

Learning Binary Decision Trees by Argmin Differentiation

Valentina Zantedeschi^{1,2} Matt J. Kusner² Vlad Niculae³

Abstract

We address the problem of learning binary decision trees that partition data for some downstream task. We propose to learn discrete parameters (i.e., for tree traversals and node pruning) and continuous parameters (i.e., for tree split functions and prediction functions) *simultaneously* using argmin differentiation. We do so by sparsely relaxing a mixed-integer program for the discrete parameters, to allow gradients to pass through the program to continuous parameters. We derive customized algorithms to efficiently compute the forward and backward passes. This means that our tree learning procedure can be used as an (implicit) layer in arbitrary deep networks, and can be optimized with arbitrary loss functions. We demonstrate that our approach produces binary trees that are competitive with existing single tree and ensemble approaches, in both supervised and unsupervised settings. Further, apart from greedy approaches (which do not have competitive accuracies), our method is faster to train than all other tree-learning baselines we compare with. The code for reproducing the results is available at <https://github.com/vzantedeschi/LatentTrees>.

1. Introduction

Learning discrete structures from unstructured data is extremely useful for a wide variety of real-world problems (Gilmer et al., 2017; Kool et al., 2018; Yang et al., 2018). One of the most computationally-efficient, easily-visualizable discrete structures that are able to represent complex functions are *binary trees*. For this reason, there has been a massive research effort on how to learn such binary trees since the early days of machine learning (Payne

¹Inria, Lille - Nord Europe research centre ²University College London, Centre for Artificial Intelligence ³Informatics Institute, University of Amsterdam. Correspondence to: V Zantedeschi <vzantedeschi@gmail.com>, M Kusner <m.kusner@ucl.ac.uk>, V Niculae <v.niculae@uva.nl>.




Figure 1: A schematic of the approach.

& Meisel, 1977; Breiman et al., 1984; Bennett, 1992; Bennett & Blue, 1996). Learning binary trees has historically been done in one of three ways. The first is via *greedy optimization*, which includes popular decision-tree methods such as classification and regression trees (CART, Breiman et al., 1984) and ID3 trees (Quinlan, 1986), among many others. These methods optimize a splitting criterion for each tree node, based on the data routed to it. The second set of approaches are based on *probabilistic relaxations* (Írsoy et al., 2012; Yang et al., 2018; Hazimeh et al., 2020), optimizing all splitting parameters at once via gradient-based methods, by relaxing hard branching decisions into branching probabilities. The third approach optimizes trees using *mixed-integer programming* (MIP, Bennett, 1992; Bennett & Blue, 1996). This jointly optimizes all continuous and discrete parameters to find globally-optimal trees.¹

Each of these approaches has their shortcomings. First, greedy optimization is generally suboptimal: tree splitting criteria are even intentionally crafted to be different than the global tree loss, as the global loss may not encourage tree growth (Breiman et al., 1984). Second, probabilistic relaxations: (a) are rarely sparse, so inputs contribute to all branches even if they are projected to hard splits after train-

¹Here we focus on learning single trees instead of tree ensembles; our work easily extends to ensembles.

ing; (b) they do not have principled ways to prune trees, as the distribution over pruned trees is often intractable. Third, MIP approaches, while optimal, are only computationally tractable on training datasets with thousands of inputs (Bertsimas & Dunn, 2017), and do not have well-understood out-of-sample generalization guarantees.

In this paper we present an approach to binary tree learning based on sparse relaxation and argmin differentiation (as depicted in Figure 1). Our main insight is that by quadratically relaxing an MIP that learns the discrete parameters of the tree (input traversal and node pruning), we can differentiate through it to simultaneously learn the continuous parameters of splitting decisions. This allows us to leverage the superior generalization capabilities of stochastic gradient optimization to learn decision splits, and gives a principled approach to learning tree pruning. Further, we derive customized algorithms to compute the forward and backward passes through this program that are much more efficient than generic approaches (Amos & Kolter, 2017; Agrawal et al., 2019a). The resulting learning method can learn trees with any differentiable splitting functions to minimize any differentiable loss, both tailored to the data at hand.

Notation. We denote scalars, vectors, and sets, as x , \mathbf{x} , and \mathcal{X} , respectively. A binary tree is a set \mathcal{T} containing nodes $t \in \mathcal{T}$ with the additional structure that each node has at most one left child and at most one right child, and each node except the unique root has exactly one parent. We denote the l_2 norm of a vector by $\|\mathbf{x}\| := (\sum_i x_i^2)^{\frac{1}{2}}$. The unique projection of point $\mathbf{x} \in \mathbb{R}^d$ onto a convex set $C \subset \mathbb{R}^d$ is $\text{Proj}_C(\mathbf{x}) := \arg \min_{\mathbf{y} \in C} \|\mathbf{y} - \mathbf{x}\|$. In particular, projection onto an interval is given by $\text{Proj}_{[a,b]}(x) = \max(a, \min(b, x))$.

2. Related Work

The paradigm of *binary tree learning* has the goal of finding a tree that iteratively splits data into meaningful, informative subgroups, guided by some criterion. Binary tree learning appears in a wide variety of problem settings across machine learning. We briefly review work in two learning settings where latent tree learning plays a key role: 1. *Classification/Regression*; and 2. *Hierarchical clustering*. Due to the generality of our setup (tree learning with arbitrary split functions, pruning, and downstream objective), our approach can be used to learn trees in any of these settings. Finally, we detail how parts of our algorithm are inspired by recent work in isotonic regression.

Classification/Regression. Decision trees for classification and regression have a storied history, with early popular methods that include classification and regression trees (CART, Breiman et al., 1984), ID3 (Quinlan, 1986), and

C4.5 (Quinlan, 1993). While powerful, these methods are greedy: they sequentially identify ‘best’ splits as those which optimize a split-specific score (often different from the global objective). As such, learned trees are likely sub-optimal for the classification/regression task at hand. To address this, Carreira-Perpinán & Tavallali (2018) proposes an alternating algorithm for refining the structure and decisions of a tree so that it is smaller and with reduced error, however still sub-optimal. Another approach is to probabilistically relax the discrete splitting decisions of the tree (İrsoy et al., 2012; Yang et al., 2018; Tanno et al., 2019). This allows the (relaxed) tree to be optimized *w.r.t.* the overall objective using gradient-based techniques, with known generalization benefits (Hardt et al., 2016; Hoffer et al., 2017). Variations on this approach aim at learning tree ensembles termed ‘decision forests’ (Kontschieder et al., 2015; Lay et al., 2018; Popov et al., 2020; Hazimeh et al., 2020). The downside of the probabilistic relaxation approach is that there is no principled way to prune these trees as inputs pass through all nodes of the tree with some probability. A recent line of work has explored mixed-integer program (MIP) formulations for learning decision trees. Motivated by the billion factor speed-up in MIP in the last 25 years, Rudin & Ertekin (2018) proposed a mathematical programming approach for learning provably optimal decision lists (one-sided decision trees; Letham et al., 2015). This resulted in a line of recent follow-up works extending the problem to binary decision trees (Hu et al., 2019; Lin et al., 2020) by adapting the efficient discrete optimization algorithm (CORELS, Angelino et al., 2017). Related to this line of research, Bertsimas & Dunn (2017) and its follow-up works (Günlük et al., 2021; Aghaei et al., 2019; Verwer & Zhang, 2019; Aghaei et al., 2020) phrased the objective of CART as an MIP that could be solved exactly. Even given this consistent speed-up all these methods are only practical on datasets with at most thousands of inputs (Bertsimas & Dunn, 2017) and with non-continuous features. Recently, Zhu et al. (2020) addressed these tractability concerns principally with a data selection mechanism that preserves information. Still, the out-of-sample generalizability of MIP approaches is not well-studied, unlike stochastic gradient descent learning.

