

Generalized and Scalable Optimal Sparse Decision Trees

by

Chudi Zhong

Department of Statistical Science
Duke University

Date: _____

Approved:

Cynthia Rudin, Supervisor

Sayan Mukherjee, Chair

Jerome Reiter

A thesis submitted in partial fulfillment of the
requirements for the degree of Master of Science
in the Department of Statistical Science
in the Graduate School of
Duke University

2020

ABSTRACT

Generalized and Scalable Optimal Sparse Decision Trees

by

Chudi Zhong

Department of Statistical Science
Duke University

Date: _____

Approved:

Cynthia Rudin, Supervisor

Sayan Mukherjee, Chair

Jerome Reiter

An abstract of a thesis submitted in partial fulfillment of the
requirements for the degree of Master of Science
in the Department of Statistical Science
in the Graduate School of
Duke University

2020

Copyright © 2020 by Chudi Zhong
All rights reserved

Abstract

Decision tree optimization is notoriously difficult from a computational perspective but essential for the field of interpretable machine learning. Despite efforts over the past 40 years, only recently have optimization breakthroughs been made that have allowed practical algorithms to find *optimal* decision trees. These new techniques have the potential to trigger a paradigm shift where it is possible to construct sparse decision trees to efficiently optimize a variety of objective functions without relying on greedy splitting and pruning heuristics that often lead to suboptimal solutions. The contribution in this work is to provide a general framework for decision tree optimization that addresses the two significant open problems in the area: treatment of imbalanced data and fully optimizing over continuous variables. We present techniques that produce optimal decision trees over a variety of objectives including F-score, AUC, and partial area under the ROC convex hull. We also introduce a scalable algorithm that produces provably optimal results in the presence of continuous variables and speeds up decision tree construction by several orders of magnitude relative to the state-of-the-art.

Acknowledgements

This work is done in collaboration with Professor Cynthia Rudin (Duke University), Professor Margo Seltzer (The University of British Columbia), Jimmy Lin, and Diane Hu. I would like to express my sincerest gratitude to all of them, for their support. I was very fortunate to have Cynthia and Margo as my advisors. Their brilliance, diligence, rigor and curiosity have been inspiring me on my road of academic research. I also want to thank Xiyang Hu for providing ideas and implementation help at the early stage of this research project.

Contents

Abstract	iv
Acknowledgements	v
List of Tables	viii
List of Figures	ix
1 Introduction	1
2 Notations and Objectives	4
2.1 Hierarchical Bound for AUC Optimization	5
2.2 Summary of Bounds	7
3 Objectives and Their Bounds for Arbitrary Monotonic Losses	9
3.1 Hierarchical objective lower bound for arbitrary monotonic losses	9
3.2 Upper bound on the number of leaves	11
3.3 Lower Bounds on Incremental Progress	12
3.4 Permutation Bound	14
3.5 Equivalent Points Bound	14
3.6 Similar Support Bound	16
4 Objectives and Their Bounds for Rank Statistics	18
4.1 Proof of Theorem 2.1	24
4.2 Proof of Theorem 2.2	24
4.3 Incremental Progress Bound to Determine Splitting for Rank Statistics	25
4.4 Equivalent points bound for rank statistics	26
4.5 Partial area under the ROCCH	27

5	Optimizing F-score with Decision Trees	28
6	Data Preprocessing Using Bucketization Sacrifices Optimality	31
7	GOSDT's DPB Algorithm	33
7.1	Support Set Identification of Nodes	34
7.2	Connection to Similar Support Bound	35
7.3	Asynchronous Bound Updates	38
7.4	Fast Selective Vector Sums	39
7.5	The GOSDT Algorithm	40
8	Comparison Between Decision Tree Methods	47
9	Experiments	49
10	Conclusion and Future Work	57
	Bibliography	58

List of Tables

8.1 Comparison of decision tree approaches 48

9.1 Comparison of different data sets 55

List of Figures

4.1	Four cases of positions of l_{i_1} and l_{i_2}	19
4.2	Four different cases of changing the rank orders after introducing a child .	20
4.3	Example derivation of the hierarchical objective lower bound for rank statistics	25
6.1	Proof by construction	32
9.1	Training ROC and test ROC of FourClass dataset	50
9.2	Trees for different objectives	50
9.3	Training accuracy vs number of leaves	51
9.4	Training time vs the number of binary features	52
9.5	Best objective score	52
9.6	Eight-leaf decision trees generated by BinOCT and GOSDT on the Monk1 dataset	53
9.7	Seven-leaf decision trees generated by DL8.5 and GOSDT on the Monk2 dataset	54
9.8	16-leaf decision trees generated by BinOCT, DL8.5, and GOSDT on the tic-tac-toe dataset	56

Chapter 1

Introduction

Decision tree learning has served as a foundation for interpretable artificial intelligence and machine learning for over half a century [MS63, PM77, Loh14]. The major approach since the 1980's has been decision tree induction, where heuristic splitting and pruning procedures grow a tree from the top down and prune it back afterwards [Qui93, BFOS84]. The problem with these methods is that they tend to produce suboptimal trees with no way of knowing how suboptimal the solution is. This leaves a gap between the performance that a decision tree *might* obtain and the performance that one actually attains, with no way to check on (or remedy) the size of this gap—and sometimes, the gap can be large.

Full decision tree optimization is NP-hard, with no polynomial-time approximation [LR76], which has led to practical challenges in proving optimality or bounding the optimality gap in a reasonable amount of time, even for small datasets. It is possible to create assumptions that reduce hard decision tree optimization problems to cases where greedy algorithms suffice, such as independence between features [KS06], which do not hold in reality. If the data can be perfectly separated with zero error, SAT solvers can be used [NIPMS18] to find optimal decision trees rapidly; however, real data is generally not separable, leaving us with no choice other than to actually solve the problem.

Decision tree optimization is amenable to branch-and-bound methods, implemented either via generic mathematical programming solvers or by customized algorithms. Solvers have been used from the 1990's [Ben92, BB96] to the present [VZ19, BCMRM18, NS⁺20, BD17, RE18, MGKS18, VBSMVB19], but these generic solvers tend to be slow. A common way to speed them up is to make approximations in preprocessing to reduce the size of the search space. For instance, “bucketization” preprocessing is used in both generic solvers [VZ19] and customized algorithms [NS⁺20] to handle continuous variables. Continuous variables pose challenges to optimality; even one continuous variable increases the number of possible splits by the number of possible values of that variable in the entire database, and each additional split leads to an exponential increase in the

size of the optimization problem. Bucketization is tempting and seems innocuous, but we prove in Chapter 6 that it sacrifices optimality.

Dynamic programming methods have been used for various decision tree optimization problems since as far back as the early 1970's [Gar72, MM73]. Of the more recent attempts at this challenging problem, [GHRS03] use a dynamic programming method for finding an optimal subtree within a predefined larger decision tree grown using standard greedy induction. Their trees inherit suboptimality from the induction procedure used to create the larger tree. The DL8 algorithm [NF07] performs dynamic programming on the space of decision trees. However, without mechanisms to reduce the size of the space and to reduce computation, the method cannot be practical. A more practical extension is the DL8.5 method [NS⁺20], which uses a hierarchical upper bound theorem to reduce the size of the search space but also uses bucketization preprocessing, which renders non-optimality; without this preprocessing or other mechanisms to reduce computation, the method suffers in computational speed.

The CORELS algorithm [ALSA⁺17a, ALSA⁺18], which is an associative classification method rather than an optimal decision tree method, breaks apart from the previous literature in that it is a custom branch-and-bound method with custom bounding theorems, its own bit-vector library, specialized data structures, and an implementation that leverages computational reuse. CORELS is able to solve problems within only a minute that using any other prior approach, might have taken weeks, or even months or years. The CORELS framework was migrated to handle full binary-split decision trees in the Optimal Sparse Decision Tree (OSDT) algorithm [HRS19], which preserved CORELS' libraries, computational reuse, and many of its theorems to reduce the size of the search space. However, OSDT solves an exponentially harder problem than that of CORELS' rule list optimization, producing scalability challenges, as we might expect.

This work addresses two fundamental limitations in existing work: unsatisfying results for imbalanced data and an inability to fully optimize over continuous variables. Thus, the first contribution of this work is *to massively generalize sparse decision tree optimization to handle a wide variety of objective functions, including weighted accuracy, balanced accuracy, F-score, AUC and partial area under the ROC convex hull*. Both CORELS and OSDT were designed to maximize

accuracy, regularized by sparsity, and neither were designed to handle other objectives. CORELS has been generalized [AFG⁺19, CR18] to handle some constraints, but not to the wide variety of different objectives one might want to handle in practice. Generalization to some objectives is straightforward (e.g., weighted accuracy) but non-trivial in cases of optimizing rank statistics (e.g., AUC), which typically require quadratic computation in the number of observations in the dataset. However, for *sparse* decision trees, this time is much less than quadratic, owing to the fact that all observations within a leaf of a tree are tied in score, and because there are a sparse number of leaves in the tree. Taking advantage of this permits us to rapidly calculate rank statistics, and thus optimize over them. The second contribution is *to present a new representation of the dynamic programming search space that exposes a high degree of computational reuse when modelling continuous features*. The new search space representation provides a useful solution to a problem identified in the CORELS paper, which is how to use “similar support” bounds in practice. A similar support bound states that if two features in the dataset are similar to each other, then bounds obtained using the first feature for a split in a tree can be leveraged to obtain bounds for the same tree, were the second feature to replace the first feature. However, if the algorithm checks the similar support bound too frequently, the bound slows the algorithm down, despite reducing the search space. Our method uses hash trees that represent similar trees using shared subtrees, which naturally accelerates the evaluation of similar trees. The implementation, coupled with a new type of incremental similar support bound, is efficient enough to handle a few continuous features by creating dummy variables for all unique split points along a feature. This permits us to obtain smaller optimality gaps and certificates of optimality for mixed binary and continuous data, for additive loss functions, several orders of magnitude more quickly than any other method that currently exists.

Our algorithm is called Generalized and Scalable Optimal Sparse Decision Trees (GOSDT, pronounced “ghost”). A chart detailing a qualitative comparison of GOSDT to previous decision tree approaches is in Chapter 8.

Chapter 2

Notations and Objectives

We denote the training dataset as $\{(x_i, y_i)\}_{i=1}^N$, where $x_i \in \{0, 1\}^M$ are binary features and $y_i \in \{0, 1\}$. For each real-valued feature, we create a split point at the mean value between every ordered pair of unique values present in the training data. Following notation of [HRS19], we represent a tree as a set of leaves; this is important because it allows us *not* to store the splits of the tree, only the conditions leading to each leaf. A leaf set $d = (l_1, l_2, \dots, l_{H_d})$ contains H_d distinct leaves, where l_i is the classification rule of the leaf i , that is, the set of conditions along the branches that lead to the leaf, and \hat{y}_i^{leaf} is the label prediction for all data in leaf i . For a tree d , we define the objective function as a combination of the loss and a penalty on the number of leaves, with regularization parameter λ :

$$R(d, \mathbf{x}, \mathbf{y}) = \ell(d, \mathbf{x}, \mathbf{y}) + \lambda H_d. \quad (2.1)$$

Let us first consider *monotonic losses* $\ell(d, \mathbf{x}, \mathbf{y})$, which are monotonically increasing in the number of false positives (FP) and the number of false negatives (FN), and thus can be expressed alternatively as $\tilde{\ell}(FP, FN)$. We will specifically consider the following objectives in our implementation. (These are negated to become losses.)

- Accuracy = $1 - \frac{FP+FN}{N}$: fraction of correct predictions.
- Balanced accuracy = $1 - \frac{1}{2}(\frac{FN}{N^+} + \frac{FP}{N^-})$: the average of true positive rate and true negative rate. Let N^+ be the number of positive samples in the training set and N^- be the number of negatives.
- Weighted accuracy = $1 - \frac{FP+\omega FN}{\omega N^++N^-}$ for a predetermined threshold ω : the cost-sensitive accuracy that penalizes more on predicting positive samples as negative.
- F-score = $1 - \frac{FP+FN}{2N^++FP-FN}$: the harmonic mean of precision and recall.

Optimizing F-score directly is difficult even for linear modeling, because it is non-convex [NCLC12]. In optimizing F-score for decision trees, it is worse – a conundrum is possible where two leaves exist, the first leaf containing a higher proportion of positives than the other leaf, yet the first is classified as negative and the second classified as positive. We discuss how this can happen in Chapter 5 and how we address it, which is to force monotonicity by sweeping across leaves from highest to lowest predictions to calculate the F-score (see Chapter 5).

We consider two objectives that are rank statistics:

- Area under the ROC convex hull (AUC_{ch}): the fraction of correctly ranked positive/negative pairs.
- Partial area under the ROC convex hull ($pAUC_{ch}$) for predetermined threshold θ : the area under the leftmost part of the ROC curve.

Some of the bounds from OSDT [HRS19] have straightforward extensions to the objectives listed above, namely the **Upper Bound on Number of Leaves** and **Leaf Permutation Bound**. However, the remainder of OSDT’s bounds do not adapt. Our new bounds are: the **Hierarchical Objective Lower Bound**, **Incremental Progress Bound to Determine Splitting**, **Lower Bound on Incremental Progress**, **Equivalent Points Bound**, **Similar Support Bound** and an **Incremental Similar Support Bound**. To focus our exposition, the derivations and bounds for balanced classification loss, weighted classification loss, and F-score loss are in Chapter 3, and derivations for AUC loss and partial AUC loss are in Chapter 4, with the exception of the hierarchical lower bound for AUC_{ch} , which appears in Section 2.1 to demonstrate how these bounds work.

2.1 Hierarchical Bound for AUC Optimization

Let us discuss objectives that are rank statistics. If a classifier creates binary (as opposed to real-valued) predictions, its ROC curve consists of only four points (0,0), (FPR, TPR), and (1,1). The AUC of a labeled tree is same as the balanced accuracy. This is because $AUC = \frac{1}{2}(\frac{TP}{N^+} \times \frac{FP}{N^-} + (\frac{TP}{N^+} + 1) \times (1 - \frac{FP}{N^-}))$, and since $TP = N^+ - FN$, we have $AUC = \frac{1}{2}(\frac{N^+ - FN}{N^+} \times \frac{FP}{N^-} + \frac{2N^+ - FN}{N^+} \times$

$\frac{N^- - FP}{N^-} = 1 - \frac{1}{2}(\frac{FP}{N^-} + \frac{FN}{N^+})$. The more interesting case is when we have real-valued predictions for each leaf and use the ROC convex hull (ROCCH), defined shortly, as the objective.

Let n_i^+ be the number of positive samples in leaf i (n_i^- is the number of negatives) and let r_i be the fraction of positives in leaf i . Let us define the area under the ROC convex hull (ROCCH) [FFHO02] for a tree. For a tree d consisting of H_d distinct leaves, $d = (l_1, \dots, l_{H_d})$, we reorder leaves according to the fraction of positives, $r_1 \geq r_2 \geq \dots \geq r_{H_d}$. For any $i = 1, \dots, H_d$, define a labeling S_i for the leaves that labels the first i leaves as positive and remaining $H_d - i$ as negative. The collection of these labelings is $\Gamma = S_0, S_1, \dots, S_{H_d}$, where each S_i defines one of the $H_d + 1$ points on the ROCCH (see e.g., [FFHO02]). The associated misranking loss is then $1 - \text{AUC}_{\text{ch}}$:

$$\ell(d, \mathbf{x}, \mathbf{y}) = 1 - \frac{1}{N^+ N^-} \sum_{i=1}^{H_d} n_i^- \left[\left(\sum_{j=1}^{i-1} n_j^+ \right) + \frac{1}{2} n_i^+ \right]. \quad (2.2)$$

Now let us derive a lower bound on the loss for trees that are incomplete, meaning that some parts of the tree are not yet fully grown. For a partially-grown tree d , the leaf set can be rewritten as $d = (d_{\text{fix}}, r_{\text{fix}}, d_{\text{split}}, r_{\text{split}}, K, H_d)$, where d_{fix} is a set of K fixed leaves that we choose not to split further and d_{split} is the set of $H_d - K$ leaves that can be further split; this notation reflects how the algorithm works, where there are multiple copies of a tree, with some nodes allowed to be split and some that are not. r_{fix} and r_{split} are fractions of positives in the leaves. If we have a new fixed d'_{fix} , which is a superset of d_{fix} , then we say d'_{fix} is a child of d_{fix} . We define $\sigma(d)$ to be all such child trees:

$$\sigma(d) = \{(d'_{\text{fix}}, r'_{\text{fix}}, d'_{\text{split}}, r'_{\text{split}}, K', H'_d) : d_{\text{fix}} \subseteq d'_{\text{fix}}\}. \quad (2.3)$$

Denote N_{split}^+ and N_{split}^- as the number of positive and negative samples captured by d_{split} respectively. Through additional splits, in the best case, d_{split} can give rise to pure leaves, where positive ratios of generated leaves are either 1 or 0. Then the top-ranked leaf could contain up to N_{split}^+ positive samples (and 0 negative samples), and the lowest-ranked leaf could capture as few as 0 positive samples and up to N_{split}^- samples. Working now with just the leaves in d_{fix} , we reorder the leaves in d_{fix} by the positive ratios (r_{fix}), such that $\forall i \in \{1, \dots, K\}, r_1 \geq r_2 \geq \dots \geq r_K$. Combining these fixed leaves with the bounds for the split leaves, we can define a lower bound on the loss as follows.

Theorem 2.1. (*Lower bound for negative AUC convex hull*) For a tree $d = (d_{\text{fix}}, r_{\text{fix}}, d_{\text{split}}, r_{\text{split}}, K, H_d)$ using AUC_{ch} as the objective, a lower bound on the loss is $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) \leq R(d, \mathbf{x}, \mathbf{y})$, where:

$$b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) = 1 - \frac{1}{N^+ N^-} \left(\sum_{i=1}^K n_i^- \left[N_{\text{split}}^+ + \left(\sum_{j=1}^{i-1} n_j^+ \right) + \frac{1}{2} n_i^+ \right] + N^+ N_{\text{split}}^- \right) + \lambda H_d. \quad (2.4)$$

This leads directly to a hierarchical lower bound for the negative of the AUC convex hull.

Theorem 2.2. (*Hierarchical objective lower bound for negative AUC convex hull*) Let $d = (d_{\text{fix}}, r_{\text{fix}}, d_{\text{split}}, r_{\text{split}}, K, H_d)$ be a tree with fixed leaves d_{fix} and $d' = (d'_{\text{fix}}, r'_{\text{fix}}, d'_{\text{split}}, r'_{\text{split}}, K', H'_d) \in \sigma(d)$ be any child tree such that its fixed leaves d'_{fix} contain d_{fix} , and $H'_d > H_d$, then $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) \leq R(d', \mathbf{x}, \mathbf{y})$.

This type of bound acts as our weapon to reduce the size of the search space: if we compare the lower bound $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y})$ for partially constructed tree d to the best current objective R^c , and find that $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) > R^c$, then there is no need to consider d or any subtree of d , as it is provably non-optimal. The hierarchical lower bound dramatically reduces the size of the search space. However, we have also an arsenal of tighter bounds at our disposal, as summarized in the next subsection.

We leave the description of partial AUC to Chapter 4. Given a parameter θ , the partial AUC of the ROCCH focuses only on the left area of the curve, consisting of the top ranked leaves, whose FPR is smaller than or equal to θ . This metric is used in information retrieval and healthcare.

