Towards distillation guarantees under algorithmic alignment

Anonymous Author(s)
Affiliation
Address
email

Abstract

Distillation is the process of condensing learnt knowledge from a large neural network trained on large datasets to a more efficient one suitable for deployment. Building on recent developments in the learning theory of distillation (Boix-Adsera, 2024), we rigorously analyze a phenomenon in which if the target class of the distillation process is algorithmically aligned with the task at hand, in terms of a linear representation hypothesis (Elhage et al., 2022), then the distillation process can be efficient. This gives rise to a novel and rigorous characterization of algorithmic alignment that could be of independent interest.

1 Introduction

An understated paradigm of modern machine learning is the incorporation of inductive biases into the learning algorithm, either through the architecture or through the optimization process. For example, the use of local shift-invariant kernels in convolutional neural networks has led to one of the most important breakthroughs of computer vision in the past century, the learning of ImageNet (Krizhevsky et al., 2012; LeCun et al., 2015; Goodfellow et al., 2016).

In recent years, as unprecedentedly large-scale computing is made possible with modern hardware, the absolute necessity of incorporating inductive bias into the model has been questioned (e.g. "transformer v.s. convolution" (Bachmann et al., 2023)). It is therefore a question of great significance whether there are inductive biases that would give tremendous benefit to be incorporated as opposed to being learned with data.

1.1 Graph machine learning

Graph machine learning is a testbed for graph-based inductive biases that may allow for exponential gains in learning efficiency. Informally, symmetry constraints of graph functions, in terms of vertex permutations, induce certain sparsity structures in the function space, making learning easier (Bietti et al., 2021; Elesedy, 2021; Tahmasebi and Jegelka, 2023). Although this is the case in specific learning settings, in general, learning graph neural networks and other equivariant networks are still hard in the worst case, requiring, for example, exponentially or superpolynomially many queries in the correlation statistical queries model of learning (Kiani et al., 2024). Understanding which settings exactly give rise to quantitative benefits for learning is an important and active area of research.

More specifically, a graph neural network (GNN) (Gilmer et al., 2017; Kipf and Welling, 2017) is a parameterization of the space of functions on graphs, potentially of different

Submitted to 39th Conference on Neural Information Processing Systems (NeurIPS 2025). Do not distribute.

sizes. A message passing neural network (MPNN) is one such example in which each node aggregates neighboring information and processes them with a neural network to form a new latent representation in each round. After a fixed number of rounds, the network outputs a learnt representation for each vertex of the graph, or combines them together to form a single representation for the whole graph, depending on the specific tasks.

1.2 Combinatorial optimization with graph ML and algorithmic alignment

One proposed area where GNNs could have strong inductive bias with the learning task is that of using neural networks to learn combinatorial optimization. It is observed (Xu et al., 2020) that the loop structure of an MPNN closely follows that of local graph algorithms, such as Bellman-Ford for shortest path. As such, Xu et al. (2020) argues that the neural network used in the aggregation operation of an MPNN only had to learn a simple function of its inputs, and not the actual for-loop structure, thereby decreasing the sample complexity of learning from supervised examples produced by such algorithms. Although the original paper provided a theoretical justification for this phenomenon through PAC learning (Valiant, 1984), a tighter analysis of what constitutes such *algorithmic alignment* has drawn many follow-up investigations (Dudzik and Veličković, 2022; Dudzik et al., 2024). Nevertheless, the idea that the learning architecture should be built to resemble a potential algorithmic paradigm, such as dynamic programming, is intuitive and has been the inspiration for many neural heuristics that are widely successful in practice (Kahng et al., 2024; Nerem et al., 2025; He and Vitercik, 2025; Gasse et al., 2019).

1.3 Contribution of this paper

The paper rigorously analyzes the advantages of employing models with high inductive bias for the right tasks. For the purpose of the workshop, we focus on the learning paradigm known as "learn first, distill later" in big data, where an enormous multipurpose network (or a foundation model) is trained on an enormous dataset of the task. Later on in production, the knowledge learnt is distilled into a more efficient models for deployment (e.g. on edge devices). Specifically, we argue that if the source class has learnt to perfectly perform some combinatorial optimization tasks on graphs of size n, if the target model is *algorithmically aligned* with an algorithm that solves this optimization (e.g. dynamic programming), in a certain sense that will be discussed, then the distillation from source to target can be efficient, in a rigorous model of learning theory known as PAC-distillation (Boix-Adsera, 2024). This is the first, as far as the authors are aware, rigorous study of distillation into graph neural network that makes use of a form of algorithmic alignment.