Hierarchical clustering. Compared to standard flat clustering, hierarchical clustering provides a structured organization of unlabeled data in the form of a tree. To learn such a clustering the vast majority of methods are greedy and work in one of two ways: 1. *Agglomerative*: a ‘bottom-up’ approach that starts each input in its own cluster and iteratively merges clusters; and 2. *Divisive*: a ‘top-down’ approach that starts with one cluster and recursively splits clusters (Zhang et al., 1997; Widjantoro et al., 2002; Krishnamurthy et al., 2012; Dasgupta, 2016; Kobren et al., 2017; Moseley & Wang, 2017). These methods suffer from similar issues as greedy approaches for classification/regression

do: they may be sub-optimal for optimizing the overall tree. Further they are often computationally-expensive due to their sequential nature. Inspired by approaches for classification/regression, recent work has designed probabilistic relaxations for learning hierarchical clusterings via gradient-based methods (Monath et al., 2019).

Our work takes inspiration from both the MIP-based and gradient-based approaches. Specifically, we frame learning the discrete tree parameters as an MIP, which we sparsely relax to allow continuous parameters to be optimized by argmin differentiation methods.

Argmin differentiation. Solving an optimization problem as a differentiable module within a parent problem tackled with gradient-based optimization methods is known as argmin differentiation, an instance of bi-level optimization (Colson et al., 2007; Gould et al., 2016). This situation arises in as diverse scenarios as hyperparameter optimization (Pedregosa, 2016), meta-learning (Rajeswaran et al., 2019), or structured prediction (Stoyanov et al., 2011; Domke, 2013; Niculae et al., 2018a). General algorithms for quadratic (Amos & Kolter, 2017) and disciplined convex programming (Section 7, Amos, 2019; Agrawal et al., 2019a;b) have been given, as well as expressions for more specific cases like isotonic regression (Djolonga & Krause, 2017). Here, by taking advantage of the structure of the decision tree induction problem, we obtain a direct, efficient algorithm.

Latent parse trees. Our work resembles but should not be confused with the latent parse tree literature in natural language processing (Yogatama et al., 2017; Liu & Lapata, 2018; Choi et al., 2018; Williams et al., 2018; Niculae et al., 2018b; Corro & Titov, 2019b;a; Maillard et al., 2019; Kim et al., 2019a;b). This line of work has a different goal than ours: to induce a tree for each individual data point (e.g., sentence). In contrast, our work aims to learn a single tree, for all instances to pass through.

Isotonic regression. Also called monotonic regression, isotonic regression (Barlow et al., 1972) constrains the regression function to be non-decreasing/non-increasing. This is useful if one has prior knowledge of such monotonicity (e.g., the mean temperature of the climate is non-decreasing). A classic algorithm is pooling-adjacent-violators (PAV), which optimizes the pooling of adjacent points that violate the monotonicity constraint (Barlow et al., 1972). This initial algorithm has been generalized and incorporated into convex programming frameworks (see Mair et al. (2009) for an excellent summary of the history of isotonic regression and its extensions). Our work builds off of the generalized PAV (GPAV) algorithm of Yu & Xing (2016).




Figure 2: A depiction of equation (1) for optimizing tree traversals given the tree parameters θ .

3. Method

Given inputs $\{\mathbf{x}_i \in \mathcal{X}\}_{i=1}^n$, our goal is to learn a latent binary decision tree \mathcal{T} with maximum depth D . This tree sends each input \mathbf{x} through branching nodes to a specific leaf node in the tree. Specifically, all branching nodes $\mathcal{T}_B \subset \mathcal{T}$ split an input \mathbf{x} by forcing it to go to its left child if $s_{\theta}(\mathbf{x}) < 0$, and right otherwise. There are three key parts of the tree that need to be identified: 1. The *continuous* parameters $\theta_t \in \mathbb{R}^d$ that describe how $s_{\theta_t}(\cdot)$ splits inputs at every node t ; 2. The *discrete* paths \mathbf{z} made by each input \mathbf{x} through the tree; 3. The *discrete* choice a_t of whether a node t should be active or pruned, i.e. inputs should reach/traverse it or not. We next describe how to represent each of these.

3.1. Tree Traversal & Pruning Programs

Let the splitting functions of a complete tree $\{s_{\theta_t} : \mathcal{X} \rightarrow \mathbb{R}\}_{t=1}^{|\mathcal{T}_B|}$ be fixed. The path through which a point \mathbf{x} traverses the tree can be encoded in a vector $\mathbf{z} \in \{0, 1\}^{|\mathcal{T}|}$, where each component z_t indicates whether \mathbf{x} reaches node $t \in \mathcal{T}$. The following integer linear program (ILP) is maximized by a vector \mathbf{z} that describes *tree traversal*:

$$\max_{\mathbf{z}} \mathbf{z}^\top \mathbf{q} \quad (1)$$

$$\begin{aligned} \text{s.t. } \forall t \in \mathcal{T} \setminus \{\text{root}\}, \quad & q_t = \min\{R_t \cup L_t\} \\ & R_t = \{s_{\theta_{t'}}(\mathbf{x}) \mid \forall t' \in A_R(t)\} \\ & L_t = \{-s_{\theta_{t'}}(\mathbf{x}) \mid \forall t' \in A_L(t)\} \\ & z_t \in \{0, 1\}, \end{aligned}$$

where we fix $q_1 = z_1 = 1$ (i.e., the root). Here $A_L(t)$ is the set of ancestors of node t whose left child must be followed to get to t , and similarly for $A_R(t)$. The quantities q_t (where $\mathbf{q} \in \mathbb{R}^{|\mathcal{T}|}$ is the tree-vectorized version of q_t) describe the ‘reward’ of sending \mathbf{x} to node t . This is based on how well the splitting functions leading to t are satisfied (q_t is positive if all splitting functions are satisfied and negative otherwise). This is visualized in Figure 2 and formalized below.

Lemma 1. For an input \mathbf{x} and binary tree \mathcal{T} with splits $\{s_{\theta_t}, \forall t \in \mathcal{T}_B\}$, the ILP in eq. (1) describes a valid tree traversal \mathbf{z} for \mathbf{x} (i.e., $z_t = 1$ for any $t \in \mathcal{T}$ if \mathbf{x} reaches t).