2.2 Summary of Bounds

Chapter 3 contains our bounds, which are the driving force behind reducing the search space. Chapter 3 contains the Hierarchical Lower Bound (Theorem 3.1) for any objective (2.1) with an arbitrary monotonic loss function. This theorem is analogous to the Hierarchical Lower Bound for AUC optimization above. Chapter 3 also contains the Objective Bound with One-Step Lookahead (Theorem 3.2), Objective Bound for Sub-Trees (Theorem 3.3), Upper Bound on the Number of Leaves (Theorem 3.4), Parent-Specific Upper Bound on the Number of Leaves (Theorem 3.5), Incremental Progress Bound to Determine Splitting (Theorem 3.6), Lower Bound on Incremental Progress (Theorem 3.7), Leaf Permutation Bound (Theorem 3.8), Equivalent Points Bound (Theorem 3.9),

and General Similar Support Bound (Theorem 3.10). As discussed, no similar support bounds have been used successfully in prior work. In Section 7.2, we show how a new *Incremental Similar Support Bound* can be implemented within our specialized DBP (*dynamic programming with bounds*) algorithm to make decision tree optimization for additive loss functions (e.g., weighted classification error) much more efficient. Bounds for AUC_{ch} and $pAUC_{ch}$ are in Chapter 4, including the powerful Equivalent Points Bound (Theorem 4.3) for AUC_{ch} and $pAUC_{ch}$ and proofs for Theorem 2.1 and Theorem 2.2.

Note that we do not use convex proxies for rank statistics, as is typically done in supervised ranking (learning-to-rank). Optimizing a convex proxy for a rank statistic can yield results that are far from optimal. (see [RW18a]). Instead, we optimize the original (exact) rank statistics, regularized by sparsity, directly on the training set.

Chapter 3

Objectives and Their Bounds for Arbitrary

Monotonic Losses

Before deriving the bounds for arbitrary monotonic losses, we first introduce more notations. As we know, a leaf set $d = (l_1, l_2, \dots, l_{H_d})$ contains H_d distinct leaves, where l_i is the classification rule of the leaf i . If a leaf is labeled, then $y_i^{(\text{leaf})}$ is the label prediction for all data in leaf i . Therefore, a labeled partially-grown tree d with the leaf set $d = (l_1, l_2, \dots, l_{H_d})$ could be rewritten as $d = (d_{\text{fix}}, \delta_{\text{fix}}, d_{\text{split}}, \delta_{\text{split}}, K, H_d)$, where $d_{\text{fix}} = (l_1, l_2, \dots, l_K)$ is a set of K fixed leaves that are not permitted to be further split, $\delta_{\text{fix}} = (y_1^{(\text{leaf})}, y_2^{(\text{leaf})}, \dots, y_K^{(\text{leaf})}) \in \{0, 1\}^K$ are the predicted labels for leaves d_{fix} , $d_{\text{split}} = (l_{K+1}, l_{K+2}, \dots, l_{H_d})$ is the set of $H_d - K$ leaves that can be further split, and $\delta_{\text{split}} = (y_{K+1}^{(\text{leaf})}, y_{K+2}^{(\text{leaf})}, \dots, y_{H_d}^{(\text{leaf})}) \in \{0, 1\}^{H_d - K}$ are the predicted labels for leaves d_{split} .

3.1 Hierarchical objective lower bound for arbitrary monotonic losses

Theorem 3.1. (*Hierarchical objective lower bound for arbitrary monotonic losses*) *Let loss functions $\ell(d, \mathbf{x}, \mathbf{y})$ be monotonically increasing in FP and FN. We now change notation of the loss to be only a function of these two quantities, written now as $\tilde{\ell}(FP, FN)$. Let $d = (d_{\text{fix}}, \delta_{\text{fix}}, d_{\text{split}}, \delta_{\text{split}}, K, H)$ be a labeled tree with fixed leaves d_{fix} , and let FP_{fix} and FN_{fix} be the false positives and false negatives of d_{fix} . Define the lower bound to the risk $R(d, \mathbf{x}, \mathbf{y})$ as follows (taking the lower bound of the split terms to be 0):*

$$R(d, \mathbf{x}, \mathbf{y}) \geq b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) = \ell(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) + \lambda H = \tilde{\ell}(FP_{\text{fix}}, FN_{\text{fix}}) + \lambda H.$$

Let $d' = (d'_{\text{fix}}, \delta'_{\text{fix}}, d'_{\text{split}}, \delta'_{\text{split}}, K', H') \in \sigma(d)$ be any child tree of d such that its fixed leaves d'_{fix} contain d_{fix} and $K' > K$ and $H' > H$. Then, $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) \leq R(d', \mathbf{x}, \mathbf{y})$.

The importance of this result is that the lower bound works for all (allowed) child trees of d . Thus, if d can be excluded via the lower bound, then all of its children can too.

Proof. Let FP_{fix} and FN_{fix} be the false positives and false negatives within leaves of d_{fix} , and let FP_{split} and FN_{split} be the false positives and false negatives within leaves of d_{split} . Similarly, denote $FP_{\text{fix}'}$ and $FN_{\text{fix}'}$ as the false positives and false negatives of d'_{fix} and let $FP_{\text{split}'}$ and $FN_{\text{split}'}$ be the false positives and false negatives of d'_{split} . Since the leaves are mutually exclusive, $FP_d = FP_{\text{fix}} + FP_{\text{split}}$ and $FN_d = FN_{\text{fix}} + FN_{\text{split}}$. Moreover, since $\tilde{\ell}(FP, FN)$ is monotonically increasing in FP and FN , we have:

$$R(d, \mathbf{x}, \mathbf{y}) = \tilde{\ell}(FP_d, FN_d) + \lambda H \geq \tilde{\ell}(FP_{\text{fix}}, FN_{\text{fix}}) + \lambda H = b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}). \quad (3.1)$$

Similarly, $R(d', \mathbf{x}, \mathbf{y}) \geq b(d'_{\text{fix}}, \mathbf{x}, \mathbf{y})$. Since $d_{\text{fix}} \subseteq d'_{\text{fix}'}$, $FP_{\text{fix}'} \geq FP_{\text{fix}}$ and $FN_{\text{fix}'} \geq FN_{\text{fix}}$, thus $\tilde{\ell}(FP_{\text{fix}'}, FN_{\text{fix}'}) \geq \tilde{\ell}(FP_{\text{fix}}, FN_{\text{fix}})$. Combined with $H' > H$, we have:

$$b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) = \tilde{\ell}(FP_{\text{fix}}, FN_{\text{fix}}) + \lambda H \leq \tilde{\ell}(FP_{\text{fix}'}, FN_{\text{fix}'}) + \lambda H' = b(d'_{\text{fix}}, \mathbf{x}, \mathbf{y}) \leq R(d', \mathbf{x}, \mathbf{y}). \quad (3.2)$$

□

Let us move to the next bound, which is the one-step lookahead. It relies on comparing the best current objective we have seen so far, denoted R^c , to the lower bound.

Theorem 3.2. (*Objective lower bound with one-step lookahead*) *Let d be a H -leaf tree with K leaves fixed and let R^c be the current best objective. If $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) + \lambda \geq R^c$, then for any child tree $d' \in \sigma(d)$, its fixed leaves d'_{fix} include d_{fix} and $H' > H$. It follows that $R(d', \mathbf{x}, \mathbf{y}) \geq R^c$.*

Proof. According to definition of the objective lower bound,

$$\begin{aligned} R(d', \mathbf{x}, \mathbf{y}) &\geq b(d'_{\text{fix}}, \mathbf{x}, \mathbf{y}) = \tilde{\ell}(FP_{\text{fix}'}, FN_{\text{fix}'}) + \lambda H' \\ &= \tilde{\ell}(FP_{\text{fix}'}, FN_{\text{fix}'}) + \lambda H + \lambda(H' - H) \\ &\geq b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) + \lambda \geq R^c, \end{aligned} \quad (3.3)$$

where on the last line we used that since H' and H are both integers, then $H' - H \geq 1$. □

According to this bound, even though we might have a tree d whose fixed leaves d_{fix} obeys lower bound $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) \leq R^c$, its child trees may still all be guaranteed to be suboptimal: if $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) + \lambda \geq R^c$, none of its child trees can ever be an optimal tree.

Theorem 3.3. (*Hierarchical objective lower bound for sub-trees and additive losses*) *Let loss functions $\ell(d, \mathbf{x}, \mathbf{y})$ be monotonically increasing in FP and FN, and let the loss of a tree d be the sum of the losses of the leaves. Let R^c be the current best objective. Let d be a tree such that the root node is split by a feature, where two sub-trees d_{left} and d_{right} are generated with H_{left} leaves for d_{left} and H_{right} leaves for d_{right} . The data captured by the left tree is $(\mathbf{x}_{\text{left}}, \mathbf{y}_{\text{left}})$ and the data captured by the right tree is $(\mathbf{x}_{\text{right}}, \mathbf{y}_{\text{right}})$. Let $b(d_{\text{left}}, \mathbf{x}_{\text{left}}, \mathbf{y}_{\text{left}})$ and $b(d_{\text{right}}, \mathbf{x}_{\text{right}}, \mathbf{y}_{\text{right}})$ be the objective lower bound of the left sub-tree and right sub-tree respectively such that $b(d_{\text{left}}, \mathbf{x}_{\text{left}}, \mathbf{y}_{\text{left}}) \leq \ell(d_{\text{left}}, \mathbf{x}_{\text{left}}, \mathbf{y}_{\text{left}}) + \lambda H_{\text{left}}$ and $b(d_{\text{right}}, \mathbf{x}_{\text{right}}, \mathbf{y}_{\text{right}}) \leq \ell(d_{\text{right}}, \mathbf{x}_{\text{right}}, \mathbf{y}_{\text{right}}) + \lambda H_{\text{right}}$. If $b(d_{\text{left}}, \mathbf{x}_{\text{left}}, \mathbf{y}_{\text{left}}) > R^c$ or $b(d_{\text{right}}, \mathbf{x}_{\text{right}}, \mathbf{y}_{\text{right}}) > R^c$ or $b(d_{\text{left}}, \mathbf{x}_{\text{left}}, \mathbf{y}_{\text{left}}) + b(d_{\text{right}}, \mathbf{x}_{\text{right}}, \mathbf{y}_{\text{right}}) > R^c$, then the tree d is not the optimal tree.*

Proof. $R(d, \mathbf{x}, \mathbf{y}) = \ell(d, \mathbf{x}, \mathbf{y}) + \lambda H = \ell(d_{\text{left}}, \mathbf{x}_{\text{left}}, \mathbf{y}_{\text{left}}) + \ell(d_{\text{right}}, \mathbf{x}_{\text{right}}, \mathbf{y}_{\text{right}}) + \lambda H_{\text{left}} + \lambda H_{\text{right}} \geq b(d_{\text{left}}, \mathbf{x}_{\text{left}}, \mathbf{y}_{\text{left}}) + b(d_{\text{right}}, \mathbf{x}_{\text{right}}, \mathbf{y}_{\text{right}})$. If $b(d_{\text{left}}, \mathbf{x}_{\text{left}}, \mathbf{y}_{\text{left}}) > R^c$ or $b(d_{\text{right}}, \mathbf{x}_{\text{right}}, \mathbf{y}_{\text{right}}) > R^c$ or $b(d_{\text{left}}, \mathbf{x}_{\text{left}}, \mathbf{y}_{\text{left}}) + b(d_{\text{right}}, \mathbf{x}_{\text{right}}, \mathbf{y}_{\text{right}}) > R^c$, then $R(d, \mathbf{x}, \mathbf{y}) > R^c$. Therefore, the tree d is not the optimal tree. \square

3.2 Upper bound on the number of leaves

Theorem 3.4. (*Upper bound on the number of leaves*) *For a dataset with M features, consider a state space of all trees. Let H be the number of leaves of tree d and let R^c be the current best objective. For all optimal tree $d^* \in \operatorname{argmin}_d R(d, \mathbf{x}, \mathbf{y})$,*

$$H^* \leq \min([\lceil R^c / \lambda \rceil, 2^M]) \quad (3.4)$$

where λ is the regularization parameter.

Proof. This bound adapts directly from OSDT [HRS19], where the proof can be found. \square

Theorem 3.5. (Parent-specific upper bound on the number of leaves) Let $d = (d_{\text{fix}}, \delta_{\text{fix}}, d_{\text{split}}, \delta_{\text{split}}, K, H)$ be a tree, $d' = (d'_{\text{fix}}, \delta'_{\text{fix}}, d'_{\text{split}}, \delta'_{\text{split}}, K', H') \in \sigma(d)$ be any child tree such that $d_{\text{fix}} \subseteq d'_{\text{fix}}$, and R^c be the current best objective. If d'_{fix} has lower bound $b(d'_{\text{fix}}, \mathbf{x}, \mathbf{y}) < R^c$, then

$$H' < \min \left(H + \left\lceil \frac{R^c - b(d_{\text{fix}}, \mathbf{x}, \mathbf{y})}{\lambda} \right\rceil, 2^M \right). \quad (3.5)$$

where λ is the regularization parameter.

Proof. This bound adapts directly from OSDT [HRS19], where the proof can be found. \square

3.3 Lower Bounds on Incremental Progress

In the implementation, Theorem 3.6 below is used to check if a leaf node within d_{split} is worth splitting. If the bound is satisfied and the leaf can be further split, then we generate new leaves and Theorem 3.7 is applied to check if this split yields new nodes or leaves that are good enough to consider in the future. Let us give an example to show how Theorem 3.6 is easier to compute than Theorem 3.7. If we are evaluating a potential split on leaf j , Theorem 3.6 requires FP_j and FN_j which are the false positives and false negatives for leaf j , but no extra information about the split we are going to make, whereas Theorem 3.7 requires that additional information. Let us work with balanced accuracy as the loss function: for Theorem 3.6 below, we would need to compute $\tau = \frac{1}{2} \left(\frac{FN_j}{N^+} + \frac{FP_j}{N^-} \right)$ but for Theorem 3.7 below we would need to calculate quantities for the new leaves we would form by splitting j into child leaves i and $i + 1$. Namely, we would need FN_i , FN_{i+1} , FP_i , and FP_{i+1} as well.

Theorem 3.6. (Incremental progress bound to determine splitting) Let $d^* = (d_{\text{fix}}, \delta_{\text{fix}}, d_{\text{split}}, \delta_{\text{split}}, K, H)$ be any optimal tree with objective R^* , i.e., $d^* \in \text{argmin}_d R(d, \mathbf{x}, \mathbf{y})$. Consider tree d' derived from d^* by deleting a pair of leaves l_i and l_{i+1} and adding their parent leaf l_j , $d' = (l_1, \dots, l_{i-1}, l_{i+2}, \dots, l_H, l_j)$. Let $\tau := \tilde{\ell}(FP_{d'}, FN_{d'}) - \tilde{\ell}(FP_{d'} - FP_{l_j}, FN_{d'} - FN_{l_j})$. Then, τ must be at least λ .

Proof. $\ell(d', \mathbf{x}, \mathbf{y}) = \tilde{\ell}(FP_{d'}, FN_{d'})$ and $\ell(d^*, \mathbf{x}, \mathbf{y}) = \tilde{\ell}(FP_{d'} + FP_{l_i} + FP_{l_{i+1}} - FP_{l_j}, FN_{d'} + FN_{l_i} + FN_{l_{i+1}} - FN_{l_j})$. The difference between $\ell(d^*, \mathbf{x}, \mathbf{y})$ and $\ell(d', \mathbf{x}, \mathbf{y})$ is maximized when

l_i and l_{i+1} correctly classify all the captured data. Therefore, τ is the maximal difference between $\ell(d', \mathbf{x}, \mathbf{y})$ and $\ell(d^*, \mathbf{x}, \mathbf{y})$. Since $\ell(d', \mathbf{x}, \mathbf{y}) - \ell(d^*, \mathbf{x}, \mathbf{y}) \leq \tau$, we can get $\ell(d', \mathbf{x}, \mathbf{y}) + \lambda(H - 1) \leq \ell(d^*, \mathbf{x}, \mathbf{y}) + \lambda(H - 1) + \tau$, that is (and remember that d^* is of size H whereas d' is of size $H - 1$), $R(d', \mathbf{x}, \mathbf{y}) \leq R(d^*, \mathbf{x}, \mathbf{y}) - \lambda + \tau$. Since d^* is optimal with respect to R , $0 \leq R(d', \mathbf{x}, \mathbf{y}) - R(d^*, \mathbf{x}, \mathbf{y}) \leq -\lambda + \tau$, thus, $\tau \geq \lambda$. \square

Hence, for a tree d , if any of its internal node contributes less than λ in loss, even though $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) \leq R^*$, it cannot be the optimal tree and none of its child tree could be the optimal tree. Thus, after evaluating tree d , we can prune it.

Theorem 3.7. (*Lower bound on incremental progress*) Let $d^* = (d_{\text{fix}}, \delta_{\text{fix}}, d_{\text{split}}, \delta_{\text{split}}, K, H)$ be any optimal tree with objective R^* , i.e., $d^* \in \operatorname{argmin}_d R(d, \mathbf{x}, \mathbf{y})$. Let d^* have leaves $d_{\text{fix}} = (l_1, \dots, l_H)$ and $\delta_{\text{fix}} = (y_1^{(\text{leaf})}, y_2^{(\text{leaf})}, \dots, y_H^{(\text{leaf})})$. For each leaf pair l_i, l_{i+1} with corresponding labels $y_i^{(\text{leaf})}, y_{i+1}^{(\text{leaf})}$ in d^* and their parent node (the leaf in the parent tree) l_j and its label $y_j^{(\text{leaf})}$, define a_i as the incremental objective of splitting l_j to get l_i, l_{i+1} : $a_i := \ell(d', \mathbf{x}, \mathbf{y}) - \ell(d^*, \mathbf{x}, \mathbf{y})$. In this case, λ provides a lower bound s.t. $a_i \geq \lambda$.

Proof. Let $d' = (d'_{\text{fix}}, \delta'_{\text{fix}}, d'_{\text{split}}, \delta'_{\text{split}}, K', H')$ be the tree derived from d^* by deleting a pair of leaves l_i and l_{i+1} , and adding their parent leaf l_j . Then,

$$\begin{aligned} R(d', \mathbf{x}, \mathbf{y}) &= \ell(d', \mathbf{x}, \mathbf{y}) + \lambda(H - 1) = a_i + \ell(d^*, \mathbf{x}, \mathbf{y}) + \lambda(H - 1) \\ &= a_i + R(d^*, \mathbf{x}, \mathbf{y}) - \lambda. \end{aligned} \tag{3.6}$$

Since $0 \leq R(d', \mathbf{x}, \mathbf{y}) - R(d^*, \mathbf{x}, \mathbf{y})$, then $a_i \geq \lambda$. \square

In the implementation, we apply both Theorem 3.6 and Theorem 3.7. If Theorem 3.6 is not satisfied, even though $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) \leq R^*$, it cannot be an optimal tree and none of its child trees could be an optimal tree. In this case, d can be pruned, as we showed before. However, if Theorem 3.6 is satisfied, we check Theorem 3.7. If Theorem 3.7 is not satisfied, then we would need to further split at least one of the two child leaves—either of the new leaves i or $i + 1$ —in order to obtain a potentially optimal tree.

3.4 Permutation Bound

Theorem 3.8. (*Leaf Permutation bound*) Let π be any permutation of $\{1, \dots, H\}$. Let $d = (d_{\text{fix}}, d_{\text{split}}, K, H)$ and $D = (D_{\text{fix}}, D_{\text{split}}, K, H)$ be trees with leaves (l_1, \dots, l_H) and $(l_{\pi(1)}, \dots, l_{\pi(H)})$ respectively, i.e., the leaves in D correspond to a permutation of the leaves in d . Then the objective lower bounds of d and D are the same and their child trees correspond to permutations of each other.