2 Preliminaries

In this section, we will discuss some of the tools and materials that will be used to arrive at our results.

2.1 Notations

In general, we denote by $\mathcal{X}$ the input set and by $\mathcal{Y}$ the label set. We focus on binary classification in this paper, so, unless otherwise stated, $\mathcal{Y} = \{0, 1\}$. For a vector in S with n entries, we denote by S_i its i-th entry. We will also write $[n] := \{1, 2, \ldots, n\}$.

In congruence with languages of computational learning theory, we follow the notations of Boix-Adsera (2024). We will usually denote by $\mathcal{C}$ the concept class (class of possible ground truths), $\mathcal{H}$ the hypothesis class (output range of the learning algorithm), the input distribution $\mathcal{D}$ that induces a distribution over sample $\mathcal{D}_c \in \mathcal{P}(\mathcal{X} \times \mathcal{Y})$ for some ground truth $c \in \mathcal{C}$. When there are many inputs $S \in \mathcal{X}^n$ for some $n \in \mathbb{N}$, we overload c to apply pointwise to each element of the vector $c(S) := [c(S_i)]_{i \in [n]}$. In the setting of distillatin, we have a source class $\mathcal{F}$ and a target class $\mathcal{H}$.

For graph-theoretic notation, we define $\mathcal{G}_n$ to be the space of all graphs on the n vertices. To make the exposition cleaner, we assume that all graphs are labeled graphs and drop the

subscript n if it is clear from context. Considering a Boolean input/output graph learning model, the input is both the initial feature vector (which encodes the initialization of some DP algorithm) and also the graph structure of $\binom{n}{2}$ bits. Assuming that the dimension of each node feature is fixed and independent of n, we let $d = O(n^2)$ denote the dimension of the input of the models.

For machine learning theoretic notation, we say that a neural network defines a latent representation $\varphi : \mathcal{X} \to \mathbb{R}^m$ defined as the concatenation of all activations in its neurons. Some authors also defined this representation as that of the penultimate layer of a deep network. The exact choice does not matter in this paper as long as they satisfy a structural assumption called the linear representation hypothesis, to be discussed later.

2.2 Set-up

PAC learning In traditional PAC (Valiant, 1984), a *concept class* (class of possible ground truth) is (ϵ, δ)-learnable in n samples if there is an algorithm $\mathcal{A}$ such that for any $\mathcal{D}$ distribution over the input and any concept in $c \in \mathcal{C}$,

$$\Pr_{S \sim \mathcal{D}^n}[\text{error}_{c,\mathcal{D}}(\mathcal{A}(S, c(S)) \leq \epsilon] \geq 1 - \delta. \tag{1}$$

Here, the error function is the 0-1 population risk: $\text{error}_{c,\mathcal{D}}(f) := \Pr_{\mathcal{D}}[f(x) \neq c(x)]$.

PAC-distillation (Boix-Adsera, 2024) PAC-distillation is a relaxation of PAC learning in which one assumes the accessibility of a successful model class $\mathcal{F}$ to train a target class $\mathcal{H}$ by finding an algorithm $\mathcal{A}$ such that for any distribution $\mathcal{D}$ on $\mathcal{X}$, any source $f \in \mathcal{F}$,

$$\Pr_{S \sim \mathcal{D}^n}[\text{error}_{f,\mathcal{D}}(\mathcal{A}(S, f) \leq \epsilon] \geq 1 - \delta. \tag{2}$$

Such an algorithm is said to (ϵ, δ) distill $\mathcal{F} \to \mathcal{H}$. Note that since the algorithm has access to the successful model f, giving a PAC distillation algorithm is easier than giving a PAC algorithm since one can just use f to query labels and simulate PAC. The advantage of this framework is 1) to sidestep some of the hardness results of PAC learning with relaxation and 2) to make use of extra natural structures in the class $\mathcal{F}$.

