
Born-Again Tree Ensembles

Thibaut Vidal¹ Maximilian Schiffer²

Abstract

The use of machine learning algorithms in finance, medicine, and criminal justice can deeply impact human lives. As a consequence, research into interpretable machine learning has rapidly grown in an attempt to better control and fix possible sources of mistakes and biases. Tree ensembles offer a good prediction quality in various domains, but the concurrent use of multiple trees reduces the interpretability of the ensemble. Against this background, we study born-again tree ensembles, i.e., the process of constructing a single decision tree of minimum size that reproduces the exact same behavior as a given tree ensemble in its entire feature space. To find such a tree, we develop a dynamic-programming based algorithm that exploits sophisticated pruning and bounding rules to reduce the number of recursive calls. This algorithm generates optimal born-again trees for many datasets of practical interest, leading to classifiers which are typically simpler and more interpretable without any other form of compromise.

1. Introduction

Tree ensembles constitute a core technique for prediction and classification tasks. Random forests (Breiman, 2001) and boosted trees (Friedman, 2001) have been used in various application fields, e.g., in medicine for recurrence risk prediction and image classification, in criminal justice for custody decisions, or in finance for credit risk evaluation. Although tree ensembles offer a high prediction quality, distorted predictions in high-stakes decisions can be exceedingly harmful. Here, interpretable machine learning models are essential to understand potential distortions and biases. Research in this domain has significantly increased (Mur-

doch et al., 2019) with numerous works focusing on the construction of optimal sparse trees (Hu et al., 2019) or on the interpretability of neural networks (Zhang et al., 2018; Melis & Jaakkola, 2018).

Currently, there exists a trade-off between the interpretability and the performance of tree (ensemble) classifiers. Single decision trees (e.g., those produced by CART) are well-known for their interpretability, whereas tree ensembles and gradient boosting approaches allow for high prediction quality but are generally more opaque and redundant. Against this background, we study born-again tree ensembles in a similar notion as born-again trees (see, Breiman & Shang, 1996), and search for a simpler classifier that faithfully reproduces the behavior of a tree ensemble.

Formally, let $(\mathbf{X}, \mathbf{y}) = \{\mathbf{x}_i, y_i\}_{i=1}^n$ be a training set in which each $\mathbf{x}_i \in \mathbb{R}^p$ is a p -dimensional numerical feature vector, and each $y_i \in \mathbb{N}$ is its associated class. Each sample of this training set has been independently drawn from an unknown distribution $(\mathcal{X}, \mathcal{Y})$. Based on this training set, a tree ensemble \mathcal{T} learns a function $F_{\mathcal{T}} : \mathcal{X} \rightarrow \mathcal{Y}$ that predicts y_i for each \mathbf{x}_i drawn from \mathcal{X} . With this notation, we state Problem 1, which is the core of our studies.

Problem 1 (Born-again tree ensemble) *Given a tree ensemble \mathcal{T} , we search for a decision tree T of **minimal size** that is **faithful** to \mathcal{T} , i.e., such that $F_T(\mathbf{x}) = F_{\mathcal{T}}(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^p$.*

We note that the condition $F_T(\mathbf{x}) = F_{\mathcal{T}}(\mathbf{x})$ applies to the entire feature space. Indeed, our goal is to faithfully reproduce the decision function of the tree ensemble for all possible inputs in \mathcal{X} . In other words, we are looking for a *new representation of the same classifier*. Problem 1 depends on the definition of a *size* metric. In this study, we refer to the size of a tree either as its *depth* (D) or its *number of leaves* (L). Additionally, we study a hierarchical objective (DL) which optimizes depth in priority and then the number of leaves. For brevity, we detail the methodology for the *depth* objective (D) in the main paper. The supplementary material contains the algorithmic adaptations needed to cover the other objectives, rigorous proofs for all theorems, as well as additional illustrations and experimental results.

Theorem 1 states the computational complexity of Problem 1.

¹Department of Computer Science, Pontifical Catholic University of Rio de Janeiro (PUC-Rio), Rio de Janeiro, Brazil. ²TUM School of Management, Technical University of Munich, Munich, Germany. Correspondence to: Thibaut Vidal <vidalt@inf.puc-rio.br>, Maximilian Schiffer <schiffer@tum.de>.

Theorem 1 *Problem 1 is NP-hard when optimizing depth, number of leaves, or any hierarchy of these two objectives.*

This result uses a direct reduction from 3-SAT. Actually, the same proof shows that the sole fact of verifying the faithfulness of a solution is NP-hard. In this work, we show that despite this intractability result, Problem 1 can be solved *to proven optimality* for various datasets of practical interest, and that the solution of this problem permits significant advances regarding tree ensemble simplification, interpretation, and analysis.

1.1. State of the Art

Our work relates to the field of interpretable machine learning, especially thinning tree ensembles and optimal decision tree construction. We review these fields concisely and refer to Guidotti et al. (2018), Murdoch et al. (2019) and Rudin (2019) for surveys and discussions on interpretable machine learning, as well as to Rokach (2016) for an overview on general work on decision forests.

Thinning tree ensembles has been studied from different perspectives and divides in two different streams, *i*) classical thinning of a tree ensemble by removing some weak learners from the original ensemble and *ii*) replacing a tree ensemble by a simpler classifier, e.g., a single decision tree.

Early works on thinning focused on finding reduced ensembles which yield a prediction quality comparable to the full ensemble (Margineantu & Dietterich, 1997). Finding such reduced ensembles has been proven to be NP-hard (Tamon & Xiang, 2000) and in some cases reduced ensembles may even outperform the full ensemble (Zhou et al., 2002). While early works proposed a static thinning, dynamic thinning algorithms that store the full ensemble but dynamically query only a subset of the trees have been investigated by Hernández-Lobato et al. (2009), Park & Furnkranz (2012), and Martínez-Muñoz et al. (2008). For a detailed discussion on this stream of research we refer to Rokach (2016), who discusses the development of ranking-based methods (see, e.g., Prodromidis et al., 1999; Caruana et al., 2004; Banfield et al., 2005; Hu et al., 2007; Partalas et al., 2010; Rokach, 2009; Zhang & Wang, 2009) and search-based methods (see, e.g., Prodromidis & Stolfo, 2001; Winderatt & Ardeshir, 2001; Zhou et al., 2002; Zhou & Tang, 2003; Rokach et al., 2006; Zhang et al., 2006).