Proof. This bound adapts directly from OSDT [HRS19], where the proof can be found. \square

Therefore, if two trees have the same leaves, up to a permutation, according to Theorem 3.8, one of them can be pruned. This bound is capable of reducing the search space by all future symmetries of trees we have already seen.

3.5 Equivalent Points Bound

As we know, for a tree $d = (d_{\text{fix}}, \delta_{\text{fix}}, d_{\text{split}}, \delta_{\text{split}}, K, H)$, the objective of this tree (and that of its children) is minimized when there are no errors in the split leaves: $FP_{\text{split}} = 0$ and $FN_{\text{split}} = 0$. In that case, the risk is equal to $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y})$. However, if multiple observations captured by a leaf in d_{split} have the same features but different labels, then no tree, including those that extend d_{split} , can correctly classify all of these observations, that is FP_{split} and FN_{split} cannot be zero. In this case, we can apply the equivalent points bound to give a tighter lower bound on the objective.

Let Ω be a set of leaves. *Capture* is an indicator function that equals 1 if x_i falls into one of the leaves in Ω , and 0 otherwise, in which case we say that $\text{cap}(x_i, \Omega) = 1$. We define a set of samples to be equivalent if they have exactly the same feature values. Let e_u be a set of equivalent points and let q_u be the minority class label that minimizes the loss among points in e_u . Note that a dataset consists of multiple sets of equivalent points. Let $\{e_u\}_{u=1}^U$ enumerate these sets.

Theorem 3.9. (*Equivalent points bound*) Let $d = (d_{\text{fix}}, \delta_{\text{fix}}, d_{\text{split}}, \delta_{\text{split}}, K, H)$ be a tree such that $l_k \in d_{\text{fix}}$ for $k \in \{1, \dots, K\}$ and $l_k \in d_{\text{split}}$ for $k \in \{K + 1, \dots, H\}$. For any tree $d' \in \sigma(d)$,

$$\ell(d', \mathbf{x}, \mathbf{y}) \geq \tilde{\ell}(FP_{\text{fix}} + FP_e, FN_{\text{fix}} + FN_e), \text{ where} \quad (3.7)$$

$$\begin{aligned}
FP_e &= \sum_{i=1}^N \sum_{u=1}^U \sum_{k=K+1}^H \text{cap}(x_i, l_k) \wedge \mathbb{1}[y_i = 0] \wedge \mathbb{1}[x_i \in e_u] \mathbb{1}[y_i = q_u] \\
FN_e &= \sum_{i=1}^N \sum_{u=1}^U \sum_{k=K+1}^H \text{cap}(x_i, l_k) \wedge \mathbb{1}[y_i = 1] \wedge \mathbb{1}[x_i \in e_u] \mathbb{1}[y_i = q_u]
\end{aligned} \tag{3.8}$$

Proof. Since $d' \in \sigma(d)$, $d' = (l'_1, \dots, l'_K, l'_{K+1}, \dots, l'_{K'}, \dots, l'_{H'})$ where we have both $l'_k \in d'_{\text{fix}}$ for $k \in \{1, \dots, K'\}$, which are the fixed leaves and also $l'_k \in d'_{\text{split}}$ for $k \in \{K' + 1, \dots, H'\}$. Note that for $k \in \{1, \dots, K\}$, $l'_k = l_k$.

Let $\Delta = d'_{\text{fix}} \setminus d_{\text{fix}}$ which are the leaves in d'_{fix} that are not in d_{fix} . Then $\ell(d', \mathbf{x}, \mathbf{y}) = \tilde{\ell}(FP_{d'}, FN_{d'}) = \tilde{\ell}(FP_{\text{fix}} + FP_{\Delta} + FP_{\text{split}'}, FN_{\text{fix}} + FN_{\Delta} + FN_{\text{split}'})$, where FP_{Δ} and FN_{Δ} are false positives and false negatives in d'_{fix} but not d_{fix} and $FP_{\text{split}'}$ and $FN_{\text{split}'}$ are false positives and false negatives in d'_{split} . For tree d' , its leaves in Δ are those indexed from K to K' . Thus, the sum over leaves of d' from K to H' includes leaves from Δ and leaves from d'_{split} .

$$\begin{aligned}
FP_{\Delta} + FP_{\text{split}'} &= \sum_{i=1}^N \sum_{u=1}^U \sum_{k=K+1}^{H'} \text{cap}(x_i, l'_k) \wedge \mathbb{1}[y_i \neq \hat{y}_k^{(\text{leaf})}] \wedge \mathbb{1}[y_i = 0] \wedge \mathbb{1}[x_i \in e_u] \\
&\geq \sum_{i=1}^N \sum_{u=1}^U \sum_{k=K+1}^{H'} \text{cap}(x_i, l'_k) \wedge \mathbb{1}[y_i = 0] \wedge \mathbb{1}[x_i \in e_u] \mathbb{1}[y_i = q_u] \\
FN_{\Delta} + FN_{\text{split}'} &= \sum_{i=1}^N \sum_{u=1}^U \sum_{k=K+1}^{H'} \text{cap}(x_i, l'_k) \wedge \mathbb{1}[y_i \neq \hat{y}_k^{(\text{leaf})}] \wedge \mathbb{1}[y_i = 1] \wedge \mathbb{1}[x_i \in e_u] \\
&\geq \sum_{i=1}^N \sum_{u=1}^U \sum_{k=K+1}^{H'} \text{cap}(x_i, l'_k) \wedge \mathbb{1}[y_i = 1] \wedge \mathbb{1}[x_i \in e_u] \mathbb{1}[y_i = q_u]
\end{aligned} \tag{3.9}$$

For $i \in \{1, \dots, N\}$, the samples in d_{split} are the same ones captured by either Δ or d'_{split} , that is $\sum_{k=K+1}^H \text{cap}(x_i, l_k) = \sum_{k=K+1}^{H'} \text{cap}(x_i, l'_k)$. Then

$$FP_{\Delta} + FP_{\text{split}'} \geq \sum_{i=1}^N \sum_{u=1}^U \sum_{k=K+1}^H \text{cap}(x_i, l_k) \wedge \mathbb{1}[y_i = 0] \wedge \mathbb{1}[x_i \in e_u] \mathbb{1}[y_i = q_u] = FP_e. \tag{3.10}$$

Similarly, $FN_{\Delta} + FN_{\text{split}'} \geq FN_e$. Therefore,

$$\begin{aligned}
\ell(d', \mathbf{x}, \mathbf{y}) &= \tilde{\ell}(FP_{\text{fix}} + FP_{\Delta} + FP_{\text{split}'}, FN_{\text{fix}} + FN_{\Delta} + FN_{\text{split}'}) \\
&\geq \tilde{\ell}(FP_{\text{fix}} + FP_e, FN_{\text{fix}} + FN_e)
\end{aligned} \tag{3.11}$$

□

3.6 Similar Support Bound

Given two trees that are exactly the same except for one internal node split by different features f_1 and f_2 , we can use the similar support bound for pruning.

Theorem 3.10. (*Similar support bound*) Define $d = (d_{\text{fix}}, \delta_{\text{fix}}, d_{\text{split}}, \delta_{\text{split}}, K, H)$ and $D = (D_{\text{fix}}, \Delta_{\text{fix}}, D_{\text{split}}, \Delta_{\text{split}}, K, H)$ to be two trees that are exactly the same except for one internal node split by different features. Let f_1 and f_2 be the features used to split that node in d and D respectively. Let t_1, t_2 be the left and right sub-trees under the node f_1 in d and let T_1, T_2 be the left and right sub-trees under the node f_2 in D . Let ω be the observations captured by only one of t_1 or T_1 , i.e.,

$$\omega := \{i : [\text{cap}(x_i, t_1) \wedge \neg \text{cap}(x_i, T_1) + \neg \text{cap}(x_i, t_1) \wedge \text{cap}(x_i, T_1)]\}. \quad (3.12)$$

Let $FP_{-\omega}$ and $FN_{-\omega}$ be the false positives and false negatives of samples except ω . The difference between the two trees' objectives is bounded as follows:

$$|R(d, \mathbf{x}, \mathbf{y}) - R(D, \mathbf{x}, \mathbf{y})| \leq \gamma, \text{ where} \quad (3.13)$$

$$\gamma := \max_{a \in \{0, \dots, |\omega|\}} [\tilde{\ell}(FP_{-\omega} + a, FN_{-\omega} + |\omega| - a)] - \tilde{\ell}(FP_{-\omega}, FN_{-\omega}). \quad (3.14)$$

Then we have

$$\left| \min_{d^+ \in \sigma(d)} R(d^+, \mathbf{x}, \mathbf{y}) - \min_{D^+ \in \sigma(D)} R(D^+, \mathbf{x}, \mathbf{y}) \right| \leq \gamma. \quad (3.15)$$

Proof. The difference between the objectives of d and D is largest when one of them correctly classifies all the data in ω but the other misclassifies all of them. If d classifies all the data corresponding to ω correctly while D misclassifies them,

$$R(d, \mathbf{x}, \mathbf{y}) - R(D, \mathbf{x}, \mathbf{y}) \geq \tilde{\ell}(FP_{-\omega}, FN_{-\omega}) - \max_{a \in \{0, \dots, |\omega|\}} [\tilde{\ell}(FP_{-\omega} + a, FN_{-\omega} + |\omega| - a)] = -\gamma. \quad (3.16)$$

We can get $R(d, \mathbf{x}, \mathbf{y}) - R(D, \mathbf{x}, \mathbf{y}) \leq \gamma$ in the same way. Therefore, $\gamma \geq R(d, \mathbf{x}, \mathbf{y}) - R(D, \mathbf{x}, \mathbf{y}) \geq -\gamma$.

Let d^* be the best child tree of d , i.e., $R(d^*, \mathbf{x}, \mathbf{y}) = \min_{d^+ \in \sigma(d)} R(d^+, \mathbf{x}, \mathbf{y})$. Let $D' \in \sigma(D)$ be its counterpart which is exactly the same except for one internal node split by a different feature.

Then, using Equation 3.16,

$$\min_{d^+ \in \sigma(d)} R(d^+, \mathbf{x}, \mathbf{y}) = R(d^*, \mathbf{x}, \mathbf{y}) \geq R(D', \mathbf{x}, \mathbf{y}) - \gamma \geq \min_{D^+ \in \sigma(D)} R(D^+, \mathbf{x}, \mathbf{y}) - \gamma. \quad (3.17)$$

Similarly, again using Equation (3.16), $\min_{D^+ \in \sigma(D)} R(D^+, \mathbf{x}, \mathbf{y}) + \gamma \geq \min_{d^+ \in \sigma(d)} R(d^+, \mathbf{x}, \mathbf{y})$. \square

Chapter 4

Objectives and Their Bounds for Rank Statistics

In this chapter, we provide proofs for Theorem 2.1 and Theorem 2.2, and adapt the Incremental Progress Bound to Determine Splitting and the Equivalent Points Bound for the objective AUC_{ch} . The Upper Bound on the Number of Leaves, Parent-Specific Upper Bound on the Number of Leaves, Lower Bound of Incremental Progress, and Permutation Bound are the same as the bounds in Chapter 3. We omit these duplicated proofs here. At the end of this chapter, we define the objective pAUC_{ch} and how we implement the derived bounds for this objective. As a reminder, we use notation $d = (d_{\text{fix}}, r_{\text{fix}}, d_{\text{split}}, r_{\text{split}}, K, H_d)$ to represent tree d .

Lemma 4.1. *Let $d = (d_{\text{fix}}, r_{\text{fix}}, d_{\text{split}}, r_{\text{split}}, K, H_d)$ be a tree. The AUC convex hull does not decrease when an impure leaf is split.*

Proof. Let l_i be the impure leaf that we intend to split, where $i \in \{1, \dots, H_d\}$. Let n_i^+ be the positive samples in l_i and n_i^- negative samples. Suppose l_i is ranked in position “pos.” If the leaf is split once, it will generate two leaves l_{i_1} and l_{i_2} such that $r_{i_1} \geq r_i \geq r_{i_2}$ without loss of generality. Let d' be the tree that consists of the leaf set $(l_1, \dots, l_{i-1}, l_{i+1}, \dots, l_{H_d}, l_{i_1}, l_{i_2})$. If $r_{i_1} = r_i = r_{i_2}$, then the rank order of leaves (according to the r_i 's) will not change, so AUC_{ch} will be unchanged after the split. Otherwise (if the rank order of leaves changes when introducing a child) we can reorder the $H + 1$ leaves, leading to the following four cases. For the new leaf set $(l_1, \dots, l_{i-1}, l_{i+1}, \dots, l_{H_d}, l_{i_1}, l_{i_2})$, either:

1. The rank of l_{i_1} is smaller than pos and the rank of l_{i_2} equals pos + 1 (requires $r_{i_1} > r_i$ and $r_i \geq r_{i_2}$);
2. The rank of l_{i_1} is smaller than pos and the rank of l_{i_2} is larger than pos + 1 (requires $r_{i_1} > r_i$ and $r_i > r_{i_2}$);
3. The rank of l_{i_1} is equal to pos and the rank of l_{i_2} is equal to pos + 1 (requires $r_{i_1} \geq r_i$ and $r_i \geq r_{i_2}$);

4. The rank of l_{i_1} is pos and the rank of l_{i_2} is larger than pos + 1 (requires $r_{i_1} \geq r_i$ and $r_i > r_{i_2}$).

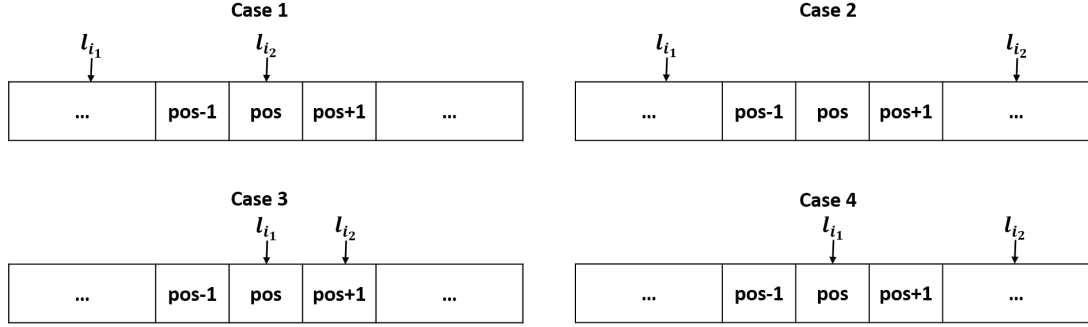


Figure 4.1: Four cases of positions of l_{i_1} and l_{i_2}

Figure 4.1 shows four cases of the positions of l_{i_1} and l_{i_2} .

Let us go through these cases in more detail.

For the new leaf set after splitting l_i , namely $(l_1, \dots, l_{i-1}, l_{i+1}, \dots, l_{H_d}, l_{i_1}, l_{i_2})$, we have that:

1. l_{i_1} has rank smaller than pos (which requires $r_{i_1} > r_i$) and l_{i_2} has rank pos + 1 (which requires $r_i \geq r_{i_2}$).

Let $A = \{l_{a_1}, l_{a_2}, \dots, l_{a_U}\}$ be a collection of leaves ranked before l_{i_1} and let $B = \{l_{b_1}, l_{b_2}, \dots, l_{b_V}\}$ be a collection of leaves ranked after l_{i_1} but before pos + 1. In this case, recalling Equation (2.2), a change in the AUC_{ch} after splitting on leaf i is due only to a subset of leaves, namely $l_{i_1}, l_{b_1}, \dots, l_{b_V}, l_{i_2}$. Then we can compute the change in the AUC_{ch} as follows:

$$\begin{aligned} \Delta_{\text{AUC}_{\text{ch}}} &= \frac{1}{N^+ N^-} \left(\frac{n_{i_1}^- n_{i_1}^+}{2} + \left(\sum_{v=1}^V n_{b_v}^- \right) n_{i_1}^+ + n_{i_2}^- \left[n_{i_1}^+ + \left(\sum_{v=1}^V n_{b_v}^+ \right) + \frac{n_{i_2}^+}{2} \right] \right. \\ &\quad \left. - n_i^- \left[\left(\sum_{v=1}^V n_{b_v}^+ \right) + \frac{n_i^+}{2} \right] \right) \end{aligned} \quad (4.1)$$

To derive the expression for $\Delta_{\text{AUC}_{\text{ch}}}$, we first sum shaded areas of rectangles and triangles under the ROC curves' convex hull for both tree d and its child tree d' , and then calculate the difference between the two shaded areas, as indicated in Figure 4.2 (a-c). This figure shows where each of the terms arises within the $\Delta_{\text{AUC}_{\text{ch}}}$: terms $n_{b_v}^- n_{i_1}^+$ and $n_{i_2}^- [n_{i_1}^+ + (\sum_{v=1}^V n_{b_v}^+)]$ come from the area of rectangles colored in dark pink in Figure 4.2 (b). Terms $\frac{n_{i_1}^- n_{i_1}^+}{2}$ and $\frac{n_{i_2}^- n_{i_2}^+}{2}$ handle the top triangles colored in light pink. Term $n_i^- (\sum_{v=1}^V n_{b_v}^+)$ represents the rectangles colored

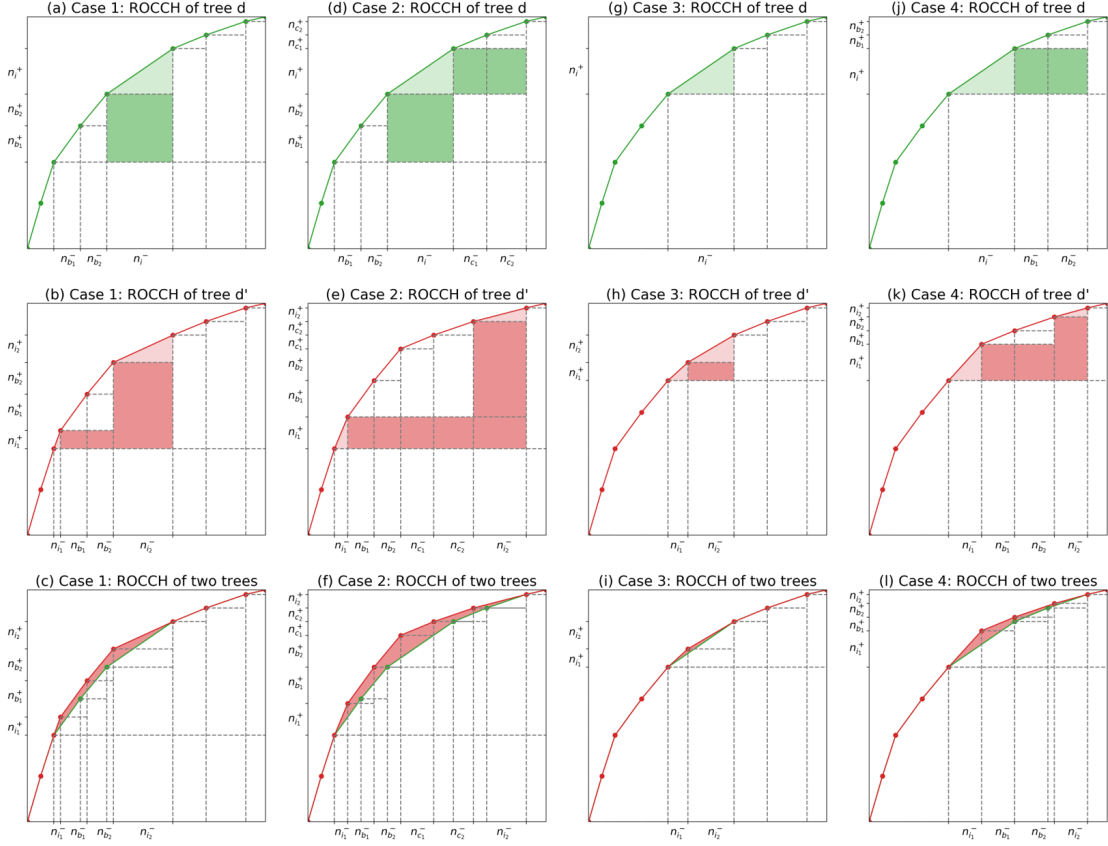


Figure 4.2: Four different cases of changing the rank orders after introducing a child. Each column represents a case. Subplots in the first row show the ROCCH of tree d and subplots in the second row indicate the ROCCH of tree d' corresponding to the different positions of two child leaves. Subplots in the third row present the change of AUC_{ch} after introducing the child leaves.

in dark green in Figure 4.2 (a) and term $\frac{n_i^- n_i^+}{2}$ deals with the triangles colored in light green. Subtracting green shaded areas from red shaded areas, we get ΔAUC_{ch} , which is represented by the (remaining) pink area in Figure 4.2 (c).

