arXiv:2309.12207v1 [cs.LG] 21 Sep 2023

Boolformer:
Symbolic Regression of Logic Functions with Transformers

Stéphane d'Ascoli*[1], Samy Bengio[2], Josh Susskind[2] and Emmanuel Abbé[1,2]

[1]EPFL
[2]Apple

Abstract

In this work, we introduce Boolformer, the first Transformer architecture trained to perform end-to-end symbolic regression of Boolean functions. First, we show that it can predict compact formulas for complex functions which were not seen during training, when provided a clean truth table. Then, we demonstrate its ability to find approximate expressions when provided incomplete and noisy observations. We evaluate the Boolformer on a broad set of real-world binary classification datasets, demonstrating its potential as an interpretable alternative to classic machine learning methods. Finally, we apply it to the widespread task of modelling the dynamics of gene regulatory networks. Using a recent benchmark, we show that Boolformer is competitive with state-of-the art genetic algorithms with a speedup of several orders of magnitude. Our code and models are available publicly.

1 Introduction

Deep neural networks, in particuler those based on the Transformer architecture [1], have lead to breakthroughs in computer vision [2] and language modelling [3], and have fuelled the hopes to accelerate scientific discovery [4]. However, their ability to perform simple logic tasks remains limited [5]. These tasks differ from traditional vision or language tasks in the combinatorial nature of their input space, which makes representative data sampling challenging.

Reasoning tasks have thus gained major attention in the deep learning community, either with explicit reasoning in the logical domain, e.g., tasks in the realm of arithmetic and algebra [6, 7], algorithmic CLRS tasks [8] or LEGO [9], or implicit reasoning in other modalities, e.g., benchmarks such as Pointer Value Retrieval [10] and Clevr [11] for vision models, or LogiQA [12] and GSM8K [13] for language models. Reasoning also plays a key role in tasks which can be tackled via Boolean modelling, particularly in the fields of biology [14] and medecine [15].

As these endeavours remain challenging for current Transformer architectures, it is natural to examine whether they can be handled more effectively with different approaches, e.g., by better exploiting the Boolean nature of the task. In particular, when learning Boolean functions with a 'classic' approach based on minimizing the training loss on the outputs of the function, Transformers learn potentially complex interpolators as they focus on minimizing the degree profile in the Fourier spectrum, which is not the type of bias desirable for generalization on domains that are not well sampled [16]. In turn, the complexity of the learned function makes its interpretability challenging. This raises the question of how to improve generalization and interpretability of such models.

In this paper, we tackle Boolean function learning with Transformers, but we rely directly on 'symbolic regression': our **Boolformer** is tasked to predict a Boolean formula, i.e., a symbolic expression of the Boolean function in terms of the three fundamental logical gates (AND, OR, NOT) such as those

---

*stephane.dascoli@gmail.com

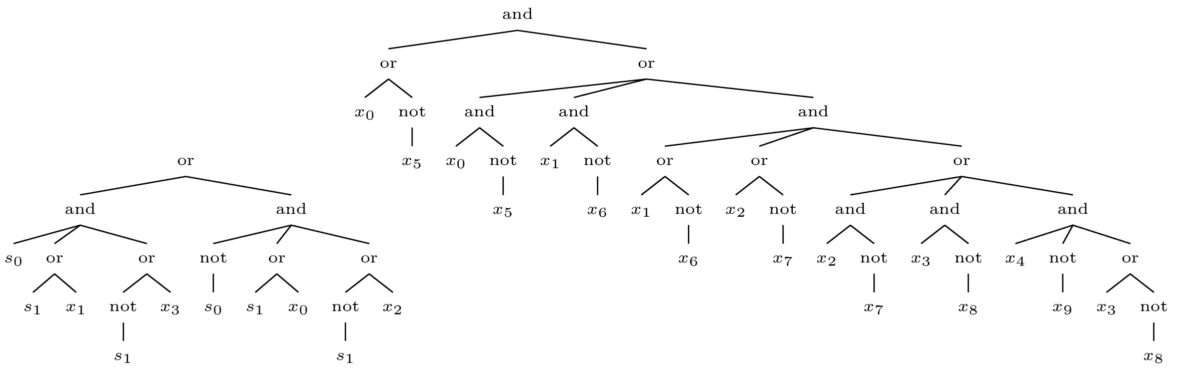
and
or
or
x_0
not
x_5
and
x_0
not
x_5
and
x_1
not
x_6
and
or
x_1
not
x_6
or
x_2
not
x_7
or
and
x_2
not
x_7
and
x_3
not
x_8
and
x_4
not
x_9
or
x_3
not
x_8
or
and
s_0
or
s_1
x_1
or
not
s_1
x_3
and
not
s_0
or
s_1
x_0
or
not
s_1
x_2

(a) Multiplexer (b) Comparator

Figure 1: **Some logical functions for which our model predicts an optimal formula.** Left: the multiplexer, a function commonly used in electronics to select one out of four sources x_0, x_1, x_2, x_3 based on two selector bits s_0, s_1. Right: given two 5-bit numbers $a = (x_0x_1x_2x_3x_4)$ and $b = (x_5x_6x_7x_7x_9)$, returns 1 if $a > b$, 0 otherwise.

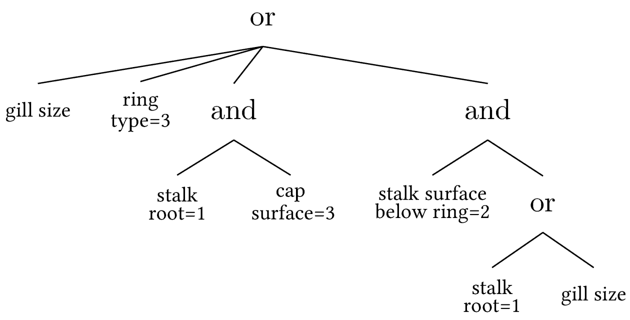
or
gill size
ring type=3
and
stalk root=1
cap surface=3
and
stalk surface below ring=2
or
stalk root=1
gill size

Figure 2: **A Boolean formula predicted to determine whether a mushroom is poisonous.** We considered the "mushroom" dataset from the PMLB database [17], and this formula achieves an F1 score of 0.96.

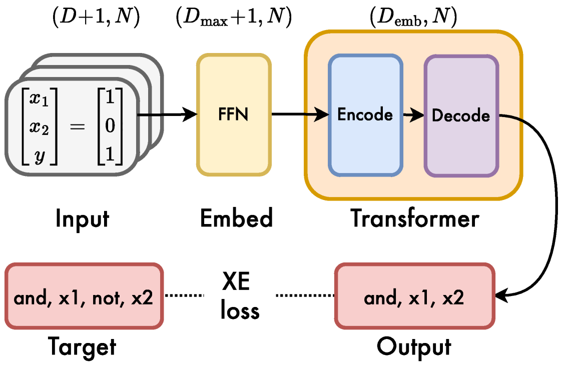
$(D+1, N)$
$(D_{\text{max}}+1, N)$
(D_{emb}, N)
$\begin{bmatrix} x_1 \\ x_2 \\ y \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$
FFN
Encode
Decode
Input
Embed
Transformer
and, x1, not, x2
XE loss
and, x1, x2
Target
Output

Figure 3: **Summary of our approach.** We feed N points $(\boldsymbol{x}, f(\boldsymbol{x})) \in \{0,1\}^{D+1}$ to a seq2seq Transformer, and supervise the prediction to f via cross-entropy loss.

of Figs. 1,2. As illustrated in Fig. 3, this task is framed as a sequence prediction problem: each training example is a synthetically generated function whose truth table is the input and whose formula is the target. By moving to this setting and controlling the data generation process, one can hope to gain both in generalization and interpretability.

We show that this approach can give surprisingly strong performance on various logical tasks both in abstract and real-world settings, and discuss how this lays the ground for future improvements and applications.

1.1 Contributions

1. We train Transformers over synthetic datasets to perform end-to-end symbolic regression of Boolean formulas and show that given the full truth table of an unseen function, the Boolformer is able to predict a compact formula, as illustrated in Fig. 1.
2. We show that Boolformer is robust to noisy and incomplete observations, by providing incomplete truth tables with flipped bits and irrelevant variables.
3. We evaluate Boolformer on various real-world binary classification tasks from the PMLB database [17] and show that it is competitive with classic machine learning approaches such as Random Forests

while providing interpretable predictions, as illustrated in Fig. 2.

4. We apply Boolformer to the well-studied task of modelling gene regulatory networks (GRNs) in biology. Using a recent benchmark, we show that our model is competitive with state-of-the-art methods with several orders of magnitude faster inference.
5. Our code and models are available publicly at the following address: `https://github.com/sdascoli/boolformer`. We also provide a pip package entitled `boolformer` for easy setup and usage.

1.2 Related work

Reasoning in deep learning Several papers have studied the ability of deep neural networks to solve logic tasks. Evans & Grefenstette [18] introduce differential inductive logic as a method to learn logical rules from noisy data, and a few subsequent works attempted to craft dedicated neural architectures to improve this ability [19–21]. Large language models (LLMs) such as ChatGPT, however, have been shown to perform poorly at simple logical tasks such as basic arithmetic [22], and tend to rely on approximations and shortcuts [23]. Although some reasoning abilities seem to emerge with scale [24], achieving holistic and interpretable reasoning in LLMs remains an open challenge.

Boolean learnability Leaning Boolean functions has been an active area in theoretical machine learning, mostly under the probably approximately correct (PAC) and statistical query (SQ) learning frameworks [25, 26]. More recently, Abbe *et al.* [27] shows that regular neural networks learn by gradually fitting monomials of increasing degree, in such a way that the sample complexity is governed by the 'leap complexity' of the target function, i.e. the largest degree jump the Boolean function sees in its Fourier decomposition. In turn, Abbe *et al.* [16] shows that this leads to a 'min-degree bias' limitation: Transformers tend to learn interpolators having least 'degree profile' in the Boolean Fourier basis, which typically lose the Boolean nature of the target and often produce complex solutions with poor out-of-distribution generalization.

Inferring Boolean formulas A few works have explored the paradigm of inferring Boolean formulas from observations using SAT solvers [28], ILP solvers [29, 30] or via LP-relaxation [31]. However, all these works predict the formulas in conjunctive or disjunctive normal forms (CNF/DNF), which typically corresponds to exponentially long formulas. In contrast, the Boolformer is biased towards predicting compact expressions[1], which is more akin to logic synthesis – the task of finding the shortest circuit to express a given function, also known as the Minimum Circuit Size Problem (MCSP). While a few heuristics (e.g. Karnaugh maps [32]) and algorithms (e.g. ESPRESSO [33]) exist to tackle the MCSP, its NP-hardness [34] remains a barrier towards efficient circuit design. Given the high resilience of computers to errors, approximate logic synthesis techniques have been introduced [35–40], with the aim of providing approximate expressions given incomplete data – this is similar in spirit to what we study in the noisy regime of Section 4.

Symbolic regression Symbolic regression (SR), i.e. the search of mathematical expression underlying a set of numerical values, is still today a rather unexplored paradigm in the ML literature. Since this search cannot directly be framed as a differentiable problem, the dominant approach for SR is genetic programming (see [41] for a recent review). A few recent publications applied Transformer-based approaches to SR [42–45], yielding comparable results but with a significant advantage: the inference time rarely exceeds a few seconds, several orders of magnitude faster than existing methods. Indeed, while the latter need to be run from scratch on each new set of observations, Transformers are trained over synthetic datasets, and inference simply consists in a forward pass.

---

[1] Consider for example the comparator of Fig. 1: since the truth table has roughly as many positive and negative outputs, the CNF/DNF involves $\mathcal{O}(2^D)$ terms where D is the number of input variables, which for $D = 10$ amounts to several thousand binary gates, versus 17 for our model.

2 Methods

Our task is to infer Boolean functions of the form $f : \{0,1\}^D \to \{0,1\}$, by predicting a Boolean formula built from the basic logical operators: AND, OR, NOT, as illustrated in Figs. 1,2. We train Transformers [1] on a large dataset of synthetic examples, following the seminal approach of [46]. For each example, the input $\mathcal{D}_{\text{fit}}$ is a set of pairs $\{(\boldsymbol{x}_i, y = f(\boldsymbol{x}_i))\}_{i=1\ldots N}$, and the target is the function f as described above. Our general method is summarized in Fig. 3. Examples are generated by first sampling a random function f, then generating the corresponding $(\boldsymbol{x}, y)$ pairs as described in the sections below.

2.1 Generating functions

We generate random Boolean formulas[2] in the form of random unary-binary trees with mathematical operators as internal nodes and variables as leaves. The procedure is detailed as follows:

1. **Sample the input dimension** D of the function f uniformly in $[1, D_{\text{max}}]$.
2. **Sample the number of active variables** S uniformly in $[1, S_{\text{max}}]$. S determines the number of variables which affect the output of f: the other variables are inactive. Then, select a set of S variables from the original D variables uniformly at random.
3. **Sample the number of binary operators** B uniformly in $[S-1, B_{\text{max}}]$ then sample B operators from {AND, OR} independently with equal probability.
4. **Build a binary tree** with those B nodes, using the sampling procedure of [46], designed to produce a diverse mix of deep and narrow versus shallow and wide trees.
5. **Negate some of the nodes** of the tree by adding NOT gates independently with probability $p_{\text{NOT}} = 1/2$.
6. **Fill in the leaves**: for each of the $B+1$ leaves in the tree, sample independently and uniformly one of the variables from the set of active variables[3].
7. **Simplify** the tree using Boolean algebra rules, as described in App. A. This greatly reduces the number of operators, and occasionally reduces the number of active variables.

Note that the distribution of functions generated in this way spans the whole space of possible Boolean functions (which is of size 2^{2^D}), but in a non-uniform fashion[4] with a bias towards a controlled depth or width. To maximize diversity, we sample large formulas (up to $B_{\text{max}} = 500$ binary gates), which are then heavily pruned in the simplification step[5]. As discussed quantitatively in App. B, the diversity of functions generated in this way is such that throughout the whole training procedure, functions of dimension $D \geq 7$ are typically encountered at most once.

