

The Computational Curse of Big Data for Bayesian Additive Regression Trees: A Hitting Time Analysis

Yan Shuo Tan¹, Omer Ronen^{*2}, Theo Saarinen^{*2}, and Bin Yu^{2,3,4}

¹Department of Statistics and Data Science, National University of Singapore

²Department of Statistics, UC Berkeley

³Department of Electrical Engineering and Computer Sciences, UC Berkeley

⁴Center for Computational Biology, UC Berkeley

July 1, 2024

Abstract

Bayesian Additive Regression Trees (BART) is a popular Bayesian non-parametric regression model that is commonly used in causal inference and beyond. Its strong predictive performance is supported by theoretical guarantees that its posterior distribution concentrates around the true regression function at optimal rates under various data generative settings and for appropriate prior choices. In this paper, we show that the BART sampler often converges slowly, confirming empirical observations by other researchers. Assuming discrete covariates, we show that, while the BART posterior concentrates on a set comprising all optimal tree structures (smallest bias and complexity), the Markov chain’s hitting time for this set increases with n (training sample size), under several common data generative settings. As n increases, the approximate BART posterior thus becomes increasingly different from the exact posterior (for the same number of MCMC samples), contrasting with earlier concentration results on the exact posterior. This contrast is highlighted by our simulations showing worsening frequentist undercoverage for approximate posterior intervals and a growing ratio between the MSE of the approximate posterior and that obtainable by artificially improving convergence via averaging multiple sampler chains. Finally, based on our theoretical insights, possibilities are discussed to improve the BART sampler convergence performance.

1 Introduction

1.1 The Rise of BART

Decision tree models such as CART (Breiman et al., 1984) are piecewise constant regression models obtained by recursively partitioning the covariate space along coordinate axes. They and their ensembles such as Random Forests (RFs) (Breiman, 2001) and Gradient Boosted Trees (GBTs) (Friedman, 2001; Chen and Guestrin, 2016) have proved to be enormously successful because of their strong predictive performance (Caruana and Niculescu-Mizil, 2006; Caruana et al., 2008; Fernández-Delgado et al., 2014). Indeed, RFs and GBTs regularly outperform even deep learning on medium-sized tabular datasets (Grinsztajn et al., 2022). Nonetheless, these tree-based methods still suffer from several notable problems: They are defined via algorithms rather than via statistical models, so it is often difficult to quantify the uncertainty of their predictions; they use greedy splitting criteria, so there is no guarantee for the optimality of the fitted model; RFs in particular grow their trees independently of each other, therefore making them statistically inefficient when fitted to data with additive structure (Tan et al., 2022a).

To address these issues, Chipman et al. (1998) proposed a Bayesian adaptation of CART (BCART) and later an ensemble of Bayesian CART trees, which they called Bayesian Additive Regression Trees (BART) (Chipman et al.,

^{*}Equal contribution, alphabetical ordering

2010). These are Bayesian non-parametric regression models, which put a prior on the space of regression functions, assume a likelihood for the observed data, and combine these to obtain a posterior. In the case of Bayesian CART and BART, priors and posteriors are supported on the subspace of functions that can be realized by decision trees (or their ensembles). Similar to Gaussian process (GP) regression, the posterior distribution can be used to provide posterior predictive credible intervals. On the other hand, unlike GP regression, there is no closed form formula for the BART posterior and one has to sample from it approximately via a Markov chain Monte Carlo (MCMC) algorithm.

BART has been shown empirically to enjoy strong predictive performance that is sometimes even superior to that of RFs and GBTs, especially after hyperparameter optimization (Hill et al., 2020). Naturally, it has become increasingly popular in diverse fields ranging from the social sciences (Green and Kern, 2010; Yeager et al., 2019) to biostatistics (Wendling et al., 2018; Starling et al., 2020) and has been particularly enthused causal inference researchers (Hill, 2011; Green and Kern, 2012; Kern et al., 2016; Dorie et al., 2019; Hahn et al., 2019). Extending and improving BART methodology remains a highly active area of research, with many variants of the algorithm proposed over the last few years (see for instance Linero (2018); Pratola (2016); Pratola et al. (2020); Luo and Pratola (2023), as well as the survey Hill et al. (2020) and the references therein.)

The strong predictive performance of BART is supported by a burgeoning body of theoretical evidence regarding the BART posterior. Most significantly, researchers have shown that the BART posterior concentrates around the true regression function used to generate the response data as n , the number of training samples, increases, with this concentration happening at optimal rates under various assumptions on the smoothness and sparsity of the regression function and for appropriate prior choices (Ročková and Saha, 2019; Ročková and van der Pas, 2020; Linero and Yang, 2018; Jeong and Rockova, 2020; Rockova and Rousseau, 2021; Castillo and Ročková, 2021). Achieving these optimal rates does not require any oracle knowledge or hyperparameter tuning—instead BART automatically adapts to the level of smoothness and sparsity, with the former even happening at a local level (Rockova and Rousseau, 2021).

1.2 Observed Poor Mixing of BART and its Significance

While there is evidence that the BART *posterior* enjoys favorable properties under a variety of settings, the fact that we can only sample approximately from the posterior via MCMC creates a gap in our understanding of how and why BART works. Specifically, if the sampler chain does not converge efficiently to the posterior distribution, that is, if it does not mix well, the output of BART *algorithm* may not enjoy the same desirable inferential properties as the BART posterior. Most popular BART implementations use remarkably similar samplers based on the original design of Chipman et al. (2010), which uses a Bayesian backfitting approach to update one tree at a time via proposed local changes to the tree nodes coupled with a Metropolis-Hastings filter. Unfortunately, as described by Hill et al. (2020), “while this algorithm is often effective, it does not always mix well.” Indeed, poor mixing for this sampler has been empirically documented by multiple sources (Chipman et al., 1998; Carnegie, 2019).

The literature contains various suggestions on how to improve the mixing time for the BART sampler. These include parallelization (Pratola et al., 2014), modifying the MCMC proposal moves (Wu et al., 2007; Pratola, 2016; Kim and Rockova, 2023), warm starts from greedily constructed tree ensembles (He and Hahn, 2021), or running multiple chains (Carnegie, 2019). Despite this interest, there has been minimal theoretical work done to quantify the mixing time and to understand why and under what settings slow mixing occurs.

1.3 Main Contributions

In this paper, we show theoretically, assuming discrete covariates, that the BART sampler often converges slowly to its posterior, confirming the empirical observations of Hill et al. (2020) and other researchers. In fact, the convergence unexpectedly becomes *worse* as n increases, in contrast with the posterior’s concentration to the true regression function becoming better.

To state our results more formally, note that a regression tree is parameterized via its *tree structure* (which features are split on and at which thresholds) and its leaf parameters (the function value on each leaf). BART combines these parameters over multiple trees to parameterize a tree ensemble (see Section 2.2.). Since the leaf parameters can be sampled in closed form conditionally on tree structures, the original BART sampler, under a slight modification, can be thought of as a Markov chain on the space of tree structures.

Data Generating Process	Allowed Moveset	Multiple Trees	Lower Bound
Additive	Full	Yes	Square root
Additive	“Grow” and “Prune”	Yes	Polynomial
Contains Pure Interaction	“Grow” and “Prune”	Yes	Square root
Root Dependence	“Grow” and “Prune”	No	Exponential

Table 1: A summary of the HPDR hitting time lower bounds provided by our paper and their dependence on the training sample size n (last column). The first two lower bounds apply to additive generative models. The third applies to a setting where the generative regression function contains a pure interaction (defined in Section 5.2.) The fourth applies to the setting where the generative function has “root dependence”, which represents a form of asymmetric dependence on the features (defined in Section 5.3.) All lower bounds apart from the fourth allow BART to use multiple trees.

We show that the BART posterior concentrates on a set comprising all optimal tree structures (i.e. those with the smallest bias and complexity), which also forms a highest posterior density region (HPDR). On the other hand, the BART sampler’s *hitting time* for this set increases with n under four common data generative settings. In other words, the sampler requires more and more steps to reach any optimal tree structure. Note that this is a frequentist analysis and requires the assumption of a generative model for the data that can and will be different from the Bayesian parameterization. Our hitting time lower bounds are summarized in the Table 1.

We complement our theoretical analysis with a comprehensive simulation study of BART involving a wide range of data-generating processes with continuous covariates. From now on, we use the term approximate posterior to refer to the distribution obtained from 1000 MCMC samples after a generous burn-in (5000 iterations as opposed to the default of 100). To create a proxy for the exact posterior, we combine samples from multiple (5) sampler chains, which is known to improve mixing (Carnegie, 2019).¹ We compare the performance of the approximate and multi-chain approximate posteriors via two metrics: (i) the RMSE of their posterior mean function from the true regression function on a held-out test set, (ii) the empirical coverage of their pointwise credible intervals for the true regression function. In both cases, the relative performance gap between the approximate and multi-chain approximate posteriors (measured as a ratio) increases as n increases. These two findings provide further evidence that the approximate BART posterior becomes increasingly and meaningfully different from the exact posterior as n increases. We also perform two other experiments to validate this claim in the setting of our theoretical lower bounds. Our theory and simulations thus echo existing advice that BART users should not blindly take its posterior credible intervals at face value and should run multiple chains whenever feasible.

Lastly, our theoretical results and their proof strategies yield insights on why the BART sampler has trouble converging, which leads us to suggest possible ways to improve its performance. Most importantly, our proof will show that a major reason why hitting times grow with training sample size is because the “temperature” of the BART sampler is inversely proportional to the training sample size. There is no need for temperature and sample size to be intrinsically tied in this manner, and we conjecture that building in more flexible temperature control, such as via simulated tempering, may help to accelerate mixing.

1.4 Prior Work on Mixing for BART

Kim and Rockova (2023) and our prior work (Ronen et al., 2022) sought to analyze mixing times for the Bayesian CART sampler. Both made the surprising discovery that its mixing time can grow exponentially in the training sample size, when the only allowed moves are growing new leaves and pruning existing ones. Kim and Rockova (2023) studied this in a one-dimensional setting and further showed that, with a more aggressive move set, Bayesian CART constrained to dyadic splits has a mixing time upper bound that is linear in the sample size. Ronen et al. (2022)’s proof strategy was to show that this Markov chain has a bottleneck state—the trivial tree comprising a single node. This bottleneck arises because when Bayesian CART makes a wrong first split followed by other informative splits, the only way to reverse the wrong first split involves pruning the informative splits, which becomes increasingly difficult as the training sample size increases.

¹Carnegie (2019) investigated the use of 1, 4, and 10 chains. Taking reference from this, we used 5 chains but did not investigate the effect of increasing the number of chains beyond this value. We believe that the optimal number could possibly vary with the sample size and DGP.

When we submitted our prior work to a peer review, some reviewers rightly pointed out that our mixing time lower bounds may be misleading to practitioners. This is because different tree structures could realize the same partition of the covariate space and hence implement the same regression function. In other words, the Bayesian CART model is not identifiable at the level of tree structures, which means that failure to mix in this space of tree structures may be benign (Redner and Walker, 1984). It does not reflect a failure to mix at the level of regression functions, let alone any degradation in the inferential properties of the algorithm’s output.




Figure 1: The Bayesian CART model is not identifiable at the level of tree structures. The two tree structures shown on the left both realize the same partition of the covariate space and even the same regression function, which is shown on the right.

A further limitation of our prior analysis is that mixing time is defined as the worst case time to convergence as we vary over all possible initial states. Since Bayesian CART is always initialized at the trivial tree, a lower bound on this worst case quantity may be unreasonably pessimistic.

These issues with relevance are resolved by considering hitting times.

2 Data Generation Models for BART and for Frequentist Analysis

2.1 Generative Model for Frequentist Analysis

We assume a d -dimensional discrete covariate space $\mathcal{X} = \{1, 2, \dots, b\}^d$ and a regression function $f^*: \mathcal{X} \rightarrow \mathbb{R}$. Suppose that we observe a training data set $\mathcal{D}_n$ comprising n independent and identically distributed tuples $(\mathbf{x}_i, y_i)$ with $\mathbf{x}_i \sim \nu$ and $y_i = f^*(\mathbf{x}_i) + \epsilon_i$ for $i = 1, 2, \dots, n$. Here, ν is a measure on $\mathcal{X}$ with full support, while ϵ is a sub-Gaussian random variable. For notational convenience, we will use $\mathbf{X}$ to denote the $n \times d$ matrix formed by stacking the covariate vectors as rows. We will also use $\mathbf{y}$, $\mathbf{f}^*$, and $\boldsymbol{\epsilon}$ to denote the n -dimensional vectors formed by stacking the responses y_i , the function values $f^*(\mathbf{x}_i)$ and the noise components ϵ_i respectively. Similarly, we will denote all vectors and matrices with boldface notation, with subscripts referencing the index of the vector. Vector coordinates will be denoted using regular font. We will denote probabilities and expectations with respect to $\mathcal{D}_n$ using $\mathbb{P}_n$ and $\mathbb{E}_n$ respectively.

2.2 Bayesian Model Specification for BART

We first describe the version of BART that we analyze in our paper, before discussing its differences with the version described by Chipman et al. (2010) and which is still most commonly used in practice.

Regression trees. A binary axis-aligned regression tree is parameterized by a tuple $(\mathfrak{T}, \boldsymbol{\mu})$. Here, $\mathfrak{T}$ refers to the *tree structure*, which specifies the topology of the tree as a rooted binary tree planar graph and, given an ordering of the graph’s vertices (e.g. via breadth-first search), specifies the splitting rule for each internal node j . Note that the splitting rule comprises a feature v_j and a threshold t_j . The leaves $\mathcal{L}_1, \mathcal{L}_2, \dots, \mathcal{L}_b$ of $\mathfrak{T}$ thus correspond to rectangular regions of the covariate space that together form a partition of the space. We let $\boldsymbol{\mu} \in \mathbb{R}^b$ be a vector of *leaf parameters*, one for each leaf of $\mathfrak{T}$. Together, $(\mathfrak{T}, \boldsymbol{\mu})$ specify a piecewise constant function g that outputs

$$g(\mathbf{x}; \mathfrak{T}, \boldsymbol{\mu}) = \mu_{l(\mathbf{x})},$$

where $l(\mathbf{x})$ is the index of the leaf containing $\mathbf{x}$.

Sum-of-trees model. Given observed data $\mathcal{D}_n$, the BART model posits $y_i = f(\mathbf{x}_i) + e_i$ for $i = 1, 2, \dots, n$, where $e_1, e_2, \dots, e_n \sim_{i.i.d.} \mathcal{N}(0, \sigma^2)$, and f is a sum of the outputs of m trees:

$$f(\mathbf{x}) = g(\mathbf{x}; \mathfrak{T}_1, \boldsymbol{\mu}_1) + g(\mathbf{x}; \mathfrak{T}_2, \boldsymbol{\mu}_2) + \dots + g(\mathbf{x}; \mathfrak{T}_m, \boldsymbol{\mu}_m).$$

We denote the ordered tuple $(\mathfrak{T}_1, \mathfrak{T}_2, \dots, \mathfrak{T}_m)$ by $\mathfrak{E}$ and call it a *tree structure ensemble* (TSE). We shall abuse notation and use $\boldsymbol{\mu}$ to refer to the concatenation of $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \dots, \boldsymbol{\mu}_m$. Note that, when conditioned on $\mathfrak{E}$ and $\mathbf{X}$, this is just a Bayesian linear regression model. To see this, let Ψ denote the $n \times b$ matrix whose columns are the indicator vectors over the training set of each leaf in $\mathfrak{E}$. We then have

$$\mathbf{y} = \Psi \boldsymbol{\mu} + \mathbf{e}. \quad (1)$$

Priors. We assume a fixed prior distribution p on $\Omega_{\text{TSE}, m}$, the space of TSEs with m trees.² Conditioned on a TSE $\mathfrak{E}$, the conditional prior distribution on the leaf parameters is an isotropic Gaussian, i.e. $p(\boldsymbol{\mu}|\mathfrak{E}) \sim \mathcal{N}(0, (\sigma^2/\lambda)\mathbf{I}_b)$, where b is the total number of leaves in all trees in $\mathfrak{E}$. Both σ^2 and λ are assumed to be fixed hyperparameters, with σ^2 taking the same value as that used in the variance of the additive noise e_i , while λ is a modulation parameter that should be set to approximately the reciprocal of the signal-to-noise ratio.

Differences with in-practice BART. Chipman et al. (1998) proposed a prior on tree structures defined in terms of a stochastic process. Starting from a single root node, the process recursively splits each node at depth d with probability $\alpha(1+d)^{-\beta}$, where α and β are hyperparameters with default values $\alpha = 0.95$ and $\beta = 2$ respectively. Features and thresholds for splits are selected uniformly at random. Chipman et al. (2010) extended this to a prior on TSEs by independence, i.e. $p(\mathfrak{E}) = \prod_{j=1}^m p(\mathfrak{T}_j)$. After rescaling the response variable to lie between -0.5 and 0.5 , the leaf parameter standard deviation is set to be $\sigma_\mu = 0.5/k\sqrt{m}$, where k is a further hyperparameter with default value $k = 2$. Finally, an inverse- χ^2 hyperprior is placed on the noise variance σ^2 and is calibrated to the observed data. This last assumption is the only way in which the BART model we study in this paper departs from that in Chipman et al. (2010). We make this change for analytical tractability and believe it to be minor, since simulations show that the posterior on σ^2 quickly converges to a fixed value and our theoretical guarantees hold for any fixed choice of σ^2 .

3 Sampling from BART via MCMC

3.1 The In-Practice BART Sampler

The sampler proposed by Chipman et al. (2010) can be described as a “Metropolis-within-Gibbs MCMC sampler” (Hill et al., 2020). More precisely, for each outer loop of the algorithm, it iterates over the tree indices $j = 1, 2, \dots, m$ and updates the j -th pair $(\mathfrak{T}_j, \boldsymbol{\mu}_j)$ using an approximate draw from the conditional distribution $p(\mathfrak{T}_j, \boldsymbol{\mu}_j | \mathfrak{E}_{-j}, \boldsymbol{\mu}_{-j}, \mathbf{y}, \sigma^2)$, where $\mathfrak{E}_{-j}$ and $\boldsymbol{\mu}_{-j}$ refer to the concatenation of current tree structures and leaf parameter vectors respectively, each

²Because the emphasis in our theory is on the dependence of hitting times on training sample size in large samples, the specific form of the prior holds no bearing on our results.

with the j -th index omitted.³ As a final step in the loop, it updates σ^2 using a draw from its full conditional distribution $p(\sigma^2|\mathfrak{E}, \boldsymbol{\mu}, \mathbf{y})$.

To describe how to sample (approximately) from $p(\mathfrak{T}_j, \boldsymbol{\mu}_j|\mathfrak{E}_{-j}, \boldsymbol{\mu}_{-j}, \mathbf{y}, \sigma^2)$, we first describe Chipman et al. (1998)’s algorithm for Bayesian CART, i.e. when the ensemble comprises a single tree. In this case, we first factorize the posterior into a conditional posterior on leaf parameters and a marginal posterior on tree structures:⁴

$$p(\mathfrak{T}, \boldsymbol{\mu}|\mathbf{y}, \sigma^2) = p(\boldsymbol{\mu}|\mathfrak{T}, \mathbf{y}, \sigma^2)p(\mathfrak{T}|\mathbf{y}, \sigma^2). \quad (2)$$

The first multiplicand on the right, $p(\boldsymbol{\mu}|\mathfrak{T}, \mathbf{y}, \sigma^2)$, is a multivariate Gaussian (with diagonal covariance) and can be sampled from directly. The second multiplicand is proportional to $p(\mathbf{y}|\mathfrak{T}, \sigma^2)p(\mathfrak{T})$, which is the product between a marginal likelihood of a Bayesian linear regression model and the prior. The marginal likelihood can be computed using standard techniques, but the posterior cannot be sampled directly. As such, a Metropolis-Hastings sampler is used with the following types of proposed moves:

1. Pick a leaf in the tree and split it (grow);
2. Pick two adjacent leaves and collapse them back into a single leaf (prune);
3. Pick an interior node and change the splitting rule (change);
4. Pick a pair of parent-child nodes that are both internal and swap their splitting rules, unless both children of the parent have the same splitting rules, in which case, swap the splitting rule of the parent with that of both children (swap).

Note that all selections in these proposed moves (of nodes, splitting rules, etc.) are made uniformly at random from all available choices.⁵ The proposed move types are chosen with probabilities π_g , π_p , π_c , and π_s respectively. Let $Q(-, -)$ denote the transition kernel of the proposal, i.e. $Q(\mathfrak{T}, \mathfrak{T}^*)$ is the probability of tree structure $\mathfrak{T}^*$ being proposed given current tree structure $\mathfrak{T}$. With $\mathfrak{T}$ and $\mathfrak{T}^*$ thus defined, the Metropolis-Hastings algorithm accepts the proposal with probability

$$\alpha(\mathfrak{T}, \mathfrak{T}^*) := \min \left\{ \frac{Q(\mathfrak{T}^*, \mathfrak{T})p(\mathfrak{T}^*|\mathbf{y}, \sigma^2)}{Q(\mathfrak{T}, \mathfrak{T}^*)p(\mathfrak{T}|\mathbf{y}, \sigma^2)}, 1 \right\}.$$

A simple reparameterization trick is used to adapt this sampler to the case when the ensemble has multiple trees. Because of the independence of the priors on different trees and the Gaussian likelihood, the conditional posterior $p(\mathfrak{T}_j, \boldsymbol{\mu}_j|\mathfrak{E}_{-j}, \boldsymbol{\mu}_{-j}, \sigma^2, \mathbf{y})$ can be rewritten in terms of the residual vector

$$\mathbf{r}_{-j} := \mathbf{y} - \sum_{k \neq j} g(\mathbf{X}; \mathfrak{T}_k, \boldsymbol{\mu}_k).$$

Specifically, we have

$$p(\mathfrak{T}_j, \boldsymbol{\mu}_j|\mathfrak{E}_{-j}, \boldsymbol{\mu}_{-j}, \sigma^2, \mathbf{y}) = p(\mathfrak{T}_j, \boldsymbol{\mu}_j|\mathbf{r}_{-j}, \sigma^2),$$

where the right-hand side is the single-tree posterior. A single Metropolis-Hastings update step as described above is performed to draw an approximate sample $(\mathfrak{T}_j, \boldsymbol{\mu}_j)$ from the conditional posterior.

3.2 The Analyzed BART Sampler

The BART sampler described above is difficult to analyze because the deterministic Gibbs outer loop makes it a time-varying Markov chain. More significantly, it is convenient in analyzing Bayesian CART to collapse the Markov chain state space by marginalizing out the leaf parameters. The collapsed state space is simply the space of tree structures, which is discrete and finite. However, we are unable to do this for BART in general because of the conditioning on the

³Note that the conditional distribution is of course also conditional on the observed covariate data $\mathbf{X}$. However, since this is always conditioned upon, we omit it from our notation to avoid clutter.

⁴Chipman et al. (1998)’s formulation of the Bayesian CART sampler marginalizes out σ^2 instead of conditioning on it. As this is no longer done for BART, we omit discussing it to avoid confusing readers.

⁵Splits that result in empty leaves are not allowed.

residuals from other trees in the inner loop. Both of these difficulties make it impossible to apply standard techniques in Markov chain theory.

To overcome this, we propose an adaptation of the sampler that brings it closer to Bayesian CART. First, we imitate (2) and factorize the posterior into a conditional posterior on leaf parameters and a marginal posterior on tree *ensemble* structures:

$$p(\mathfrak{E}, \boldsymbol{\mu} | \mathbf{y}) = p(\boldsymbol{\mu} | \mathfrak{E}, \mathbf{y}) p(\mathfrak{E} | \mathbf{y}).$$

The conditional posterior on leaf parameters is still a multivariate Gaussian and can be sampled from directly, while the marginal posterior on tree ensemble structures remains proportional to the product of the marginal likelihood of a Bayesian linear regression model and the prior: $p(\mathfrak{E} | \mathbf{y}) \propto p(\mathbf{y} | \mathfrak{E}) p(\mathfrak{E})$ (see (1)). To sample from this marginal posterior, we run Metropolis-Hastings MCMC similarly to before. However, instead of cycling deterministically through the trees in an inner loop as before, we pick a tree index uniformly at random. We propose an updated tree using the same transition kernel $Q(-, -)$, but write the acceptance probability in terms of the full marginal posterior instead of conditioning on the residuals from other trees. For further clarity, the algorithm is summarized in pseudocode as Algorithm 1.

We denote the transition kernel of the sampler using $P(-, -)$, and, to avoid confusion with randomness arising from sampling the training set, we will denote all probabilities and expectations with respect to the algorithmic randomness using P and E respectively.

Algorithm 1 BART sampler.

- 1: **BART**($\mathcal{D}_n$: data, m : no. of trees, σ^2 : guess for noise variance, λ : guess for reciprocal SNR, $\boldsymbol{\pi}$: proposal probabilities, p_{TSE} : TSE prior, t_{max} : no. of sampler iterations)
 - 2: Initialize $\mathfrak{T}_1, \mathfrak{T}_2, \dots, \mathfrak{T}_m$ as trivial trees.
 - 3: **for** $t = 1, 2, \dots, t_{max}$
 - 4: Sample $k \sim \text{Unif}(\{1, 2, \dots, m\})$.
 - 5: Propose $\mathfrak{T}^* \sim Q(\mathfrak{T}_k, \mathfrak{T}^*)$.
 - 6: Set $\alpha(\mathfrak{T}_k, \mathfrak{T}^*) = \min \left\{ \frac{Q(\mathfrak{T}^*, \mathfrak{T}_k) p(\mathfrak{T}_1, \dots, \mathfrak{T}_{k-1}, \mathfrak{T}^*, \mathfrak{T}_{k+1}, \dots, \mathfrak{T}_m | \mathbf{y})}{Q(\mathfrak{T}_k, \mathfrak{T}^*) p(\mathfrak{T}_1, \dots, \mathfrak{T}_{k-1}, \mathfrak{T}_k, \mathfrak{T}_{k+1}, \dots, \mathfrak{T}_m | \mathbf{y})}, 1 \right\}$.
 - 7: Set $\mathfrak{T}_k = \mathfrak{T}^*$ with probability $\alpha(\mathfrak{T}_k, \mathfrak{T}^*)$.
-

4 BIC for BART and Posterior Concentration on Optimal TSEs

The goal of this section is to first show how to quantify the bias and complexity of a TSE separately and then jointly via BIC. We will then show that, as a function of TSEs, the posterior probability $p(\mathfrak{E} | \mathbf{y})$ concentrates on the set of TSEs with zero bias and the lowest possible complexity, and are therefore minimizers of BIC. As such, as argued in the introduction, the highest posterior density region contains all of the most desirable TSEs. This implies that lower bounds on the hitting times of this region reflect computational drawbacks of practical consequence.

4.1 Measuring Bias and Complexity for TSEs

We first discuss how to quantify the bias and complexity of a TSE.

Partitions. A cell $\mathcal{C}$ is a rectangular region of $\mathcal{X}$, i.e.

$$\mathcal{C} = \{\mathbf{x} \in \mathcal{X} : a_i < x_i \leq b_i \text{ for } i = 1, \dots, d\},$$

with lower and upper limits a_i and b_i respectively in coordinate i for $i = 1, 2, \dots, d$. A *partition* is a collection of disjoint cells $\mathcal{C}_1, \dots, \mathcal{C}_b$ whose union is the whole space $\mathcal{X}$. Every tree structure $\mathfrak{T}$ induces a partition $\mathcal{P}$ via its leaves. Not only is $\mathcal{P}$ a sufficient statistic for $\mathfrak{T}$, it also completely characterizes the bias and complexity of the resulting data model conditioned on $\mathfrak{T}$. Indeed, this data model is just Bayesian linear regression on the indicator functions on the leaves of $\mathfrak{T}$. Since the functions are orthogonal, the degrees of freedom of the regression is equal to the size of the partition.