Simplifying Equation (4.1), we get

$$\Delta AUC_{ch} = \frac{1}{N^+ N^-} \left(n_{i_1}^+ \left(\sum_{v=1}^V n_{b_v}^- \right) + n_{i_2}^- \left(\sum_{v=1}^V n_{b_v}^+ \right) - n_i^- \left(\sum_{v=1}^V n_{b_v}^+ \right) \right) + \frac{n_i^- n_{i_1}^+ + 2n_{i_2}^- n_{i_1}^+ + n_{i_2}^- n_{i_2}^+ - n_i^- n_i^+}{2} \quad (4.2)$$

Recall that $n_i^- = n_{i_1}^- + n_{i_2}^-$ and $n_i^+ = n_{i_1}^+ + n_{i_2}^+$. Then, simplifying,

$$\Delta_{\text{AUC}_{\text{ch}}} = \frac{1}{N^+N^-} \left(n_{i_1}^+ \left(\sum_{v=1}^V n_{b_v}^- \right) - n_{i_1}^- \left(\sum_{v=1}^V n_{b_v}^+ \right) + \frac{n_{i_1}^+ n_{i_2}^- - n_{i_1}^- n_{i_2}^+}{2} \right). \quad (4.3)$$

Since $r_{i_1} > r_{b_1} > r_{b_2} > \dots > r_{b_V}$, $\forall v \in \{1, 2, \dots, V\}$, $\frac{n_{i_1}^+}{n_{i_1}^+ + n_{i_1}^-} > \frac{n_{b_v}^+}{n_{b_v}^+ + n_{b_v}^-}$. Then we can get $n_{i_1}^+ n_{b_v}^- > n_{b_v}^+ n_{i_1}^-$. Hence $n_{i_1}^+ (\sum_{v=1}^V n_{b_v}^-) - n_{i_1}^- (\sum_{v=1}^V n_{b_v}^+) > 0$. Similarly, because $r_{i_1} > r_{i_2}$, $\frac{n_{i_1}^+}{n_{i_1}^+ + n_{i_1}^-} > \frac{n_{i_2}^+}{n_{i_2}^+ + n_{i_2}^-}$. Then $n_{i_1}^+ n_{i_2}^- > n_{i_1}^- n_{i_2}^+$. Therefore, $\Delta_{\text{AUC}_{\text{ch}}} > 0$.

2. l_{i_1} has a ranking smaller than pos (which requires $r_{i_1} > r_i$) and l_{i_2} has a ranking larger than pos + 1 (which requires $r_i > r_{i_2}$).

Let $A = \{l_{a_1}, \dots, l_{a_U}\}$ be a collection of leaves that ranked before l_{i_1} and $B = \{l_{b_1}, \dots, l_{b_V}\}$ be a collection of leaves that ranked after l_{i_1} but before pos + 1, and $C = \{l_{c_1}, \dots, l_{c_W}\}$ be a collection of leaves that ranked after pos + 1 but before the rank of l_{i_2} . In this case, the change is caused by $l_{i_1}, l_{b_1}, \dots, l_{b_V}, l_{c_1}, \dots, l_{c_W}, l_{i_2}$. Then we can compute the change in the AUC_{ch} as follows:

$$\begin{aligned} \Delta_{\text{AUC}_{\text{ch}}} = & \frac{1}{N^+N^-} \left(\frac{n_{i_1}^- n_{i_1}^+}{2} + \left(\sum_{v=1}^V n_{b_v}^- \right) n_{i_1}^+ + \left(\sum_{w=1}^W n_{c_w}^- \right) n_{i_1}^+ + n_{i_2}^- \left[n_{i_1}^+ + \left(\sum_{v=1}^V n_{b_v}^+ \right) \right. \right. \\ & \left. \left. + \left(\sum_{w=1}^W n_{c_w}^+ \right) + \frac{n_{i_2}^+}{2} \right] - n_{i_1}^- \left[\left(\sum_{v=1}^V n_{b_v}^+ \right) + \frac{n_{i_1}^+}{2} \right] - \left(\sum_{w=1}^W n_{c_w}^- \right) n_{i_1}^+ \right) \end{aligned} \quad (4.4)$$

Similar to the derivation proposed in case 1, we first sum shaded areas of rectangles and triangles under the ROC curves' convex hull for both tree d and its child tree d' and then calculate the difference between two shaded areas as indicated in Figure 4.2 (d-f). These three subfigures show where each of the terms arises within the $\Delta_{\text{AUC}_{\text{ch}}}$: terms $n_{b_v}^- n_{i_1}^+$, $n_{c_w}^- n_{i_1}^+$, and $n_{i_2}^- [n_{i_1}^+ + (\sum_{v=1}^V n_{b_v}^+) + (\sum_{w=1}^W n_{c_w}^+)]$ come from the area of rectangles colored in dark pink in Figure 4.2 (e). Terms $\frac{n_{i_1}^- n_{i_1}^+}{2}$ and $\frac{n_{i_2}^- n_{i_2}^+}{2}$ handle the top triangles colored in light pink. Terms $n_{i_1}^- n_{b_v}^+$ and $n_{i_1}^+ n_{c_w}^-$ represent the rectangles colored in dark green in Figure 4.2 (d) and term $\frac{n_{i_1}^- n_{i_1}^+}{2}$ deals with the triangles colored in light green. Subtracting green shaded areas from red shaded areas, we can get $\Delta_{\text{AUC}_{\text{ch}}}$, which is represented by the light red area in Figure 4.2 (f).

Recall that $n_i^- = n_{i_1}^- + n_{i_2}^-$ and $n_i^+ = n_{i_1}^+ + n_{i_2}^+$. Then, simplifying Equation (4.4), we get

$$\begin{aligned}
\Delta_{\text{AUC}_{\text{ch}}} &= \frac{1}{N^+ N^-} \left(\left(\sum_{v=1}^V n_{b_v}^- \right) n_{i_1}^+ + \left(\sum_{w=1}^W n_{c_w}^- \right) n_{i_1}^+ + \left(\sum_{v=1}^V n_{b_v}^+ \right) n_{i_2}^- + \left(\sum_{w=1}^W n_{c_w}^+ \right) n_{i_2}^- \right. \\
&\quad \left. - \left(\sum_{v=1}^V n_{b_v}^+ \right) n_{i_1}^- - \left(\sum_{w=1}^W n_{c_w}^- \right) n_{i_1}^+ + \frac{n_{i_1}^- n_{i_1}^+ + 2n_{i_2}^- n_{i_1}^+ + n_{i_2}^- n_{i_2}^+ - n_{i_1}^- n_{i_1}^+}{2} \right) \\
&= \frac{1}{N^+ N^-} \left(\left(\sum_{v=1}^V n_{b_v}^- \right) n_{i_1}^+ - \left(\sum_{v=1}^V n_{b_v}^+ \right) n_{i_1}^- + \left(\sum_{w=1}^W n_{c_w}^+ \right) n_{i_2}^- - \left(\sum_{w=1}^W n_{c_w}^- \right) n_{i_2}^+ \right. \\
&\quad \left. + \frac{n_{i_2}^- n_{i_1}^+ - n_{i_1}^- n_{i_2}^+}{2} \right)
\end{aligned} \tag{4.5}$$

Since $r_{i_1} > r_{b_1} > \dots > r_{b_V}$, $\forall v \in \{1, \dots, V\}$, $\frac{n_{i_1}^+}{n_{i_1}^+ + n_{i_1}^-} > \frac{n_{b_v}^+}{n_{b_v}^+ + n_{b_v}^-}$. Then we get $n_{i_1}^+ n_{b_v}^- > n_{b_v}^+ n_{i_1}^-$. Thus, $(\sum_{v=1}^V n_{b_v}^-) n_{i_1}^+ - (\sum_{v=1}^V n_{b_v}^+) n_{i_1}^- > 0$. Similarly, since $r_{c_1} > \dots > r_{c_W} > r_{i_2}$, $\forall w \in \{1, \dots, W\}$, $n_{c_w}^+ n_{i_2}^- > n_{i_2}^+ n_{c_w}^-$. Thus, $(\sum_{w=1}^W n_{c_w}^+) n_{i_2}^- - (\sum_{w=1}^W n_{c_w}^-) n_{i_2}^+ > 0$. Moreover, because $r_{i_1} > r_{i_2}$, $n_{i_2}^- n_{i_1}^+ > n_{i_1}^- n_{i_2}^+$. Hence, $\Delta_{\text{AUC}_{\text{ch}}} > 0$

3. l_{i_1} has a ranking same as pos (which requires $r_{i_1} \leq r_i$) and l_{i_2} has a ranking pos + 1 (which requires $r_i \leq r_{i_2}$).

In this case, the change of AUC_{ch} is caused by l_{i_1} and l_{i_2} . Then we compute the change in the AUC_{ch} as follows:

$$\Delta_{\text{AUC}_{\text{ch}}} = \frac{1}{N^+ N^-} \left(\frac{n_{i_1}^- n_{i_1}^+}{2} + n_{i_2}^- n_{i_1}^+ + \frac{n_{i_2}^- n_{i_2}^+}{2} - \frac{n_{i_1}^- n_{i_1}^+}{2} \right) \tag{4.6}$$

We derive the expression in the similar way as case 1 and case 2. Term $n_{i_2}^- n_{i_1}^+$ comes from the area of rectangle colored in dark pink in Figure 4.2 (h) and terms $\frac{n_{i_1}^- n_{i_1}^+}{2}$ and $\frac{n_{i_2}^- n_{i_2}^+}{2}$ handle the top triangles colored in light pink. Term $\frac{n_{i_1}^- n_{i_1}^+}{2}$ deals with the top triangle colored in light green in Figure 4.2 (g). Subtracting green shaded areas from pink shaded areas, we get $\Delta_{\text{AUC}_{\text{ch}}}$, which is represented by the (remaining) pink area in Figure 4.2 (i).

Recall $n_i^- = n_{i_1}^- + n_{i_2}^-$ and $n_i^+ = n_{i_1}^+ + n_{i_2}^+$. Simplifying Equation (4.6), we get

$$\Delta_{\text{AUC}_{\text{ch}}} = \frac{1}{N^+ N^-} \left(\frac{n_{i_1}^+ n_{i_2}^- - n_{i_2}^+ n_{i_1}^-}{2} \right) \tag{4.7}$$

Since $r_{i_1} > r_{i_2}$, $n_{i_1}^+ n_{i_2}^- - n_{i_2}^+ n_{i_1}^-$. Therefore, $\Delta_{\text{AUC}_{\text{ch}}} > 0$.

4. l_{i_1} has a ranking same as pos (which requires $r_{i_1} \leq r_i$) and l_{i_2} has a ranking larger than pos + 1 (which requires $r_i > r_{i_2}$).

Let $A = \{l_{a_1}, \dots, l_{a_V}\}$ be a collection of leaves that ranked before l_{i_1} and $B = \{l_{b_1}, \dots, l_{b_V}\}$ be a collection of leaves that ranked after l_{i_1} but before l_{i_2} . In this case the change of AUC_{ch} is caused by $l_{i_1}, l_{b_1}, \dots, l_{b_V}, l_{i_2}$. Then we can compute the change as follows:

$$\begin{aligned} \Delta_{\text{AUC}_{\text{ch}}} = & \frac{1}{N^+N^-} \left(\frac{n_{i_1}^- n_{i_1}^+}{2} + \left(\sum_{v=1}^V n_{b_v}^- \right) n_{i_1}^+ + n_{i_2}^- \left[n_{i_1}^+ + \left(\sum_{v=1}^V n_{b_v}^+ \right) + \frac{n_{i_2}^+}{2} \right] \right. \\ & \left. - \frac{n_i^- n_i^+}{2} - \left(\sum_{v=1}^V n_{b_v}^- \right) n_i^+ \right) \end{aligned} \quad (4.8)$$

The Figure 4.2 (j-l) show where each of the terms arises within the $\Delta_{\text{AUC}_{\text{ch}}}$: terms $n_{b_v}^- n_{i_1}^+$ and $n_{i_2}^- [n_{i_1}^+ + (\sum_{v=1}^V n_{b_v}^+)]$ come from the area of rectangles colored in dark pink in Figure 4.2 (k) and terms $\frac{n_{i_1}^- n_{i_1}^+}{2}$ and $\frac{n_{i_2}^- n_{i_2}^+}{2}$ handle triangles colored in light pink. Term $n_{b_v}^- n_i^+$ represents the rectangle colored in dark green in Figure 4.2 (j) and term $\frac{n_i^- n_i^+}{2}$ deals with the triangle colored in light green. Subtracting green shaded areas from pink shaded areas, we get $\Delta_{\text{AUC}_{\text{ch}}}$, which is represented by the (remaining) pink area in Figure 4.2 (l).

Recall $n_i^- = n_{i_1}^- + n_{i_2}^-$ and $n_i^+ = n_{i_1}^+ + n_{i_2}^+$. Simplifying Equation (4.8), we get

$$\begin{aligned} \Delta_{\text{AUC}_{\text{ch}}} = & \frac{1}{N^+N^-} \left(\left(\sum_{v=1}^V n_{b_v}^- \right) n_{i_1}^+ + \left(\sum_{v=1}^V n_{b_v}^+ \right) n_{i_2}^- - \left(\sum_{v=1}^V n_{b_v}^- \right) n_i^+ \right. \\ & \left. + \frac{n_{i_1}^- n_{i_1}^+ + 2n_{i_2}^- n_{i_1}^+ + n_{i_2}^- n_{i_2}^+ - n_i^- n_i^+}{2} \right) \\ = & \frac{1}{N^+N^-} \left(\left(\sum_{v=1}^V n_{b_v}^+ \right) n_{i_2}^- - \left(\sum_{v=1}^V n_{b_v}^- \right) n_{i_2}^+ + \frac{n_{i_2}^- n_{i_1}^+ - n_{i_1}^- n_{i_2}^+}{2} \right) \end{aligned} \quad (4.9)$$

Since $r_{b_1} > \dots > r_{b_V} > r_{i_2}$, $\forall v \in \{1, \dots, V\}$, $n_{b_v}^+ n_{i_2}^- > n_{i_2}^+ n_{b_v}^-$. Thus, $(\sum_{v=1}^V n_{b_v}^+) n_{i_2}^- > (\sum_{v=1}^V n_{b_v}^-) n_{i_2}^+$. Since $r_{i_1} > r_{i_2}$, $n_{i_1}^+ n_{i_2}^- - n_{i_2}^+ n_{i_1}^-$. Therefore, $\Delta_{\text{AUC}_{\text{ch}}} > 0$.

Therefore, once an impure leaf is split, the AUC_{ch} doesn't decrease. If the split is leading to the change of the rank order of leaves, then AUC_{ch} increases. \square

4.1 Proof of Theorem 2.1

Proof. Let tree $d = (d_{\text{fix}}, r_{\text{fix}}, d_{\text{split}}, r_{\text{split}}, K, H_d)$ and leaves of tree d are mutually exclusive. According to Lemma 4.1, for leaves that can be further split, we can do splits to increase the AUC_{ch} . In the best possible hypothetical case, all new generated leaves are pure (contain only positive or negative samples). In this case, AUC_{ch} is maximized for d_{split} . In this case, we will show that $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) = 1 - \frac{1}{N^+N^-} \left(\sum_{i=1}^K n_i^- \left[N_{\text{split}}^+ + \left(\sum_{j=1}^{i-1} n_j^+ \right) + \frac{1}{2}n_i^+ \right] + N^+N_{\text{split}}^- \right) + \lambda H_d$ as defined by Equation (2.4). Then $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) \leq R(d, \mathbf{x}, \mathbf{y})$.

To derive the expression for $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y})$, we sum areas of rectangles and triangles under the ROC curve's convex hull. Figure 4.3 shows where each of the terms arises within this sum: the first term in the sum, which is $n_i^- N_{\text{split}}^+$, comes from the area of the lower rectangle of the ROC curve's convex hull, colored in green. This rectangle arises from the block of N_{split}^+ positives at the top of the ranked list (within the split leaves, whose hypothetical predictions are 1). The $n_i^- \left(\sum_{j=1}^{i-1} n_j^+ \right)$ term handles the areas of the growing rectangles, colored in blue in Figure 4.3. The areas of the triangles comprising the top of the ROCCH account for the third term $\frac{1}{2}n_i^- n_i^+$. The final term $N^+N_{\text{split}}^-$ comes from the rectangle on the right, colored red, stemming from the split observations within a hypothetical purely negative leaf.

□

4.2 Proof of Theorem 2.2

Proof. According to Theorem 2.1, $R(d', \mathbf{x}, \mathbf{y}) \geq b(d'_{\text{fix}}, \mathbf{x}, \mathbf{y})$ and $R(d, \mathbf{x}, \mathbf{y}) \geq b(d_{\text{fix}}, \mathbf{x}, \mathbf{y})$. Since $d_{\text{fix}} \subseteq d'_{\text{fix}}$, leaves in d_{split} but not in d'_{split} can be further split to generate pure leaves for tree d but fixed for tree d' . Based on Lemma 4.1 and $H_d \leq H_{d'}$, $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) \leq b(d'_{\text{fix}}, \mathbf{x}, \mathbf{y})$. Therefore, $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) \leq R(d', \mathbf{x}, \mathbf{y})$.

□

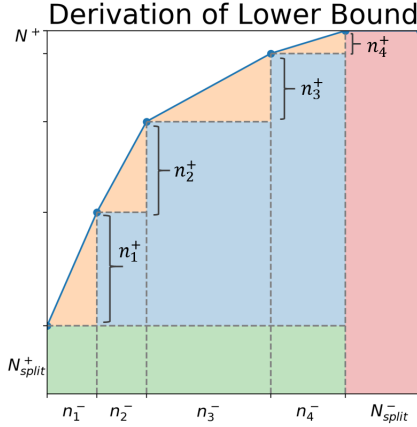


Figure 4.3: Example derivation of the hierarchical objective lower bound for rank statistics. N^+ at top left is the number of positive samples in the dataset. N_{split}^+ at bottom left is the number of positive samples (N_{split}^- is the number of negative samples) captured by leaves that can be further split.