To represent Boolean formulas as sequences fed to the Transformer, we enumerate the nodes of the trees in prefix order, i.e., direct Polish notation as in [46]: operators and variables are represented as single autonomous tokens, e.g. $[\texttt{AND}, x_1, \texttt{NOT}, x_2]$[6]. The inputs are embedded using $\{\texttt{0}, \texttt{1}\}$ tokens.

2.2 Generating inputs

Once the function f is generated, we sample N points $\boldsymbol{x}$ uniformly in the Boolean hypercube and compute the corresponding outputs $y = f(\boldsymbol{x})$. Optionally, we may flip the bits of the inputs and outputs independently with probability σ_{flip}; we consider the two following setups.

---

[2] A Boolean formula is a tree where input bits can appear more than once, and differs from a Boolean circuit, which is a directed graph which can feature cycles, but where each input bit appears once at most.

[3] The first S variables are sampled without replacement in order for all the active variables to appear in the tree.

[4] More involved generation procedures, e.g. involving Boolean circuits, could be envisioned as discussed in Sec. 5, but we leave this for future work.

[5] The simplification leads to a non-uniform distribution of number of operators as shown in App. A.

[6] We discard formulas which require more than 200 tokens.

Noiseless regime The noiseless regime, studied in Sec. 3, is defined as follows:

- **Noiseless data:** there is no bit flipping, i.e. $\sigma_{\text{flip}} = 0$.
- **Full support:** all the input bits affect the output, i.e. $S = D$.
- **Full observability:** the model has access to the whole truth table of the Boolean function, i.e. $N = 2^D$. Due to the quadratic length complexity of Transformers, this limits us to rather small input dimensions, i.e. $D_{\text{max}} = 10$.

Noisy regime In the noisy regime, studied in Sec. 4, the model must determine which variables affect the output, while also being able to cope with corruption of the inputs and outputs. During training, we vary the amount of noise for each sample so that the model can handle a variety of noise levels:

- **Noisy data:** the probability of each bit (both input and output) being flipped σ_{flip} is sampled uniformly in $[0, 0.1]$.
- **Partial support:** the model can handle high-dimensional functions, $D_{\text{max}} = 120$, but the number of active variables is sampled uniformly in $[0, 6]$. All the other variables are inactive.
- **Partial observability:** a subset of the hypercube is observed: the number of input points N is sampled uniformly in $[30, 300]$, which is typically much smaller that 2^D. Additionally, instead of sampling uniformly (which would cause distribution shifts if the inputs are not uniformly distributed at inference), we generate the input points via a random walk in the hypercube. Namely, we sample an initial point $\boldsymbol{x}_0$ then construct the following points by flipping independently each coordinate with a probability γ_{expl} sampled uniformly in $[0.05, 0.25]$.

2.3 Model

Embedder Our model is provided N input points $(\boldsymbol{x}, y) \in \{0, 1\}^{D+1}$, each of which is represented by $D + 1$ tokens of dimension D_{emb}. As D and N become large, this would result in very long input sequences (e.g. 10^4 tokens for $D = 100$ and $N = 100$) which challenge the quadratic complexity of Transformers. To mitigate this, we introduce an embedder to map each input pair $(\boldsymbol{x}, y)$ to a single embedding, following [44]. The embedder pads the empty input dimensions to D_{max}, enabling our model to handle variable input dimension, then concatenates all the tokens and feeds the $(D_{\text{max}} + 1)D_{\text{emb}}$-dimensional result into a 2-layer fully-connected feedforward network (FFN) with ReLU activations, which projects down to dimension D_{emb}. The resulting N embeddings of dimension D_{emb} are then fed to the Transformer.

Transformer We use a sequence-to-sequence Transformer architecture [1] where both the encoder and the decoder use 8 layers, 16 attention heads and an embedding dimension of 512, for a total of around 60M parameters (2M in the embedder, 25M in the encoder and 35M in the decoder). A notable property of this task is the permutation invariance of the N input points. To account for this invariance, we remove positional embeddings from the encoder. The decoder uses standard learnable positional encodings.

2.4 Training and evaluation

Training We optimize a cross-entropy loss with the Adam optimizer and a batch size of 128, warming up the learning rate from 10^{-7} to 2×10^{-4} over the first 10,000 steps, then decaying it using a cosine anneal for the next 300,000 steps, then restarting the annealing cycle with a damping factor of 3/2. We do not use any regularization, either in the form of weight decay or dropout. We train our models on around 30M examples; on a single NVIDIA A100 GPU with 80GB memory and 8 CPU cores, this takes about 3 days.

Inference At inference time, we find that beam search is the best decoding strategy in terms of diversity and quality. In most results presented in this paper, we use a beam size of 10. One major advantage here is that we have an easy criterion to rank candidates, which is how well they fit the input data – to assess this, we use the fitting error defined in the following section. Note that when the data is noiseless, the model will often find several candidates which perfectly fit the inputs, as shown in App. G: in this case, we select the shortest formula, i.e. that with smallest number of gates.

Evaluation Given a set of input-output pairs $\mathcal{D}$ generated by a target function $f_\star$, we compute the error of a predicted function f as $\epsilon_\mathcal{D} = \frac{1}{|\mathcal{D}|}\sum_{(\boldsymbol{x},y)\in\mathcal{D}} 1[f(\boldsymbol{x}) = f_\star(\boldsymbol{x})]$. We can then define:

- **Fitting error:** error obtained when re-using the points used to predict the formula, $\mathcal{D} = \mathcal{D}_\text{fit}$
- **Fitting accuracy:** defined as 1 if the fitting error is strictly equal to 0, and 0 otherwise.
- **Test error:** error obtained when sampling points uniformly at random in the hypercube outside of $\mathcal{D}_\text{fit}$. Note that we can only assess this in the noisy regime, where the model observes a subset of the hypercube.
- **Test accuracy:** defined as 1 if the test error is strictly equal to 0, and 0 otherwise.

3 Noiseless regime: finding the shortest formula

We begin by the noiseless regime (see Sec. 2.2). This setting is akin to logic synthesis, where the goal is to find the shortest formula that implements a given function.

In-domain performance In Fig. 4, we report the performance of the model when varying the number of input bits and the number of operators of the ground truth. Metrics are averaged over 10^4 samples from the random generator; as demonstrated in App. B, these samples have typically not been seen during training for $D \geq 7$.
We observe that the model is able to recover the target function with high accuracy in all cases, even for $D \geq 7$ where memorization is impossible. We emphasize however that these results only quantify the performance of our model on the distribution of functions it was trained on, which is highly-nonuniform in the 2^{2^D}-dimensional space of Boolean functions. We give a few examples of success and failure cases below.

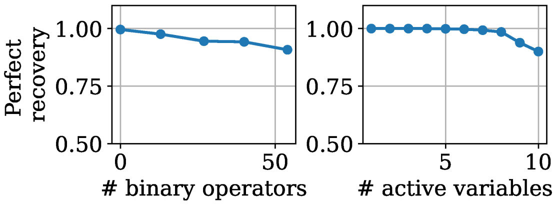

Figure 4: **Our model is able to recover the formula of unseen functions with high accuracy.** We report the fitting error and accuracy of our model when varying the number of binary gates and input bits. Metrics are averaged over 10k samples from the random function generator.

Success and failure cases In Fig. 1, we show two examples of Boolean functions where our model successfully predicts a compact formula for: the 4-to-1 multiplexer (which takes 6 input bits) and the 5-bit comparator (which takes 10 input bits). In App. D, we provide more examples: addition and multiplication, as well as majority and parity functions. By increasing the dimensionality of each problem up to the point of failure, we show that in all cases our model typically predicts exact and compact formulas as long as the function can be expressed with less than 50 binary gates (which is the largest size seen during training, as larger formulas exceed the 200 token limit) and fails beyond.

Hence, the failure point depends on the intrinsic difficulty of the function: for example, Boolformer can predict an exact formula for the comparator function up to $D = 10$, but only $D = 6$ for multiplication, $D = 5$ for majority and $D = 4$ for parity as well as typical random functions (whose outputs are independently sampled from $\{0, 1\}$). Parity functions are well-known to be the most difficult functions to learn for SQ models due to their leap-complexity, are also the hardest to learn in our framework because they require the most operators to be expressed (the XOR operator being excluded in this work).

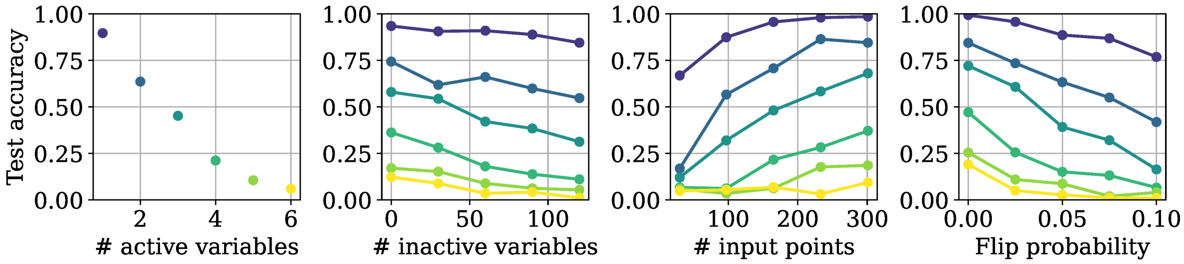

Figure 5: **Our model is robust to data incompleteness, bit flipping and noisy variables.** We display the error and accuracy of our model when varying the four factors of difficulty described in Sec. 2. The colors depict different number of active variables, as shown in the first panel. Metrics are averaged over 10k samples from the random generator.

4 Noisy regime: applications to real-world data

We now turn to the noisy regime, which is defined at the end of Sec. 2.2. We begin by examining in-domain performance as before, then present two real-world applications: binary classification and gene regulatory network inference.

4.1 Results on noisy data

In Fig. 5, we show how the performance of our model depends on the various factors of difficulty of the problem. The different colors correspond to different numbers of active variables, as shown in the leftmost panel: in this setting with multiple sources of noise, we see that accuracy drops much faster with the number of active variables than in the noiseless setting.

As could be expected, performance improves as the number of input points N increases, and degrades as the amount of random flipping and the number of inactive variables increase. However, the influence of the two latter parameters is rather mild, signalling that our model has an excellent ability to identify the support of the function and discard noise.

4.2 Application to interpretable binary classification

In this section, we show that our noisy model can be applied to binary classification tasks, providing an interpretable alternative to classic machine learning methods on tabular data.

Method We consider the tabular datasets from the Penn Machine Learning Benchmark (PMLB) from [17]. These encapsulate a wide variety of real-world problems such as predicting chess moves, toxicity of mushrooms, credit scores and heart diseases. Since our model can only take binary features as input, we discard continuous features, and binarize the categorical features with $C > 2$ classes into C binary variables. Note that this procedure can greatly increase the total number of features – we only keep datasets for which it results in less than 120 features (the maximum our model can handle). We randomly sample 25% of the examples for testing and report the F1 score obtained on this held out set.

We compare our model with two classic machine learning methods: logistic regression and random forests, using the default hyperparameters from `sklearn`. For random forests, we test two values for the number of estimators: 1 (in which case we obtain a simple decision tree as for the boolformer) and 100.

Results Results are reported in Fig. 6, where for readability we only display the datasets where the RandomForest with 100 estimators achieves an F1 score above 0.75. The performance of the Boolformer is similar on average to that of logistic regression: logistic regression typically performs better on "hard" datasets where there is no exact logical rule, for example medical diagnosis tasks such as `heart_h`, but worse on logic-based datasets where the data is not linearly separable such as `xd6`.

The F1 score of our model is slightly below that of a random forest of 100 trees, but slightly above that of the random forest with a single tree. This is remarkable considering that the Boolean formula it outputs only contains a few dozen nodes at most, whereas the trees of random forest use up to several hundreds. We show an example of a Boolean formula predicted for the mushroom toxicity dataset in Fig. 2, and a more extensive collection of formulas in App. E.

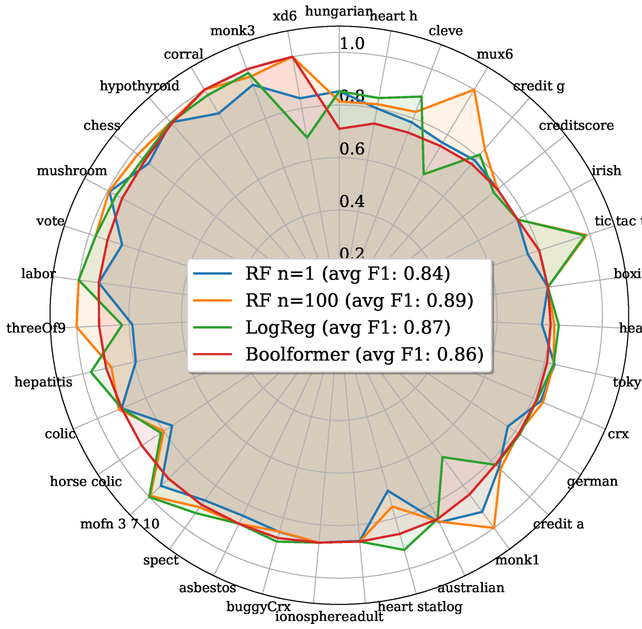

Figure 6: **Our model is competitive with classic machine learning methods while providing highly interpretable results.** We display the F1 score obtained on various binary classification datasets from the Penn Machine Learning Benchmark [17]. We compare the F1 score of the Boolformer with random forests (using 1 and 100 estimators) and logistic regression, using the default settings of `sklearn`, and display the average F1 score of each method in the legend.

4.3 Inferring Boolean networks: application to gene regulatory networks

A Boolean network is a dynamical system composed of D bits whose transition from one state to the next is governed by a set of D Boolean functions[7]. These types of networks have attracted a lot of attention in the field of computational biology as they can be used to model gene regulatory networks (GRNs) [47] – see App. F for a brief overview of this field. In this setting, each bit represents the (discretized) expression of a gene (on or off) and each function represents the regulation of a gene by the other genes. In this section, we investigate the applicability of our symbolic regression-based approach to this task.