In practice, $\mathcal{F}$ can be thought of as pre-trained complex neural networks that have achieved low errors on some tasks, and the target class $\mathcal{H}$ can be understood as function classes with inductive bias that can more efficiently represent the ground truth, for example, invariant neural networks such as convolutional neural nets (CNNs) or graph neural nets (GNNs). Distillation then asks if there are efficient algorithms to find a good representation of the ground truth in the target class.

The design of $\mathcal{F}$ and $\mathcal{H}$ should be taken with great care so that the distillation of PAC is not trivial. For example, if $\mathcal{F} \subseteq \mathcal{H}$ then the learning algorithm can just return the second argument. Or, if $\mathcal{H}$ admits efficient approximations of functions in $\mathcal{F}$, then returning the approximation also solves the problem.

Linear representation hypothesis (LRH) (Elhage et al., 2022) is the main structural assumption on the source class $\mathcal{F}$. It can be stated as follows:

Definition 1 (τ-LRH)**.** *Fix a source neural network $f \in \mathcal{F}$ and let $\varphi : \mathcal{X} \to \mathbb{R}^m$ be the latent representation of f. Let $\mathcal{Z}$ be a set of functions $z : \mathcal{X} \to \mathbb{R}$. For any $\tau > 0$, we say that f satisfies τ -LRH for features $\mathcal{Z}$ if for all $z \in \mathcal{Z}$, there exists $w \in \mathbb{R}^m$ such that $\|w\| \leq \tau$ and $\langle w, \varphi(x) \rangle = z(x)$ for all $x \in \mathcal{X}$.*

In essence, Z is the set of intermediate computations or high-level features of the target class. To give a taste of the results that can be obtained from this framework, we restate a distillation theorem.

Theorem 1 (Theorem 3.6 of (Boix-Adsera, 2024))**.** *Let $\mathcal{F}$ be the set of neural networks f that implicitly compute a decision tree $T : \{0,1\}^d \to \{0,1\}$ of depth r and size s such that f satisfies τ-LRH for features $\mathcal{Z}_T := \{\text{AND}_S : S \text{ is a path starting at the root of } T\}$. Let $\mathcal{H}$ be the set of decision trees with depth r and size s. Then for any $\epsilon, \delta \in (0,1)$, there is an algorithm that (ϵ, δ)-distills from $\mathcal{F}$ to $\mathcal{H}$ that runs in polynomial time in $d, m, 1/\epsilon, s, 2^r, \log(1/\delta), \tau$ and B and takes polynomially many samples in $1/\epsilon, s, \log(d/\delta), \log(\tau B)$ where $B \geq \max_x \|\varphi(x)\|$.*

This is a remarkable result, since learning a decision tree in the PAC framework is conjectured to take $d^{\Omega(r)}$ time, but PAC distillation takes only $\text{poly}(d, 2^r)$ time.

Algorithmic alignment In this paper, we propose that a trained neural network that satisfies τ -LRH for some features Z is an example of algorithmic alignment. More general alignments, which will be defined later, can also lead to efficient distillation bounds.

3 Main results

All graphs discussed in this section are labeled and assumed to have n vertices for some $n \in \mathbb{N}$.

We first give some definitions specific to our settings:

Definition 2 (Neural networks that computes graph algorithms)**.** *A neural network $\nu : \mathcal{G} \times \mathcal{X} \to \mathcal{Y}$* computes graph algorithm *$\mathcal{A}$ if ν agrees with A on all inputs of size n. It is* efficient *if it can be evaluated in polynomial time in n.*

Definition 3 (Local-iteration algorithm)**.** *For any function $f : \{0,1\}^n \times \mathcal{G} \times [n] \to \{0,1\}$ let $\mathcal{A}^l[f] : \{0,1\}^n \times \mathcal{G} \to \{0,1\}$ be a graph input algorithm that computes:*

Algorithm 1: Local-iteration algorithm $\mathcal{A}^l[f]$

Input: Initialization vector INIT, Graph $G = (V, E)$
Output: $\{0,1\}$ classification

$h_{v,0} \leftarrow$ Initialize for each $v \in V$ with INIT
for $t \in [T]$ **do**
  for $v \in V$ **do**
    $h_{v,t} \leftarrow f(\{\{h_{u,t-1}\}\}_{u \in N(v)}, G, v)$ ▷ f `can select the neighbors using` G `and` v
return $h_{n,T}$.