4.3 Incremental Progress Bound to Determine Splitting for Rank Statistics

Theorem 4.2. (*Incremental progress bound to determine splitting for rank statistics*) Let $d^* = (d_{\text{fix}}, r_{\text{fix}}, d_{\text{split}}, r_{\text{split}}, K, H_d)$ be any optimal tree with objective R^* . Consider tree d' derived from d^* by deleting a pair of leaves l_i and l_{i+1} and adding their parent leaf l_j , $d' = (l_1, \dots, l_{i-1}, l_{i+2}, \dots, l_{H_d}, l_j)$. Let n_j^+ be the number of positive samples (n_j^- be the number of negative samples) in leaf j . Calculate $\ell(d', \mathbf{x}, \mathbf{y})$ as in Equation (2.2). Define d'_{-j} as the tree d' after dropping leaf j , adding two hypothetical pure leaves (i.e., one has all positives and the other has all negatives), and reordering the remaining $H_d - 2$ leaves based on the fraction of positives. Then, we can calculate the loss of the tree d'_{-j} as

$$\ell(d'_{-j}, \mathbf{x}, \mathbf{y}) = 1 - \frac{1}{N^+ N^-} \left(\sum_{i=1}^{H_d-2} n_i^- \left[n_j^+ + \left(\sum_{v=1}^{i-1} n_v^+ \right) + \frac{1}{2} n_i^+ \right] + N^+ n_j^- \right). \quad (4.10)$$

Let $\tau := \ell(d', \mathbf{x}, \mathbf{y}) - \ell(d'_{-j}, \mathbf{x}, \mathbf{y})$. Then τ must be at least λ .

Proof. By the way $\ell(d'_{-j}, \mathbf{x}, \mathbf{y})$ is defined (using the same counting argument we used in the proof of Theorem 2.1), it is a lower bound for $\ell(d^*, \mathbf{x}, \mathbf{y})$. These two quantities are equal when the

split leaves are all pure. So, we have $\ell(d'_{-j}, \mathbf{x}, \mathbf{y}) \leq \ell(d^*, \mathbf{x}, \mathbf{y})$. Since $\ell(d', \mathbf{x}, \mathbf{y}) - \ell(d^*, \mathbf{x}, \mathbf{y}) \leq \ell(d', \mathbf{x}, \mathbf{y}) - \ell(d'_{-j}, \mathbf{x}, \mathbf{y}) = \tau$, we can get $\ell(d', \mathbf{x}, \mathbf{y}) + \lambda(H_d - 1) \leq \ell(d^*, \mathbf{x}, \mathbf{y}) + \lambda(H_d - 1) + \tau$, that is (and remember that d^* is of size H_d whereas d' is of size $H_d - 1$), $R(d', \mathbf{x}, \mathbf{y}) \leq R(d^*, \mathbf{x}, \mathbf{y}) - \lambda + \tau$. Since d^* is optimal with respect to R , then $0 \leq R(d', \mathbf{x}, \mathbf{y}) - R(d^*, \mathbf{x}, \mathbf{y}) \leq -\lambda + \tau$, thus $\tau \geq \lambda$. \square

Similar to the Incremental Progress Bound to Determine Splitting for arbitrary monotonic losses, for a tree d , if any of its internal node contributes less than λ in loss, it is not the optimal tree. Thus, after evaluating tree d , we can prune it.

4.4 Equivalent points bound for rank statistics

Similar to the equivalent points bound for arbitrary monotonic losses, for a tree $d = (d_{\text{fix}}, r_{\text{fix}}, d_{\text{split}}, r_{\text{split}}, K, H_d)$, the objective of this tree and its children is minimized when leaves that can be further split to generate pure leaves. In the case when it is possible to split the data into pure leaves, the risk could be equal to $b(d_{\text{fix}}, \mathbf{x}, \mathbf{y})$. However, if multiple observations captured by a leaf in d_{split} have the same features but different labels, then no tree, including those that extend d_{split} , can correctly classify all of these observations; that is, leaves in d_{split} can not generate pure leaves. In this case, we can leverage these equivalent points to give a tighter lower bound for the objective. We use the same notation for capture, set of equivalent points as in Chapter 3, and minority class label is simply the label that fewer samples have.

Let $d = (d_{\text{fix}}, r_{\text{fix}}, d_{\text{split}}, r_{\text{split}}, K, H_d)$ be a tree. Data in leaves from d_{split} can be separated into U equivalence classes. For $u = 1, \dots, U$, let $N_u = \sum_{i=1}^N \text{cap}(x_i, d_{\text{split}}) \wedge \mathbb{1}[x_i \in e_u]$, that is, the number of samples captured by d_{split} belonging to equivalence set u , and let $\delta_u = \sum_{i=1}^N \text{cap}(x_i, d_{\text{split}}) \wedge \mathbb{1}[x_i \in e_u] \mathbb{1}[y_i = q_u]$ be the number of minority-labeled samples captured by d_{split} in equivalence point set u . Then we define \tilde{r}_u as the classification rule we would make on each equivalence class separately (if we were permitted):

$$\tilde{r}_u = \begin{cases} \frac{\delta_u}{N_u} & \text{if } \delta_u \geq 0 \text{ and } q_u = 1 \\ \frac{N_u - \delta_u}{N_u} & \text{if } \delta_u \geq 0 \text{ and } q_u = 0. \end{cases}$$

Let us combine this with d_{fix} to get a bound. We order the combination of leaves in d_{fix} and equivalence classes $u \in \{1, \dots, U\}$ by the fraction of positives in each, $\text{Sort}(r_1, \dots, r_i, \dots, r_K, \tilde{r}_1, \dots, \tilde{r}_u, \dots, \tilde{r}_U)$ from highest to lowest. Let us re-index these sorted $K + U$ elements by index \tilde{i} . Denote $n_{\tilde{i}}^+$ as the number of positive samples in either the leaf or equivalence class corresponding to \tilde{i} . (Define $n_{\tilde{i}}^-$ to be the number of negative samples analogously). We define our new, tighter, lower bound as:

$$b(d_{\text{equiv}}, \mathbf{x}, \mathbf{y}) = 1 - \frac{1}{N^+ N^-} \sum_{\tilde{i}=1}^{K+U} n_{\tilde{i}}^- \left[\left(\sum_{j=1}^{\tilde{i}-1} n_j^+ \right) + \frac{1}{2} n_{\tilde{i}}^+ \right] + \lambda H_d. \quad (4.11)$$

Theorem 4.3. (*Equivalent points bound for rank statistics*) For a tree $d = (d_{\text{fix}}, r_{\text{fix}}, d_{\text{split}}, r_{\text{split}}, K, H_d)$, Let tree $d' = (d'_{\text{fix}}, r'_{\text{fix}}, d'_{\text{split}}, r'_{\text{split}}, K', H_{d'}) \in \sigma(d)$ be any child tree such that its fixed leaves d'_{fix} contain d_{fix} and $H_{d'} \geq H_d$. Then $b(d_{\text{equiv}}, \mathbf{x}, \mathbf{y}) \leq R(d', \mathbf{x}, \mathbf{y})$, where $b(d_{\text{equiv}}, \mathbf{x}, \mathbf{y})$ is defined in Equation (4.11).

Proof. The proof is similar to the proof of Theorem 2.1 and Theorem 2.2. □

4.5 Partial area under the ROCCH

We discuss the partial area under the ROC convex hull in this section. The ROCCH for a decision tree is defined in Section 2.1, where leaves are rank-ordered by the fraction of positives in the leaves. Given a threshold θ , the partial area under the ROCCH (pAUC_{ch}) looks at only the leftmost part of the ROCCH, that is focusing on the top ranked leaves. This measure is important for applications such as information retrieval and maintenance [see, e.g., RW18a]. In our implementation, all bounds derived for the objective AUC_{ch} can be adapted directly for pAUC_{ch} , where all terms are calculated only for false positive rates smaller or equal to θ .

In our code, we implement all of the rank statistics bounds, with one exception for the partial AUCCH – the equivalent points bound. We do not implement the equivalent points bound for partial AUCCH since the pAUC_{ch} statistic is heavily impacted by the leaves with high fraction of positives, which means that the leaves being repeatedly calculated for the objective tend not to be impure, and thus the equivalence points bound is less effective.

Chapter 5

Optimizing F-score with Decision Trees

For a labeled tree $d = (d_{\text{fix}}, \delta_{\text{fix}}, d_{\text{split}}, \delta_{\text{split}}, K, H)$, the F-score loss is defined as

$$\ell(d, \mathbf{x}, \mathbf{y}) = \frac{FP + FN}{2N^+ + FP - FN}. \quad (5.1)$$

For objectives like accuracy, balanced accuracy and weighted accuracy, the loss of a tree is the sum of loss in each leaf. For F-score loss, however, FP and FN appear in both numerator and denominator, thus the loss no longer can be calculated using a sum over the leaves.

Lemma 5.1. *The label of a single leaf depends on the labels of other leaves when optimizing F-score loss.*

Proof. Let l_1, \dots, l_{H_d} be the leaves of tree d . Suppose $H_d - 1$ leaves are labeled. Let FP_{H_d-1} and FN_{H_d-1} be the number of false positives and false negatives of these $H_d - 1$ leaves, respectively. Let $A = FP_{H_d-1} + FN_{H_d-1}$ and $B = 2N^+ + FP_{H_d-1} - FN_{H_d-1}$, and by these definitions, we will always have $A \leq B$. Let $n_{H_d}^+$ be the number of positive samples in leaf H_d and $n_{H_d}^-$ be the number of negative samples. The leaf's predicted label can be either positive or negative. The loss of the tree depends on this predicted label as follows:

$$\text{If } \hat{y}_{H_d}^{(\text{leaf})} = 1, \text{ there can be only false positives, thus } \ell(d, \mathbf{x}, \mathbf{y}) = \frac{A + n_{H_d}^-}{B + n_{H_d}^-}. \quad (5.2)$$

$$\text{If } \hat{y}_{H_d}^{(\text{leaf})} = 0, \text{ there can be only false negatives, thus } \ell(d, \mathbf{x}, \mathbf{y}) = \frac{A + n_{H_d}^+}{B - n_{H_d}^+}. \quad (5.3)$$

Calculating loss (5.3) minus loss (5.2):

$$\frac{A + n_{H_d}^+}{B - n_{H_d}^+} - \frac{A + n_{H_d}^-}{B + n_{H_d}^-} = \frac{(A + n_{H_d}^+) \times (B + n_{H_d}^-) - (A + n_{H_d}^-) \times (B - n_{H_d}^+)}{(B - n_{H_d}^+)(B + n_{H_d}^-)}. \quad (5.4)$$

Denote Δ as the numerator of (5.4), that is

$$(A + n_{H_d}^+) \times (B + n_{H_d}^-) - (A + n_{H_d}^-) \times (B - n_{H_d}^+).$$

Then we can get

$$\Delta = AB + An_{H_d}^- + Bn_{H_d}^+ + n_{H_d}^+ n_{H_d}^- - AB + An_{H_d}^+ - Bn_{H_d}^- + n_{H_d}^- n_{H_d}^+ \quad (5.5)$$

$$= 2n_{H_d}^+ n_{H_d}^- + An_{H_d}^+ + Bn_{H_d}^+ + An_{H_d}^- - Bn_{H_d}^-. \quad (5.6)$$

The value of Δ depends on A , B , $n_{H_d}^+$ and $n_{H_d}^-$. Hence, in order to minimize the loss, the predicted label of leaf H_d is 0 if $\Delta \leq 0$ and 1 otherwise. Therefore, the predicted label of a single leaf depends on A and B , which depend on the labels of the other samples, as well as the positive and negative samples captured by that leaf. \square

Theorem 5.2. (*Optimizing F-score Poses a Unusual Challenge*) Let l_1, \dots, l_{H_d} be the leaves of tree d and let N^+ be the number of positive samples in the dataset. Let Γ_1 and Γ_2 be two predicted labelings for the first $H_d - 1$ leaves. Leaf H_d has a fixed predicted label. Suppose the loss for the F-score (Equation (5.1)) of the first $H_d - 1$ leaves based on labeling method Γ_1 is smaller than the loss based on labeling method Γ_2 (where in both cases, leaf H_d has the same predicted label). It is not guaranteed that the F_1 loss of the tree d based on the first labeling Γ_1 is always smaller than the loss based on the second labeling Γ_2 .

Proof. Let $FP_{H_d-1}^{(1)}$ and $FN_{H_d-1}^{(1)}$ be the number of false positives and number of false negatives for the first $H_d - 1$ leaves from the labeling method Γ_1 and similarly define $FP_{H_d-1}^{(2)}$ and $FN_{H_d-1}^{(2)}$ for labeling method Γ_2 . Denote $A_1 = FP_{H_d-1}^{(1)} + FN_{H_d-1}^{(1)}$ and $B_1 = 2N^+ + FP_{H_d-1}^{(1)} - FN_{H_d-1}^{(1)}$. Similarly, denote $A_2 = FP_{H_d-1}^{(2)} + FN_{H_d-1}^{(2)}$ and $B_2 = 2N^+ + FP_{H_d-1}^{(2)} - FN_{H_d-1}^{(2)}$. As we know from the assumptions of the theorem, $\frac{A_1}{B_1} \leq \frac{A_2}{B_2}$.

Denote FP_{H_d} and FN_{H_d} be the number of false positives and number of false negatives of the last leaf l_{H_d} .

Let $\ell^{(1)}(d, \mathbf{x}, \mathbf{y})$ and $\ell^{(2)}(d, \mathbf{x}, \mathbf{y})$ be the loss of the tree d based on two different predicted labelings of the leaves.

Suppose the predicted label of leaf l_{H_d} is 1. (An analogous result holds when the predicted label of leaf l_{H_d} is 0.) Then $\ell^{(1)}(d, \mathbf{x}, \mathbf{y}) = \frac{A_1 + FP_{H_d}}{B_1 + FP_{H_d}}$ and $\ell^{(2)}(d, \mathbf{x}, \mathbf{y}) = \frac{A_2 + FP_{H_d}}{B_2 + FP_{H_d}}$. Let Δ be the

numerator of $\ell^{(2)}(d, \mathbf{x}, \mathbf{y}) - \ell^{(1)}(d, \mathbf{x}, \mathbf{y})$.

$$\begin{aligned}\Delta &= (A_2 + FP_{H_d})(B_1 + FP_{H_d}) - (A_1 + FP_{H_d})(B_2 + FP_{H_d}) \\ &= A_2B_1 - A_1B_2 + (A_2 - B_2 + B_1 - A_1)FP_{H_d}.\end{aligned}\tag{5.7}$$

Since $\frac{A_1}{B_1} \leq \frac{A_2}{B_2}$, $A_2B_1 \geq A_1B_2$, that is, the first two terms together are nonnegative. Meanwhile, $A_1 \leq B_1$ and $A_2 \leq B_2$. Thus, Δ could be negative or positive. Therefore, even though the labeling method Γ_1 leads to smaller loss for the first $H_d - 1$ leaves and the label of the last leaf depends on the label of previous $H_d - 1$ leaves, it is not guaranteed that the loss of the tree is smaller than that based on Γ_2 . It is easy to construct examples of A_1, A_2, B_1, B_2 , and FP_{H_d} where the result is either positive or negative, as desired. \square

Lemma 5.1 and Theorem 5.2 indicate that optimizing F-score loss is much harder than other arbitrary monotonic losses such as balanced accuracy and weighted accuracy. Thus, we simplify the labeling step by incorporating a parameter ω at each leaf node s.t. l_i is labeled as 0 if $\omega n_i^+ \leq n_i^-$ and 1 otherwise $\forall i \in \{1, \dots, H_d\}$.

Chapter 6

Data Preprocessing Using Bucketization

Sacrifices Optimality

As discussed, a preprocessing step common to DL8.5 [NS⁺20] and BinOct [VZ19] reduces the search space, but as we will prove, also sacrifices accuracy; we refer to this preprocessing as *bucketization*.

Definition: The *bucketization* preprocessing step proceeds as follows. Order the observations according to any feature j . For any two neighboring positive observations, no split can be considered between them. For any two neighboring negative observations, no split can be considered between them. All other splits are permitted. While bucketization may appear innocuous, we prove it sacrifices optimality.

Theorem 6.1. *The maximum training accuracy for a decision tree on a dataset preprocessed with bucketization can be lower (worse) than the maximum accuracy for the same dataset without bucketization.*

The proof is by construction. In Figure 6.1 we present a data set such that optimal training with bucketization cannot produce an optimal value of the objective. In particular, the optimal accuracy without bucketization is 93.5%, whereas the accuracy of training with bucketization is 92.2%. These numbers were obtained using BinOCT, DL8.5, and GOSDT. Remember, the algorithms provide a proof of optimality; these accuracy values are optimal, and the same values were found by all three algorithms.

Our dataset is two-dimensional. We expect the sacrifice to become worse with higher dimensions.

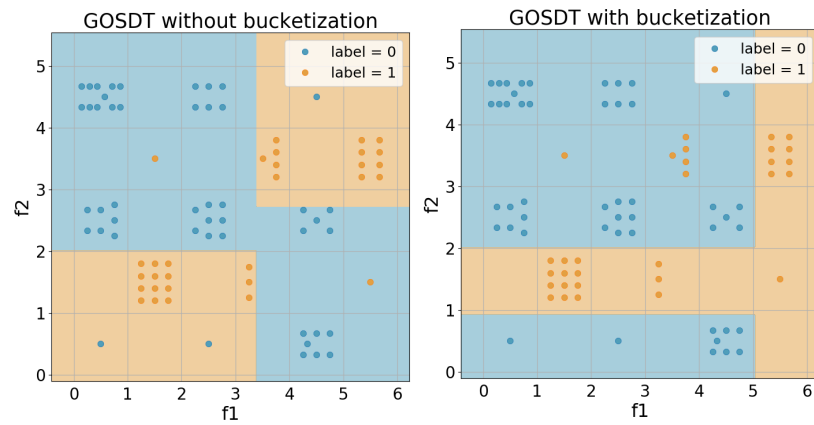


Figure 6.1: Proof by construction. Bucketization leads to suboptimality.

Chapter 7

GOSDT's DPB Algorithm

GOSDT uses an algorithm we call *dynamic programming with bounds* (DPB), which provides a dramatic run time improvement for additive loss functions. Like other dynamic programming approaches, we decompose a problem into smaller child problems that can be solved either recursively through a function call or in parallel by delegating work to a separate thread. For decision tree optimization, these subproblems are the possible left and right branches of a decision node.

GOSDT maintains two primary data structures: a *priority queue* to schedule problems to solve and a *dependency graph* to store problems and their dependency relationships. We define a dependency relationship $dep(p_\pi, p_c)$ between problems p_π and p_c if and only if the solution of p_π depends on the solution of p_c . Each p_c is further specified as p_l^j or p_r^j indicating that it is the left or right branch produced by splitting on feature j .

Sections 7.1, 7.2, and 7.3 highlight two key differences between GOSDT and DL8.5 (since DL 8.5 also uses a form of dynamic programming): (1) DL8.5 uniquely identifies a problem p by the *Boolean assertion* that is a conjunctive clause of all splitting conditions in its ancestry, while GOSDT represents a problem p by the *samples that satisfy the Boolean assertion*. This is elaborated in Section 7.1. (2) DL8.5 uses blocking recursive invocations to execute problems, while GOSDT uses a *priority queue* to schedule problems for later. This is elaborated in Section 7.3. Section 7.4 presents additional performance optimizations. Section 7.5 presents Algorithm 1: a high-level summary of GOSDT. Note that GOSDT operates on weighted, additive, non-negative loss functions. For ease of notation, we use classification error as the loss in our exposition.