Proof. Assume without loss of generality that $s_{\theta_t}(\mathbf{x}) \neq 0$ for all $t \in \mathcal{T}$ (if this does not hold for some node t simply add a small constant ϵ to $s_{\theta_t}(\mathbf{x})$). Recall that for any non-leaf node t' in the fixed binary tree \mathcal{T} , a point \mathbf{x} is sent to the left child if $s_{\theta_{t'}}(\mathbf{x}) < 0$ and to the right child otherwise. Following these splits from the root, every point \mathbf{x} reaches a unique leaf node $f_{\mathbf{x}} \in \mathcal{F}$, where \mathcal{F} is the set of all leaf nodes. Notice first that both $\min L_{f_{\mathbf{x}}} > 0$ and $\min R_{f_{\mathbf{x}}} > 0$.² This is because $\forall t' \in A_L(f_{\mathbf{x}})$ it is the case that $s_{\theta_{t'}}(\mathbf{x}) < 0$, and $\forall t' \in A_R(f_{\mathbf{x}})$ we have that $s_{\theta_{t'}}(\mathbf{x}) > 0$. This is due to the definition of $f_{\mathbf{x}}$ as the leaf node that \mathbf{x} reaches (i.e., it traverses to $f_{\mathbf{x}}$ by following the aforementioned splitting rules). Further, for any other leaf $f \in \mathcal{F} \setminus \{f_{\mathbf{x}}\}$ we have that either $\min L_f < 0$ or $\min R_f < 0$. This is because there exists a node t' on the path to f (and so t' is either in $A_L(f)$ or $A_R(f)$) that does not agree with the splitting rules. For this node it is either the case that (i) $s_{\theta_{t'}}(\mathbf{x}) > 0$ and $t' \in A_L(f)$, or that (ii) $s_{\theta_{t'}}(\mathbf{x}) < 0$ and $t' \in A_R(f)$. In the first case $\min L_f < 0$, in the second $\min R_f < 0$. Therefore we have that $q_f < 0$ for all $f \in \mathcal{F} \setminus \{f_{\mathbf{x}}\}$, and $q_{f_{\mathbf{x}}} > 0$. In order to maximize the objective $\mathbf{z}^\top \mathbf{q}$ we will have that $z_{f_{\mathbf{x}}} = 1$ and $z_f = 0$. Finally, let $\mathcal{N}_{f_{\mathbf{x}}}$ be the set of non-leaf nodes visited by \mathbf{x} on its path to $f_{\mathbf{x}}$. Now notice that the above argument applied to leaf nodes can be applied to nodes at any level of the tree: $q_t > 0$ for $t \in \mathcal{N}_{f_{\mathbf{x}}}$ while $q_{t'} < 0$ for $t' \in \mathcal{T} \setminus \mathcal{N}_{f_{\mathbf{x}}} \cup f_{\mathbf{x}}$. Therefore $\mathbf{z}_{\mathcal{N}_{f_{\mathbf{x}}} \cup f_{\mathbf{x}}} = 1$ and $\mathbf{z}_{\mathcal{T} \setminus \mathcal{N}_{f_{\mathbf{x}}} \cup f_{\mathbf{x}}} = 0$, completing the proof. \square

This solution is unique so long as $s_{\theta_t}(\mathbf{x}_i) \neq 0$ for all $t \in \mathcal{T}, i \in \{1, \dots, n\}$ (i.e., $s_{\theta_t}(\mathbf{x}_i) = 0$ means equal preference to split \mathbf{x}_i left or right). Further the integer constraint on z_{it} can be relaxed to an interval constraint $z_{it} \in [0, 1]$ w.l.o.g. This is because if $s_{\theta_t}(\mathbf{x}_i) \neq 0$ then $z_t = 0$ if and only if $q_t < 0$ and $z_t = 1$ when $q_t > 0$ (and $q_t \neq 0$).

While the above program works for any fixed tree, we would like to be able to also learn the structure of the tree itself. We do so by learning a pruning optimization variable $a_t \in \{0, 1\}$, indicating if node $t \in \mathcal{T}$ is active (if 1) or pruned (if 0). We adapt eq. (1) into the following pruning-aware mixed integer program (MIP) considering all inputs $\{\mathbf{x}_i\}_{i=1}^n$:

$$\begin{aligned} & \max_{\mathbf{z}_1, \dots, \mathbf{z}_n, \mathbf{a}} \sum_{i=1}^n \mathbf{z}_i^\top \mathbf{q}_i - \frac{\lambda}{2} \|\mathbf{a}\|^2 & (2) \\ \text{s.t. } & \forall i \in [n], a_t \leq a_{p(t)}, \quad \forall t \in \mathcal{T} \setminus \{\text{root}\} \\ & z_{it} \leq a_t \\ & z_{it} \in [0, 1], a_t \in \{0, 1\} \end{aligned}$$

²We use the convention that $\min \emptyset = \infty$ (i.e., for paths that go down the rightmost or leftmost parts of the tree).




Figure 3: Routing of point \mathbf{x}_i at node t , without pruning, without (left) or with (right) quadratic relaxation. Depending on the decision split $s_{\theta_t}(\cdot)$, \mathbf{x}_i reaches node t 's left child l (right child r respectively) if $z_{il} > 0$ ($z_{ir} > 0$). The relaxation makes \mathbf{z}_i continuous and encourages points to have a margin ($|s_{\theta_t}| > 0.5$).

with $\|\cdot\|$ denoting the l_2 norm. We remove the first three constraints in eq. (1) as they are a deterministic computation independent of $\mathbf{z}_1, \dots, \mathbf{z}_n, \mathbf{a}$. We denote by $p(t)$ the parent of node t . The new constraints $a_t \leq a_{p(t)}$ ensure that child nodes t are pruned if the parent node $p(t)$ is pruned, hence enforcing the tree hierarchy, while the other new constraint $z_{it} \leq a_t$ ensures that no point \mathbf{x}_i can reach node t if node t is pruned, resulting in losing the associated reward q_{it} . Overall, the problem consists in a trade-off, controlled by hyper-parameter $\lambda \in \mathbb{R}$, between minimizing the number of active nodes through the pruning regularization term (since $a_t \in \{0, 1\}$, $\|\mathbf{a}\|^2 = \sum_t a_t^2 = \sum_t \mathbb{I}(a_t = 1)$) while maximizing the reward from points traversing the nodes.

3.2. Learning Tree Parameters

We now consider the problem of learning the splitting parameters θ_t . A natural approach would be to do so in the MIP itself, as in the optimal tree literature. However, this would severely restrict allowed splitting functions as even linear splitting functions can only practically run on at most thousands of training inputs (Bertsimas & Dunn, 2017). Instead, we propose to learn s_{θ_t} via gradient descent.

