

# MAPTree: Beating “Optimal” Decision Trees with Bayesian Decision Trees

Colin Sullivan\*, Mo Tiwari\*, Sebastian Thrun

Department of Computer Science, Stanford University  
colins26@stanford.edu, motiwari@stanford.edu, thrun@stanford.edu

## Abstract

Decision trees remain one of the most popular machine learning models today, largely due to their out-of-the-box performance and interpretability. In this work, we present a Bayesian approach to decision tree induction via maximum a posteriori inference of a posterior distribution over trees. We first demonstrate a connection between maximum a posteriori inference of decision trees and AND/OR search. Using this connection, we propose an AND/OR search algorithm, dubbed MAPTree, which is able to recover the maximum a posteriori tree. Lastly, we demonstrate the empirical performance of the maximum a posteriori tree both on synthetic data and in real world settings. On 16 real world datasets, MAPTree either outperforms baselines or demonstrates comparable performance but with much smaller trees. On a synthetic dataset, MAPTree also demonstrates greater robustness to noise and better generalization than existing approaches. Finally, MAPTree recovers the maximum a posteriori tree faster than existing sampling approaches and, in contrast with those algorithms, is able to provide a certificate of optimality. The code for our experiments is available at <https://github.com/ThrunGroup/maptree>.

## Introduction

Decision trees are amongst the most widely used machine learning models today due to their empirical performance, generality, and interpretability. A decision tree is a binary tree in which each internal node corresponds to an *if/then/else* comparison on a feature value; a label for a datapoint is produced by determining the corresponding leaf node into which it falls. The predicted label is usually the majority vote (respectively, mean) of the label of training datapoints at the leaf node in classification (respectively, regression).

Despite recent advances in neural networks, decision trees remain a popular choice amongst machine learning practitioners. Decision trees form the backbone of more complex ensemble models such as Random Forest (Breiman 2001) and XGBoost (Chen and Guestrin 2016), which have been the leading models in many machine learning competitions and often outperform neural networks on tabular data (Grin-

sztajn, Oyallon, and Varoquaux 2022). Decision trees naturally work with complex data where the features can be of mixed data types, e.g., binary, categorical, or continuous. Furthermore, decision trees are highly interpretable and the prediction-generating process can be inspected, which can be a necessity in domains such as law and healthcare. Furthermore, inference in decision trees is highly efficient as it relies only on efficient feature value comparisons. Given decision trees’ popularity, an improvement upon existing decision tree approaches would have widespread impact.

**Contributions:** In this work, we:

- Formalize a connection between maximum a posteriori inference of Bayesian Classification and Regression Trees (BCART) and AND/OR search problems,
- Propose an algorithm, dubbed MAPTree, for search on AND/OR graphs that recovers the maximum a posteriori tree of the BCART posterior over decision trees,
- Demonstrate that MAPTree is significantly faster than previous sampling-based approaches,
- Demonstrate that the tree recovered by MAPTree either a) outperforms current state-of-the-art algorithms in performance, or b) demonstrates comparable performance but with smaller trees, and
- Provide a heavily optimized C++ implementation that is also callable from Python for practitioners.

## Related Work

In this work, we focus on the construction of individual decision trees. We compare our proposed algorithm with four main classes of prior algorithms: greedy algorithms, “Optimal” Decision Trees (ODTs), “Optimal” Sparse Decision Trees (OSDTs), and sampling-based approaches.

The most popular method for constructing decision trees is a greedy approach that recursively splits nodes based on a heuristic such as Gini impurity or entropy (in classification) or mean-squared error (in regression) (Quinlan 1986). However, individual decision trees constructed in this manner often overfit the training data; ensemble methods such as Random Forest and XGBoost attempt to ameliorate overfitting but are significantly more complex than a single decision tree (Breiman 2001; Chen and Guestrin 2016).

So-called “optimal” decision trees reformulate the problem of decision tree induction as a global optimization prob-

\*These authors contributed equally.

lem, i.e., to find the tree that maximizes global objective function, such as training accuracy, of a given maximum depth (Bertsimas and Dunn 2017; Verwer and Zhang 2019; Verhaeghe et al. 2020; Aglin, Nijssen, and Schaus 2020). Though this problem is NP-Hard in general (Hyafil and Rivest 1976), existing approaches can find the global optimum of shallow trees (depth  $\leq 5$ ) on medium-sized datasets with thousands of datapoints and tens of features. The original ODT approaches were based on mixed integer programming or binary linear program formulations (Verhaeghe et al. 2020; Nijssen and Fromont 2007; Bertsimas and Dunn 2017; Verwer and Zhang 2019). Other work attempts to improve upon these methods using caching branch-and-bound search (Aglin, Nijssen, and Schaus 2020), constraint programming with AND/OR search (Verhaeghe et al. 2020), or dynamic programming with bounds (van der Linden, de Weerd, and Demirović 2022). ODTs have been shown to outperform their greedily constructed counterparts with smaller trees (Verhaeghe et al. 2020; Verwer and Zhang 2019) but still suffer from several drawbacks. First, choosing the maximum depth hyperparameter is nontrivial, even with cross-validation, and the maximum depth cannot be set too large as the runtime of these algorithms scales exponentially with depth. Furthermore, ODTs often suffer from overfitting, especially when the maximum depth is set too large. Amongst ODT approaches, Verhaeghe et al. (2020) formulates the search for an optimal decision tree in terms of an AND/OR graph and is most similar to ours, but still suffers from the aforementioned drawbacks. Additionally, many ODT algorithms exhibit poor anytime behavior (Kiossou et al. 2023). Optimal sparse decision trees attempt to adapt ODT approaches to train smaller and sparser trees by incorporating a sparsity penalty in their objectives. As a result, OSDTs are smaller and less prone to overfitting than ODTs (Hu, Rudin, and Seltzer 2019; Lin et al. 2020). These approaches, however, often underfit the data (Hu, Rudin, and Seltzer 2019; Lin et al. 2020).