Partition ensemble models (PEMs). A TSE $\mathfrak{E}$ induces an ensemble of partitions $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_m$. Indeed the data model conditioned on $\mathfrak{E}$ is still a Bayesian linear regression on the indicator functions of all leaves in $\mathfrak{E}$. However, these indicators are no longer orthogonal, which means that different ensembles, with possibly different numbers of leaves, can give rise to the same subspace of regressors, making $\mathfrak{E}$ not identifiable from data. To avoid this issue, we directly consider the subspace of the regression function. Formally, let $\mathbb{V} \subset L^2(\mathcal{X}^d, \nu)$ be the subspace spanned by indicators of the cells in $\mathcal{P}_j$ for $j = 1, 2, \dots, m$. We call this the *partition ensemble model* (PEM) associated to the TSE $\mathfrak{E}$ and indicate this association via the mapping $\mathbb{V} = \mathcal{F}(\mathfrak{E})$.

Measuring bias. Let $\Pi_{\mathfrak{E}}$ denote orthogonal projection onto $\mathcal{F}(\mathfrak{E})$ in $L^2(\nu)$. We define the squared (mangitude of the) bias of $\mathfrak{E}$ with respect to a regression function f as

$$\text{Bias}^2(\mathfrak{E}; f) := \int (f - \Pi_{\mathfrak{E}}[f])^2 d\nu. \quad (3)$$

This is precisely the squared bias of Bayesian linear regression on $\mathfrak{E}$ if we ignore the regularization effect from the leaf parameter priors, which is inconsequential in large sample sizes.

Measuring complexity. When conditioned on a TSE $\mathfrak{E}$ and ignoring regularization from leaf parameter priors, the degrees of freedom of the resulting Bayesian linear regression model is just the dimension of $\mathcal{F}(\mathfrak{E})$. We denote this by $\text{df}(\mathfrak{E})$ and use it as a measure of complexity of $\mathfrak{E}$. Note that this definition does not depend on the covariate distribution ν (see Lemma J.2 in the appendix.)

Function dimension and optimal sets. Excessive complexity leads to overfitting and is hence undesirable. To quantify the excess, we first define the *m-ensemble dimension* of a regression function f as

$$\text{dim}_m(f) := \min\{\text{df}(\mathfrak{E}) : f \in \mathcal{F}(\mathfrak{E}) \text{ and } \mathfrak{E} \in \Omega_{\text{TSE}, m}\}. \quad (4)$$

In large sample sizes n , which is the setting we are concerned with, the TSEs that result in the smallest MSE must be bias-free. We hence define the set of optimal TSEs in $\Omega_{\text{TSE}, m}$ to be the minimizers of (4). More generally, we define a series of nested sets with increasing levels of suboptimality tolerance via:

$$\text{OPT}_m(f, k) := \{\mathfrak{E} \in \Omega_{\text{TSE}, m} : f \in \mathcal{F}(\mathfrak{E}) \text{ and } \text{df}(\mathfrak{E}) \leq \text{dim}_m(f) + k\}.$$

4.2 BIC and BART Posterior Concentration

The Bayesian information criterion (BIC) (Schwarz, 1978) of a TSE $\mathfrak{E}$ is given by

$$\text{BIC}(\mathfrak{E}) = \frac{\mathbf{y}^T(\mathbf{I} - \mathbf{P}_{\mathfrak{E}})\mathbf{y}}{\sigma^2} + \text{df}(\mathfrak{E}) \log n + \log(2\pi\sigma^2)n.$$

Here, $\mathbf{P}_{\mathfrak{E}}$ refers to projection onto $\mathcal{F}(\mathfrak{E})$ with respect to the empirical norm $\|\cdot\|_n$ (realized as a matrix.) Ignoring the effect of the noise vector for now, we see that the first term, divided by the sample size n , is an estimate for the squared bias. Meanwhile, the second term directly measures the model complexity. Hence, BIC quantifies the quality of a TSE by accounting for both bias and complexity. Indeed, under our data generative model (Section 2.1) we have the following concentration lemma:

Proposition 4.1 (Concentration of BIC differences). *Consider two TSEs $\mathfrak{E}$ and $\mathfrak{E}'$ and denote the difference in their BIC values as $\Delta\text{BIC}(\mathfrak{E}, \mathfrak{E}') = \text{BIC}(\mathfrak{E}) - \text{BIC}(\mathfrak{E}')$. Then for any $0 < \delta < 1$, with probability at least $1 - \delta$ with respect to $\mathbb{P}_n$, we have*

$$\Delta\text{BIC}(\mathfrak{E}, \mathfrak{E}') = \frac{n}{\sigma^2} (\text{Bias}^2(\mathfrak{E}; f^*) - \text{Bias}^2(\mathfrak{E}'; f^*)) + O\left(\sqrt{n \log(1/\delta)} + \log(1/\delta)\right). \quad (5)$$

If furthermore, both TSEs have the same bias, i.e. $\Pi_{\mathfrak{E}}[f^] = \Pi_{\mathfrak{E}'}[f^*]$, then we have*

$$\Delta\text{BIC}(\mathfrak{E}, \mathfrak{E}') = \log n (\text{df}(\mathfrak{E}) - \text{df}(\mathfrak{E}')) + O(\log(1/\delta)). \quad (6)$$

From this proposition, we also see that $\text{OPT}_m(f^*, k)$ for $k = 0, 1, 2, \dots$ are just sublevel sets of BIC when n is large enough. We next show that BIC is closely connected to the log marginal likelihood for TSEs as follows:

Proposition 4.2 (Log marginal likelihood and BIC). *Consider a TSE $\mathfrak{E}$. Then for any $0 < \delta < 1$, there is a minimal sample size N so that for all $n \geq N$, with probability at least $1 - \delta$ with respect to $\mathbb{P}_n$, the log marginal likelihood satisfies*

$$\log p(\mathbf{y}|\mathfrak{E}) = -\frac{\text{BIC}(\mathfrak{E})}{2} + O(1).$$

Consequently, the log marginal posterior also satisfies

$$\log p(\mathfrak{E}|\mathbf{y}) = -\frac{\text{BIC}(\mathfrak{E})}{2} - \log p(\mathbf{y}) + O(1).$$

This almost linear relationship implies that $\text{OPT}_m(f, k)$ for $k = 0, 1, 2, \dots$ are also superlevel sets of the marginal posterior. In other words, they form highest posterior density regions (HPDR). As advertised, these will be the target sets of our hitting time analysis. Finally, combining the previous two propositions gives the following result on posterior concentration.

Proposition 4.3 (BART posterior concentration). *For any $0 < \delta, \epsilon < 1$, there is a minimal sample size N so that for all $n \geq N$, with probability at least $1 - \delta$ with respect to $\mathbb{P}_n$, the marginal posterior measure on $\Omega_{\text{TSE}, m}$ satisfies*

$$p(\text{OPT}_m(f^*, 0) \mid \mathbf{y}) > 1 - \epsilon.$$

5 Hitting Time Lower Bounds for BART MCMC

Our primary findings consist of lower bounds on HPDR hitting times for BART, explored across four distinct settings for BART and the data generating process (DGP). As discussed in the previous section, these regions also comprise sublevel sets of BIC. We first define hitting times in a general setting.

Hitting times. Let (X_t) be a discrete time Markov chain on a finite state space Ω . Let $\mathcal{A} \subset \Omega$ be a subset. The *hitting time* of $\mathcal{A}$ is defined as:

$$\tau_{\mathcal{A}} := \min\{t \geq 0: X_t \in \mathcal{A}\}.$$

Note that this is a random variable and that it, in principle, depends on the initial state X_0 . In our analysis, the initial state is always chosen to be an ensemble of trivial trees and so will not be referenced in the notation to avoid clutter.

5.1 Square Root and Polynomial Lower Bounds for Additive Models

Our first two lower bounds are for the setting where the DGP is an additive model.

Theorem 5.1 (Lower bound for additive model). *Let f^* be an additive function, i.e.*

$$f^*(\mathbf{x}) = f_1(x_1) + f_2(x_2) + \dots + f_{m'}(x_{m'})$$

*with $m' \geq 2$. Suppose $x_1, x_2, \dots, x_{m'}$ are independent. Suppose $m \leq m'$, and we make arbitrary choices for the other BART hyperparameters $\sigma^2, \lambda, \pi, p_{\text{TSE}}$. Then for any $0 < \delta < 1$, there is a minimal sample size N so that for all $n \geq N$, with probability at least $1 - \delta$ with respect to $\mathbb{P}_n$, the Markov chain induced by **BART**($\mathcal{D}_n, m, \sigma^2, \lambda, \pi, p_{\mathfrak{E}}, -$) satisfies*

$$E\{\tau_{\text{OPT}_m(f^*, (q_{\max}-2)(q_{\min}-2)-2)}\} = \Omega(n^{1/2}),$$

where $q_{\max} = \max_{1 \leq i \leq m'} \dim_1(f_i)$, $q_{\min} = \min_{1 \leq i \leq m'} \dim_1(f_i)$.

If furthermore, $m < m'$ and we disallow “change” and “swap” moves, i.e. $\pi_c = \pi_s = 0$, then we have⁶

$$E\{\tau_{\text{OPT}_m(f^*, q_{\max}-q_{\min}-1)}\} = \Omega(n^{q_{\min}/2-1}).$$

⁶Recall that $\text{OPT}_m(f^*, 0) \subset \text{OPT}_m(f^*, k)$ for $k \geq 0$. Since $\text{OPT}_m(f^*, k) = \emptyset$ for negative k , these statements can be interpreted as being meaningless unless $q_{\max}$ and $q_{\min}$ satisfy the relevant constraints.

Additive models are natural generalizations of linear models that have been widely studied in statistics and machine learning (Hastie and Tibshirani, 1986; Hastie et al., 2009). When fitted to real world datasets, they often enjoy good prediction accuracy. Hence, we view an additive generative model to be a natural class of functions for our study.

Furthermore, several works deriving consistency guarantees for frequentist greedy decision trees and random forests have used additive models as a generative function class, as the assumption of additivity helps to circumvent some of the practical and theoretical difficulties arising from greedy splitting (Scornet et al., 2015; Klusowski, 2021). On the other hand, Tan et al. (2022a) showed generalization lower bounds for decision trees in this setting, with the recommendation that models with multiple trees fit additive models better than those comprising a single tree (see also Tan et al. (2022b).) Indeed, given the assumptions of Theorem 5.1, we have $\dim_l(f^*) > \dim_m(f^*)$ for any $l < m$, while $\dim_l(f^*) = \dim_m(f^*)$ for any $l \geq m$ (see Proposition F.7.) In other words, a minimum BIC value is achievable if and only if the number of trees in the BART model is larger than or equal to the number of components in the additive model.

However, Theorem 5.1 tells us that even when the BART model is correctly specified and contains an efficient representation of f^* , i.e. when $m' = m$, the BART sampler may fail to reach a TSE implementing such a representation within a reasonable time. This adds another perspective to recent arguments that there may be value in overparameterization (see for instance Bartlett et al. (2020).) Allowing for more trees than the number of additive components may empower the sampler with more freedom of navigation to avoid potential computational bottlenecks. We confirm this conjecture empirically in our simulations section.

5.2 Square Root Lower Bound for Pure Interactions

Our next lower bound is for the setting when the data generating process has a pure interaction, which we define as follows. Let x_i and x_j be two features, i.e. components of $\mathbf{x} \sim \nu$. We say that they form a *pure interaction* with respect to a regression function f^* if they are jointly dependent with the response y , but are separately conditionally independent of y for any conditioning set of indices I unless it includes the other’s index. Mathematically, we can write this as follows:

- $(x_i, x_j) \not\perp\!\!\!\perp y$;
- $x_i \perp\!\!\!\perp y \mid x_I$ and $x_j \perp\!\!\!\perp y \mid x_I$ for any $I \subset \{1, 2, \dots, d\}$ such that $i, j \notin I$.

A canonical example of a pure interaction is the exclusive-or (XOR) function over binary features with a uniform distribution, i.e. $f(\mathbf{x}) = x_1 x_2$ for $\mathcal{X} = \{-1, 1\}^d$. This function is well-known to be difficult to learn using either CART (Syrkanis and Zampetakis, 2020; Mazumder and Wang, 2024) or neural networks (Abbe et al., 2022). As such, it is perhaps unsurprising that the BART sampler also experiences difficulties in this setting.

Theorem 5.2 (Lower bound for pure interaction). *Let f^* contain a pure interaction. Suppose we disallow “change” moves, i.e. $\pi_c = 0$, and we make arbitrary choices for all other BART hyperparameters $m, \sigma^2, \lambda, \pi_g, \pi_p, \pi_s, p_{TSE}$. Then for any $0 < \delta < 1$, there is a minimal sample size N so that for all $n \geq N$, with probability at least $1 - \delta$ with respect to $\mathbb{P}_n$, the Markov chain induced by $\mathbf{BART}(\mathcal{D}_n, m, \sigma^2, \lambda, \boldsymbol{\pi}, p_{\mathcal{E}}, -)$ satisfies*

$$E\{\tau_{\text{OPT}_m(f^*, \infty)}\} = \Omega(n^{1/2}),$$

where $\text{OPT}_m(f^*, \infty) := \cup_{k=0}^{\infty} \text{OPT}_m(f^*, k)$.

Note that the suboptimality gap for Theorem 5.2 is much wider than that for Theorem 5.1. Indeed, it implies that the only TSEs that are reachable within $o(n^{1/2})$ iterations of the sampler have nonzero bias. As such, the MSE of the BART sampler output does not even converge to zero with the training sample size, unless we allow for $\Omega(n^{1/2})$ iterations of the sampler.

5.3 Exponential Lower Bound for Bayesian CART

Our final hitting time lower bound shows that the HPDR hitting time for Bayesian CART can be exponential in the training sample size. This complements and improves our results in Ronen et al. (2022), which provided an exponential

lower bound for mixing time for Bayesian CART. While the previous result relied on extremely weak assumptions, the improved version requires a new assumption on the asymmetry of the regression function f^* in terms of its dependence on different features. Specifically, we say that a regression function f^* has *root dependence* if there exists a feature x_i and threshold t such that:

- $\text{Corr}^2(y, \mathbf{1}\{x_i \leq t\}) > 0$;
- (i, t) does not occur as a root split on any tree structure $\mathfrak{T} \in \text{OPT}_1(f, 0)$.

An example of such a function is the “staircase” function $f^*(\mathbf{x}) = \sum_{j=1}^s \prod_{k=1}^j \mathbf{1}\{x_k > 1\}$. Any feature x_i for $2 \leq i \leq s$ satisfies the above two properties. On the other hand, additive functions on independent features do not satisfy these properties.

Theorem 5.3 (Lower bound for Bayesian CART with root dependence). *Suppose f^* has root dependence. Suppose $m = 1$ and that we disallow “change” and “swap” moves, i.e. $\pi_c = \pi_s = 0$. Suppose we make arbitrary choices for all other BART hyperparameters $\sigma^2, \lambda, \pi_g, \pi_p, p_{TSE}$. Then the Markov chain induced by $\mathbf{BART}(\mathcal{D}_n, 1, \sigma^2, \lambda, \pi, p_{\mathfrak{E}}, -)$ satisfies*

$$\liminf_{n \rightarrow \infty} \mathbb{E}_n \left\{ \frac{\log E\{\tau_{\text{OPT}_1(f^*, 0)}\}}{n} \right\} \geq \frac{1}{2\sigma^2} \left(\int (f^*)^2 d\nu - \left(\int f^* d\nu \right)^2 \right).$$

Remark 5.4. Our hitting time lower bounds directly imply mixing time lower bounds in the space of PEMs, which, as argued in Section 4, are identifiable from data. Since the bias and degrees of freedom of a TSE $\mathfrak{E}$ is defined in terms of its associated PEM $\mathcal{F}(\mathfrak{E})$, $\text{OPT}_m(f^*, k)$ is the preimage under $\mathcal{F}$ of a set $\widehat{\text{OPT}}_m(f^*, k)$ in the space of PEMs. Hence, a hitting time lower bound for $\text{OPT}_m(f^*, k)$ is simultaneously a hitting time lower bound for $\widehat{\text{OPT}}_m(f^*, k)$ when considering the induced Markov chain on the space of PEMs. It is easy to see that this is a lower bound for the mixing time.

Remark 5.5. For the sake of narrative clarity, we have not tried to optimize the suboptimality gaps (i.e. k in $\text{OPT}_m(f^*, k)$) in our lower bounds. Note also that we have not attempted to investigate the dependence of our lower bounds on other data or algorithmic hyperparameters. These of course influence the minimum sample size N as well as the hidden constant factor in the Big-Omega notation of our lower bounds.

6 Hitting Time Lower Bounds via Barrier Sets

In this section, we briefly outline the proof strategy we use to derive our hitting time lower bounds. Our first move is to make use of the standard interpretation of a symmetric Markov chain as a random walk on a network, whose vertices comprise the states of the Markov chain and whose edges comprise pairs of states with positive transition probability. We next notice that hitting times are closely related to escape probabilities, which can be interpreted as voltages on the network. Voltages can then be calculated using standard network simplification techniques. Putting these ingredients together creates the following recipe for deriving hitting time lower bounds:

Proposition 6.1 (Recipe for hitting time lower bounds). *Let $\mathfrak{E}_0, \mathfrak{E}_1, \mathfrak{E}_2, \dots$ denote the Markov chain induced by a run of $\mathbf{BART}(\mathcal{D}_n, m, \sigma^2, \lambda, \pi, p_{\mathfrak{E}}, -)$ for any fixed dataset $\mathcal{D}_n$ and any choice of hyperparameters. Let $\mathfrak{E}_{\text{bad}} \in \Omega_{TSE, m}$ be a TSE such that, for some $0 < \delta < 1$, with probability at least $1 - \delta$ with respect to $\mathbb{P}_n$,*

$$P\{\tau_{\mathfrak{E}_{\text{bad}}} < \tau_{\text{OPT}_m(f^*, k)}\} = \Omega(1).$$

Let $\mathcal{B} \subset \Omega_{TSE, m}$ be a subset such that every path from $\mathfrak{E}_{\text{bad}}$ to $\text{OPT}_m(f^, k)$ intersects $\mathcal{B}$. Then with probability at least $1 - 2\delta$ with respect to $\mathbb{P}_n$, the hitting time of $\text{OPT}_m(f^*, k)$ satisfies*

$$E\{\tau_{\text{OPT}_m(f^*, k)}\} = \Omega\left(\exp\left(\frac{1}{2} \min_{\mathfrak{E} \in \mathcal{B}} \Delta\text{BIC}(\mathfrak{E}, \mathfrak{E}_{\text{bad}})\right)\right).$$




Figure 2: Visual illustration of Proposition 6.1. The chain is initialized at $\mathfrak{E}_0$ and with positive probability hits a suboptimal TSE $\mathfrak{E}_{\text{bad}}$ before $\text{OPT}_m(f^*, k)$. This causes the chain to get stuck, as it can only reach $\text{OPT}_m(f^*, k)$ by passing through an “impassable” barrier set $\mathcal{B}$.

In other words, our hitting time lower bounds follow immediately once we are able to identify (i) a suboptimal TSE $\mathfrak{E}_{\text{bad}}$ that is reachable by the sampler before it hits $\text{OPT}_m(f^*, k)$ and (ii) a barrier $\mathcal{B}$ that separates $\mathfrak{E}_{\text{bad}}$ from $\text{OPT}_m(f^*, k)$, i.e. is a vertex cutset, and which has higher BIC, therefore acting as an impassable “barrier”. We now briefly describe our choices for each of the three different lower bound settings we have considered in this paper.

Additive models. For simplicity, we only discuss the case of $m' = m$ (when the number of components is equal to the number of trees.) For $i = 1, 2, \dots, m$, denote $q_i = \dim_1(f_i)$, and let $0 = \xi_{i,0} < \xi_{i,1} < \dots < \xi_{i,q_i} = b$ denote the *knots* of f_i , i.e. the values for which $f_i(\xi_{i,j}) \neq f_i(\xi_{i,j+1})$, together with the endpoints.⁷ Without loss of generality, assume that $f_1, \dots, f_m$ are ordered in descending order of their 1-ensemble dimension, i.e. $q_1 \geq q_2 \geq \dots \geq q_m$.

We now define $\mathfrak{E}_{\text{bad}}$ and a “bad set” $\mathcal{A}$. To this end, we define a collection of partition models $\mathbb{V}_1, \mathbb{V}_2, \dots, \mathbb{V}_m$ (spans of indicators in a single partition) as follows. First, for $i = 3, 4, \dots, m$, $j = 1, 2, \dots, q_i$, define the cells $\mathcal{L}_{i,j} := \{\mathbf{x} : \xi_{i,j-1} < x_i \leq \xi_{i,j}\}$, and set $\mathbb{V}_i = \text{span}(\{\mathbf{1}_{\mathcal{L}_{i,j}} : j = 1, 2, \dots, q_i\})$, i.e. for each i , $\mathbb{V}_i$ contains splits only on feature i and only at the knots of f_i . To introduce inefficiency, we define each of $\mathbb{V}_1$ and $\mathbb{V}_2$ to have splits on both features 1 and 2. This construction is demonstrated in Figure 3. The formal details are fairly involved and will be deferred to the appendix. We define $\mathcal{A}$ via

$$\mathcal{A} := \{(\mathfrak{T}_1, \mathfrak{T}_2, \dots, \mathfrak{T}_m) : \mathcal{F}(\mathfrak{T}_i) = \mathbb{V}_i \text{ for } i = 1, 2, \dots, m\},$$

and set $\mathcal{B}$ to be the *outer boundary* of $\mathcal{A}$, i.e.

$$\mathcal{B} = \{\mathfrak{E} \in \Omega_{\text{TSE},m} : \mathfrak{E} \text{ has an edge to } \mathcal{A} \text{ and } \mathfrak{E} \notin \mathcal{A}\}.$$

Finally, we pick $\mathfrak{E}_{\text{bad}}$ to be a particular element of $\mathcal{A}$ whose precise construction will be detailed in the Appendix H. Therein, we will also show that

- $\mathfrak{E}$ has zero bias for all $\mathfrak{E} \in \mathcal{A}$;
- $\mathcal{A} \cap \text{OPT}_m(f^*, (q_{\max} - 2)(q_{\min} - 2) - 2) = \emptyset$;
- $\min_{\mathfrak{E} \in \mathcal{B}} \Delta \text{BIC}(\mathfrak{E}, \mathfrak{E}_{\text{bad}}) \geq \log n - O(1)$.

⁷ $\dim_1(f_i)$ is simply the number of constant pieces of f_i or alternatively, one larger than the number of knots of f_i .




Figure 3: Visual illustration of the construction for $\mathfrak{E}_{\text{bad}}$ used in the proof of Theorem 5.1. The left panel displays the function $f_1(x_1) + f_2(x_2)$ together with all knots of f_1 and f_2 . We define $\mathfrak{E}_{\text{bad}}$ so that its first two trees $\mathfrak{T}_1$ and $\mathfrak{T}_2$ induce the partitions $\mathbb{V}_1$ and $\mathbb{V}_2$ respectively. These combine for a total of 13 leaves. On the other hand, an optimal TSE will instead make use of $\mathbb{V}_1^*$ and $\mathbb{V}_2^*$, which combine for a total of only 8 leaves. Nonetheless, we still have $f_1 + f_2 \in \mathbb{V}_1^* + \mathbb{V}_2^*$.

Pure interactions. We first define the notion of *reachability* as follows: Given $\mathfrak{E}, \mathfrak{E}' \in \Omega_{\text{TSE}, m}$, we say that $\mathfrak{E} \preceq \mathfrak{E}'$ if $\mathfrak{E}$ and $\mathfrak{E}'$ are connected by an edge and if either $\text{Bias}^2(\mathfrak{E}; f^*) > \text{Bias}^2(\mathfrak{E}'; f^*)$ or $\text{Bias}^2(\mathfrak{E}; f^*) = \text{Bias}^2(\mathfrak{E}'; f^*)$ and $\text{df}(\mathfrak{E}) \geq \text{df}(\mathfrak{E}')$. Note that, because we allow only “grow” and “prune” moves, for adjacent $\mathfrak{E}$ and $\mathfrak{E}'$, $\mathcal{F}(\mathfrak{E})$ and $\mathcal{F}(\mathfrak{E}')$ are nested subspaces, so that $\text{Bias}^2(\mathfrak{E}; f^*) = \text{Bias}^2(\mathfrak{E}'; f^*)$ if and only if $\Pi_{\mathfrak{E}}[f^*] = \Pi_{\mathfrak{E}'}[f^*]$. We say that $\mathfrak{E}$ is reachable from $\mathfrak{E}'$, denoted $\mathfrak{E} \succeq \mathfrak{E}'$, if there is a sequence of TSEs $\mathfrak{E} = \mathfrak{E}^0, \mathfrak{E}^1, \dots, \mathfrak{E}^k = \mathfrak{E}'$ such that $\mathfrak{E}^i \preceq \mathfrak{E}^{i+1}$ for $i = 0, 1, \dots, k-1$.

Without loss of generality, let (x_1, x_2) be a pure interaction for f^* . Let $\mathfrak{E}_{\text{bad}}$ be any TSE such that

- $\mathfrak{E}_{\text{bad}}$ is reachable from $\mathfrak{E}_\emptyset$;
- There does not exist $\mathfrak{E} \in \Omega_{\text{TSE}, m}$ such that $\mathfrak{E}$ is reachable from $\mathfrak{E}_{\text{bad}}$ but $\mathfrak{E}_{\text{bad}}$ is not reachable from $\mathfrak{E}$.

Note that such a TSE exists because $\Omega_{\text{TSE}, m}$ is finite and $\succeq$ is a partial ordering on this space. We set $\mathcal{A}$ to be the equivalence class of $\mathfrak{E}_{\text{bad}}$ under $\succeq$ and set $\mathcal{B}$ to be the outer boundary of $\mathcal{A}$. We will show in the appendix that no TSE in $\mathcal{A}$ makes a split on either x_1 or x_2 , which implies that $\mathcal{A} \cap \text{OPT}_m(f^*, \infty) = \emptyset$. We will also show that $\min_{\mathfrak{E} \in \mathcal{B}} \Delta \text{BIC}(\mathfrak{E}, \mathfrak{E}_{\text{bad}}) \geq \log n - O(1)$.

Bayesian CART. Without loss of generality, let x_1 be the feature that gives f^* root dependence. By assumption, there is a threshold t such that splitting the trivial tree on x_1 at t gives a decrease in squared bias. We set

$$\mathcal{A} = \{\mathfrak{T} \in \Omega_{\text{TSE}, 1} : \mathfrak{T} \text{ has root split on } x_1 \text{ at } t\},$$

and

$$\mathfrak{T}_{\text{bad}} = \arg \min \{\text{BIC}(\mathfrak{E}) : \mathfrak{E} \in \mathcal{A}\}.$$

Note that the outer boundary of $\mathcal{A}$ is a singleton set comprising the trivial tree $\mathfrak{T}_\emptyset$. By assumption, we have $\mathcal{A} \cap \text{OPT}_1(f^*, 0) = \emptyset$. We will show in the appendix that

$$\Delta \text{BIC}(\mathfrak{T}_\emptyset, \mathfrak{T}_{\text{bad}}) \geq \frac{n}{\sigma^2} \left(\int (f^*)^2 d\nu - \left(\int f^* d\nu \right)^2 \right) - o(n).$$

From these constructions, we also see that the reason why hitting times grow with training sample size is because the barrier sets become increasingly difficult to pass through. Heuristically, we can say that this is because the intrinsic “temperature” of the BART sampler is inversely proportional to the training sample size.