In the remainder of the paper, we will assume that we are given the stopping time l *a priori*. For k-local algorithms, $l = k$, while for more general algorithms, l can depend on n. Unless otherwise stated, we focus on the former case.

We consider the following intuitive form of algorithmic alignment that was proposed in the seminal paper of Xu et al. (2020):

Definition 4 (Local-iteration alignment)**.** *Fix a source neural network $\nu \in \mathcal{F}$ and let $\varphi : \mathcal{X} \to \mathbb{R}^m$ be the latent representation of ν. Let Z be a set of functions $z : \{0,1\}^n \times \mathcal{G} \to \{0,1\}$. For any $\tau > 0$, we say that ν satisfies τ-local-iteration alignment for features Z if for all $z \in Z$, there exists a $w \in \mathbb{R}^m$ with $\|w\| \leq \tau$ and $\langle w, \varphi \rangle = \mathcal{A}^l[z]$.*

Finally, we define a decision tree:

Definition 5. *A decision tree $T : \{0,1\}^d \to \{0,1\}$ is a labeled rooted binary tree with leaf labeled 0 or 1 and internal node labeled by its input variables $x_1 \dots x_d$. Each input takes a path by evaluating the input variable to arrive at the leaf that is the output of the tree.*

3.1 Concept class, source class and target class

Our concept class (the collection of possible ground truths) will be:

$$\mathcal{C}_{s,r} = \{\mathcal{A}[T] \mid \ T \text{ is a decision tree with depth } s \text{ size } r\} \tag{3}$$

To define the source class, we first need to give the set of features that is supposedly linearly represented by our source functions. This is done analogously to (Boix-Adsera, 2024): given a decision tree T, the root-prefix path conjunctions form T's features:

$$Z'_T := \left\{ \bigwedge_{i=1}^{l} p_i \mid p_i \text{ is a literal } (x_j \text{ or } \neg x_j \text{ for some } j) \text{ s.t. } (p_1, \dots, p_l) \text{ is a path from root} \right\} \tag{4}$$

Given these features for decision trees, the features for local-iteration algorithms are:

$$Z_T := \left\{A^T[b] \mid b \in Z'_T\right\} \tag{5}$$

We postulate that such loops over prefix path conjunctions are simply representable by the neural network's latent representation.

The source class (the collection of neural networks that have successfully learnt some graph algorithms) will be:

$$\mathcal{F}^{\tau}_{s,r} = \{\text{neural networks implicitly computing } \mathcal{A}[T] \text{ for some decision tree } T \\ \text{that also satisfied } \tau\text{-local-iteration alignment for features } Z_T\}$$

And finally, the target class is the same as the concept class. In practice, this is a subset of graph neural networks, and the distillation process can be understood as distilling from learnt neural networks to graph neural networks.

3.2 GNNs are more efficient than decision trees

In the following, we state a simple fact that makes for a good exercise:

Lemma 1. *There exists a simple decision tree $T : \{0,1\}^n \times \{0,1\}^{\binom{n}{2}} \to \{0,1\}$ that can be evaluated in polynomial time in n such that while $\mathcal{A}[T]$ can be evaluated in polyonomial time, it cannot be represented by decision trees of polynomial size.*

Proof sketch. Consider the 2-reachability DP. In other words, given the adjacency matrix of a graph on $n \geq 2$ vertices, is there a path of length at most 2 that connects the vertex labeled 1 and n? The full proof is in the Appendix. □

This fact means that without the for loop structure, one cannot just convert the concept class $\mathcal{C}_{s,r}$ into the class of efficient decision trees. This forms a type of algorithmic mismatch which would be resolved in the next part, using graph neural networks.

3.3 GNNs can be distilled from learnt and aligned neural networks efficiently

We are now ready to state the main theorem.