Another class of approaches, called Bayesian Classification and Regression Trees (BCART), introduce a posterior over tree structures given the data and sample trees from this posterior. Initially, BCART methods were observed to generate better trees than greedy methods (Denison, Mallick, and Smith 1998). Many variations to the BCART methodology were developed using sampling methods based on Markov-Chain Monte Carlo (MCMC), such as Metropolis-Hastings (Pratola 2016) and others (Geels, Pratola, and Herbei 2022; Lakshminarayanan, Roy, and Teh 2013). These methods, however, often suffer from exponentially long mixing times in practice and become stuck in local minima (Kim and Rockova 2023). In one study, the posterior over trees was represented as a lattice over itemsets (Nijssen 2008). This approach discovered the maximum a posteriori tree within the hypothesis space of decision trees. However, this approach required enumerating and storing the entire space of decision trees and therefore placed stringent constraints on the search space of possible trees, based on leaf node support and maximum depth. Our method utilises the same posterior over tree structures introduced by BCART. In contrast with prior work, however, we are able to recover

the provably maximum a posteriori tree from this posterior in the unconstrained setting.

## Preliminaries and Notation

In this paper, we focus on the binary classification task, though our techniques extend to multi-class classification and regression. We also focus on binary datasets, as is common in the decision tree literature (Verhaeghe et al. 2020; Nijssen 2008; Nijssen and Fromont 2007) since many datasets can be binarized via bucketing, one-hot encoding, and other techniques.

**General notation:** We assume we are given a binary dataset  $\mathcal{X} \in \{0, 1\}^{N \times F}$  with  $N$  samples,  $F$  features, and associated binary labels  $\mathcal{Y} \in \{0, 1\}^N$ . We let  $[u] := \{1, \dots, u\}$ ,  $\mathcal{I} \subseteq [N]$  the indices of a subsample of the dataset, and  $(x_i, y_i)$  denote the  $i$ th sample and its label. We define  $\mathcal{X}|_{\mathcal{I}} := \{x_i : i \in \mathcal{I}\} \subset \mathcal{X}$ ,  $\mathcal{Y}|_{\mathcal{I}} := \{y_i : i \in \mathcal{I}\} \subset \mathcal{Y}$ , and  $\mathcal{I}|_{f=k} := \{i : i \in \mathcal{I} \text{ and } (x_i)_f = k\}$ , for  $k \in \{0, 1\}$ . Finally, we let  $c^k(\mathcal{I})$  be the count of points in  $\mathcal{I}$  with label  $k \in \{0, 1\}$ , i.e.,  $c^k(\mathcal{I}) = |\{i : i \in \mathcal{I} \text{ and } y_i = k\}|$  and  $\mathcal{V}(\mathcal{I})$  be the set of nontrivial features splits of the samples in  $\mathcal{I}$ , i.e., the set of features such that neither  $\mathcal{I}|_{f=0}$  nor  $\mathcal{I}|_{f=1}$  is nonempty.

**Tree notation:** We let  $T = \{n_1, n_2, \dots, n_{M+L}\}$  be a binary classification tree represented as a collection of its nodes and use  $n$  to refer to a node in  $T$ ,  $m$  to refer to one of the  $M$  internal nodes in  $T$ , and  $l$  to refer to one of the  $L$  leaf nodes in  $T$ . Furthermore, we use  $\mathcal{I}(n)$  to denote the indices of the samples in  $\mathcal{X}$  that reach node  $n$  in  $T$ , namely  $\{i : x_i \in \text{space}(n)\}$ , where  $\text{space}(n)$  is the subset of feature space that reaches node  $n$  in  $T$ . We also use  $c_l^k$  to denote the count of points assigned to leaf  $l$  with label  $k \in \{0, 1\}$  (i.e.,  $c_l^k = c^k(\mathcal{I}(l))$ ),  $T_{\text{internal}} = \{m_1, m_2, \dots, m_M\} \subset T$  to denote the set of internal nodes in tree  $T$ , and  $T_{\text{leaves}} = \{l_1, l_2, \dots, l_L\} \subset T$  is the set of all leaf nodes in tree  $T$ . Finally, we use  $d(n)$  to denote the depth of node  $n$  in  $T$ .

## AND/OR Graph Search

We briefly recapitulate the concept of AND/OR graphs and a search algorithm for AND/OR graphs, AO\*. AND/OR graph search can be viewed as a generalization of the shortest path problem that allows nodes consisting of independent subproblems to be decomposed and solved separately. Thus, a solution of an AND/OR graph is not a path but rather a subgraph  $\mathcal{S}$  with cost, denoted  $\text{cost}(\mathcal{S})$ , equal to the sum across the costs of its edges. AND/OR graphs contain two types of nodes: terminal nodes and nonterminal nodes. Non-terminal nodes can be further subdivided into AND nodes and OR nodes, with a special OR node designated as the *root* or *start* node  $r$ . For a given AND/OR graph  $\mathcal{G}$ , a *solution graph*  $\mathcal{S}$  on an AND/OR graph is a connected subset of nodes of  $\mathcal{G}$  in which:

1.  $r \in \mathcal{S}$ ,
2. for every AND node  $a \in \mathcal{S}$ , all the immediate children of  $a$  are also in  $\mathcal{S}$ , and
3. for every non-terminal OR node  $o \in \mathcal{S}$  exactly one of  $o$ 's children is also in  $\mathcal{S}$ .




Figure 1: An example (general) AND/OR graph, with AND nodes drawn as squares, and OR nodes drawn as solid circles, and terminal nodes drawn as dashed circles. The minimal cost solution is highlighted in red and has cost  $0 + 0 + 3 + 4 + 1 + 2 = 10$ . This diagram demonstrates an AND/OR graph where the root node  $r$  is an AND node; in MAPTree, the root node is an OR node.

Intuitively, the children of an AND node  $a$  represent subtasks that must all be solved for  $a$  to be satisfied (e.g., simultaneous prerequisites), and the children of an OR node  $o$  represent mutually exclusive satisfying choices.