In their seminal work about born-again trees, Breiman & Shang (1996) were the first to introduce a thinning problem that aimed at replacing a tree ensemble by a newly constructed simpler classifier. Here, they used a tree ensemble to create a data set which is then used to build a born-again tree with a prediction accuracy close to the accuracy of the tree ensemble. Ensuing work followed three different concepts. Meinshausen (2010) introduced the concept of

node harvesting, i.e., reducing the number of decision nodes to generate an interpretable tree. Recent works along this line used tree space prototypes to sparsen a tree (Tan et al., 2016) or rectified decision trees that use hard and soft labels (Bai et al., 2019). Friedman & Popescu (2008) followed a different concept and proposed a linear model to extract rules from a tree ensemble, which can then be used to rebuild a single tree. Similarly, Sirikulviriyaya & Sinthupinyo (2011) focused on deducing rules from a random forest, while Hara & Hayashi (2016) focused on rule extraction from tree ensembles via bayesian model selection, and Molias et al. (2019) used a local-based, path-oriented similarity metric to select rules from a tree ensemble. Recently, some works focused on directly extracting a single tree from a tree ensemble based on stabilized but yet heuristic splitting criteria (Zhou & Hooker, 2016), genetic algorithms (Vandewiele et al., 2017), or by actively sampling training points (Bastani et al., 2017a;b). All of these works focus on the creation of sparse decision trees that remain interpretable but can be used to replace a tree ensemble while securing a similar prediction performance. However, these approaches do not guarantee faithfulness, such that the new classifier is not guaranteed to retain the same decision function and prediction performance.

In the field of neural networks, related studies were done on *model compression* (Bucilua et al., 2006). The proposed approaches often use knowledge distillation, i.e., using a high-capacity teacher to train a compact student with similar knowledge (see, e.g., Hinton et al., 2015). Recent works focused on creating soft decision trees from a neural network (Frosst & Hinton, 2017), decomposing the gradient in knowledge distillation (Furlanello et al., 2018), deriving a class of models for self-explanatory neural networks (Melis & Jaakkola, 2018), or specified knowledge representations in high conv-layers for interpretable convolutional neural networks (Zhang et al., 2018). Focusing on feed-forward neural networks, Frankle & Carbin (2018) proposed pruning techniques that identify subnetworks which perform close to the original network. Clark et al. (2019) studied born-again multi task networks for natural language processing, while Kisamori & Yamazaki (2019) focused on synthesizing an interpretable simulation model from a neural network. As neural networks are highly non-linear and even less transparent than tree ensembles, all of these approaches remain predominantly heuristic and faithfulness is typically not achievable.

Optimal decision trees. Since the 1990's, some works focused on constructing decision trees based on mathematical programming techniques. Bennett (1992) used linear programming to construct trees with linear combination splits and showed that this technique performs better than conventional univariate split algorithms. Bennett & Blue (1996)

focused on building global optimal decision trees to avoid overfitting, while Nijssen & Fromont (2007) presented an exact algorithm to build a decision tree for specific depth, accuracy, and leaf requirements. Recently, Bertsimas & Dunn (2017) presented a mixed integer programming formulation to construct optimal classification trees. On a similar note, Günlük et al. (2018) presented an integer programming approach for optimal decision trees with categorical data, and Verwer & Zhang (2019) presented a binary linear program for optimal decision trees. Hu et al. (2019) presented a scalable algorithm for optimal sparse binary decision trees. While all these works show that decision trees are in general amenable to be built with optimization techniques, none of these works focused on constructing born-again trees that match the accuracy of a given tree ensemble.

Summary. Thinning problems have been studied for both tree ensembles and neural networks in order to derive interpretable classifiers that show a similar performance than the aforementioned algorithms. However, all of these works embed heuristic construction techniques or an approximative objective, such that the resulting classifiers do not guarantee a behavior and prediction performance equal to the original tree ensemble or neural network. These approaches appear to be plausible for born-again neural networks, as neural networks have highly non-linear structures that cannot be easily captured in an optimization approach. In contrast, work in the field of building optimal decision trees showed that the construction of decision trees is generally amenable for optimization based approaches. Nevertheless, these works focused so far on constructing sparse or optimal trees that outperform heuristically created trees, such that the question whether one could construct an optimal decision tree that serves as a born-again tree ensemble remains open. Answering this question and discussing some of its implications is the focus of our study.

1.2. Contributions

With this work, we revive the concept of born-again tree ensembles and aim to construct a single —minimum-size— tree that faithfully reproduces the decision function of the original tree ensemble. More specifically, our contribution is fourfold. First, we formally define the problem of constructing optimal born-again tree ensembles and prove that this problem is NP-hard. Second, we highlight several properties of this problem and of the resulting born-again tree. These findings allow us to develop a dynamic-programing based algorithm that solves this problem efficiently and constructs an optimal born-again tree out of a tree ensemble. Third, we discuss specific pruning strategies for the born-again tree that allow to reduce redundancies that cannot be identified in the original tree ensemble. Fourth, besides providing theoretical guarantees, we present numerical studies which

allow to analyze the characteristics of the born-again trees in terms of interpretability and accuracy. Further, these studies show that our algorithm is amenable to a wide range of real-world data sets.

We believe that our results and the developed algorithms open a new perspective in the field of interpretable machine learning. With this approach, one can construct simple classifiers that bear all characteristics of a tree ensemble. Besides interpretability gains, this approach casts a new light on tree ensembles and highlights new structural properties.

2. Fundamentals

In this section, we introduce some fundamental definitions. Afterwards, we discuss a worst-case bound on the depth of an optimal born-again tree.

Tree ensemble. We define a tree ensemble \mathcal{T} as a set of trees $t \in \mathcal{T}$ with weights w_t . For any sample \mathbf{x} , the tree ensemble returns the majority vote of its trees: $F_{\mathcal{T}}(\mathbf{x}) = \text{WEIGHTED-MAJORITY}\{(F_t(\mathbf{x}), w_t)\}_{t \in \mathcal{T}}$ (ties are broken in favor of the smaller index).

Cells. Let H_j be the set of all split levels (i.e., hyperplanes) extracted from the trees for each feature j . We can partition the feature space \mathbb{R}^p into cells $\mathcal{S}_{\text{ELEM}} = \{1, \dots, |H_1| + 1\} \times \dots \times \{1, \dots, |H_p| + 1\}$ such that each cell $\mathbf{z} = (z_1, \dots, z_p) \in \mathcal{S}_{\text{ELEM}}$ represents the box contained between the $(z_j - 1)^{\text{th}}$ and z_j^{th} hyperplanes for each feature $j \in \{1, \dots, p\}$. Cells such that $z_j = 1$ (or $z_j = |H_j| + 1$) extend from $-\infty$ (or to ∞ , respectively) along dimension j . We note that the decision function of the tree ensemble $F_{\mathcal{T}}(\mathbf{z})$ is constant in the interior of each cell \mathbf{z} , allowing us to exclusively use the hyperplanes of $\{H_j\}_{j=1}^d$ to construct an optimal born-again tree.