7 Theoretical Limitations and Future Work

In this section, we detail the limitations of our theoretical results, which naturally suggest directions for future work.

Difference from in-practice BART. The three differences between the version of BART we analyzed and BART as used in practice are that:

- We assume a fixed noise parameter σ^2 instead of putting a prior on it;
- At each iteration, we pick a random tree to update instead of cycling deterministically through the trees;
- We change the rejection probability in the Metropolis filter to be in terms of the marginal posteriors on TSEs instead of being conditional posteriors on the updated tree.

We believe that these differences to be relatively minor and that they may even improve the mixing properties of BART, albeit at the cost of higher computational complexity per iteration.⁸

Failure to address MSE and coverage directly. Let $\mathfrak{E}_0, \mathfrak{E}_1, \mathfrak{E}_2, \dots$ denote the Markov chain induced by a run of BART. For $j = 0, 1, 2, \dots$, let h_j be a draw from the conditional posterior on regression functions, $p(f|\mathfrak{E}_j, \mathbf{y})$. The BART algorithm returns the collection

$$\{h_{t_{\text{burn-in}}+1}, h_{t_{\text{burn-in}}+2}, \dots, h_{t_{\text{max}}}\}, \quad (7)$$

where $t_{\text{burn-in}}$ is the number of burn-in iterations to be discarded. The final fitted function is the mean of this collection, while approximate credible prediction intervals are derived from quantiles. While our hitting time lower bounds imply that the Markov chain on TSEs fails to converge and hence that the $p(f|\mathfrak{E}_j, \mathbf{y})$ ’s are separately suboptimal, this does not preclude the *mean* of (7) having good MSE performance. It also does not address coverage of the approximate credible intervals.

Restriction to discrete covariates. We assumed that our covariate distribution was discrete, i.e. $\mathcal{X} = \{1, 2, \dots, b\}^d$. Many real datasets, of course, contain continuous features. In this case, in-practice BART computes a grid of quantile values for each continuous feature, and selects a split threshold only from among these values. Although this effectively makes the covariate space discrete, it also means that the space varies with the training sample size n . Our proof relies heavily on the uniform concentration of all node and split-based quantities across a finite set and so does not automatically generalize to this setting.

Failure to address dependence on other DGP parameters. We did not analyze the dependence of our hitting time lower bounds on d , the dimension of the feature space, b , the number of categories for each discrete feature, s , the sparsity of the true regression function, and ν , the covariate distribution.

Failure to address dependence on other algorithmic parameters. We did not analyze the dependence of our hitting time lower bounds on $p_{\mathfrak{E}}$, the prior on $\Omega_{\text{TSE}, m}$, the move probabilities $\pi_g, \pi_p, \pi_c, \pi_s$, and the variance parameters σ^2 and λ .

⁸Computing the marginal posterior involves solving a d -dimensional linear regression whereas computing the conditional posterior involves solving a univariate linear regression.

Asymptotic nature of results. The hitting time lower bounds hold only when the training sample size is larger than a minimum number N . This number depends on the DGP and algorithmic parameters, and if tracked, can be exponentially large in some of them.

8 Simulations

In this section, we describe the results of a simulation study designed to bridge some of the theoretical limitations raised in the previous section. Specifically, we directly study the RMSE of the posterior mean, with respect to the true regression function, and the empirical coverage of pointwise posterior credible intervals for the true regression function. We do this for the output of the original BART algorithm, investigating how they vary according to various DGP and algorithmic parameters. Our experiments show the following:

- When the data is generated from an additive model, RMSE and coverage for BART improve with the number of trees, even when the number of trees is larger than the number of additive components;
- When using Bayesian CART, the root split is often chosen suboptimally and yet is rarely reversed. The probability that this root split is reversed decreases as the training sample size increases;
- Across a wide variety of real and synthetic DGPs, RMSE and coverage improve when averaging the results of multiple BART sampler chains. The relative improvement gap becomes increasingly pronounced as the training sample size increases.

The first and second results validate our hitting time lower bounds (Theorem 5.1 and 5.3 respectively) and suggest that they hold for the original BART algorithm and for reasonable training sample sizes. Since averaging results from multiple chains should not make a difference if each chain is well-mixed, the third result suggests that the tendency of mixing and hitting times to increase with the number of training samples is a fairly general phenomenon that holds across a wide variety of DGPs.

Code availability. All the code necessary to reproduce the experiments in this section is publicly available at github.com/theo-s/bart-hitting-time-sims. The computing infrastructure used was a Linux cluster managed by Department of Statistics at UC Berkeley. Most runs of the simulation used a single 24-core node with 128 GB of RAM, while the larger datasets required a large-memory node with 792 GB RAM and 96 cores.

Algorithm settings and hyperparameters. We use the `dbarts` R package (Dorie, 2022) with almost all hyperparameters kept at their default values. In particular, $\pi_g = \pi_p = 0.25$, $\pi_c = 0.1$, $\pi_s = 0.4$, and the number of posterior samples is $ndpost=1000$. The only exception is that we increase the number of burn-in iterations from 100 to 5000 ($nskip=5000$), in order to highlight that mixing does not occur within a reasonable number of iterations. The responses are centered and scaled to have variance one. The prior for the noise variance is calibrated using the residuals of a linear model fit on all of the predictors.⁹ Unless otherwise noted in the experiment description, the number of trees was kept equal to 200.

Data. For each experiment, we generate a training dataset $\mathcal{D}_n$ comprising n i.i.d. tuples $(\mathbf{x}_i, y_i)$, where $y_i = f^*(\mathbf{x}_i) + \epsilon_i$, $\epsilon \sim \mathcal{N}(0, \eta^2)$. f^* , η^2 , and the covariate distribution will be varied from experiment to experiment.

Evaluation. Both RMSE and empirical coverage are calculated over an independent test set consisting of 1000 points drawn from the same distribution. The results are averaged over 100 experimental replicates, with error bars representing $\pm 1.96SE$. To compute empirical coverage, we count the proportion of the test set points whose credible interval for the function value contains the ground truth.

⁹This is standard for most software implementations of the algorithm.

8.1 Experiment 1: Under an Additive Model, More Trees Improves Performance

Our first experiment studies how RMSE and empirical coverage depends on the interaction of the number of trees specified in the BART model and the number of components in an additive DGP. We chose the regression function to be one of the two forms described below. We varied the number of trees in the grid $\{1, 2, \dots, 10\}$ and the training sample size in the grid $\{10K, 20K, 50K, 100K\}$.

DGP. We let $\mathbf{x}_i \in \mathbb{R}^{10}$, $\mathbf{x}_i \sim \mathcal{N}(0, \Sigma)$, where we set $\Sigma_{ii} = 1$, $\Sigma_{ij} = 0.01$ for $i \neq j$. We let $\eta^2 = 2$, and take f^* to be an additive function with 5 components taking one the following two functional forms:

1. *Linear:* $f^*(\mathbf{x}) = 0.2x_1 - x_2 + 0.6x_3 - 0.9x_4 + 0.85x_5$;
2. *Smooth:* $f^*(\mathbf{x}) = 0.2x_1^2 - x_2 + 0.6 \cos(x_3) - 0.9|x_4|^{1/2} + 0.85 \sin(x_5)$.

Results: The results are displayed in Figure 4 and show that the RMSE for BART decreases while the empirical coverage increases as the number of trees increases, even when the number of trees is larger than the number of additive components. This trend holds over both functional forms and is consistent across different choices of training sample sizes. As discussed in previous sections, the BART posterior already achieves the maximum goodness of fit (i.e. BIC is minimized) when the number of trees is equal to the number of additive components, implying that further improvement in RMSE and empirical coverage as we increase the number of trees arises purely from better mixing. This thus corroborates Theorem 5.1. Furthermore, the decrease in RMSE (and increase in empirical coverage) has a larger slope for larger training sample sizes, indicating that the improvement in mixing is more pronounced in these settings.




Figure 4: When fitted to an additive model, RMSE for BART decreases and coverage increases as the number of trees increases, even when the number of trees is larger than the number of additive components (i.e. $m > 5$.) This suggests that more trees leads to better mixing and thereby corroborates Theorem 5.1. This trend holds over different functional forms and several choices of training sample sizes. Both RMSE and coverage are calculated on an independent test set and are averaged over 100 experimental replicates, with error bars representing $\pm 1.96\text{SE}$.

8.2 Experiment 2: Root Split Gets Stuck on Suboptimal Feature

In the second experiment, we fit Bayesian CART (BART with $n_{tree} = 1$) to DGPs in which the optimal root split is known and study the root split behavior of the sampler across its iterations. We chose the regression function to be one of the two forms described below and varied the training sample size in $\{100, 10K\}$. For each of these choices, we ran the sampler for 500 iterations and counted (i) the percentage of iterations for which the root split was correct as well as (ii) the number of total changes to the root split across the 500 iterations.¹⁰ We repeated this process 100 times.

DGP. We let $\mathbf{x}_i \in \mathbb{R}^{10}$, $\mathbf{x}_i \sim \mathcal{N}(0, \Sigma)$, where we set $\Sigma_{ii} = 1$, $\Sigma_{ij} = 0.01$ for $i \neq j$. We set $\eta = 1$ and take f^* to be a piecewise constant function with a single optimal tree structure, taking one of the following forms (see also Figure 5):

$$1. \text{ DGP 1: } f^*(\mathbf{x}) = \begin{cases} 12, & \text{if } x_1 > .1 \wedge x_2 > .7 \\ 9, & \text{if } x_1 > .1 \wedge x_2 \leq .7 \\ 3, & \text{if } x_1 \leq .1 \wedge x_3 > -.2 \\ 6, & \text{if } x_1 \leq .1 \wedge x_3 \leq -.2 \end{cases}$$

$$2. \text{ DGP 2: } f^*(\mathbf{x}) = \begin{cases} 12, & \text{if } x_1 > .1 \wedge x_2 > .7 \wedge x_4 > .1 \\ 9, & \text{if } x_1 > .1 \wedge x_2 > .7 \wedge x_4 \leq .1 \\ 3, & \text{if } x_1 > .1 \wedge x_2 \leq .7 \wedge x_5 > .1 \text{ or } x_1 \leq .1 \wedge x_3 > -.2 \wedge x_6 > .2 \\ 4, & \text{if } x_1 > .1 \wedge x_2 \leq .7 \wedge x_5 \leq .1 \text{ or } x_1 \leq .1 \wedge x_3 > -.2 \wedge x_6 \leq .2 \\ 2, & \text{if } x_1 \leq .1 \wedge x_3 \leq -.2 \wedge x_7 > .15 \\ 1, & \text{if } x_1 \leq .1 \wedge x_3 \leq -.2 \wedge x_7 \leq .15 \end{cases}$$

In both cases the optimal first split is on x_1 .




Figure 5: The two regression functions used in Experiment 2 to explore root split behavior of Bayesian CART.

Results. The distributions of both quantities across the 100 replicates are displayed in Figure 6. They show that, as the training sample size increases, the percentage of correct first splits does not converge to 1, but instead seems to develop a bimodal distribution. Moreover, the number of changes to the root split decreases. This suggests that the sampler becomes more likely to get stuck in a set comprising trees with an incorrect root split, which agrees with Theorem 5.3. What is especially notable is that while Theorem 5.3 requires a restriction to only “grow” and “prune” moves, these results do not have such a restriction.

¹⁰We did not discard any burn-in samples for this investigation.




Figure 6: When the DGP has a unique optimal root split, as the training sample size increases, the Bayesian CART sampler becomes more likely to get stuck in a set comprising trees with incorrect root splits. For each of 100 experimental replicates, we initialize a sampler at the trivial tree and run it for 500 iterations. We then calculate (i) the percentage of sampler iterations with a correct root split and (ii) the total number of changes to the root split. The top and bottom panels plot the distribution of the first and second quantities respectively. As the training sample size increases, the distribution for (i) seems to become bimodal around 0 and 1. Meanwhile, (ii) becomes more concentrated around 0.

8.3 Experiment 3: Multiple BART Chains Improves Mixing

In the final experiment, we investigate the effect of training sample size on the mixing performance of BART as we fit it to a variety of DGPs, described below. We varied the number of training samples in the grid $\{1K, 10K, 20K, 50K, 100K\}$ and, for each dataset, varied the number of chains for the BART sampler in $\{1, 2, 5, 10\}$. For each number of chains, we divided a fixed budget of 1000 posterior samples evenly amongst the chains and obtained RMSE and empirical coverage values by averaging the posterior samples thus obtained. Note that we still use the same number of burn-in iterations for each chain.

DGP. We study the following DGPs:

1. *Low* (Lei and Candès, 2021): $f^*(\mathbf{x}) = g(x_1)g(x_2)$, $g(x) = \frac{2}{1+\exp\{-12(x-0.5)\}}$. We let $\mathbf{x}_i \in \mathbb{R}^{10}$, $\mathbf{x}_i \sim \mathcal{N}(0, \Sigma)$, where we set $\Sigma_{ii} = 1$, $\Sigma_{ij} = 0.01$ for $i \neq j$.
2. *High* (Lei and Candès, 2021): $f^*(\mathbf{x}) = g(x_1)g(x_2)$, $g(x) = \frac{2}{1+\exp\{-12(x-0.5)\}}$. We let $\mathbf{x}_i \in \mathbb{R}^{100}$, $\mathbf{x}_i \sim \mathcal{N}(0, \Sigma)$, where we set $\Sigma_{ii} = 1$, $\Sigma_{ij} = 0.01$ for $i \neq j$.
3. *Local Sparse Spiky* (Behr et al., 2021): $f^*(\mathbf{x}) = 2 \cdot \mathbf{1}\{x_1 < 0, x_3 > 0\} - 3 \cdot \mathbf{1}\{x_5 > 0, x_6 > 1\} + 8 \cdot \mathbf{1}\{x_3 < 1.5, x_5 < 1\}$. We let $\mathbf{x}_i \in \mathbb{R}^{10}$, $\mathbf{x}_i \sim \mathcal{N}(0, \Sigma)$, where we set $\Sigma_{ii} = 1$, $\Sigma_{ij} = 0.01$ for $i \neq j$.
4. *Piecewise Linear* (Künzel et al., 2019): $f^*(\mathbf{x}) = \begin{cases} \beta_1^T \mathbf{x}, & \text{if } x_{20} < -.4 \\ \beta_2^T \mathbf{x}, & \text{if } x_{20} < .4 \\ \beta_3^T \mathbf{x}, & \text{otherwise} \end{cases}$, where $\beta_1, \beta_2, \beta_3 \sim \text{Unif}([-15, 15]^{20})$.

We let $\mathbf{x}_i \in \mathbb{R}^{20}$, $\mathbf{x}_i \sim \mathcal{N}(0, \Sigma)$, where we set $\Sigma_{ii} = 1$, $\Sigma_{ij} = 0.01$ for $i \neq j$.

5. *Sum*: $f^*(\mathbf{x}) = \sum_{j=1}^{10} x_j$. We let $\mathbf{x}_i \in \mathbb{R}^{20}$, $\mathbf{x}_i \sim \mathcal{N}(0, \Sigma)$, where we set $\Sigma_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0.01 & \text{if } i = j + 10 \\ 0 & \text{otherwise} \end{cases}$.

6. *Tree*: $f^*(\mathbf{x}) = \mathfrak{T}(\mathbf{x})$, where $\mathfrak{T}$ is a decision tree function fitted to a standard Gaussian response vector, using the CART algorithm with maximal depth of 7. We let $\mathbf{x}_i \in \mathbb{R}^{10}$, $\mathbf{x}_i \sim \mathcal{N}(0, \Sigma)$, where we set $\Sigma_{ii} = 1$, $\Sigma_{ij} = 0.01$ for $i \neq j$.

We set η such that *Tree* and *Piecewise* have signal to noise ratios of 1, *High* and *Low* have signal to noise ratios of .5, and *Local Sparse Spiky* and *Sum* have signal to noise ratios of .75.

Results. The results are displayed in Figure 7. Note that instead of plotting RMSE, we have chosen to plot relative RMSE, which measures the ratio between the RMSE obtained from multiple chains and that obtained for a single chain for a given data setting. The results show that both RMSE and empirical coverage improves, often quite significantly, as we add multiple chains. This means that different chains give rise to significantly different distributions, implying that the BART sampler has not mixed even after the large number of burn-in iterations. Furthermore, as the number of training samples increases, the relative performance gap between a single chain and multiple chains increases. Our results therefore provide evidence that the tendency of HPDR hitting times to grow with training sample size is consistent across a wider range of DGPs than was studied theoretically.

Real data simulations. We also performed a similar experiment with a number of benchmark datasets. Results and further description for this experiment are provided in Appendix K.




Figure 7: The RMSE (top panel) and empirical coverage (bottom panel) of BART improve if we average posterior samples from multiple sampler chains, given a fixed total budget of posterior samples. This improvement means that the BART sampler has not mixed ever after 5000 burn-in iterations. Furthermore, the relative performance gaps increases with the number of training samples, providing evidence that the tendency of HPDR hitting time to grow with training sample size is consistent across a wide range of DGPs. Both RMSE and coverage are calculated on an independent test set and are averaged over 100 experimental replicates, with error bars representing $\pm 1.96SE$.

9 Discussion

It is widely accepted that Chipman et al. (2010)’s BART sampler often has issues with mixing and that there is much room for computational improvement. In this paper, we vastly improve upon our earlier work (Ronen et al., 2022) to provide theoretical computational lower bounds for a slightly modified version of the BART sampler. Ours is the first work to analyze BART rather than Bayesian CART. Furthermore, we create a new framework for analysis by studying hitting times of HPDRs instead of mixing times, which resolves issues of identifiability and leads to more meaningful computational lower bounds. We derive these lower bounds under four different, fairly realistic algorithmic and data generative settings and show in all cases that they grow with the training sample size. We complement our theoretical results with a simulation study that validates our results and also suggests that our central thesis, that BART mixing and hitting times increase with the number of training samples, is a fairly general phenomenon that holds across a wide variety of DGPs. We also argue that this is due to the unnecessarily coarse way in which the BART sampler relates temperature and training sample size.

Our results give theoretical and empirical support to some of the choices that BART practitioners often already make. Specifically, they support the use of more trees in the BART ensemble, and they support the use of multiple BART sampler chains. How to select the optimal number of trees and chains is an intriguing and important question and will be left to future work. In addition, our results advocate for the design of better BART samplers and suggest possible approaches for improvement. First, we believe that there is great potential in exploring various forms of temperature control. This could take the form of simulated annealing (Van Laarhoven et al., 1987), or simulated tempering (Marinari and Parisi, 1992), which has previously been explored for Bayesian CART (Angelopoulos and Cussens, 2005), but has yet to be adapted to BART. Second, we believe that the proposal distribution should favor more promising split directions, instead of being uniform at random. There is now a vast literature on how using gradient and even Hessian information can help to accelerate MCMC in continuous state spaces (see for instance Neal et al. (2011)), and there is recent work in extending this to discrete spaces (Zanella, 2020). Third, we believe instead of constraining the proposal distribution to “local” moves, it could benefit from incorporating moves that alter tree structures more drastically. This has been explored somewhat by Kim and Rockova (2023).

Acknowledgements

We would like to thank Antonio Linero and Jungguem Kim for insightful discussions. We would also like to thank the Statistical Computing Facility (SCF) of the Department of Statistics at University of California, Berkeley for computing support, and in particular, Jacob Steinhardt for access to his group’s large-memory nodes within the SCF Linux cluster. We gratefully acknowledge partial support from NSF TRIPODS Grant 1740855, DMS-1613002, 1953191, 2015341, 2209975, 20241842, IIS 1741340, ONR grant N00014-17-1-2176, NSF grant 2023505 on Collaborative Research: Foundations of Data Science Institute (FODSI), NSF grant MC2378 to the Institute for Artificial CyberThreat Intelligence and OperationN (ACTION), the NSF and the Simons Foundation for the Collaboration on the Theoretical Foundations of Deep Learning through awards DMS-2031883 and 814639, and a Weill Neurohub grant. YT was partially supported by NUS Start-up Grant A-8000448-00-00.

References

- Abbe, E., Adsera, E. B., and Misiakiewicz, T. (2022). The merged-staircase property: a necessary and nearly sufficient condition for sgd learning of sparse functions on two-layer neural networks. In *Conference on Learning Theory*, pages 4782–4887. PMLR.
- Agarwal, A., Tan, Y. S., Ronen, O., Singh, C., and Yu, B. (2022). Hierarchical shrinkage: Improving the accuracy and interpretability of tree-based models. In *International Conference on Machine Learning*, pages 111–135. PMLR.
- Angelopoulos, N. and Cussens, J. (2005). Tempering for Bayesian CART. In *Proceedings of the 22nd International Conference on Machine Learning*, pages 17–24.
- Bartlett, P. L., Long, P. M., Lugosi, G., and Tsigler, A. (2020). Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070.
- Behr, M., Wang, Y., Li, X., and Yu, B. (2021). Provable boolean interaction recovery from tree ensemble obtained via random forests. *arXiv preprint arXiv:2102.11800*.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.
- Breiman, L., Friedman, J., Olshen, R., and Stone, C. J. (1984). *Classification and regression trees*. Chapman and Hall/CRC.
- Carnegie, N. B. (2019). Comment: Contributions of model features to BART causal inference performance using ACIC 2016 competition data. *Statistical Science*, 34(1):90 – 93.
- Caruana, R., Karampatziakis, N., and Yessenalina, A. (2008). An empirical evaluation of supervised learning in high dimensions. In *Proceedings of the 25th International Conference on Machine Learning*, pages 96–103.
- Caruana, R. and Niculescu-Mizil, A. (2006). An empirical comparison of supervised learning algorithms. In *Proceedings of the 23rd International Conference on Machine Learning*, pages 161–168.
- Castillo, I. and Ročková, V. (2021). Uncertainty quantification for Bayesian CART. *The Annals of Statistics*, 49(6):3482–3509. Publisher: Institute of Mathematical Statistics.
- 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*, pages 785–794.
- 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.
- Chipman, H. A., George, E. I., and McCulloch, R. E. (2010). BART: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298.
- Dorie, V. (2022). dbarts: Discrete Bayesian additive regression trees sampler. R package version 0.9-22.
- Dorie, V., Hill, J., Shalit, U., Scott, M., and Cervone, D. (2019). Automated versus do-it-yourself methods for causal inference: Lessons learned from a data analysis competition. *Statistical Science*, 34(1):43–68.
- Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004). Least angle regression. *The Annals of Statistics*, 32(2):407–499.
- Fernández-Delgado, M., Cernadas, E., Barro, S., and Amorim, D. (2014). Do we need hundreds of classifiers to solve real world classification problems? *The Journal of Machine Learning Research*, 15(1):3133–3181.
- Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, 29(5):1189–1232.

- Green, D. P. and Kern, H. L. (2010). Modeling heterogeneous treatment effects in large-scale experiments using Bayesian additive regression trees. In *The Annual Summer Meeting of the Society of Political Methodology*, pages 100–110.
- Green, D. P. and Kern, H. L. (2012). Modeling heterogeneous treatment effects in survey experiments with Bayesian additive regression trees. *Public Opinion Quarterly*, 76(3):491–511.
- Grinsztajn, L., Oyallon, E., and Varoquaux, G. (2022). Why do tree-based models still outperform deep learning on typical tabular data? *Advances in Neural Information Processing Systems*, 35:507–520.
- Hahn, P. R., Dorie, V., and Murray, J. S. (2019). Atlantic causal inference conference (acic) data analysis challenge 2017. *arXiv preprint arXiv:1905.09515*.
- Hastie, T. and Tibshirani, R. (1986). Generalized additive models. *Statistical Science*, 1(3):297–318.
- Hastie, T., Tibshirani, R., Friedman, J. H., and Friedman, J. H. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, volume 2. Springer.
- He, J. and Hahn, P. R. (2021). Stochastic tree ensembles for regularized nonlinear regression. *Journal of the American Statistical Association*, pages 1–20.
- Hill, J., Linero, A., and Murray, J. (2020). Bayesian additive regression trees: A review and look forward. *Annual Review of Statistics and Its Application*, 7(1).
- Hill, J. L. (2011). Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240.
- Jeong, S. and Rockova, V. (2020). The art of BART: On flexibility of Bayesian forests. *arXiv preprint arXiv:2008.06620*, 3(69):146.
- Kern, H. L., Stuart, E. A., Hill, J., and Green, D. P. (2016). Assessing methods for generalizing experimental impact estimates to target populations. *Journal of Research on Educational Effectiveness*, 9(1):103–127.
- Kim, J. and Rockova, V. (2023). On mixing rates for Bayesian CART. *arXiv preprint arXiv:2306.00126*.
- Klusowski, J. M. (2021). Universal consistency of decision trees in high dimensions. *arXiv preprint arXiv:2104.13881*.
- Künzel, S. R., Sekhon, J. S., Bickel, P. J., and Yu, B. (2019). Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165.
- Lei, L. and Candès, E. J. (2021). Conformal inference of counterfactuals and individual treatment effects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.
- Levin, D. A., Peres, Y., and Wilmer, E. L. (2006). *Markov chains and mixing times*. American Mathematical Society.
- Linero, A. R. (2018). Bayesian regression trees for high-dimensional prediction and variable selection. *Journal of the American Statistical Association*, 113(522):626–636.
- Linero, A. R. and Yang, Y. (2018). Bayesian regression tree ensembles that adapt to smoothness and sparsity. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(5):1087–1110.
- Luo, H. and Pratola, M. T. (2023). Sharded Bayesian additive regression trees. *arXiv preprint arXiv:2306.00361*.
- Marinari, E. and Parisi, G. (1992). Simulated tempering: a new monte carlo scheme. *Europhysics Letters*, 19(6):451.
- Mazumder, R. and Wang, H. (2024). On the convergence of CART under sufficient impurity decrease condition. *Advances in Neural Information Processing Systems*, 36.