To do so, we must be able to compute the partial derivatives $\partial \mathbf{a} / \partial \mathbf{q}$ and $\partial \mathbf{z} / \partial \mathbf{q}$. However, the solutions of eq. (2) are piecewise-constant, leading to null gradients. To avoid this, we relax the integer constraint on \mathbf{a} to the interval $[0, 1]$ and add quadratic regularization $-1/4 \sum_i (\|\mathbf{z}_i\|^2 + \|1 - \mathbf{z}_i\|^2)$. The regularization term for \mathbf{z} is symmetric so that it shrinks solutions toward even left-right splits (see Figure 3 for a visualization, and the supplementary for a more detailed justification with an empirical evaluation of the approximation gap). Rearranging and negating the objective yields

$$\begin{aligned} & T_\lambda(\mathbf{q}_1, \dots, \mathbf{q}_n) = & (3) \\ & \arg \min_{\mathbf{z}_1, \dots, \mathbf{z}_n, \mathbf{a}} \lambda/2 \|\mathbf{a}\|^2 + 1/2 \sum_{i=1}^n \|\mathbf{z}_i - \mathbf{q}_i - 1/2\|^2 \\ \text{s.t. } & \forall i \in [n], a_t \leq a_{p(t)}, \quad \forall t \in \mathcal{T} \setminus \{\text{root}\} \\ & z_{it} \leq a_t \\ & z_{it} \in [0, 1], a_t \in [0, 1]. \end{aligned}$$

Algorithm 1 Pruning via isotonic optimization

```

initial partition  $\mathcal{G} \leftarrow \{\{1\}, \{2\}, \dots\} \subset 2^{\mathcal{T}}$ 
for  $G \in \mathcal{G}$  do
   $a_G \leftarrow \arg \min_a \sum_{t \in G} g_t(a)$       {Eq. (7), Prop. 2}
end for
repeat
   $t_{\max} \leftarrow \arg \max_t \{a_t : a_t > a_{p(t)}\}$ 
  merge  $G' \leftarrow G' \cup G$  where  $G' \ni t_{\max}$  and  $G \ni p(t_{\max})$ .
  update  $a_G \leftarrow \arg \min_a \sum_{t \in G} g_t(a)$       {Eq. (7), Prop. 2}
until no more violations  $a_t > a_{p(t)}$  exist
    
```

The regularization makes the objective strongly convex. It follows from convex duality that T_λ is Lipschitz continuous (Zalinescu, 2002, Corollary 3.5.11). By Rademacher’s theorem (Borwein & Lewis, 2010, Theorem 9.1.2), T_λ is thus differentiable almost everywhere. Generic methods such as OptNet (Amos & Kolter, 2017) could be used to compute the solution and gradients. However, by using the tree structure of the constraints, we next derive an efficient specialized algorithm. The main insight, shown below, reframes pruned binary tree learning as isotonic optimization.

Proposition 1. Let $\mathcal{C} = \{\mathbf{a} \in \mathbb{R}^{|\mathcal{T}|} : a_t \leq a_{p(t)} \text{ for all } t \in \mathcal{T} \setminus \{\text{root}\}\}$. Consider \mathbf{a}^* the optimal value of

$$\arg \min_{\mathbf{a} \in \mathcal{C} \cap [0,1]^{|\mathcal{T}|}} \sum_{t \in \mathcal{T}} \left(\frac{\lambda}{2} a_t^2 + \sum_{i: a_t < q_{it} + 1/2} (a_t - q_{it} - 1/2)^2 \right). \quad (4)$$

Define³ \mathbf{z}_i^* such that $z_{it}^* = \text{Proj}_{[0, a_t]}(q_{it} + 1/2)$. Then, $T_\lambda(\mathbf{q}_1, \dots, \mathbf{q}_n) = \mathbf{z}_1^*, \dots, \mathbf{z}_n^*, \mathbf{a}^*$.

Proof. The constraints and objective of eq. (3) are separable, so we may push the minimization *w.r.t.* \mathbf{z} inside the objective:

$$\arg \min_{\mathbf{a} \in \mathcal{C} \cap [0,1]^{|\mathcal{T}|}} \frac{\lambda}{2} \|\mathbf{a}\|^2 + \sum_{t \in \mathcal{T}} \sum_{i=1}^n \min_{0 \leq z_{it} \leq a_t} 1/2 (z_{it} - q_{it} - 1/2)^2. \quad (5)$$

Each inner minimization $\min_{0 \leq z_{it} \leq a_t} 1/2 (z_{it} - q_{it} - 1/2)^2$ is a one-dimensional projection onto box constraints, with solution $z_{it}^* = \text{Proj}_{[0, a_t]}(q_{it} + 1/2)$. We use this to eliminate \mathbf{z} , noting that each term $1/2 (z_{it}^* - q_{it} - 1/2)^2$ yields

$$\begin{cases} 1/2 (q_{it} + 1/2)^2, & q_{it} < -1/2 \\ 0, & -1/2 \leq q_{it} < a_t - 1/2 \\ 1/2 (a_t - q_{it} - 1/2)^2, & q_{it} \geq a_t - 1/2. \end{cases} \quad (6)$$

The first two branches are constants *w.r.t.* a_t . The problems in eq. (3) and eq. (4) differ by a constant and have the same constraints, therefore they have the same solutions. \square

³Here $\text{Proj}_{\mathcal{S}}(x)$ is the projection of x onto set \mathcal{S} . If \mathcal{S} are box constraints, projection amounts to clipping.




Figure 4: One step in the tree isotonic optimization process from Algorithm 1. Shaded tubes depict the partition groups \mathcal{G} , and the emphasized arc is the largest among the violated inequalities. The group of t_{\max} merges with the group of its parent, and the new value for the new, larger group a_G is computed.

Efficiently inducing trees as isotonic optimization.

From Proposition 1, notice that eq. (4) is an instance of tree-structured isotonic optimization: the objective decomposes over *nodes*, and the inequality constraints correspond to edges in a rooted tree:

$$\arg \min_{\mathbf{a} \in \mathcal{C}} \sum_{t \in \mathcal{T}} g_t(a_t), \quad \text{where} \quad (7)$$

$$g_t(a_t) = \frac{\lambda}{2} a_t^2 + \sum_{i: a_t \leq q_{it} + 1/2} 1/2 (a_t - q_{it} - 1/2)^2 + \iota_{[0,1]}(a_t).$$

where $\iota_{[0,1]}(a_t) = \infty$ if $a_t \notin [0, 1]$ and 0 otherwise. This problem can be solved by a generalized *pool adjacent violators* (PAV) algorithm: Obtain a tentative solution by ignoring the constraints, then iteratively remove violating edges $a_t > a_{p(t)}$ by *pooling together* the nodes at the end points. Figure 4 provides an illustrative example of the procedure. At the optimum, the nodes are organized into a partition $\mathcal{G} \subset 2^{\mathcal{T}}$, such that if two nodes t, t' are in the same group $G \in \mathcal{G}$, then $a_t = a_{t'} := a_G$.

When the inequality constraints are the edges of a rooted tree, as is the case here, the PAV algorithm finds the optimal solution in at most $|\mathcal{T}|$ steps, where each involves updating the a_G value for a newly-pooled group by solving a one-dimensional subproblem of the form (Yu & Xing, 2016)⁴

$$a_G = \arg \min_{a \in \mathbb{R}} \sum_{t \in G} g_t(a), \quad (8)$$

resulting in Algorithm 1. It remains to show how to solve eq. (8). The next result, proved in the supplementary, gives an exact and efficient solution, with an algorithm that requires finding the nodes with highest q_{it} (*i.e.*, the nodes which \mathbf{x}_i is most highly encouraged to traverse).

⁴Compared to Yu & Xing (2016), our tree inequalities are in the opposite direction. This is equivalent to a sign flip of parameter \mathbf{a} , *i.e.*, to selecting the *maximum* violator rather than the minimum one at each iteration.




Figure 5: Comparison of Algorithm 1’s running time (ours) with the running time of cvxpylayers (Agrawal et al., 2019a) for tree depth $D=3$, varying n .