Theorem 2. *For any $\epsilon, \delta \in (0,1)$, there is an algorithm that (ϵ,δ)-distills from $\mathcal{F}^{\tau}_{s,r}$ to $\mathcal{C}_{s,r}$ and runs in polynomial time in $n, m, 1/\epsilon, s, 2^{r^l}, \log(1/\delta), \tau, B$ and takes a polynomial sample in $1/\epsilon, s, \log(d/\delta), \log(\tau B)$ from $\mathcal{D}$ where $B = \max_x \|\varphi(x)\|$.*

The proof follows from (Boix-Adsera, 2024) and uses Algorithm 2, which first builds a set of conjunctions that is a superset of all root-prefix conjunctions in the true tree and then stitches these conjunctions together efficiently using a DP.

Algorithm 2: GNN distillation algorithm

Input: Neural network ν, representation φ, random samples from $\mathcal{D}$, depth bound $R \in \mathbb{N}$, error parameters $\epsilon, \delta > 0$.

Output: GNN that computes $\mathcal{A}^l[\widehat{T}]$

$\mathcal{S}_0 \leftarrow \{\emptyset\}$ **for** $i = 1 \in [R]$ **do**

  $\mathcal{P}_{i-1} \leftarrow \left\{S \in \mathcal{S}_{i-1} \text{ s.t } \textsc{LinearProbe}(\mathcal{A}^l[\bigwedge_{p_i \in S} p_i], \varphi, B, \tau, 2^{-i^l-3}, \frac{\delta}{2|\mathcal{S}_{i-1}|R}) = \text{true}\right\}$

  $\mathcal{S}_i \leftarrow \bigcup_{S \in \mathcal{P}_{i1}} \bigcup_{j=1}^{d} \{S \cup x_l, S \cup \neg x_l\}$

$\mathcal{S} \leftarrow \bigcup_{j=1}^{R} \mathcal{S}_j$

$\hat{v}_{S'} \leftarrow \mathbb{E}_x[\bigwedge_{p_i \in S'} p_i(x)(2\nu(x)-1)] \pm \epsilon/s$, for each $S' \in \mathcal{S}^l$

return $\text{argmax}_T \text{val}(T, \hat{v})$ where T is over decision trees with $Z'_T \subseteq \{\bigwedge_{p_i \in S} p_i \mid S \in \mathcal{S}\}$

In Algorithm 2,

1. The LinearProbe$(g, \varphi, B, \tau, \epsilon, \delta, \mathcal{D})$ subroutine comes from Lemma 3.7 of (Boix-Adsera, 2024) that runs in polynomial time and draws polynomially many samples to return true w.p. $\geq 1-\delta$ if there is a $w \in \mathbb{R}^m$ with $\|w\| \leq \tau$ and $\mathbb{E}_x[(\langle w, \varphi \rangle - g)^2] \leq 2$ and return false w.p. $\geq 1-\delta$ if for all such w, the expectation is at least 2ϵ.
2. The maximization over decision trees in the final step is done with a DP (Guijarro et al., 1999; Mehta and Raghavan, 2002) and the function val is chosen such that maximizing it corresponds to maximizing the 0-1 risk.

We defer the correctness proof Theorem 2 to the Appendix.

4 Conclusion and discussion

In this workshop paper, we extend the work of Boix-Adsera (2024) in PAC-distillation to characterize learning models that have built-in algorithmic alignment properties, such as GNNs for dynamic programming. We showed that while there are some DP algorithms whose inner function is a small decision tree, the whole DP itself cannot be represented by an efficient decision tree – a case of misalignment. On the other hand, the local iteration structure of a GNN with decision tree aggregation allows for efficient distillation from a large, learnt neural network that exhibits a certain kind of linear representability.

Although this marks the first work in this direction, there is much room for improvement with this workshop version of the paper. For brevity of exposition, we adapt the analysis of (Boix-Adsera, 2024) as is and naively, under the assumption that the outer loop of the DP has a constant number of iterations and that the size of the innermost decision tree is small. A more complicated and nuanced analysis greatly reduces these requirements.

Finally, the purpose of this workshop paper is to introduce the distillation pipeline, which has shown great promise in rigorously studying algorithmic alignment. Besides a more detailed analysis, future directions include the statistical learning theory questions of sample complexity under distillation or the study of pretrained large language models, whose finetuning might be thought of as a form of distillation.

References

Gregor Bachmann, Sotiris Anagnostidis, and Thomas Hofmann. Scaling MLPs: A tale of inductive bias. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=R45A8eKcax.