- Nash, W. J., Sellers, T. L., Talbot, S. R., Cawthorn, A. J., and Ford, W. B. (1994). The population biology of abalone (haliotis species) in tasmania. i. blacklip abalone (h. rubra) from the north coast and islands of bass strait. *Sea Fisheries Division, Technical Report*, 48:p411.
- Neal, R. M. et al. (2011). MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2(11):2.
- Pace, K. and Barry, R. (1997). Sparse spatial autoregressions. *Statistics & Probability Letters*, 33(3):291–297.
- Pratola, M. T. (2016). Efficient metropolis–hastings proposal mechanisms for Bayesian regression tree models. *Bayesian analysis*, 11(3):885–911.
- Pratola, M. T., Chipman, H. A., Gattiker, J. R., Higdon, D. M., McCulloch, R., and Rust, W. N. (2014). Parallel Bayesian Additive Regression Trees. *Journal of Computational and Graphical Statistics*, 23(3):830–852. Publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/10618600.2013.841584>.
- Pratola, M. T., Chipman, H. A., George, E. I., and McCulloch, R. E. (2020). Heteroscedastic BART via multiplicative regression trees. *Journal of Computational and Graphical Statistics*, 29(2):405–417.
- Redner, R. A. and Walker, H. F. (1984). Mixture densities, maximum likelihood and the EM algorithm. *SIAM Review*, 26(2):195–239.
- Rockova, V. and Rousseau, J. (2021). Ideal Bayesian spatial adaptation. *arXiv preprint arXiv:2105.12793*.
- Ročková, V. and Saha, E. (2019). On theory for BART. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2839–2848. PMLR.
- Ročková, V. and van der Pas, S. (2020). Posterior concentration for Bayesian regression trees and forests. *The Annals of Statistics*, 48(4):2108–2131.
- Romano, J. D., Le, T. T., La Cava, W., Gregg, J. T., Goldberg, D. J., Chakraborty, P., Ray, N. L., Himmelstein, D., Fu, W., and Moore, J. H. (2021). PMLB v1.0: an open-source dataset collection for benchmarking machine learning methods. *Bioinformatics*, 38(3):878–880.
- Ronen, O., Saarinen, T., Tan, Y. S., Duncan, J., and Yu, B. (2022). A mixing time lower bound for a simplified version of BART. *arXiv preprint arXiv:2210.09352*.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464.
- Scornet, E., Biau, G., and Vert, J.-P. (2015). Consistency of random forests. *Annals of Statistics*, 43(4):1716–1741.
- Starling, J. E., Murray, J. S., Carvalho, C. M., Bukowski, R. K., and Scott, J. G. (2020). BART with targeted smoothing: An analysis of patient-specific stillbirth risk. *The Annals of Applied Statistics*, 14(1):28–50.
- Syrkanis, V. and Zampetakis, M. (2020). Estimation and inference with trees and forests in high dimensions. In *Conference on Learning Theory*, pages 3453–3454. PMLR.
- Tan, Y. S., Agarwal, A., and Yu, B. (2022a). A cautionary tale on fitting decision trees to data from additive models: generalization lower bounds. In *International Conference on Artificial Intelligence and Statistics*, pages 9663–9685. PMLR.
- Tan, Y. S., Singh, C., Nasser, K., Agarwal, A., and Yu, B. (2022b). Fast interpretable greedy-tree sums (FIGS). *arXiv preprint arXiv:2201.11931*.
- Van Laarhoven, P. J., Aarts, E. H., van Laarhoven, P. J., and Aarts, E. H. (1987). *Simulated Annealing*. Springer.
- Vershynin, R. (2018). *High-Dimensional Probability: An Introduction with Applications in Data Science*, volume 47. Cambridge University Press.

- Wainwright, M. J. (2019). *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, volume 48. Cambridge University Press.
- Wendling, T., Jung, K., Callahan, A., Schuler, A., Shah, N. H., and Gallego, B. (2018). Comparing methods for estimation of heterogeneous treatment effects using observational data from health care databases. *Statistics in medicine*, 37(23):3309–3324.
- Wu, Y., Tjelmeland, H., and West, M. (2007). Bayesian CART: Prior specification and posterior simulation. *Journal of Computational and Graphical Statistics*, 16(1):44–66.
- Yeager, D. S., Hanselman, P., Walton, G. M., Murray, J. S., Crosnoe, R., Muller, C., Tipton, E., Schneider, B., Hulleman, C. S., Hinojosa, C. P., et al. (2019). A national experiment reveals where a growth mindset improves achievement. *Nature*, 573(7774):364–369.
- Zanella, G. (2020). Informed proposals for local MCMC in discrete spaces. *Journal of the American Statistical Association*, 115(530):852–865.

A Literature summary for BART posterior concentration results

We list here all results on posterior concentration that we are aware of:

- Posterior concentration at the optimal $n^{-\alpha/(2\alpha+p)}$ rate for BCART and BART when the regression function is Hölder α -smooth, $0 < \alpha \leq 1$ (Ročková and Saha, 2019).
- Posterior concentration at the optimal $n^{-\alpha/(2\alpha+s)}$ rate for BCART and BART when the regression function is Hölder α -smooth, $0 < \alpha \leq 1$, and s -sparse (Ročková and van der Pas, 2020).
- Posterior concentration at the almost optimal $n^{-\alpha/(2\alpha+1)} \log n$ rate for BART when the regression function is an additive sum of univariate components, each of which is Hölder α -smooth, $0 < \alpha \leq 1$ (Ročková and van der Pas, 2020).
- Extensions of the above results to $\alpha > 1$, provided BART is modified to allow for “soft” splits (Linero and Yang, 2018).
- Posterior concentration at the optimal rate (up to log factors) for BART when the regression function is piecewise heterogeneous anisotropic Hölder smooth (Jeong and Rockova, 2020; Rockova and Rousseau, 2021).

B Proofs for Section 4.2

We first introduce some additional notation that will be used throughout the rest of the appendix. We denote $D_f = \|f^*\|_\infty$ and let D_ϵ denote the sub-Gaussian parameter of the noise random variable ϵ (Wainwright, 2019). C will be used to denote a universal constant (i.e. not depending on any parameters) that may change from line to line. For simplicity, many of our results are written using big- O notation, which will indicate leading order dependence on the sample size n as well as the error probability δ .

B.1 Main Proofs

Lemma B.1 (Concentration of empirical risk difference). *Let $\mathfrak{E}$ and $\mathfrak{E}'$ be two partition ensemble models. Then for any $0 < \delta < 1$, with probability at least $1 - \delta$, the risk difference between them satisfies*

$$\left| \mathbf{y}^T (\mathbf{P}_{\mathfrak{E}} - \mathbf{P}_{\mathfrak{E}'}) \mathbf{y} - n \int (\Pi_{\mathfrak{E}} f^*)^2 - (\Pi_{\mathfrak{E}'} f^*)^2 d\nu \right| = O\left(\sqrt{n \log(1/\delta)}\right). \quad (8)$$

If furthermore, $\Pi_{\mathfrak{E}} f^* = \Pi_{\mathfrak{E}'} f^*$, then the above bound can be improved to

$$|\mathbf{y}^T (\mathbf{P}_{\mathfrak{E}} - \mathbf{P}_{\mathfrak{E}'}) \mathbf{y}| = O(\log(1/\delta)). \quad (9)$$

Proof of Proposition 4.1 Since

$$\Delta \text{BIC}(\mathfrak{E}, \mathfrak{E}') = \frac{\mathbf{y}^T (\mathbf{P}_{\mathfrak{E}} - \mathbf{P}_{\mathfrak{E}'}) \mathbf{y}}{\sigma^2} + (\text{df}(\mathfrak{E}) - \text{df}(\mathfrak{E}')) \log n,$$

the desired concentration follows immediately from Lemma B.1. $\square$

Lemma B.2 (Log marginal likelihood formula). *Let $\mathfrak{E}$ be a tree ensemble structure (TSE), and let Ψ be an $n \times b$ matrix whose columns comprise the indicators of the leaves in $\mathfrak{E}$. The log marginal likelihood satisfies*

$$\begin{aligned} -2 \log p(\mathbf{y} | \mathbf{X}, \mathfrak{E}) &= n \log(2\pi\sigma^2) + \log \det \left(\lambda^{-1} \Psi^T \Psi + \mathbf{I} \right) \\ &\quad + \frac{1}{\sigma^2} \left(\|\Psi \hat{\boldsymbol{\mu}}_{LS} - \mathbf{y}\|_2^2 + \hat{\boldsymbol{\mu}}_{LS}^T \Psi^T \left(\mathbf{I} - \Psi \left(\Psi^T \Psi + \lambda \mathbf{I} \right)^{-1} \Psi^T \right) \Psi \hat{\boldsymbol{\mu}}_{LS} \right), \end{aligned} \quad (10)$$

where

$$\hat{\boldsymbol{\mu}}_{LS} := \underset{\boldsymbol{\mu}}{\operatorname{argmin}} \|\boldsymbol{\Psi}\boldsymbol{\mu} - \mathbf{y}\|_2^2$$

is the solution to the least squares problem.

Lemma B.3 (Concentration of log-determinant). *With the notation of Lemma B.2, denote $\hat{\boldsymbol{\Sigma}} := \frac{1}{n}\boldsymbol{\Psi}^T\boldsymbol{\Psi}$ and $\boldsymbol{\Sigma} := \mathbb{E}\{\hat{\boldsymbol{\Sigma}}\}$. Then for any $0 < \delta < 1$, for $n \geq \max\{64m^{3/2}s_{\min}^{-2}\log(2\operatorname{df}(\mathfrak{E})/\delta), 4\lambda s_{\min}^{-1}\}$, with probability at least $1 - \delta$, we have*

$$\left| \log \det \left(\lambda^{-1}\boldsymbol{\Psi}^T\boldsymbol{\Psi} + \mathbf{I} \right) - \operatorname{df}(\mathfrak{E}) \log(n) - \sum_{i=1}^{\operatorname{df}(\mathfrak{E})} \log(s_i/\lambda) \right| = O\left(\sqrt{\frac{\log(1/\delta)}{n}}\right), \quad (11)$$

where $s_1, s_2, \dots, s_{\operatorname{df}(\mathfrak{E})}$ denote the values of the nonzero eigenvalues of $\boldsymbol{\Sigma}$.

Lemma B.4 (Concentration of error term). *With the notation of Lemma B.3, for any $0 < \delta < 1$, when $n \geq 16m^{3/2}s_{\min}^{-2}\log(2\operatorname{df}(\mathfrak{E})/\delta)$, where $s_{\min}$ is the minimum nonzero eigenvalue of $\boldsymbol{\Sigma}$, with probability at least $1 - \delta$, we have*

$$\hat{\boldsymbol{\mu}}_{LS}^T \boldsymbol{\Psi}^T \left(\mathbf{I} - \boldsymbol{\Psi} \left(\boldsymbol{\Psi}^T \boldsymbol{\Psi} + \lambda \mathbf{I} \right)^{-1} \boldsymbol{\Psi}^T \right) \boldsymbol{\Psi} \hat{\boldsymbol{\mu}}_{LS} = O(1). \quad (12)$$

Proof of Proposition 4.2. Starting with equation (10) from Lemma B.2, plug in equations (11) and (12) from Lemma B.2 and Lemma B.3 respectively. Notice that $\boldsymbol{\Psi}\hat{\boldsymbol{\mu}}_{LS} = \mathbf{P}_{\mathfrak{E}}\mathbf{y}$, so that

$$\|\boldsymbol{\Psi}\hat{\boldsymbol{\mu}}_{LS} - \mathbf{y}\|_2^2 = \|(\mathbf{I} - \mathbf{P}_{\mathfrak{E}})\mathbf{y}\|_2^2 = \mathbf{y}^T(\mathbf{I} - \mathbf{P}_{\mathfrak{E}})\mathbf{y}.$$

This completes the proof. $\square$

Proof of Proposition 4.3. Let $\mathfrak{E}^* \in \operatorname{OPT}_m(f^*, 0)$. We will show that for any $\mathfrak{E} \notin \operatorname{OPT}_m(f^*, 0)$, there is some $N_{\mathfrak{E}}$ such that

$$p(\mathfrak{E}|\mathbf{y}) \leq \frac{\epsilon p(\mathfrak{E}^*|\mathbf{y})}{|\Omega_{\text{TSE},m}|} \quad (13)$$

with probability at least $1 - \delta/|\Omega_{\text{TSE},m}|$ for all $n \geq N_{\mathfrak{E}}$. If this is true, set $N = \max_{\mathfrak{E} \notin \operatorname{OPT}_m(f^*, 0)} N_{\mathfrak{E}}$ and take $n \geq N$. On the intersection of all these events, we have

$$p(\operatorname{OPT}_m(f^*, 0)^c | \mathbf{y}) = \sum_{\mathfrak{E} \notin \operatorname{OPT}_m(f^*, 0)} p(\mathfrak{E}|\mathbf{y}) \leq \sum_{\mathfrak{E} \notin \operatorname{OPT}_m(f^*, 0)} \frac{\epsilon p(\mathfrak{E}^*|\mathbf{y})}{|\Omega_{\text{TSE},m}|} \leq \epsilon.$$

To prove (13), fix $\mathfrak{E}$. Using Proposition 4.2 and Proposition 4.1, we get a probability $1 - \delta/|\Omega_{\text{TSE},m}|$ event on which

$$\begin{aligned} & \log p(\mathfrak{E}^*|\mathbf{y}) - \log p(\mathfrak{E}|\mathbf{y}) \\ &= \begin{cases} \frac{n}{2\sigma^2} (\operatorname{Bias}^2(\mathfrak{E}; f^*)) + O\left(\sqrt{n \log(1/\delta)} + \log(1/\delta)\right) & \text{if } \operatorname{Bias}^2(\mathfrak{E}; f^*) \neq 0 \\ \frac{\log n}{2} (\operatorname{df}(\mathfrak{E}) - \operatorname{df}(\mathfrak{E}^*)) + O(\log(1/\delta)) & \text{otherwise.} \end{cases} \end{aligned}$$

In either case, we get $\frac{p(\mathfrak{E}|\mathbf{y})}{p(\mathfrak{E}^*|\mathbf{y})} \rightarrow 0$ as $n \rightarrow \infty$. $\square$

B.2 Further Details

Proof of Lemma B.1. Recall that $\Pi_{\mathfrak{E}}$ refers to orthogonal projection onto $\mathcal{F}(\mathfrak{E})$ in $L^2(\nu)$, while $\mathbf{P}_{\mathfrak{E}}$ refers to orthogonal projection onto $\mathcal{F}(\mathfrak{E})$ with respect to the empirical norm $\|\cdot\|_n$. With this in mind, decompose $y =$

$\epsilon + (f^*(x) - \Pi_{\mathfrak{E}} f^*(x)) + \Pi_{\mathfrak{E}} f^*(x)$ and write this in vector form as $\mathbf{y} = \epsilon + (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*) + \mathbf{f}_{\mathfrak{E}}^*$. Since $\Pi_{\mathfrak{E}} f^* \in \mathcal{F}(\mathfrak{E})$, we have $\mathbf{P}_{\mathfrak{E}} \mathbf{f}_{\mathfrak{E}}^* = \mathbf{f}_{\mathfrak{E}}^*$. We can then therefore expand the quadratic form $\mathbf{y}^T \mathbf{P}_{\mathfrak{E}} \mathbf{y}$ as follows:

$$\begin{aligned} \mathbf{y}^T \mathbf{P}_{\mathfrak{E}} \mathbf{y} &= \epsilon^T \mathbf{P}_{\mathfrak{E}} \epsilon + (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{P}_{\mathfrak{E}} (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*) + 2\epsilon^T \mathbf{P}_{\mathfrak{E}} (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*) + (\mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{P}_{\mathfrak{E}} \mathbf{f}_{\mathfrak{E}}^* + 2\epsilon^T \mathbf{P}_{\mathfrak{E}} \mathbf{f}_{\mathfrak{E}}^* + 2(\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{P}_{\mathfrak{E}} \mathbf{f}_{\mathfrak{E}}^* \\ &= \epsilon^T \mathbf{P}_{\mathfrak{E}} \epsilon + (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{P}_{\mathfrak{E}} (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*) + 2\epsilon^T \mathbf{P}_{\mathfrak{E}} (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*) + (\mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{f}_{\mathfrak{E}}^* + 2\epsilon^T \mathbf{f}_{\mathfrak{E}}^* + 2(\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{f}_{\mathfrak{E}}^*. \end{aligned} \quad (14)$$

Note that ϵ , $(f^*(x) - \Pi_{\mathfrak{E}} f^*(x))$, and $\Pi_{\mathfrak{E}} f^*(x)$ are uncorrelated random variables, with ϵ being also independent of the other two variables. This implies that the third, fifth and sixth terms in (14) have zero mean. On the other hand, because of finite sample fluctuations, $(\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)$ and $\mathbf{f}_{\mathfrak{E}}^*$, are not necessarily orthogonal as vectors.

To bound the expectation of (14), first observe that $f^* - \Pi_{\mathfrak{E}} f^*$ and $\Pi_{\mathfrak{E}} f^*$ are bounded random variables and thus have both standard deviation and sub-Gaussian norm bounded by D_f (Wainwright, 2019). We then compute

$$\begin{aligned} &\mathbb{E}\{\epsilon^T \mathbf{P}_{\mathfrak{E}} \epsilon + (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{P}_{\mathfrak{E}} (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*) + 2\epsilon^T \mathbf{P}_{\mathfrak{E}} (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*) + (\mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{f}_{\mathfrak{E}}^* + 2\epsilon^T \mathbf{f}_{\mathfrak{E}}^* + 2(\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{f}_{\mathfrak{E}}^*\} \\ &= \mathbb{E}\{\epsilon^T \mathbf{P}_{\mathfrak{E}} \epsilon\} + \mathbb{E}\{(\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{P}_{\mathfrak{E}} (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)\} + \mathbb{E}\{(\mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{f}_{\mathfrak{E}}^*\} \\ &= \text{tr}\{\mathbf{P}_{\mathfrak{E}}\}(\text{Var}\{\epsilon\} + \text{Var}\{f^* - f_{\mathfrak{E}}^*\}) + n \int (\Pi_{\mathfrak{E}} f^*)^2 d\nu \\ &\leq C(D_{\epsilon} + D_f) \text{df}(\mathfrak{E}) + n \int (\Pi_{\mathfrak{E}} f^*)^2 d\nu. \end{aligned} \quad (15)$$

Next, we bound the fluctuations of each term in (14) separately. Using the Hanson-Wright inequality (Wainwright, 2019; Vershynin, 2018), we get $1 - \delta$ probability events over which

$$|\epsilon^T \mathbf{P}_{\mathfrak{E}} \epsilon - \mathbb{E}\{\epsilon^T \mathbf{P}_{\mathfrak{E}} \epsilon\}| \leq CD_{\epsilon} \max\left\{\log(1/\delta), \sqrt{\text{df}(\mathfrak{E}) \log(1/\delta)}\right\}, \quad (16)$$

and

$$|(\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{P}_{\mathfrak{E}} (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*) - \mathbb{E}\{(\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{P}_{\mathfrak{E}} (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)\}| \leq CD_f \max\left\{\log(1/\delta), \sqrt{\text{df}(\mathfrak{E}) \log(1/\delta)}\right\}. \quad (17)$$

Using Hoeffding's inequality (Wainwright, 2019), we have further $1 - \delta$ probability events over which

$$\left|(\mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{f}_{\mathfrak{E}}^* - n \int (\Pi_{\mathfrak{E}} f^*)^2 d\nu\right| \leq CD_f^2 \sqrt{n \log(1/\delta)}, \quad (18)$$

$$|\epsilon^T \mathbf{f}_{\mathfrak{E}}^*| \leq CD_f D_{\epsilon} \sqrt{n \log(1/\delta)}, \quad (19)$$

$$|(\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{f}_{\mathfrak{E}}^*| \leq CD_f^2 \sqrt{n \log(1/\delta)}. \quad (20)$$

For the third term in (14), we use Cauchy-Schwarz followed by Young's inequality to get

$$\begin{aligned} 2|\epsilon^T \mathbf{P}_{\mathfrak{E}} (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)| &\leq 2(\epsilon^T \mathbf{P}_{\mathfrak{E}} \epsilon)^{1/2} ((\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{P}_{\mathfrak{E}} (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*))^{1/2} \\ &\leq \epsilon^T \mathbf{P}_{\mathfrak{E}} \epsilon + (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{P}_{\mathfrak{E}} (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*). \end{aligned} \quad (21)$$

Conditioning on all the events guaranteeing (16) to (20) and plugging in these bounds together with (15) into (14), we get

$$\begin{aligned} &\left|\mathbf{y}^T \mathbf{P}_{\mathfrak{E}} \mathbf{y} - n \int (\Pi_{\mathfrak{E}} f^*)^2 d\nu\right| \\ &\leq C\left((D_{\epsilon} + D_f) \text{df}(\mathfrak{E}) + (D_f + D_{\epsilon}) \log(1/\delta) + D_f(D_f + D_{\epsilon}) \sqrt{n \log(1/\delta)} + (D_f + D_{\epsilon}) \sqrt{\text{df}(\mathfrak{E}) \log(1/\delta)}\right) \\ &= O\left(\sqrt{n \log(1/\delta)} + \log(1/\delta)\right). \end{aligned}$$

Repeating the same argument for $\mathfrak{E}'$ and adjusting δ so that the intersection of all events conditioned on has probability at least $1 - \delta$ completes the proof of (8).

To prove (9), observe that under the additional assumption, we can cancel terms in (14) to get

$$\mathbf{y}^T (\mathbf{P}_{\mathfrak{E}} - \mathbf{P}_{\mathfrak{E}'}) \mathbf{y} = \epsilon^T (\mathbf{P}_{\mathfrak{E}} - \mathbf{P}_{\mathfrak{E}'}) \epsilon + (\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*)^T (\mathbf{P}_{\mathfrak{E}} - \mathbf{P}_{\mathfrak{E}'})(\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*) + 2\epsilon^T (\mathbf{P}_{\mathfrak{E}} - \mathbf{P}_{\mathfrak{E}'})(\mathbf{f}^* - \mathbf{f}_{\mathfrak{E}}^*).$$

Applying (21) followed by (16) and (17) completes the proof.

Proof of Lemma B.2. Recall that the full log likelihood satisfies

$$p(\mathbf{y}|\mathbf{X}, \mathfrak{E}, \boldsymbol{\mu}) = (2\pi\sigma^2)^{-n/2} \exp\left(-\frac{\|\boldsymbol{\Psi}\boldsymbol{\mu} - \mathbf{y}\|_2^2}{2\sigma^2}\right),$$

while conditioned on $\mathfrak{E}$, the prior on $\boldsymbol{\mu}$ satisfies

$$p(\boldsymbol{\mu}|\mathfrak{E}) = (2\pi\sigma^2\lambda^{-1})^{-b/2} \exp\left(-\frac{\lambda\|\boldsymbol{\mu}\|_2^2}{2\sigma^2}\right),$$

where b is the number of columns in $\boldsymbol{\Psi}$. Hence

$$p(\mathbf{y}|\mathbf{X}, \mathfrak{E}, \boldsymbol{\mu})p(\boldsymbol{\mu}|\mathfrak{E}) = (2\pi\sigma^2)^{-n/2} (2\pi\sigma^2\lambda^{-1})^{-b/2} \exp\left(-\frac{1}{2\sigma^2} \left(\|\boldsymbol{\Psi}\boldsymbol{\mu} - \mathbf{y}\|_2^2 + \lambda\|\boldsymbol{\mu}\|_2^2\right)\right). \quad (22)$$

Consider the orthogonal decomposition

$$\|\boldsymbol{\Psi}\boldsymbol{\mu} - \mathbf{y}\|_2^2 = \|\boldsymbol{\Psi}\hat{\boldsymbol{\mu}}_{LS} - \mathbf{y}\|_2^2 + \|\boldsymbol{\Psi}(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{LS})\|_2^2.$$

We next add the second term on the right to the exponent in the prior and complete the square:

$$\begin{aligned} & \|\boldsymbol{\Psi}(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{LS})\|_2^2 + \lambda\|\boldsymbol{\mu}\|_2^2 \\ &= \boldsymbol{\mu}^T (\boldsymbol{\Psi}^T \boldsymbol{\Psi} + \lambda \mathbf{I}) \boldsymbol{\mu} - 2(\boldsymbol{\Psi}\hat{\boldsymbol{\mu}}_{LS})^T \boldsymbol{\mu} + \hat{\boldsymbol{\mu}}_{LS}^T \boldsymbol{\Psi}^T \boldsymbol{\Psi} \hat{\boldsymbol{\mu}}_{LS} \\ &= (\boldsymbol{\mu} - \boldsymbol{\mu}_0)^T (\boldsymbol{\Psi}^T \boldsymbol{\Psi} + \lambda \mathbf{I}) (\boldsymbol{\mu} - \boldsymbol{\mu}_0) - \hat{\boldsymbol{\mu}}_{LS}^T \boldsymbol{\Psi}^T \boldsymbol{\Psi} (\boldsymbol{\Psi}^T \boldsymbol{\Psi} + \lambda \mathbf{I})^{-1} \boldsymbol{\Psi}^T \boldsymbol{\Psi} \hat{\boldsymbol{\mu}}_{LS} + \hat{\boldsymbol{\mu}}_{LS}^T \boldsymbol{\Psi}^T \boldsymbol{\Psi} \hat{\boldsymbol{\mu}}_{LS}. \end{aligned} \quad (23)$$

where

$$\boldsymbol{\mu}_0 = (\boldsymbol{\Psi}^T \boldsymbol{\Psi} + \lambda \mathbf{I})^{-1} \boldsymbol{\Psi}^T \boldsymbol{\Psi} \hat{\boldsymbol{\mu}}_{LS}.$$

The constant term in (23) is

$$\begin{aligned} & -\hat{\boldsymbol{\mu}}_{LS}^T \boldsymbol{\Psi}^T \boldsymbol{\Psi} (\boldsymbol{\Psi}^T \boldsymbol{\Psi} + \lambda \mathbf{I})^{-1} \boldsymbol{\Psi}^T \boldsymbol{\Psi} \hat{\boldsymbol{\mu}}_{LS} + \hat{\boldsymbol{\mu}}_{LS}^T \boldsymbol{\Psi}^T \boldsymbol{\Psi} \hat{\boldsymbol{\mu}}_{LS} \\ &= \hat{\boldsymbol{\mu}}_{LS}^T \boldsymbol{\Psi}^T \left(\mathbf{I} - \boldsymbol{\Psi} (\boldsymbol{\Psi}^T \boldsymbol{\Psi} + \lambda \mathbf{I})^{-1} \boldsymbol{\Psi}^T \right) \boldsymbol{\Psi} \hat{\boldsymbol{\mu}}_{LS}. \end{aligned} \quad (24)$$

Plugging (24) back into (22) and integrating, we get

$$\begin{aligned} p(\mathbf{y}|\mathbf{X}, \mathfrak{E}) &= \int p(\mathbf{y}|\mathbf{X}, \mathfrak{E}, \boldsymbol{\mu}) p(\boldsymbol{\mu}|\mathfrak{E}) d\boldsymbol{\mu} \\ &= (2\pi\sigma^2)^{-n/2} \exp\left(-\frac{\|\boldsymbol{\Psi}\hat{\boldsymbol{\mu}}_{LS} - \mathbf{y}\|_2^2 + \hat{\boldsymbol{\mu}}_{LS}^T \boldsymbol{\Psi}^T \left(\mathbf{I} - \boldsymbol{\Psi} (\boldsymbol{\Psi}^T \boldsymbol{\Psi} + \lambda \mathbf{I})^{-1} \boldsymbol{\Psi}^T \right) \boldsymbol{\Psi} \hat{\boldsymbol{\mu}}_{LS}}{2\sigma^2}\right) \\ &\quad \cdot \int (2\pi\sigma^2\lambda^{-1})^{-b/2} \exp\left(-\frac{(\boldsymbol{\mu}^T - \boldsymbol{\mu}_0)^T (\boldsymbol{\Psi}^T \boldsymbol{\Psi} + \lambda \mathbf{I}) (\boldsymbol{\mu}^T - \boldsymbol{\mu}_0)}{2\sigma^2}\right) d\boldsymbol{\mu}. \end{aligned} \quad (25)$$

By a change of variables, the integral can be computed as

$$\int (2\pi\sigma^2\lambda^{-1})^{-b/2} \exp\left(-\frac{(\boldsymbol{\mu}^T - \boldsymbol{\mu}_0)^T (\lambda^{-1} \boldsymbol{\Psi}^T \boldsymbol{\Psi} + \mathbf{I}) (\boldsymbol{\mu}^T - \boldsymbol{\mu}_0)}{2\sigma^2\lambda^{-1}}\right) d\boldsymbol{\mu} = \det(\lambda^{-1} \boldsymbol{\Psi}^T \boldsymbol{\Psi} + \mathbf{I})^{-1/2}.$$

Plugging this back into (25) and taking logarithms yields (10).

