

POSTERIOR CONCENTRATION FOR BAYESIAN REGRESSION TREES AND FORESTS

BY VERONIKA ROČKOVÁ AND^{*,‡}, STÉPHANIE VAN DER PAS^{†,§}

Revision 29 May 2019

First Draft: 29 August 2017

University of Chicago[‡]

Leiden University[§]

Since their inception in the 1980's, regression trees have been one of the more widely used non-parametric prediction methods. Tree-structured methods yield a histogram reconstruction of the regression surface, where the bins correspond to terminal nodes of recursive partitioning. Trees are powerful, yet susceptible to over-fitting. Strategies against overfitting have traditionally relied on pruning greedily grown trees. The Bayesian framework offers an alternative remedy against overfitting through priors. Roughly speaking, a good prior charges smaller trees where overfitting does not occur. While the consistency of random histograms, trees and their ensembles has been studied quite extensively, the theoretical understanding of the Bayesian counterparts has been missing. In this paper, we take a step towards understanding why/when do Bayesian trees and forests not overfit. To address this question, we study the speed at which the posterior concentrates around the true smooth regression function. We propose a *spike-and-tree* variant of the popular Bayesian CART prior and establish new theoretical results showing that regression trees (and forests) (a) are capable of recovering smooth regression surfaces (with smoothness not exceeding one), achieving optimal rates up to a log factor, (b) can adapt to the unknown level of smoothness and (c) can perform effective dimension reduction when $p > n$. These results provide a piece of missing theoretical evidence explaining why Bayesian trees (and additive variants thereof) have worked so well in practice.

1. Non-parametric Regression Setup. The remarkable empirical success of Bayesian tree-based regression [15, 18, 14] has raised considerable interest in understanding why and when these methods produce good results. Despite their extensive use in practice, theoretical justifications have, thus

* veronika.rockova@chicagobooth.edu;

This work was supported by the James S. Kemper Foundation Faculty Research Fund at the University of Chicago Booth School of Business.

† svdpas@math.leidenuniv.nl

Keywords and phrases: Additive Regression, Asymptotic Minimality, BART, Bayesian CART, Posterior Concentration, Recursive Partitioning, Regression Trees

far, been unavailable. To narrow this yawning gap, we consider the fundamental problem of making inference about an unknown regression function.

Our setup consists of the nonparametric regression model

$$(1.1) \quad Y_i = f_0(\mathbf{x}_i) + \varepsilon_i, \quad \varepsilon_i \stackrel{iid}{\sim} \mathcal{N}(0, 1),$$

where output variables $\mathbf{Y}^{(n)} = (Y_1, \dots, Y_n)'$ are related in a stochastic fashion to a set of p potential covariates $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})'$, $1 \leq i \leq n$. We assume that the covariates $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})'$ are fixed and have been rescaled so that $\mathbf{x}_i \in [0, 1]^p$. The true unknown regression surface $f_0(\mathbf{x})$ will be assumed to be smooth, possibly involving only a small fraction q_0 of the p potential covariates.

In the absence of a parametric model, a natural strategy to estimate the unknown regression function is by partitioning the covariate space into cells and then estimating the function locally within each cell from available observations. Such strategies yield histogram reconstructions of the regression surface and have been analyzed theoretically by multiple authors [34, 19, 41]. Regression trees [9] are among the most popular data-dependent histogram methods, where the partitioning scheme is obtained through nested parallel axis splitting. Trees are an integral constituent of ensemble methods that aggregate single tree learners into forests to boost prediction [8]. Tree-based regression, either single or ensemble, is arguably one of the most popular machine learning tools today. In particular, Bayesian variants of these methods (Bayesian CART and BART) [14, 18] have earned a prominent role as one of the top machine learners. While consistency results for classical trees and random forests have been available [25, 26, 21, 9, 4, 5, 36, 45], theory on the also very widely used Bayesian counterparts is non-existent. Our goal in this paper is to provide the first frequentist optimality results for *Bayesian* trees, and their ensembles.

Most of the work on Bayesian nonparametric regression has revolved around Gaussian processes [43, 3, 46]. While there are multiple results on recursive partitioning (or histogram) priors for Bayesian density estimation [10, 37, 33, 11] (or non-linear autoregression [23]), the literature on Bayesian regression histograms is far more deserted. One fundamental contribution is due to Coram and Lalley [16], who showed consistency of Bayesian binary regression with uniform mixture priors on step functions and with one predictor. More recently, van der Pas and Ročková [42] considered a similar setup for estimating step mean functions in Gaussian regression, again with a single predictor. This paper goes far beyond that framework, addressing (a) the full-fledged high-dimensional setup with a diverging number of potential covariates, (b) tree ensembles for additive regression.

The purpose of this paper is to study the rate of convergence of posterior distributions induced by step function priors on the regression surface when $p > n$. The speed of convergence is measured by the size of the smallest shrinking ball around f_0 that contains most of the posterior probability. In pioneering works, Ghosal, Ghosh and Van der Vaart [23] and Shen and Wasserman [38] obtained rates of convergence for infinite-dimensional parametric models with iid observations in terms of the size of the model (measured by the metric entropy) and concentration rate of the prior around f_0 . These results were later extended to infinite-dimensional models that are *not iid* by Ghosal and Van der Vaart [24]. Their general conceptual framework serves as an umbrella for our development.

1.1. *Our Contributions.* We initially assume that f_0 is Hölder continuous (with smoothness not exceeding one) and may depend only on a small fraction of q_0 predictors. The optimal rate of estimation of a q_0 -variable function, which is known to be α -smooth, is $n^{-\alpha/(2\alpha+q_0)}$ [39]. Our first result shows that, with suitable regularization priors, single Bayesian regression trees achieve this minimax rate (up to a log factor). In other words, the posterior behaves nearly as well as if we knew α and the number of active covariates q_0 , concentrating at a rate that only depends on the number of active predictors. This is the first optimality result for Bayesian regression trees, demonstrating their adaptability and reluctance to overfit in high-dimensional scenarios with $p > n$. The regularization is achieved through our proposed *spike-and-tree* prior, a new variant of the Bayesian CART prior for dimension reduction and model-free variable selection. Going further, we show that Bayesian *additive regression trees* also achieve (near) minimax-rate optimal performance when approximating a single smooth function. Finally, the tree ensembles are *also* shown to be certifiably optimal when the true function is an actual sum of smooth functions, again concentrating at a near minimax rate.