7.1 Support Set Identification of Nodes

GOSDT leverages the Equivalent Points Bound in Theorem 3.9 to save space and computational time. To use this bound, we find the unique values of $\{x_1, \dots, x_N\}$, denoted by $\{z_1, \dots, z_U\}$, so that each x_i equals one of the z_u 's. We store fractions z_u^+ and z_u^- , which are the fraction of positive and negative samples corresponding to each z_u .

$$\begin{aligned} Z &= \{z_u : z_u \in \text{unique}(X), 1 \leq u \leq U \leq N\} \\ z_u^- &= \frac{1}{N} \sum_{i=1}^N \mathbb{1}[y_i = 0 \wedge x_i = z_u] \\ z_u^+ &= \frac{1}{N} \sum_{i=1}^N \mathbb{1}[y_i = 1 \wedge x_i = z_u]. \end{aligned}$$

Define a as a Boolean assertion that is a conjunctive clause of conditions on the features of z_u (e.g., a is true if and only if the first feature of z_u is 1 and the second feature of z_u is 0). We define *support set* s_a as the set of z_u that satisfy assertion a :

$$s_a = \{z_u : a(z_u) = \text{True}, 1 \leq u \leq U\}. \quad (7.1)$$

We implement each s_a as a bit-vector of length U and use it to *uniquely identify a problem* p in the dependency graph. The bits $\{s_{au}\}_u$ of s_a are defined as follows.

$$s_{au} = 1 \iff z_u \in s_a. \quad (7.2)$$

With each p , we track a set of values including its *lower bound* (lb) and *upper bound* (ub) of the optimal objective classifying its support set s_a .

In contrast, DL8.5 identifies each problem p using Boolean assertion a rather than s_a : this is an important difference between GOSDT and DL8.5 because many assertions a could correspond to the same support set s_a . However, given the same objective and algorithm, an optimal decision tree for any problem p depends only on the support set s_a . It does not depend on a if one already knows s_a . That is, if two different trees both contained a leaf capturing the same set of samples s , the set of possible child trees for that leaf is the same between the two trees. In GOSDT, we solve problems identified by s . That is, we solve all problems identified by any a that produces s .

7.2 Connection to Similar Support Bound

No previous approach has been able to implement a similar support bound effectively. We provide GOSDT's new form of similar support bound, called the *incremental similar support bound*. There are two reasons why this bound works where other attempts have failed: (1) This bound works on *partial trees*. One does not need to have constructed a full tree to use it. (2) The bound takes into account the hierarchical objective lower bound, and hence leverages the way we search the space within the DBP algorithm. In brief, the bound effectively removes many similar trees from the search region by looking at only one of them (which does not need to be a full tree). We consider weighted, additive, non-negative loss functions: $\ell(d, \mathbf{x}, \mathbf{y}) = \sum_i \text{weight}_i \text{loss}(x_i, y_i)$. Define the maximum weighted loss: $\ell^{\max} = \max_{x,y} [\text{weight}(x, y) \times \text{loss}(x, y)]$.

Theorem 7.1. (*Incremental Similar Support Bound*) Consider two trees $d = (d_{\text{fix}}, d_{\text{split}}, K, H)$ and $D = (D_{\text{fix}}, D_{\text{split}}, K, H)$ that differ only by their root node (hence they share the same K and H values). Further, the root nodes between the two trees are similar enough that the support going to the left and right branches differ by at most ω fraction of the observations. (That is, there are ωN observations that are captured either by the left branch of d and right branch of D or vice versa.) Define $S_{\text{uncertain}}$ as the maximum of the support within d_{split} and D_{split} : $S_{\text{uncertain}} = \max(\text{supp}(d_{\text{split}}), \text{supp}(D_{\text{split}}))$. For any child tree $d' \in \sigma(d)$ grown from d (grown from the nodes in d_{split} , that would not be excluded by the hierarchical objective lower bound) and for any child tree $D' \in \sigma(D)$ grown from D (grown from nodes in D_{split} , not excluded by the hierarchical objective lower bound), we have:

$$|R(d', \mathbf{x}, \mathbf{y}) - R(D', \mathbf{x}, \mathbf{y})| \leq (\omega + 2S_{\text{uncertain}})\ell^{\max}.$$

Proof. We will proceed in three steps. The first step is to show that

$$R(d, \mathbf{x}, \mathbf{y}) - R(D, \mathbf{x}, \mathbf{y}) \leq \omega \ell^{\max}.$$

The second step is to show:

$$R(d, \mathbf{x}, \mathbf{y}) \leq R(d', \mathbf{x}, \mathbf{y}) + S_{\text{uncertain}} \ell^{\max},$$

for all feasible children d' of d . The same bound will hold for D and any of its children D' . The third step is to show

$$R(d', \mathbf{x}, \mathbf{y}) \leq R(d, \mathbf{x}, \mathbf{y}) + S_{\text{uncertain}} \ell^{\max},$$

which requires different logic than the proof of Step 2. Together, Steps 2 and 3 give

$$|R(d', \mathbf{x}, \mathbf{y}) - R(d, \mathbf{x}, \mathbf{y})| \leq S_{\text{uncertain}} \ell^{\max}.$$

From here, we use the triangle inequality and the bounds from the three steps to obtain the final bound:

$$\begin{aligned} & |R(d', \mathbf{x}, \mathbf{y}) - R(D', \mathbf{x}, \mathbf{y})| \\ &= |R(d', \mathbf{x}, \mathbf{y}) - R(d, \mathbf{x}, \mathbf{y}) + R(d, \mathbf{x}, \mathbf{y}) - R(D, \mathbf{x}, \mathbf{y}) + R(D, \mathbf{x}, \mathbf{y}) - R(D', \mathbf{x}, \mathbf{y})| \\ &\leq |R(d', \mathbf{x}, \mathbf{y}) - R(d, \mathbf{x}, \mathbf{y})| + |R(d, \mathbf{x}, \mathbf{y}) - R(D, \mathbf{x}, \mathbf{y})| + |R(D, \mathbf{x}, \mathbf{y}) - R(D', \mathbf{x}, \mathbf{y})| \\ &\leq S_{\text{uncertain}} \ell^{\max} + \omega \ell^{\max} + S_{\text{uncertain}} \ell^{\max}, \end{aligned}$$

which is the statement of the theorem. Let us now go through the three steps.

First step: Define “move” as the set of indices of the observations that either go down the left branch of the root of d and the right of D , or that go down the right of d and the left of D . The remaining data will be denoted “/move.” These remaining data points will be classified the same way by both d and D . The expression above follows from the additive form of the objective R :

$$R(d, \mathbf{x}, \mathbf{y}) = \ell(d, \mathbf{x}^{\text{move}}, \mathbf{y}^{\text{move}}) + \ell(d, \mathbf{x}'^{\text{move}}, \mathbf{y}'^{\text{move}}) + \lambda H,$$

$$R(D, \mathbf{x}, \mathbf{y}) = \ell(D, \mathbf{x}^{\text{move}}, \mathbf{y}^{\text{move}}) + \ell(D, \mathbf{x}'^{\text{move}}, \mathbf{y}'^{\text{move}}) + \lambda H,$$

and since $\ell(d, \mathbf{x}'^{\text{move}}, \mathbf{y}'^{\text{move}}) = \ell(D, \mathbf{x}'^{\text{move}}, \mathbf{y}'^{\text{move}})$ since this just considers overlapping leaves, we have:

$$|R(d, \mathbf{x}, \mathbf{y}) - R(D, \mathbf{x}, \mathbf{y})| \leq |\ell(d, \mathbf{x}^{\text{move}}, \mathbf{y}^{\text{move}}) - \ell(D, \mathbf{x}^{\text{move}}, \mathbf{y}^{\text{move}})| \leq \omega \ell^{\max}.$$

(For the last inequality, the maximum is attained when one of $\ell(d, \mathbf{x}^{\text{move}}, \mathbf{y}^{\text{move}})$ and $\ell(D, \mathbf{x}^{\text{move}}, \mathbf{y}^{\text{move}})$ is zero and the other attains its maximum possible value.)

Second step: Recall that d' is a child of d so that d'_{fix} contains d_{fix} . Let us denote the leaves in d'_{fix} that are not in d_{fix} by $d'_{\text{fix}}/d_{\text{fix}}$. Then,

$$\begin{aligned} R(d', \mathbf{x}, \mathbf{y}) &= \ell(d'_{\text{fix}}, \mathbf{x}, \mathbf{y}) + \ell(d'_{\text{split}}, \mathbf{x}, \mathbf{y}) + \lambda H_{d'} \\ &= \ell(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) + \ell(d'_{\text{fix}}/d_{\text{fix}}, \mathbf{x}, \mathbf{y}) + \ell(d'_{\text{split}}, \mathbf{x}, \mathbf{y}) + \lambda H_{d'}. \end{aligned}$$

Adding $\ell(d_{\text{split}}, \mathbf{x}, \mathbf{y})$ to both sides,

$$\begin{aligned} R(d', \mathbf{x}, \mathbf{y}) + \ell(d_{\text{split}}, \mathbf{x}, \mathbf{y}) &= \ell(d_{\text{fix}}, \mathbf{x}, \mathbf{y}) + \ell(d_{\text{split}}, \mathbf{x}, \mathbf{y}) + \ell(d'_{\text{fix}}/d_{\text{fix}}, \mathbf{x}, \mathbf{y}) \\ &\quad + \ell(d'_{\text{split}}, \mathbf{x}, \mathbf{y}) + \lambda H_{d'} \\ &= R(d, \mathbf{x}, \mathbf{y}) + \ell(d'_{\text{fix}}/d_{\text{fix}}, \mathbf{x}, \mathbf{y}) + \ell(d'_{\text{split}}, \mathbf{x}, \mathbf{y}) + \lambda(H_{d'} - H) \\ &\geq R(d, \mathbf{x}, \mathbf{y}), \end{aligned} \tag{7.3}$$

since the terms we removed were all nonnegative. Now,

$$\begin{aligned} R(d, \mathbf{x}, \mathbf{y}) &\leq R(d', \mathbf{x}, \mathbf{y}) + \ell(d_{\text{split}}, \mathbf{x}, \mathbf{y}) \\ &\leq R(d', \mathbf{x}, \mathbf{y}) + \text{supp}(d_{\text{split}})\ell^{\max} \\ &\leq R(d', \mathbf{x}, \mathbf{y}) + S_{\text{uncertain}}\ell^{\max}. \end{aligned}$$

Third step: Here we will use the hierarchical objective lower bound. We start by noting that since we have seen d , we have calculated its objective $R(d, \mathbf{x}, \mathbf{y})$, and it must be as good or worse than the current best value that we have seen so far (or else it would have replaced the current best). So $R^c \leq R(d, \mathbf{x}, \mathbf{y})$. The hierarchical objective lower bound (Theorem 3.1) would be violated if the following holds. This expression states that $b(d', \mathbf{x}, \mathbf{y})$ is worse than $R(d, \mathbf{x}, \mathbf{y})$ (which is worse than the current best), which means we would have already excluded d' from consideration:

$$R^c \leq R(d, \mathbf{x}, \mathbf{y}) < b(d', \mathbf{x}, \mathbf{y}) = R(d', \mathbf{x}, \mathbf{y}) - \ell(d'_{\text{split}}, \mathbf{x}, \mathbf{y}).$$

This would be a contradiction. Thus, the converse holds:

$$R(d', \mathbf{x}, \mathbf{y}) - \ell(d'_{\text{split}}, \mathbf{x}, \mathbf{y}) = b(d', \mathbf{x}, \mathbf{y}) \leq R(d, \mathbf{x}, \mathbf{y}).$$

Thus,

$$R(d', \mathbf{x}, \mathbf{y}) \leq R(d, \mathbf{x}, \mathbf{y}) + \ell(d'_{\text{split}}, \mathbf{x}, \mathbf{y}).$$

Now to create an upper bound for $\ell(d'_{\text{split}}, \mathbf{x}, \mathbf{y})$ as ℓ^{\max} times the support of d'_{split} . Since d' is a child of d , its support on the split leaves is less than or equal to that of d , $\text{supp}(d'_{\text{split}}) \leq \text{supp}(d_{\text{split}})$. Thus, $\ell(d'_{\text{split}}, \mathbf{x}, \mathbf{y}) \leq \ell^{\max} \text{supp}(d_{\text{split}}) \leq \ell^{\max} S_{\text{uncertain}}$. Hence, we have the result for the final step of the proof, namely:

$$R(d', \mathbf{x}, \mathbf{y}) \leq R(d, \mathbf{x}, \mathbf{y}) + S_{\text{uncertain}} \ell^{\max}.$$

□

Unlike the similar support bounds in CORELS and OSDT, which require pairwise comparisons of problems, this incremental similar support bound emerges from the support set representation. The descendants $\sigma(d)$ and $\sigma(D)$ share many of the same support sets. Because of this, the shared components of their upper and lower bounds are updated simultaneously (to similar values). This bound is helpful when our data contains continuous variables: if a split at value v was already visited, then splits at values close to v can reuse prior computation to immediately produce tight upper and lower bounds.

7.3 Asynchronous Bound Updates

GOSDT computes the objective values hierarchically by defining the minimum objective $R^*(p)$ of problem p as an aggregation of minimum objectives over the child problems p_l^j and p_r^j of p for $1 \leq j \leq M$.

$$R^*(p) = \min_j (R^*(p_l^j) + R^*(p_r^j)). \quad (7.4)$$

Since DL8.5 computes each $R^*(p_l^j)$ and $R^*(p_r^j)$ in a blocking call to the child problems p_l^j and p_r^j , it necessarily computes $R^*(p_l^j) + R^*(p_r^j)$ after solving p_l^j and p_r^j , which is a disadvantage. In contrast, GOSDT computes bounds over $R^*(p_l^j) + R^*(p_r^j)$ that are available before knowing the exact values of $R^*(p_l^j)$ and $R^*(p_r^j)$. The bounds over $R^*(p_l^j)$ and $R^*(p_r^j)$ are solved asynchronously and possibly in parallel. If for some j and j' the bounds imply that $R^*(p_l^j) + R^*(p_r^j) > R^*(p_l^{j'}) +$

$R^*(p_r^{j'})$, then we can conclude that p 's solution no longer depends on p_l^j and p_r^j . Since GOSDT executes asynchronously, it can draw this conclusion and focus on $R^*(p_l^{j'}) + R^*(p_r^{j'})$ without fully solving $R^*(p_l^j) + R^*(p_r^j)$.

To encourage this type of bound update, GOSDT uses the priority queue to send high-priority signals to each parent p_π of p when an update is available, prompting a recalculation of $R^*(p_\pi)$ using Equation 7.4.

7.4 Fast Selective Vector Sums

New problems (i.e., p_l and p_r) require initial upper and lower bounds on the optimal objective. We define the initial lower bound lb and upper bound ub for a problem p identified by support set s as follows, for $1 \leq u \leq U$:

$$z_u^{\min} = \min(z_u^-, z_u^+) \quad (7.5)$$

This is the fraction of minority class samples in equivalence class u .

$$lb = \lambda + \sum_u s_u z_u^{\min} \quad (7.6)$$

This is a basic equivalence points bound, that predicts all minority class equivalence points incorrectly.

$$ub = \lambda + \min\left(\sum_u s_u z_u^-, \sum_u s_u z_u^+\right). \quad (7.7)$$

This upper bound comes from a baseline algorithm of predicting all one class.

We use the *prefix sum trick* in order to speed up computations of sums of a subset of elements of a vector. That is, for any vector z^{vec} (e.g., z^{\min} , z^- , z^+), we want to compute a sum of a subsequence of z^{vec} : we precompute (during preprocessing) a prefix sum vector $z^{\text{cumulative}}$ of the vector z^{vec} defined by the cumulative sum $z_u^{\text{cumulative}} = \sum_{j=1}^u z_j^{\text{vec}}$. During the algorithm, to sum over ranges of contiguous values in z^{vec} , over some indices a through b , we now need only take $z^{\text{cumulative}}[b] - z^{\text{cumulative}}[a - 1]$. This reduces a linear time calculation to constant time. This fast sum is leveraged over calculations with the support sets of input features—for example, quickly determining the difference in support sets between two features.

7.5 The GOSDT Algorithm

Algorithm 1 constructs and optimizes problems in the *dependency graph* such that, upon completion, we can extract the optimal tree by traversing the dependency graph by greedily choosing the split with the lowest objective value. We present the key components of this algorithm, highlighting the differences between GOSDT and DL8.5. Note that all features have been binarized prior to executing the algorithm.

Lines 8 to 11: Remove an item from the queue. If its bounds are equal, no further optimization is possible and we can proceed to the next item on the queue.

Lines 12 to 18: Construct new subproblems, p_l, p_r by splitting on feature j . Use lower and upper bounds of p_l and p_r to compute new bounds for p . *We key the problems by the bit vector corresponding to their support set s .* Keying problems in this way avoids processing the same problem twice; other dynamic programming implementations, such as DL8.5, will process the same problem multiple times.

Lines 19 to 22: Update p with lb' and ub' computed using Equation 7.4 and propagate that update to all ancestors of p by enqueueing them with high priority (p will have multiple parents if two different conjunctions produce the same support set). This triggers the ancestor problems to recompute their bounds using Equation 7.4. *The high priority ensures that ancestral updates occur before processing p_l and p_r .* This scheduling is one of the key differences between GOSDT and DL8.5; by eagerly propagating bounds up the dependency tree, GOSDT prunes the search space more aggressively.

Lines 25 to 31: Enqueue p_l and p_r only if the interval between their lower and upper bounds overlaps with the interval between p 's lower and upper bounds. This ensures that eventually the lower and upper bounds of p converge to a single value.

Lines 33 to 41: Define **FIND_OR_CREATE_NODE**, which constructs (or finds an existing) problem p corresponding to support set s , initializing the lower and upper bound and checking if p is eligible to be split, using the Incremental Progress Bound to Determine Splitting (Theorem 3.6) and the Lower Bound on Incremental Progress (Theorem 3.7) (these bounds are checked in subroutine *fails_bounds*). *get_lower_bound* returns lb using Equation 7.6. *get_upper_bound* returns

ub using Equation 7.7.