Benchmark We use the recent benchmark for GRN inference introduced by [48]. This benchmark compares 5 methods for Boolean network inference on 30 Boolean networks inferred from biological data, with sizes ranging from 16 to 64 genes, and assesses both dynamical prediction (how well the model predicts the dynamics of the network) and structural prediction (how well the model predicts the Boolean functions compared to the ground truth). Structural prediction is framed as the binary classification task of predicting whether variable i influences variable j, and can hence be evaluated by many binary classification metrics; we report here the structural F1 and the AUROC metrics which are the most holistic, and defer other metrics to App. F.

Method Our model predicts each component f_i of the Boolean network independently, by taking as input the whole state of the network at times $[0 \dots t-1]$ and as output the state of the ith bit at times $[1 \dots t]$. Once each component has been predicted, we can build a causal influence graph, where an arrow connects node i to node j if j appears in the update equation of i: an example is shown in Fig. 7c.

---

[7] The i-th function f_i takes as input the state of the D bits at time t and returns the state of the i-th bit at time $t+1$.

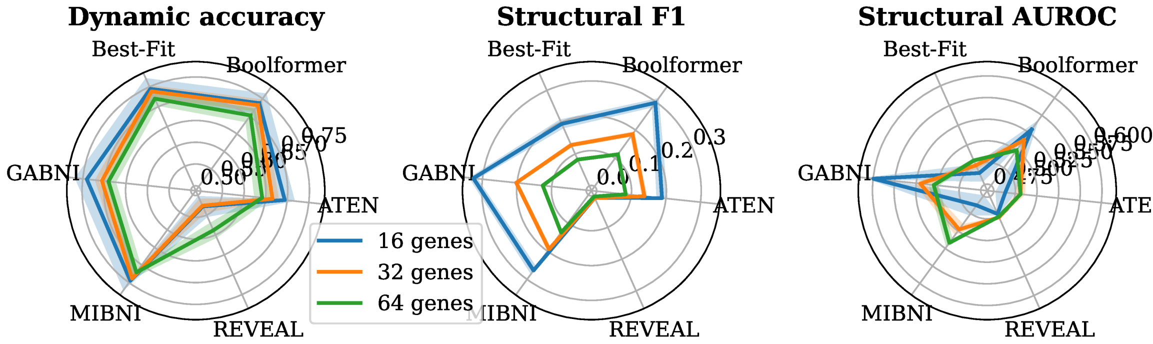

(a) Dynamic and structural metrics

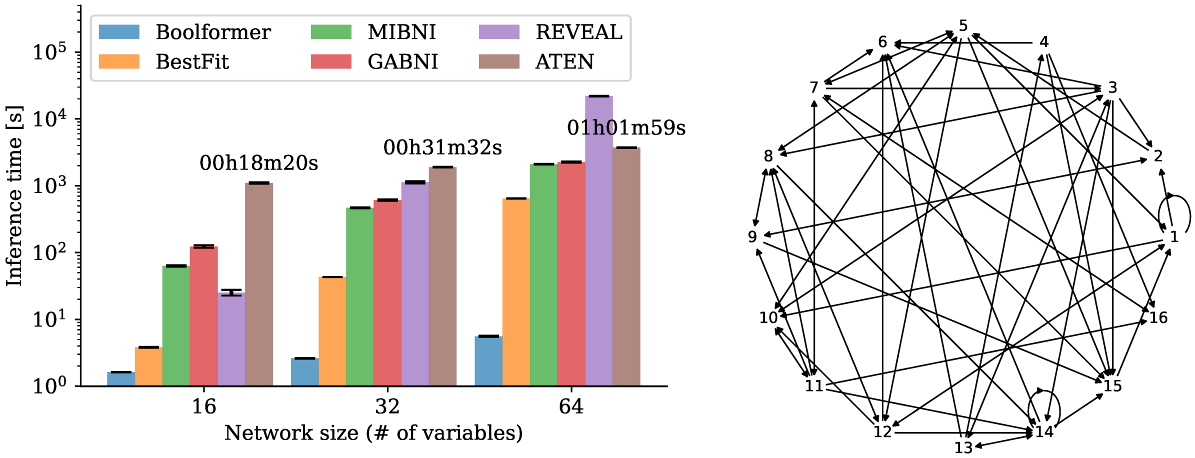

(b) Average inference time

(c) Example of a GRN inferred

Figure 7: **Our model is competitive with state-of-the-art methods for GRN inference with orders of magnitude faster inference.** (a) We compare the ability of our model to predict the next states (dynamic accuracy) and the influence graph (structural accuracy) with that of other methods using a recent benchmark [48] – more details in Sec. 4.3. (b) Average inference time of the various methods. (c) From the Boolean formulas predicted, one can construct an influence graph where each node represents a gene, and each arrow signals that one gene regulates another.

Note that since the dynamics of the Boolean network tend to be slow, an easy way to get rather high dynamical accuracy would be to simply predict the trivial fixed point $f_i = x_i$. In concurrent approaches, the function set explored excludes this solution; in our case, we simply mask the ith bit from the input when predicting f_i.

Results We display the results of our model on the benchmark in Fig. 7a. The Boolformer performs on par with the SOTA algorithms, GABNI [49] and MIBNI [50]. A striking feature of our model is its inference speed, displayed in Fig. 7b: a few seconds, against up to an hour for concurrent approaches, which mainly rely on genetic programming. Note also that our model predicts an interpretable Boolean function, whereas the other SOTA methods (GABNI and MIBNI) simply pick out the most important variables and the sign of their influence.

5 Discussion and limitations

In this work, we have shown that Transformers excel at symbolic regression of logical functions, both in the noiseless setup where they could potentially provide valuable insights for circuit design, and in the real-world setup of binary classification where they can provide interpretable solutions. Their ability to infer GRNs several orders of magnitude faster than existing methods offers the promise of many other exciting applications in biology, where Boolean modelling plays a key role [15]. There are however several limitations in our current approach, which open directions for future work.

First, due to the quadratic cost of self-attention, the number of input points is limited to a thousand during training, which limits the model's performance on high-dimensional functions and large datasets (although the model does exhibit some length generalization abilities at inference, as shown in App. C). One could address this shortcoming with linear attention mechanisms [51, 52], at the risk of degrading performance[8].

Second, the logical functions which our model is trained on do not include the XOR gate explicitly, limiting both the compactness of the formulas it predicts and its ability to express complex formulas such as parity functions. The reason for this limitation is that our generation procedure relies on expression simplification, which requires rewriting the XOR gate in terms of AND, OR and NOT. We leave it as a future work to adapt the generation of simplified formulas containing XOR gates, as well as operators with higher parity as in [40].

Third, the simplicity of the formulas predicted is limited in two additional ways: our model only handles (i) single-output functions – multi-output functions are predicted independently component-wise and (ii) gates with a fan-out of one[9]. As a result, our model cannot reuse intermediary results for different outputs or for different computations within a single output[10]. One could address this either by post-processing the generated formulas to identify repeated substructures, or by adapting the data generation process to support multi-output functions (a rather easy extension) and cyclic graphs (which would require more work).

Finally, this paper mainly focused on investigating concrete applications and benchmarks to motivate the potential and development of Boolformers. In future research, we will tackle various theoretical aspects of this paradigm, such as the model simplicity bias, the sample complexity and the 'generalization on the unseen' [27] of the Boolformer, comparing with other existing methods for Boolean learning.

Acknowledgements

We thank Philippe Schwaller, Geemi Wellawatte, Enric Boix-Adsera, Alexander Mathis and François Charton for insightful discussions. We also thank Russ Webb, Samira Abnar and Omid Saremi for valuable thoughts and feedback on this work. SD acknowledges funding from the EPFL AI4science program.

---

[8] We hypothesize that full attention span is particularly important in this specific task: the attention maps displayed in App. H are visually quite dense and high-rank matrices.

[9] Note that although the fan-in is fixed to 2 during training, it is easy to transform the predictions to larger fan-in by merging ORs and ANDs together.

[10] Consider the D-parity: one can build a formula with only $3(n-1)$ binary AND-OR gates by storing $D-1$ intermediary results: $a_1 = XOR(x_1, x_2), a_2 = XOR(a_1, x_3), \dots, a_{n-1} = XOR(a_{D-2}, x_D)$. Our model needs to recompute these intermediary values, leading to much larger formulas, e.g. 35 binary gates instead of 9 for the 4-parity as illustrated in App. D.

References

1. Vaswani, A. *et al. Attention is all you need* in *Advances in neural information processing systems* (2017), 5998–6008.
2. Dosovitskiy, A. *et al.* An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
3. Brown, T. *et al. Language Models are Few-Shot Learners* in *Advances in Neural Information Processing Systems* (eds Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. & Lin, H.) **33** (Curran Associates, Inc., 2020), 1877–1901. `https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf`.
4. Jumper, J. *et al.* Highly accurate protein structure prediction with AlphaFold. *Nature* **596,** 583–589 (2021).
5. Delétang, G. *et al.* Neural networks and the chomsky hierarchy. *arXiv preprint arXiv:2207.02098* (2022).
6. Saxton, D., Grefenstette, E., Hill, F. & Kohli, P. Analysing mathematical reasoning abilities of neural models. *arXiv preprint arXiv:1904.01557* (2019).
7. Lewkowycz, A. *et al.* Solving quantitative reasoning problems with language models. *Advances in Neural Information Processing Systems* **35,** 3843–3857 (2022).
8. Veličković, P. *et al. The CLRS algorithmic reasoning benchmark* in *International Conference on Machine Learning* (2022), 22084–22102.
9. Zhang, Y. *et al.* Unveiling transformers with lego: a synthetic reasoning task. *arXiv preprint arXiv:2206.04301* (2022).
10. Zhang, C., Raghu, M., Kleinberg, J. & Bengio, S. Pointer value retrieval: A new benchmark for understanding the limits of neural network generalization. *arXiv preprint arXiv:2107.12580* (2021).
11. Johnson, J. *et al. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning* in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2017), 2901–2910.
12. Liu, J. *et al.* Logiqa: A challenge dataset for machine reading comprehension with logical reasoning. *arXiv preprint arXiv:2007.08124* (2020).
13. Cobbe, K. *et al.* Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168* (2021).
14. Wang, R.-S., Saadatpour, A. & Albert, R. Boolean modeling in systems biology: an overview of methodology and applications. *Physical biology* **9,** 055001 (2012).
15. Hemedan, A. A., Niarakis, A., Schneider, R. & Ostaszewski, M. Boolean modelling as a logic-based dynamic approach in systems medicine. *Computational and Structural Biotechnology Journal* **20,** 3161–3172 (2022).
16. Abbe, E. *et al.* Learning to reason with neural networks: Generalization, unseen data and boolean measures. *arXiv preprint arXiv:2205.13647* (2022).
17. Olson, R. S., La Cava, W., Orzechowski, P., Urbanowicz, R. J. & Moore, J. H. PMLB: a large benchmark suite for machine learning evaluation and comparison. *BioData Mining* **10,** 36. `https://doi.org/10.1186/s13040-017-0154-4` (Dec. 2017).
18. Evans, R. & Grefenstette, E. Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research* **61,** 1–64 (2018).
19. Ciravegna, G. *et al.* Logic explained networks. *Artificial Intelligence* **314,** 103822 (2023).
20. Shi, S. *et al. Neural logic reasoning* in *Proceedings of the 29th ACM International Conference on Information & Knowledge Management* (2020), 1365–1374.

21. Dong, H. *et al.* Neural logic machines. *arXiv preprint arXiv:1904.11694* (2019).
22. Jelassi, S. *et al.* Length Generalization in Arithmetic Transformers. *arXiv preprint arXiv:2306.15400* (2023).
23. Liu, B., Ash, J. T., Goel, S., Krishnamurthy, A. & Zhang, C. Transformers learn shortcuts to automata. *arXiv preprint arXiv:2210.10749* (2022).
24. Wei, J. *et al. Emergent Abilities of Large Language Models* 2022. arXiv: `2206.07682 [cs.CL]`.
25. Hellerstein, L. & Servedio, R. A. On pac learning algorithms for rich boolean function classes. *Theoretical Computer Science* **384,** 66–76 (2007).
26. Reyzin, L. Statistical queries and statistical algorithms: Foundations and applications. *arXiv preprint arXiv:2004.00557* (2020).
27. Abbe, E., Boix-Adsera, E. & Misiakiewicz, T. SGD learning on neural networks: leap complexity and saddle-to-saddle dynamics. *arXiv preprint arXiv:2302.11055* (2023).
28. Narodytska, N., Ignatiev, A., Pereira, F., Marques-Silva, J. & Ras, I. *Learning Optimal Decision Trees with SAT.* in *Ijcai* (2018), 1362–1368.
29. Wang, T. & Rudin, C. Learning optimized Or's of And's. *arXiv preprint arXiv:1511.02210* (2015).
30. Su, G., Wei, D., Varshney, K. R. & Malioutov, D. M. Interpretable two-level boolean rule learning for classification. *arXiv preprint arXiv:1511.07361* (2015).
31. Malioutov, D. M., Varshney, K. R., Emad, A. & Dash, S. Learning interpretable classification rules with boolean compressed sensing. *Transparent Data Mining for Big and Small Data,* 95–121 (2017).
32. Karnaugh, M. The map method for synthesis of combinational logic circuits. *Transactions of the American Institute of Electrical Engineers, Part I: Communication and Electronics* **72,** 593–599 (1953).
33. Rudell, R. L. & Sangiovanni-Vincentelli, A. Multiple-valued minimization for PLA optimization. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **6,** 727–750 (1987).
34. Murray, C. D. & Williams, R. R. On the (non) NP-hardness of computing circuit complexity. *Theory of Computing* **13,** 1–22 (2017).
35. Scarabottolo, I., Ansaloni, G. & Pozzi, L. *Circuit carving: A methodology for the design of approximate hardware* in *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)* (2018), 545–550.
36. Venkataramani, S., Sabne, A., Kozhikkottu, V., Roy, K. & Raghunathan, A. *SALSA: Systematic logic synthesis of approximate circuits* in *Proceedings of the 49th Annual Design Automation Conference* (2012), 796–801.
37. Venkataramani, S., Roy, K. & Raghunathan, A. *Substitute-and-simplify: A unified design paradigm for approximate and quality configurable circuits* in *2013 Design, Automation & Test in Europe Conference & Exhibition (DATE)* (2013), 1367–1372.
38. Boroumand, S., Bouganis, C.-S. & Constantinides, G. A. *Learning boolean circuits from examples for approximate logic synthesis* in *Proceedings of the 26th Asia and South Pacific Design Automation Conference* (2021), 524–529.
39. Oliveira, A. & Sangiovanni-Vincentelli, A. Learning complex boolean functions: Algorithms and applications. *Advances in Neural Information Processing Systems* **6** (1993).
40. Rosenberg, G. *et al.* Explainable AI using expressive Boolean formulas. *arXiv preprint arXiv:2306.03976* (2023).
41. La Cava, W. *et al.* Contemporary symbolic regression methods and their relative performance. *arXiv preprint arXiv:2107.14351* (2021).
42. Biggio, L., Bendinelli, T., Neitz, A., Lucchi, A. & Parascandolo, G. *Neural Symbolic Regression that Scales* 2021. arXiv: `2106.06427 [cs.LG]`.