1.2. *Notation.* The notation \lesssim will be used to denote inequality up to a constant, $a \asymp b$ denotes $a \lesssim b$ and $b \lesssim a$ and $a \vee b$ denotes $\max\{a, b\}$. The ε -covering number of a set Ω for a semimetric d , denoted by $N(\varepsilon, \Omega, d)$, is the minimal number of d -balls of radius ε needed to cover set Ω . We denote by $\phi(\cdot; \sigma^2)$ the normal density with zero mean and variance σ^2 and by $P_f^n = \bigotimes P_{f,i}$ the n -fold product measure of the n independent observations under (1.1) with a regression function f . By $\mathbb{P}_n^{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{x}_i}$ we denote the empirical distribution of the observed covariates and let $\|\cdot\|_n$ denote the empirical norm on $L_2(\mathbb{P}_n^{\mathbf{x}})$. With $\|\cdot\|_2$ we denote the standard Euclidean norm. For $\mathbf{x} \in \mathbb{R}^p$ we denote with $\mathbf{x}_{\mathcal{S}}$ the subvector of \mathbf{x} indexed by $\mathcal{S} \subseteq$

$\{1, \dots, p\}$. With \mathcal{C}^r we denote r -times continuously differentiable functions on $[0, 1]$.

1.3. *Outline.* We outline our goals and strategy in Section 2. We then review several useful concepts for analyzing recursive partitioning schemes in Section 3. In Section 4, we state our first main result on the posterior concentration for Bayesian CART. In Section 5, we develop tools for analyzing Bayesian additive regression trees and show their optimal posterior concentration in non-additive regression. Section 6 presents the final development of our theory concerning the recovery of an additive regression function with additive trees. We conclude with a discussion in Section 7. The proofs of our main theorems are presented in Sections 8, 9 and 10.

2. Background. In this section we lay down rudiments of our modus operandi. Our setup comprises a sequence of statistical experiments with observations $\mathbf{Y}^{(n)} = (Y_1, \dots, Y_n)'$ and models P_f^n defined in (1.1). Each model P_f^n is parametrized by a regression function $f : [0, 1]^p \rightarrow \mathbb{R}$ that lives in an infinite-dimensional space \mathcal{F} endowed with a prior distribution. With adequate priors, the posterior can exhibit nice frequentist properties, which get passed onto its location/scale summary measures. One such property is the ability to pile up in shrinking neighborhoods around the true regression function f_0 . The speed at which the shrinking occurs is the posterior concentration rate and assesses the quality of the posterior beyond just the mere fact that it is consistent. In our setup, we investigate such concentration properties in terms of $\|\cdot\|_n$ neighborhoods of f_0 , where

$$\|f - f_0\|_n^2 = \frac{1}{n} \sum_{i=1}^n [f(\mathbf{x}_i) - f_0(\mathbf{x}_i)]^2.$$

The key to our approach will be drawing upon the foundational posterior concentration theory for non-iid observations, laid down in the seminal paper by Ghosal and Van der Vaart [24].

Our results are obtained under a unifying hat of a general result which requires three conditions to hold. Namely, suppose that for a sequence $\varepsilon_n^2 \rightarrow 0$ such that $n\varepsilon_n^2$ is bounded away from zero and sets $\mathcal{F}_n \subset \mathcal{F}$ we have

$$(2.1) \quad \sup_{\varepsilon > \varepsilon_n} \log N\left(\frac{\varepsilon}{36}, \{f \in \mathcal{F}_n : \|f - f_0\|_n < \varepsilon\}, \|\cdot\|_n\right) \leq n\varepsilon_n^2$$

$$(2.2) \quad \Pi(f \in \mathcal{F} : \|f - f_0\|_n \leq \varepsilon_n) \geq e^{-dn\varepsilon_n^2}$$

$$(2.3) \quad \Pi(\mathcal{F} \setminus \mathcal{F}_n) = o(e^{-(d+2)n\varepsilon_n^2})$$

for some $d > 2$. Then it follows from Theorem 4 in [24] that the posterior distribution concentrates at the rate ε_n^2 , i.e.

$$(2.4) \quad \Pi \left(f \in \mathcal{F} : \|f_0 - f\|_n > M_n \varepsilon_n \mid \mathbf{Y}^{(n)} \right) \rightarrow 0$$

in $P_{f_0}^n$ -probability, as $n \rightarrow \infty$, for any $M_n \rightarrow \infty$.

The conditions of Ghosal and Van der Vaart [24] provide a very general recipe for showing posterior concentration in infinite-dimensional models. Our goal in this paper is to obtain tailored statements for Bayesian regression trees and forests. The major challenge will be (a) designing a sequence of approximating spaces (a sieve) $\mathcal{F}_n \subset \mathcal{F}$ and (b) endowing \mathcal{F} with a prior distribution such that the three conditions hold simultaneously for ε_n as small as possible. To this end, we will build on, and develop, tools for an analysis of recursive partitioning schemes.

Throughout this work, we assume that the true regression function f_0 is Hölder continuous and the smoothness parameter does not exceed one, (or an additive composition thereof), in the sense made precise below.

DEFINITION 2.1. *With \mathcal{H}_p^α we denote the space of uniformly α -Hölder continuous functions, i.e.*

$$\mathcal{H}_p^\alpha = \left\{ f : [0, 1]^p \rightarrow \mathbb{R}; \|f\|_{\mathcal{H}^\alpha} \equiv \sup_{\mathbf{x}, \mathbf{y} \in [0, 1]^p} \frac{|f(\mathbf{x}) - f(\mathbf{y})|}{\|\mathbf{x} - \mathbf{y}\|_2^\alpha} < \infty \right\},$$

where $\alpha \in (0, 1]$ and where $\|f\|_{\mathcal{H}^\alpha}$ is the Hölder coefficient.

The assumption $\alpha \leq 1$ is standard in the study of piecewise constant estimators and priors, see for example [24, 37, 22]. The reason for this limitation is that step functions are relatively rough; e.g. the approximation error of histograms for functions that are smoother than Lipschitz is at least of the order $1/K$, where K is the number of bins. The number of steps required to approximate a smooth function well is thus too large, creating a costly bias-variance tradeoff.