One of the most popular AND/OR graph search algorithms is AO\* (Mahanti and Bagchi 1985, 1983). The AO\* algorithm explores potential paths in an AND/OR graph in a best-first fashion, guided by a heuristic. When a new node is explored, its children are revealed and the cost for that node and all of its ancestors is updated; the search then continues. This process is repeated until the the root node is marked as solved, indicating that no immediately accessible nodes could lead to an increase in heuristic value. The AO\* algorithm is guaranteed to find the minimal cost solution if the heuristic is *admissible*, i.e., the heuristic estimate of cost is always less than or equal to the actual cost of a node. For more details on the AO\* algorithm, we refer the reader to (Mahanti and Bagchi 1985). An example AND/OR graph is given in Figure 1 with its minimal cost solution shown in red.

**Additional AND/OR graph notation:** In addition to the notation defined above, we use  $t$  to refer to a terminal node. When searching over an AND/OR graph, we use  $\mathcal{G}$  to refer to the implicit (entire) AND/OR graph and  $\mathcal{G}' \subset \mathcal{G}$  to explicit (explored) AND/OR graph, as in prior work.

### Bayesian Classification and Regression Trees (BCART)

Bayesian Decision Trees are a family of statistical models of decision trees introduced in Chipman, George, and McCulloch (1998) and Denison, Mallick, and Smith (1998). A Bayesian Decision Tree (BDT) is a pair  $(T, \Theta)$  where  $T$  is a tree and  $\Theta = (\theta_{l_1}, \theta_{l_2}, \dots, \theta_{l_L})$  parameterizes the independent probability distributions over labels in the leaf nodes of tree  $T$ . We are interested in the binary classification setting, where each  $\theta_l$  parameterizes a Bernoulli distribution  $\text{Ber}(\theta_l)$  with  $\theta_l \in [0, 1]$ . We denote by  $\text{Beta}(\rho^1, \rho^0)$  the Beta distribution with parameters  $\rho^1, \rho^0 \in \mathbb{R}^+$  and by  $B(c^1, c^0)$  the

Beta function.

We note that a BDT's tree  $T$  partitions the data such that the sample subsets  $\mathcal{I}(l_1), \mathcal{I}(l_2), \dots, \mathcal{I}(l_L)$  fall into leaves  $l_1, l_2, \dots, l_L$ . Furthermore, a BDT defines a probability distribution over the respective labels occurring in their leaves: each label in leaf  $l$  is sampled from  $\text{Ber}(\theta_l)$ . Every BDT therefore induces a likelihood function, given in Theorem 1.

**Theorem 1.** *The likelihood of a BDT  $(T, \Theta)$  generating labels  $\mathcal{Y}$  given features  $\mathcal{X}$  is*

$$P(\mathcal{Y}|\mathcal{X}, T, \Theta) = \prod_{l \in \mathcal{T}_{\text{leaves}}} \prod_{i \in \mathcal{I}(l)} \theta_l^{y_i} (1 - \theta_l)^{1 - y_i} \quad (1)$$

$$= \prod_{l \in \mathcal{T}_{\text{leaves}}} \theta_l^{c_l^1} (1 - \theta_l)^{c_l^0} \quad (2)$$

The specific formulation of BCART also assumes a prior distribution over  $\Theta$ , i.e., that  $\theta \sim \text{Beta}(\rho^1, \rho^0)$  for each  $\theta \in \Theta$ . With this assumption, we can derive the likelihood function  $P(\mathcal{Y}|\mathcal{X}, T)$ ; see Theorem 2.

**Theorem 2.** *Assume that each  $\theta \sim \text{Beta}(\rho^1, \rho^0)$  for each  $\theta \in \Theta$ . Then the likelihood of a tree  $T$  generating labels  $\mathcal{Y}$  given features  $\mathcal{X}$  is*

$$P(\mathcal{Y}|\mathcal{X}, T) = \prod_{l \in \mathcal{T}_{\text{leaves}}} \frac{B(c_l^1 + \rho^1, c_l^0 + \rho^0)}{B(\rho^1, \rho^0)} \quad (3)$$

Theorems 1 and 2 are proven in the appendices; we note they have been observed in different forms in prior work (Chipman, George, and McCulloch 1998).

For notational convenience, we define a leaf count likelihood function  $\ell_{\text{leaf}}(c^1, c^0)$  for integers  $c^1$  and  $c^0$ :

$$\ell_{\text{leaf}}(c^1, c^0) := \frac{B(c^1 + \rho^1, c^0 + \rho^0)}{B(\rho^1, \rho^0)} \quad (4)$$

and we can rewrite Equation 3 as

$$P(\mathcal{Y}|\mathcal{X}, T) = \prod_{l \in \mathcal{T}_{\text{leaves}}} \ell_{\text{leaf}}(c_l^1, c_l^0) \quad (5)$$

In this work, we utilize the original prior over trees from (Chipman, George, and McCulloch 1998), given in Definition 3.

**Definition 3.** The original BCART prior distribution over trees is

$$P(T|\mathcal{X}) = \left( \prod_{l \in \mathcal{T}_{\text{leaves}}} p_{\text{leaf}}(d(l), \mathcal{I}(l)) \right) \times \left( \prod_{m \in \mathcal{T}_{\text{internal}}} p_{\text{inner}}(d(m), \mathcal{I}(m)) \right)$$

where

$$p_{\text{leaf}}(d, \mathcal{I}) = \begin{cases} 1, & \mathcal{V}(\mathcal{I}) = \emptyset \\ 1 - p_{\text{split}}(d), & \mathcal{V}(\mathcal{I}) \neq \emptyset \end{cases} \quad (6)$$




Figure 2: Example of the defined BCART AND/OR graph  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$ . OR nodes are represented as circles with solid borders, terminal nodes as circles with dashed borders, and AND nodes as squares. In this dataset, two feature splits are possible at the root node ( $f_0$  and  $f_1$ ) and no further splits are possible at deeper nodes. The best solution on this AND/OR graph is highlighted in red and corresponds with a stump which splits the root node, corresponding to the entire dataset, on feature  $f_1$ .