43. Valipour, M., You, B., Panju, M. & Ghodsi, A. SymbolicGPT: A Generative Transformer Model for Symbolic Regression. *arXiv preprint arXiv:2106.14131* (2021).

44. Kamienny, P.-A., d'Ascoli, S., Lample, G. & Charton, F. *End-to-end Symbolic Regression with Transformers* in *Advances in Neural Information Processing Systems* (eds Oh, A. H., Agarwal, A., Belgrave, D. & Cho, K.) (2022). `https://openreview.net/forum?id=GoOuIrDHG_Y`.

45. Tenachi, W., Ibata, R. & Diakogiannis, F. I. Deep symbolic regression for physics guided by units constraints: toward the automated discovery of physical laws. *arXiv preprint arXiv:2303.03192* (2023).

46. Lample, G. & Charton, F. Deep learning for symbolic mathematics. *arXiv preprint arXiv:1912.01412* (2019).

47. Zhao, M., He, W., Tang, J., Zou, Q. & Guo, F. A comprehensive overview and critical evaluation of gene regulatory network inference technologies. *Briefings in Bioinformatics* **22,** bbab009 (2021).

48. Pušnik, Ž., Mraz, M., Zimic, N. & Moškon, M. Review and assessment of Boolean approaches for inference of gene regulatory networks. *Heliyon,* e10222 (2022).

49. Barman, S. & Kwon, Y.-K. A Boolean network inference from time-series gene expression data using a genetic algorithm. *Bioinformatics* **34,** i927–i933 (2018).

50. Barman, S. & Kwon, Y.-K. A novel mutual information-based Boolean network inference method from time-series gene expression data. *PloS one* **12,** e0171097 (2017).

51. Choromanski, K. *et al.* Rethinking attention with performers. *arXiv preprint arXiv:2009.14794* (2020).

52. Wang, S., Li, B. Z., Khabsa, M., Fang, H. & Ma, H. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768* (2020).

53. Singh, N. & Vidyasagar, M. bLARS: An algorithm to infer gene regulatory networks. *IEEE/ACM transactions on computational biology and bioinformatics* **13,** 301–314 (2015).

54. Haury, A.-C., Mordelet, F., Vera-Licona, P. & Vert, J.-P. TIGRESS: trustful inference of gene regulation using stability selection. *BMC systems biology* **6,** 1–17 (2012).

55. Huynh-Thu, V. A., Irrthum, A., Wehenkel, L. & Geurts, P. Inferring regulatory networks from expression data using tree-based methods. *PloS one* **5,** e12776 (2010).

56. Adabor, E. S. & Acquaah-Mensah, G. K. Restricted-derestricted dynamic Bayesian Network inference of transcriptional regulatory relationships among genes in cancer. *Computational biology and chemistry* **79,** 155–164 (2019).

57. Huynh-Thu, V. A. & Geurts, P. dynGENIE3: dynamical GENIE3 for the inference of gene networks from time series expression data. *Scientific reports* **8,** 3384 (2018).

58. Liang, S., Fuhrman, S. & Somogyi, R. *Reveal, a general reverse engineering algorithm for inference of genetic network architectures* in *Biocomputing* **3** (1998).

59. Lähdesmäki, H., Shmulevich, I. & Yli-Harja, O. On learning gene regulatory networks under the Boolean network model. *Machine learning* **52,** 147 (2003).

60. Shi, N., Zhu, Z., Tang, K., Parker, D. & He, S. ATEN: And/Or tree ensemble for inferring accurate Boolean network topology and dynamics. *Bioinformatics* **36,** 578–585 (2020).

A Expression simplification

The data generation procedure heavily relies on expression simplification. This is of utmost importance for four reasons:

- It reduces the output expression length and hence memory usage as well as increasing speed
- It improves the supervision by reducing expressions to a more canonical form, easier to guess for the model
- It increases the effective diversity of the beam search, by reducing the number of equivalent expressions generated
- It encourages the model to output the simplest formula, which is a desirable property.

We use the package `boolean.py`[11] for this, which is considerably faster than `sympy` (the function `simplify_logic` of the latter has exponential complexity, and is hence only implemented for functions with less then 9 input variables).

Empirically, we found the following procedure to be optimal in terms of average length obtained after simplification:

1. Preprocess the formula by applying basic logical equivalences: double negation elimination and De Morgan's laws.
2. Parse the formula with `boolean.py` and run the `simplify()` method *twice*
3. Apply once again the first step

Note that this procedure drastically reduces the number of operators and renders the final distribution highly nonuniform, as shown in Fig. 8.

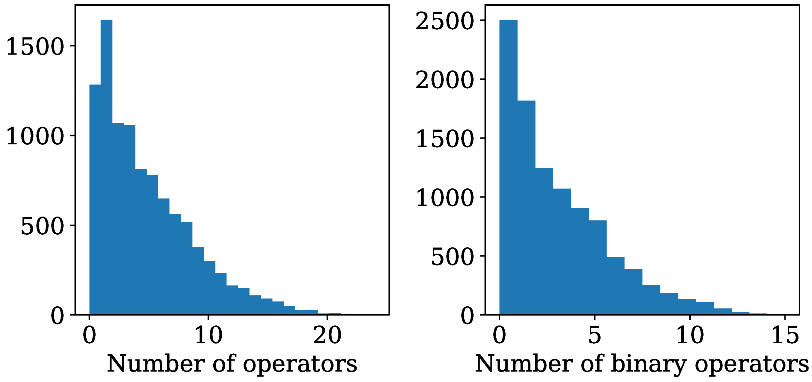

Figure 8: Distribution of number of operators after expression simplification. The initial number of binary operators is sampled uniformly in [1, 1000]. The total number of examples is 10^4.

B Does the Boolformer memorize?

One natural question is whether our model simply performs memorization on the training set. Indeed, the number of possible functions of D variables is finite, and equal to 2^{2^D}.

Let us first assume naively that our generator is uniform in the space of boolean functions. Since $2^{2^4} \simeq 6 \times 10^4$ (which is smaller than the number of examples seen during training) and $2^{2^5} \simeq 5.10^9$

---

[11] https://github.com/bastikr/Boolean.py

(which is much larger), one could conclude that for $D \leq 4$, all functions are memorized, whereas for $D > 4$, only a small subset of all possible functions are seen, hence memorization cannot occur.

However, the effective number of unique functions seen during training is actually smaller because our generator of random functions is nonuniform in the space of boolean functions. In this case, for which value of D does memorization become impossible? To investigate this question, for each $D < D_{\text{max}}$, we sample $\min\left(2^{2^D}, 100\right)$ unique functions from our random generator, and count how many times their exact truth table is encountered over an epoch (300,000 examples).

Results are displayed in Fig. 9. As expected, the average number of occurences of each function decays exponentially fast, and falls to zero for $D = 7$, meaning that each function is typically unique for $D \geq 7$. Hence, memorization cannot occur for $D \geq 7$. Yet, as shown in Fig. 4, our model achieves excellent accuracies even for functions of 10 variables, which excludes memorization as a possible explanation for the ability of our model to predict logical formulas accurately.

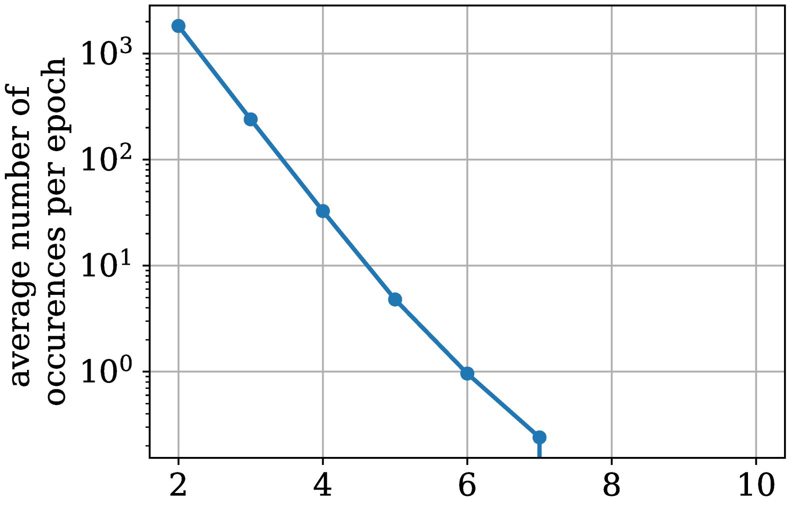

Figure 9: **Functions with 7 or more variables are typically never seen more than once during training.** We display the average number of times functions of various input dimensionalities are seen during an epoch (300,000 examples). For each point on the curve, the average is taken over $\min(2^{2^D}), 100)$ unique functions.

C Length generalization

In this section we examine the ability of our model to length generalize. In this setting, there are two types of generalization one can define: generalization in terms of the number of inputs N, or in terms of the number of active variables S[12]. We examine length generalization in the noisy setup (see Sec. 2.2), because in the noiseless setup, the model already has access to all the truth table (increasing N does not bring any extra information), and all the variables are active (we cannot increase S as it is already equal to D).

C.1 Number of inputs

Since the input points fed to the model are permutation invariant, our model does not use any positional embeddings. Hence, not only can our model handle $N > N_{\text{max}}$, but performance often continues to improve beyond N_{max}, as we show for two datasets extracted from PMLB [17] in Fig. 10.

---

[12] Note that our model cannot generalize to a problem of higher dimensionality D than seen during training, as its vocabulary only contains the names of variables ranging from x_1 to $x_{D_{\text{max}}}$.

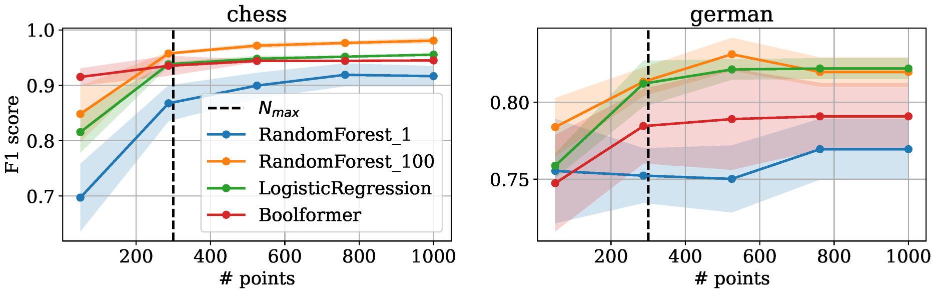

Figure 10: **Our model can length generalize in terms of sequence length.** We test a model trained with $N_{\text{max}} = 300$ on the `chess` and `german` datasets of PMLB. Results are averaged over 10 random samplings of the input points, with the shaded areas depicting the standard deviation.

C.2 Number of variables

To assess whether our model can infer functions which contain more active variables than seen during training, we evaluated a model trained on functions with up to 6 active variables on functions with 7 or more active variables. We provided the model with the truth table of two very simply functions: the OR and AND of the first $S \geq 7$ variables. We observe that the model succeeds for $S = 7$, but fails for $S \geq 8$, where it only includes the first 7 variables in the OR / AND. Hence, the model can length generalize to a small extent in terms of number of active variables, but less easily than in terms of number of inputs. We hypothesize that proper length generalization could be achieved by "priming", i.e. adding even a small number of "long" examples, as performed in [22].

D Formulas predicted for logical circuits

In Figs. 11 and 12, we show examples of some common arithmetic and logical formulas predicted by our model in the noiseless regime, with a beam size of 100. In all cases, we increase the dimensionality of the problem until the failure point of the Boolformer.

E Formulas predicted for PMLB datasets

In Fig. 13, we report a few examples of boolean formulas predicted for the PMLB datasets in Fig. 6. In each case, we also report the F1 scores of logistic regression and random forests with 100 estimators.