In some applications, it is reasonable to expect that the regression function f_0 depends only on a small fraction of input covariates. For a set of indices $\mathcal{S} \subseteq \{1, \dots, p\}$, we define

$$\mathcal{C}(\mathcal{S}) = \{f : [0, 1]^p \rightarrow \mathbb{R}; f \text{ is constant in directions } \{1, \dots, p\} \setminus \mathcal{S}\}.$$

We will consider two estimation regimes:

(R1) Regime 1: f_0 is α -Hölder continuous and depends on an unknown subset \mathcal{S}_0 of $|S_0| = q_0$ covariates, i.e. $f_0 \in \mathcal{H}_p^\alpha \cap \mathcal{C}(\mathcal{S}_0)$.

(R2) Regime 2: f_0 is an aggregate of T_0 α^t -Hölder continuous functions f_0^t , $1 \leq t \leq T_0$, each depending on an unknown subset \mathcal{S}_0^t of $|\mathcal{S}_0^t| = q_0^t$ covariates, i.e. $f_0(\mathbf{x}) = \sum_{t=1}^{T_0} f_0^t(\mathbf{x})$, where $f_0^t \in \mathcal{H}_p^{\alpha^t} \cap \mathcal{C}(\mathcal{S}_0^t)$.

For an estimation procedure to be successful in Regime 1, it needs to be doubly adaptive (with respect to α and q_0). We will show that both single trees and tree ensembles adapt accordingly, performing at a near-minimax rate. Regime 2 makes the performance discrepancies more apparent, where the additive structure of f_0 is appreciated by tree ensembles, which can achieve a faster convergence rate than single trees. A variant of Regime 2 was studied by [46] who derived minimax rates for estimating additive smooth functions and showed optimal concentration of additive Gaussian processes. Here, we approximate f_0 with step functions and their aggregates. We limit considerations to step functions that are underpinned by recursive partitioning schemes.

3. On Recursive Partitions. Tree-based regression consists of first finding an underlying partitioning scheme that hierarchically subdivides a dataset into more homogeneous subsets, and then learning a piecewise constant function on that partition. This section serves to review several useful concepts for analyzing such nested partitioning rules that will be instrumental in our analysis.

3.1. General Partitions. Given $K \in \mathbb{N}$, we define a K -partition of $[0, 1]^p$ as a sequence $\boldsymbol{\Omega} = \{\Omega_k\}_{k=1}^K$ of K contiguous rectangles $\Omega_k \subset [0, 1]^p$, where $\cup_{k=1}^K \Omega_k = [0, 1]^p$. Sufficient conditions for consistency of regression histograms have traditionally revolved around two requirements on the partitioning cells Ω_k (Section 6.3 in [20]). The first one pertains to cell counts, where Ω_k 's should be large enough to capture a sufficient number of points to render local estimation meaningful. The second one pertains to cell sizes, where Ω_k 's should be small enough to detect local changes in the regression surface. We borrow some of these concepts for our theoretical analysis.

For the first requirement, let us formalize the notion of the cell size in terms of the empirical measure induced by observations $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$. For each cell Ω_k , we define by

$$(3.1) \quad \mu(\Omega_k) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}(\mathbf{x}_i \in \Omega_k)$$

the proportion of observations falling inside Ω_k . Throughout this work, we focus on partitions whose boxes can adaptively stretch and shrink, allowing the splits to arrange themselves in a data-dependent way [40]. The simplest

data-adaptive partition is based on statistically equivalent blocks [1, 20], where all cells have approximately the same number of points, i.e. $\mu(\Omega_k) \asymp 1/K$. We deviate from such a strict rule by allowing for imbalance and define the so called *valid* partitions.

DEFINITION 3.1. (*Valid Partitions*) Denote by $\Omega = \{\Omega_k\}_{k=1}^K$ a partition of $[0, 1]^p$. We say that Ω is valid if

$$(3.2) \quad \mu(\Omega_k) \geq \bar{C}^2/n \quad \text{for all } k = 1, \dots, K$$

for some constant $\bar{C}^2 \in \mathbb{N} \setminus \{0\}$.

Valid partitions have non-empty cells, where the allocation *does not* need to be balanced. In *balanced partitions* (introduced in van der Pas and Ročková [42]), the cells satisfy $\frac{C_{min}^2}{K} \leq \mu(\Omega_k) \leq \frac{C_{max}^2}{K}$ for some $C_{min} < 1 < C_{max}$. One prominent example of such balanced partitions is the median tree (or a k - d tree) [2], which will be discussed in the next section and will be used as a benchmark tree approximation towards establishing Condition 2.2.

For the second requirement, let us formalize the notion of the cell size in terms of the local spread of the data. To this end, we introduce the partition *diameter* [44, 27].

DEFINITION 3.2. (*Diameter*) Denote by $\Omega = \{\Omega_k\}_{k=1}^K$ a partition of $[0, 1]^p$ and by $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ a collection of data points in $[0, 1]^p$. For an index set $\mathcal{S} \subseteq \{1, \dots, p\}$, we define a diameter of Ω_k w.r.t. \mathcal{S} as

$$\text{diam}(\Omega_k; \mathcal{S}) = \max_{\mathbf{x}, \mathbf{y} \in \Omega_k \cap \mathcal{X}} \|\mathbf{x}_{\mathcal{S}} - \mathbf{y}_{\mathcal{S}}\|_2.$$

and with $\text{diam}(\Omega; \mathcal{S}) = \sqrt{\sum_{k=1}^K \mu(\Omega_k) \text{diam}^2(\Omega_k; \mathcal{S})}$ we define a diameter of the entire partition Ω w.r.t. \mathcal{S} .

The diameter of Ω_k corresponds to the largest $\|\cdot\|_2$ distance between \mathcal{S} -coordinate projections of two points that fall inside Ω_k . This is one of the more flexible notions of a cell diameter, which takes into account the data itself, not just the physical cell size. Traditionally, bias and convergence rate analysis of tree-based estimators have been characterized in terms of the cell diameters. The rate depends on how fast the diameters shrink as we move down the tree: the more rapidly, the better. As will be seen later in Section 3.3, controlling the diameter will be essential for obtaining tight bounds on the approximation error.

3.2. *Tree Partitions.* We are ultimately interested in partitions that can be obtained with *nested* axis-parallel splits. Such partitions can be represented by a tree diagram, a hierarchical arrangement of nodes. We will focus on binary tree partitions, where each split yields two children nodes. Namely, starting from a parent node $[0, 1]^p$, a binary tree partition is grown by successively applying a splitting rule on a chosen internal node, say $\Omega_k^* \subset [0, 1]^p$. The node Ω_k^* is split into two cells by a perpendicular bisection along one of the p coordinates, say j , at a chosen observed value $\tau \in \mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$. The newborn cells $\{\mathbf{x} \in \Omega_k^* : x_j \leq \tau\}$ and $\{\mathbf{x} \in \Omega_k^* : x_j > \tau\}$ constitute two rectangular regions of $[0, 1]^p$, which can be split further (to become internal nodes) or end up being terminal nodes. The terminal cells after $K - 1$ splits then yield a box-shaped (tree) partition $\{\Omega_k\}_{k=1}^K$.