$$p_{\text{inner}}(d, \mathcal{I}) = \begin{cases} 0, & \mathcal{V}(\mathcal{I}) = \emptyset \\ p_{\text{split}}(d)/|\mathcal{V}(\mathcal{I})|, & \mathcal{V}(\mathcal{I}) \neq \emptyset \end{cases} \quad (7)$$

and

$$p_{\text{split}}(d) = \alpha(1+d)^{-\beta} \quad (8)$$

Intuitively,  $p_{\text{split}}(d)$  is the prior probability of any node splitting and is allocating equally amongst valid splits. This choice of prior,  $P(T|\mathcal{X})$ , combined with the likelihood function in Equation 5 induces the posterior distribution over trees  $P(T|\mathcal{Y}, \mathcal{X})$ :

$$P(T|\mathcal{Y}, \mathcal{X}) \propto P(\mathcal{Y}|\mathcal{X}, T)P(T|\mathcal{X}) \quad (9)$$

Throughout our analysis, we treat the dataset  $(\mathcal{X}, \mathcal{Y})$  as fixed.

### Connecting BCART with AND/OR Graphs

Given a dataset  $(\mathcal{X}, \mathcal{Y})$ , we will now construct a special AND/OR graph  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$ . We will then show that a minimal cost solution graph on  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$  corresponds directly with the maximum a posteriori tree given our choice of prior distributions  $P(T|\mathcal{X})$  and  $P(\Theta)$ . Using this construction, the problem of finding the maximum a posteriori tree of our posterior is reduced to that of finding the minimum cost solution graph on  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$ .

**Definition 4** (BCART AND/OR graph  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$ ). Given a dataset  $(\mathcal{X}, \mathcal{Y})$ , construct the AND/OR graph  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$  as follows:

1. For every possible subset  $\mathcal{I} \subset [N]$  and depth  $d \in \{0, \dots, F\}$ , create an OR node  $o_{\mathcal{I},d}$ .
2. For every OR node  $o_{\mathcal{I},d}$  created in Step 1, create a terminal node  $t_{\mathcal{I},d}$  and draw an edge from  $o_{\mathcal{I},d}$  to  $t_{\mathcal{I},d}$  with cost  $\text{cost}(o_{\mathcal{I},d}, t_{\mathcal{I},d}) = -\log p_{\text{leaf}}(d, \mathcal{I}) - \log \ell_{\text{leaf}}(c^1(\mathcal{I}), c^0(\mathcal{I}))$ .




Figure 3: Example map between an example solution of the AND/OR graph  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$  depicted in Figure 2 and its corresponding binary classification tree. We see that the resulting tree is a stump which splits on feature  $f_1$  at the root.

3. For every OR node  $o_{\mathcal{I},d}$  created in Step 1, create  $F$  AND nodes  $a_{\mathcal{I},d,1}, \dots, a_{\mathcal{I},d,F}$  and draw an edge from  $o_{\mathcal{I},d}$  to each  $a_{\mathcal{I},d,f}$  with cost  $\text{cost}(o_{\mathcal{I},d}, a_{\mathcal{I},d,f}) = -\log p_{\text{inner}}(d)$ .
4. For every pair  $a_{\mathcal{I},d,f}$  and  $o_{\mathcal{I}',d+1}$  where  $\mathcal{I}|_{f=k} = \mathcal{I}'$  for some  $f \in [F]$  and  $k \in \{0, 1\}$ , draw an edge from  $a_{\mathcal{I},d,f}$  to  $o_{\mathcal{I}',d+1}$  with cost  $\text{cost}(a_{\mathcal{I},d,f}, o_{\mathcal{I}',d+1}) = 0$ .
5. Let  $o_{[n],0}$ , the OR node representing all sample indices, be the unique root node  $r$  of  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$ .
6. Remove all OR nodes representing empty subsets and their neighbors.
7. Remove all nodes not connected to the root node  $r$ .

We note that  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$  contains  $F \times 2^n$  OR Nodes,  $F \times 2^n$  terminal nodes (one for each OR Node), and  $F^2 \times 2^n$  AND nodes ( $F$  for each OR Node) and so is finite.

Intuitively, each OR node  $o_{\mathcal{I},d}$  in  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$  corresponds with the subproblem of discovering a maximum a posteriori subtree starting from depth  $d$  and over the subset of samples  $\mathcal{I}$  from dataset  $\mathcal{X}, \mathcal{Y}$ . Each AND node  $a_{\mathcal{I},d,f}$  then represents the same subproblem but given that a decision was already made to split on feature  $f$  at the root node of this subtree. A valid solution graph on  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$  corresponds with a binary classification tree  $T$  on the dataset  $(\mathcal{X}, \mathcal{Y})$  and the value of a solution is related to the posterior probability of  $T$  given by  $P(T|\mathcal{Y}, \mathcal{X})$ . We formalize these properties in Theorems 5 and 6.

**Theorem 5.** *Every solution graph on AND/OR graphs induces a unique binary decision tree. Furthermore, every decision tree can be represented as a unique solution graph under this correspondence. Thus, there is natural bijection between solution graphs on  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$  and binary decision trees.*

**Theorem 6.** *Under the natural bijection described in Theorem 5, given a solution graph  $\mathcal{S}$  and its corresponding tree  $T$ , we have that  $\text{cost}(\mathcal{S}) = -\log P(T, \mathcal{Y}|\mathcal{X})$ . Therefore the minimal cost solution over  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$  corresponds with a maximum a posteriori tree.*

The bijection constructed in Theorems 5 and 6 is depicted in Figure 3. Due to space constraints, we defer a formal description of this bijection to Appendix ??.