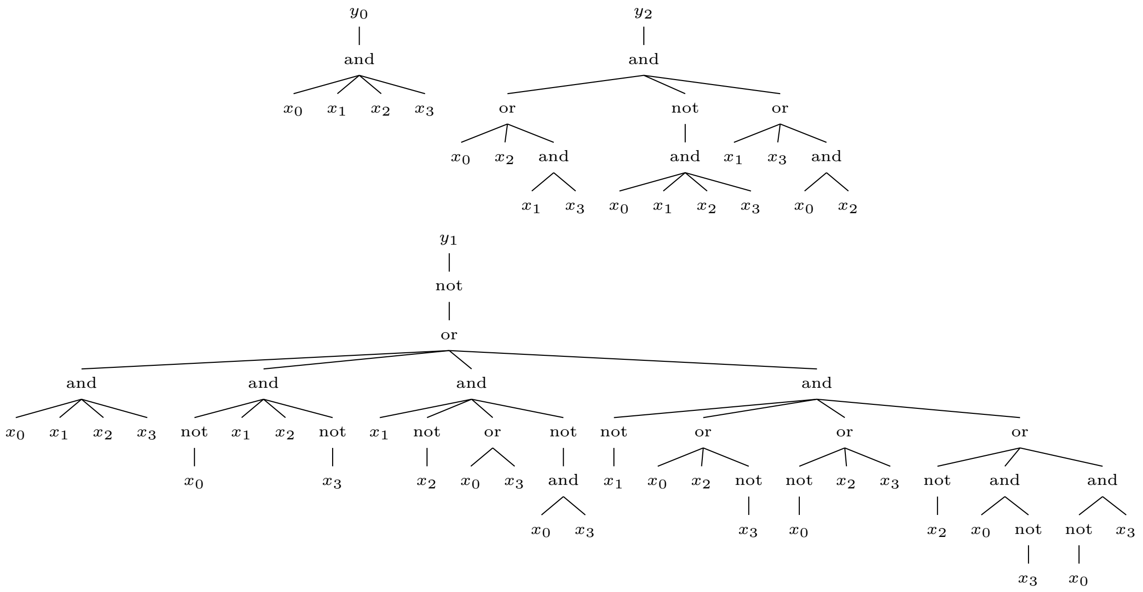
y_0
and
x_0 x_1 x_2 x_3
y_2
and
or not or
x_0 x_2 and and x_1 x_3 and
x_1 x_3 x_0 x_1 x_2 x_3 x_0 x_2
y_1
not
or
and and and and
x_0 x_1 x_2 x_3 not x_1 x_2 not x_1 not or not not or or or
x_0 x_3 x_2 x_0 x_3 and x_1 x_0 x_2 not not x_2 x_3 not and and
x_0 x_3 x_3 x_0 x_2 x_0 not not x_3
x_3 x_0

(a) Addition of four bits: $y_0y_1y_2 = x_0 + x_1 + x_2 + x_3$. All formulas are correct.

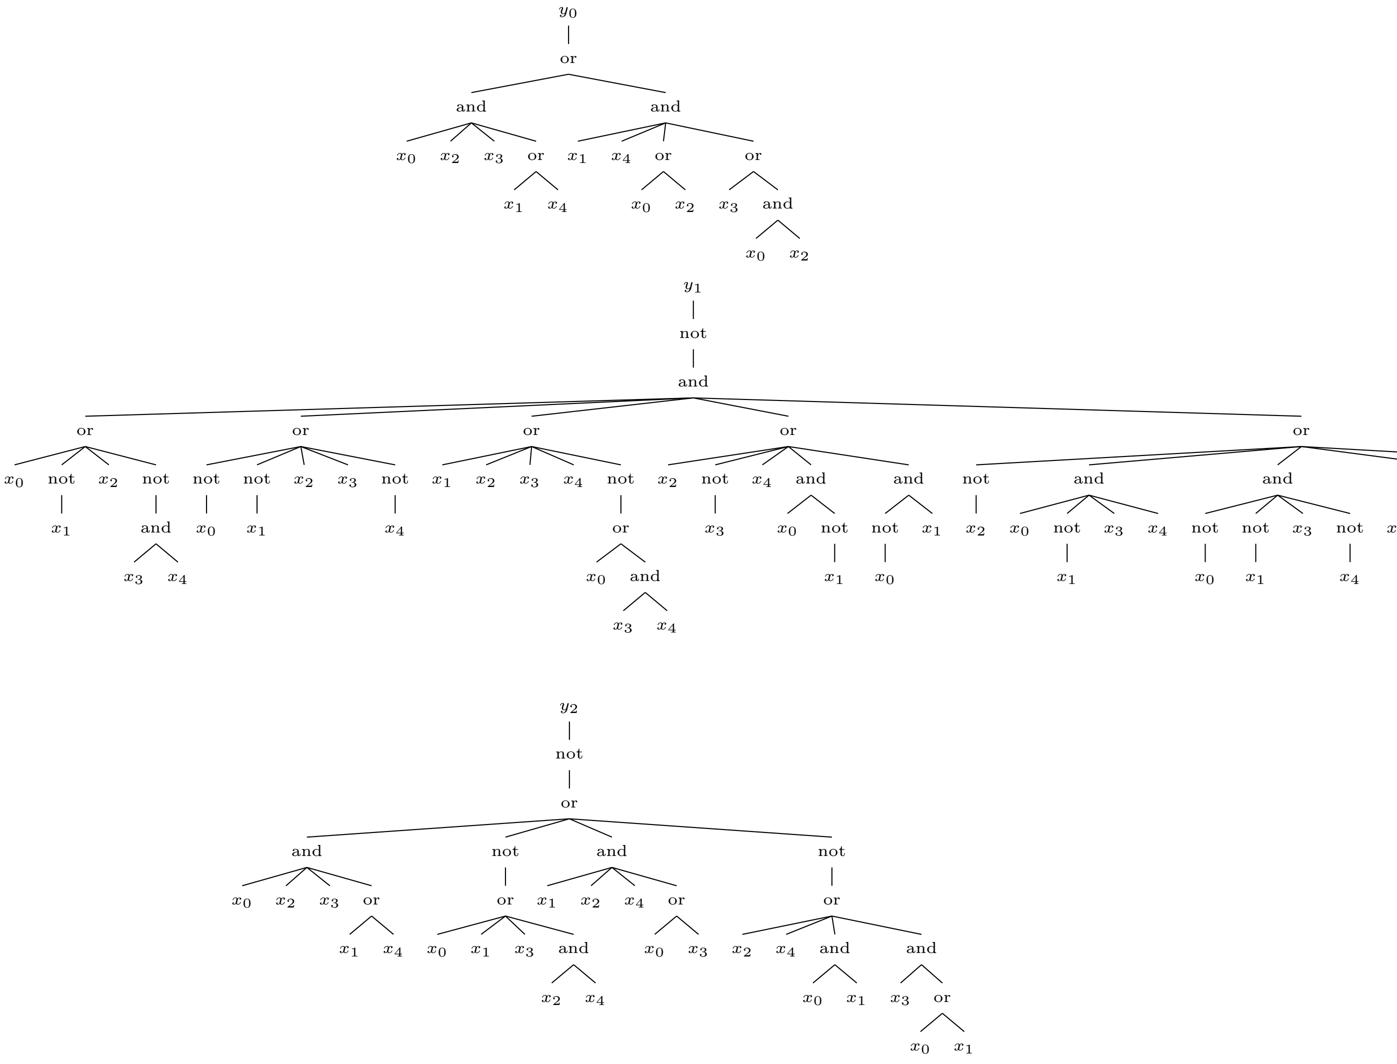
y_0
or
and and
x_0 x_2 x_3 or x_1 x_4 or or
x_1 x_4 x_0 x_2 x_3 and
x_0 x_2
y_1
not
and
or or or or or
x_0 not x_2 not not not x_2 x_3 not x_1 x_2 x_3 x_4 not x_2 not x_4 and and not and and
x_1 and x_0 x_1 x_4 or x_3 x_0 not not x_1 x_2 x_0 not x_3 x_4 not not x_3 not
x_3 x_4 x_0 and x_1 x_0 x_1 x_0 x_1 x_4
x_3 x_4
y_2
not
or
and not and not
x_0 x_2 x_3 or or x_1 x_2 x_4 or or
x_1 x_4 x_0 x_1 x_3 and x_0 x_3 x_2 x_4 and and
x_2 x_4 x_0 x_1 x_3 or
x_0 x_1

(b) Addition of five bits: $y_0y_1y_2 = x_0 + x_1 + x_2 + x_3 + x_4$. The first formula is correct, the second gets 20% error, the last 3% error.

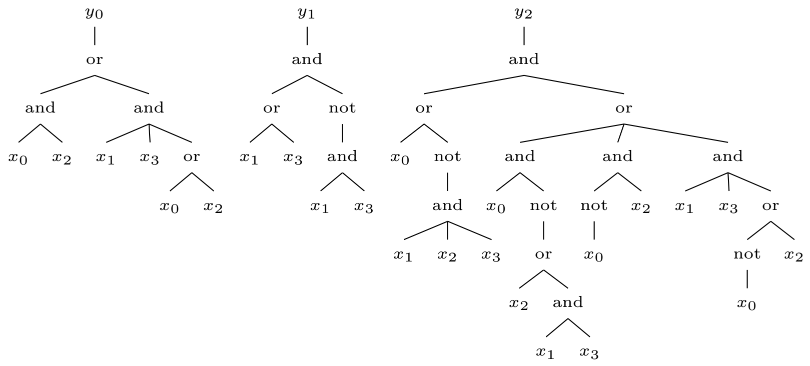
y_0 y_1 y_2
or and and
and and or not or or
x_0 x_2 x_1 x_3 or x_1 x_3 and x_0 not and and and
x_0 x_2 x_1 x_3 and x_0 not not x_2 x_1 x_3 or
x_1 x_2 x_3 or x_0 not x_2
x_2 and x_0
x_1 x_3

(c) Addition of two 2-bit numbers: $y_0y_1y_2 = (x_0x_1) + (x_2x_3)$. All formulas are correct.

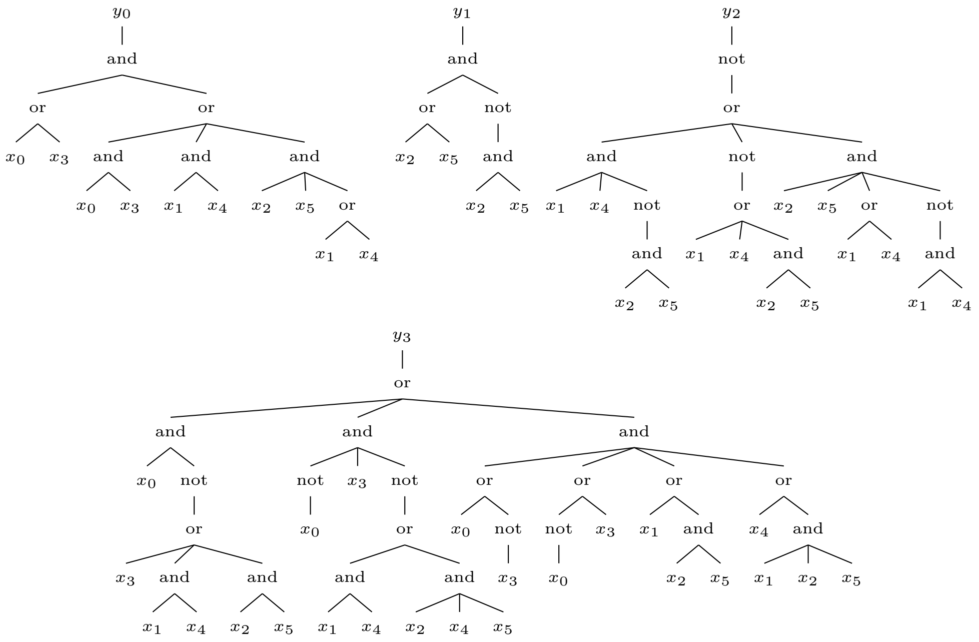
y_0 y_1 y_2
and and not
or or or not or
x_0 x_3 and and and x_2 x_5 and and not and
x_0 x_3 x_1 x_4 x_2 x_5 or x_2 x_5 x_1 x_4 not or x_2 x_5 or not
x_1 x_4 and x_1 x_4 and x_1 x_4 and
x_2 x_5 x_2 x_5 x_1 x_4
y_3
or
and and and
x_0 not not x_3 not or or or or
or x_0 or x_0 not not x_3 x_1 and x_4 and
x_3 and and and and x_3 x_0 x_2 x_5 x_1 x_2 x_5
x_1 x_4 x_2 x_5 x_1 x_4 x_2 x_4 x_5

(d) Addition of two 3-bit numbers: $y_0y_1y_2y_3 = (x_0x_1x_2) + (x_3x_4x_5)$. All formulas are correct, except y_3 which gets an error of 3%.

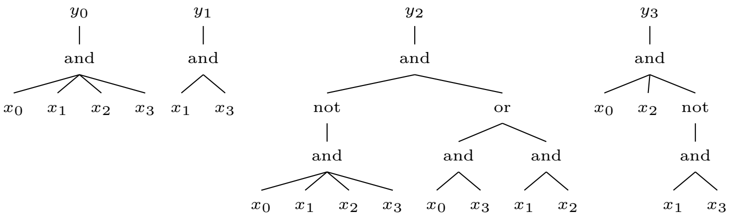
y_0 y_1 y_2 y_3
and and and and
x_0 x_1 x_2 x_3 x_1 x_3 not or x_0 x_2 not
and and and and
x_0 x_1 x_2 x_3 x_0 x_3 x_1 x_2 x_1 x_3

(e) Multiplication of two 2-bit numbers: $y_0y_1y_2y_3 = (x_0x_1) \times (x_2x_3)$. All formulas are correct.