Unlike *dyadic* trees, where the threshold τ is preset at a midpoint of the rectangle Ω_k along the j^{th} direction, we allow for splits at available observations \mathcal{X} . Such data-dependent splits are integral to Bayesian CART and BART implementations [15, 18, 14]. With more opportunities for each split, many more tree topologies can be generated. However, since the tree partitions are arranged in a nested fashion, their combinatorial complexity is not too large (as shown in Lemma 3.1 below).

We will denote each tree-structured K -partition by $\mathcal{T}^K = \{\Omega_k\}_{k=1}^K$. With \mathcal{V}_S^K we denote a *family of valid tree partitions* of $[0, 1]^p$, obtained by splitting $K - 1$ times at observed values in \mathcal{X} along each coordinate direction inside $\mathcal{S} \subseteq \{1, \dots, p\}$ at least once. In particular, each tree $\mathcal{T} \in \mathcal{V}_S^K$ is valid according to Definition 3.1 and uses up all covariates in \mathcal{S} for splits, where \mathcal{S} can be regarded as an index set of active predictors. We will refer to the partitioning number $\Delta(\mathcal{V}_S^K)$ (in a similar vein as in [34]) as the overall number of distinct partitions of \mathcal{X} that can be induced by members of the partitioning family \mathcal{V}_S^K .

LEMMA 3.1. *Denote by $\mathcal{S} \subseteq \{1, \dots, p\}$ an index set of active covariates. Let \mathcal{V}_S^K denote the set of valid tree partitions obtained with $K - 1$ splits. Then*

$$\Delta(\mathcal{V}_S^K) = \frac{|\mathcal{S}|^{K-1} n!}{(n - K + 1)!}.$$

PROOF. This follows from the recursive formula $\Delta(\mathcal{V}_S^K) = \Delta(\mathcal{V}_S^{K-1})(n - K + 2)|\mathcal{S}|$, where we use the fact that there are $\Delta(\mathcal{V}_S^{K-1})$ possible trees with $K - 1$ cells which have altogether $n - K + 2$ potential next splits along one of the $|\mathcal{S}|$ coordinates. \square

Lemma 3.1 will be fundamental for understanding the combinatorial richness of trees and for obtaining bounds on the covering numbers towards

establishing Condition (2.1).

REMARK 3.1. (*The k-d tree partition*) We now pause to revisit one of the most popular space partitioning structures, the k-d tree partition [2]. Such a partition $\widehat{\mathcal{T}}$ is constructed by cycling over coordinate directions in \mathcal{S} , where all nodes at the same level are split along the same axis. For a given direction $j \in \mathcal{S}$, each internal node, say Ω_k^* , will be split at a median of the point set (along the j^{th} axis). This split will pass $\lfloor \mu(\Omega_k^*)n/2 \rfloor$ and $\lceil \mu(\Omega_k^*)n/2 \rceil$ observations onto its two children, thereby roughly halving the number of points. After s rounds of splits on each variable, all K terminal nodes have at least $\lfloor n/K \rfloor$ observations, where $K = 2^{s|\mathcal{S}|}$. The k-d tree partitions are thus balanced in light of Definition 2.4 of van der Pas and Ročková [42]. While k-d trees are not adaptive to intrinsic dimensionality of data (in comparison with more flexible partitions such as random projections trees [44]), the diameters of k-d tree partitions can shrink fast, as long as the number of directions is not too large. In particular, $\text{diam}(\widehat{\mathcal{T}}; \mathcal{S}) \leq \sqrt{|\mathcal{S}|/K^{1/(2|\mathcal{S}|)}}$ for $K = 2^{s|\mathcal{S}|}$ for some $s \geq 1$ according to Proposition 6 of [44]. The k-d tree construction will be instrumental in establishing Condition (2.2).

3.3. *On Tree-structured Step Functions.* The second critical ingredient in building a tree regressor is learning the piecewise function on a given partition. In this section, we describe some facts about the approximating properties of such tree-structured step functions (further referred to as trees). The understanding of how well we can approximate a smooth regression surface will be vital for establishing Condition (2.2).

For a family of valid tree partitions $\mathcal{V}_{\mathcal{S}}^K$, we denote by

$$(3.3) \quad \mathcal{F}(\mathcal{V}_{\mathcal{S}}^K) = \left\{ f_{\mathcal{T}, \boldsymbol{\beta}} : [0, 1]^p \rightarrow \mathbb{R}; f_{\mathcal{T}, \boldsymbol{\beta}}(\mathbf{x}) = \sum_{k=1}^K \beta_k \mathbb{I}_{\Omega_k}(\mathbf{x}); \mathcal{T} \in \mathcal{V}_{\mathcal{S}}^K, \boldsymbol{\beta} \in \mathbb{R}^K \right\}$$

the set of all step functions supported by members of the tree partitioning family $\mathcal{V}_{\mathcal{S}}^K$. Each function is underpinned by a valid tree partition $\mathcal{T} = \{\Omega_k\}_{k=1}^K \in \mathcal{V}_{\mathcal{S}}^K$ and a vector of step heights $\boldsymbol{\beta} = (\beta_1, \dots, \beta_K)' \in \mathbb{R}^K$.

The cell diameters oversee how closely one can approximate $f \in \mathcal{H}_p^\alpha \cap \mathcal{C}(\mathcal{S})$ with $f_{\mathcal{T}, \boldsymbol{\beta}} \in \mathcal{F}(\mathcal{V}_{\mathcal{S}}^K)$ and it is desirable that they decay fast with K . Ideally, the approximation error should be no slower than $q^\gamma / K^{1/q}$ for some $\gamma > 0$, where $q = |\mathcal{S}|$. To get an intuitive insight into this requirement, consider the following perfectly regular partition: a q -dimensional ‘‘chessboard’’ that splits $[0, 1]^q$ into $K = s^q$ cubes of length $1/s = 1/K^{1/q}$. The maximal interpoint distance inside each cube will be at most $\sqrt{q}/K^{1/q}$. This

partition is, however, not adaptive and thereby less practical. It turns out that, under a mild requirement on the spread of the data points \mathcal{X} , the fast diameter decay is *also* guaranteed by the data-adaptive k - d trees mentioned in Remark 3.1. The “mild requirement” is formalized in our definition below.