Proposition 2. *The solution to the one-dimensional problem in eq. (8) for any G is given by*

$$\arg \min_{a \in \mathbb{R}} \sum_{t \in G} g_t(a) = \text{Proj}_{[0,1]}(a(k^*)) \quad (9)$$

$$\text{where } a(k^*) := \frac{\sum_{(i,t) \in S(k^*)} (q_{it} + 1/2)}{\lambda|G| + k^*},$$

$S(k) = \{j^{(1)}, \dots, j^{(k)}\}$ is the set of indices $j = (i, t) \in \{1, \dots, n\} \times G$ into the k highest values of \mathbf{q} , i.e., $q_{j^{(1)}} \geq q_{j^{(2)}} \geq \dots \geq q_{j^{(m)}}$, and k^* is the smallest k satisfying $a(k) > q_{j^{(k+1)}} + 1/2$.

Figure 5 compares our algorithm to a generic differentiable solver (details and additional comparisons in supplementary material).

Backward pass and efficient implementation details.

Algorithm 1 is a sequence of differentiable operations that can be implemented *as is* in automatic differentiation frameworks. However, because of the prominent loops and indexing operations, we opt for a low-level implementation as a C++ extension. Since the \mathbf{q} values are constant *w.r.t.* \mathbf{a} , we only need to sort them once as preprocessing, resulting in a overall time complexity of $\mathcal{O}(n|\mathcal{T}| \log(n|\mathcal{T}|))$ and space complexity of $\mathcal{O}(n|\mathcal{T}|)$. For the backward pass, rather than relying on automatic differentiation, we make two remarks about the form of \mathbf{a} . Firstly, its elements are organized in groups, i.e., $a_t = a_{t'} = a_G$ for $\{t, t'\} \subset G$. Secondly, the value a_G inside each group depends only on the optimal support set $S_G^* := S(k^*)$ as defined for each subproblem by Proposition 2. Therefore, in the forward pass, we must store only the node-to-group mappings and the sets S_G^* . Then, if G contains t ,

$$\frac{\partial a_t^*}{\partial q_{it'}} = \begin{cases} \frac{1}{\lambda|G| + k^*}, & 0 < a_t^* < 1, \text{ and } (i, t') \in S_G^*, \\ 0, & \text{otherwise.} \end{cases}$$

As T_λ is differentiable almost everywhere, these expressions yield the unique Jacobian at all but a measure-zero set of points, where they yield one of the Clarke generalized

Algorithm 2 Learning with decision tree representations.

```

initialize neural network parameters  $\phi, \theta$ 
repeat
  sample batch  $\{\mathbf{x}_i\}_{i \in B}$ 
  induce traversals:
   $\{\mathbf{z}_i\}_{i \in B}, \mathbf{a} = T_\lambda(\{q_\theta(\mathbf{x}_i)\})$  {Algorithm 1; differentiable}
  update parameters using  $\nabla_{\{\theta, \phi\}} \ell(f_\phi(\mathbf{x}_i, \mathbf{z}_i))$  {autograd}
until convergence or stopping criterion is met
    
```

Jacobians (Clarke, 1990). We then rely on automatic differentiation to propagate gradients from \mathbf{z} to \mathbf{q} , and from \mathbf{q} to the split parameters θ : since \mathbf{q} is defined element-wise via min functions, the gradient propagates through the minimizing path, by Danskin’s theorem (Proposition B.25, Bertsekas, 1999; Danskin, 1966).

3.3. The Overall Objective

We are now able to describe the overall optimization procedure that simultaneously learns tree parameters: (a) input traversals $\mathbf{z}_1, \dots, \mathbf{z}_n$; (b) tree pruning \mathbf{a} ; and (c) split parameters θ . Instead of making predictions using heuristic scores over the training points assigned to a leaf (e.g., majority class), we learn a prediction function $f_\phi(\mathbf{z}, \mathbf{x})$ that minimizes an arbitrary loss $\ell(\cdot)$ as follows:

$$\min_{\theta, \phi} \sum_{i=1}^n \ell(f_\phi(\mathbf{x}_i, \mathbf{z}_i)) \quad (10)$$

where $\mathbf{z}_1, \dots, \mathbf{z}_n, \mathbf{a} := T_\lambda(q_\theta(\mathbf{x}_1), \dots, q_\theta(\mathbf{x}_n))$.

This corresponds to embedding a decision tree as a layer of a deep neural network (using \mathbf{z} as intermediate representation) and optimizing the parameters θ and ϕ of the whole model by back-propagation. In practice, we perform mini-batch updates for efficient training; the procedure is sketched in Algorithm 2. Here we define $q_\theta(\mathbf{x}_i) := \mathbf{q}_i$ to make explicit the dependence of \mathbf{q}_i on θ .

4. Experiments

In this section we showcase our method on both: (a) *Classification/Regression* for tabular data, where tree-based models have been demonstrated to have superior performance over MLPs (Popov et al., 2020); and (b) *Hierarchical clustering* on unsupervised data. Our experiments demonstrate that our method leads to predictors that are competitive with state-of-the-art tree-based approaches, scaling better with the size of datasets and generalizing to many tasks. Further we visualize the trees learned by our method and how sparsity is easily adjusted by tuning the hyper-parameter λ .

Architecture details. We use a linear function or a multi-layer perceptron (L fully-connected layers with ELU activation (Clevert et al., 2016) and dropout) for $f_\phi(\cdot)$ and choose between linear or linear followed by ELU splitting

Table 1: Results on tabular regression datasets. We report average and standard deviations over 4 runs of MSE and **bold** best results, and those within a standard deviation from it, for each family of algorithms (single tree or ensemble). For the single tree methods we additionally report the average training times (s).

	METHOD	YEAR	MICROSOFT	YAHOO
Single	CART	96.0 ± 0.4	0.599 ± 1e-3	0.678 ± 17e-3
	ANT	77.8 ± 0.6	0.572 ± 2e-3	0.589 ± 2e-3
	Ours	77.9 ± 0.4	0.573 ± 2e-3	0.591 ± 1e-3
Ens.	NODE	76.2 ± 0.1	0.557 ± 1e-3	0.569 ± 1e-3
	XGBoost	78.5 ± 0.1	0.554 ± 1e-3	0.542 ± 1e-3
Time	CART	23s	20s	26s
	ANT	4674s	1457s	4155s
	ours	1354s	1117s	825s

functions $s_\theta(\cdot)$ (we limit the search for simplicity, there are no restrictions except differentiability).

4.1. Supervised Learning on Tabular Datasets

Our first set of experiments is on supervised learning with heterogeneous tabular datasets, where we consider both regression and binary classification tasks. We minimize the Mean Square Error (MSE) on regression datasets and the Binary Cross-Entropy (BCE) on classification datasets. We compare our results with tree-based architectures, which either train a single or an ensemble of decision trees. Namely, we compare against the greedy CART algorithm (Breiman et al., 1984) and two optimal decision tree learners: OPTREE with local search (Optree-LS, Dunn, 2018) and a state-of-the-art optimal tree method (GOSDT, Lin et al., 2020). We also consider three baselines with probabilistic routing: deep neural decision trees (DNDT, Yang et al., 2018), deep neural decision forests (Kontschieder et al., 2015) configured to jointly optimize the routing and the splits and to use an ensemble size of 1 (NDF-1), and adaptive neural networks (ANT, Tanno et al., 2019). As for the ensemble baselines, we compare to NODE (Popov et al., 2020), the state-of-the-art method for training a forest of differentiable oblivious decision trees on tabular data, and to XGBoost (Chen & Guestrin, 2016), a scalable tree boosting method. We carry out the experiments on the following datasets. **Regression:** *Year* (Bertin-Mahieux et al., 2011), Temporal regression task constructed from the Million Song Dataset; *Microsoft* (Qin & Liu, 2013), Regression approach to the MSLR-Web10k Query-URL relevance prediction for learning to rank; *Yahoo* (Chapelle & Chang, 2011), Regression approach to the C14 learning-to-rank challenge. **Binary classification:** *Click*, Link click prediction based on the KDD Cup 2012 dataset, encoded and subsampled following Popov et al. (2020); *Higgs* (Baldi et al., 2014), prediction of Higgs boson-producing events.