Regions. We define a *region* of the feature space as a pair $(\mathbf{z}^L, \mathbf{z}^R) \in \mathcal{S}_{\text{ELEM}}^2$ such that $\mathbf{z}^L \leq \mathbf{z}^R$. Region $(\mathbf{z}^L, \mathbf{z}^R)$ encloses all cells \mathbf{z} such that $\mathbf{z}^L \leq \mathbf{z} \leq \mathbf{z}^R$. Let $\mathcal{S}_{\text{REGIONS}}$ be the set of all regions. An optimal born-again tree T for a region $(\mathbf{z}^L, \mathbf{z}^R)$ is a tree of minimal size such that $F_T(\mathbf{x}) = F_{\mathcal{T}}(\mathbf{x})$ within this region.




Figure 1. Example of a cell, region and splitting hyperplane

Figure 1 depicts a cell and a region on a two-dimensional feature space. We also provide a more extensive example of the born-again tree generation process in the supplementary material. The number of cells and regions increases rapidly with the number of hyperplanes and features, formally:

$$|\mathcal{S}_{\text{ELEM}}| = \prod_{j=1}^p (|H_j| + 1) \quad (1)$$

$$|\mathcal{S}_{\text{REGION}}| = \prod_{j=1}^p \frac{(|H_j| + 1)(|H_j| + 2)}{2}. \quad (2)$$

Moreover, Theorem 2 gives initial bounds on the size of the born-again decision tree.

Theorem 2 *The depth of an optimal born-again tree T satisfies $\Phi(T) \leq \sum_{t \in \mathcal{T}} \Phi(t)$, where $\Phi(t)$ represents the depth of a tree t . This bound is tight.*

This bound corresponds to a worst case behavior which is usually attained only on purposely-designed pathological cases. As highlighted in our computational experiments, the average tree depth remains generally lower than this analytical worst case. Beyond interpretability benefits, the tree depth represents the number of sequential operations (hyperplane comparisons) needed to determine the class of a given sample during the test stage. Therefore, an optimized born-again tree is not only more interpretable, but it also requires less test effort, with useful applications for classification in embarked systems, typically occurring within limited time and processing budgets.

3. Methodology

In this section, we introduce a dynamic programming (DP) algorithm which optimally solves Problem 1 for many data sets of practical interest. Let $\Phi(\mathbf{z}^L, \mathbf{z}^R)$ be the depth of an optimal born-again decision tree for a region $(\mathbf{z}^L, \mathbf{z}^R) \in \mathcal{S}_{\text{REGION}}$. We can then limit the search to optimal born-again trees whose left and right sub-trees represent optimal born-again trees for the respective sub-regions. Hence, we can recursively decompose a larger problem into subproblems using

$$\Phi(\mathbf{z}^L, \mathbf{z}^R) = \begin{cases} 0 & \text{if } \text{ID}(\mathbf{z}^L, \mathbf{z}^R) \\ \min_{1 \leq j \leq p} \left\{ \min_{z_j^L \leq l < z_j^R} \left\{ 1 + \max\{\Phi(\mathbf{z}^L, \mathbf{z}_{jl}^R), \Phi(\mathbf{z}_{jl}^L, \mathbf{z}^R)\} \right\} \right\}, & \end{cases} \quad (3)$$

in which $\text{ID}(\mathbf{z}^L, \mathbf{z}^R)$ takes value TRUE if and only if all cells \mathbf{z} such that $\mathbf{z}^L \leq \mathbf{z} \leq \mathbf{z}^R$ admit the same weighted majority class. In this equation, $\mathbf{z}_{jl}^R = \mathbf{z}^R + \mathbf{e}_j(l - z_j^R)$ represents the ‘‘top right’’ corner of the left region obtained in the

subdivision, and $\mathbf{z}_{jl}^L = \mathbf{z}^L + \mathbf{e}_j(l + 1 - z_j^L)$ is the ‘‘bottom left’’ corner of the right region obtained in the subdivision.

While Equation (3) bears the main rationale of our algorithm, it suffers in its basic state from two main weaknesses that prevent its translation into an efficient algorithm: firstly, each verification of the first condition (i.e., the base case) requires evaluating whether $\text{ID}(\mathbf{z}^L, \mathbf{z}^R)$ is true and possibly requires the evaluation of the majority class on an exponential number of cells *if done brute force*. Secondly, the recursive call considers all possible hyperplanes within the region to find the minimum over $j \in \{1, \dots, p\}$ and $z_j^L \leq l < z_j^R$. In the following, we propose strategies to mitigate both drawbacks.

To avoid the evaluation of $\text{ID}(\mathbf{z}^L, \mathbf{z}^R)$ by inspection, we integrate this evaluation within the recursion to profit from the memory structures of the DP algorithm. With these changes, the recursion becomes:

$$\begin{aligned} \Phi(\mathbf{z}^L, \mathbf{z}^R) = & \quad (4) \\ \min_j \left\{ \min_{z_j^L \leq l < z_j^R} \left\{ \mathbb{1}_{jl}(\mathbf{z}^L, \mathbf{z}^R) + \max\{\Phi(\mathbf{z}^L, \mathbf{z}_{jl}^R), \Phi(\mathbf{z}_{jl}^L, \mathbf{z}^R)\} \right\} \right\} \\ \text{where } \mathbb{1}_{jl}(\mathbf{z}^L, \mathbf{z}^R) = & \begin{cases} 0 & \text{if } \Phi(\mathbf{z}^L, \mathbf{z}_{jl}^R) = \Phi(\mathbf{z}_{jl}^L, \mathbf{z}^R) = 0 \\ & \text{and } F_{\mathcal{T}}(\mathbf{z}^L) = F_{\mathcal{T}}(\mathbf{z}^R); \\ 1 & \text{otherwise.} \end{cases} \end{aligned}$$

To limit the number of recursive calls, we can filter out for each dimension j any hyperplane $l \in \{1, \dots, |H_j|\}$ such that $F_{\mathcal{T}}(\mathbf{z}) = F_{\mathcal{T}}(\mathbf{z} + \mathbf{e}_j)$ for all \mathbf{z} such that $\mathbf{z}_j = l$, and exploit two additional properties of the problem.