DEFINITION 3.3. Denote by $\widehat{\mathcal{T}} = \{\widehat{\Omega}_k\}_{k=1}^K \in \mathcal{V}_{\mathcal{S}}^K$ the k - d tree where $\mathcal{S} \subseteq \{1, \dots, p\}$ and $K = 2^s |\mathcal{S}|$. We say that a dataset \mathcal{X} is (M, \mathcal{S}) -regular if

$$(3.4) \quad \max_{1 \leq k \leq K} \text{diam}(\widehat{\Omega}_k; \mathcal{S}) < M \sum_{k=1}^K \mu(\widehat{\Omega}_k) \text{diam}(\widehat{\Omega}_k; \mathcal{S})$$

for some large enough constant $M > 0$ and all $s \in \mathbb{N} \setminus \{0\}$.

The definition states that in a regular dataset, the maximal diameter in the k - d tree partition should not be much larger than a “typical” diameter. This condition ensures that, as more and more data points are collected, the data conform to a structure that does not permit outliers in active directions \mathcal{S} . For example, a fixed design on a regular grid (including directions \mathcal{S}) will satisfy (3.4). Our notion of regularity goes farther by allowing (a) the predictors to be correlated, (b) the points to scatter unevenly and/or close to a lower-dimensional manifold. The manifold, however, should be varying in active directions \mathcal{S} and do so sufficiently *smoothly* (or be monotone) so that the cells in the k - d tree do not contain isolated clouds of points. For example, data points arising from distributions with atomic marginals (in active directions) would violate regularity.

As will be seen in the following lemma, for regular datasets, k - d trees have a faster diameter decay (roughly halving the cell diameters after one round of $q = |\mathcal{S}|$ splits), thereby attaining better approximation error. The following lemma is an important element of our proof skeleton, showing that there exists a tree (a k - d tree) that approximates well.

LEMMA 3.2. Assume $f \in \mathcal{H}_p^\alpha \cap \mathcal{C}(\mathcal{S})$, where $|\mathcal{S}| = q$ and $\alpha \leq 1$, and that \mathcal{X} is (M, \mathcal{S}) -regular. Then there exists a tree-structured step function $f_{\widehat{\mathcal{T}}, \widehat{\beta}} \in \mathcal{V}_{\mathcal{S}}^K$ for some $K = 2^{s_q}$ with $s \in \mathbb{N} \setminus \{0\}$ such that

$$(3.5) \quad \|f - f_{\widehat{\mathcal{T}}, \widehat{\beta}}\|_n \leq \|f\|_{\mathcal{H}^\alpha} C M^\alpha q / K^{\alpha/q}$$

for some $C > 0$.

PROOF. Supplemental Material (Section 3). □

Controlling the approximation error is only one of the critical aspects in our theoretical study. The second will be monitoring the complexity of our approximating function classes. Finding the right balance between the two will be instrumental for obtaining optimal performance.

Now that we have developed the necessary tools, we can dive into the main results of the paper.

4. Adaptive Dimension Reduction with Trees. In Regime 1, we assume that the target regression surface $f_0 \in \mathcal{H}_p^\alpha \cap \mathcal{C}(\mathcal{S}_0)$, although initially conceived as a function of $\mathbf{x} = (x_1, \dots, x_p)'$, in fact depends only on a small fraction of q_0 features $\mathbf{x}_{\mathcal{S}_0}$, where $\mathcal{S}_0 \subset \{1, \dots, p\}$. If an oracle were able to isolate \mathcal{S}_0 , the L_2 minimax rate would improve from $n^{-\alpha/(2\alpha+p)}$ to $n^{-\alpha/(2\alpha+q_0)}$ and it would be the fastest rate possible. When no such oracle information is available, [46] characterized the minimax rate as follows: $\varepsilon_n^2 = \lambda^2 (\sqrt{n}\lambda)^{-4\alpha/(2\alpha+q_0)} + \frac{q_0}{n} \log\left(\frac{p}{q_0}\right)$ where λ is the isotropic Hölder norm. The first term is the classical minimax risk for estimating a q_0 -dimensional smooth function. The second term is the penalty incurred by variable selection uncertainty. While the number of features p is not forbidden from growing to infinity much faster than n , we keep in mind that consistent estimation will only be possible in sparse contexts where q_0 is small relative to p and n (in which case the complexity penalty will be dominated by the first term in the minimax rate).

4.1. Spike-and-Tree Priors. Bayesian regression tree implementations that do not induce sparsity (when in fact present) are unfit for inference in high-dimensional setups. In particular, the traditional Bayesian CART prior [15] grows trees by splitting each node, indexed by the depth level d , with a probability $g(d) = \gamma/(1+d)^\alpha$, where $\alpha > 1, \gamma \in (0, 1)$ are tuning parameters. The splitting variable is picked from $\{x_1, \dots, x_p\}$ *uniformly* (i.e. with a fixed probability $\theta_i = 1/p$). Recently, [29] proposed an adaptive variant of this prior by placing a sparsity-inducing Dirichlet prior on the splitting proportions θ_i . This prior uses fewer variables in the tree construction and thereby is more reluctant to overfit. In another popular Bayesian CART implementation, [18] suggest directly placing a prior on K and a conditionally uniform prior on tree topologies with K bottom leaves. Again, in its original form, this prior will likely suffer from the curse of dimensionality, failing to harvest the intrinsic lower-dimensional structure. Here, we propose a fix to this problem. To make the Bayesian CART prior of [18] appropriate for high-dimensional setups, we propose a *spike-and-tree* variant by injecting one more layer: a complexity prior over the active set of predictors.

