I|! \approx \sqrt{2\pi |I|} \left(\frac{|I|}{e}\right)^{|I|}$$

$$\text{order}! \approx \sqrt{2\pi \text{order}} \left(\frac{\text{order}}{e}\right)^{\text{order}}$$

$$(|I| - \text{order})! \approx \sqrt{2\pi(|I| - \text{order})} \left(\frac{|I| - \text{order}}{e} \right)^{|I| - \text{order}}$$

Substituting these into the binomial coefficient formula, we get:

$$\binom{|I|}{\text{order}} \approx \frac{\sqrt{2\pi|I|} \left(\frac{|I|}{e} \right)^{|I|}}{\sqrt{2\pi \text{order}} \left(\frac{\text{order}}{e} \right)^{\text{order}} \cdot \sqrt{2\pi(|I| - \text{order})} \left(\frac{|I| - \text{order}}{e} \right)^{|I| - \text{order}}}$$

$$\binom{|I|}{\text{order}} \approx \frac{\sqrt{2\pi|I|} \left(\frac{|I|}{e} \right)^{|I|}}{\sqrt{2\pi \text{order}} \sqrt{2\pi(|I| - \text{order})} \cdot \left(\frac{\text{order}}{e} \right)^{\text{order}} \cdot \left(\frac{|I| - \text{order}}{e} \right)^{|I| - \text{order}}}$$

$$\binom{|I|}{\text{order}} \approx \frac{\sqrt{2\pi|I|}}{\sqrt{2\pi \text{order}} \sqrt{2\pi(|I| - \text{order})}} \cdot \frac{\left(\frac{|I|}{e} \right)^{|I|}}{\left(\frac{\text{order}}{e} \right)^{\text{order}} \cdot \left(\frac{|I| - \text{order}}{e} \right)^{|I| - \text{order}}}$$

We can simplify the first term as follows:

$$\frac{\sqrt{2\pi|I|}}{\sqrt{2\pi \text{order}} \cdot \sqrt{2\pi(|I| - \text{order})}} = \frac{\sqrt{2\pi|I|}}{\sqrt{(2\pi)^2 \cdot \text{order} \cdot (|I| - \text{order})}} = \frac{\sqrt{2\pi|I|}}{2\pi \sqrt{\text{order} \cdot (|I| - \text{order})}}$$

Next, the exponential term simplifies as follows:

$$\frac{\left(\frac{|I|}{e} \right)^{|I|}}{\left(\frac{\text{order}}{e} \right)^{\text{order}} \cdot \left(\frac{|I| - \text{order}}{e} \right)^{|I| - \text{order}}} = \frac{|I|^{|I|} \cdot e^{-|I|}}{\text{order}^{\text{order}} \cdot e^{-\text{order}} \cdot (|I| - \text{order})^{|I| - \text{order}} \cdot e^{-(|I| - \text{order})}}$$

Upon canceling the exponential terms $e^{-|I|}$, the expression reduces to:

$$\frac{|I|^{|I|}}{\text{order}^{\text{order}} \cdot (|I| - \text{order})^{|I| - \text{order}}}$$

The simplified combined expression is:

$$\binom{|I|}{\text{order}} \approx \frac{\sqrt{2\pi|I|}}{2\pi \sqrt{\text{order} \cdot (|I| - \text{order})}} \cdot \frac{|I|^{|I|}}{\text{order}^{\text{order}} \cdot (|I| - \text{order})^{|I| - \text{order}}}$$

Since we are concerned with the asymptotic behavior, we can safely ignore the first term and focus on the latter, thereby demonstrating that the number of combinations in the CatInt grows exponentially.

$$\mathcal{O}(\text{Comb}(|I|, \text{order})) \approx \frac{|I|^{|I|}}{\text{order}^{\text{order}} \cdot (|I| - \text{order})^{|I| - \text{order}}}$$

Statements and Declarations

Availability and Reproducibility

The code for the HorNets is freely available via the following GitHub link: <https://github.com/bkoloski/hornets>. To ensure reproducibility and replicability, we provide a Singularity container environment, which integrates the code and provides a benchmarking environment.