### MAPTree

Theorems 5 and 6 imply that it is sufficient to find the minimum cost solution graph on  $\mathcal{G}_{\mathcal{X},\mathcal{Y}}$  to recover the MAP

---

**Algorithm 1:** MAPTree

**Input:** Root OR Node  $r$ , cost function  $\text{cost}$ , and heuristic function  $h$  for AND/OR graph  $\mathcal{G}$   
**Output:** Solution graph  $\mathcal{S}$

- 1:  $\mathcal{G}' := \{r\}$
- 2:  $\mathcal{E} := \emptyset$
- 3:  $LB[r] := h(r)$
- 4:  $UB[r] := \infty$
- 5: **while**  $LB[r] < UB[r]$  and time remaining **do**
- 6:    $o := \text{findNodeToExpand}(r, \text{cost}, \mathcal{E}, LB, UB)$
- 7:   Let  $t$  be the terminal node child of  $o$
- 8:    $\mathcal{E} := \mathcal{E} \cup \{o\}$
- 9:    $\mathcal{G}' := \mathcal{G}' \cup \{t\}$
- 10:   Let  $\{a_1, \dots, a_F\}$  be the AND node children of  $o$
- 11:   **for all**  $a_f \in \{a_1, \dots, a_F\}$  **do**
- 12:     Let  $\{o_{f=0}, o_{f=1}\}$  be the OR node children of  $a_f$
- 13:      $LB[o_{f=0}] := h(o_{f=0})$
- 14:      $LB[o_{f=1}] := h(o_{f=1})$
- 15:      $v_{f=0}^{(lb)} = \text{cost}(a_f, o_{f=0}) + h(o_{f=0})$
- 16:      $v_{f=1}^{(lb)} = \text{cost}(a_f, o_{f=1}) + h(o_{f=1})$
- 17:      $LB[a_f] := v_{f=0} + v_{f=1}$
- 18:      $\mathcal{G}' := \mathcal{G}' \cup \{a, o_{f=0}, o_{f=1}\}$
- 19:   **end for**
- 20:    $\text{updateLowerBounds}(o, \text{cost}, LB)$
- 21:    $\text{updateUpperBounds}(o, \text{cost}, UB)$
- 22: **end while**
- 23: **return**  $\text{getSolution}(r, \text{cost}, UB)$

---

tree under the BCART posterior. In this section, we introduce MAPTree, an AND/OR search algorithm that finds a minimal cost solution on  $\mathcal{G}_{\mathcal{X}, \mathcal{Y}}$ . MAPTree is shown in Algorithm 1.

A key component of MAPTree is the Perfect Split Heuristic  $h$  that guides the search, presented in Definition 7.

**Definition 7** (Perfect Split Heuristic). For OR node  $o_{\mathcal{I}, d}$  with terminal node child  $t_{\mathcal{I}, d}$ , let

$$h(o_{\mathcal{I}, d}) = -\max\{ \quad (10)$$

$$\log \ell_{\text{leaf}}(c^1(\mathcal{I}), c^0(\mathcal{I})), \quad (11)$$

$$\log p_{\text{split}}(d, \mathcal{I}) \quad (12)$$

$$+ \log \ell_{\text{leaf}}(c^1(\mathcal{I}), 0) \quad (13)$$

$$+ \log \ell_{\text{leaf}}(0, c^0(\mathcal{I})) \} \quad (14)$$

and for AND node  $a_{\mathcal{I}, d, f}$  with OR node children  $o_{\mathcal{I}|f=0, d+1}$  and  $o_{\mathcal{I}|f=1, d+1}$ , let

$$h(a_{\mathcal{I}, d, f}) = h(o_{\mathcal{I}|f=0, d+1}) + h(o_{\mathcal{I}|f=1, d+1}) \quad (15)$$

Intuitively, the Perfect Split Heuristic describes the negative log posterior probability of the best potential subtree rooted at the given OR node  $o_{\mathcal{I}, d}$ : one that perfectly classifies the data in a single additional split. The heuristic guides the search away from subproblems that are too deep or for which the labels have already been poorly divided. We prove that this heuristic is a lower bound (admissible) and consistent in later sections.

---

**Algorithm 2:** getSolution

**Input:** ORNode  $o$ , cost function  $\text{cost}$ , and upper bounds  $UB$   
**Output:** Solution graph  $\mathcal{S}$

- 1: Let  $\{a_1, \dots, a_F\}$  be the AND node children of  $o$
- 2: Let  $t$  be the terminal node child of  $o$
- 3:  $a_{f^*} := \arg \min_{c \in \{a_1, \dots, a_F\}} (\text{cost}(o, c) + UB[c])$
- 4:  $\mathcal{S} := \{o\}$
- 5: **if**  $\text{cost}(o, t) + UB[t] \leq \text{cost}(o, a_{f^*}) + UB[a_{f^*}]$  **then**
- 6:    $\mathcal{S} := \mathcal{S} \cup \{t\}$
- 7: **else**
- 8:   Let  $o_{f^*=0}, o_{f^*=1}$  be the children of  $a_{f^*}$
- 9:    $\mathcal{S} := \mathcal{S} \cup \{a_{f^*}\}$
- 10:    $\mathcal{S}_0 := \text{getSolution}(o_{f^*=0})$
- 11:    $\mathcal{S}_1 := \text{getSolution}(o_{f^*=1})$
- 12:    $\mathcal{S} := \mathcal{S} \cup \mathcal{S}_0 \cup \mathcal{S}_1$
- 13: **end if**
- 14: **return**  $\mathcal{S}$

---



---

**Algorithm 3:** findNodeToExpand

**Input:** Root node  $r$ , cost function  $\text{cost}$ , set of expanded nodes  $\mathcal{E}$ , lower bounds  $LB$  and upper bounds  $UB$   
**Output:** ORNode  $o$

- 1:  $o := r$
- 2: **while**  $o \in \mathcal{E}$  **do**
- 3:   Let  $\{a_1, \dots, a_F\}$  be the children of  $o$
- 4:    $a^* := \arg \min_{c \in \{a_1, \dots, a_F\}} (\text{cost}(o, c) + LB[c])$
- 5:   Let  $o_0^*, o_1^*$  be the children of  $a^*$
- 6:   **if**  $UB[o_0^*] - LB[o_0^*] > UB[o_1^*] - LB[o_1^*]$  **then**
- 7:      $o := o_0^*$
- 8:   **else**
- 9:      $o := o_1^*$
- 10:   **end if**
- 11: **end while**
- 12: **return**  $o$

---

**Analysis of MAPTree**

We now introduce several key properties of MAPTree. In particular, we show that (1) the Perfect Split Heuristic is consistent and therefore also admissible, (2) MAPTree finds the maximum a posteriori tree of the BCART posterior upon completion, and (3) upon early termination, MAPTree returns the minimum cost solution within the explored explicit graph  $\mathcal{G}'$ . Theorems 8 - 12 and Corollary 11 are proven in Appendix ??.