Theorem 3 *Let $(\mathbf{z}^L, \mathbf{z}^R)$ and $(\bar{\mathbf{z}}^L, \bar{\mathbf{z}}^R)$ be two regions such that $\mathbf{z}^L \leq \bar{\mathbf{z}}^L \leq \bar{\mathbf{z}}^R \leq \mathbf{z}^R$, then $\Phi(\bar{\mathbf{z}}^L, \bar{\mathbf{z}}^R) \leq \Phi(\mathbf{z}^L, \mathbf{z}^R)$.*

Theorem 3 follows from the fact that any feasible tree satisfying $F_T(\mathbf{x}) = F_{\mathcal{T}}(\mathbf{x})$ on a region $(\mathbf{z}^L, \mathbf{z}^R)$ also satisfies this condition for any subregion $(\bar{\mathbf{z}}^L, \bar{\mathbf{z}}^R)$. Therefore, $\Phi(\bar{\mathbf{z}}^L, \bar{\mathbf{z}}^R)$ constitutes a lower bound of $\Phi(\mathbf{z}^L, \mathbf{z}^R)$. Combining this bound with Equation (4), we get

$$\begin{aligned} & \max\{\Phi(\mathbf{z}^L, \mathbf{z}_{jl}^R), \Phi(\mathbf{z}_{jl}^L, \mathbf{z}^R)\} \\ & \leq \Phi(\mathbf{z}^L, \mathbf{z}^R) \\ & \leq \mathbb{1}_{jl}(\mathbf{z}^L, \mathbf{z}^R) + \max\{\Phi(\mathbf{z}^L, \mathbf{z}_{jl}^R), \Phi(\mathbf{z}_{jl}^L, \mathbf{z}^R)\} \end{aligned}$$

for each $j \in \{1, \dots, p\}$ and $z_j^L \leq l < z_j^R$.

This result will be fundamental to use bounding techniques and therefore save numerous recursions during the DP algorithm. With Theorem 4, we can further reduce the number of candidates in each recursion.

Theorem 4 *Let $j \in \{1, \dots, p\}$ and $l \in \{z_j^L, \dots, z_j^R - 1\}$.*

- If $\Phi(\mathbf{z}^L, \mathbf{z}_{j_l}^R) \geq \Phi(\mathbf{z}_{j_l}^L, \mathbf{z}^R)$ then $\forall l' > l$

$$\mathbb{1}_{j_l}(\mathbf{z}^L, \mathbf{z}^R) + \max\{\Phi(\mathbf{z}^L, \mathbf{z}_{j_l}^R), \Phi(\mathbf{z}_{j_l}^L, \mathbf{z}^R)\}$$

$$\leq \mathbb{1}_{j_{l'}}(\mathbf{z}^L, \mathbf{z}^R) + \max\{\Phi(\mathbf{z}^L, \mathbf{z}_{j_{l'}}^R), \Phi(\mathbf{z}_{j_{l'}}^L, \mathbf{z}^R)\}$$
- If $\Phi(\mathbf{z}^L, \mathbf{z}_{j_l}^R) \leq \Phi(\mathbf{z}_{j_l}^L, \mathbf{z}^R)$ then $\forall l' < l$

$$\mathbb{1}_{j_l}(\mathbf{z}^L, \mathbf{z}^R) + \max\{\Phi(\mathbf{z}^L, \mathbf{z}_{j_l}^R), \Phi(\mathbf{z}_{j_l}^L, \mathbf{z}^R)\}$$

$$\leq \mathbb{1}_{j_{l'}}(\mathbf{z}^L, \mathbf{z}^R) + \max\{\Phi(\mathbf{z}^L, \mathbf{z}_{j_{l'}}^R), \Phi(\mathbf{z}_{j_{l'}}^L, \mathbf{z}^R)\}.$$

Based on Theorem 4, we can discard all hyperplane levels $l' > l$ in Equation (4) if $\Phi(\mathbf{z}^L, \mathbf{z}_{j_l}^R) \geq \Phi(\mathbf{z}_{j_l}^L, \mathbf{z}^R)$. The same argument holds when $\Phi(\mathbf{z}^L, \mathbf{z}_{j_l}^R) \leq \Phi(\mathbf{z}_{j_l}^L, \mathbf{z}^R)$ with $l' < l$. We note that the two cases of Theorem 4 are not mutually exclusive. No other recursive call is needed for the considered feature when an equality occurs. Otherwise, at least one case holds, allowing us to search the range $l \in \{z_j^L, \dots, z_j^R - 1\}$ in Equation (4) by binary search with only $O(\log(z_j^R - z_j^L))$ subproblem calls.

General algorithm structure. The DP algorithm presented in Algorithm 1 capitalizes upon all the aforementioned properties. It is initially launched on the region representing the complete feature space, by calling BORN-AGAIN($\mathbf{z}^L, \mathbf{z}^R$) with $\mathbf{z}^L = (1, \dots, 1)^\top$ and $\mathbf{z}^R = (|H_1| + 1, \dots, |H_p| + 1)^\top$.

Firstly, the algorithm checks whether it attained a base case in which the region $(\mathbf{z}^L, \mathbf{z}^R)$ is restricted to a single cell (Line 1). If this is the case, it returns an optimal depth of zero corresponding to a single leaf, otherwise it tests whether the result of the current subproblem defined by region $(\mathbf{z}^L, \mathbf{z}^R)$ is not yet in the DP memory (Line 2). If this is the case, it directly returns the known result.

Past these conditions, the algorithm starts enumerating possible splits and opening recursions to find the minimum of Equation (4). By Theorem 4 and the related discussions, it can use a binary search for each feature to save many possible evaluations (Lines 9 and 10). By Theorem 3, the exploitation of lower and upper bounds on the optimal solution value (Lines 7, 9, 20, and 21) allows to stop the iterative search whenever no improving solution can exist. Finally, the special case of Lines 13 and 14 covers the case in which $\Phi(\mathbf{z}^L, \mathbf{z}_{j_l}^R) = \Phi(\mathbf{z}_{j_l}^L, \mathbf{z}^R) = 0$ and $F_{\mathcal{T}}(\mathbf{z}^L) = F_{\mathcal{T}}(\mathbf{z}^R)$, corresponding to a homogeneous region in which all cells have the same majority class. As usual in DP approaches, our algorithm memorizes the solutions of sub-problems and reuses them in future calls (Lines 15, 17, and 26).

We observe that this algorithm maintains the optimal solution of each subproblem in memory, but not the solution itself in order to reduce memory consumption. Retrieving the solution after completing the DP can be done with a simple inspection of the final states and solutions, as detailed in the supplementary material.

The maximum number of possible regions is $|\mathcal{S}_{\text{REGION}}| = \prod_j \frac{(|H_j|+1)(|H_j|+2)}{2}$ (Equation 2) and each call to BORN-AGAIN takes up to $O(\sum_j \log |H_j|)$ elementary operations due to Theorem 4, leading to a worst-case complexity of $O(|\mathcal{S}_{\text{REGION}}| \sum_j \log |H_j|)$ time for the overall recursive algorithm. Such an exponential complexity is expectable for an NP-hard problem. Still, as observed in our experiments, the number of regions explored with the bounding strategies is much smaller in practice than the theoretical worst case.