Acknowledgements

The authors acknowledge the financial support from the Slovenian Research and Innovation Agency through research the funding schemes P2-0103 and PR-12394.

Footnotes

¹ PolyClip – an acronym for clipped polynomials.

² Note that there exists a scenario where -1, 1 and 0 are outputs of polyClip, enabling direct logical interpretation. In practice, such behavior can be forced with rounding/discretization.

³ Computing cardinality of the types of real-life data sets considered is space-wise negligible – should exact cardinality computation become a memory bottleneck, probabilistic counting can be employed.

⁴ Note that one-hot encoding can be used to obtain a binary space from arbitrary categorical inputs.

⁵ For high-dimensional real-life data sets, even third-order interactions can be problematic.

⁶ Note that sampling higher order interactions exhaustively is computationally intractable, even for moderately sized data sets (e.g., hundred features).

⁷ $\mathcal{O}(\text{Comb}(|I|, \text{order}))$, this can be shown via Sterling's approximation – the proof can be found in the Appendix.

References

1. [△]Dong S, Wang P, Abbas K (2021). "A survey on deep learning and its applications". *Computer Science Review*. 40: 100379.

2. ^a ^bBader S, Hf6lldobler S, Mayer-Eichberger V (2007). "Extracting propositional rules from feed-forward neural networks: a new decompositional approach". In: *Proceedings of the 3rd International Conference on Neural-Symbolic Learning and Reasoning - Volume 230. NeSy'07, Aachen, DEU: CEUR-WS.org, pp. 4-9.*
3. [^]Al Iqbal R. "Eclectic extraction of propositional rules from neural networks". In: *14th International Conference on Computer and Information Technology (ICCIIT 2011), IEEE, 2011, pp. 234-239.*
4. [^]Labaf M, Hitzler P, Evans AB (2017). "Propositional rule extraction from neural networks under background knowledge." In: *NeSy.*
5. ^a ^bAmizadeh S, Palangi H, Polozov A, Huang Y, Koishida K (2020). "Neuro-symbolic visual reasoning: Disentangling 'visual' from 'reasoning'." In: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event. Proceedings of Machine Learning Research, 119: 279-290.*
6. [^]Manhaeve R, Dumancic S, Kimmig A, Demeester T, Raedt LD. "DeepProbLog: Neural Probabilistic Logic Programming". In: *Bengio S, Wallach HM, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R, editors. Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada. 2018. p. 3753-3763. Available from: <https://dblp.org/rec/conf/nips/ManhaeveDKDR18.bib>.*
7. [^]Towell GG, Shavlik JW (1994). "Knowledge-based artificial neural networks". *Artificial intelligence. 70 (1-2): 119-165.*
8. [^]Rastegari M, Ordonez V, Redmon J, Farhadi A (2016). "Xnor-net: Imagenet classification using binary convolutional neural networks". In: *European Conference on Computer Vision, pp. 525-542. Springer.*
9. [^]Courbariaux M, Bengio Y, David J (2015). "BinaryConnect: Training Deep Neural Networks with binary weights during propagations". In: *Cortes C, Lawrence ND, Lee DD, Sugiyama M, Garnett R, editors. Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada, pp. 3123-3131.*
10. [^]Moss DJ, Nurvitadhi E, Sim J, Mishra A, Marr D, Subhaschandra S, Leong PH. "High performance binary neural networks on the Xeon + FPGA platform". In: *2017 27th International Conference on Field Programmable Logic and Applications (FPL). IEEE; 2017. p. 1-4.*
11. [^]Sun X, Yin S, Peng X, Liu R, Seo J-s, Yu S (2018). "XNOR-RRAM: A scalable and parallel resistive synaptic architecture for binary neural networks". In: *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE), pp. 1423-1428. IEEE.*