Bayesian models for feature selection have traditionally involved a hierarchy of priors over subset sizes $q = |\mathcal{S}|$ and subsets $\mathcal{S} \subseteq \{1, \dots, p\}$ [13, 12]. Instead of modeling the mean outcome as a linear functional of active predictors $\mathbf{x}_{\mathcal{S}}$, here we grow a tree from $\mathbf{x}_{\mathcal{S}}$. We begin by treating q_0 as unknown with an exponentially decaying prior [13]

$$(T1) \quad \pi(q) \propto c^{-q} p^{-aq}, \quad q = 0, 1, \dots, p, \quad \text{for some } a, c > 0.$$

Next, given the dimensionality q , we assume that all $\binom{p}{q}$ subsets \mathcal{S} of $q = |\mathcal{S}|$ covariates are a-priori equally likely, i.e.

$$(T2) \quad \pi(\mathcal{S} | q) = 1 / \binom{p}{q}.$$

Given q and the feature set \mathcal{S} , a tree is grown by splitting at least once on every coordinate inside \mathcal{S} , all the way down to K terminal nodes. If we knew α and q_0 , the optimal choice of K would be $K \asymp n^{q_0/(2\alpha+q_0)}$, for which the actual minimax rate could be achieved. For the more practical case when α and q_0 are both unknown, we shall assume that K arrived from a prior $\pi(k | q)$. As noted by [16], the tail behavior of $\pi(k | q)$ is critical for controlling the vulnerability/resilience to overfitting. We consider the Poisson distribution (constrained to $\mathbb{N} \setminus \{0\}$), as suggested by [18] in their Bayesian CART implementation. Namely, for $k \in \mathbb{N} \setminus \{0\}$ we have

$$(T3) \quad \pi(K) = \frac{\lambda^K}{(e^\lambda - 1)K!}, \quad K = 1, 2, \dots, \quad \text{for some } \lambda \in \mathbb{R}.$$

For its practical implementation, one would truncate its support to the maximum number of splits that can be made with n observations. When λ is small, (T3) is concentrated on models with smaller complexity where overfitting does not occur. Increasing λ leads to smearing the prior mass over partitions with more jumps. Similar complexity priors with an exponential decay $e^{-Ck \log k}$ have been deployed previously in nonparametric problems [28, 16, 31, 18, 42].

Given q, \mathcal{S} and K , we assign a uniform prior over *valid* tree topologies $\mathcal{T} = \{\Omega_k\}_{k=1}^K \in \mathcal{V}_{\mathcal{S}}^K$, i.e.

$$(T4) \quad \pi(\mathcal{T} | \mathcal{S}, K) = \frac{1}{\Delta(\mathcal{V}_{\mathcal{S}}^K)} \mathbb{I}(\mathcal{T} \in \mathcal{V}_{\mathcal{S}}^K).$$

Similar constraints on trees where each terminal node is assigned a minimal number of data points have been implemented in stochastic search algorithms [18]. At the very least, we can choose $\bar{C} = 1$ in (3.2), merely

requiring that the cells be non-empty. Finally, given the partition of size K , we assign an iid Gaussian prior on the step heights (similarly as in [15])

$$(T5) \quad \pi(\boldsymbol{\beta} | K) = \prod_{k=1}^K \phi(\beta_k; 1).$$

REMARK 4.1. *The name spike-and-tree prior deserves a bit of explanation. It follows from the fact that (T1) and (T2) will be satisfied if each covariate has a prior probability $\theta \sim \mathcal{B}(1, p^u)$ of contributing to the mean regression surface for some $u > 1$ [13]. Endowing each covariate x_j with a Bernoulli indicator γ_j , where $\Pi(\gamma_j = 1 | \theta) = \theta$, and building a tree on $\mathcal{S} = \{j : \gamma_j = 1\}$, one obtains a mixture prior on $f(\mathbf{x})$ that pertains to spike-and-slab variable selection. Here, the slab is a tree prior built on active covariates rather than an independent product prior on active regression coefficients. This hierarchical construction has distinct advantages for variable selection. In linear regression, it is customary to select variables by thresholding marginal posterior inclusion probabilities $\Pi(\gamma_j = 1 | \mathbf{Y}^{(n)})$. These will be available also under our spike-and-tree construction. Inspecting the posterior probabilities $\Pi(\gamma_j = 1 | \mathbf{Y}^{(n)})$ obtained with our hierarchical tree prior will be a new avenue for conducting variable selection in Bayesian CART and BART, an alternative to [7]. Thus, our prior has important practical implications for performing principled model-free variable selection.*

4.2. *Posterior Concentration for Bayesian CART.* The difficulty in properly analyzing Bayesian CART stems from the combinatorial richness of the prior that makes it less tractable analytically. By building on our developments from previous sections, we are now fully equipped to present the first theoretical result concerning this method.

The following theorem solidifies the optimality properties of Bayesian CART by showing that under the hierarchical prior (T1)-(T5), the posterior adapts to both smoothness and sparsity, concentrating at the (near) minimax rate that depends only on the number of strong covariates regardless of how many noise variables are present. The near-minimaxity refers to an additional log factor. The result holds for sparse (high-dimensional) regimes, where p can be potentially much larger than n and where $q_0 \lesssim \log^{1/2} n$. We will denote by

$$\mathcal{F}_{\mathcal{T}} = \bigcup_{q=0}^{\infty} \bigcup_{K=1}^{\infty} \bigcup_{\mathcal{S}: |\mathcal{S}|=q} \mathcal{F}(\mathcal{V}_{\mathcal{S}}^K)$$

the collection of all tree-structured step functions (with various tree sizes and split subsets) that can be obtained by partitioning \mathcal{X} .

THEOREM 4.1. *Assume $f_0 \in \mathcal{H}_p^\alpha \cap \mathcal{C}(\mathcal{S}_0)$ with $0 < \alpha \leq 1$ and $0 < q_0 = |\mathcal{S}_0|$ such that $q_0 \|f_0\|_{\mathcal{H}^\alpha} \lesssim \log^{1/2} n$ and $\|f_0\|_\infty \lesssim \log^{1/2} n$. Moreover, we assume that $\log p \lesssim n^{q_0/(2\alpha+q_0)}$ and that \mathcal{X} is (M, \mathcal{S}_0) -regular. We endow \mathcal{F}_T with priors (T1)-(T5). Then with $\varepsilon_n = n^{-\alpha/(2\alpha+q_0)} \log^{1/2} n$ we have*

$$\Pi \left(f \in \mathcal{F}_T : \|f_0 - f\|_n > M_n \varepsilon_n \mid \mathbf{Y}^{(n)} \right) \rightarrow 0,$$

for any $M_n \rightarrow \infty$ in P_0^n -probability, as $n, p \rightarrow \infty$.