Alberto Bietti, Luca Venturi, and Joan Bruna. On the sample complexity of learning under invariance and geometric stability. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS '21, Red Hook, NY, USA, 2021. Curran Associates Inc. ISBN 9781713845393.

Enric Boix-Adsera. Towards a theory of model distillation, 2024. URL https://arxiv.org/abs/2403.09053.

Andrew J Dudzik and Petar Veličković. Graph neural networks are dynamic programmers. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 20635–20647. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/8248b1ded388fcdbbd121bcdfea3068c-Paper-Conference.pdf.

Andrew Joseph Dudzik, Tamara von Glehn, Razvan Pascanu, and Petar Veličković. Asynchronous algorithmic alignment with cocycles. In Soledad Villar and Benjamin Chamberlain, editors, *Proceedings of the Second Learning on Graphs Conference*, volume 231 of *Proceedings of Machine Learning Research*, pages 3:1–3:17. PMLR, 27–30 Nov 2024. URL https://proceedings.mlr.press/v231/dudzik24a.html.

Bryn Elesedy. Provably strict generalisation benefit for invariance in kernel methods. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS '21, Red Hook, NY, USA, 2021. Curran Associates Inc. ISBN 9781713845393.

Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish, Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. Toy Models of Superposition, September 2022. URL http://arxiv.org/abs/2209.10652. arXiv:2209.10652 [cs].

Maxime Gasse, Didier Chételat, Nicola Ferroni, Laurent Charlin, and Andrea Lodi. *Exact combinatorial optimization with graph convolutional neural networks*. Curran Associates Inc., Red Hook, NY, USA, 2019.

Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, page 1263–1272. JMLR.org, 2017.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. Deep learning (book). In *MIT Press*, 2016.

David Guijarro, Víctor Lavín, and Vijay Raghavan. Exact learning when irrelevant variables abound. *Inf. Process. Lett.*, 70(5):233–239, June 1999. ISSN 0020-0190. doi: 10.1016/S0020-0190(99)00063-0. URL https://doi.org/10.1016/S0020-0190(99)00063-0.

Yu He and Ellen Vitercik. Primal-dual neural algorithmic reasoning. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=iBpkzB5LEr.

Andrew B. Kahng, Robert R. Nerem, Yusu Wang, and Chien-Yi Yang. Nn-steiner: a mixed neural-algorithmic approach for the rectilinear steiner minimum tree problem. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI'24/IAAI'24/EAAI'24. AAAI Press, 2024. ISBN 978-1-57735-887-9. doi: 10.1609/aaai.v38i12.29200. URL https://doi.org/10.1609/aaai.v38i12.29200.

Bobak Kiani, Thien Le, Hannah Lawrence, Stefanie Jegelka, and Melanie Weber. On the hardness of learning under symmetries. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=ARPrtuzAnQ.

Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=SJU4ayYgl.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 25. Curran Associates, Inc., 2012.

Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553): 436–444, 2015.

Dinesh P. Mehta and Vijay Raghavan. Decision tree approximations of boolean functions. *Theor. Comput. Sci.*, 270(1-2):609–623, 2002. URL https://doi.org/10.1016/S0304-3975(01)00011-1.

Robert R. Nerem, Samantha Chen, Sanjoy Dasgupta, and Yusu Wang. Graph neural networks extrapolate out-of-distribution for shortest paths, 2025. URL https://arxiv.org/abs/2503.19173.

Behrooz Tahmasebi and Stefanie Jegelka. The exact sample complexity gain from invariances for kernel regression. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=6iouUxI45W.

Robert Tarjan. Depth-first search and linear graph algorithms. In *12th Annual Symposium on Switching and Automata Theory (swat 1971)*, pages 114–121, 1971. doi: 10.1109/SWAT.1971.10.

L. G. Valiant. A theory of the learnable. In *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, STOC '84, page 436–445, New York, NY, USA, 1984. Association for Computing Machinery. ISBN 0897911334. doi: 10.1145/800057.808710. URL https://doi.org/10.1145/800057.808710.

Keyulu Xu, Jingling Li, Mozhi Zhang, Simon S. Du, Ken ichi Kawarabayashi, and Stefanie Jegelka. What can neural networks reason about? In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=rJxbJeHFPS.