12. [^]Lodhi H. Deep relational machines. In: *International Conference on Neural Information Processing*. Springer; 2013. p. 212–219.
13. [^]Dash T, Srinivasan A, Vig L, Orhobor OI, King RD. "Large-scale assessment of deep relational machines". In: *International Conference on Inductive Logic Programming*. Springer; 2018. p. 22–37.
14. ^{a, b}Manhaeve R, Dumancic S, Kimmig A, Demeester T, De Raedt L. "DeepProbLog: Neural Probabilistic Logic Programming". In: Bengio S, Wallach H, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R, editors. *Advances in Neural Information Processing Systems*. 2018; 31. Available from: https://proceedings.nips.cc/paper_files/paper/2018/file/dc5d637ed5e62c36ecb73b654b05ba2a-Paper.pdf.
15. [^]Lavrač N, Kralj B, Robnik-Šikonja M (2020). "Propositionalization and embeddings: two sides of the same coin". *Machine Learning*. 109 (7): 1465–1507. doi:[10.1007/s10994-020-05890-8](https://doi.org/10.1007/s10994-020-05890-8).
16. ^{a, b}Cropper A, Dumančić S (2022). "Inductive logic programming at 30: a new introduction". *Journal of Artificial Intelligence Research*. 74: 765–850.
17. [^]Li Q, Huang S, Hong Y, Chen Y, Wu YN, Zhu S (2020). "Closed loop neural-symbolic learning via integrating neural perception, grammar parsing, and symbolic reasoning". In: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13–18 July 2020, Virtual Event. Proceedings of Machine Learning Research*. 119: 5884–5894. Available from: <https://dblp.org/rec/conf/icml/LiHHCWZ20.bib>.
18. [^]Dong H, Mao J, Lin T, Wang C, Li L, Zhou D (2019). "Neural logic machines". In: *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019*. Available from: <https://dblp.org/rec/conf/iclr/DongMLWLZ19.bib>.
19. [^]Škrlj B, Kralj J, Lavrač N (2020). "Embedding-based Silhouette community detection". *Machine Learning*. 109(11): 2161–2193. doi:[10.1007/s10994-020-05882-8](https://doi.org/10.1007/s10994-020-05882-8).
20. ^{a, b}Beck F, Frcnkranz J (2021). "An empirical investigation into deep and shallow rule learning". *Frontiers in Artificial Intelligence*. 4: 689398.
21. [^]Qiao L, Wang W, Lin B (2021). "Learning accurate and interpretable decision rule sets from neural networks". In: *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2–9, 2021*, pp. 4303–4311. Available from: <https://dblp.org/rec/conf/aaai/QiaoWL21.bib>.

22. [△]Shi S, Xie Y, Wang Z, Ding B, Li Y, Zhang M (2022). "Explainable neural rule learning". In: *Proceedings of the ACM Web Conference 2022*. WWW '22, pp. 3031–3041. Association for Computing Machinery, New York, NY, USA. doi:[10.1145/3485447.3512023](https://doi.org/10.1145/3485447.3512023).
23. [△]Dierckx L, Veroneze R, Nijssen S (2023). "RL-Net: Interpretable Rule Learning with Neural Networks". In: *Advances in Knowledge Discovery and Data Mining: 27th Pacific-Asia Conference on Knowledge Discovery and Data Mining, PAKDD 2023, Osaka, Japan, May 25–28, 2023, Proceedings, Part I*. Berlin, Heidelberg: Springer. pp. 95–107. doi:[10.1007/978-3-031-33374-3_8](https://doi.org/10.1007/978-3-031-33374-3_8).
24. [△]Pogancic MV, Paulus A, Musil V, Martius G, Rolinek M (2020). "Differentiation of blackbox combinatorial solvers". In: *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26–30, 2020*. Available from: <https://dblp.org/rec/conf/iclr/PogancicPMMR20.bib>.
25. [△]Dumancic S, Guns T, Meert W, Blockeel H (2019). "Learning relational representations with auto-encoding logic programs". In: Kraus S, editor. *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10–16, 2019*, pp. 6081–6087. doi:[10.24963/ijcai.2019/842](https://doi.org/10.24963/ijcai.2019/842).
26. [△]Ellis K, Wong C, Nye M, Sablé-Meyer M, Morales L, Hewitt L, Cary L, Solar-Lezama A, Tenenbaum JB (2021). "Dreamcoder: bootstrapping inductive program synthesis with wake-sleep library learning". In: *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation. PLDI 2021*, pp. 835–850. Association for Computing Machinery, New York, NY, USA. doi:[10.1145/3453483.3454080](https://doi.org/10.1145/3453483.3454080).
27. [△]Ribeiro MT, Singh S, Guestrin C (2016). "Why should I trust you?": Explaining the predictions of any classifier. In: Krishnapuram B, Shah M, Smola AJ, Aggarwal CC, Shen D, Rastogi R, editors. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13–17, 2016*, pp. 1135–1144. doi:[10.1145/2939672.2939778](https://doi.org/10.1145/2939672.2939778).
28. [△]Lundberg SM, Lee S. A unified approach to interpreting model predictions. In: Guyon I, Luxburg U, Bengio S, Wallach HM, Fergus R, Vishwanathan SVN, Garnett R, editors. *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4–9, 2017, Long Beach, CA, USA*. 2017. p. 4765–4774.
29. [△]Slack D, Hilgard S, Jia E, Singh S, Lakkaraju H (2020). "Fooling LIME and SHAP: Adversarial attacks on post hoc explanation methods". In: *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pp. 180–186.