For all tasks, we follow the preprocessing and task setup from (Popov et al., 2020). All datasets come with train-

Table 2: Results on tabular classification datasets. We report average and standard deviations over 4 runs of error rate. Best result for each family of algorithms (single tree or ensemble) are in **bold**. Experiments are run on a machine with 16 CPUs and 64GB of RAM, with a training time limit of 3 days. We denote methods that exceed this memory and training time as OOM and OOT, respectively. For the single tree methods we additionally report the average training times (s) when available.

	METHOD	CLICK	HIGGS
Single Tree	GOSDT	OOM	OOM
	OPTREE-LS	OOT	OOT
	DNDT	0.4866 ± 1e-2	OOM
	NDF-1	0.3344 ± 5e-4	0.2644 ± 8e-4
	CART	0.3426 ± 11e-3	0.3430 ± 8e-3
	ANT	0.4035 ± 0.1150	0.2430 ± 6e-3
	Ours	0.3340 ± 3e-4	0.2201 ± 3e-4
Ens.	NODE	0.3312 ± 2e-3	0.210 ± 5e-4
	XGBoost	0.3310 ± 2e-3	0.2334 ± 1e-3
Time	DNDT	681s	-
	NDF-1	3768s	43593s
	CART	3s	113s
	ANT	75600s	62335s
	ours	524s	18642s

ing/test splits. We make use of 20% of the training set as validation set for selecting the best model over training and for tuning the hyperparameters. We tune the hyperparameters for all methods. and optimize eq. (10) and all neural network methods (DNDT, NDF, ANT and NODE) using the Quasi-Hyperbolic Adam (Ma & Yarats, 2019) stochastic gradient descent method. Further details are provided in the supplementary. Tables 1 and 2 report the obtained results on the regression and classification datasets respectively.⁵ Unsurprisingly, ensemble methods outperform single-tree ones on all datasets, although at the cost of being harder to visualize/interpret. Our method has the advantage of (a) generalizing to any task; (b) outperforming or matching all single-tree methods; (c) approaching the performance of ensemble-based methods; (d) scaling well with the size of datasets. These experiments show that our model is also significantly faster to train, compared to its differentiable tree counterparts NDF-1 and ANT, while matching or beating the performance of these baselines, and it generally provides the best trade-off between time complexity and accuracy over all datasets (visualizations of this trade-off are reported in the supplementary material). Further results on smaller datasets are available in the supplementary material to provide a comparison with optimal tree baselines.

4.2. Self-Supervised Hierarchical Clustering

To show the versatility of our method, we carry out a second set of experiments on hierarchical clustering tasks. Inspired by the recent success of self-supervised learning approaches

⁵GOSDT/DNDT/Optree-LS/NDF are for classification only.

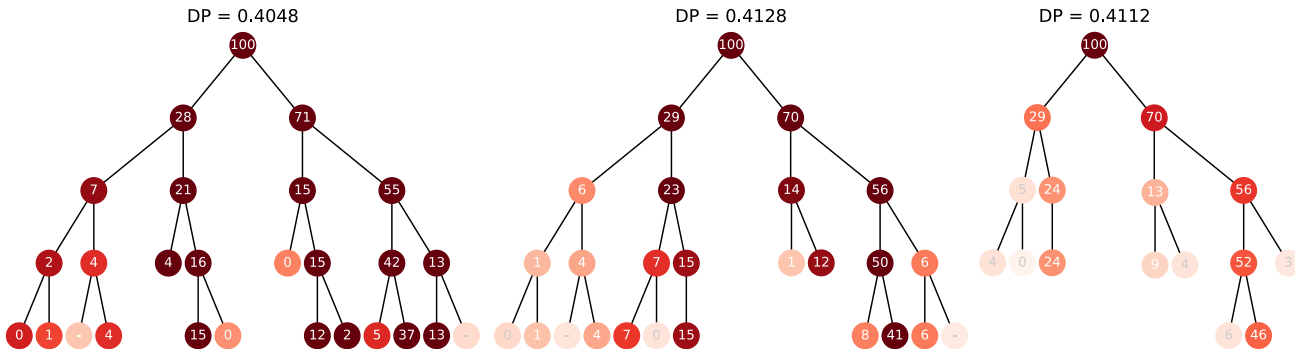


Figure 6: *Glass* tree routing distribution, in rounded percent of dataset, for λ left-to-right in $\{1, 10, 100\}$. The larger λ , the more nodes are pruned. We report dendrogram purity (DP) and represent the nodes by the percentage of points traversing them (normalized at each depth level) and with a color intensity depending on their a_t (the darker the higher a_t). The empty nodes are labeled by a dash – and the inactive nodes ($a_t = 0$) have been removed.

Table 3: Results for hierarchical clustering. We report average and standard deviations of *dendrogram purity* over four runs.

METHOD	GLASS	COVTYPE
Ours	0.468 ± 0.028	0.459 ± 0.008
gHHC	0.463 ± 0.002	0.444 ± 0.005
HKMeans	0.508 ± 0.008	0.440 ± 0.001
BIRCH	0.429 ± 0.013	0.440 ± 0.002

(Lan et al., 2019; He et al., 2020), we learn a tree for hierarchical clustering in a self-supervised way. Specifically, we regress a subset of input features from the remaining features, minimizing the MSE. This allows us to use eq. (10) to learn a hierarchy (tree). To evaluate the quality of the learned trees, we compute their dendrogram purity (DP; Monath et al., 2019). DP measures the ability of the learned tree to separate points from different classes, and corresponds to the expected purity of the least common ancestors of points of the same class.

We experiment on the following datasets: *Glass* (Dua & Graff, 2017): glass identification for forensics, and *Covtype* (Blackard & Dean, 1999; Dua & Graff, 2017): cartographic variables for forest cover type identification. For *Glass*, we regress features ‘Refractive Index’ and ‘Sodium,’ and for *Covtype* the horizontal and vertical ‘Distance To Hydrology.’ We split the datasets into training/validation/test sets, with sizes 60%/20%/20%. Here we only consider linear f_ϕ . As before, we optimize Problem 10 using the Quasi-Hyperbolic Adam algorithm and tune the hyper-parameters using the validation reconstruction error.

As baselines, we consider: BIRCH (Zhang et al., 1997) and Hierarchical KMeans (HKMeans), the standard methods for performing clustering on large datasets; and the recently proposed gradient-based Hyperbolic Hierarchical Clustering (gHHC, Monath et al., 2019) designed to construct trees in hyperbolic space. Table 3 reports the dendrogram purity scores for all methods. Our method yields results




Figure 7: Percentage of active nodes during training as a function of the number of epochs on *Glass*, $D = 6$, $\lambda = 10$.

comparable to all baselines, even though not specifically tailored to hierarchical clustering.