A Appendix: Omitted proofs

A.1 Proof for Lemma 1

Proof. Consider the combinatorial problem of deciding, for a labeled graph, whether the first and last vertex is connected with a path of length at most 2, or 2-reachability.

The classic dynamic programming (DP) algorithm for this problem runs in time $O(n)$ (Tarjan, 1971).

However, any decision tree that correctly solves this problem on all labeled graphs of size n must have exponential size. To see this, we bound the number of leaves of a correct tree (which in turn bounds the order of its size since a decision tree is binary).

Consider the subset of graphs on the n vertices labeled by $[n]$ where the only possible edges are $(1, n)$ and $(1, v)$, (v, n) for all $v \in V \backslash \{1, n\}$. There are $2^{2(n-2)+1}$ such graphs. Among them, graphs that fail to have a path of size at most 2 between 1 and n does not have the $(1, n)$ edge and for each other v, have one of the 3 configurations out of 4 possible choices of presence/absence of the pair $(1, v)$, (v, n). This counts to $(3/4)^{n-2}/2$ fraction of the total number of graphs.

Now, each 0-leaf (leaf that outputs 0 for the DT) of a correct DT on these inputs fixes a certain presence/absence of some edges on the path from the DT's root to it. Once certain variables are fixed, all other variables are free to range between 0 and 1 and the output of the DT is still 0. This means that $(1, n)$ must always be included in the fixed variables, and so is at least one in each pair $(1, v)$, (v, n). Thus, each 0-leaf accounts for at most a fraction of $2^{-(n-1)}$ of the total number of graphs.

Therefore, the number of leaves must be at least $(3/4)^{n-2}/2/2^{-(n-1)}$, which is exponential in n. □

A.2 Proof for Theorem 2

We first set up additional notation, in line with the setup in (Boix-Adsera, 2024).

Recall that the input space is $\mathcal{X} = \{0, 1\}^d$. Some bits of the input are from the initial feature vector, while other bits are used to encode the graph structure. Literals are of the form x_i or $\neg x_i$ for some $i \in [d]$. A clause S consists of literals $p_1, \ldots, p_{|S|}$ and we define $\mathrm{AND}_S(x) := \bigwedge_{p \in S} p$. S is a non-degenerate k-clause if $|S| = k$ and each variable appears at most once in S (otherwise, AND_S will always be false).

Recall that given a decision tree T, we defined Z'_T as the collection of all AND_S functions where S range over all paths that start at the root (root-prefix paths), including the trivial path of length 0 that always evaluates to 1; and Z_T as the collection of all $\mathcal{A}^l[\mathrm{AND}_S]$ functions.

We will show that with probability at least $1 - \delta/2$,

1. For any root-prefix paths S in the true tree where $|S| \leq R$, $S \in \mathcal{S}$
2. $|S_i| \leq \mathrm{poly}(2^{\Theta(i^l)}, \tau, B, d)$

The first statement follows from Lemma 3.7 of (Boix-Adsera, 2024) because it is guaranteed that all root-prefix paths S of length at most R of the true tree have their $\mathcal{A}^l[\mathrm{AND}_S]$ checked with the LinearProbe subroutine and thus accepted into $\mathcal{S}$ with high probability. Assume to the contrary that there is a path $(p_1, \ldots, p_s)$ that was not checked by LinearProbe, this means that the prefix $(p_1, \ldots, p_{s-1})$ was also not checked or was checked but failed. The latter is not possible because we assumed that the source class satisfies τ-local-iteration-alignment, so the prefix was simply not checked. Inductively, this means that $\emptyset \ni \mathcal{S}_0$ is not checked, which is impossible and hence a contradiction.

The second statement follows from Lemma 3.8 in (Boix-Adsera, 2024) applied to a collection of k^l-clauses. This reduction is possible because unrolling the for loops of $\mathcal{A}^l[\mathrm{AND}_S]$ gives a conjunction of size k^T. This can be seen by writing the AND_S as a multilinear boolean

polynomial and noticing that the loop over vertices in V is done in parallel, while degrees are only added in the outermost loop over T runs. This is a very rough treatment of the feature set, since in fact most of the k^l-clauses are degenerate due to unrolling. A more careful analysis would bring down the final complexity even more, but for the purpose of the extended abstract, we opt to present this weaker result and delay the strongest bounds to the full paper.