**Theorem 8** (Consistency of the Perfect Split Heuristic). *The Perfect Split Heuristic in Definition 7 is consistent, i.e., for any OR node  $o$  with children  $\{t, a_1, \dots, a_F\}$ :*

$$h(o) \leq \min_{c \in \{t, a_1, \dots, a_F\}} \text{cost}(o, c) + h(c) \quad (16)$$

and for any AND node  $a$  with children  $\{o_0, o_1\}$ :

$$h(a) \leq \sum_{c \in \{o_0, o_1\}} \text{cost}(a, c) + h(c) \quad (17)$$

**Theorem 9** (Finiteness of MAPTree). *Algorithm 1 always terminates.*

---

**Algorithm 4:** updateLowerBounds

---

**Input:** ORNode  $l$ , cost function  $\text{cost}$ , lower bounds  $LB$ 

```

1:  $\mathcal{V} = \{l\}$ 
2: while  $|\mathcal{V}| > 0$  do
3:   Remove a node  $o$  from  $\mathcal{V}$  with maximal depth
4:   Let  $\{a_1, \dots, a_F\}$  be the AND node children of  $o$ 
5:   Let  $t$  be the terminal node child of  $o$ 
6:    $v_{\text{split}}^{(lb)} = \min_{c \in \{a_1, \dots, a_F\}} (\text{cost}(o, c) + LB[c])$ 
7:    $v^{(lb)} = \min\{v_{\text{split}}^{(lb)}, \text{cost}(o, t)\}$ 
8:   if  $v^{(lb)} > LB[o]$  then
9:      $LB[o] := v^{(lb)}$ 
10:    Add all parents of  $o$  to  $\mathcal{V}$ 
11:   end if
12: end while

```

---



---

**Algorithm 5:** updateUpperBounds

---

**Input:** ORNode  $l$ , cost function  $\text{cost}$ , upper bounds  $UB$ 

```

1:  $\mathcal{V} = \{l\}$ 
2: while  $|\mathcal{V}| > 0$  do
3:   Remove a node  $o$  from  $\mathcal{V}$  with maximal depth
4:   Let  $\{a_1, \dots, a_F\}$  be the AND node children of  $o$ 
5:   Let  $t$  be the terminal node child of  $o$ 
6:    $v_{\text{split}}^{(ub)} = \min_{c \in \{a_1, \dots, a_F\}} (\text{cost}(o, c) + UB[c])$ 
7:    $v^{(ub)} = \min\{v_{\text{split}}^{(ub)}, \text{cost}(o, t)\}$ 
8:   if  $v^{(ub)} < UB[o]$  then
9:      $UB[o] := v^{(ub)}$ 
10:    Add all parents of  $o$  to  $\mathcal{V}$ 
11:   end if
12: end while

```

---

**Theorem 10** (Correctness of MAPTree). *When Algorithm 1 does not terminate early due to the time remaining condition, it always outputs a minimal cost solution on  $\mathcal{G}_{\mathcal{X}, \mathcal{Y}}$  upon completion.*

**Corollary 11.** *Consider the tree induced by the output of Algorithm 1 under the natural bijection described in Section . By Theorems 5 and 6, this tree is the maximum a posteriori tree  $\arg \max_T P(T|\mathcal{X}, \mathcal{Y})$ .*

**Theorem 12** (Anytime optimality of MAPTree). *Upon early termination, Algorithm 1 outputs the minimal cost solution across the explicit subgraph  $\mathcal{G}'$  of already explored nodes.*

## Experiments

We evaluate the performance of MAPTree in multiple settings. In all experiments in this section, we set  $\alpha = 0.95$  and  $\beta = 0.5$ . We find that our results are not highly dependent on the choices of  $\alpha$  and  $B$ ; see Appendix ??.

In the first setting, we compare the efficiency of MAPTree to the Sequential Monte Carlo (SMC) and Markov-Chain Monte Carlo (MCMC) baselines from Lakshminarayanan, Roy, and Teh (2013) and Chipman, George, and McCul-

loch (1998), respectively. In the second setting, we create a synthetic dataset in which the true labels are generated by a randomly generated tree and measure generalization performance with respect to training dataset size. In the third setting, we measure the generalization accuracy, log likelihood, and tree size of models generated by MAPTree and baseline algorithms across all 16 datasets from the CP4IM dataset repository (Guns, Nijssen, and De Raedt 2011).

## Speed Comparisons against MCMC and SMC

We first compare the performance of MAPTree with the SMC and MCMC baselines from Lakshminarayanan, Roy, and Teh (2013) and Chipman, George, and McCulloch (1998), respectively, on all 16 binary classification datasets from the CP4IM dataset repository (Guns, Nijssen, and De Raedt 2011). We note that all three methods, given infinite exploration time, should recover the maximum a posteriori tree from the BCART posterior. However, it has been observed that the mixing times for Markov-Chain-based methods, such as the MCMC and SMC baselines, is exponential in the depth of the data-generating tree (Kim and Rockova 2023). Furthermore, the SMC and MCMC methods are unable to determine when they have converged, nor can they provide a certificate of optimality upon convergence.