PROOF. Section 8. □

REMARK 4.2. *It is useful to note that Theorem 4.1 holds also when $p < n$ and when q_0 is fixed as $n \rightarrow \infty$. When q_0 is fixed, however, the assumptions $\|f_0\|_\infty \lesssim \log^{1/2} n$ and $q_0 \|f_0\|_{\mathcal{H}^\alpha} \lesssim \log^{1/2} n$ can be omitted. The first assumption is needed here to make sure that the step sizes of an approximating k - d tree are well behaved when $q_0 \rightarrow \infty$. The result holds for a bit slower rate $\varepsilon_n = n^{-\alpha/(2\alpha+q_0)} \log^\beta n$ with $\beta > 1/2$ under slightly relaxed assumptions $q_0 \|f_0\|_{\mathcal{H}^\alpha} \lesssim \log^\beta n$ and $\|f_0\|_\infty \lesssim \log^\beta n$.*

The assumption of a regular design is an inevitable consequence of treating \mathbf{x} 's as fixed. As noted by [6] in their study of random forests, point-wise consistency results have been complicated by the difficulty in controlling local (cell) diameters. The regularity assumption guarantees this control and is apt to be satisfied for most realizations of \mathbf{x} from reasonable distributions on $[0, 1]^p$. The following Corollary certifies that Bayesian CART, under a suitable complexity prior on the number of terminal nodes, is reluctant to overfit. This is seen from the behavior of the posterior, which concentrates on values K that are only a constant multiple larger than the optimal oracle value $n^{q_0/(2\alpha+q_0)}$.

COROLLARY 4.1. *(Bayesian regression trees do not overfit.) Under the assumptions of Theorem 4.1 with $0 < \alpha \leq 1$ we have*

$$\Pi \left(K > C_k n^{q_0/(2\alpha+q_0)} \mid \mathbf{Y}^{(n)} \right) \rightarrow 0$$

in P_0^n -probability for a suitable constant $C_k > 0$.

PROOF. This statement follows from Lemma 1 of [24] and holds upon the satisfaction of the condition $\Pi(K > C_k n^{q_0/(2\alpha+q_0)}) = o(e^{-(d+2)n\varepsilon_n^2})$ for some suitably large $d > 2$. This is shown in Section 8.3. □

Corollary 4.1 also reveals a fundamental limitation of trees (step functions) in recovering smoother functions. To see this, consider $f_0(x) = x$ which possesses Hölder smoothness 1 and, by Corollary 4.1, will thus be approximated by trees with at most $n^{1/3}$ leaves (up to multiplicative constants) with high posterior probability. However, f_0 is also in \mathcal{C}^2 and the approximation error by a regular histogram with $n^{1/3}$ leaves will be at least of the order $n^{-1/3}$ which is too large to achieve the minimax rate of $n^{-5/2}$ over \mathcal{C}^2 .

REMARK 4.3. (*Bayesian CART à la Chipman, George and McCulloch [15]*) A closer look at the proof reveals that Theorem 4.1 also holds for priors $\pi(k|q) \lesssim \pi(k|q) \lesssim e^{-ck \log(n)^\eta}$ for some $0 < \eta < 1$. At the lower end are the complexity priors deployed in similar contexts [37, 42]. The Bayesian CART prior of [15], which is deployed in BART (and discussed in Section 4.1) with $\alpha = 0$ is at the upper end. To see this, note that $\alpha = 0$ corresponds to a homogeneous Galton-Watson (GW) process, where the number of terminal nodes K satisfies $\Pi(K > k) \leq e^{-(k-0.5) \log[(1-\gamma)/\gamma]}$. With $\gamma = c/n$ for some $c < n$ we have $\Pi(K > k) \propto e^{-ck \log n}$. Going further, the tail bound for the heterogeneous GW process (with $\alpha \neq 0$) satisfies $\Pi(K > k) \leq e^{-C_K k \log k}$ for some $C_K > 0$ under a suitable modification of the split probability (i.e. $\pi(d) \propto a^d$ for some $1/n < a \leq 1/2$), as shown formally in Ročková and Saha [35]. They also show that the Bayesian CART posterior à la Chipman et al. [15] concentrates at the optimal rate when S_0 is known. Endowed with the spike-and-slab wrapper, Theorem 4.1 can be thus extended to the actual Bayesian CART prior deployed in BART.

5. Tree Ensembles. Combining multiple trees through additive aggregation has proved to be remarkably effective for enhancing prediction [8, 14]. This section offers new theoretical insights into the mechanics behind the Bayesian variants of such tree ensemble methods. Our approach rests on a detailed analysis of the *collective behavior* of partitioning cells generated by individual trees. We will see that the overall performance is affected not only by the quality of single trees but also how well they can collaborate [9].

5.1. *Bayesian Additive Regression Trees.* Additive regression trees grow an ensemble predictor by binding together T tree-shaped regressors. For subsets $\mathcal{S} = \{\mathcal{S}^1, \dots, \mathcal{S}^T\}$ and tree sizes $\mathbf{K} = (K^1, \dots, K^T)'$, we define a sum-of-trees model (forest) as

$$(5.1) \quad f_{\mathcal{E}, \mathcal{B}}(\mathbf{x}) = \sum_{t=1}^T f_{\mathcal{T}^t, \beta^t}(\mathbf{x}) = \sum_{t=1}^T \sum_{k=1}^{K^t} \beta_k^t \mathbb{I}(\mathbf{x} \in \Omega_k^t),$$

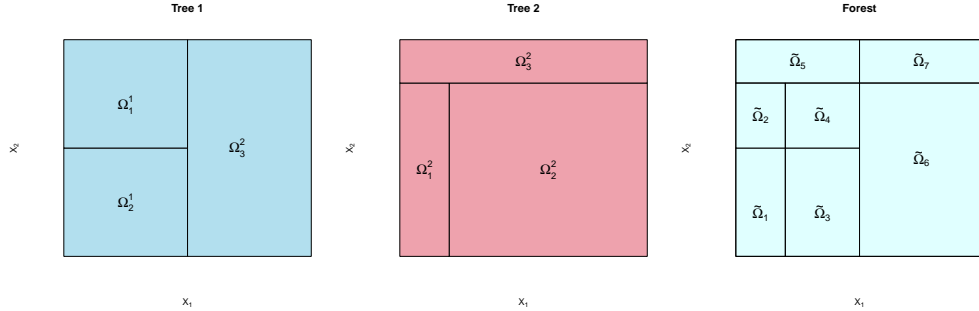


FIG 1. *Ensemble of tree partitions and their global partition.*

where $f_{\mathcal{T}^t, \beta^t} \in \mathcal{F}(\mathcal{V}_{\mathcal{S}^t}^{K^t})$ is a tree learner associated with step sizes $\beta^t \in \mathbb{R}^{K^t}$. Each learner is allowed to use different splitting variables \mathcal{S}^t and different number of bottom nodes K^t . With $\mathcal{E} = \{\mathcal{T}^1, \dots, \mathcal{T}^T\}$ we denote an ensemble of tree partitions and with $\mathcal{B} = (\beta^{1'}, \dots, \beta^{T'})' \in \mathbb{R}^{TK}$ the terminal node parameters, where $\bar{K} = \frac{1}{T} \sum_{t=1}^T K^t$.