Finally, it is left to show that the dynamic programming algorithm in the final step is correct in stitching up the innermost decision tree.

For a fixed candidate decision tree $\tilde{T}$, define $\mathcal{T}[\tilde{T}]$ as the decision tree equivalence of $\mathcal{A}[T]$ created by unrolling $\mathcal{A}[T]$.

Denote by $\text{Leaves}(\mathcal{T}[\tilde{T}])$ the set of all leaf clauses of $\mathcal{T}[\tilde{T}]$ and $\mathcal{T}[\tilde{T}](S)$ the output at $S \in \text{Leaves}(\mathcal{T}[\tilde{T}])$. Note that if $\text{Leaves}(\tilde{T}) \in \mathcal{S}$ then $\text{Leaves}(\mathcal{T}[\tilde{T}]) \in \mathcal{S}^l$. Define $v_S = \mathbb{E}_x[\text{AND}_S(x)(2\nu(x) - 1)]$ for some $S \in \mathcal{S}^l$. Thus,

$$\text{val}(\tilde{T}, v) := \sum_{S \in \text{Leaves}(\mathcal{T}[\tilde{T}])} v_S(2\mathcal{T}[\tilde{T}](S) - 1) \tag{6}$$

$$= \mathbb{E}_x \left[\sum_{S \in \text{Leaves}(\mathcal{T}[\tilde{T}])} \text{AND}_S(x)(2\nu(x) - 1)(2\mathcal{T}[\tilde{T}](S) - 1) \right]. \tag{7}$$

Because for each input x, there is a unique path through $\mathcal{T}[\tilde{T}]$, there is a unique $S \in \text{Leaves}(\mathcal{T}[\tilde{T}])$ such that $\text{AND}_S(x) = 1$ and for the remaining S, $\text{AND}_S(x) = 0$. Furthermore, when $\text{AND}_S(x) = 1$, $\mathcal{T}[\tilde{T}](S) = \mathcal{T}[\tilde{T}](x) = \mathcal{A}^l[\tilde{T}](x)$. Therefore:

$$\text{val}(\tilde{T}, v) = 2 \Pr_x[\mathcal{A}^l[\tilde{T}](x) = \nu(x)] - 1. \tag{8}$$

Therefore, maximizing val over $\tilde{T}$ is equivalent to maximizing the 0-1 risk of $\mathcal{A}^l[\tilde{T}]$.

In our algorithm, we use Hoeffding inequality to approximate val with random sampling. Note that this step requires approximating $|\mathcal{S}|^l$ entries of v naively and runs in time $\text{poly}(|\mathcal{S}|^l, dl, m)$ where $m = \text{poly}(1/\epsilon, \log(|S|^l/\delta))$ is the number of draws to obtain the Hoeffding bound. When choosing $R = r$, the size of the true innermost decision tree, the number of leaves of $\mathcal{T}[T]$ is of order $2^{O(lr^l)}$ based on the previous bound on $|\mathcal{S}|$. Finally, we can run the dynamic program that computes for each $S \in \mathcal{S}$, and each tree size $s' = 0..s$, the best subtree of size s' rooted at the end of the clause S.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The paper is about proving a distillation results for graph learning architecture under some algorithmic alignment assumption.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The discussion part of the paper highlights limitations and future directions.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Sketch proofs are provided with enough details so as to not repeat existing works extensively.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer:[NA]

Justification: There are no experiments in this paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [NA]

   Justification: There are no data or code in this paper.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [NA]

   Justification: There are no experiments in this paper.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [NA]

   Justification: There are no experiments in this paper.

   Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [NA]

   Justification: There are no experiments in this paper.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: Yes

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [NA]

Justification: The paper is a mathematical and theoretical study in the theory of computation and does not carry extra societal impacts that are worth highlighting.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: No such risk

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All original papers are properly cited.

Guidelines:

- The answer NA means that the paper does not use existing assets.

- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [NA]

    Justification: No new assets

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: No crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: No such risks

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: LLM are used only for editing

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.