Algorithm 1 GOSDT(R, x, y, λ)

```
1: input:  $R, Z, z^-, z^+, \lambda$  // risk, samples, labels, regularizer
2:  $Q = \emptyset$  // priority queue
3:  $G = \emptyset$  // dependency graph
4:  $s_0 \leftarrow \{1, \dots, 1\}$  // bit-vector of 1's of length  $U$ 
5:  $p_0 \leftarrow \text{FIND\_OR\_CREATE\_NODE}(G, s_0)$  // node for root
6:  $Q.\text{push}(s_0)$  // add to priority queue
7: while  $p_0.lb \neq p_0.ub$  do
8:    $s \leftarrow Q.\text{pop}()$  // index of problem to work on
9:    $p \leftarrow G.\text{find}(s)$  // find problem to work on
10:  if  $p.lb = p.ub$  then
11:    continue // problem already solved
12:   $(lb', ub') \leftarrow (\infty, \infty)$  // very loose starting bounds
13:  for each feature  $j \in [1, M]$  do
14:     $s_l, s_r \leftarrow \text{split}(s, j, Z)$  // create children if they don't exist
15:     $p_l^j \leftarrow \text{FIND\_OR\_CREATE\_NODE}(G, s_l)$ 
16:     $p_r^j \leftarrow \text{FIND\_OR\_CREATE\_NODE}(G, s_r)$ 
17:     $lb' \leftarrow \min(lb', p_l^j.lb + p_r^j.lb)$ 
18:     $ub' \leftarrow \min(ub', p_l^j.ub + p_r^j.ub)$ 
```

Algorithm 1 continue

```
19:   if  $p.lb \neq lb'$  or  $p.ub \neq ub'$  then
20:      $(p.lb, p.ub) \leftarrow (lb', ub')$ 
21:     for  $p_\pi \in G.parent(p)$  do
22:        $Q.push(p_\pi.id, priority = 1)$  // propagate information upwards
23:     if  $p.lb \geq p.ub$  then
24:       continue // problem solved just now
           // loop, enqueue all children that are dependencies
25:     for each feature  $j \in [1, M]$  do
26:       repeat line 14-16 // fetch  $p_l^j$  and  $p_r^j$  in case of update from other thread
27:        $lb' \leftarrow p_l^j.lb + p_r^j.lb$ 
28:        $ub' \leftarrow p_l^j.ub + p_r^j.ub$ 
29:       if  $lb' < ub'$  and  $lb' \leq p.ub$  then
30:          $Q.push(s_l, priority = 0)$ 
31:          $Q.push(s_r, priority = 0)$ 
32: return
```

```
33: subroutine FIND_OR_CREATE_NODE(G,s)
34:   if  $G.find(s) = NULL$  //  $p$  not yet in dependency graph
35:      $p.id \leftarrow s$  // identify  $p$  by  $s$ 
36:      $p.lb \leftarrow get\_lower\_bound(s, Z, z^-, z^+)$ 
37:      $p.ub \leftarrow get\_upper\_bound(s, Z, z^-, z^+)$ 
38:     if  $fails\_bounds(p)$  then
39:        $p.lb = p.ub$  // no more splitting allowed
40:      $G.insert(p)$  // put  $p$  in dependency graph
41:   return  $G.find(s)$ 
```

Algorithm 2 $\text{get_lower_bound}(s, Z, z^-, z^+) \rightarrow lb$

input: s, Z, z^-, z^+ // support, unique sample set, negative sample set, positive sample set**output:** lb // Risk lower bound*// Compute the risk contributed if applying a class to every equivalent class independently***for** each equivalence class $u \in [1, U]$ **define***// Values provided in Z*

$$z_u^- = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[y_i = 0 \wedge x_i = z_u]$$

$$z_u^+ = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[y_i = 1 \wedge x_i = z_u]$$

// Risk of assigning a class to equivalence class u

$$z_u^{\min} = \min(z_u^-, z_u^+)$$

// Add all risks for each class u and a single λ (lower bound of the complexity penalty)

$$lb \leftarrow \lambda + \sum_u s_u z_u^{\min}$$

return lb

Algorithm 3 $\text{get_upper_bound}(s, Z, z^-, z^+) \rightarrow ub$

input: s, Z, z^-, z^+ // support, unique sample set, negative sample set, positive sample set**output:** ub // Risk upper bound*// Compute the risk contributed if applying a single class to all samples in s* **for** each equivalence class $u \in [1, U]$ **define***// Add up the positive and negative class weights under equivalence class u*

$$z_u^- = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[y_i = 0 \wedge x_i = z_u]$$

$$z_u^+ = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[y_i = 1 \wedge x_i = z_u]$$

*// Total the positive and negatives over all classes u , choosing the smaller total as the**misclassification and add a single λ for the complexity penalty of a leaf*

$$ub \leftarrow \lambda + \min(\sum_u s_u z_u^-, \sum_u s_u z_u^+)$$

return ub

Algorithm 4 $\text{fails_bounds}(p) \rightarrow v$

input: p // current problem

output: v // boolean indicating valid problem

// If this expression is true then the lower bound on incremental accuracy is crossed by

// all descendents. This works since $ub - lb$ is an upperbound on incremental accuracy

// for any descent

$\text{incremental_accuracy} \leftarrow p.ub - p.lb \leq \lambda$

// If expression is true then the lower bound on leaf classification accuracy is crossed

$\text{leaf_accuracy} \leftarrow p.ub \leq 2\lambda$

if $(\text{incremental_accuracy} = \text{True}) \vee (\text{leaf_accuracy} = \text{True})$ **then**

return True

return False

Algorithm 5 $\text{split}(s, j, Z) \rightarrow s_l, s_r$

input: s, j, Z // support set, feature index, unique sample set

output: s_l, s_r // left and right splits

// Create the left key which is the subset of s such that feature j tests negative

$s_l = \{\mathbb{1}[s_u = 1 \wedge Z_{u,j} = 0] \mid 1 \leq u \leq U\}$

// Create the right key which is the subset of s such that feature j tests positive

$s_r = \{\mathbb{1}[s_u = 1 \wedge Z_{u,j} = 1] \mid 1 \leq u \leq U\}$

return s_l, s_r

Algorithm 6 $\text{extract}(t) \rightarrow s$ // Extract optimal tree after running the algorithm

input: s // Key of the problem from which we want to build a tree

output: t // Optimal tree

$p \leftarrow \text{FIND_OR_CREATE_NODE}(G, s)$ // Find the node associated to this key

$t \leftarrow \text{None}$ // Create a null tree

$\text{base_bound} \leftarrow p.\text{ub}$ // The risk if we end this node as a leaf

$\text{base_prediction} \leftarrow 0$ // The prediction if we end this node as a leaf

$\text{split_bound} \leftarrow \infty$ // The risk if we split this node

$\text{split_feature} \leftarrow 0$ // The index of the feature we should use to split this node

for each feature $j \in [1, M]$ **do** // Check all possible features

$s_l, s_r \leftarrow \text{split}(s, j, Z)$ // Key of the the children for this split

$p_l^j \leftarrow \text{FIND_OR_CREATE_NODE}(G, s_l)$ // Find left child

$p_r^j \leftarrow \text{FIND_OR_CREATE_NODE}(G, s_r)$ // Find right child

 // Check if the risk of this split is better than the best split risk so far

if $p_l^j.\text{ub} + p_r^j.\text{ub} < \text{split_bound}$ **then**

$\text{split_bound} \leftarrow p_l^j.\text{ub} + p_r^j.\text{ub}$ // Update the best split risk

$\text{split_feature} \leftarrow j$ // Best feature index to split

 // Calculate the total positive and negative weights of each equivalence class

for each equivalence class $u \in [1, U]$ **define**

 // Values come from equivalence class matrix Z as seen in Algorithm 3

$z_u^- = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[y_i = 0 \wedge x_i = z_u]$ // total negatives

$z_u^+ = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[y_i = 1 \wedge x_i = z_u]$ // total positives

 // Select only the positive and negative weights captured by s

$\text{negatives} \leftarrow \sum_u s_u z_u^-$

$\text{positives} \leftarrow \sum_u s_u z_u^+$

Algorithm 6 continue

```
// Set the leaf prediction based on class with the higher selected total weight  
if negatives < positives then  
    // Leaf predicts the majority class as 1 since positive weights are higher  
    base_prediction.pred  $\leftarrow$  1  
  
// Base case: If the risk of remaining as a leaf is better than splitting, remain as leaf  
if base_bound  $\leq$  split_feature  
    // Construct and return a leaf node  
    t.type  $\leftarrow$  leaf  
    t.prediction  $\leftarrow$  base_prediction  
    return t  
  
// Recursive case: One of the splits performs better than the leaf  
// Generate left and right splits based on best split feature  
sl, sr  $\leftarrow$  split(s, split_feature, Z)  
  
// Recurse onto child keys to create left and right subtrees  
tl  $\leftarrow$  extract(sl)  
tr  $\leftarrow$  extract(sr)  
  
// Construct and return a split node containing the left and right subtrees  
t.type  $\leftarrow$  tree  
t.split  $\leftarrow$  split_feature  
t.left  $\leftarrow$  tl  
t.right  $\leftarrow$  tr  
return t
```

Chapter 8

Comparison Between Decision Tree Methods

Table 8.1 demonstrate differences between GOSDT and previous decision tree approaches.

Table 8.1: Comparison of decision tree approaches. **Green** is a comparative advantage. **Red** is a comparative disadvantage. **Blue** highlights dynamic programming-based methods. White is neutral.

Attributes	GOSDT	DL8	DL8.5	BinOCT	DTC	CART
Built-in Preprocessor	Yes	No	No	Yes	No	No
Preprocessing Strategy	All values	Weka Discretization	Bucketization	Bucketization	None	None
Preprocessing Preserves Optimality?	Yes	No	No	No	N/A	N/A
Optimization Strategy	DPB	DP	DPB	ILP	Greedy Search	Greedy Search
Optimization Preserves Optimality?	Yes	Yes	Yes	Yes	No	No
Applies Hierarchical Upper bound to Reduce Search Space?	Yes	No	Yes	N/A	N/A	N/A
Uses Support Set Node Identifiers?	Yes	No	No	N/A	N/A	N/A
Can use Multiple Cores?	Yes	No	No	Yes	No	No
Can prune using updates from partial evaluation of subproblem?	Yes	No	No	Depends on (generic) solver	N/A	N/A
Strategy for preventing overfitting	Penalize Leaves	Structural Constraints	Structural Constraints	Structural Constraints	MDL Criteria	Structural Constraints
Can we modify this to use regularization?	N/A	Yes	Yes	Yes	N/A	No
Does it address class imbalance?	Yes	Maybe	Maybe	No	Maybe	Maybe

Chapter 9

Experiments

GOSDT’s novelty lies in its ability to optimize a large class of objective functions and its ability to efficiently handle continuous variables without sacrificing optimality. Thus, our evaluation results:

- 1) Demonstrate our ability to optimize over a large class of objectives (AUC in particular),
- 2) Show that GOSDT outperforms other approaches in producing models that are both accurate and sparse,
- 3) Show how GOSDT scales in its handling of continuous variables, relative to other methods, and
- 4) Show time-to-optimality results.

We used a total of 9 data sets: Five of them are from the UCI Machine Learning Repository [DKT17], (*monk1*, *monk2*, *monk3*, *tic-tac-toe*, *car evaluation*), one from LIBSVM [CL11], (*FourClass*). The other three datasets are the ProPublica recidivism dataset [LMKA16] (*COMPAS*), the Fair Isaac credit risk data sets [FGI⁺18] (*FICO*), and the mobile advertisement data sets [WRDV⁺17] (*coupon*). We predict which individuals are arrested within two years of release ($N = 5,020$) on the recidivism data set, whether an individual will default on a loan for the FICO dataset, and whether a customer is going to accept a coupon for a bar considering demographic and contextual attributes. Table 9.1 presents a summary of the datasets.

The experiments for optimizing rank statistics were run on a personal laptop with a 2.6GHz i5-7300U processor and 8GB of RAM. All other experiments were run on a 2.30 GHz (30 MB L3 cache) Intel Xeon E7-4870 v2 processor with 60 cores across 4 NUMA nodes. We disabled hyper-threading. The server has 512 GB RAM uniformly distributed (128 GB each) across the four NUMA nodes. The host OS is Ubuntu 16.04.6 LTS. We set a 5-minute time limit on all experiments, unless otherwise stated. All algorithms that support multi-threading are modified to run sequentially.

Optimizing Many Different Objectives: We use the *FourClass* dataset [CL11] to show optimal decision trees corresponding to different objectives. Figures 9.1 and 9.2 show the training ROC

and test ROC of decision trees generated for six different objectives and optimal trees for accuracy, AUC and partial area under ROC convex hull. Optimizing different objectives produces different trees with different FP and FN . (No other decision tree method is designed to fully optimize any objective except accuracy, so there is no comparison to other methods.)

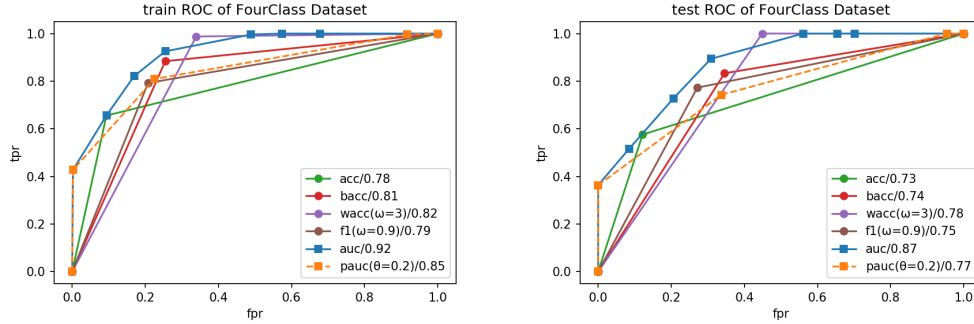


Figure 9.1: Training ROC and test ROC of FourClass dataset ($\lambda = 0.01$, time limit=30min). A/B in the legend at the bottom right of each subfigures shows the objective and its parameter/area under the ROC.

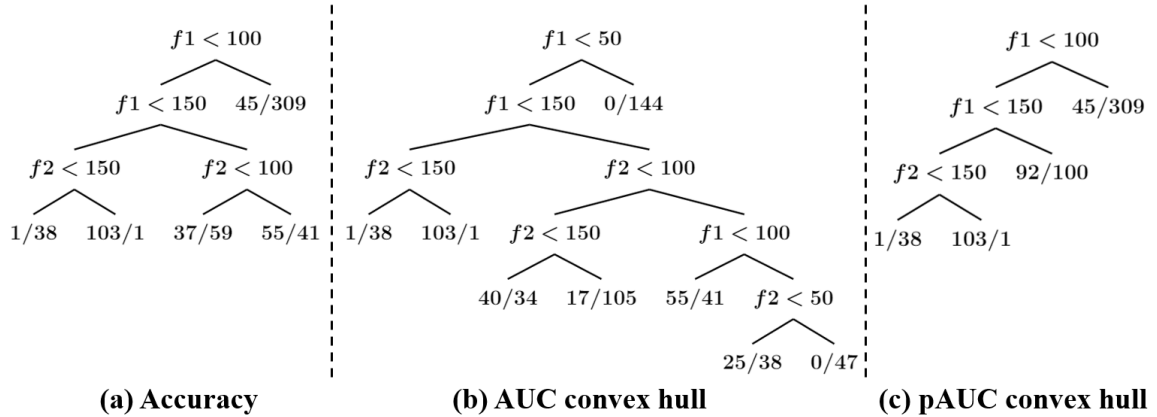


Figure 9.2: Trees for different objectives. A/B in the leaf node represents the number of positive/negative samples amongst the total samples captured in that leaf.

Binary Datasets: Accuracy vs Sparsity: We compare models produced from BinOCT [VZ19], CART [BFOS84], DL8.5 [NS⁺20], OSDT [HRS19], and GOSDT. For each method, we select hyperparameters to produce trees of varying numbers of leaves and plot training accuracy against sparsity (number of leaves). GOSDT directly optimizes the trade-off between training accuracy and number of leaves, producing points on the efficient frontier. Figure 9.3 shows (1) that GOSDT typically produces excellent training accuracy with a reasonable number of leaves, and (2) that we

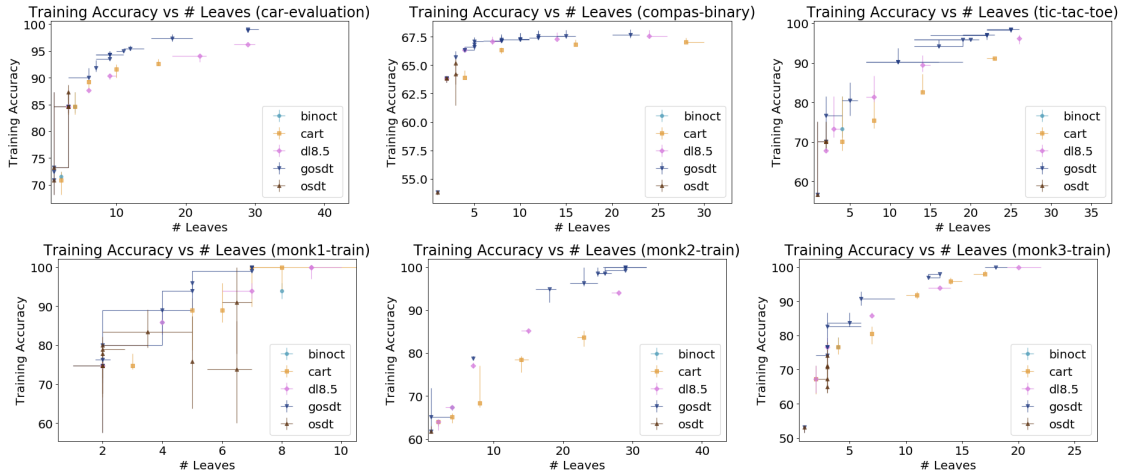


Figure 9.3: Training accuracy achieved by BinOCT, CART, DL8.5, GOSDT, and OSDT as a function of the number of leaves.

can quantify how close to optimal the other methods are.

Continuous Datasets: Slowdown vs Thresholds: We preprocessed by encoding continuous variables as a set of binary variables, using all possible thresholds. We then reduced the number of binary variables by combining sets of k variables for increasing values of k . Binary variables were ordered, firstly, in order of their thresholds, and secondly, in order of the indices of the continuous variable they correspond to. We measure the slow-down introduced by increasing the number of binary variables relative to each algorithm’s own best time. Figure 9.4 shows these results for CART, DL8.5, GOSDT, and OSDT. As the number of features increases (i.e., k approaches 1), GOSDT typically slows down less than DL8.5 and OSDT. This relatively smaller slowdown allows GOSDT to handle more thresholds introduced by continuous variables.

Time to Optimality: We run *OSDT*, *GOSDT*, and *PyGOSDT* with a regularization coefficient of 0.005. Note GOSDT is implemented in C++ and run as a native executable and PyGOSDT is an early Python implementation of GOSDT. For each run we track the progress of the algorithm by plotting the minimum objective score seen so far. Once the algorithm terminates or reaches a time limit of 30 minutes, the values are stored in a file. Figure 9.5 shows the different behaviors between GOSDT, PyGOSDT, and OSDT. In general, both PyGOSDT and GOSDT complete their certificate of optimality earlier than OSDT.

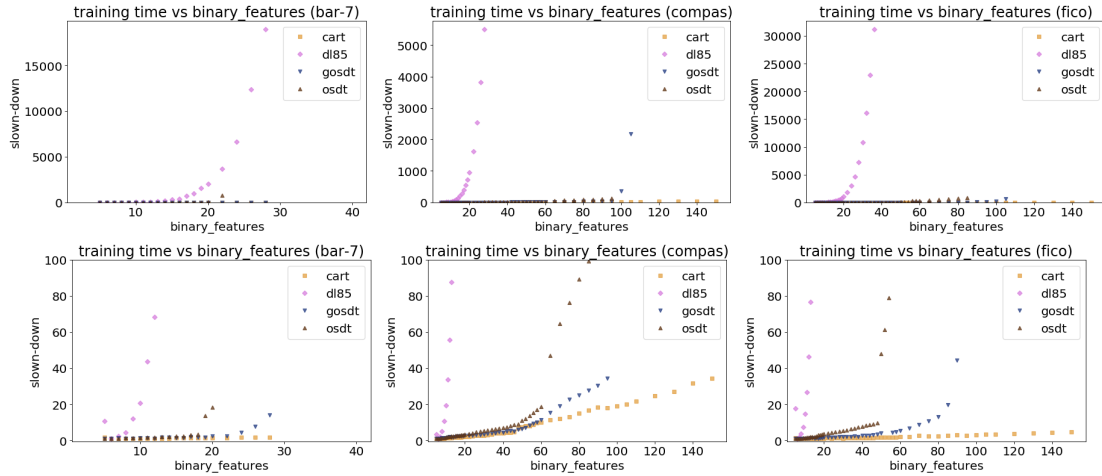


Figure 9.4: Training time of BinOCT, CART, DL8.5, GOSDT, and OSDT as a function of the number of binary features used to encode the continuous dataset ($\lambda = 0.3125$ or max depth = 5).

Note that PyGOSDT’s implementation does not include high-priority bound updates. This causes PyGOSDT to maintain a higher objective score before making a sharp drop upon completion (with a certificate of optimality). GOSDT, on the other hand, behaves similarly to OSDT because both algorithms aggressively prioritize lowering the best observed objective score. We observe that under the tic-tac-toe data set this appears to be less advantageous. While PyGOSDT’s progress initially appears less promising, it completed remarkably faster than both GOSDT and OSDT. This suggests that optimal prioritization is dependent on specifics of the optimization problem.