Tree Pruning The hyper-parameter λ in Problem 3 controls how aggressively the tree is pruned, hence the amount of tree splits that are actually used to make decisions. This is a fundamental feature of our framework as it allows to smoothly trim the portions of the tree that are not necessary for the downstream task, resulting in lower computing and memory demands at inference. In Figure 6, we study the effects of pruning on the tree learned on *Glass* with a depth fixed to $D = 4$. We report how inputs are distributed over the learned tree for different values of λ . We notice that increasing λ effectively prune nodes and entire portions of the tree, without significantly impact performance (as measured by dendrogram purity).

To look into the evolution of pruning during training, we further plot the % of active (unpruned) nodes within a training epoch in Figure 7. We observe that (a) it appears possible to increase and decrease this fraction through training (b) the fraction seems to stabilize in the range 45%-55% after a few epochs.




Figure 8: Class routing distributions on *Glass*, with distributions normalized over each depth level. Trees were trained with optimal hyper-parameters (depth $D = 5$), but we plot nodes up to $D = 3$ for visualization ease. Empty nodes are labeled by a dash $-$.

Class Routing To gain insights on the latent structure learned by our method, we study how points are routed through the tree, depending on their class. The *Glass* dataset is particularly interesting to analyze as its classes come with an intrinsic hierarchy, e.g., with superclasses *Window* and *NonWindow*. Figure 8 reports the class routes for four classes. As the trees are constructed without supervision, we do not expect the structure to exactly reflect the class partition and hierarchy. Still, we observe that points from the same class or super-class traverse the tree in a similar way. Indeed, trees for class *Build* 8(a) and class *Vehicle* 8(b) that both belong to the *Window* super-class, share similar paths, unlike the classes *Containers* 8(c) and *Headlamps* 8(d).

5. Discussion

In this work we have presented a new optimization approach to learn trees for a variety of machine learning tasks. Our method works by sparsely relaxing a ILP for tree traversal and pruning, to enable simultaneous optimization of these parameters, alongside splitting parameters and downstream functions via argmin differentiation. Our approach nears or improves upon recent work in both supervised learning and hierarchical clustering. We believe there are many exciting avenues for future work. One particularly interesting direction would be to unify recent advances in tight relaxations of nearest neighbor classifiers with this approach to learn efficient neighbor querying structures such as ball trees. Another idea is to adapt this method to learn instance-specific trees such as parse trees.

Acknowledgements

The authors are thankful to Mathieu Blondel and Caio Corro for useful suggestions. They acknowledge the support of Microsoft AI for Earth grant in the form of Microsoft Azure computing credits. VN acknowledges support from the European Research Council (ERC StG DeepSPIN 758969) and the Fundação para a Ciência e Tecnologia through contract UIDB/50008/2020 while at the Instituto de Telecomunicações.

References

- Aghaei, S., Azizi, M. J., and Vayanos, P. Learning optimal and fair decision trees for non-discriminative decision-making. In *Proc. of AAAI*, 2019.
- Aghaei, S., Gomez, A., and Vayanos, P. Learning optimal classification trees: Strong max-flow formulations. *preprint arXiv:2002.09142*, 2020.
- Agrawal, A., Amos, B., Barratt, S., Boyd, S., Diamond, S., and Kolter, Z. Differentiable convex optimization layers. In *Proc. of NeurIPS*, 2019a.
- Agrawal, A., Barratt, S., Boyd, S., Busseti, E., and Moursi, W. M. Differentiating through a cone program. *Journal of Applied and Numerical Optimization*, 2019b. ISSN 25625527. doi: 10.23952/jano.1.2019.2.02.
- Amos, B. *Differentiable Optimization-Based Modeling for Machine Learning*. PhD thesis, Carnegie Mellon University, May 2019.
- Amos, B. and Kolter, J. Z. OptNet: Differentiable optimization as a layer in neural networks. In *Proc. of ICML*, 2017.
- Angelino, E., Larus-Stone, N., Alabi, D., Seltzer, M., and Rudin, C. Learning certifiably optimal rule lists for categorical data. *Journal of Machine Learning Research*, 2017.
- Baldi, P., Sadowski, P., and Whiteson, D. Searching for exotic particles in high-energy physics with deep learning. *Nature Communications*, 2014.
- Barlow, R., Bartholomev, D., Brenner, J., and Brunk, H. Statistical inference under order restrictions: The theory and application of isotonic regression, 1972.
- Bennett, K. P. Decision tree construction via linear programming. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 1992.
- Bennett, K. P. and Blue, J. A. Optimal decision trees. *Rensselaer Polytechnic Institute Math Report*, 1996.

- Bertin-Mahieux, T., Ellis, D. P., Whitman, B., and Lamere, P. The million song dataset. In *Proc. of ISMIR*, 2011.
- Bertsekas, D. P. *Nonlinear Programming*. Athena Scientific Belmont, 1999.
- Bertsimas, D. and Dunn, J. Optimal classification trees. *Machine Learning*, 2017.
- Blackard, J. A. and Dean, D. J. Comparative accuracies of artificial neural networks and discriminant analysis in predicting forest cover types from cartographic variables. *Computers and Electronics in Agriculture*, 1999.
- Borwein, J. and Lewis, A. S. *Convex analysis and nonlinear optimization: theory and examples*. Springer Science & Business Media, 2010.
- Breiman, L., Friedman, J., Stone, C. J., and Olshen, R. A. *Classification and Regression Trees*. CRC press, 1984.
- Carreira-Perpinán, M. A. and Tavallali, P. Alternating optimization of decision trees, with application to learning sparse oblique trees. In *Proc. of NeurIPS*, 2018.
- Chapelle, O. and Chang, Y. Yahoo! learning to rank challenge overview. In *Proc. of the learning to rank challenge*, 2011.
- Chen, T. and Guestrin, C. XGBoost: A scalable tree boosting system. In *Proc. of KDD*, 2016.
- Choi, J., Yoo, K. M., and Lee, S.-g. Learning to compose task-specific tree structures. In *Proc. of AAAI*, 2018.
- Clarke, F. H. *Optimization and nonsmooth analysis*. SIAM, 1990.
- Clevert, D.-A., Unterthiner, T., and Hochreiter, S. Fast and accurate deep network learning by exponential linear units (ELUs). *Proc. of ICLR*, 2016.
- Colson, B., Marcotte, P., and Savard, G. [An overview of bilevel optimization](#). *Annals of operations research*, 2007.
- Corro, C. and Titov, I. Learning latent trees with stochastic perturbations and differentiable dynamic programming. In *Proc. of ACL*, 2019a.
- Corro, C. and Titov, I. Differentiable Perturb-and-Parse: Semi-Supervised Parsing with a Structured Variational Autoencoder. In *Proc. of ICLR*, 2019b.
- Danskin, J. M. The theory of max-min, with applications. *SIAM Journal on Applied Mathematics*, 1966.
- Dasgupta, S. A cost function for similarity-based hierarchical clustering. In *Proc. of STOC*, 2016.
- Djlonga, J. and Krause, A. Differentiable learning of submodular models. In *Proc. of NeurIPS*, 2017.
- Domke, J. Learning graphical model parameters with approximate marginal inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2013.
- Dua, D. and Graff, C. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Dunn, J. W. *Optimal trees for prediction and prescription*. PhD thesis, Massachusetts Institute of Technology, 2018.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In *Proc. of ICML*, 2017.
- Gould, S., Fernando, B., Cherian, A., Anderson, P., Cruz, R. S., and Guo, E. On differentiating parameterized argmin and argmax problems with application to bi-level optimization. *CoRR*, abs/1607.05447, 2016.
- Günlük, O., Kalagnanam, J., Menickelly, M., and Scheinberg, K. Optimal decision trees for categorical data via integer programming. *Journal of Global Optimization*, 2021.
- Hardt, M., Recht, B., and Singer, Y. Train faster, generalize better: Stability of stochastic gradient descent. In *Proc. of ICML*, 2016.
- Hazimeh, H., Ponomareva, N., Mol, P., Tan, Z., and Mazumder, R. The tree ensemble layer: Differentiability meets conditional computation. In *Proc. of ICML*, 2020.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *Proc. of CVPR*, 2020.
- Hoffer, E., Hubara, I., and Soudry, D. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. In *Proc. of NeurIPS*, 2017.
- Hu, X., Rudin, C., and Seltzer, M. Optimal sparse decision trees. In *Proc. of NeurIPS*, 2019.
- İrsoy, O., Yıldız, O. T., and Alpaydın, E. Soft decision trees. In *Proc. of ICPR*, 2012.
- Kim, Y., Dyer, C., and Rush, A. Compound probabilistic context-free grammars for grammar induction. In *Proc. of ACL*, 2019a.
- Kim, Y., Rush, A., Yu, L., Kuncoro, A., Dyer, C., and Melis, G. Unsupervised recurrent neural network grammars. In *Proc. of NAACL-HLT*, 2019b.
- Kobren, A., Monath, N., Krishnamurthy, A., and McCallum, A. A hierarchical algorithm for extreme clustering. In *Proc. of KDD*, 2017.