30. ^{a, b, c}Arik SÖ, Pfister T (2021). "TabNet: Attentive Interpretable Tabular Learning". In: Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2–9, 2021, pp. 6679–6687. <https://dblp.org/rec/conf/aaai/ArikP21.bib>.
31. [^]Škrlić B, Džeroski S, Lavrač N, Petković M (2020). "Feature importance estimation with self-attention networks". In: Proceedings of the European Conference on Artificial Intelligence, pp. 1491–1498.
32. [^]Liu Z, Wang Y, Vaidya S, Ruehle F, Halverson J, Soljačić M, Hou TY, Tegmark M (2024). "Kan: Kolmogorov-arnold networks". ArXiv preprint. [abs/2404.19756](https://arxiv.org/abs/2404.19756). Available from: <https://arxiv.org/abs/2404.19756>.
33. [^]Cai W, Jiang J, Wang F, Tang J, Kim S, Huang J (2024). "A survey on mixture of experts". arXiv preprint [arXiv:2407.06204](https://arxiv.org/abs/2407.06204). Available from: <https://arxiv.org/abs/2407.06204>.
34. [^]Salakhutdinov R. Deep learning. In: Macskassy SA, Perlich C, Leskovec J, Wang W, Ghani R, editors. The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, New York, NY, USA - August 24 - 27, 2014. 2014. p. 1973. doi:[10.1145/2623330.2630809](https://doi.org/10.1145/2623330.2630809).
35. ^{a, b}Le TT, Fu W, Moore JH (2020). "Scaling tree-based automated machine learning to biomedical big data with a feature set selector". *Bioinformatics*. 36 (1): 250–256.
36. [^]Li J, Cheng K, Wang S, Morstatter F, Trevino RP, Tang J, Liu H (2018). "Feature selection: A data perspective". *ACM Computing Surveys (CSUR)*. 50 (6): 94.
37. ^{a, b, c, d, e, f, g, h}Hoshida Y, Brunet J-P, Tamayo P, Golub TR, Mesirov JP. Subclass mapping: Identifying common subtypes in independent disease data sets. *PLOS ONE*. 2007;2(11):1–8. doi:[10.1371/journal.pone.0001195](https://doi.org/10.1371/journal.pone.0001195).
38. [^]Singh D, Febbo PG, Ross K, Jackson DG, Manola J, Ladd C, Tamayo P, Renshaw AA, D'Amico AV, Richie JP, et al. Gene expression correlates of clinical prostate cancer behavior. *Cancer cell*. 2002; 1(2): 203–209.
39. [^]Haslinger C, Schweifer N, Stilgenbauer S, Dohner H, Lichter P, Kraut N, Stratowa C, Abseher R (2004). "Microarray gene expression profiling of B-cell chronic lymphocytic leukemia subgroups defined by genomic aberrations and VH mutation status". *Journal of Clinical Oncology*. 22 (19): 3937–3949.
40. [^]Bhattacharjee A, Richards WG, Staunton J, Li C, Monti S, Vasa P, Ladd C, Beheshti J, Bueno R, Gillette M, et al. Classification of human lung carcinomas by mRNA expression profiling reveals distinct adenocarcinoma subclasses. *Proceedings of the National Academy of Sciences*. 98(24):13790–13795 (2001).