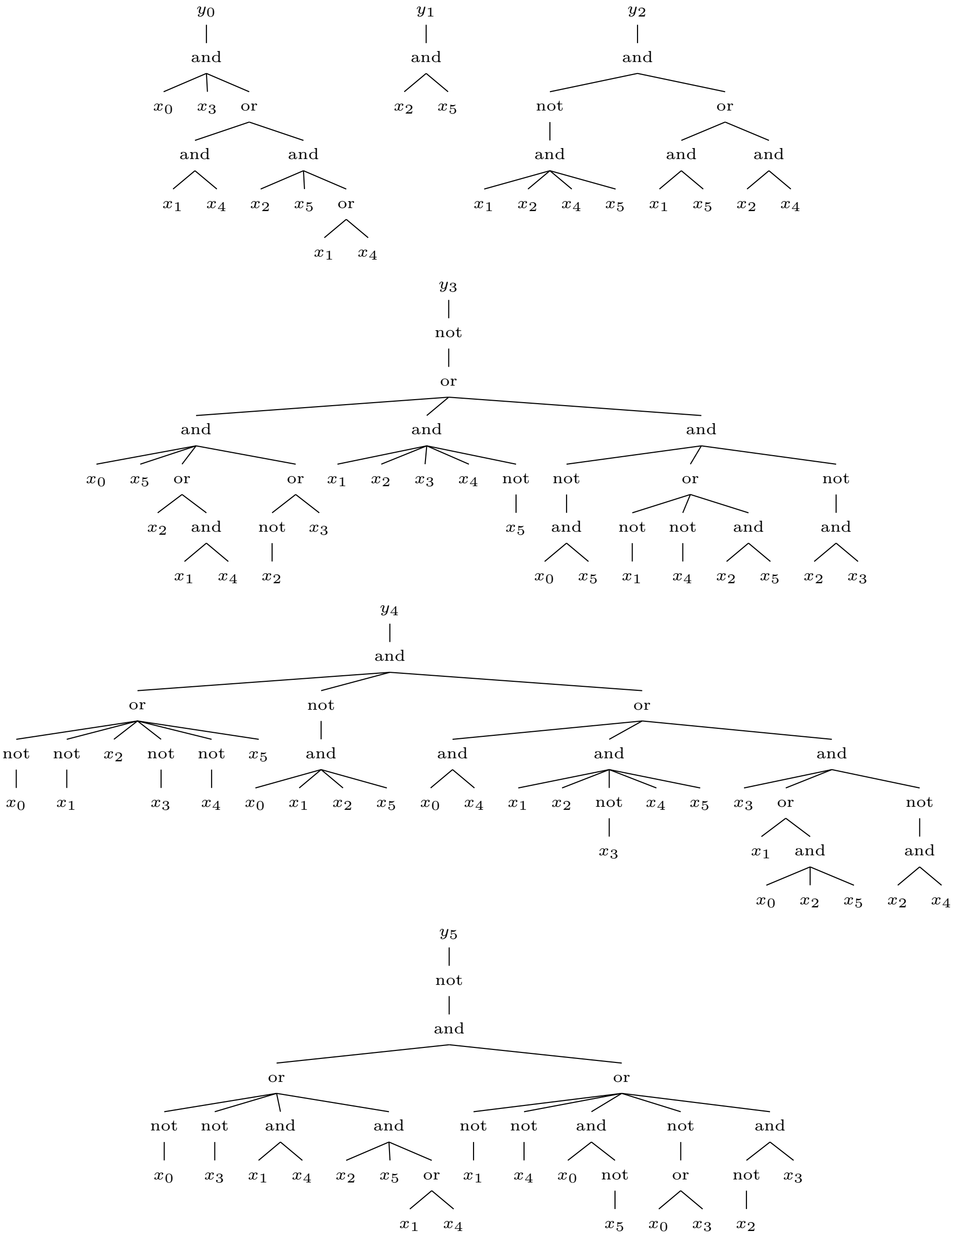
y_0 y_1 y_2
and and and
x_0 x_3 or x_2 x_5 not or
and and and and and
x_1 x_4 x_2 x_5 or x_1 x_2 x_4 x_5 x_1 x_5 x_2 x_4
x_1 x_4
y_3
not
or
and and and
x_0 x_5 or or x_1 x_2 x_3 x_4 not not or not
x_2 and not x_3 x_5 and not not and and
x_1 x_4 x_2 x_0 x_5 x_1 x_4 x_2 x_5 x_2 x_3
y_4
and
or not or
not not x_2 not not x_5 and and and and
x_0 x_1 x_3 x_4 x_0 x_1 x_2 x_5 x_0 x_4 x_1 x_2 not x_4 x_5 x_3 or not
x_3 x_1 and and
x_0 x_2 x_5 x_2 x_4
y_5
not
and
or or
not not and and not not and not and
x_0 x_3 x_1 x_4 x_2 x_5 or x_1 x_4 x_0 not or not x_3
x_1 x_4 x_5 x_0 x_3 x_2

(f) Multiplication of two 3-bit numbers: $y_0y_1y_2y_3y_4y_5 = (x_0x_1x_2) \times (x_3x_4x_5)$. All formulas are correct, except y_4 which gets an error of 5%.

Figure 11: **Some arithmetic formulas predicted by our model.**

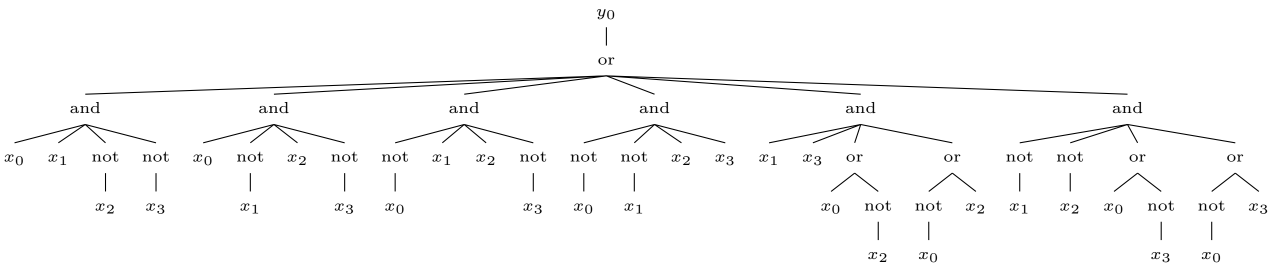

(a) 4-parity: 0% error.

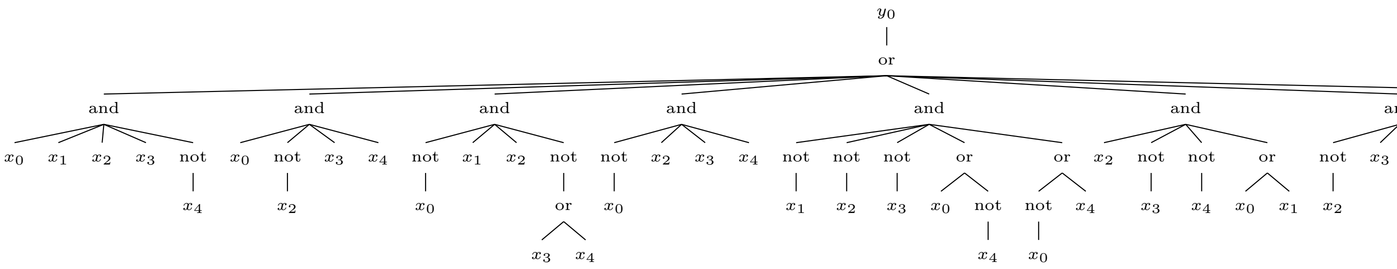

(b) 5-parity: 28% error.

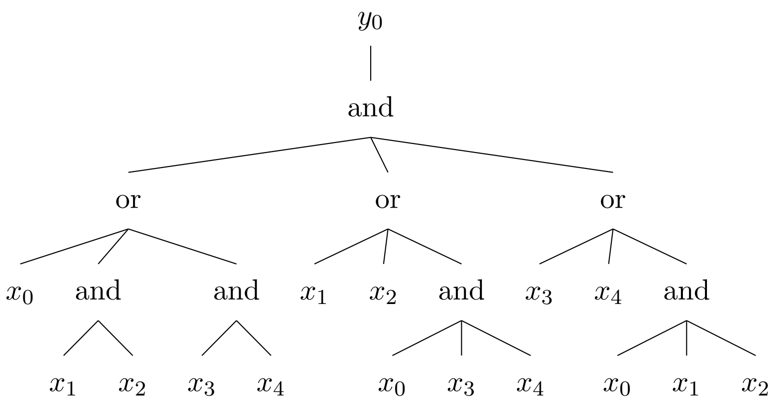

(c) 5-majority: 0% error.

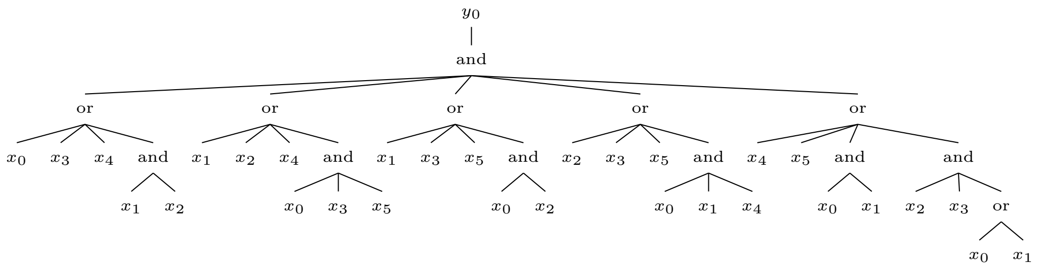

(d) 6-majority: 6% error.

Figure 12: **Some logical functions predicted by our model.**

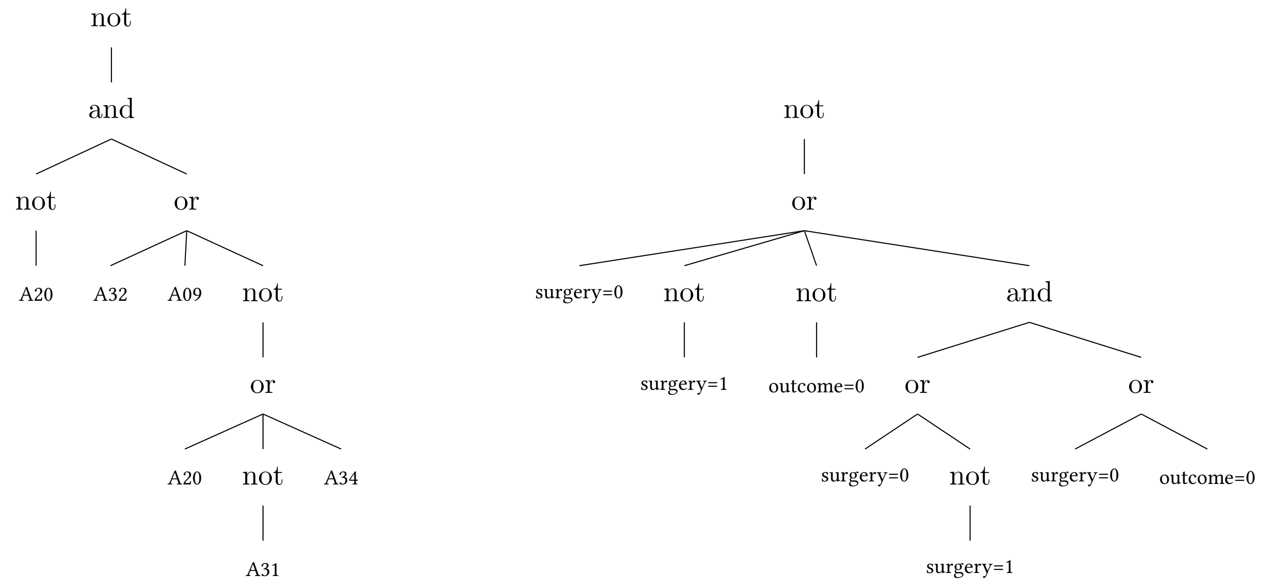
not
and
not
or
A20
A32
A09
not
or
A20
not
A34
A31
not
or
surgery=0
not
not
and
surgery=1
outcome=0
or
or
surgery=0
not
surgery=0
outcome=0
surgery=1

(a) chess. F1: 0.947. LogReg: 0.958. RF: 0.987.

(b) horse colic. F1: 0.900. LogReg: 0.822. RF: 0.861.

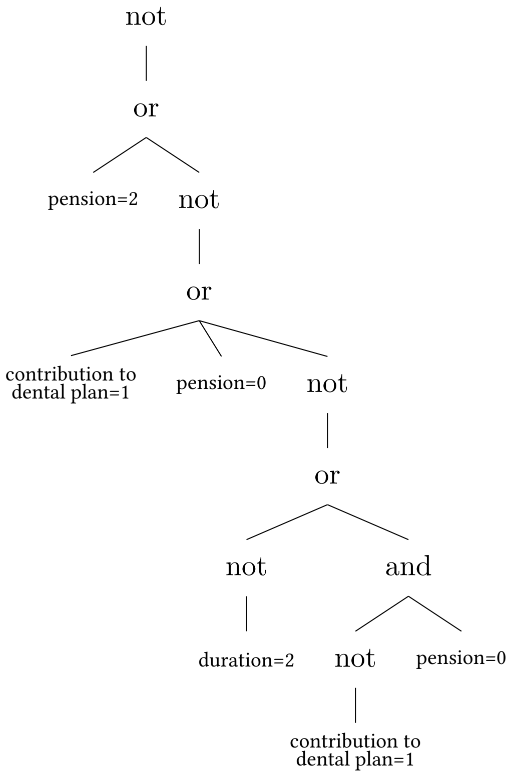
not
or
pension=2
not
or
contribution to dental plan=1
pension=0
not
or
not
and
duration=2
not
pension=0
contribution to dental plan=1

(c) labor. F1: 0.960. LogReg: 1.000. RF: 1.000.