In our experiments, we modify the hyperparameters of each algorithm and measure the training time and log posterior of the data under the output tree (Figure 4). In 12 of the 16 datasets in Figure ??, MAPTree outperforms SMC and MCMC and is able to find trees with higher log posterior faster than the baseline algorithms. Furthermore, in 5 of the 16 datasets, MAPTree converges to the provably optimal tree, i.e., the maximum a posteriori tree of the BCART posterior.

## Fitting a Synthetic Dataset

We measure the generalization performance of MAPTree and various other baseline algorithms as a function of training dataset size on tree-generated data.

**Synthetic Data:** We construct a synthetic dataset where labels are generated by a randomly generated tree. We first construct a random binary tree structure as specified in Devroye and Kruszewski (1995) via recursive random divisions of the available internal nodes to the left or right subtree. Next, features are selected for each internal node uniformly at random such that no internal node splits on the same feature as its ancestors. Lastly, labels are assigned to the leaf nodes in alternating fashion so as to avoid compression of the underlying tree structure. Individual datapoints with 40 features are then sampled with each feature drawn i.i.d. from  $\text{Ber}(1/2)$ , and their labels are determined by following the generated tree to a leaf node. We repeat this process 20 times, generating 20 datasets for 20 random trees. We also randomly flip  $\epsilon$  of the training data labels, with  $\epsilon$  ranging from 0 to 0.25 to simulate label noise.

In our experiments, MAPTree generates trees which outperform both the greedy, top-down approaches and ODT methods in test accuracy for various training dataset sizes and values of label corruption proportion  $\epsilon$ ; the results

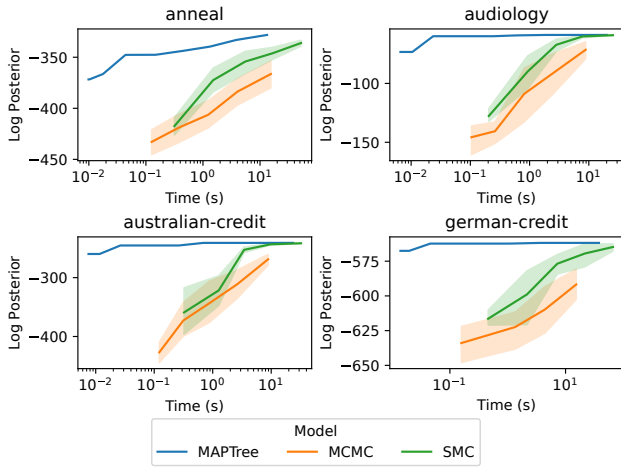


Figure 4: Comparison of MAPTree, SMC, and MCMC on 4 datasets (results for an additional 12 datasets are presented in the appendix). Curves are created by modifying the hyperparameters for each algorithm and measuring training time and log posterior of the data under the tree. Higher and further left is better, i.e., better log posteriors in less time. In 12 of the 16 datasets, MAPTree outperforms SMC and MCMC and is able to find trees with higher log posterior faster than the baseline algorithms. Furthermore, in 5 of the 16 datasets, MAPTree converges to the provably optimal tree, i.e., the MAP tree. 95% confidence intervals are derived by bootstrapping the results of 10 random seeds and time is averaged across the 10 seeds.

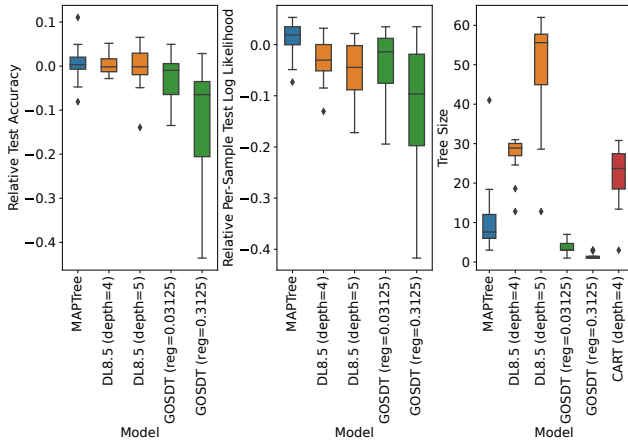


Figure 5: Box-and-whisker plot of stratified 10-fold relative test accuracy, relative per-sample log likelihood, and size of the output tree for MAPTree and various baseline algorithms for each of the 16 CP4IM datasets. Test accuracy and log likelihood are relative to that of CART (max depth 4). Higher is better for the left and center plots and lower is better for the right plot. Against all baseline algorithms, MAPTree either a) performs better in test accuracy or log likelihood, or b) performs comparably in test accuracy and log likelihood but produces smaller trees.

are presented in Figure ?? in Appendix ?? due to space constraints. We note that though some baseline algorithms demonstrate comparable performance at a single noise level, no baseline algorithm demonstrates test accuracy comparable to MAPTree across all noise levels. We also emphasize that MAPTree requires no hyperparameter tuning, whereas we experimented with various values of hyperparameters for the baseline algorithms in which performance was highly dependent on hyperparameter values (e.g., DL8.5 and GOSDT); see Appendix ??.

### Accuracy, Likelihood, and Size Comparisons on Real World Benchmarks

We also compare the accuracy, test log likelihood, and sizes of trees generated by MAPTree and baseline algorithms on 16 real world binary classification datasets from the CP4IM dataset repository (Guns, Nijssen, and De Raedt 2011) (Figure 5). Against all baseline algorithms, MAPTree either a) performs better in test accuracy or log likelihood, or b) performs comparably in test accuracy and log likelihood but produces smaller trees.

### Discussion and Conclusions

We presented MAPTree, an algorithm which provably finds the maximum a posteriori tree of the BCART posterior for a given dataset. Our algorithm is inspired by best-first-search algorithms over AND/OR graphs and the observation that the search problem for trees can be framed as a search problem over an appropriately constructed AND/OR graph.