Proof of Lemma B.3. Let $\hat{s}_1, \hat{s}_2, \dots, \hat{s}_b$ be the eigenvalues of $\hat{\Sigma}$. Using Lemma B.5, we have $\hat{s}_i = 0$ for any $i > \text{df}(\mathfrak{E})$. We may therefore compute

$$\begin{aligned} \log \det \left(\lambda^{-1} \Psi^T \Psi + \mathbf{I} \right) &= \sum_{i=1}^{\text{df}(\mathfrak{E})} \log(n\lambda^{-1}\hat{s}_i + 1) \\ &= \sum_{i=1}^{\text{df}(\mathfrak{E})} \log\left(\frac{\hat{s}_i + \lambda/n}{s_i}\right) + \text{df}(\mathfrak{E}) \log n + \sum_{i=1}^{\text{df}(\mathfrak{E})} \log(s_i/\lambda). \end{aligned}$$

It remains to bound the first term. To this end, we first condition on the $1 - \delta$ probability event guaranteed by Lemma B.6. Then, we observe that

$$\begin{aligned} \left| \frac{\hat{s}_i + \lambda/n}{s_i} - 1 \right| &\leq \frac{1}{s_i} \left(|\hat{s}_i - s_i| + \frac{\lambda}{n} \right) \\ &\leq \frac{1}{s_{\min}} \left(\|\hat{\Sigma} - \Sigma\| + \frac{\lambda}{n} \right) \\ &\leq \frac{1}{s_{\min}} \left(\max \left\{ \sqrt{\frac{4m^{3/2} \log(2\text{df}(\mathfrak{E})/\delta)}{n}}, \frac{4\sqrt{m} \log(2\text{df}(\mathfrak{E})/\delta)}{n} \right\} + \frac{\lambda}{n} \right). \end{aligned} \quad (26)$$

Here, the second inequality makes use of Weyl's inequality.

Recall the elementary inequality

$$|\log x| \leq 2|x - 1|$$

for $0 < x < 1/2$. Using this together with (26), we get

$$\sum_{i=1}^{\text{df}(\mathfrak{E})} \log\left(\frac{\hat{s}_i + \lambda/n}{s_i}\right) = O\left(\sqrt{\frac{\log(1/\delta)}{n}}\right)$$

when $n \geq \max\{64m^{3/2}s_{\min}^{-2} \log(2\text{df}(\mathfrak{E})/\delta), 4\lambda s_{\min}^{-1}\}$. □

Proof of Lemma B.4. Recall that $\Psi \hat{\mu}_{LS} = \mathbf{P}_{\mathfrak{E}} \mathbf{y}$. We thus rewrite and bound the error term as

$$\begin{aligned} \hat{\mu}_{LS}^T \Psi^T \left(\mathbf{I} - \Psi \left(\Psi^T \Psi + \lambda \mathbf{I} \right)^{-1} \Psi^T \right) \Psi \hat{\mu}_{LS} &= \mathbf{y}^T \mathbf{P}_{\mathfrak{E}} \mathbf{M} \mathbf{P}_{\mathfrak{E}} \mathbf{y} \\ &\leq \|\mathbf{P}_{\mathfrak{E}} \mathbf{y}\|_2^2 \|\mathbf{M}\|, \end{aligned} \quad (27)$$

where

$$\mathbf{M} = \mathbf{P}_{\mathfrak{E}} - \Psi \left(\Psi^T \Psi + \lambda \mathbf{I} \right)^{-1} \Psi^T.$$

Using similar arguments as in the proof of Lemma B.1, we bound

$$\begin{aligned} \|\mathbf{P}_{\mathfrak{E}} \mathbf{y}\|_2^2 &= (\mathbf{f}_{\mathfrak{E}}^*)^T \mathbf{f}_{\mathfrak{E}}^* + \epsilon^T \mathbf{P}_{\mathfrak{E}} \epsilon \\ &\leq D_f^2 \left(n + \sqrt{n \log(1/\delta)} \right) + D_{\epsilon}^2 \left(n + \sqrt{\text{df}(\mathfrak{E}) \log(1/\delta)} \right). \end{aligned} \quad (28)$$

Meanwhile, the nonzero eigenvalues of $\mathbf{M}$ are of the form

$$1 - \frac{\hat{s}_i}{\hat{s}_i + \lambda/n} = \frac{\lambda}{n\hat{s}_i + \lambda}$$

for $i = 1, 2, \dots, \text{df}(\mathfrak{E})$. These can be further bounded as

$$\begin{aligned} \frac{\lambda}{n\hat{s}_i + \lambda} &\leq \frac{\lambda}{n(s_i - |\hat{s}_i - s_i|)} \\ &\leq \frac{\lambda}{n(s_{\min} - \|\hat{\Sigma} - \Sigma\|)}. \end{aligned} \quad (29)$$

Taking $n \geq 16m^{3/2}s_{\min}^{-2} \log(2\text{df}(\mathfrak{E})/\delta)$ and conditioning on the $1 - \delta$ probability event guaranteed by Lemma B.6, we can further bound (29) by $\lambda/2ns_{\min}$, which gives $\|\mathbf{M}\| = O(n^{-1})$. Combining this with (28) and plugging them back into (27) finishes the proof. $\square$

Lemma B.5 (Rank of empirical covariance matrix). *With the notation of Lemma B.3, we have $\text{rank}(\hat{\Sigma}) \leq \text{rank}(\Sigma)$.*

Proof. Let $\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_n$ denote the rows of Ψ , noting that they are i.i.d. random vectors. Note that if $\Sigma \mathbf{v} = 0$ for some $\mathbf{v}$, this implies that $\text{Cov}\{\mathbf{v}^T \mathbf{l}\} = \mathbf{v}^T \Sigma \mathbf{v} = 0$, and so $\mathbf{v}^T \mathbf{l} \equiv 0$ as a random variable. In particular, we have $\mathbf{v}^T \mathbf{l}_i = 0$ for $i = 1, 2, \dots, n$, and we also get $\hat{\Sigma} \mathbf{v} = 0$. As such, the nullspace for Σ is contained within the nullspace for $\hat{\Sigma}$. The conclusion follows. $\square$

Lemma B.6 (Concentration of empirical covariance matrix). *With the notation of Lemma B.3, for any $0 < \delta < 1$, with probability at least $1 - \delta$, we have*

$$\left\| \frac{1}{n} \Psi^T \Psi - \Sigma \right\| \leq \max \left\{ \sqrt{\frac{4m^{3/2} \log(2\text{df}(\mathfrak{E})/\delta)}{n}}, \frac{4\sqrt{m} \log(2\text{df}(\mathfrak{E})/\delta)}{n} \right\}.$$

Proof. Let $\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_n$ denote the rows of Ψ as before. Since each point can only be contained in a single leaf on each tree, we have $\|\mathbf{l}_j\|_2 = \sqrt{m}$, while Σ also satisfies $\|\Sigma\| \leq m$. Using Corollary 6.20¹¹ in Wainwright (2019), we therefore have

$$\left\| \frac{1}{n} \Psi^T \Psi - \Sigma \right\| \leq 2\text{df}(\mathfrak{E}) \exp\left(-\frac{nt^2}{2\sqrt{m}(m+t)}\right)$$

for any $t > 0$. Rearranging this equation completes the proof. $\square$

C Background on Markov Chains

For the whole of this section, let $X_0, X_1, \dots$ be an irreducible and aperiodic discrete time Markov chain on a finite state space Ω , with stationary distribution π .

C.1 Networks and Voltages

Harmonic functions. Let P be the transition kernel of (X_t) . We call a function $h: \Omega \rightarrow \mathbb{R}$ harmonic for P at a state x if

$$h(x) = \sum_{y \in \Omega} P(x, y) h(y). \quad (30)$$

Lemma C.1 (Uniqueness of harmonic extensions, Proposition 9.1. in Levin et al. (2006)). *Let $\mathcal{A} \subset \Omega$ be a subset of the state space. Let $h_{\mathcal{A}}: \mathcal{A} \rightarrow \mathbb{R}$ be a function defined on $\mathcal{A}$. The function $h: \Omega \rightarrow \mathbb{R}$ defined by $h(x) := \mathbb{E}\{h_{\mathcal{A}}(X_{\tau_{\mathcal{A}}}) | X_0 = x\}$ is the unique extension of $h_{\mathcal{A}}$ such that $h(x) = h_{\mathcal{A}}(x)$ for all $x \in \mathcal{A}$ and h is harmonic for P at all $x \in \Omega \setminus \mathcal{A}$.*

¹¹While Corollary 6.20 is stated with the assumption that the rows have mean zero, the proof in Wainwright (2019) illustrates that this is unnecessary.

The values of a harmonic function can be computed by solving the system of linear equations given by (30) for each $x \in \Omega \setminus \mathcal{A}$. This is hard to do directly by hand for complicated state spaces, but when the Markov chain is symmetric, i.e. the stationary distribution satisfies $\pi(x)P(x, y) = \pi(y)P(y, x)$ for any x, y , we can use several operations to *simplify the state space while preserving the values of the harmonic function on the remaining state space*. Indeed, harmonic functions are equivalent to voltages on electrical circuits, and it is well-known how to simplify circuits in order to calculate voltages:

1. (Gluing) Points on the circuit with the same voltage can be joined.
2. (Series law) Two resistors in series with resistances r_1 and r_2 can be merged with the new resistor having resistance $r_1 + r_2$.
3. (Parallel law) Two resistors in parallel with resistances r_1 and r_2 can be merged with the new resistor having resistance $1/(1/r_1 + 1/r_2)$.

We make this connection rigorous by introducing the following definitions and via the subsequent lemmas.

Networks, conductance, resistance. A *network* (Ω, c) is a tuple comprising a finite state space Ω and a symmetric function $c: \Omega \times \Omega \rightarrow \mathbb{R}_+$ called the *conductance*. The *resistance* function is defined as $r(x, y) = \frac{1}{c(x, y)}$, and can take the value of positive infinity. We say that $\{x, y\}$ is an *edge* in the network if $c(x, y) > 0$. Any network (Ω, c) has an associated Markov chain whose transition probabilities are defined by $P(x, y) = \frac{c(x, y)}{\sum_{z \in \Omega} c(x, z)}$.

Voltage and current flow. We say that a function is harmonic on the network if it is harmonic with respect to P . Given $a, z \in \Omega$, a *voltage* W between a and z is a function that is harmonic on $\Omega \setminus \{a, z\}$. The *current flow* $I: \Omega \times \Omega \rightarrow \mathbb{R}$ associated with W is defined as $I(x, y) = c(x, y)(W(x) - W(y))$. The strength of the current flow is defined as

$$\|I\| := \sum_{y \in \Omega} c(a, y)(W(a) - W(y)).$$

Effective resistance and effective conductance. The *effective resistance* between a and z is defined as

$$R(a \leftrightarrow z) := \frac{W(a) - W(z)}{\|I\|},$$

noting that this is independent of the choice of W by the uniqueness property in Lemma C.1. The *effective conductance* between a and z is defined as $C(a \leftrightarrow z) := 1/R(a \leftrightarrow z)$.

Lemma C.2 (Network simplification rules). *Consider a network (Ω, c) . Define the following operations that each produces a modified network (Ω', c') .*

1. (Gluing) Given $u, v \in \Omega$, define $\Omega' := \Omega \setminus \{v\}$ and

$$c'(x, y) = \begin{cases} c(x, y) & x, y \neq u, \\ c(x, u) + c(x, v) & x \neq u, y = u, \\ c(u, y) + c(v, y) & x = u, y \neq u \\ c(u, v) + c(u, y) + c(u, x) & x = y = u. \end{cases}$$

2. (Parallel and series laws) Given $u, v, w \in \Omega$ with $c(v, x) = 0$ for all $x \notin \{u, w\}$, define $\Omega' := \Omega \setminus \{v\}$ and

$$c'(x, y) = \begin{cases} c(u, w) + \frac{c(u, v)c(v, w)}{c(u, v) + c(v, w)} & (x, y) = (u, w) \text{ or } (w, u), \\ c(x, y) & \text{otherwise.} \end{cases}$$

Consider a function h that is harmonic on $\Omega \setminus \mathcal{A}$. The following hold:

1. If we glue states $u, v \in \Omega$ such that $h(u) = h(v)$, then h remains harmonic on $\Omega' \setminus \mathcal{A}$ with respect to the modified transition matrix P' .
2. If we apply the parallel and series laws to $u, v, w \in \Omega$ with $v \notin \mathcal{A}$, then h remains harmonic on $\Omega' \setminus \mathcal{A}$ with respect to the modified transition matrix P' .

Furthermore, if h is a voltage between two states $a, z \in \Omega$, then applying the operations does not change the effective conductance and resistance between them.

Proof. The first statement is obvious as the mean value equation for harmonic functions can be repeated almost verbatim. For the second statement, we just have to check the mean value equation for h at u . This is equivalent to the equation

$$\sum_{x \in \Omega'} c'(u, x)(h(x) - h(u)) = 0. \quad (31)$$

First note that under the original network, our assumption on $c(v, x)$ and the mean value property at v gives

$$c(v, w)(h(w) - h(v)) = c(u, v)(h(v) - h(u)). \quad (32)$$

Next, we compute

$$\begin{aligned} c'(u, w)(h(w) - h(u)) &= \left(c(u, w) + \frac{c(u, v)c(v, w)}{c(u, v) + c(v, w)} \right) (h(w) - h(u)) \\ &= c(u, w)(h(w) - h(u)) + \frac{c(u, v)c(v, w)}{c(u, v) + c(v, w)} ((h(w) - h(v)) + (h(v) - h(u))) \\ &= c(u, w)(h(w) - h(u)) + \frac{c(u, v)^2 + c(u, v)c(v, w)}{c(u, v) + c(v, w)} (h(v) - h(u)) \\ &= c(u, w)(h(w) - h(u)) + c(v, w)(h(v) - h(u)), \end{aligned} \quad (33)$$

where the third equality follows from (32). We therefore have

$$\sum_{x \in \Omega'} c'(u, x)(h(x) - h(u)) = \sum_{x \in \Omega} c(u, x)(h(x) - h(u)),$$

and the mean value equation at u for the original network implies (31).

Finally, to conclude invariance of effective conductance, we observe that it is defined in terms of voltages and current flows. We have already shown that voltages are unchanged, so we just need to argue that the strength of the current flow is similarly unchanged. This is immediate whenever $a \notin \{u, w\}$. When $a = u$, this follows from (33). $\square$

Lemma C.3 (Rayleigh's Monotonicity Law, Theorem 9.12 in Levin et al. (2006)). *Given a network with two resistance functions $r, r' : \Omega \times \Omega \rightarrow \mathbb{R}_+ \cup \{\infty\}$, suppose that $r \leq r'$ pointwise. Then we have*

$$R(a \leftrightarrow z; r) \leq R(a \leftrightarrow z; r').$$

C.2 Hitting Precedence Probabilities

Hitting precedence probabilities. Let $\mathcal{A}, \mathcal{B} \subset \Omega$ be disjoint subsets. The *hitting precedence probability* of $\mathcal{A}$ relative to $\mathcal{B}$ is defined as the following function on Ω :

$$\text{HPP}(x; \mathcal{A}, \mathcal{B}) := P\{\tau_{\mathcal{A}} < \tau_{\mathcal{B}} \mid X_0 = x\}.$$

Lemma C.4. $\text{HPP}(-; \mathcal{A}, \mathcal{B})$ is a harmonic function on $\Omega \setminus (\mathcal{A} \cup \mathcal{B})$.

Proof. Write $h_{\mathcal{A} \cup \mathcal{B}}(z) = \mathbf{1}\{z \in \mathcal{A}\}$. It is easy to see that

$$h_{\mathcal{A} \cup \mathcal{B}}(X_{\tau_{\mathcal{A} \cup \mathcal{B}}}) = \begin{cases} 1 & \tau_{\mathcal{A}} < \tau_{\mathcal{B}} \\ 0 & \tau_{\mathcal{A}} > \tau_{\mathcal{B}}. \end{cases}$$

Hence,

$$\mathbb{E}\{h_{\mathcal{A} \cup \mathcal{B}}(X_{\tau_{\mathcal{A} \cup \mathcal{B}}}) | X_0 = x\} = \mathbb{P}\{\tau_{\mathcal{A}} < \tau_{\mathcal{B}} \mid X_0 = x\} = \text{HPP}(x; \mathcal{A}, \mathcal{B}).$$

By Lemma C.1, the left hand side is a harmonic function on $\Omega \setminus (\mathcal{A} \cup \mathcal{B})$. $\square$

Lemma C.5 (HPP from a bottleneck state). *Given a network (Ω, c) with states a, x, z and such that $c(u, z) = 0$ for all $u \notin \{x, z\}$. Let $x = x_0, x_1, \dots, x_k = a$ be any sequence of states such that there is some $\rho > 0$ for which $c(x_{i-1}, x_i) \geq \rho^{-1}c(x, z)$ for $i = 1, \dots, k$. Then*

$$\mathbb{P}\{\tau_a < \tau_z | X_0 = x\} \geq \frac{1}{k\rho + 1}.$$

Proof. Let (Ω, c') be the modified network in which we set

$$c'(u, v) = \begin{cases} c(u, v) & \{u, v\} \in \{\{x_{i-1}, x_i\} : i = 1, \dots, k\} \cup \{\{x, z\}\} \\ 0 & \text{otherwise.} \end{cases}$$

Then we have

$$\begin{aligned} \mathbb{P}\{\tau_a < \tau_z | X_0 = x\} &= \frac{r(x, z)}{R(a \leftrightarrow x; r) + r(x, z)} \\ &\geq \frac{r(x, z)}{R(a \leftrightarrow x; r') + r(x, z)}, \end{aligned} \tag{34}$$

where the equality uses Lemma C.6 and the inequality uses Lemma C.3. Next, using the series law from Lemma C.2, we have

$$\begin{aligned} R(a \leftrightarrow x; r') &= \sum_{i=1}^k r(x_{i-1}, x_i) \\ &\leq k\rho r(x, z). \end{aligned}$$

Plugging this into (34) and cancelling $r(x, z)$ in the numerator and denominator completes the proof. $\square$

Lemma C.6. *Given a network (Ω, c) with states a, x, z and such that $c(u, z) = 0$ for all $u \notin \{x, z\}$. Then we have*

$$P\{\tau_a < \tau_z \mid X_0 = x\} = \frac{C(a \leftrightarrow x)}{C(a \leftrightarrow x) + c(x, z)}.$$

Proof. Let $h(y) := \mathbb{P}\{\tau_a < \tau_z | X_0 = y\}$, and note that h is a voltage between a and z . As such, we have

$$R(a \leftrightarrow z) = \frac{h(a) - h(z)}{\|I\|} = \frac{1}{\|I\|}.$$

On the other hand, h is also a voltage between a and x on the reduced state space $\Omega \setminus \{z\}$, which gives

$$R(a \leftrightarrow x) = \frac{h(a) - h(x)}{\|I\|}.$$

Finally, by the series law, we have

$$R(a \leftrightarrow z) = R(a \leftrightarrow x) + r(x, z).$$

Putting everything together, we get

$$\begin{aligned}
P\{\tau_a < \tau_z | X_0 = x\} &= 1 - (h(a) - h(x)) \\
&= 1 - R(a \leftrightarrow x) \|I\| \\
&= 1 - \frac{R(a \leftrightarrow x)}{R(a \leftrightarrow z)} \\
&= \frac{r(x, z)}{R(a \leftrightarrow x) + r(x, z)} \\
&= \frac{C(a \leftrightarrow x)}{C(a \leftrightarrow x) + c(x, z)},
\end{aligned}$$

as we wanted. $\square$

D Proof of Theorem 5.2

We will first present the proofs for Theorem 5.2 and Theorem 5.3 because they are relatively simple compared to that of Theorem 5.1. For convenience, we repeat the relevant constructions and definitions here.

Reachability. Given $\mathfrak{E}, \mathfrak{E}' \in \Omega_{\text{TSE}, m}$, we say that $\mathfrak{E} \preceq \mathfrak{E}'$ if $\mathfrak{E}$ and $\mathfrak{E}'$ are connected by an edge, and if either $\text{Bias}^2(\mathfrak{E}; f^*) > \text{Bias}^2(\mathfrak{E}'; f^*)$ or $\text{Bias}^2(\mathfrak{E}; f^*) = \text{Bias}^2(\mathfrak{E}'; f^*)$ and $\text{df}(\mathfrak{E}) \geq \text{df}(\mathfrak{E}')$. Note that, because we allow only “grow” and “prune” moves, for adjacent $\mathfrak{E}$ and $\mathfrak{E}'$, $\mathcal{F}(\mathfrak{E})$ and $\mathcal{F}(\mathfrak{E}')$ are nested subspaces, so that $\text{Bias}^2(\mathfrak{E}; f^*) = \text{Bias}^2(\mathfrak{E}'; f^*)$ if and only if $\Pi_{\mathcal{F}(\mathfrak{E})}[f^*] = \Pi_{\mathcal{F}(\mathfrak{E}')}[f^*]$. We say that $\mathfrak{E}$ is reachable from $\mathfrak{E}'$, denoted $\mathfrak{E} \succeq \mathfrak{E}'$, if there is a sequence of TSEs $\mathfrak{E} = \mathfrak{E}^0, \mathfrak{E}^1, \dots, \mathfrak{E}^k = \mathfrak{E}'$ such that $\mathfrak{E}^i \preceq \mathfrak{E}^{i+1}$ for $i = 0, 1, \dots, k-1$.

Set-up. Without loss of generality, let (x_1, x_2) be a pure interaction for f^* . Let $\mathfrak{E}_{\text{bad}}$ be any TSE such that

- $\mathfrak{E}_{\text{bad}}$ is reachable from $\mathfrak{E}_\emptyset$;
- There does not exist $\mathfrak{E} \in \Omega_{\text{TSE}, m}$ such that $\mathfrak{E}$ is reachable from $\mathfrak{E}_{\text{bad}}$ but $\mathfrak{E}_{\text{bad}}$ is not reachable from $\mathfrak{E}$.

Note that such a TSE exists because $\Omega_{\text{TSE}, m}$ is finite and $\succeq$ is a partial ordering on this space. By definition, there exists a sequence of TSEs $\mathfrak{E}_\emptyset = \mathfrak{E}^0, \mathfrak{E}^1, \dots, \mathfrak{E}^k = \mathfrak{E}_{\text{bad}}$ such that $\mathfrak{E}^i \preceq \mathfrak{E}^{i+1}$ for $i = 0, 1, \dots, k-1$. We set $\mathcal{A}$ to be the equivalence class of $\mathfrak{E}_{\text{bad}}$ under $\succeq$ and set $\mathcal{B}$ to be the outer boundary of $\mathcal{A}$. For n large enough, using Proposition 4.2 and Proposition 4.1, there is a $1 - \delta/2$ event over which, for $i = 1, 2, \dots, k$,

$$\begin{aligned}
&\log p(\mathfrak{E}^i | \mathbf{y}) - \log p(\mathfrak{E}^{i-1} | \mathbf{y}) \\
&= \begin{cases} \frac{n}{2\sigma^2} (\text{Bias}^2(\mathfrak{E}^{i-1}; f^*) - \text{Bias}^2(\mathfrak{E}^i; f^*)) + O\left(\sqrt{n \log(k/\delta)}\right) & \text{if } \text{Bias}^2(\mathfrak{E}^{i-1}; f^*) > \text{Bias}^2(\mathfrak{E}^i; f^*) \\ \frac{\log n}{2} (\text{df}(\mathfrak{E}^{i-1}) - \text{df}(\mathfrak{E}^i)) + O(\log(k/\delta)) & \text{otherwise.} \end{cases}
\end{aligned}$$

In either case, we get

$$\frac{p(\mathfrak{E}^i | \mathbf{y})}{p(\mathfrak{E}^{i-1} | \mathbf{y})} = \Omega(1). \tag{35}$$

Using Proposition 4.1 again, there is a further $1 - \delta/2$ probability event over which $\Delta \text{BIC}(\mathfrak{E}, \mathfrak{E}_{\text{bad}})$ satisfies either (5) or (6) simultaneously for all $\mathfrak{E} \in \mathcal{B}$ (after dividing the δ that appears in the formulas by $|\mathcal{B}|$). Condition on these two events.

Hitting precedence probability lower bound. Using Lemma D.1 and Lemma D.2, we see that $\mathfrak{E}^i \notin \text{OPT}_m(f^*, \infty)$ for $i = 0, 1, \dots, k$. We therefore have

$$\begin{aligned} P\{\tau_{\mathfrak{E}_{\text{bad}}} < \tau_{\text{OPT}_m(f^*, \infty)}\} &\geq P\{\mathfrak{E}_i = \mathfrak{E}^i \text{ for } i = 1, 2, \dots, k\} \\ &= \prod_{i=1}^k P(\mathfrak{E}^{i-1}, \mathfrak{E}^i). \end{aligned} \quad (36)$$

It suffices to show that $P(\mathfrak{E}^{i-1}, \mathfrak{E}^i)$ is bounded from below by a constant. To see this, we note that

$$\begin{aligned} P(\mathfrak{E}^{i-1}, \mathfrak{E}^i) &= Q(\mathfrak{E}^{i-1}, \mathfrak{E}^i) \min\left\{\frac{Q(\mathfrak{E}^i, \mathfrak{E}^{i-1})p(\mathfrak{E}^i|\mathbf{y})}{Q(\mathfrak{E}^{i-1}, \mathfrak{E}^i)p(\mathfrak{E}^{i-1}|\mathbf{y})}, 1\right\} \\ &= \Omega\left(\min\left\{\frac{p(\mathfrak{E}^i|\mathbf{y})}{p(\mathfrak{E}^{i-1}|\mathbf{y})}\right\}, 1\right) \\ &= \Omega(1), \end{aligned} \quad (37)$$

where the first two inequalities follow because the proposal distributions do not depend on the training sample size n , while the last equality follows from equation (35).

BIC lower bound. Consider $\mathfrak{E}' \in \mathcal{B}$. By definition of $\mathcal{B}$, there exists $\mathfrak{E} \in \mathcal{A}$ such that $\mathfrak{E}$ and $\mathfrak{E}'$ are connected by an edge, but $\mathfrak{E} \not\preceq \mathfrak{E}'$. This implies that either $\text{Bias}^2(\mathfrak{E}; f^*) < \text{Bias}^2(\mathfrak{E}'; f^*)$ or $\text{Bias}^2(\mathfrak{E}; f^*) = \text{Bias}^2(\mathfrak{E}'; f^*)$ and $\text{df}(\mathfrak{E}; f^*) < \text{df}(\mathfrak{E}'; f^*)$. Since $\mathfrak{E}$ and $\mathfrak{E}_{\text{bad}}$ are mutually reachable, we have $\text{Bias}^2(\mathfrak{E}; f^*) = \text{Bias}^2(\mathfrak{E}_{\text{bad}}; f^*)$ and $\text{df}(\mathfrak{E}) = \text{df}(\mathfrak{E}_{\text{bad}})$. We therefore conclude that $\mathfrak{E}'$ either has a larger squared bias or larger degrees of freedom compared to $\mathfrak{E}_{\text{bad}}$. Applying equations (5) and (6) and taking the minimum sample size N large enough gives

$$\Delta\text{BIC}(\mathfrak{E}, \mathfrak{E}_{\text{bad}}) \geq \log n + O(\log(|\mathcal{B}|/\delta)). \quad (38)$$

Conclusion. Applying Proposition 6.1 with equations (37) and (38), we get a $1 - 2\delta$ probability event over which

$$E\{\tau_{\text{OPT}_m(f^*, \infty)}\} = \Omega(n^{1/2}). \quad (39)$$

Lemma D.1. For $i = 0, 1, \dots, k$, no tree in $\mathfrak{E}^i$ contains a split on either x_1 or x_2 .