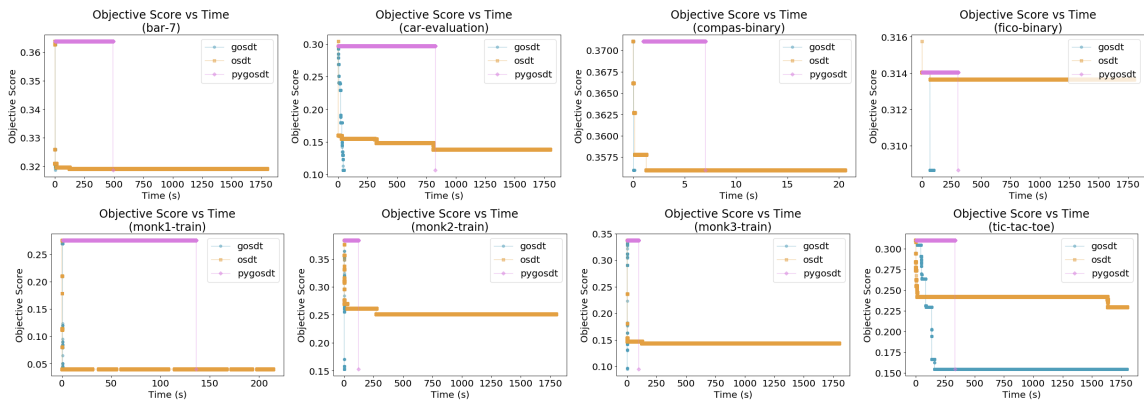
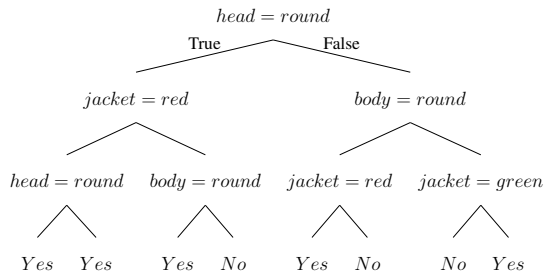
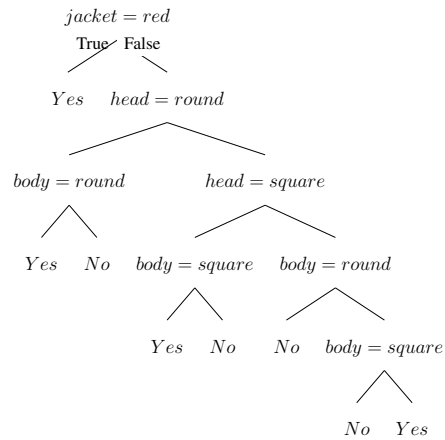


Figure 9.5: Best objective score of OSDT, GOSDT, and PyGOSDT over the course of their run time. ($\lambda = 0.005$)

Optimal Trees: We present some of the trees that achieved the peak median accuracy from Figures 9.3 and 9.4. Figure 9.6 shows a comparison between the results of training BinOCT (a) and GOSDT (b) on the Monk 1 data set. GOSDT is able to produce a model with 20% higher accuracy than BinOCT even though both trees have 8 leaves. Figure 9.7 shows a comparison between DL8.5 (a) and GOSDT (b) on the Monk 2 data set. GOSDT is able to produce a model with 3% higher accuracy than DL8.5 even though both trees have 7 leaves. Figure 9.8 shows a comparison between BinOCT (a), DL8.5 (b), and GOSDT (c) on the Tic-Tac-Toe data set. GOSDT is able to produce a model with higher accuracy than both BinOCT and DL8.5 when all trees have 16 leaves.

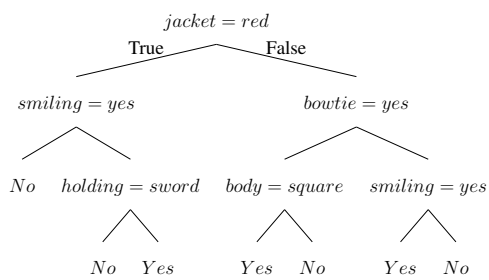


(a) BinOCT (training accuracy: 90.9%, test accuracy: 84.0%)

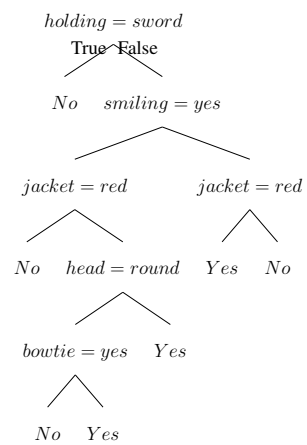


(b) GOSDT (training accuracy: 100%, test accuracy: 100%)

Figure 9.6: Eight-leaf decision trees generated by BinOCT and GOSDT on the Monk1 dataset. The tree generated by BinOCT includes two useless splits (the head=round in the bottom left), while GOSDT avoids this problem. BinOCT is 91% accurate, GOSDT is 100% accurate.



(a) DL8.5 (training accuracy: 76.3%, test accuracy: 70.6%)

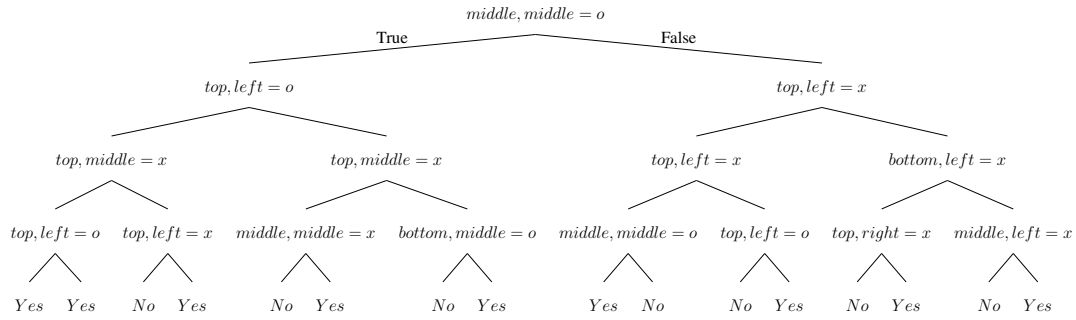


(b) GOSDT (training accuracy: 79.3%, test accuracy: 73.5%)

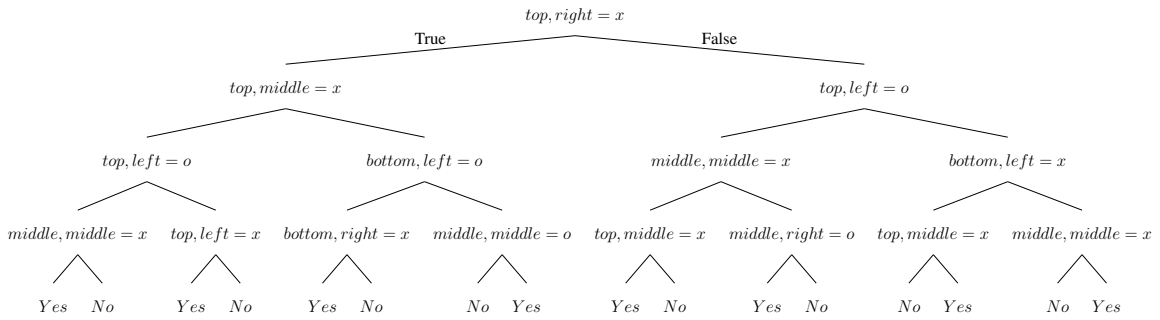
Figure 9.7: Seven-leaf decision trees generated by DL8.5 and GOSDT on the Monk2 dataset. With the same number of leaves, DL8.5 is 76.3% accurate, GOSDT is 79.3% accurate.

Table 9.1: Comparison of different data sets (and their preprocessed derivatives). **Categorical Features** denote the number of features that are inherently categorical or represent discrete ranges of a continuous feature. These features generally have a low cardinality. **Continuous Features** denote the number of features that take on many integer or real number values. These features generally have a high cardinality. **Encoded Binary Features** denote the number of binary features required to encode all categorical and continuous values observed in the dataset. This value is approximately the total cardinality of all categorical and continuous features.

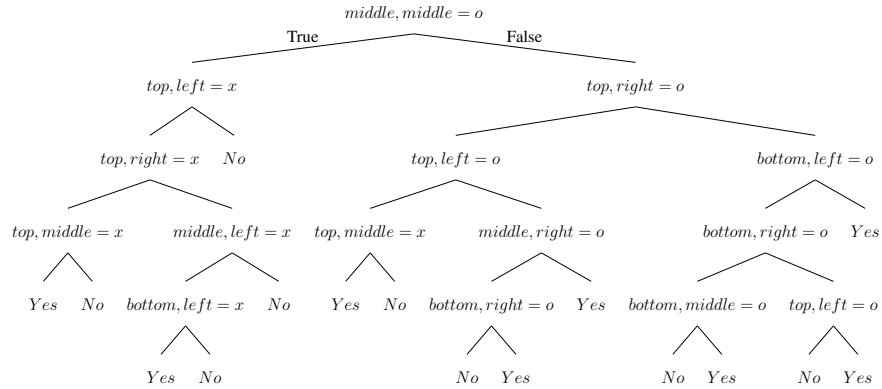
Data Set	Samples	Categorical Features	Continuous Features	Encoded Binary Features
monk 1	124	11	0	11
monk 2	169	11	0	11
monk 3	122	11	0	11
car evaluation	1729	15	0	15
tic-tac-toe	958	18	0	18
coupon	12684	26	0	129
bar-7	1913	5	0	14
Four Class	862	0	2	345
COMPAS	12382	2	20	647
COMPAS-2016	5020	4	2	85
COMPAS-binary	6907	12	0	12
FICO	1000	0	23	1407
FICO-binary	10459	17	0	17



(a) BinOCT (training accuracy: 82.4%, test accuracy 34.4%)



(b) DL8.5 (training accuracy: 86.9%, test accuracy: 61.5%)



(c) GOSDT (training accuracy: 90.9%, test accuracy: 70.8%)

Figure 9.8: 16-leaf decision trees generated by BinOCT, DL8.5, and GOSDT on the tic-tac-toe dataset. The tree generated by BinOCT includes some useless splits such as *top, left = o* on the bottom left and *middle, middle = o* near the center of the bottom row. These extra splits repeat earlier decisions from the tree, so they are clearly useless and lead to at least one empty leaf. DL8.5 also prefers to generate complete binary trees. GOSDT is more effective in generating sparse trees. With the same number of leaves, BinOCT is 82.4% accurate, DL8.5 is 86.9% accurate, and GOSDT is 90.9% accurate.

Chapter 10

Conclusion and Future Work

GOSDT (and related methods) differ fundamentally from other types of machine learning algorithms. Unlike neural networks and typical decision tree methods, it provides a proof of optimality for highly non-convex problems. Unlike support vector machines, ensemble methods, and neural networks again, it produces sparse interpretable models without using convex proxies—it solves the exact problem of interest in an efficient and provably optimal way. As usual, statistical learning theoretic guarantees on test error are the tightest for simpler models (sparse models) with the lowest empirical error on the training set, hence the choice of accuracy, regularized by sparsity (simplicity). If GOSDT is stopped early, it reports an optimality gap, which allows the user to assess whether the tree is sufficiently close to optimal, retaining learning theoretic guarantees on test performance. GOSDT is most effective for datasets with a small or medium number of features. The algorithm scales well with the number of observations, easily handling tens of thousands.

There are many avenues for future work. Since GOSDT provides a framework to handle objectives that are monotonic in FP and FN , one could create many more objectives than we have enumerated here. Going beyond these objectives to handle other types of monotonicity, fairness, ease-of-use, and cost-related soft and hard constraints are natural extensions. There are many avenues to speed up GOSDT related to exploration of the search space, garbage collection, and further bounds.

Bibliography

- [AFG⁺19] Ulrich Aïvodji, Julien Ferry, Sébastien Gambs, Marie-José Huguet, and Mohamed Siala. Learning Fair Rule Lists. *arXiv e-prints*, page arXiv:1909.03977, Sep 2019.
- [ALSA⁺17a] E. Angelino, N. Larus-Stone, D. Alabi, M. Seltzer, and C. Rudin. Learning certifiably optimal rule lists for categorical data. In *Proc. ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2017.
- [ALSA⁺17b] Elaine Angelino, Nicholas Larus-Stone, Daniel Alabi, Margo Seltzer, and Cynthia Rudin. Learning certifiably optimal rule lists. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 35–44. ACM, 2017.
- [ALSA⁺18] Elaine Angelino, Nicholas Larus-Stone, Daniel Alabi, Margo Seltzer, and Cynthia Rudin. Learning certifiably optimal rule lists for categorical data. *Journal of Machine Learning Research*, 18(234):1–78, 2018.
- [BB96] K. P. Bennett and J. A. Blue. Optimal decision trees. Technical report, R.P.I. Math Report No. 214, Rensselaer Polytechnic Institute, 1996.
- [BCMRM18] Rafael Blanquero, Emilio Carrizosa, Cristina Molero-Río, and Dolores Romero Morales. Optimal randomized classification trees. August 2018.
- [BD17] Dimitris Bertsimas and Jack Dunn. Optimal classification trees. *Machine Learning*, 106(7):1039–1082, 2017.
- [Ben92] Kristin Bennett. Decision tree construction via linear programming. In *Proceedings of the 4th Midwest Artificial Intelligence and Cognitive Science Society Conference, Utica, Illinois*, 1992.
- [BFOS84] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification and Regression Trees*. Wadsworth, 1984.
- [Bre17] Leo Breiman. *Classification and regression trees*. Routledge, 2017.
- [CBHK02] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- [CC08] David A Cieslak and Nitesh V Chawla. Learning decision trees for unbalanced data. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 241–256. Springer, 2008.
- [CL11] Chih-Chung Chang and Chih-Jen Lin. Libsvm: a library for support vector machines, 2011. software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [CR18] C. Chen and C. Rudin. An optimization approach to learning falling rule lists. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.

- [DKT17] Dua Dheeru and Efi Karra Taniskidou. UCI machine learning repository, 2017.
- [FFHO02] César Ferri, Peter Flach, and José Hernández-Orallo. Learning decision trees using the area under the roc curve. In *ICML*, volume 2, pages 139–146, 2002.
- [FGI⁺18] FICO, Google, Imperial College London, MIT, University of Oxford, UC Irvine, and UC Berkeley. Explainable Machine Learning Challenge. <https://community.fico.com/s/explainable-machine-learning-challenge>, 2018.
- [Gar72] M. R. Garey. Optimal binary identification procedures. *SIAM Journal on Applied Mathematics*, 23(2):173–186, 1972.
- [GHR03] M. Garofalakis, D. Hyun, R. Rastogi, and K. Shim. Building decision trees with constraints. *Data Mining and Knowledge Discovery*, 7:187–214, 2003.
- [HRS19] Xiyang Hu, Cynthia Rudin, and Margo Seltzer. Optimal sparse decision trees. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2019.
- [KS06] Adam R. Klivans and Rocco A. Servedio. Toward attribute efficient learning of decision lists and parities. *Journal of Machine Learning Research*, 7:587–602, 2006.
- [KWS14] Bartosz Krawczyk, Michał Woźniak, and Gerald Schaefer. Cost-sensitive decision tree ensembles for effective imbalanced classification. *Applied Soft Computing*, 14:554–562, 2014.
- [LCCC10] Wei Liu, Sanjay Chawla, David A Cieslak, and Nitesh V Chawla. A robust decision tree algorithm for imbalanced data sets. In *Proceedings of the 2010 SIAM International Conference on Data Mining*, pages 766–777. SIAM, 2010.
- [LMKA16] J. Larson, S. Mattu, L. Kirchner, and J. Angwin. How we analyzed the COMPAS recidivism algorithm. *ProPublica*, 2016.
- [Loh14] Wei-Yin Loh. Fifty years of classification and regression trees. *International Statistical Review*, 82(3):329–348, 2014.
- [LR76] Hyafil Laurent and Ronald L Rivest. Constructing optimal binary decision trees is np-complete. *Information processing letters*, 5(1):15–17, 1976.
- [MGKS18] Matt Menickelly, Oktay Günlük, Jayant Kalagnanam, and Katya Scheinberg. Optimal decision trees for categorical data via integer programming. *Preprint at arXiv:1612.03225*, January 2018.
- [MM73] W. S. Meisel and D. Michalopoulos. A partitioning algorithm with application in pattern classification and the optimization of decision tree. *IEEE Trans. Comput., C-22*, pages 93–103, 1973.
- [MS63] James N Morgan and John A Sonquist. Problems in the analysis of survey data, and a proposal. *Journal of the American statistical association*, 58(302):415–434, 1963.

- [NCLC12] Ye Nan, Kian Ming Chai, Wee Sun Lee, and Hai Leong Chieu. Optimizing f-measure: A tale of two approaches. *arXiv preprint arXiv:1206.4625*, 2012.
- [NF07] Siegfried Nijssen and Elisa Fromont. Mining optimal decision trees from itemset lattices. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 530–539. ACM, 2007.
- [NIPMS18] Nina Narodytska, Alexey Ignatiev, Filipe Pereira, and Joao Marques-Silva. Learning optimal decision trees with SAT. In *Proc. International Joint Conferences on Artificial Intelligence (IJCAI)*, pages 1362–1368, 2018.
- [NS⁺20] Siegfried Nijssen, Pierre Schaus, et al. Learning optimal decision trees using caching branch-and-bound search. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.
- [PM77] Harold J. Payne and William S. Meisel. An algorithm for constructing optimal binary decision trees. *IEEE Transactions on Computers*, C-26(9):905–916, 1977.
- [Qui86] J. Ross Quinlan. Induction of decision trees. *Machine learning*, 1(1):81–106, 1986.
- [Qui93] J. R. Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, 1993.
- [Qui14] J Ross Quinlan. *C4. 5: programs for machine learning*. Elsevier, 2014.
- [RE18] Cynthia Rudin and Seyda Ertekin. Learning customized and optimized lists of rules with mathematical programming. *Mathematical Programming C (Computation)*, 10:659–702, 2018.
- [RW18a] Cynthia Rudin and Yining Wang. Direct learning to rank and rerank. In *International Conference on Artificial Intelligence and Statistics*, pages 775–783, 2018.
- [RW18b] Cynthia Rudin and Yining Wang. Direct learning to rank and rerank. *arXiv preprint arXiv:1802.07400*, 2018.
- [VBSMVB19] Matheus Guedes Vilas Boas, Haroldo Gambini Santos, Luiz Henrique de Campos Merschmann, and Greet Vanden Berghe. Optimal decision trees for the algorithm selection problem: integer programming based approaches. *International Transactions in Operational Research*, 2019.
- [VC07] Florian Verhein and Sanjay Chawla. Using significant, positively associated and relatively class correlated rules for associative classification of imbalanced datasets. In *Seventh IEEE International Conference on Data Mining (ICDM 2007)*, pages 679–684. IEEE, 2007.
- [VZ19] Sicco Verwer and Yingqian Zhang. Learning optimal classification trees using a binary linear program formulation. In *33rd AAAI Conference on Artificial Intelligence*, 2019.
- [WRDV⁺17] T. Wang, C. Rudin, F. Doshi-Velez, Y. Liu, E. Klampfl, and P. MacNeille. A Bayesian framework for learning rule sets for interpretable classification. *Journal of Machine Learning Research*, 18(70):1–37, 2017.