Algorithm 1 BORN-AGAIN($\mathbf{z}^L, \mathbf{z}^R$)

```

1: if ( $\mathbf{z}^L = \mathbf{z}^R$ ) return 0
2: if ( $\mathbf{z}^L, \mathbf{z}^R$ ) exists in memory then
3:   return MEMORY( $\mathbf{z}^L, \mathbf{z}^R$ )
4: end if
5: UB  $\leftarrow$   $\infty$ 
6: LB  $\leftarrow$  0
7: for  $j = 1$  to  $p$  and LB < UB do
8:   (LOW, UP)  $\leftarrow$  ( $z_j^L, z_j^R$ )
9:   while LOW < UP and LB < UB do
10:     $l \leftarrow \lfloor (\text{LOW} + \text{UP})/2 \rfloor$ 
11:     $\Phi_1 \leftarrow$  BORN-AGAIN( $\mathbf{z}^L, \mathbf{z}^R + \mathbf{e}_j(l - z_j^R)$ )
12:     $\Phi_2 \leftarrow$  BORN-AGAIN( $\mathbf{z}^L + \mathbf{e}_j(l + 1 - z_j^L), \mathbf{z}^R$ )
13:    if ( $\Phi_1 = 0$ ) and ( $\Phi_2 = 0$ ) then
14:      if  $f(\mathbf{z}^L, \mathcal{T}) = f(\mathbf{z}^R, \mathcal{T})$  then
15:        MEMORIZE( $(\mathbf{z}^L, \mathbf{z}^R), 0$ ) and return 0
16:      else
17:        MEMORIZE( $(\mathbf{z}^L, \mathbf{z}^R), 1$ ) and return 1
18:      end if
19:    end if
20:    UB  $\leftarrow$   $\min\{\text{UB}, 1 + \max\{\Phi_1, \Phi_2\}\}$ 
21:    LB  $\leftarrow$   $\max\{\text{LB}, \max\{\Phi_1, \Phi_2\}\}$ 
22:    if ( $\Phi_1 \geq \Phi_2$ ) then UP  $\leftarrow$   $l$ 
23:    if ( $\Phi_1 \leq \Phi_2$ ) then LOW  $\leftarrow$   $l + 1$ 
24:  end while
25: end for
26: MEMORIZE( $(\mathbf{z}^L, \mathbf{z}^R), \text{UB}$ ) and return UB
    
```

4. Computational Experiments

The goal of our computational experiments is fourfold:

1. Evaluating the computational performance of the proposed DP algorithm as a function of the data set characteristics, e.g., the size metric in use, the number of trees in the original ensemble, and the number of samples and features in the datasets.
2. Studying the structure and complexity of the born-again trees for different size metrics.
3. Measuring the impact of a simple pruning strategy applied on the resulting born-again trees.
4. Proposing and evaluating a fast heuristic algorithm to find faithful born-again trees.

The DP algorithm was implemented in C++ and compiled with GCC 9.2.0 using flag `-O3`, whereas the original random forests were generated in Python (using scikit-learn v0.22.1). All our experiments were run on a single thread of an Intel(R) Xeon(R) CPU E5-2620v4 2.10GHz, with 128GB of available RAM, running CentOS v7.7. In the remainder of this section, we discuss the preparation of the data and then describe each experiment. Detailed computational results, data, and source codes are available in the supplementary material and at the following address: <https://github.com/vidalt/BA-Trees>.

4.1. Data Preparation

We focus on a set of six datasets from the UCI machine learning repository [UCI] and from previous work by Smith et al. (1988) [SmithEtAl] and Hu et al. (2019) [HuEtAl] for which using random forests (with ten trees) showed a significant improvement upon stand-alone CART. The characteristics of these datasets are summarized in Table 1: number of samples n , number of features p , number of classes K and class distribution CD. To obtain discrete numerical features, we used one-hot encoding on categorical data and binned continuous features into ten ordinal scales. Then, we generated training and test samples for all data sets using a ten-fold cross validation. Finally, for each fold and each dataset, we generated a random forest composed of ten trees with a maximum depth of three (i.e., eight leaves at most), considering $p/2$ random candidate features at each split. This random forest constitutes the input to our DP algorithm.

Table 1. Characteristics of the data sets

Data set	n	p	K	CD	Src.
BC: Breast-Cancer	683	9	2	65-35	UCI
CP: COMPAS	6907	12	2	54-46	HuEtAl
FI: FICO	10459	17	2	52-48	HuEtAl
HT: HTRU2	17898	8	2	91-9	UCI
PD: Pima-Diabetes	768	8	2	65-35	SmithEtAl
SE: Seeds	210	7	3	33-33-33	UCI

4.2. Computational Effort

In a first analysis, we evaluate the computational time of Algorithm 1 for different data sets and size metrics. Figure 2 reports the results of this experiment as a box-whisker plot, in which each box corresponds to ten runs (one for each fold) and the whiskers extend to 1.5 times the interquartile range. Any sample beyond this limit is reported as outlier and noted with a “o”. D denotes a depth-minimization objective, whereas L refers to the minimization of the number of leaves, and DL refers to the hierarchical objective which prioritizes the smallest depth, and then the smallest number of leaves. As can be seen, constructing a born-again tree with




Figure 2. Computational times to construct a born-again tree from a random forest with 10 trees and depth 3, for each objective (D/L/DL) and data set

objective D yields significantly lower computational times compared to using objectives L and DL. Indeed, the binary search technique resulting from Theorem 4 only applies to objective D, leading to a reduced number of recursive calls in this case compared to the other algorithms.

In our second analysis, we focus on the FICO case and randomly extract subsets of samples and features to produce smaller data sets. We then measure the computational effort of Algorithm 1 for metric D (depth optimization) as a function of the number of features ($p \in \{2, 3, 5, 7, 10, 12, 15, 17\}$), the number of samples ($n \in \{250, 500, 750, 1000, 2500, 5000, 7500, 10459\}$), and the number of trees in the original random forest ($T \in \{3, 5, 7, 10, 12, 15, 17, 20\}$). Figure 3 reports the results of this experiment. Each boxplot corresponds to ten runs, one for each fold.

We observe that the computational time of the DP algorithm is strongly driven by the number of features, with an exponential growth relative to this parameter. This result is in line with the complexity analysis of Section 3. The number of trees influences the computational time significantly less. Surprisingly, the computational effort of the algorithm actually decreases with the number of samples. This is due to the fact that with more sample information, the decisions of the individual trees of the random forest are less varied, leading to fewer distinct hyperplanes and therefore to fewer possible states in the DP.

4.3. Complexity of the Born-Again Trees

We now analyze the depth and number of leaves of the born-again trees for different objective functions and datasets in Table 2.

As can be seen, the different objectives can significantly




Figure 3. Growth of the computational time (in milliseconds) of Algorithm 1 as a function of the number of samples, features and trees. In each experiment, the parameters which are not under scrutiny are fixed to their baseline values of $n = 2.5 \times 10^3$, $p = 10$ and $T = 10$.