Proof. Suppose otherwise. By changing the labeling of x_1 and x_2 if necessary, there exists

$$i := \min\{1 \leq j \leq l: \mathfrak{E}_j \text{ contains split on } x_1\}. \quad (40)$$

Since only “grow” and “prune” moves are allowed, $\mathfrak{E}_i$ is obtained from $\mathfrak{E}_{i-1}$ via a “grow” move that splits a leaf node $\mathcal{L}$ into:

$$\mathcal{L}_L := \{\mathbf{x} \in \mathcal{L}: x_i \leq t\}, \quad \mathcal{L}_R := \{\mathbf{x} \in \mathcal{L}: x_i > t\}.$$

Define the function $\psi := \mathbf{1}_{\mathcal{L}_L} - \mathbf{1}_{\mathcal{L}_R}$. Then the span of $\{\mathbf{1}_{\mathcal{L}_L}, \mathbf{1}_{\mathcal{L}_R}\}$ is the same as that of $\{\mathbf{1}_{\mathcal{L}}, \psi\}$, which implies that

$$\mathcal{F}(\mathfrak{E}_i) = \text{span}\{\mathcal{F}(\mathfrak{E}_{i-1}), \psi\}. \quad (41)$$

Furthermore, we have $\psi \notin \mathcal{F}(\mathfrak{E}_{i-1})$ since all functions in $\mathcal{F}(\mathfrak{E}_{i-1})$ do not depend on x_1 . This implies that $\text{df}(\mathfrak{E}_i) = \text{df}(\mathfrak{E}_{i-1}) + 1$. On the other hand, since $x_i \perp\!\!\!\perp y \mid \mathbf{x} \in \mathcal{L}$, we have $\psi \perp y$, which means that $\text{Bias}^2(\mathfrak{E}_i; f^*) = \text{Bias}^2(\mathfrak{E}_{i-1}; f^*)$. As such, we have $\mathfrak{E}_{i-1} \not\preceq \mathfrak{E}_i$, which gives a contradiction. $\square$

Lemma D.2. For any $\mathfrak{E} \in \Omega_{\text{TSE}, m}$, if $\mathfrak{E}$ does not contain splits on both x_1 and x_2 , then $\text{Bias}^2(\mathfrak{E}; f^*) > 0$.

Proof. Since $(x_1, x_2) \not\perp\!\!\!\perp y$, there exist values (a_1, a_2) , (b_1, b_2) and $(c_3, c_4, \dots, c_d)$ such that $f^*(a_1, a_2, c_3, \dots, c_d) \neq f^*(b_1, b_2, c_3, \dots, c_d)$. On the other hand, all functions in $\mathcal{F}(\mathfrak{E})$ are constant with respect to x_1 and x_2 , so that $f^* \notin \mathcal{F}(\mathfrak{E})$. $\square$

E Proof of Theorem 5.3

Set-up. For convenience, we repeat the relevant construction here. Without loss of generality, let x_1 be the feature that gives f^* root dependence. By assumption, there is a threshold t such that splitting the trivial tree on x_1 at t gives a decrease in squared bias. We set

$$\mathcal{A} = \{\mathfrak{T} \in \Omega_{\text{TSE},1} : \mathfrak{T} \text{ has root split on } x_1 \text{ at } t\},$$

and

$$\mathfrak{T}_{\text{bad}} = \arg \min \{\text{BIC}(\mathfrak{E}) : \mathfrak{E} \in \mathcal{A}\}.$$

Note that the outer boundary of $\mathcal{A}$ is a singleton set comprising the trivial tree $\mathfrak{T}_\emptyset$. By assumption, we have $\mathcal{A} \cap \text{OPT}_1(f^*, 0) = \emptyset$.

We continue this construction by denoting $\mathfrak{T}_0 = \mathfrak{T}_\emptyset$ and letting $\mathfrak{T}_1$ be a tree structure comprising a single root split at (x_1, t) . By adding the splits in $\mathfrak{T}_{\text{bad}}$ iteratively, we get a sequence $\mathfrak{T}_1, \mathfrak{T}_2, \dots, \mathfrak{T}_k$ of nested tree structures with $\mathfrak{T}_k = \mathfrak{T}_{\text{bad}}$. By Proposition 4.1 and Proposition 4.2, for any training sample size n large enough, there is a $1 - \delta$ probability event with respect to $\mathbb{P}_n$ such that

$$\log p(\mathfrak{T}_j | \mathbf{y}) - \log p(\mathfrak{T}_0 | \mathbf{y}) = \frac{n}{2\sigma^2} (\text{Bias}^2(\mathfrak{T}_0; f^*) - \text{Bias}^2(\mathfrak{T}_j; f^*)) + O\left(\sqrt{n \log(2k/\delta)}\right) \quad (42)$$

for $j = 1, 2, \dots, k$. Condition on this event.

We note the following:

$$\text{Bias}^2(\mathfrak{T}_0; f^*) = \int (f^*)^2 d\nu - \left(\int f^* d\nu \right)^2 \quad (43)$$

$$\text{Bias}^2(\mathfrak{T}_k; f^*) = 0 \quad (44)$$

$$\text{Bias}^2(\mathfrak{T}_j; f^*) \leq \text{Bias}^2(\mathfrak{T}_1; f^*) < \text{Bias}^2(\mathfrak{T}_0; f^*) \quad \text{for } j = 1, 2, \dots, k. \quad (45)$$

Here, the last statement follows from the fact that

$$\frac{\text{Bias}^2(\mathfrak{T}_0; f^*) - \text{Bias}^2(\mathfrak{T}_1; f^*)}{\text{Bias}^2(\mathfrak{T}_0; f^*)} = \text{Corr}^2(f^*(\mathbf{x}), \mathbf{1}\{x_1 \leq t\}) > 0,$$

and because $\mathfrak{T}_j$ is a refinement of $\mathfrak{T}_1$ for each $j = 1, 2, \dots, k$.

Hitting precedence probability lower bound. We define a conductance function on $\Omega_{\text{TSE},1}$ via

$$c(\mathfrak{T}, \mathfrak{T}') = p(\mathfrak{T} | \mathbf{y}) P(\mathfrak{T}, \mathfrak{T}'),$$

where P is the transition kernel of the BART sampler. It is clear that the Markov chain for Bayesian CART is equivalent to the Markov chain associated with the network $(\Omega_{\text{TSE},1}, c)$. By ignoring quantities that do not depend on the training sample size n , we get

$$\begin{aligned} c(\mathfrak{T}, \mathfrak{T}') &= \min\{p(\mathfrak{T} | \mathbf{y}) Q(\mathfrak{T}, \mathfrak{T}'), p(\mathfrak{T}' | \mathbf{y}) Q(\mathfrak{T}', \mathfrak{T})\} \\ &\asymp \min\{p(\mathfrak{T} | \mathbf{y}), p(\mathfrak{T}' | \mathbf{y})\}. \end{aligned}$$

Consider the set $\mathcal{C} := \Omega \setminus (\mathcal{A} \cup \{\mathfrak{T}_\emptyset\})$. Then $\text{OPT}_1(f^*, 0) \subset \mathcal{C}$ by assumption, and we have $\tau_{\mathcal{C}} \leq \tau_{\text{OPT}_1(f^*, 0)}$. This means that it suffices to consider

$$P\{\tau_{\mathfrak{T}_{\text{bad}}} < \tau_{\mathcal{C}}\} = \text{HPP}(\mathfrak{T}_\emptyset; \mathcal{C}, \mathfrak{T}_{\text{bad}})$$

and to bound it from below. To calculate values for the harmonic function $\text{HPP}(-; \mathcal{C}, \mathfrak{T}_{\text{bad}})$, we use Lemma C.2 to glue all states in $\mathcal{C}$ together without changing the values the function values. We will abuse notation and denote the new glued state using $\mathcal{C}$ while continuing to use c to denote the conductance function on the new state space.

Notice that the only edge from $\mathcal{C}$ connects to $\mathfrak{T}_\emptyset$. As such, we can compute

$$\begin{aligned} c(\mathfrak{T}_\emptyset, \mathcal{C}) &\leq \sum_{\mathfrak{T} \sim \mathfrak{T}_\emptyset} \min\{p(\mathfrak{T}_\emptyset|\mathbf{y})Q(\mathfrak{T}_\emptyset, \mathfrak{T}), p(\mathfrak{T}|\mathbf{y})Q(\mathfrak{T}, \mathfrak{T}_\emptyset)\} \\ &\leq p(\mathfrak{T}_\emptyset|\mathbf{y}) \sum_{\mathfrak{T} \sim \mathfrak{T}_\emptyset} Q(\mathfrak{T}_\emptyset, \mathfrak{T}) \\ &\leq p(\mathfrak{T}_\emptyset|\mathbf{y}). \end{aligned}$$

This implies, using equations (42) and (45), that we can bound the ratios of conductances for $j = 1, 2, \dots, k$ as:

$$\begin{aligned} \log \frac{c(\mathfrak{T}_j, \mathfrak{T}_{j-1})}{c(\mathfrak{T}_0, \mathcal{C})} &\gtrsim \min\{\log p(\mathfrak{T}_j|\mathbf{y}) - \log p(\mathfrak{T}_0|\mathbf{y}), \log p(\mathfrak{T}_{j-1}|\mathbf{y}) - \log p(\mathfrak{T}_0|\mathbf{y})\} \\ &\geq \begin{cases} \frac{n}{2\sigma^2} (\text{Bias}^2(\mathfrak{T}_0; f^*) - \text{Bias}^2(\mathfrak{T}_{j-1}; f^*)) + O\left(\sqrt{n \log(2k/\delta)}\right) & j \geq 2 \\ 0 & j = 1. \end{cases} \end{aligned} \quad (46)$$

By (45), there is some minimum training sample size N so that for all $n \geq N$, the right hand side side of (46) is nonnegative. In this case, the assumptions of Lemma C.5 hold, and we get

$$P\{\tau_{\mathfrak{T}_{\text{bad}}} < \tau_{\mathcal{C}}\} = \Omega(1) \quad (47)$$

as desired.

BIC lower bound. From equations (43) and (44), we have

$$\begin{aligned} \Delta \text{BIC}(\mathfrak{T}_\emptyset, \mathfrak{T}_{\text{bad}}) &= \frac{n}{\sigma^2} (\text{Bias}^2(\mathfrak{T}_\emptyset; f^*) - \text{Bias}^2(\mathfrak{T}_{\text{bad}}; f^*)) + O\left(\sqrt{n \log(2k/\delta)}\right) \\ &= \frac{n}{\sigma^2} \left(\int (f^*)^2 d\nu - \left(\int f^* d\nu \right)^2 \right) + O\left(\sqrt{n \log(2k/\delta)}\right). \end{aligned} \quad (48)$$

Conclusion. Applying Proposition 6.1 with equations (47) and (48), we get a $1 - 2\delta$ probability event over which

$$E\{\tau_{\text{OPT}_1(f^*, 0)}\} = \Omega\left(\exp\left(\frac{n}{2\sigma^2} \left(\int (f^*)^2 d\nu - \left(\int f^* d\nu \right)^2 \right) + O\left(\sqrt{n \log(2k/\delta)}\right)\right)\right).$$

Taking logarithms, dividing by n and applying Markov's inequality gives

$$\mathbb{E}_n \left\{ \frac{\log E\{\tau_{\text{OPT}_1(f^*, 0)}\}}{n} \right\} \geq \frac{(1 - \delta)(1 - O(n^{-1/2}))}{2\sigma^2} \left(\int (f^*)^2 d\nu - \left(\int f^* d\nu \right)^2 \right).$$

Letting $\delta \rightarrow 0$ and taking $n \rightarrow \infty$ finishes the proof.

F Splitting Rules, Local Decision Stumps and Coverage

Before proving Theorem 5.1, we first introduce some required machinery.

Local decision stump basis. Let $\mathfrak{T}$ be a tree structure. Let (v_j, τ_j) , $j = 1, \dots, l$ denote the splits (or splitting rules) on $\mathfrak{T}$ (the labels of its internal nodes). Every node on the tree corresponds to rectangular region $\mathfrak{t} \subset \mathcal{X}$ that is obtained by recursively partitioning the covariate space using the splits further up the tree. If $\mathfrak{t}$ is an internal node, it has two children nodes denoted $\mathfrak{t}_L$ and $\mathfrak{t}_R$ defined by

$$\mathfrak{t}_L := \{\mathbf{x} \in \mathfrak{t} : x_v \leq \tau\}$$

$$\mathbf{t}_R := \{\mathbf{x} \in \mathbf{t} : x_v > \tau\}$$

where (v, τ) is the split on $\mathbf{t}$. For each internal node $(\mathbf{t}_j, v_j, \tau)$, define a *local decision stump function*

$$\psi_j(\mathbf{x}) := \frac{\nu(\mathbf{t}_R)\mathbf{1}\{\mathbf{x} \in \mathbf{t}_L\} - \nu(\mathbf{t}_L)\mathbf{1}\{\mathbf{x} \in \mathbf{t}_R\}}{\sqrt{\nu(\mathbf{t}_L)\nu(\mathbf{t}_R)}}, \quad (49)$$

where $\mathbf{t}_L$ and $\mathbf{t}_R$ denote the children of $\mathbf{t}$. It is easy to check (see for instance Agarwal et al. (2022)) that $\psi_1, \psi_2, \dots, \psi_l$ are orthogonal, and that, together with the constant function $\psi_0 \equiv 1$, form a basis for $\mathcal{F}(\mathfrak{T})$. Using this basis makes it more convenient to analyze the difference between $\mathcal{F}(\mathfrak{T})$ and $\mathcal{F}(\mathfrak{T}')$ when $\mathfrak{T}'$ is obtained from $\mathfrak{T}$ via a “grow” move. Indeed, let ψ_{l+1} denote the local decision stump corresponding to the new split. We then have $\mathcal{F}(\mathfrak{T}') = \mathcal{F}(\mathfrak{T}) \oplus \text{span}(\psi_{l+1})$.

Now consider a TSE $\mathfrak{E} = (\mathfrak{T}_1, \mathfrak{T}_2, \dots, \mathfrak{T}_m)$. We may also write a basis for $\mathcal{F}(\mathfrak{E})$ by concatenating the bases $\{\psi_{i,1}, \psi_{i,2}, \dots, \psi_{i,l_i}\}$ for each tree $\mathfrak{T}_i$, together with the constant function. If $\mathfrak{E}'$ is obtained from TSE via a “grow” move, we likewise have the property $\mathcal{F}(\mathfrak{E}') = \mathcal{F}(\mathfrak{E}) \oplus \text{span}(\psi')$, where ψ' is the local decision stump corresponding to the new split. On the other hand, the basis functions from different trees need not be orthogonal to each other. To regain orthogonality, we use the following lemma:

Lemma F.1 (Conditions for orthogonality). *Suppose $\nu = \nu_1 \times \nu_2 \times \dots \times \nu_d$ is a product measure on $\mathcal{X}$. Let $\mathfrak{T}_1$ and $\mathfrak{T}_2$ be two trees, and let $I_1, I_2 \subset \{1, 2, \dots, d\}$ be two disjoint subsets of indices such that $\mathfrak{T}_1$ and $\mathfrak{T}_2$ contain splits only on features in I_1 and I_2 respectively. Then the local decision stumps for both trees, $\{\psi_{1,1}, \psi_{1,2}, \dots, \psi_{1,l_1}\}$ and $\{\psi_{2,1}, \psi_{2,2}, \dots, \psi_{2,l_2}\}$, are orthogonal to each other.*

Proof. Consider two stumps from different trees: ψ_{1,k_1} and ψ_{2,k_2} . Under the assumption of a product measure, $\{x_i : i \in I_1\}$ is independent of $\{x_i : i \in I_2\}$. Since ψ_{1,k_1} is a function of the first set of variables and ψ_{2,k_2} is a function of the second set, they are thus independent of each other. We therefore have

$$\int \psi_{1,k_1} \psi_{2,k_2} d\nu = \int \psi_{1,k_1} d\nu_{I_1} \int \psi_{2,k_2} d\nu_{I_2} = 0,$$

where the second equality follows from the fact that all local decision stumps have mean zero. $\square$

Lemma F.2 (Existence of informative split). *For any finite contiguous subset of integers I , let $g : I \rightarrow \mathbb{R}$ be non-constant. Let ν be any measure on I . Then there exists a split at a threshold t with associated decision stump ψ such that t is a knot for g and*

$$\left(\int \phi g d\nu \right)^2 > 0.$$

As such, there is a sequence of recursive splits with associated local decision stumps $\psi_1, \psi_2, \dots, \psi_q$, where q is the number of knots of g , such that

- $g \in \text{span}(\psi_1, \psi_2, \dots, \psi_q)$;
- $(\int \psi_i g d\nu)^2 > 0$ for $i = 1, 2, \dots, q$;
- ψ_i splits on a knot of g for $i = 1, 2, \dots, q$.

Proof. Let $i_1, i_2, \dots, i_k$ denote the knots of g , and let $\tilde{\psi}_j$ denote the decision stump functions corresponding to a split at threshold $x = i_j$ for $j = 1, 2, \dots, k$, using the formula (49). Let $\tilde{\psi}_0 \equiv 1$ denote the constant function as usual. Then it is easy to see that $g \in \text{span}(\tilde{\psi}_0, \tilde{\psi}_1, \dots, \tilde{\psi}_k)$. As such, if $g \perp \tilde{\psi}_k$ for $k > 1$, then $g \in \text{span}(\tilde{\psi}_0)$, i.e. g is a constant function. To conclude the second statement, we apply the first part recursively to leaves obtained by making each split. $\square$

Lemma F.3 (Formula for decrease in bias). *Suppose $\mathfrak{E}'$ is obtained from $\mathfrak{E}$ via a “grow” move. Let ψ be the local decision stump associated with the new split. Let $\phi := \frac{\psi - \Pi_{\mathcal{F}(\mathfrak{E})}[\psi]}{\|\psi - \Pi_{\mathcal{F}(\mathfrak{E})}[\psi]\|_{L^2(\nu)}}$. Then for any regression function f , we have*

$$\text{Bias}^2(\mathfrak{E}'; f) = \text{Bias}^2(\mathfrak{E}; f) - \left(\int \phi f d\nu \right)^2.$$

Proof. We have

$$\mathcal{F}(\mathfrak{E}') = \mathcal{F}(\mathfrak{E}) \oplus \text{span}(\psi') = \mathcal{F}(\mathfrak{E}) \oplus \text{span}(\phi),$$

with the last expression comprising an orthogonal decomposition. As such, we have

$$\begin{aligned} \text{Bias}^2(\mathfrak{E}; f) - \text{Bias}^2(\mathfrak{E}'; f) &= \|f - \Pi_{\mathcal{F}(\mathfrak{E})}[f]\|_{L^2(\nu)}^2 - \|f - \Pi_{\mathcal{F}(\mathfrak{E})}[f] - \Pi_{\text{span}(\phi)}[f]\|_{L^2(\nu)}^2 \\ &= \|\Pi_{\text{span}(\phi)}[f]\|_{L^2(\nu)}^2 \\ &= \left(\int \phi f d\nu \right)^2 \end{aligned}$$

as we wanted. $\square$

Coverage. We first introduce some useful notation. Given a coordinate index i , let $\mathbf{x}^{-i} \in \{1, 2, \dots, b\}^{d-1}$, $x_i \in \{1, 2, \dots, b\}$. Combining these, we let $(\mathbf{x}^{-i}, x_i) \in \{1, 2, \dots, b\}^d$ have i -th coordinate equal to x_i and all other coordinates given by $\mathbf{x}^{-i}$. Also use $\mathbf{e}_i \in \mathbb{R}^d$ to denote the i -th coordinate vector. Given a real-valued function f defined on $\{1, 2, \dots, b\}^d$, we say that $\mathbf{x} \in \{1, 2, \dots, b\}^d$ is a *jump location* for f with respect to feature i if $f(\mathbf{x}) \neq f(\mathbf{x} + \mathbf{e}_i)$. Let $\mathbb{V}$ be a PEM. For each feature $i = 1, 2, \dots, d$ and $t = 1, 2, \dots, b$, the *coverage* of $\mathbb{V}$ of the split (i, t) is defined as

$$\text{coverage}(i, t; \mathbb{V}) := \left\{ \mathbf{x}^{-i} \in \{1, 2, \dots, b\}^{d-1} : \exists f \in \mathbb{V}, (\mathbf{x}^{-i}, t) \text{ is a jump location for } f \text{ with respect to feature } i \right\}. \quad (50)$$

If $\text{coverage}(i, t; \mathbb{V}) = \{1, 2, \dots, b\}^{d-1}$, we say that $\mathbb{V}$ has *full coverage* of the split (i, t) . Note that given a collection of PEMs $\mathbb{V}_1, \mathbb{V}_2, \dots, \mathbb{V}_m$, we have

$$\text{coverage}(i, t; \mathbb{V}_1 + \mathbb{V}_2, \dots, \mathbb{V}_m) = \bigcup_{j=1}^m \text{coverage}(i, t; \mathbb{V}_j).$$

Lemma F.4 (Zero bias requires full coverage of all knots). *Suppose $f^* \in \mathbb{V}$, where $f^*(\mathbf{x}) = f_1(x_1) + f_2(x_2) + \dots + f_{m'}(x_{m'})$ for some univariate functions $f_1, f_2, \dots, f_{m'}$. Then for any feature $1 \leq i \leq m'$ and any knot t of f_i , i.e. a value for which $f_i(t) \neq f_i(t+1)$, $\mathbb{V}$ has full coverage of the split (i, t) .*

Proof. For any $\mathbf{x}^{-i} \in \{1, 2, \dots, b\}^{d-1}$, we have

$$f^*((\mathbf{x}^{-i}, t+1)) - f^*((\mathbf{x}^{-i}, t)) = f_i(t+1) - f_i(t) \neq 0,$$

so that $(\mathbf{x}^{-i}, t)$ is a jump location for f^* with respect to feature i . Hence, if $f^* \in \mathbb{V}$, it provides the desired function in the definition (50). $\square$

Lemma F.5 (Full coverage implies inclusion of grid cells). *Suppose $\mathbb{V} = \mathcal{F}(\mathfrak{T})$ for a single tree $\mathfrak{T}$. Suppose $\mathbb{V}$ has full coverage of split $(i, \xi_{i,1}), (i, \xi_{i,2}), \dots, (i, \xi_{i,q_i})$ for $i = 1, 2, \dots, k$. For any choice of $1 \leq j_i \leq q_i$, $i = 1, 2, \dots, k$, denote the cell*

$$\mathcal{C} := \{\mathbf{x} : \xi_{i,j_i-1} < x_i \leq \xi_{i,j_i} \text{ for } i = 1, 2, \dots, k\}.$$

We then have $\mathbf{1}_{\mathcal{C}} \in \mathbb{V}$.

Proof. Let $\mathbf{x} \in \mathcal{C}$ be any point, and let $\mathcal{L}(\mathbf{x})$ be the leaf of $\mathfrak{T}$ containing $\mathbf{x}$. We claim that $\mathcal{L}(\mathbf{x}) \subset \mathcal{C}$. Suppose not, then there exists a coordinate direction i in which $\mathcal{L}(\mathbf{x})$ exceeds $\mathcal{C}$. By reordering if necessary, we may thus assume that $(\mathbf{x}^{-i}, \xi_{i,j_i} + 1) \in \mathcal{L}(\mathbf{x})$. For any $f \in \mathbb{V}$, we may write $f = a_0 \mathbf{1}_{\mathcal{L}(\mathbf{x})} + \sum_{l=1}^L a_l \mathbf{1}_{\mathcal{L}_l}$, where $\mathcal{L}_1, \mathcal{L}_2, \dots, \mathcal{L}_L$ are leaves in $\mathfrak{T}$. We then have

$$f(\mathbf{x}^{-i}, \xi_{i,j_i} + 1) = a_0 = f(\mathbf{x}^{-i}, \xi_{i,j_i}),$$

which contradicts our assumption that $\mathbb{V}$ has full coverage of the split at (i, ξ_{i,j_i}) . We thus have

$$\mathcal{C} = \bigcup_{\mathbf{x} \in \mathcal{C}} \mathcal{L}(\mathbf{x}) = \bigcup_{\mathcal{L}_l \subset \mathcal{C}} \mathcal{L}_l.$$

Since the right hand side is union over a collection of disjoint sets, we have $\mathbf{1}_{\mathcal{C}} = \sum_{\mathcal{L}_l \subset \mathcal{C}} \mathbf{1}_{\mathcal{L}_l} \in \mathbb{V}$. $\square$

Lemma F.6 (Sufficient condition for lack of coverage). *Let $\mathfrak{T}$ be a tree structure. Consider a split (i, t) , and let $\mathfrak{t}$ be any node in $\mathfrak{T}$. Suppose the following hold:*

- $(\mathbf{x}^{-i}, t) \in \mathfrak{t}$ for some $\mathbf{x}^{-i} \in \{1, 2, \dots, b\}^{d-1}$;
- No ancestor or descendant of $\mathfrak{t}$, including $\mathfrak{t}$ itself, uses the splitting rule (i, t) ;

Then if we let $\mathfrak{t}^{-i}$ denote the projection of $\mathfrak{t}$ onto all but the i -th coordinate, we have

$$\text{coverage}(i, t; \mathcal{F}(\mathfrak{T})) \cap \mathfrak{t}^{-i} = \emptyset.$$

Proof. Let $\mathbf{z}^{-i} \in \mathfrak{t}^{-i}$ be any point. By the first assumption, t is within the bounds of $\mathfrak{t}$ along direction i , so $(\mathbf{z}^{-i}, t) \in \mathfrak{t}$. Let $\mathcal{L}_0$ be the leaf node containing $(\mathbf{z}^{-i}, t)$. By the second assumption, $(\mathbf{z}^{-i}, t+1) \in \mathcal{L}_0$, otherwise some parent of $\mathcal{L}$ would have made the split (i, t) . To show that $\mathbf{z}^{-i} \notin \text{coverage}(i, t; \mathcal{F}(\mathfrak{T}))$, it suffices to show that $f(\mathbf{z}^{-i}, t+1) = f(\mathbf{z}^{-i}, t)$ for all $f \in \mathcal{F}(\mathfrak{T})$. We may write $f = \sum_{j=0}^k a_j \mathbf{1}_{\mathcal{L}_j}$ where $\mathcal{L}_1, \mathcal{L}_2, \dots, \mathcal{L}_j$ are other leaves in $\mathfrak{T}$. We then have $f(\mathbf{z}^{-i}, t+1) = a_0 = f(\mathbf{z}^{-i}, t)$ as we wanted. $\square$

Proposition F.7 (Dimension of additive functions). *Suppose $f^* \in \mathbb{V}$, where $f^*(\mathbf{x}) = f_1(x_1) + f_2(x_2) + \dots + f_m(x_m)$ for some univariate functions $f_1, f_2, \dots, f_m$. Then for any $l < m$ we have $\dim_l(f^*) > \dim_m(f^*)$, while for any $l \geq m$, we have $\dim_l(f^*) = \dim_m(f^*)$.*

Proof. Let $\mathbb{V} = \mathcal{F}(\mathfrak{E})$ for some $\mathfrak{E}$ with l trees, and suppose $f^* \in \mathbb{V}$. For $i = 1, 2, \dots, m$, let $(i, \xi_{i,1}), (i, \xi_{i,2}), \dots, (i, \xi_{i,q_i-1})$ denote the knots of f_i , where $q_i = \dim_1(f_i)$. We claim that $\mathfrak{E}$ must contain a node with splitting rule $(i, \xi_{i,j})$ for every choice of i and j . If not, then $\mathbb{V}$ does not cover $(i, \xi_{i,j})$, contradicting Lemma F.4.