41. [△]Spira A, Beane JE, Shah V, Steiling K, Liu G, Schembri F, Gilman S, Dumas Y-M, Calner P, Sebastiani P, et al. Airway epithelial gene expression in the diagnostic evaluation of smokers with suspect lung cancer. *Nature Medicine*. 13(3):361–366 (2007).
42. [△]Bajwa G, DeBerardinis RJ, Shao B, Hall B, Farrar JD, Gill MA (2016). "Cutting edge: critical role of glycolysis in human plasmacytoid dendritic cell antiviral responses". *The Journal of Immunology*. 196 (5): 2004–2009.
43. [△]Weinstein JN, Collisson EA, Mills GB, Shaw KRM, Ozenberger BA, Ellrott K, Shmulevich I, Sander C, Stuart JM, Network CGAR, et al. The cancer genome atlas pan-cancer analysis project. *Nature genetics*. 45 (10):1113 (2013).
44. ^{a, b}Grinsztajn L, Oyallon E, Varoquaux G (2022). "Why do tree-based models still outperform deep learning on typical tabular data?" *Advances in neural information processing systems*. 35: 507–520.
45. [△]Cortes C, Vapnik V (1995). "Support-vector networks". *Machine Learning*. 20 (3): 273–297.
46. [△]Breiman L. *Classification and regression trees*. Wadsworth International Group. 2017.
47. [△]Chen T, Guestrin C (2016). "Xgboost: A scalable tree boosting system". In: Krishnapuram B, Shah M, Sasmola AJ, Aggarwal CC, Shen D, Rastogi R, editors. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13–17, 2016*, pp. 785–794. doi:10.1145/2939672.2939785.
48. [△]Lemhadri I, Ruan F, Abraham L, Tibshirani R (2021). "Lassonet: A neural network with feature sparsity". *J. Mach. Learn. Res.*. 22: 127:1–127:29.
49. [△]Demšar J (2006). "Statistical comparisons of classifiers over multiple data sets". *Journal of Machine Learning Research*. 7 (Jan): 1–30.
50. [△]Benavoli A, Corani G, Demšar J, Zaffalon M (2017). "Time for a change: a tutorial for comparing multiple classifiers through bayesian analysis". *Journal of Machine Learning Research*. 18 (77): 1–36.
51. [△]Garcez Ad, Lamb LC (2023). "Neurosymbolic AI: The 3rd wave". *Artificial Intelligence Review*. 56 (11): 12387–12406.
52. [△]Marconato E, Teso S, Vergari A, Passerini A. Not all neuro-symbolic concepts are created equal: Analysis and mitigation of reasoning shortcuts. In: Oh A, Naumann T, Globerson A, Saenko K, Hardt M, Levine S, editors. *Advances in Neural Information Processing Systems*. 2023;36.
53. [△]Malik A, Patel P, Ehsan L, Guleria S, Hartka T, Adewole S, Syed S (2021). "Ten simple rules for engaging with artificial intelligence in biomedicine". *PLoS Computational Biology*. 17 (2): e1008531. Public Library of Science San Francisco, CA USA.

54. [△]Wong F, Collins JJ (2023). "‘Explainable’ AI identifies a new class of antibiotics". *Nature*. 6887: 8.

55. [△]Wong F, Zheng EJ, Valeri JA, Donghia NM, Anahtar MN, Omori S, Li A, Cubillos-Ruiz A, Krishnan A, Jin W, et al. (2024). "Discovery of a structural class of antibiotics with explainable deep learning". *Nature*. 626 (7997): 177–185.

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.