Table 2. Depth and number of leaves of the born-again trees

Data set	D		L		DL	
	Depth	# Leaves	Depth	# Leaves	Depth	# Leaves
BC	12.5	2279.4	18.0	890.1	12.5	1042.3
CP	8.9	119.9	8.9	37.1	8.9	37.1
FI	8.6	71.3	8.6	39.2	8.6	39.2
HT	6.0	20.2	6.3	11.9	6.0	12.0
PD	9.6	460.1	15.0	169.7	9.6	206.7
SE	10.2	450.9	13.8	214.6	10.2	261.0
Avg.	9.3	567.0	11.8	227.1	9.3	266.4

influence the outcome of the algorithm. For several data sets, the optimal depth of the born-again tree is reached with any objective, as an indirect consequence of the minimization of the number of leaves. In other cases, however, prioritizing the minimization of the number of leaves may generate 50% deeper trees for some data sets (e.g., PD). The hierarchical objective DL succeeds in combining the benefits of both objectives. It generates a tree with minimum depth and with a number of leaves which is usually close to the optimal one from objective L.

4.4. Post-Pruned Born-Again Trees

Per definition, the born-again tree reproduces the same exact behavior as the majority class of the original ensemble classifier *on all regions* of the feature space \mathcal{X} . Yet, some regions of \mathcal{X} may not contain any training sample, either due to data scarcity or simply due to incompatible feature values (e.g., “sex = MALE” and “pregnancy = TRUE”). These regions may also have non-homogeneous majority classes from the tree ensemble viewpoint due to the combinations of decisions from multiple trees. The born-again tree, however, is agnostic to this situation and imitates the original classifi-

cation within all the regions, leading to some splits which are mere artifacts of the ensemble’s behavior but never used for classification.

To circumvent this issue, we suggest to apply a simple post-pruning step to eliminate inexpressive tree sub-regions. We therefore verify, from bottom to top, whether both sides of each split contain at least one training sample. Any split which does not fulfill this condition is pruned and replaced by the child node of the branch that contains samples. The complete generation process, from the original random forest to the *pruned* born-again tree is illustrated in Figure 4. In this simple example, it is noteworthy that the born-again tree uses an optimal split at the root node which is different from all root splits in the ensemble. We also clearly observe the role of the post-pruning step, which contributes to eliminate a significant part of the tree.

To observe the impact of the post-pruning on a larger range of datasets, Table 3 reports the total number of leaves of the random forests, as well as the average depth and number of leaves of the born-again trees before and after post-pruning. As previously, the results are averaged over the ten folds. As can be seen, post-pruning significantly reduces the size of the born-again trees, leading to a final number of leaves which is, on average, smaller than the total number of leaves in the original tree ensemble. This indicates a significant gain of simplicity and interpretability.

However, post-pruning could cause a difference of behavior between the original tree ensemble classifier and the final pruned born-again tree. To evaluate whether this filtering had any significant impact on the classification performance of the born-again tree, we finally compare the out-of-sample accuracy (Acc.) and F1 score of the three classifiers in Table 4.

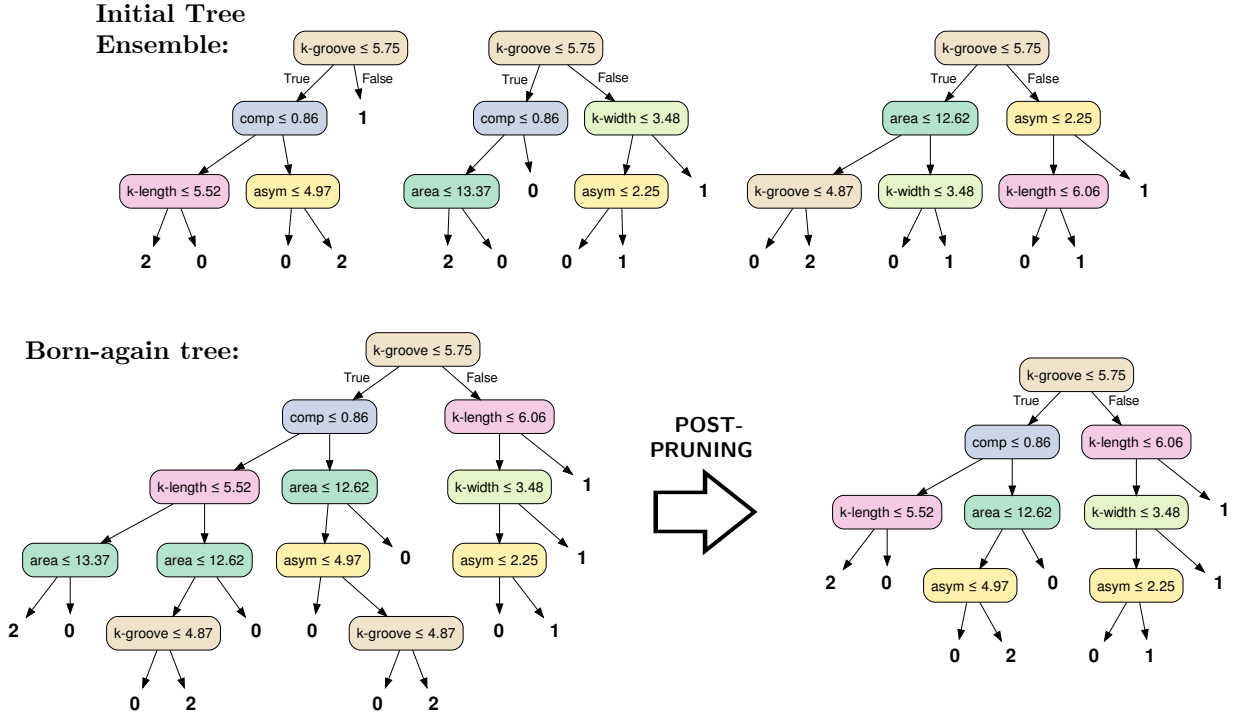

 Figure 4. Example of a post-pruned born-again tree on the *Seeds* data set

Table 3. Comparison of depth and number of leaves

Data set	Random Forest		Born-Again		BA + Pruning	
	# Leaves	Depth	# Leaves	Depth	# Leaves	Depth
BC	61.1	12.5	2279.4	9.1	35.9	9.1
CP	46.7	8.9	119.9	7.0	31.2	7.0
FI	47.3	8.6	71.3	6.5	15.8	6.5
HT	42.6	6.0	20.2	5.1	13.2	5.1
PD	53.7	9.6	460.1	9.4	79.0	9.4
SE	55.7	10.2	450.9	7.5	21.5	7.5
Avg.	51.2	9.3	567.0	7.4	32.8	7.4

Table 4. Accuracy and F1 score comparison

Data set	Random Forest		Born-Again		BA + Pruning	
	Acc.	F1	Acc.	F1	Acc.	F1
BC	0.953	0.949	0.953	0.949	0.946	0.941
CP	0.660	0.650	0.660	0.650	0.660	0.650
FI	0.697	0.690	0.697	0.690	0.697	0.690
HT	0.977	0.909	0.977	0.909	0.977	0.909
PD	0.746	0.692	0.746	0.692	0.750	0.700
SE	0.790	0.479	0.790	0.479	0.790	0.481
Avg.	0.804	0.728	0.804	0.728	0.803	0.729

First of all, the results of Table 4 confirm the faithfulness of our algorithm, as they verify that the prediction quality of the random forests and the born-again tree ensembles are identical. This was expected *per definition* of Problem 1. Furthermore, only marginal differences were observed between the out-of-sample performance of the born-again tree with pruning and the other classifiers. For the considered datasets, pruning contributed to eliminate inexpressive regions of the tree without much impact on classification performance.