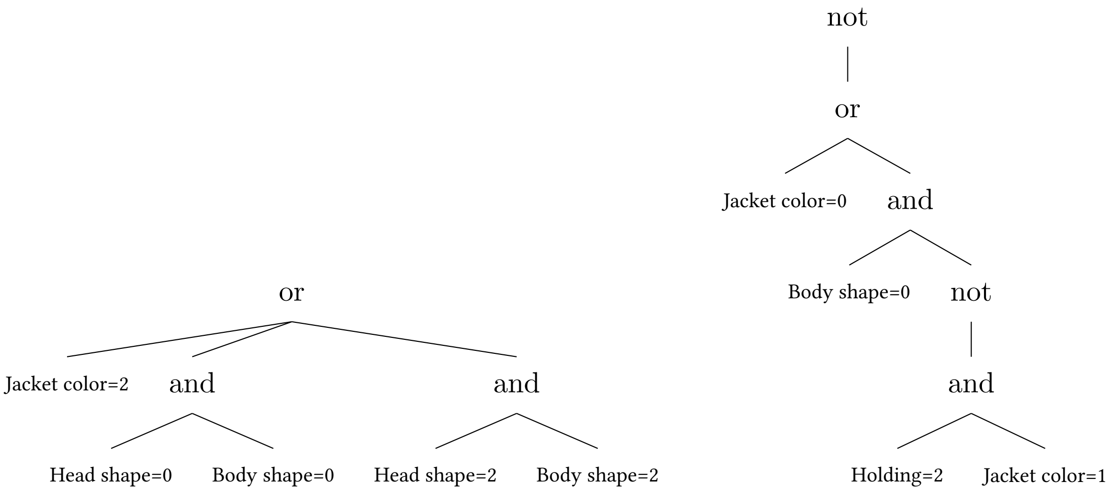
or
Jacket color=2
and
and
Head shape=0
Body shape=0
Head shape=2
Body shape=2
not
or
Jacket color=0
and
Body shape=0
not
and
Holding=2
Jacket color=1

(d) monk1. F1: 0.915. LogReg: 0.732. RF: 1.000.

(e) monk3. F1: 1.000. LogReg: 0.985. RF: 0.993.

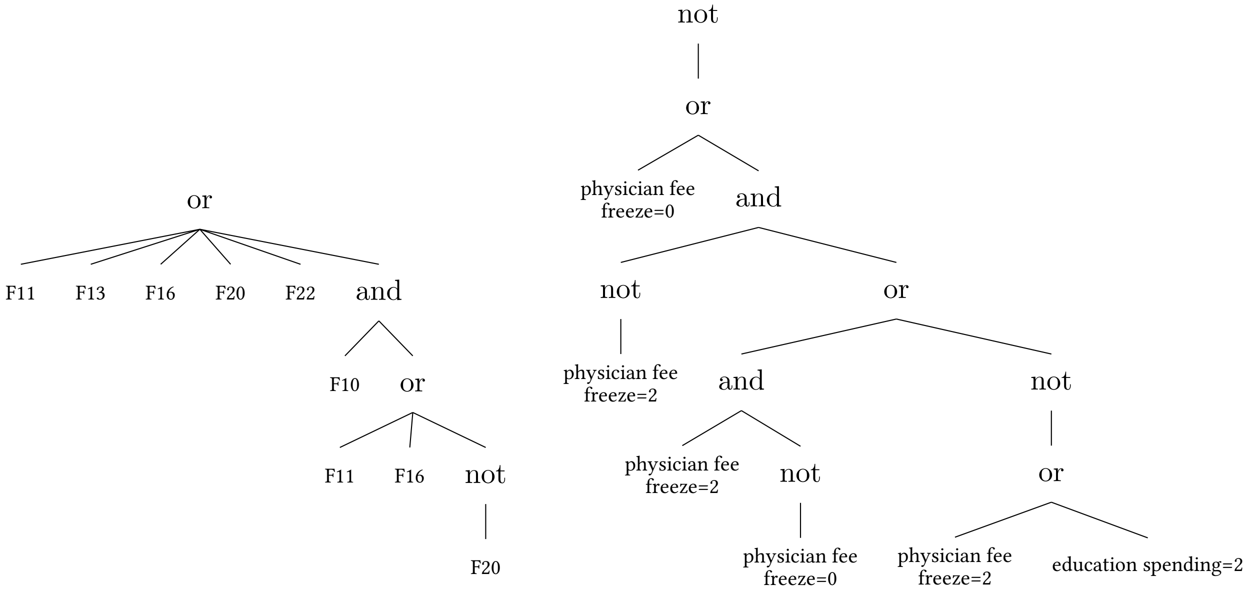
or
F11
F13
F16
F20
F22
and
F10
or
F11
F16
not
F20
not
or
physician fee freeze=0
and
not
physician fee freeze=2
or
and
physician fee freeze=2
not
physician fee freeze=0
not
or
physician fee freeze=2
education spending=2

(f) spect. F1: 0.919. LogReg: 0.930. RF: 0.909.

(g) vote. F1: 0.971. LogReg: 0.974. RF: 0.974.

Figure 13: **Some logical formulas predicted by our noisy model for some binary classification PMLB datasets.** In each case, we report the name of the dataset and the F1 score of the Boolformer, logistic regression and random forest in the caption.

F Additional results on gene regulatory network inference

In this section, we give an brief overview of the field of GRN inference and present additional results using our Boolformer.

F.1 A brief overview of GRNs

Inferring the behavior of GRNs is a central problem in computational biology, which consists in deciphering the activation or inhibition of one gene by another gene from a set of noisy observations. This task is very challenging due to the low signal-to-noise ratios recorded in biological systems, and the difficulty to obtain temporal ordering and ground truth networks.

GRN algorithms can infer relationships between the genes based on static observations [53–55], or on input time-series recordings [56, 57], and can either infer correlational relationships, i.e. undirected graphs, or causal relationships, i.e. directed graphs – the latter being more useful, but harder to obtain.

We focus on methods which model the dynamics of GRNs via Boolean networks: REVEAL [58], Best-Fit [59], MIBNI [50], GABNI [49] and ATEN [60]. We evaluate our approach on the recent benchmark from [48].

F.2 Additional results

The benchmark studied in the main text assesses both dynamical prediction (how well the model predicts the dynamics of the network) and structural prediction (how well the model predicts the Boolean functions compared to the ground truth). Structural prediction is framed as the binary classification task of predicting whether variable i influences variable j, and can hence be evaluated by several binary classification metrics, defined below[13]:

$$\mathrm{Acc} = \frac{\mathrm{TP}+\mathrm{TN}}{\mathrm{TP}+\mathrm{TN}+\mathrm{FP}+\mathrm{FN}}, \quad \mathrm{Pre} = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FP}}, \quad \mathrm{Rec} = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FN}}, \quad F1 = 2\frac{\mathrm{Pre}\cdot\mathrm{Rec}}{\mathrm{Pre}+\mathrm{Rec}},$$

$$\mathrm{MCC} = \frac{\mathrm{TP}\cdot\mathrm{TN}-\mathrm{FP}\cdot\mathrm{FN}}{\sqrt{(\mathrm{TP}+\mathrm{FP})(\mathrm{TP}+\mathrm{FN})(\mathrm{TN}+\mathrm{FP})(\mathrm{TN}+\mathrm{FN})}}, \quad \mathrm{AUROC} = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FN}} + \frac{\mathrm{TN}}{\mathrm{TN}+\mathrm{FP}} - 1.$$

We report these metrics in Fig. 14.

---

[13] The authors of the benchmark consider the two latter to be the best metrics to give a comprehensive view on the classifier performance for this task.

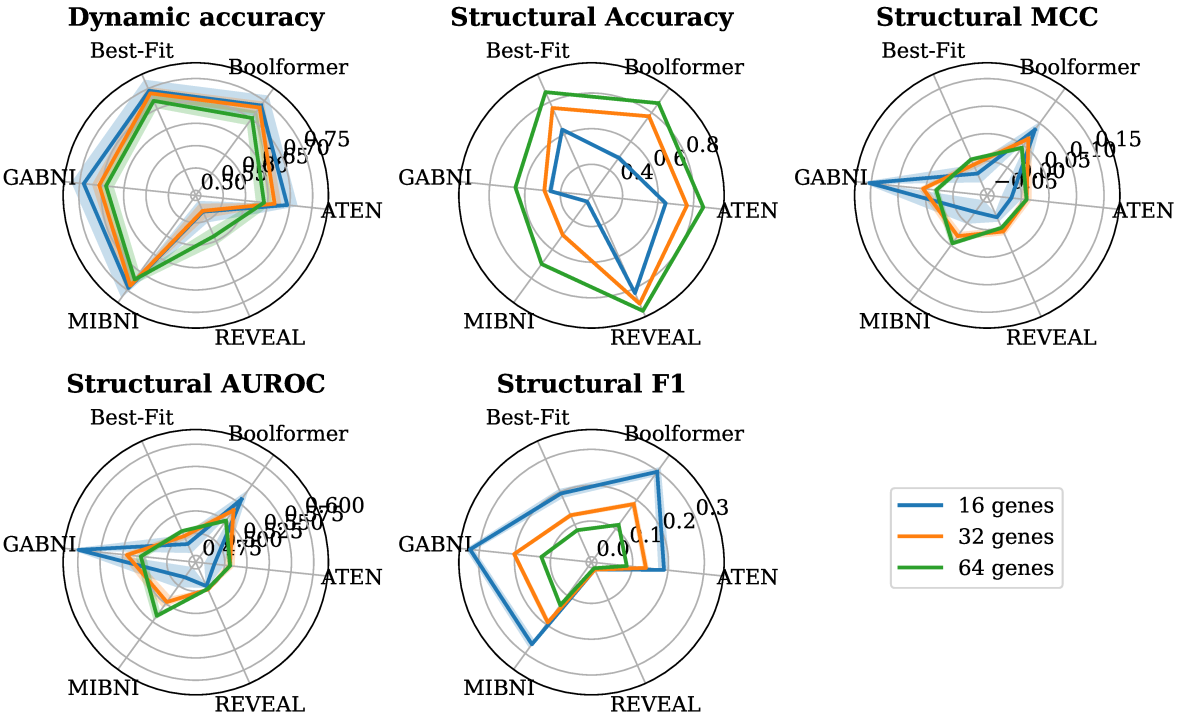

Figure 14: **Binary classification metrics used in the gene regulatory network benchmark.** The competitors and metrics are taken from the recent benchmark of [48], and described in Sec. 4.3.

G Exploring the beam candidates

In this section, we explore the beam candidates produced by the Boolformer. In Fig. 15, we show the 8 top-ranked candidates when predicting a simple logic function, the 2-comparator. We see that all candidates perfectly match the ground truth, but have different structure.

H Attention maps

In Fig. 16, we show the attention maps produced by our model when presented three truth tables: (a) that of the 4-digit multiplier, (b) that of the 4-parity function and (c) a random truth table. Each panel corresponds to a different layer and head of the model.

Recall that each of the N inputs to the transformer is the embedding of an $(\mathbf{x}, y)$ pair; in this case, we ordered the embeddings according to the binary number they form, i.e. from left to right: 0000, 0001, 0010, ..., 1111. We see highly structured patterns emerging, especially for the first two functions which are non-random.

For example, for the 4-digit multiplier, some attention heads have hadamard-like structure (e.g. head 5 of layer 6), some have block-structured checkboard patterns (e.g. head 12 of layer 4), and many heads put most attention weight on the final input, 1111, which is more informative (e.g. head 6 of layer 3).

For the parity function, we see a particularly interesting shape emerge in several heads (e.g. head 2 of layer 4), and observe that many heads perform anti-diagonal attention (e.g. head 4 of layer 6).

I Embeddings

In this section, we show that the model learns a compressed representation of the hypercube which conserves distances, both qualitatively in Fig. 17, and quantitatively in Fig. 18, where we plot the L2 distance in embedding space against the hamming distance in input space, showing a linear relationship.

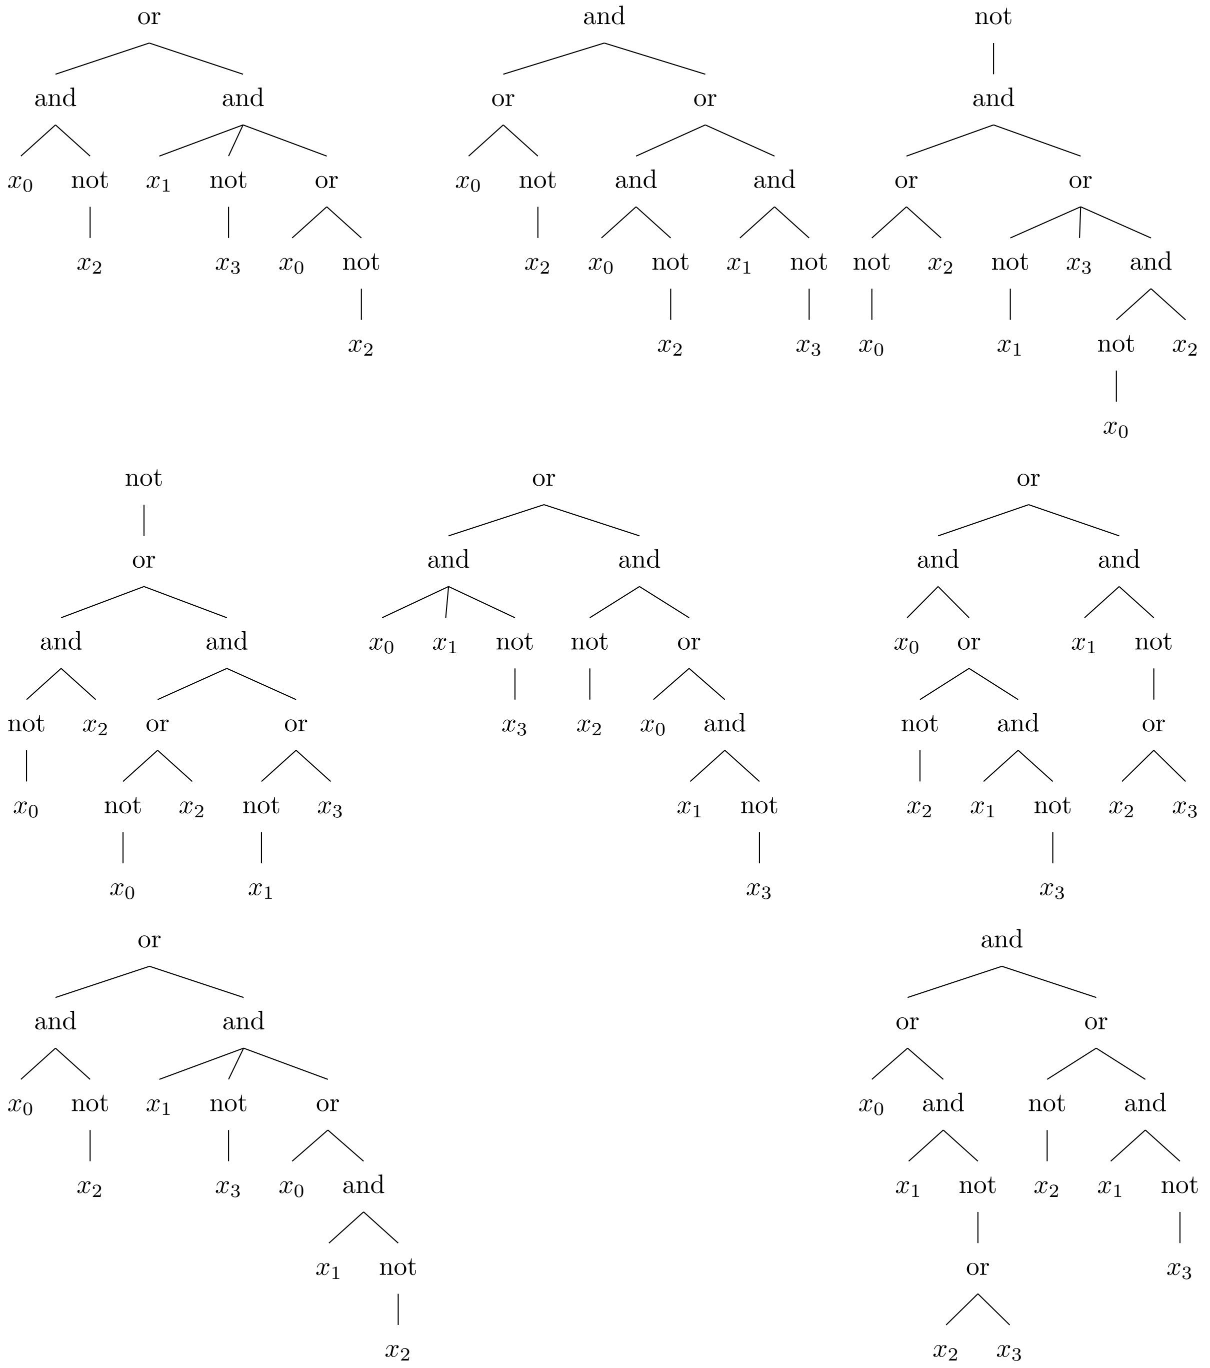
or
and and
x_0 not x_1 not or
x_2 x_3 x_0 not
x_2
and
or or
x_0 not and and
x_2 x_0 not x_1 not
x_2 x_3
not
and
or or
not x_2 not x_3 and
x_0 x_1 not x_2
x_0
not
or
and and
not x_2 or or
x_0 not x_2 not x_3
x_0 x_1
or
and and
x_0 x_1 not not or
x_3 x_2 x_0 and
x_1 not
x_3
or
and and
x_0 or x_1 not
not and or
x_2 x_1 not x_2 x_3
x_3
or
and and
x_0 not x_1 not or
x_2 x_3 x_0 and
x_1 not
x_2
and
or or
x_0 and not and
x_1 not x_2 x_1 not
or x_3
x_2 x_3