- Kontschieder, P., Fiterau, M., Criminisi, A., and Rota Bulo, S. Deep neural decision forests. In *Proc. of ICCV*, 2015.
- Kool, W., van Hoof, H., and Welling, M. Attention, learn to solve routing problems! In *Proc. of ICLR*, 2018.
- Krishnamurthy, A., Balakrishnan, S., Xu, M., and Singh, A. Efficient active algorithms for hierarchical clustering. In *Proc. of ICML*, 2012.
- Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., and Soricut, R. Albert: A lite bert for self-supervised learning of language representations. In *Proc. of ICLR*, 2019.
- Lay, N., Harrison, A. P., Schreiber, S., Dawer, G., and Barbu, A. Random hinge forest for differentiable learning. *preprint arXiv:1802.03882*, 2018.
- Letham, B., Rudin, C., McCormick, T. H., Madigan, D., et al. Interpretable classifiers using rules and bayesian analysis: Building a better stroke prediction model. *The Annals of Applied Statistics*, 2015.
- Lin, J., Zhong, C., Hu, D., Rudin, C., and Seltzer, M. Generalized and scalable optimal sparse decision trees. In *Proc. of ICML*, 2020.
- Liu, Y. and Lapata, M. Learning structured text representations. *TACL*, 2018.
- Ma, J. and Yarats, D. Quasi-hyperbolic momentum and adam for deep learning. *Proc. of ICLR*, 2019.
- Maillard, J., Clark, S., and Yogatama, D. Jointly learning sentence embeddings and syntax with unsupervised tree-lstms. *Natural Language Engineering*, 2019.
- Mair, P., Hornik, K., and de Leeuw, J. Isotone optimization in r: pool-adjacent-violators algorithm (pava) and active set methods. *Journal of Statistical Software*, 2009.
- Monath, N., Zaheer, M., Silva, D., McCallum, A., and Ahmed, A. Gradient-based hierarchical clustering using continuous representations of trees in hyperbolic space. In *Proc. of KDD*, 2019.
- Moseley, B. and Wang, J. Approximation bounds for hierarchical clustering: Average linkage, bisecting k-means, and local search. In *Proc. of NeurIPS*, 2017.
- Niculae, V., Martins, A. F., Blondel, M., and Cardie, C. SparseMAP: Differentiable sparse structured inference. In *Proc. of ICML*, 2018a.
- Niculae, V., Martins, A. F., and Cardie, C. Towards dynamic computation graphs via sparse latent structure. In *Proc. of EMNLP*, 2018b.
- Payne, H. J. and Meisel, W. S. An algorithm for constructing optimal binary decision trees. *IEEE Transactions on Computers*, 1977.
- Pedregosa, F. Hyperparameter optimization with approximate gradient. In *Proc. of ICML*, 2016.
- Popov, S., Morozov, S., and Babenko, A. Neural oblivious decision ensembles for deep learning on tabular data. *Proc. of ICLR*, 2020.
- Qin, T. and Liu, T. Introducing LETOR 4.0 datasets. *CoRR*, abs/1306.2597, 2013.
- Quinlan, J. R. Induction of decision trees. *Machine Learning*, 1986.
- Quinlan, J. R. C4. 5: Programs for machine learning. 1993.
- Rajeswaran, A., Finn, C., Kakade, S., and Levine, S. Meta-learning with implicit gradients. In *Proc. of NeurIPS*, 2019.
- Rudin, C. and Ertekin, Ş. Learning customized and optimized lists of rules with mathematical programming. *Mathematical Programming Computation*, 2018.
- Stoyanov, V., Ropson, A., and Eisner, J. Empirical risk minimization of graphical model parameters given approximate inference, decoding, and model structure. In *Proc. of AISTATS*, 2011.
- Tanno, R., Arulkumaran, K., Alexander, D., Criminisi, A., and Nori, A. Adaptive neural trees. In *Proc. of ICML*, 2019.
- Verwer, S. and Zhang, Y. Learning optimal classification trees using a binary linear program formulation. In *Proc. of AAAI*, 2019.
- Widyantoro, D. H., Ioerger, T. R., and Yen, J. An incremental approach to building a cluster hierarchy. In *Proc. of ICDM*, 2002.
- Williams, A., Drozdov, A., and Bowman, S. R. Do latent tree learning models identify meaningful structure in sentences? *TACL*, 2018.
- Yang, Y., Morillo, I. G., and Hospedales, T. M. Deep neural decision trees. *preprint arXiv:1806.06988*, 2018.
- Yogatama, D., Blunsom, P., Dyer, C., Grefenstette, E., and Ling, W. Learning to compose words into sentences with reinforcement learning. In *Proc. of ICLR*, 2017.
- Yu, Y.-L. and Xing, E. P. Exact algorithms for isotonic regression and related. In *Journal of Physics: Conference Series*, 2016.

Zalinescu, C. *Convex analysis in general vector spaces*.
World Scientific, 2002.

Zhang, T., Ramakrishnan, R., and Livny, M. Birch: A
new data clustering algorithm and its applications. *Data
Mining and Knowledge Discovery*, 1997.

Zhu, H., Murali, P., Phan, D. T., Nguyen, L. M., and
Kalagnanam, J. A scalable mip-based method for learning
optimal multivariate decision trees. In *Proc. of NeurIPS*,
2020.