4.5. Heuristic Born-Again Trees

As Problem 1 is NP-hard, the computational time of our algorithm will eventually increase exponentially with the number of features (see Figure 3). This is due to the increasing number of recursions, and to the challenge of testing homogeneity for regions without exploring all cells. Indeed, even proving that a given region is homogeneous (i.e., that it contains cells of the same class) remains NP-hard, although it is solvable in practice using integer programming techniques. Accordingly, we take a first step towards scalable heuristic algorithms in the following. We therefore explain how our born-again tree construction algorithm can be modified to preserve the faithfulness guarantee while achieving only a heuristic solution in terms of size.

We made the following adaptations to Algorithm 1 to derive its heuristic counterpart. For each considered region

(z^L, z^R) , we proceed as follows.

1. Instead of evaluating all possible splits and opening recursions, we randomly select $N_c = 1000$ cells in the region and pick the splitting hyperplane that maximizes the information gain.
2. If all these cells belong to the same class, we rely on an integer programming solver to prove whether the region is homogeneous or not. If the region is homogeneous, we define a leaf. Otherwise, we have detected a violating cell, and continue splitting until all regions are homogeneous to guarantee faithfulness.

With these adaptations, the heuristic algorithm finds born-again trees that are guaranteed to be faithful but not necessarily minimal in size. Table 5 compares the computational time of the optimal born-again tree algorithm using objective D “ $T_D(s)$ ”, objective L “ $T_L(s)$ ” with that of the heuristic algorithm “ $T_H(s)$ ”. It also reports the percentage gap of the heuristic tree depth “ $Gap_D(\%)$ ” and number of leaves “ $Gap_L(\%)$ ” relative to the optimal solution values of each objective.

Table 5. Computational time and optimality gap of the heuristic born-again tree algorithm

Data set	$T_D(s)$	$T_L(s)$	$T_H(s)$	$Gap_D(\%)$	$Gap_L(\%)$
BC	39.09	1381.45	1.14	44.80	48.37
CP	0.01	0.10	0.04	0.00	4.85
FI	0.05	0.23	0.03	0.00	1.79
HT	0.01	0.01	0.01	8.33	10.92
PD	0.91	17.95	0.19	44.79	25.63
SE	0.37	5.96	0.24	37.25	29.03
Avg.	6.74	234.28	0.28	22.53	20.10

As visible in these experiments, the CPU time of the heuristic algorithm is significantly smaller than that of the optimal method, at the expense of an increase in tree depth and number of leaves, by 22.53% and 20.10% on average, respectively. To test the limits of this heuristic approach, we also verified that it could run in a reasonable amount of time (faster than a minute) on larger datasets such as Ionosphere, Spambase, and Miniboone (the latter with over 130,000 samples and 50 features).

5. Conclusions

In this paper, we introduced an efficient algorithm to transform a random forest into a single, smallest possible, decision tree. Our algorithm is optimal, and provably returns a single tree with the same decision function as the original tree ensemble. In brief, we obtain a different representation of the same classifier, which helps us to analyze random forests from a different angle. Interestingly, when investigating the structure of the results, we observed that born-again

decision trees contain many inexpressive regions designed to faithfully reproduce the behavior of the original ensemble, but which do not contribute to effectively classify samples. It remains an interesting research question to properly understand the purpose of these regions and their contribution to the generalization capabilities of random forests. In a first simple experiment, we attempted to apply post-pruning on the resulting tree. Based on our experiments on six structurally different datasets, we observed that this pruning does not diminish the quality of the predictions but significantly simplifies the born-again trees. Overall, the final pruned trees represent simple, interpretable, and high-performance classifiers, which can be useful for a variety of application areas.

As a perspective for future work, we recommend to progress further on solution techniques for the born-again tree ensembles problem, proposing new optimal algorithms to effectively handle larger datasets as well as fast and accurate heuristics. Heuristic upper bounds can also be jointly exploited with mathematical programming techniques to eliminate candidate hyperplanes and recursions. Another interesting research line concerns the combination of the dynamic programming algorithm for the construction of the born-again tree with active pruning during construction, leading to a different definition of the recursion and to different base-case evaluations. Finally, we recommend to pursue the investigation of the structural properties of tree ensembles in light of this new born-again tree representation.

Acknowledgements

The authors gratefully thank Toni Pacheco, doctorate student from the computer science department of PUC-Rio, for his extensive help in the methodology design, implementation, and numerical analyses during the paper submission and revision periods. They also express their gratitude to the editors and referees for their insightful recommendations, as well as to Simone Barbosa, Quentin Capart, and Artur Pessoa for rich scientific discussions. This research has been partially supported by CAPES, CNPq [grant number 308528/2018-2] and FAPERJ [grant number E-26/202.790/2019] in Brazil.

References

- Bai, J., Li, Y., Li, J., Jiang, Y., and Xia, S. Rectified decision trees: Towards interpretability, compression and empirical soundness. *arXiv preprint arXiv:1903.05965*, 2019.
- Banfield, R. E., Hall, L. O., Bowyer, K. W., and Kegelmeyer, W. P. Ensemble diversity measures and their application to thinning. *Information Fusion*, 6(1):49–62, 2005.