Construct $\mathfrak{E}$ via a sequence of grow moves in arbitrary order, thereby deriving a sequence $\mathfrak{E}_0 = \mathfrak{E}^0, \mathfrak{E}^1, \dots, \mathfrak{E}^k$. Consider the first time $t_{i,j}$ for which a split using rule $(i, \xi_{i,j})$ is added to the TSE. Let $\psi_{i,j}$ denote the local decision stump associated with the split. Since $\mathcal{F}(\mathfrak{E}^{t_{i,j}-1})$ has zero coverage of $(i, \xi_{i,j})$, we have $\text{df}(\psi_{i,j}) = \text{df}(\mathfrak{E}_{i,j}) + 1$. Adding this up over all splits gives the lower bound $\text{df}(\mathbb{V}) \geq \sum_{i=1}^m q_i - m$. By putting all splits on feature i on the i -th tree for $i = 1, 2, \dots, m$, we see that this lower bound is achievable whenever $l \geq m$, thereby giving

$$\dim_l(f^*) = \sum_{i=1}^m q_i - m = \dim_m(f^*).$$

When $l < m$, then by the pigeonhole principle, there exists splits on different features (i, s) and (j, t) that occur on the same tree. We claim that this implies that either there exists two linearly independent local decision stumps ϕ and ψ splitting on (i, s) or the same applies to (j, t) . Suppose not, then the unique local decision stump splitting on (i, s) must not depend on any other feature, while the same applies to that splitting on (j, t) . However, because nodes depend on the feature used in the root split, this cannot be simultaneously true if they both correspond to splits on the same tree. By repeating the argument above, we therefore get $\text{df}(\mathfrak{E}) \geq \sum_{i=1}^m q_i - m + 1$. Since this holds for any $\mathfrak{E}$, we have

$$\dim_l(f^*) \geq \sum_{i=1}^m q_i - m + 1 > \dim_m(f^*)$$

as we wanted. $\square$

G Proof of Theorem 5.1 Part 1

Set-up. For $i = 1, 2, \dots, m'$, denote $q_i = \dim_1(f_i)$, and let $0 = \xi_{i,0} < \xi_{i,1} < \dots < \xi_{i,q_i} = b$ denote the *knots* of f_j , i.e. the values for which $f_i(\xi_{i,j}) \neq f_i(\xi_{i,j+1})$, together with the endpoints.¹² Without loss of generality, assume that $f_1, \dots, f_{m'}$ are ordered in descending order of their 1-ensemble dimension, i.e. $q_1 \geq q_2 \geq \dots \geq q_{m'}$.

We now define $\mathfrak{E}_{\text{bad}}$ and a “bad set” $\mathcal{A}$. To this end, we define a collection of partition models $\mathbb{V}_1, \mathbb{V}_2, \dots, \mathbb{V}_m$ (spans of indicators in a single partition) as follows. First, for $i = 3, 4, \dots, m-1$, $j = 1, 2, \dots, q_i$, define the cells

¹² $\dim_1(f_i)$ is simply the number of constant pieces of f_i or alternatively, one larger than the number of knots of f_i .

$\mathcal{L}_{i,j} := \{\mathbf{x}: \xi_{i,j-1} < x_i \leq \xi_{i,j}\}$, and set $\mathbb{V}_i = \text{span}(\{\mathbf{1}_{\mathcal{L}_{i,j}}: j = 1, 2, \dots, q_i\})$, i.e. for each i , $\mathbb{V}_i$ contains splits only on feature i , and only at the knots of f_i . This also implies that $f_i \in \mathbb{V}_i$. Next, for the remaining component functions, $i = m, m+1, \dots, m'$, $j_i = 1, 2, \dots, q_i$, define the cells

$$\mathcal{L}_{m,j_m,j_{m+1},\dots,j_{m'}} := \{\mathbf{x}: \xi_{i,j_i-1} < x_i \leq \xi_{i,j_i} \text{ for } i = m, m+1, \dots, m'\}, \quad (51)$$

and set $\mathbb{V}_m$ to be the span of their indicators. Observe that $\mathbb{V}_m$ comprises a grid contains splits only on features $m, m+1, \dots, m'$, and only at the knots of $f_m, f_{m+1}, \dots, f_{m'}$ respectively. This also implies $f_m, f_{m+1}, \dots, f_{m'} \in \mathbb{V}_m$.

Finally, to introduce inefficiency, we define each of $\mathbb{V}_1$ and $\mathbb{V}_2$ to have splits on both features 1 and 2. This construction is fairly involved and will be detailed in full in the next section. For now, it suffices to assume that $f_1 + f_2 \in \mathbb{V}_1 + \mathbb{V}_2$, which implies that $f^* \in \mathbb{V}_1 + \mathbb{V}_2 + \dots + \mathbb{V}_m$. Define $\mathcal{A}$ via

$$\mathcal{A} := \{(\mathfrak{T}_1, \mathfrak{T}_2, \dots, \mathfrak{T}_m): \mathcal{F}(\mathfrak{T}_i) = \mathbb{V}_i \text{ for } i = 1, 2, \dots, m\}.$$

It is clear that for any $\mathfrak{E} \in \mathcal{A}$, we have $\mathcal{F}(\mathfrak{E}) = \mathbb{V}_1 + \mathbb{V}_2 + \dots + \mathbb{V}_m$, so that $\mathcal{A}$ comprises a collection of TSEs with zero bias. We will set $\mathfrak{E}_{\text{bad}}$ to be a particular element of $\mathcal{A}$. We set $\mathcal{B}$ to be the outer boundary of $\mathcal{A}$.

Hitting precedence probability lower bound. *Step 1: Construction of path.* We construct a path in $\Omega_{\text{TSE},m}$ comprising $\mathfrak{E}_0 = \mathfrak{E}^0, \mathfrak{E}^1, \dots, \mathfrak{E}^s = \mathfrak{E}_{\text{bad}}$ such that $\mathfrak{E}^{i+1}$ is obtained from $\mathfrak{E}^i$ via a “grow” move for $i = 0, 1, \dots, s-1$. To do this, we first apply Lemma F.2 to obtain, for each feature $j = 1, 2, \dots, m'$, a sequence of recursive splits $\psi_{j,1}, \psi_{j,2}, \dots, \psi_{j,q_j-1}$ such that

- $f_j \in \text{span}(\psi_{j,1}, \psi_{j,2}, \dots, \psi_{j,q_j-1})$;
- $(\int f_j \psi_{j,l} d\nu_j)^2 > 0$ for $l = 1, 2, \dots, q_j - 1$;
- $\psi_{j,l}$ splits on a knot of f_j for $l = 1, 2, \dots, q_j - 1$.

Note that for convenience, we have identified the splits with their associated local decision stumps. We break up the path into m segments, $0 = s_0 < s_1 < \dots < s_m = k$, with the j -th segment comprising $\mathfrak{E}^{s_{j-1}+1}, \mathfrak{E}^{s_{j-1}+2}, \dots, \mathfrak{E}^{s_j}$, possessing the following desired properties:

- During this segment, the j -th tree is grown from the empty tree $\mathfrak{T}_0$ to its final state $\mathfrak{T}_j^*$, while no other trees are modified;
- $\mathcal{F}(\mathfrak{T}_j^*) = \mathbb{V}_j$;
- $\text{Bias}^2(\mathfrak{E}^{i-1}; f^*) > \text{Bias}^2(\mathfrak{E}^i; f^*)$ for $i = s_{j-1} + 1, s_{j-1} + 2, \dots, s_j$.

For $j = 3, 4, \dots, m-1$, the splits $\psi_{j,1}, \psi_{j,2}, \dots, \psi_{j,q_j-1}$ immediately yield a sequence of trees $\mathfrak{T}_0 = \mathfrak{T}_j^0, \mathfrak{T}_j^1, \dots, \mathfrak{T}_j^{q_j-1} = \mathfrak{T}_j^*$ such that each $\mathfrak{T}_j^l$ is obtained from $\mathfrak{T}_j^{l-1}$ via adding the split $\psi_{j,l}$. Next, assuming correctness of the construction up to $\mathfrak{E}^{s_{j-1}}$, $\psi_{j,l}$ is orthogonal to $\mathcal{F}(\mathfrak{E}^{s_{j-1}})$ and also to the previous splits $\psi_{j,1}, \psi_{j,2}, \dots, \psi_{j,l-1}$, by Lemma F.1. This allows us to apply Lemma F.3 to get

$$\text{Bias}^2(\mathfrak{E}^{s_{j-1}+i}; f) = \text{Bias}^2(\mathfrak{E}^{s_{j-1}+i-1}; f) - \left(\int \psi_i f^* d\nu \right)^2.$$

We then notice that by the independence of different features,

$$\left(\int \psi_i f^* d\nu \right)^2 = \left(\int \psi_i f_j d\nu \right)^2 > 0.$$

To construct $\mathfrak{T}_m^*$, fix $0 \leq p < m' - m$, and assume that we have constructed a tree $\mathfrak{T}_{m,p}^*$ such that the leaves of $\mathfrak{T}_{m,p}^*$ comprise the collection (see equation (51)):

$$\{\mathcal{L}_{m,j_m,j_{m+1},\dots,j_{m+p}}: 1 \leq j_{m+i} \leq q_{m+i}, i = 0, 1, 2, \dots, p\}$$

We now construct a sequence $\mathfrak{T}_{m,p}^* = \mathfrak{T}_{m,p}^0, \mathfrak{T}_{m,p}^1, \dots, \mathfrak{T}_{m,p}^r = \mathfrak{T}_{m,p+1}^*$ by looping over the leaves, and for each leaf $\mathcal{L} = \mathcal{L}_{m,j_m,j_{m+1},\dots,j_{m+p}}$, iteratively adding the splits $\tilde{\psi}_1, \tilde{\psi}_2, \dots, \tilde{\psi}_{q_{m+p+1}-1}$, where $\tilde{\psi}_l = \nu(\mathcal{L})^{-1/2} \psi_{m+p+1,l} \mathbf{1}_{\mathcal{L}}$ for $l = 1, 2, \dots, q_{m+p+1} - 1$. Note that these are orthonormal. When adding the split $\tilde{\psi}_l$, the decrease in squared bias can thus be computed as

$$\text{Bias}^2(\mathfrak{T}_{m,p}^{i-1}) - \text{Bias}^2(\mathfrak{T}_{m,p}^i) = \left(\int \tilde{\psi}_l f^* d\nu \right)^2 = \nu(\mathcal{L}) \left(\int \psi_{m+p+1,l} f_j d\nu \right)^2 > 0. \quad (52)$$

We lift the sequence of tree structures to a sequence of TSEs, using Lemma F.1 and Lemma F.3 as before to translate (52) to be in terms of the sequence of TSEs.

It remains to define $\mathbb{V}_1$ and $\mathbb{V}_2$ and construct $\mathfrak{T}_1^*$ and $\mathfrak{T}_2^*$. Let a_i be the index of the knot ξ_{i,a_i} forming the threshold for ψ_{i,q_i} for $i = 1, 2$. For $k = 0, 1, \dots, q_1 - 1$, define $b_k = k$ for $k < a_1$ and $b_k = k + 1$ for $k \geq a_1$. Similarly, for $l = 0, 1, \dots, q_2 - 1$, define $c_l = l$ for $l < a_2$ and $c_l = l + 1$ for $l \geq a_2$. Set $\mathbb{V}_1$ to be the span of indicators of the cells

$$\mathcal{L}_{1,k,l} := \{ \mathbf{x} : \xi_{1,b_{k-1}} < x_1 \leq \xi_{1,b_k} \text{ and } \xi_{2,c_{l-1}} < x_2 \leq \xi_{2,c_l} \}$$

for $k = 1, 2, \dots, q_1 - 1$ and $l = 1, 2$. Next, define $d_0 = 0, d_1 = a_1, d_2 = q_1, e_1 = 0, e_2 = a_2, e_3 = q_2$. We set $\mathbb{V}_2$ to be the span of indicators of the cells

$$\mathcal{L}_{2,k,l} := \{ \mathbf{x} : \xi_{1,d_{k-1}} < x_1 \leq \xi_{1,d_k} \text{ and } \xi_{2,e_{l-1}} < x_2 \leq \xi_{2,e_l} \}$$

for $k, l = 1, 2$. We construct $\mathfrak{T}_1^*$ similarly to $\mathfrak{T}_m^*$, through recursive partitioning on x_1 and x_2 , but without making the final split on both features. Using the same argument as before, we obtain the desired sequence of trees. Finally, to construct $\mathfrak{T}_2^*$, we simply make the omitted splits on the new tree. More precisely, we define the splits

$$\psi_1 = \frac{\nu\{x_1 > \xi_{1,a_1}\} \mathbf{1}\{x_1 \leq \xi_{1,a_1}\} - \nu\{x_1 \leq \xi_{1,a_1}\} \mathbf{1}\{x_1 > \xi_{1,a_1}\}}{\sqrt{\nu\{x_1 \leq \xi_{1,a_1}\} \nu\{x_1 > \xi_{1,a_1}\}}}$$

and $\psi_2 = \nu\{x_1 \leq \xi_{1,a_1}\}^{1/2} \phi \mathbf{1}\{x \leq \xi_{1,a_1}\}$, $\psi_3 = \nu\{x_1 > \xi_{1,a_1}\}^{1/2} \phi \mathbf{1}\{x > \xi_{1,a_1}\}$, where

$$\phi = \frac{\nu\{x_2 > \xi_{2,a_2}\} \mathbf{1}\{x_2 \leq \xi_{2,a_2}\} - \nu\{x_2 \leq \xi_{2,a_2}\} \mathbf{1}\{x_2 > \xi_{2,a_2}\}}{\sqrt{\nu\{x_2 \leq \xi_{2,a_2}\} \nu\{x_2 > \xi_{2,a_2}\}}}.$$

It is clear that these, together with the constant function, span $\mathbb{V}_2$. Furthermore, it is easy to see that we have

$$\begin{aligned} \frac{\psi_1 - \Pi_{\mathbb{V}_1}[\psi_1]}{\|\psi_1 - \Pi_{\mathbb{V}_1}[\psi_1]\|_{L^2(\nu)}} &= \psi_{1,q_1}, \\ \frac{\psi_2 - \Pi_{\mathbb{V}_1 \oplus \text{span}(\psi_1)}[\psi_2]}{\|\psi_2 - \Pi_{\mathbb{V}_1 \oplus \text{span}(\psi_1)}[\psi_2]\|_{L^2(\nu)}} &= \nu\{x_1 \leq \xi_{1,a_1}\}^{-1/2} \psi_{2,q_2} \mathbf{1}\{x \leq \xi_{1,a_1}\}, \\ \frac{\psi_3 - \Pi_{\mathbb{V}_1 \oplus \text{span}(\psi_1, \psi_2)}[\psi_3]}{\|\psi_3 - \Pi_{\mathbb{V}_1 \oplus \text{span}(\psi_1, \psi_2)}[\psi_3]\|_{L^2(\nu)}} &= \nu\{x_1 > \xi_{1,a_1}\}^{-1/2} \psi_{2,q_2} \mathbf{1}\{x > \xi_{1,a_1}\}. \end{aligned}$$

By construction, we have

$$\left(\int \psi_{1,q_1} f_1 d\nu_1 \right)^2, \left(\int \psi_{2,q_2} f_2 d\nu_2 \right)^2 > 0.$$

Applying Lemma F.3 then shows that the desired property for the 2nd segment of TSEs is satisfied. It is clear from the construction that $\psi_{j1}, \psi_{j2}, \dots, \psi_{jq_{j-1}} \in \mathbb{V}_1$ and that $\psi_{jq_j} \in \mathbb{V}_2$ for $j = 1, 2$. This implies that $f_1 + f_2 \in \mathbb{V}_1 + \mathbb{V}_2$ as desired.

Step 2: Disjointness from optimal set. We have already proved, in this construction, that every “grow” move adds a split that decreases bias, and therefore must be linearly independent from the existing PEM. In other words, we have $\text{df}(\mathfrak{E}^{i+1}) = \text{df}(\mathfrak{E}^i)$ for $i = 0, 1, \dots, k - 1$. Expanding this gives $\text{df}(\mathfrak{E}_{\text{bad}}) = k + 1 = \sum_{j=1}^m (s_j - s_{j-1}) + 1$. Since each split increments the number of leaf nodes by one, each $s_j - s_{j-1} + 1$ is equal to the number of cells in $\mathbb{V}_j$, which

we now compute as follows. For $j = 3, 4, \dots, m-1$, we have $s_j - s_{j-1} = \dim_1(f_j) - 1$, while for $j = m$, we have $s_m - s_{m-1} = \prod_{l=m}^{m'} \dim_1(f_l) - 1$. Likewise, we have $s_1 = (\dim_1(f_1) - 1)(\dim_1(f_2) - 1) - 1$ and $s_2 - s_1 = 3$. Putting these together gives

$$\text{df}(\mathfrak{E}_{\text{bad}}) = \sum_{j=3}^{m-1} \dim_1(f_j) + \prod_{l=m}^{m'} \dim_1(f_l) + (\dim_1(f_1) - 1)(\dim_1(f_2) - 1) + 2 - m.$$

Let us now show that this is suboptimal, by constructing a TSE $\mathfrak{E}_{\text{good}}$ with zero bias but fewer degrees of freedom. To do so, we simply set $\mathfrak{E}_{\text{good}} = (\mathfrak{T}'_1, \mathfrak{T}'_2, \mathfrak{T}^*_3, \dots, \mathfrak{T}^*_m)$, where $\mathfrak{T}^*_j$ has the same structure as in $\mathfrak{T}_{\text{bad}}$ for $j = 3, 4, \dots, m$. On the other hand, we define $\mathfrak{T}'_1$ using $\psi_{1,1}, \psi_{1,2}, \dots, \psi_{1,q_1}$ and $\mathfrak{T}'_2$ using $\psi_{2,1}, \psi_{2,2}, \dots, \psi_{2,q_2}$. By assumption, we have $f_j \in \mathcal{F}(\mathfrak{T}'_j)$ for $j = 1, 2$, which yields unbiasedness. The degrees of freedom is given by

$$\text{df}(\mathfrak{E}_{\text{good}}) = \sum_{j=1}^{m-1} \dim_1(f_j) + \prod_{l=m}^{m'} \dim_1(f_l) - m. \quad (53)$$

Taking the difference gives

$$\begin{aligned} \text{df}(\mathfrak{E}_{\text{bad}}) - \text{df}(\mathfrak{E}_{\text{good}}) &= (\dim_1(f_1) - 1)(\dim_1(f_2) - 1) + 2 - \dim_1(f_1) - \dim_1(f_2) \\ &= (\dim_1(f_1) - 2)(\dim_1(f_2) - 2) - 1, \end{aligned}$$

which is strictly larger than 0 if $\dim_1(f_1), \dim_1(f_2) > 3$. Combining this with the fact that $\text{Bias}^2(\mathfrak{E}^i) > 0$ for all $i < s$, we therefore have $\mathfrak{E}^i \notin \text{OPT}_m(f^*, k)$ for all $i = 0, 1, \dots, s$, with $k = (\dim_1(f_1) - 2)(\dim_1(f_2) - 2) - 2$.

Step 3: Conclusion. Using Proposition 4.1 and Proposition 4.2, for n large enough, there is a $1 - \delta/2$ event over which

$$\log p(\mathfrak{E}^i | \mathbf{y}) - \log p(\mathfrak{E}^{i-1} | \mathbf{y}) = \frac{n}{2\sigma^2} (\text{Bias}^2(\mathfrak{E}^{i-1}; f^*) - \text{Bias}^2(\mathfrak{E}^i; f^*)) + O\left(\sqrt{n \log(s/\delta)}\right). \quad (54)$$

Conditioning on this set, we get $\frac{p(\mathfrak{E}^i | \mathbf{y})}{p(\mathfrak{E}^{i-1} | \mathbf{y})} = \Omega(1)$ for $i = 1, 2, \dots, s$ for all n large enough. We may then repeat the calculations in the proof of Theorem 5.2, specifically equations (36) and (37), to get

$$P\{\tau_{\mathfrak{E}_{\text{bad}}} < \tau_{\text{OPT}_m(f^*, k)}\} \geq P\{\mathfrak{E}_i = \mathfrak{E}^i \text{ for } i = 1, 2, \dots, s\} = \Omega(1).$$

BIC lower bound. Consider $\mathfrak{E}' = (\mathfrak{T}'_1, \mathfrak{T}'_2, \dots, \mathfrak{T}'_m) \in \mathcal{B}$. By definition of $\mathcal{B}$, there exists $\mathfrak{E} = (\mathfrak{T}_1, \mathfrak{T}_2, \dots, \mathfrak{T}_m) \in \mathcal{A}$ such that $\mathfrak{E}'$ is obtained from $\mathfrak{E}$ via a “grow”, “prune”, “change”, or “swap” move. We now consider each type of move and show that either $\text{Bias}^2(\mathfrak{E}'; f^*) > 0$ or $\text{df}(\mathfrak{E}') > \text{df}(\mathfrak{E}) = \text{df}(\mathfrak{E}_{\text{bad}})$. To prove this, we make use of a few key observations.

- Since a move only affects one tree with index i_0 , we have $\mathfrak{T}'_i = \mathfrak{T}_i$ and $\mathcal{F}(\mathfrak{T}'_i) = \mathbb{V}_i$ for all $i \neq i_0$;
- Every split $(i, \xi_{i,j})$ necessary for $\mathcal{F}(\mathfrak{E}')$ to be unbiased (see Lemma F.4) has full coverage in a single tree and has zero coverage in all other trees;
- Since no move is allowed to result in an empty leaf node, if $\mathfrak{t}$ is an internal node in $\mathfrak{T}_{i_0}$ or $\mathfrak{T}'_{i_0}$ with a split (i, ξ) , no ancestor or descendent of $\mathfrak{t}$ makes the same split (i, ξ) .

For convenience, we denote $\mathbb{V}_{-i_0} = \mathbb{V}_1 + \dots + \mathbb{V}_{i_0-1} + \mathbb{V}_{i_0+1} + \dots + \mathbb{V}_m$.

Case 1: “grow” move. Let ψ denote the local decision stump associated with the new split, and let $\mathcal{L}$ denote the leaf that is split. As shown earlier, we have $\psi \perp \mathbb{V}_{i_0}$. Next, for any boundary point $(\mathbf{x}^{-j}, \xi)$ of $\mathcal{L}$ in direction j (that is not an a boundary point of the entire space $\mathcal{X}$), ψ has a jump location at $(\mathbf{x}^{-j}, \xi)$ with respect to feature j . On the other hand, by our construction, this means that (j, ξ) is a split fully covered by $\mathbb{V}_{i_0}$ and has zero coverage in $\mathbb{V}_{-i_0}$. By Lemma F.4, this means $\psi \notin \mathbb{V}_{-i_0}$. Together, this implies that $\psi \notin \mathcal{F}(\mathfrak{E})$ and $\text{df}(\mathfrak{E}') = \text{df}(\mathfrak{E}) + 1$ as desired.

Case 2: “prune” move. Suppose the pruned split is (k, t) and occurs on a leaf $\mathcal{L}$. Let $\mathcal{L}^{-k}$ denote the projection of the leaf onto all but the k -th coordinate. Applying the third observation above together with Lemma F.6 gives

$$\text{coverage}(k, t; \mathcal{F}(\mathfrak{E}')) \cap \mathcal{L}^{-k} = \text{coverage}(k, t; \mathcal{F}(\mathfrak{T}'_{i_0})) \cap \mathcal{L}^{-k} = \emptyset.$$

Lemma F.4 then implies that $\text{Bias}^2(\mathfrak{E}'; f^*) > 0$.

Case 3: “change” move. Suppose the “changed” split occurs on a node t and is with respect to a feature k at threshold $\xi_{k,j}$. Then since no descendant of t makes the same split, if we let t^{-k} denote the projection of t onto all but the k -th coordinate, then as before, we have

$$\text{coverage}(k, t; \mathcal{F}(\mathfrak{E}')) \cap t^{-k} = \text{coverage}(k, t; \mathcal{F}(\mathfrak{T}'_{i_0})) \cap t^{-k} = \emptyset,$$

and $\text{Bias}^2(\mathfrak{E}'; f^*) > 0$.

Case 4: “swap” move. The swap move can either be performed on a pair of parent-child nodes, or on a parent with both of its children, if both children have the same splitting rule. In the latter case, $\mathcal{F}(\mathfrak{T}') = \mathcal{F}(\mathfrak{T})$, contradicting our assumption that $\mathfrak{E}' \notin \mathcal{A}$. In the former case, let t denote the parent, and let t_L and t_R denote its two children. Suppose without loss of generality that the splitting rules (j, s) and (k, t) , of t and t_L respectively, are to be swapped. Then by construction of $\mathfrak{T}_{i_0}$, t must be a knot for f_k . By the third observation, no ancestor of t_L uses the splitting rule (k, t) , which implies that a descendent of t_R must split on (k, t) . However, this would mean that this swap move is not allowed, giving a contradiction.

Finally, using Proposition 4.1 and Proposition 4.2, for n large enough, there is a $1 - \delta/2$ event over which

$$\begin{aligned} \Delta \text{BIC}(\mathfrak{E}', \mathfrak{E}) &= \begin{cases} \frac{n}{\sigma^2} \text{Bias}^2(\mathfrak{E}'; f^*) + O\left(\sqrt{n \log(|\mathcal{B}|/\delta)}\right) & \text{if } \text{Bias}^2(\mathfrak{E}'; f^*) > 0 \\ \log n(\text{df}(\mathfrak{E}') - \text{df}(\mathfrak{E}_{\text{bad}})) + O(\log(|\mathcal{B}|/\delta)) & \text{otherwise.} \end{cases} \end{aligned} \quad (55)$$

Condition further on this event.

Conclusion. Applying Proposition 6.1 together with equations (54) and (55), while taking n large enough, we get a $1 - 2\delta$ probability event over which

$$E\{\tau_{\text{OPT}_m(f^*, k)}\} = \Omega(n^{1/2}),$$

where $k = (\dim_1(f_1) - 2)(\dim_1(f_2) - 2) - 1$.