Figure 15: **Beam search reveals equivalent formulas**. We show the first 8 beam candidates for the 2-comparator, which given two 2-bit numbers $a = (x_0 x_1)$ and $b = (x_2 x_3)$, returns 1 if $a > b$, 0 otherwise. All candidates perfectly match the ground truth.

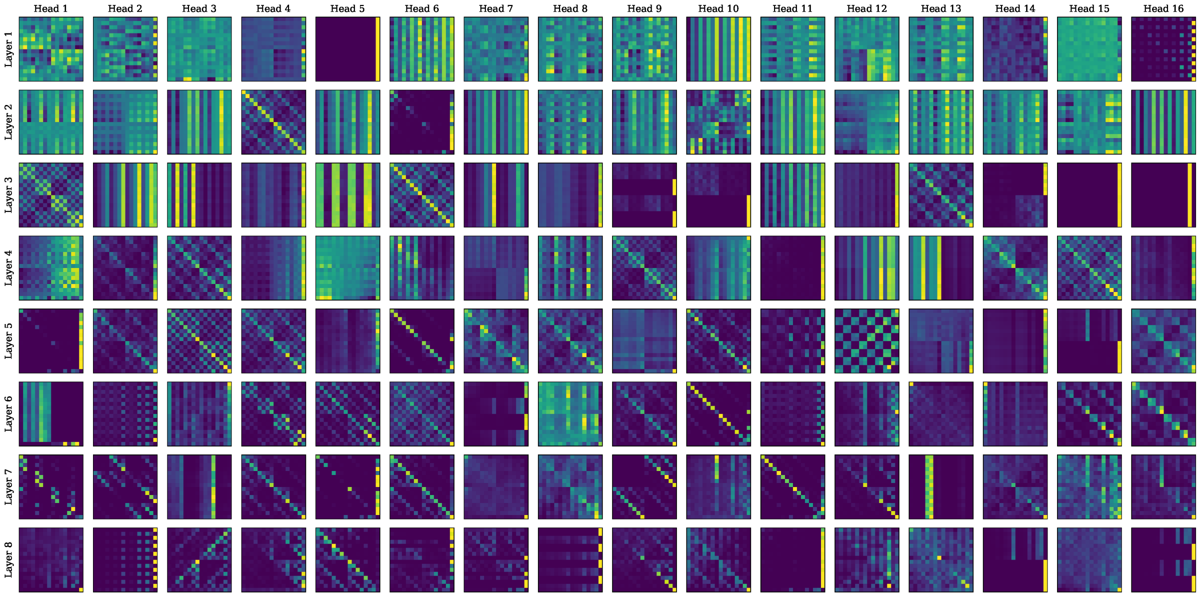

(a) 4-digit multiplier

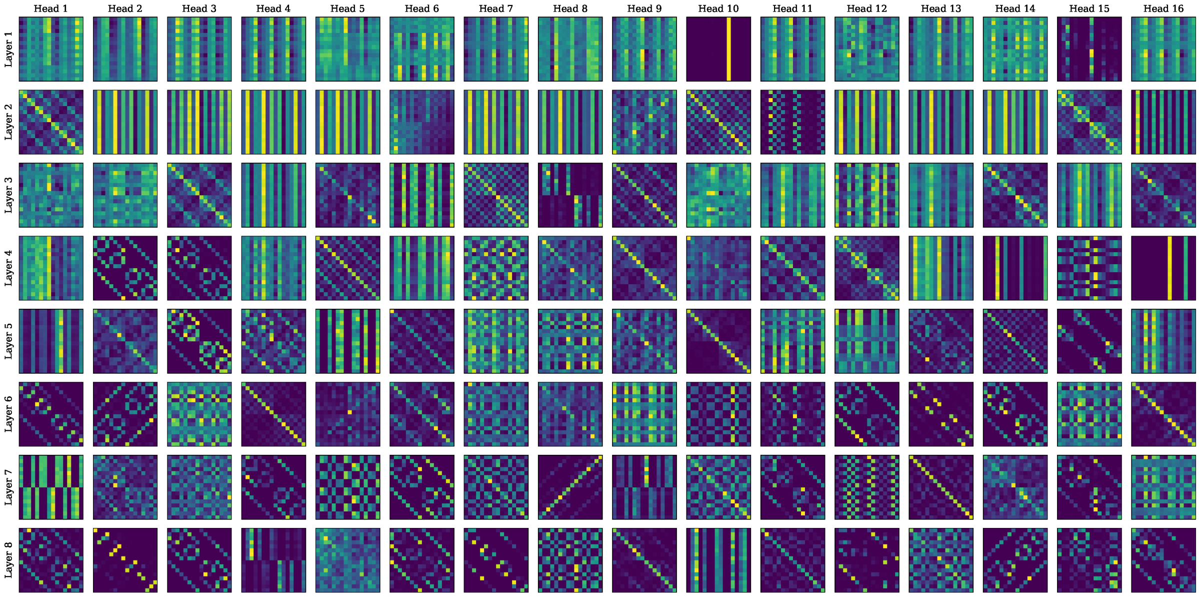

(b) 4-parity

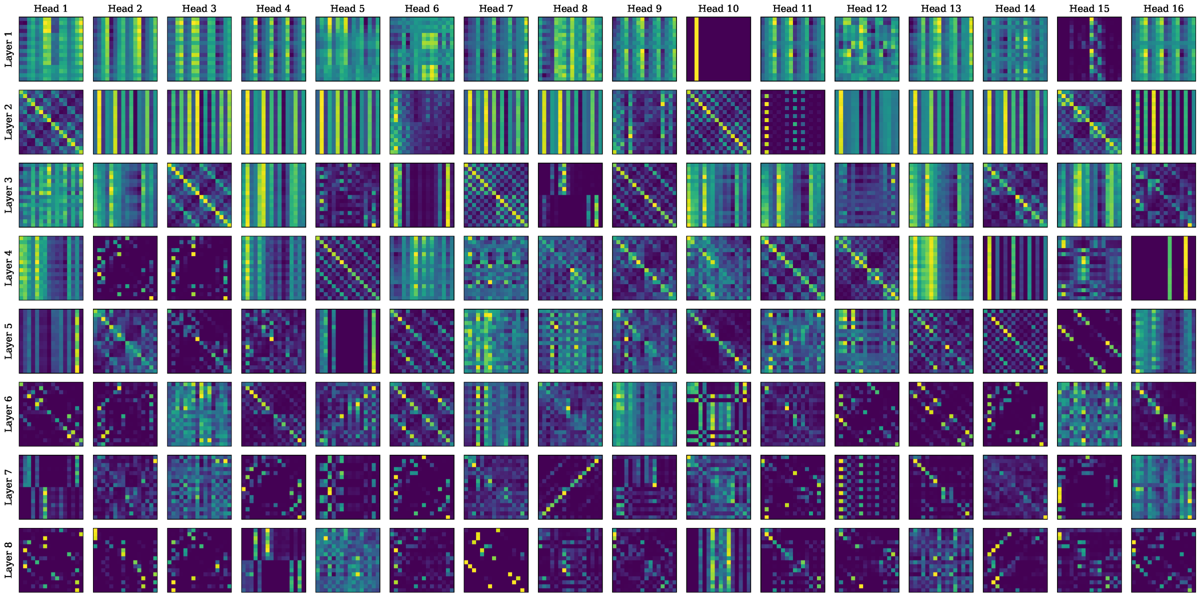

(c) 4d random data

Figure 16: **The attention maps reveal intricate analysis.** See Sec. H for more details on this figure.

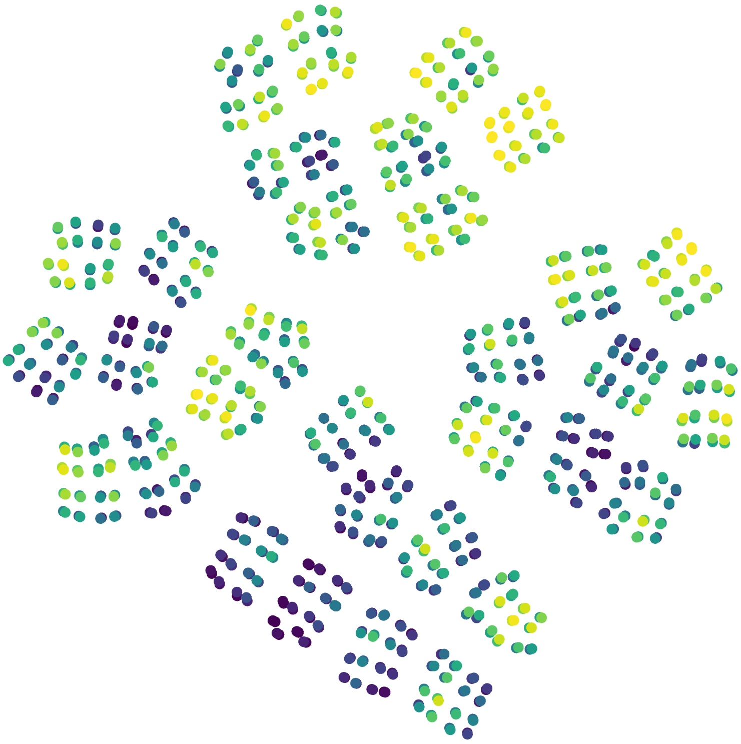

Figure 17: **T-SNE representation of the embeddings.** We fed the 1024 input combinations of the 10-dimensional hypercube to the embedder, and colored them according to the number of 1's, from 0 (blue, which corresponds to 0000000000) to 10 (yellow, which corresponds to 1111111111).

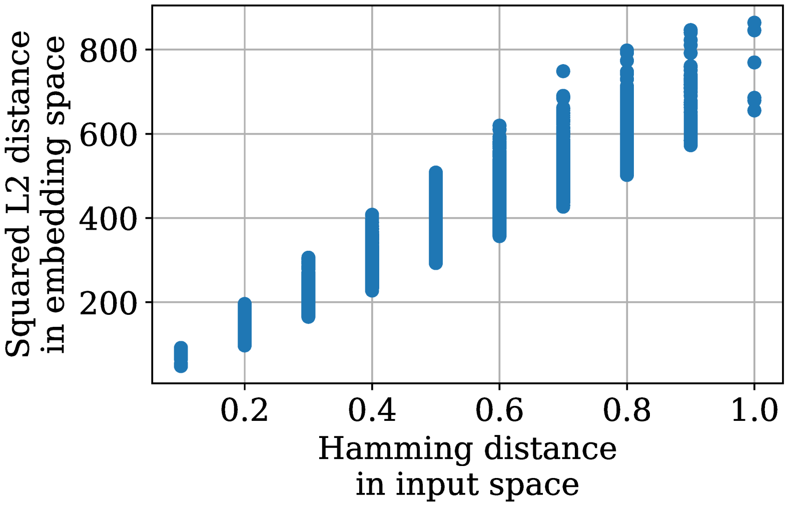

Figure 18: **The embedder conserves distances.** We plot the squared L2 distance of the embeddings of all points in the 10-dimensional hypercube against the hamming distance in the input space.