MAPTree outperforms thematically similar approaches such as SMC- and MCMC-based algorithms, finding higher log-posterior trees faster, and is able to determine when it has converged to the maximum a posteriori tree, unlike prior work. MAPTree also outperforms greedy, ODT, and ODST construction methods in test accuracy on the synthetic dataset constructed in Section . Furthermore, on many real world benchmark datasets, MAPTree either a) demonstrates better generalization performance, or b) demonstrates comparable generalization performance but with smaller trees.

A limitation of MAPTree is that it constructs a potentially large AND/OR graph, which consumes a significant amount of memory. We leave optimizations that may permit MAPTree to run on huge datasets to future work. Nonetheless, with the optimizations presented in Section , we find that MAPTree was performant enough to run on the CP4IM benchmark datasets used in evaluation of previous ODT benchmarks.

### Acknowledgements

We would like to thank the anonymous reviewers and Area Chair for their reviews and helpful feedback.

M. T. was funded by a J.P. Morgan AI Fellowship, a Stanford Interdisciplinary Graduate Fellowship, a Stanford Data Science Scholarship, and an Oak Ridge Institute for Science and Engineering Fellowship.

## References

- Aglin, G.; Nijssen, S.; and Schaus, P. 2020. Learning Optimal Decision Trees Using Caching Branch-and-Bound Search. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04): 3146–3153. Number: 04.
- Bertsimas, D.; and Dunn, J. 2017. Optimal classification trees. *Machine Learning*, 106.
- Breiman, L. 2001. Random Forests. *Machine Learning*, 45(1): 5–32.
- Chen, T.; and Guestrin, C. 2016. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, KDD '16, 785–794. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-4232-2.
- Chipman, H. A.; George, E. I.; and McCulloch, R. E. 1998. Bayesian CART Model Search. *Journal of the American Statistical Association*, 93(443): 935–948. Publisher: Taylor & Francis.
- Denison, D. G. T.; Mallick, B. K.; and Smith, A. F. M. 1998. A Bayesian CART Algorithm. *Biometrika*, 85(3): 363–377.
- Devroye, L.; and Kruszewski, P. 1995. The Botanical Beauty of Random Binary Trees. In *International Symposium Graph Drawing and Network Visualization*.
- Geels, V.; Pratola, M. T.; and Herbei, R. 2022. The Taxicab Sampler: MCMC for Discrete Spaces with Application to Tree Models. *Journal of Statistical Computation and Simulation*, 1–22. Publisher: Taylor & Francis.
- Grinsztajn, L.; Oyallon, E.; and Varoquaux, G. 2022. Why do Tree-Based Models Still Outperform Deep Learning on Tabular Data?
- Guns, T.; Nijssen, S.; and De Raedt, L. 2011. Itemset mining: A constraint programming perspective. *Artificial Intelligence*, 175(12): 1951–1983.
- Hu, X.; Rudin, C.; and Seltzer, M. 2019. Optimal Sparse Decision Trees. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Hyafil, L.; and Rivest, R. L. 1976. Constructing optimal binary decision trees is NP-complete. *Information Processing Letters*, 5(1): 15–17.
- Kim, J.; and Rockova, V. 2023. On Mixing Rates for Bayesian CART. ArXiv:2306.00126 [math, stat].
- Kiossou, H.; Schaus, P.; Nijssen, S.; and Houndji, V. R. 2023. Time Constrained DL8.5 Using Limited Discrepancy Search. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2022, Grenoble, France, September 19–23, 2022, Proceedings, Part V*, 443–459. Berlin, Heidelberg: Springer-Verlag. ISBN 978-3-031-26418-4.
- Lakshminarayanan, B.; Roy, D. M.; and Teh, Y. W. 2013. Top-down particle filtering for bayesian decision trees. In *Proceedings of the 30th international conference on international conference on machine learning - volume 28*, ICML'13, III–280–III–288. JMLR.org. Place: Atlanta, GA, USA.
- Lin, J.; Zhong, C.; Hu, D.; Rudin, C.; and Seltzer, M. 2020. Generalized and scalable optimal sparse decision trees. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *ICML'20*, 6150–6160. JMLR.org.
- Mahanti, A.; and Bagchi, A. 1983. Admissible heuristic search in and/or graphs. *Theoretical Computer Science*, 24(2): 207–219. Publisher: Elsevier.
- Mahanti, A.; and Bagchi, A. 1985. AND/OR graph heuristic search methods. *Journal of the ACM*, 32(1): 28–51.
- Nijssen, S. 2008. Bayes optimal classification for decision trees. In *Proceedings of the 25th international conference on Machine learning*, ICML '08, 696–703. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-60558-205-4.
- Nijssen, S.; and Fromont, E. 2007. Mining optimal decision trees from itemset lattices. In *Knowledge discovery and data mining*.
- Pratola, M. T. 2016. Efficient Metropolis–Hastings Proposal Mechanisms for Bayesian Regression Tree Models. *Bayesian Analysis*, 11(3): 885–911. Publisher: International Society for Bayesian Analysis.
- Quinlan, J. R. 1986. Induction of Decision Trees. *Machine Learning*, 1(1): 81–106.
- van der Linden, J.; de Weerd, M.; and Demirović, E. 2022. Fair and optimal decision trees: A dynamic programming approach. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in neural information processing systems*, volume 35, 38899–38911. Curran Associates, Inc.
- Verhaeghe, H.; Nijssen, S.; Pesant, G.; Quimper, C.-G.; and Schaus, P. 2020. Learning optimal decision trees using constraint programming. *Constraints*, 25(3): 226–250.
- Verwer, S.; and Zhang, Y. 2019. Learning optimal classification trees using a binary linear program formulation. In *Proceedings of the thirty-third AAAI conference on artificial intelligence and thirty-first innovative applications of artificial intelligence conference and ninth AAAI symposium on educational advances in artificial intelligence*, AAAI'19/IAAI'19/EAAI'19. AAAI Press. ISBN 978-1-57735-809-1. Place: Honolulu, Hawaii, USA Number of pages: 8 tex.articleno: 200.