H Proof of Theorem 5.1 Part 2

Set-up. We first use the same definition of $\mathfrak{E}_{\text{good}} = (\mathfrak{T}_1, \mathfrak{T}_2, \dots, \mathfrak{T}_m)$ as in the previous section (see the paragraph immediately preceding equation (53)). To define $\mathfrak{E}_{\text{bad}} = (\mathfrak{T}_1^*, \mathfrak{T}_2^*, \dots, \mathfrak{T}_m^*)$, we start with $\mathfrak{E}_{\text{good}}$ and simply swap the roles of f_1 and $f_{m'}$. More precisely, we let $\mathfrak{T}_j^* = \mathfrak{T}_j$ for $j = 2, 3, \dots, m-1$. We define $\mathfrak{T}_1^*$ as comprising the local decision stumps $\phi_{m',1}, \phi_{m',2}, \dots, \phi_{m',q_{m'}}$, which were defined in the proof of the hitting precedence probability lower bound in the previous section. Define $\mathbb{V}_j = \mathcal{F}(\mathfrak{T}_j^*)$ for $j = 1, 2, \dots, m-1$. We then define $\mathbb{V}_m$ to be a grid on features $\{1, m, m+1, \dots, m'-1\}$, or in other words, $\mathcal{F}(\mathfrak{T}_m^*)$ is the span of indicators of the cells

$$\mathcal{L}_{m,j_m,j_{m+1},\dots,j_{m'}} := \{\mathbf{x}: \xi_{i,j_{i-1}} < x_i \leq \xi_{i,j_i} \text{ for } i = 1, m, m+1, \dots, m'-1\}, \quad (56)$$

as we vary $i = 1, m, m+1, \dots, m'-1$ and $j_i = 1, 2, \dots, q_i$. We construct $\mathfrak{T}_m^*$ such that $\mathcal{F}(\mathfrak{T}_m^*)$, in the manner described in the proof of the hitting precedence probability lower bound in the previous section. Here, recall that $\dim_1(f_1) > \dim_1(f_m)$, which will introduce suboptimality into $\mathfrak{E}_{\text{bad}}$.

Notice that $f^* \in \mathbb{V} := \mathbb{V}_1 + \mathbb{V}_2 + \dots + \mathbb{V}_m$ via the same argument as in the previous section. Next, we define the set $\mathcal{A}$ via

$$\mathcal{A} := \{(\mathfrak{E}'_1, \mathfrak{E}', \dots, \mathfrak{E}'): \mathcal{F}(\mathfrak{E}'_j) \supseteq \mathbb{V}_j \text{ for } j = 1, 2, \dots, m\}.$$

We define $\mathcal{B}$ to be the outer boundary of $\mathcal{A}$.

Hitting precedence probability lower bound. This follows the proof of the hitting precedence probability lower bound in the previous section almost exactly. First, using the same construction, there is path in $\Omega_{\text{TSE},m}$ comprising $\mathfrak{E}_\emptyset = \mathfrak{E}^0, \mathfrak{E}^1, \dots, \mathfrak{E}^s = \mathfrak{E}_{\text{bad}}$ such that $\mathfrak{E}^{i+1}$ is obtained from $\mathfrak{E}^i$ via a “grow” move and

$$\text{Bias}^2(\mathfrak{E}^i; f^*) > \text{Bias}^2(\mathfrak{E}^{i-1}; f^*)$$

for $i = 0, 1, \dots, s-1$.

Next, we compute

$$\text{df}(\mathfrak{E}_{\text{bad}}) = \sum_{j=2}^{m-1} \dim_1(f_j) + \dim_1(f_{m'}) + \dim_1(f_1) \prod_{j=m}^{m'-1} \dim_1(f_j) - m.$$

Taking the difference between this and (53) gives

$$\begin{aligned} \text{df}(\mathfrak{E}_{\text{bad}}) - \text{df}(\mathfrak{E}_{\text{good}}) &= \dim_1(f_{m'}) - \dim_1(f_1) + \dim_1(f_1) \prod_{j=m}^{m'-1} \dim_1(f_j) - \prod_{j=m}^{m'} \dim_1(f_j) \\ &= (\dim_1(f_1) - \dim_1(f_{m'})) \prod_{j=m}^{m'-1} \dim_1(f_j), \end{aligned} \quad (57)$$

which is strictly larger than 0 if $\dim_1(f_1) > \dim_1(f_{m'})$. Combining this with the fact that $\text{Bias}^2(\mathfrak{E}^i) > 0$ for all $i < s$, we therefore have $\mathfrak{E}^i \notin \text{OPT}_m(f^*, k-1)$ where k is the quantity on the right hand side of (57).

Finally, using Proposition 4.1 and Proposition 4.2, for n large enough, there is a $1 - \delta/2$ event over which

$$\log p(\mathfrak{E}^i | \mathbf{y}) - \log p(\mathfrak{E}^{i-1} | \mathbf{y}) = \frac{n}{2\sigma^2} (\text{Bias}^2(\mathfrak{E}^{i-1}; f^*) - \text{Bias}^2(\mathfrak{E}^i; f^*)) + O\left(\sqrt{n \log(s/\delta)}\right). \quad (58)$$

Conditioning on this set, we get $\frac{p(\mathfrak{E}^i | \mathbf{y})}{p(\mathfrak{E}^{i-1} | \mathbf{y})} = \Omega(1)$ for $i = 1, 2, \dots, s$ for all n large enough. We may then repeat the calculations in the proof of Theorem 5.2, specifically equations (36) and (37), to get

$$P\{\tau_{\mathfrak{E}_{\text{bad}}} < \tau_{\text{OPT}_m(f^*, k)}\} \geq P\{\mathfrak{E}_i = \mathfrak{E}^i \text{ for } i = 1, 2, \dots, s\} = \Omega(1). \quad (59)$$

BIC lower bound. Consider $\mathfrak{E}' = (\mathfrak{T}'_1, \mathfrak{T}'_2, \dots, \mathfrak{T}'_{m'}) \in \mathcal{B}$. By definition of $\mathcal{B}$, there exists $\mathfrak{E} = (\mathfrak{T}_1, \mathfrak{T}_2, \dots, \mathfrak{T}_m) \in \mathcal{A}$ such that $\mathfrak{E}'$ is obtained from $\mathfrak{E}$ via a “prune” move (a “grow” move will now break the defining constraint of $\mathcal{A}$.) We now show that either $\text{Bias}^2(\mathfrak{E}'; f^*) > 0$ or $\text{df}(\mathfrak{E}') \geq \text{df}(\mathfrak{E}_{\text{bad}}) + \min_{1 \leq i \leq m'} \dim_1(f_i) - 2$.

Let $\mathfrak{T}_{i_0}$ be the tree that is pruned, and supposed the prune split is (j, t) , occurring on a node $\mathfrak{t}$. Since $\mathcal{F}(\mathfrak{T}'_{i_0}) \not\supseteq \mathbb{V}_{i_0}$, (j, t) must be on a feature split on in $\mathfrak{T}_{i_0}^*$ and on a knot $\xi_{j,k}$ for f_j . Furthermore, because only “grow” and “prune” moves are allowed, and this is the first time $\mathcal{F}(\mathfrak{T}'_{i_0}) \not\supseteq \mathbb{V}_{i_0}$, $\mathfrak{T}_{i_0}^*$ must be a subtree of $\mathfrak{T}_{i_0}$, and $\mathfrak{t}$ is an internal node of $\mathfrak{T}_{i_0}^*$. Using Lemma F.6, we have $\text{coverage}(j, t; \mathcal{F}(\mathfrak{T}'_{i_0})) \cap \mathfrak{t}^{-j} = \emptyset$.

Suppose that $\text{Bias}^2(\mathfrak{E}'; f^*) = 0$, i.e. $f^* \in \mathcal{F}(\mathfrak{E}')$. Let $I_1, I_2, \dots, I_m$ be the subsets of feature indices split on in trees $\mathfrak{T}_1^*, \mathfrak{T}_2^*, \dots, \mathfrak{T}_m^*$ respectively. We claim that there exists some tree $\mathfrak{T}_i$, $i \neq i_0$ such that for all choices of coordinates $\mathbf{x}_{I_i}$ in the index set I_i , there exists a choice of coordinates $\mathbf{z}_{-I_i \cup \{j\}}$ in the index set $\{1, 2, \dots, b\} \setminus (I_i \cup \{j\})$ such that $(\mathbf{x}_{I_i}, \mathbf{z}_{-I_i \cup \{j\}}) \in \text{coverage}(j, t; \mathcal{F}(\mathfrak{T}_i))$. This claim is proved as Lemma H.1 below. Assuming it for now, let $\mathcal{L}_1, \mathcal{L}_2, \dots, \mathcal{L}_L$ be the leaves of $\mathfrak{T}_i^*$. Consider one such leaf $\mathcal{L}_l$. Since $\mathcal{L}_l$ does not depend on the features in I_i , we may pick $(\mathbf{x}_{I_i}, \mathbf{z}_{-I_i}) \in \text{coverage}(j, t; \mathcal{F}(\mathfrak{T}_i))$ such that $(\mathbf{x}_{I_i}, \mathbf{z}_{-I_i}) \in \mathcal{L}_l$. By Lemma F.6, $\mathcal{L}_l$ in $\mathfrak{T}_i$ must contain a descendant that uses the splitting rule (j, t) . In particular, $\mathcal{L}_l$ is split in $\mathfrak{T}_i$, and this split can be represented by a local decision stump ψ_l . In this manner, we obtain $\psi_1, \psi_2, \dots, \psi_L$. Denote $\mathbb{U} = \text{span}(\psi_1, \psi_2, \dots, \psi_{L-1})$. We claim that $\mathbb{U} \cap \mathbb{V} = \emptyset$. To see this, take any $f = \sum_{l=1}^{L-1} a_l \mathbf{1}_{\mathcal{L}_l} \in \mathbb{U}$ and assume that $f \neq 0$. First note that $f \perp \mathbb{V}_i$ by orthogonality of local decision stump features from a single tree. Furthermore, f has jump locations with respect to features in I_i . But no function in $\mathbb{V}_{-i}$ depends on features in I_i , which means that $f \notin \mathbb{V}_{-i}$. This gives $f \notin \mathbb{V}$.

Let ψ' denote the local decision stump associated to the pruned split, and let $\mathbb{V}'$ denote its orthogonal complement in $\mathbb{V}$. We have $\mathbb{V}' \subset \mathcal{F}(\mathfrak{E}')$ and have further shown that $\mathbb{U} \subset \mathcal{F}(\mathfrak{E}')$. We therefore have

$$\begin{aligned} \text{df}(\mathfrak{E}') &= \dim(\mathcal{F}(\mathfrak{E}')) \\ &\geq \dim(\mathbb{V}') + \dim(\mathbb{U}) \\ &= \text{df}(\mathfrak{E}_{\text{bad}}) - 1 + L - 1 \\ &\geq \text{df}(\mathfrak{E}_{\text{bad}}) + \min_{1 \leq i \leq m'} \dim_1(f_i) - 2. \end{aligned}$$

Here, the first inequality follows from the trivial intersection of $\mathbb{V}'$ and $\mathbb{U}$.

Using Proposition 4.1 and Proposition 4.2, there is a $1 - \delta/2$ event over which

$$\begin{aligned} \Delta \text{BIC}(\mathfrak{E}', \mathfrak{E}) &= \begin{cases} \frac{n}{\sigma^2} \text{Bias}^2(\mathfrak{E}'; f^*) + O\left(\sqrt{n \log(|\mathcal{B}|/\delta)}\right) & \text{if } \text{Bias}^2(\mathfrak{E}'; f^*) > 0 \\ \log n(\text{df}(\mathfrak{E}') - \text{df}(\mathfrak{E}_{\text{bad}})) + O\left(\sqrt{\log(|\mathcal{B}|/\delta)}\right) & \text{otherwise.} \end{cases} \end{aligned} \quad (60)$$

Condition further on this event.

Conclusion. Applying Proposition 6.1 together with equations (59) and (60), while taking n large enough, we get a $1 - 2\delta$ probability event over which

$$E\{\tau_{\text{OPT}_m(f^*, k)}\} = \Omega\left(n^{a/2-1}\right),$$

where $k = \max_{1 \leq i \leq m'} \dim_1(f_i) - \min_{1 \leq i \leq m'} \dim_1(f_i) - 1$ and $a = \min_{1 \leq i \leq m'} \dim_1(f_i)$.

Lemma H.1 (Existence of tree covering a cylinder). *There exists some tree $\mathfrak{T}_i$, $i \neq i_0$ such that for all choices of coordinates $\mathbf{x}_{I_i}$ in the index set I_i , there exists a choice of coordinates $\mathbf{z}_{-I_i \cup \{j\}}$ in the index set $\{1, 2, \dots, b\} \setminus (I_i \cup \{j\})$ such that $(\mathbf{x}_{I_i}, \mathbf{z}_{-I_i \cup \{j\}}) \in \text{coverage}(j, t; \mathcal{F}(\mathfrak{T}_i))$.*

Proof. Let $r_1, r_2, \dots, r_m$ be a permutation of $\{1, 2, \dots, m\}$, with $r_1 = i_0$ (this implies that $j \in I_{r_1}$). For $k = 2, 3, \dots, m-1$, set $J_k = \cup_{i \geq k} I_{r_i}$, and set $\mathbb{V}^k = \mathcal{F}(\mathfrak{T}_{r_k}) + \mathcal{F}(\mathfrak{T}_{r_{k+1}}) + \dots + \mathcal{F}(\mathfrak{T}_{r_m})$.

We make the following observation: For some k , suppose there exists a cylinder set $\mathcal{C}_k \subset \text{coverage}(j, t; \mathbb{V}^k)$ that does not depend on any feature in J_k . Suppose that there exists some $\mathbf{x}_{I_{r_k}}$ in the projection of $\mathcal{C}$ to coordinates in I_{r_k} such that $(\mathbf{x}_{I_{r_k}}, \mathbf{z}_{-I_i \cup \{j\}}) \notin \text{coverage}(j, t; \mathcal{F}(\mathfrak{T}_{r_k}))$ for all choices of $\mathbf{z}_{-I_i \cup \{j\}}$. Then taking the intersection of $\mathcal{C}_k$ and $\{\mathbf{x}_{I_{r_k}}\} \times \{1, 2, \dots, b\}^{-I_{r_k} \cup \{j\}}$ gives a cylinder set $\mathcal{C}_{k+1}$ that does not depend on any feature in J_{k+1} and such that $\mathcal{C}_{k+1} \subset \text{coverage}(j, t; \mathbb{V}^{k+1})$.

Now, since $\mathfrak{t}^{-j} \cap \text{coverage}(j, t; \mathcal{F}(\mathfrak{T}_{r_1})) = \emptyset$, we have $\mathfrak{t}^{-j} \subset \text{coverage}(j, t; \mathbb{V}^2)$. Since $\mathfrak{t}^{-j}$ is an internal node of $\mathfrak{T}_{r_1}^*$, it is a cylinder set that does not depend on any feature in J_2 . By applying the above observation inductively on $k = 2, 3, \dots, m-1$, we obtain the statement of the lemma. $\square$

I Proof of Proposition 6.1

Proposition 6.1 will follow almost immediately from the following more general result.

Proposition I.1 (General statement for recipe). *Let $X_0, X_1, \dots$ be an irreducible and aperiodic discrete time Markov chain on a finite state space Ω , with stationary distribution π . Let $x \in \Omega$ be a state and $\mathcal{C} \subset \Omega$ be a subset such that their hitting times from an initial state $x_0 \in \Omega$ satisfy*

$$P\{\tau_x < \tau_{\mathcal{C}} \mid X_0 = x_0\} \geq c$$

for some constant c . Let $\mathcal{B} \subset \Omega$ be a subset such that every path from x to $\mathcal{C}$ intersects $\mathcal{B}$. Then the hitting time of $\mathcal{C}$ satisfies

$$E\{\tau_{\mathcal{C}} \mid X_0 = x_0\} \geq \frac{c\pi(x)}{\pi(\mathcal{B})}.$$

Proof. By conditioning on the event $\{\tau_x < \tau_{\mathcal{C}}\}$, we calculate

$$\begin{aligned} E\{\tau_{\mathcal{C}} \mid X_0 = x_0\} &\geq E\{\tau_{\mathcal{C}} \mathbf{1}\{\tau_x < \tau_{\mathcal{C}}\} \mid X_0 = x_0\} \\ &= E\{\tau_{\mathcal{C}} \mid X_0 = x_0, \tau_x < \tau_{\mathcal{C}}\} P\{\tau_x < \tau_{\mathcal{C}} \mid X_0 = x_0\}. \end{aligned}$$

The second multiplicand on the right is lower bounded by c by assumption, so we just need to bound the first one. Using the strong Markov property, we first lower bound this as:

$$\begin{aligned} E\{\tau_{\mathcal{C}} \mid X_0 = x_0, \tau_x < \tau_{\mathcal{C}}\} &= E\{\tau_{\mathcal{C}} \mid X_0 = x\} + E\{\tau_x \mid \tau_x < \tau_{\mathcal{C}}\} \\ &\geq E\{\tau_{\mathcal{C}} \mid X_0 = x\}. \end{aligned} \tag{61}$$

Let Z denote the number of times (X_t) returns to x before hitting $\mathcal{B}$. Then, assuming that $X_0 = x$, we have the inequalities

$$\tau_{\mathcal{C}} \geq \tau_{\mathcal{B}} \geq Z + 1.$$

Note that $Z + 1$ is a geometric random variable with success probability

$$p = \{\tau_x^+ > \tau_{\mathcal{B}} \mid X_0 = x\},$$

where τ_x^+ is the first return time to x , i.e.

$$\min\{t > 0: X_t = x\}.$$

We therefore continue (61) to get

$$E\{\tau_{\mathcal{C}} \mid X_0 = x\} \geq \frac{1}{P\{\tau_x^+ > \tau_{\mathcal{B}} \mid X_0 = x\}}. \tag{62}$$

We next write this probability in terms of another random variable W , which we define to be the number of visits to states in $\mathcal{B}$ before returning to x , when the chain is started at $X_0 = x$. We then have

$$P\{\tau_x^+ > \tau_{\mathcal{B}} \mid X_0 = x\} = P\{W \geq 1\} \leq E\{W\}. \tag{63}$$

To bound this expectation, for each $y \in \mathcal{B}$, let W_y denote the number of visits to y before returning to x , and observe that $W = \sum_{y \in \mathcal{B}} W_y$. Let π denote the unique stationary distribution of (X_t) . Using Lemma I.2, we then have $E\{W_y\} \leq \pi(y)/\pi(x)$. Adding up these inequalities, we get

$$E\{W\} = \sum_{y \in \mathcal{B}} E\{W_y\} \leq \frac{\pi(y)}{\pi(x)} = \frac{\pi(\mathcal{B})}{\pi(x)}. \tag{64}$$

Combining equations (62), (63), and (64) completes the proof. $\square$

Proof of Proposition 6.1. It is clear that the Markov chain induced by a run of the BART sampler is irreducible and aperiodic, with stationary distribution given by the marginal posterior $p(\mathfrak{E}|\mathbf{y})$. We hence use Proposition I.1 to get

$$E\{\tau_{\text{OPT}_m(f^*, k)}\} = \Omega\left(\frac{p(\mathfrak{E}_{\text{bad}}|\mathbf{y})}{p(\mathcal{B}|\mathbf{y})}\right).$$

We now compute

$$\begin{aligned} \frac{p(\mathfrak{E}_{\text{bad}}|\mathbf{y})}{p(\mathcal{B}|\mathbf{y})} &\geq \frac{1}{|\mathcal{B}|} \min_{\mathfrak{E} \in \mathcal{B}} \frac{p(\mathfrak{E}_{\text{bad}}|\mathbf{y})}{p(\mathfrak{E}|\mathbf{y})} \\ &= \frac{1}{|\mathcal{B}|} \exp\left(\min_{\mathfrak{E} \in \mathcal{B}} \{\log p(\mathfrak{E}_{\text{bad}}|\mathbf{y}) - \log p(\mathfrak{E}|\mathbf{y})\}\right) \\ &\geq \frac{1}{|\mathcal{B}|} \exp\left(\frac{1}{2} \min_{\mathfrak{E} \in \mathcal{B}} \Delta\text{BIC}(\mathfrak{E}, \mathfrak{E}_{\text{bad}}) - O\left(\sqrt{\log(|\mathcal{B}|/\delta)}\right)\right) \\ &= \Omega\left(\exp\left(\frac{1}{2} \min_{\mathfrak{E} \in \mathcal{B}} \Delta\text{BIC}(\mathfrak{E}, \mathfrak{E}_{\text{bad}})\right)\right), \end{aligned}$$

where the second inequality follows from Proposition 4.2 and holds with probability at least $1 - \delta$. $\square$

Lemma I.2 (Bounding number of visits). *Let $X_0, X_1, \dots$ be an irreducible and aperiodic discrete time Markov chain on a finite state space Ω , with stationary distribution π . For any two states $x, y \in \Omega$, we have*

$$\mathbb{E}\{\text{number of visits to } y \text{ before returning to } x | X_0 = x\} = \frac{\pi(y)}{\pi(x)}. \quad (65)$$

Proof. Fix x and denote the quantity on the left side of equation (65) by $\tilde{\pi}(y)$. We may then rewrite equations (1.25) and (1.26) of Levin et al. (2006) in our notation as follows:

$$\begin{aligned} \pi(y) &= \frac{\tilde{\pi}(y)}{E\{\tau_x^+ | X_0 = x\}}, \\ \pi(x) &= \frac{1}{E\{\tau_x^+ | X_0 = x\}}, \end{aligned}$$

where τ_x^+ is the first return time to x . Taking the ratio of the two equations completes the proof. $\square$

J Invariance of PEM Dimension to Change of Measure

Lemma J.1 (Characterization of subspace dimension). *Let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ be vectors in an inner product space. Let $\mathbf{G}$ be the Gram matrix of these vectors, in other words, its (i, j) entry satisfies $G_{ij} = \langle \mathbf{v}_i, \mathbf{v}_j \rangle$ for $1 \leq i, j \leq n$. Then the dimension of the subspace spanned by $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ is equal to the number of nonzero eigenvalues of $\mathbf{G}$.*

Proof. By restricting to a linearly independent set, it suffices to show that $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ are linearly independent if and only if $\mathbf{G}$ is invertible. For the forward direction, suppose $\mathbf{G}$ is not invertible, then there exists a vector of coefficients $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n)$ such that $\boldsymbol{\alpha}^T \mathbf{G} \boldsymbol{\alpha} = 0$. But in that case, we have

$$0 = \boldsymbol{\alpha}^T \mathbf{G} \boldsymbol{\alpha} = \left\langle \sum_{i=1}^n \alpha_i \mathbf{v}_i, \sum_{i=1}^n \alpha_i \mathbf{v}_i \right\rangle.$$

By definition of the inner product, this means that $\sum_{i=1}^n \alpha_i \mathbf{v}_i = 0$, contradicting linear independence. The reverse direction is similar. $\square$

Lemma J.2 (Invariance of dimension to covariate distribution). *Let ν and ν' be two measures on a compact covariate space $\mathcal{X}$ that are absolutely continuous with respect to each other. Let $v_1, v_2, \dots, v_n \in L^2(\mathcal{X}, \nu)$. Then $v_1, v_2, \dots, v_n$ span the same subspace in both $L^2(\mathcal{X}, \nu)$ and $L^2(\mathcal{X}, \nu')$.*

Proof. By restricting to a linearly independent set, it suffices to show that $v_1, v_2, \dots, v_n$ are linearly independent in $L^2(\mathcal{X}, \nu)$ if and only if they are linearly independent in $L^2(\mathcal{X}, \nu')$. By definition of absolute continuity, there exists a constant $c > 0$ such that

$$c^{-1} \int_{\mathcal{X}} f(x) d\nu(x) \leq \int_{\mathcal{X}} f(x) d\nu'(x) \leq c \int_{\mathcal{X}} f(x) d\nu(x)$$

for any $f \in L^2(\mathcal{X}, \nu)$. Let $\mathbf{G}$ and $\mathbf{G}'$ be their Gram matrices in $L^2(\mathcal{X}, \nu)$ and $L^2(\mathcal{X}, \nu')$ respectively. Let $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n)$ be any vector of coefficients. Then we have

$$\begin{aligned} \boldsymbol{\alpha}^T \mathbf{G} \boldsymbol{\alpha} &= \int_{\mathcal{X}} \left(\sum_{i=1}^n \alpha_i v_i(x) \right)^2 d\nu(x) \\ &\leq c \int_{\mathcal{X}} \left(\sum_{i=1}^n \alpha_i v_i(x) \right)^2 d\nu'(x) \\ &= c \boldsymbol{\alpha}^T \mathbf{G}' \boldsymbol{\alpha}. \end{aligned}$$

Similarly, we get $\boldsymbol{\alpha}^T \mathbf{G}' \boldsymbol{\alpha} \leq c \boldsymbol{\alpha}^T \mathbf{G} \boldsymbol{\alpha}$. $\square$

This shows that whenever the covariate distribution ν has full support, it is equivalent to the uniform measure.

K Real Data Simulations

We further investigated the effect of training sample size on the mixing performance of BART as we fit it to a number of real world regression datasets. We studied several datasets used in the random forest paper along with three of the largest non-redundant datasets from the PMLB benchmark. Details on the datasets are provided below in Table 2.

Name	Samples	Features
Breast tumor (Romano et al., 2021)	116640	9
California housing (Pace and Barry, 1997)	20640	8
Echo months (Romano et al., 2021)	17496	9
Satellite image (Romano et al., 2021)	6435	36
Abalone (Nash et al., 1994)	4177	8
Diabetes (Efron et al., 2004)	442	10

Table 2: Real world datasets studied.

Furthermore, a heuristic estimate of the signal-to-noise ratio present in each dataset can be obtained by inspecting the simulation results in Tan et al. (2022b). California housing and Satellite image have relatively high SNR (the R^2 of random forest is larger than 0.8), whereas all other datasets have relatively low SNR (the R^2 of random forest is lower than 0.6.)

For each dataset we set aside 15% of the overall dataset as a test set. The remaining data is used as the training data set and is subsampled to create a variety of sample sizes. For each subsample proportion, we sample a random set of the training data and fit the BART algorithm on this subset of the training data with 1, 2, and 5 chains. This is repeated for 100 Monte Carlo iterations for each sample size and the RMSE on the test set is evaluated each time. The remainder of the simulation set-up is the same as that of Experiment 3, as described in Section 8.

Results. The results are displayed in Figure 8. Note that instead of plotting RMSE, we have chosen to plot relative RMSE, which measures the ratio between the RMSE obtained from multiple chains and that obtained for a single chain for a given data setting. Since we are working with real datasets, error is measured with respect to the observed responses, rather than a true regression function. In almost all of the real data sets, we see that increasing the number of chains in the BART algorithm consistently decreases the RMSE of the predictions, and the performance gap grows with the training sample size. This provides further evidence that the poor mixing performance of BART applies to real world datasets. On the other hand, the effect seems to be much less significant compared to that for the simulated datasets studied in Experiment 3. This is unsurprising as the RMSE for real datasets incorporates and hence is inflated by aleatoric uncertainty in the responses. Notably, in the data sets with higher SNR (California housing and Satellite image), the performance gap remains significant.




Figure 8: RMSE for the predictions from BART run with several different numbers of chains relative to the RMSE obtained from running BART with a single chain. The X axis displays the percentage of the training data that is sampled for each Monte Carlo replication. The RMSE is calculated over an independent test set consisting of 15% of the overall data.

L Additional results for Experiment 3

In order to further explore the effect of the number of chains on mixing performance, we also repeated Experiment 3 using 20 chains. The results are shown in Figure 9




Figure 9: The RMSE (top panel) and empirical coverage (bottom panel) of BART improve if we average posterior samples from multiple sampler chains, given a fixed total budget of posterior samples. The relative performance gaps increases with the number of training samples, providing evidence that the tendency of HPDR hitting time to grow with training sample size is consistent across a wide range of DGPs. The biggest improvement seems to occur when increasing the number of chains from 1 to 2 and there seems to be diminishing returns thereafter. Both RMSE and coverage are calculated on an independent test set and are averaged over 100 experimental replicates, with error bars representing $\pm 1.96SE$.