- Bastani, O., Kim, C., and Bastani, H. Interpretability via model extraction. *arXiv preprint arXiv:1706.09773*, 2017a.
- Bastani, O., Kim, C., and Bastani, H. Interpreting blackbox models via model extraction. *arXiv preprint arXiv:1705.08504*, 2017b.
- Bennett, K. Decision tree construction via linear programming. In *Proceedings of the 4th Midwest Artificial Intelligence and Cognitive Science Society Conference, Utica, Illinois*, 1992.
- Bennett, K. and Blue, J. Optimal decision trees. Technical report, Rensselaer Polytechnique Institute, 1996.
- Bertsimas, D. and Dunn, J. Optimal classification trees. *Machine Learning*, 106(7):1039–1082, 2017.
- Breiman, L. Random forests. *Machine Learning*, 45(1): 5–32, 2001.
- Breiman, L. and Shang, N. Born again trees. Technical report, University of California Berkeley, 1996.
- Buciluă, C., Caruana, R., and Niculescu-Mizil, A. Model compression. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2006.
- Caruana, R., Niculescu-Mizil, A., Crew, G., and Ksikes, A. Ensemble selection from libraries of models. In *Proceedings of the twenty-first International Conference on Machine Learning*, pp. 18, 2004.
- Clark, K., Luong, M.-T., Khandelwal, U., Manning, C. D., and Le, Q. V. Bam! born-again multi-task networks for natural language understanding. *arXiv preprint arXiv:1907.04829*, 2019.
- Frankle, J. and Carbin, M. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018.
- Friedman, J. Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5):1189–1232, 2001.
- Friedman, J. H. and Popescu, B. E. Predictive learning via rule ensembles. *The Annals of Applied Statistics*, 2(3): 916–954, 2008.
- Frosst, N. and Hinton, G. Distilling a neural network into a soft decision tree. *arXiv preprint arXiv:1711.09784*, 2017.
- Furlanello, T., Lipton, Z. C., Tschannen, M., Itti, L., and Anandkumar, A. Born again neural networks. *arXiv preprint arXiv:1805.04770*, 2018.
- Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., and Pedreschi, D. A survey of methods for explaining black box models. *ACM Computing Surveys (CSUR)*, 51(5):1–42, 2018.
- Günlük, O., Kalagnanam, J., Menickelly, M., and Scheinberg, K. Optimal decision trees for categorical data via integer programming. *arXiv preprint arXiv:1612.03225*, 2018.
- Hara, S. and Hayashi, K. Making tree ensembles interpretable: A bayesian model selection approach. *arXiv preprint arXiv:1606.09066*, 2016.
- Hernández-Lobato, D., Martínez-Muoz, G., and Suárez, A. Statistical instance-based pruning in ensembles of independent classifiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(2):364–369, 2009.
- Hinton, G., Vinyals, O., and Dean, J. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Hu, Q., Yu, D., Xie, Z., and Li, X. Eros: Ensemble rough subspaces. *Pattern recognition*, 40(12):3728–3739, 2007.
- Hu, X., Rudin, C., and Seltzer, M. Optimal sparse decision trees. In *Advances in Neural Information Processing Systems*, 2019.
- Kisamori, K. and Yamazaki, K. Model bridging: To interpretable simulation model from neural network. *arXiv preprint arXiv:1906.09391*, 2019.
- Margineantu, D. and Dietterich, T. Pruning adaptive boosting. In *Proceedings of the Fourteenth International Conference Machine Learning*, 1997.
- Martínez-Muñoz, G., Hernández-Lobato, D., and Suárez, A. An analysis of ensemble pruning techniques based on ordered aggregation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(2):245–259, 2008.
- Meinshausen, N. Node harvest. *The Annals of Applied Statistics*, pp. 2049–2072, 2010.
- Melis, D. A. and Jaakkola, T. Towards robust interpretability with self-explaining neural networks. In *Advances in Neural Information Processing Systems*, 2018.
- Mollas, I., Tsoumakas, G., and Bassiliades, N. Lionforests: Local interpretation of random forests through path selection. *arXiv preprint arXiv:1911.08780*, 2019.
- Murdoch, W., Singh, C., Kumbier, K., Abassi-Asl, R., and Yu, B. Interpretable machine learning: definitions, methods, and applications. *arXiv preprint arXiv:1901.04592v1*, 2019.

- Nijssen, S. and Fromont, E. Mining optimal decision trees from itemset lattices. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2007.
- Park, S. and Furnkranz, J. Efficient prediction algorithms for binary decomposition techniques. *Data Mining and Knowledge Discovery*, 24(1):40–77, 2012.
- Partalas, I., Tsoumakas, G., and Vlahavas, I. An ensemble uncertainty aware measure for directed hill climbing ensemble pruning. *Machine Learning*, 81(3):257–282, 2010.
- Prodromidis, A. L. and Stolfo, S. J. Cost complexity-based pruning of ensemble classifiers. *Knowledge and Information Systems*, 3(4):449–469, 2001.
- Prodromidis, A. L., Stolfo, S. J., and Chan, P. K. Effective and efficient pruning of meta-classifiers in a distributed data mining system. *Knowledge Discovery and Data Mining Journal*, 32, 1999.
- Rokach, L. Collective-agreement-based pruning of ensembles. *Computational Statistics & Data Analysis*, 53(4):1015–1026, 2009.
- Rokach, L. Decision forest: Twenty years of research. *Information Fusion*, 27:111–125, 2016.
- Rokach, L., Maimon, O., and Arbel, R. Selective voting—getting more for less in sensor fusion. *International Journal of Pattern Recognition and Artificial Intelligence*, 20(03):329–350, 2006.
- Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.
- Sirikulviriyaya, N. and Sinthupinyo, S. Integration of rules from a random forest. In *International Conference on Information and Electronics Engineering*, volume 6, 2011.
- Smith, J., Everhart, J., Dickson, W., Knowler, W., and Johannes, R. Using the ADAP learning algorithm to forecast the onset of diabetes mellitus. In *Proceedings of the Annual Symposium on Computer Applications in Medical Care*, pp. 261–265. IEEE Computer Society Press, 1988.
- Tamon, C. and Xiang, J. On the boosting pruning problem. In *Proceedings of the 11th European Conference on Machine Learning*, 2000.
- Tan, H. F., Hooker, G., and Wells, M. T. Tree space prototypes: Another look at making tree ensembles interpretable. *arXiv preprint arXiv:1611.07115*, 2016.
- Vandewiele, G., Lannoye, K., Janssens, O., Ongenaes, F., De Turck, F., and Van Hoecke, S. A genetic algorithm for interpretable model extraction from decision tree ensembles. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 2017.
- Verwer, S. and Zhang, Y. Learning optimal classification trees using a binary linear program formulation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019.
- Windeatt, T. and Ardeshtir, G. An empirical comparison of pruning methods for ensemble classifiers. In *International Symposium on Intelligent Data Analysis*, 2001.
- Zhang, H. and Wang, M. Search for the smallest random forest. *Statistics and its Interface*, 2(3):381, 2009.
- Zhang, Q., Nian Wu, Y., and Zhu, S.-C. Interpretable convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- Zhang, Y., Burer, S., and Street, W. N. Ensemble pruning via semi-definite programming. *Journal of Machine Learning Research*, 7(Jul):1315–1338, 2006.
- Zhou, Y. and Hooker, G. Interpreting models via single tree approximation. *arXiv preprint arXiv:1610.09036*, 2016.
- Zhou, Z., Wu, J., and Tang, W. Ensembling neural networks: many could be better than all. *Artificial Intelligence*, 137:239–263, 2002.
- Zhou, Z.-H. and Tang, W. Selective ensemble of decision trees. In *International Workshop on Rough Sets, Fuzzy Sets, Data Mining, and Granular-Soft Computing*, 2003.