Sum-of-trees models offer an improved representation flexibility by chopping up the predictor space into more refined segmentations. These segmentations are obtained by superimposing multiple tree partitions, yielding what we define below as a *global* partition.

DEFINITION 5.1. *For a partition ensemble \mathcal{E} , we define a global partition*

$$\tilde{\mathcal{T}}(\mathcal{E}) = \{\tilde{\Omega}_k\}_{k=1}^{K(\mathcal{E})}$$

as the partition obtained by merging all cuts in $\mathcal{T}^1, \dots, \mathcal{T}^T$. We refer to $\tilde{\Omega}_k$'s as global cells in the ensemble.

The concept of the global partition can be better understood from Figure 1, where splits from $T = 2$ trees (each having 3 leaves) are merged to obtain a global partition with $K(\mathcal{E}) = 7$ global cells. Generally, the global partition \mathcal{E} itself is *not* necessarily a tree and can have as many as $K(\mathcal{E}) \leq \prod_{t=1}^T K^t$ cells. This upper bound can be attained if each trees splits $K^t - 1$ times on a single variable, where each tree uses a different one. Without loss of generality, the global partition will be assumed to have non-empty cells. This requirement can be met by merging vacuous cells with their nonempty neighbors.

Bayesian additive regression trees were conceived as a collection of weak learners that capture *different* aspects of the predictor space [14]. To char-

acterize the amount of diversity/correlation between trees in the ensemble, we introduce the so-called stretching matrix.

DEFINITION 5.2. *For a partition ensemble $\mathcal{E} = \{\mathcal{T}^1, \dots, \mathcal{T}^T\}$, we define the stretching matrix $\mathbf{A}(\mathcal{E}) = (a_{ij})_{i,j}$ as follows: for each $1 \leq i \leq K(\mathcal{E})$ and $1 \leq j \leq T\bar{K}$ we have*

$$(5.2) \quad a_{ij} = \mathbb{I}\{\tilde{\Omega}_i \cap \Omega_m^t \neq \emptyset\},$$

where $1 \leq t \leq T$ and $1 \leq m \leq K^t$ are such that $j = \sum_{l=1}^{t-1} K^l + m$ and where Ω_m^t is the m^{th} (local) cell in the t^{th} tree \mathcal{T}^t .

Each row of the stretching matrix $\mathbf{A}(\mathcal{E})$ corresponds to one global cell and each column to one local cell. The row entries sum to T , indicating which local cells overlap with that global cell (as shown in Figure 2 for partitions from Figure 1). To further characterize the pattern of overlap between trees, we introduce the $(T\bar{K}) \times (T\bar{K})$ Gram matrix

$$(5.3) \quad \tilde{\mathbf{A}}(\mathcal{E}) \equiv \mathbf{A}(\mathcal{E})' \mathbf{A}(\mathcal{E}) = (\tilde{a}_{ij})_{i,j}.$$

The off-diagonal elements measure the “similarity” between local cells, say Ω_j^t and Ω_k^u , in terms of the number of global cells that they share. More formally, let $r[\tilde{\mathcal{T}}(\mathcal{E}), V] = |\{\tilde{\Omega}_k : \tilde{\Omega}_k \cap V \neq \emptyset\}|$ be the restricted cell count [34], measuring the number of global cells that intersect with a compact set V . For $i = \sum_{s=1}^{t-1} K^s + j$ and $l = \sum_{s=1}^{u-1} K^s + k$ we can write $\tilde{a}_{il} = r[\tilde{\mathcal{T}}(\mathcal{E}), \Omega_j^t \cap \Omega_k^u]$. Small off-diagonal entries indicate less overlap, where the individual trees capture more diverse aspects of the predictor space. The diagonal elements, on the other hand, quantify the “persistence” of each local cell, say Ω_k^u , counting the number of global cells it stretches over. More formally, for $i = \sum_{s=1}^{u-1} K^s + k$ we have $\tilde{a}_{ii} = r[\tilde{\mathcal{T}}(\mathcal{E}), \Omega_k^u]$. The amount of diversity (or tree dis-similarity) in the ensemble can be quantified with eigenvalues of $\tilde{\mathbf{A}}(\mathcal{E})$. We denote by $\lambda_{\min}(\mathcal{E})$ (resp. $\lambda_{\max}(\mathcal{E})$) the minimal (resp. maximal) singular values of $\mathbf{A}(\mathcal{E})$ (i.e. square roots of extremal nonzero eigenvalues of $\tilde{\mathbf{A}}(\mathcal{E})$). If some trees in the ensemble are redundant, the conditioning number $\kappa(\mathcal{E}) \equiv \lambda_{\max}(\mathcal{E})/\lambda_{\min}(\mathcal{E})$ will be large. The idea of diversifying trees was originally introduced by Breiman [8] via subsampling. One could, in principle, impose a restriction on $\lambda_{\min}(\mathcal{E})$ in the prior to encourage the trees to collaborate and diversify. However, this is *not* required for our theoretical study. We will focus on the so-called valid ensembles which consist of valid trees.

$$\mathbf{A}(\mathcal{E}) = \begin{matrix} & \Omega_1^1 & \Omega_2^1 & \Omega_3^1 & \Omega_1^2 & \Omega_2^2 & \Omega_3^2 \\ \tilde{\Omega}_1 & \begin{pmatrix} 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix} \end{matrix}, \quad \mathbf{A}(\mathcal{E})' \mathbf{A}(\mathcal{E}) = \left(\begin{array}{ccc|ccc} 3 & 0 & 0 & 1 & 1 & 1 \\ 0 & 2 & 0 & 1 & 1 & 0 \\ 0 & 0 & 2 & 0 & 1 & 1 \\ \hline 1 & 1 & 0 & 2 & 0 & 0 \\ 1 & 1 & 1 & 0 & 3 & 0 \\ 1 & 0 & 1 & 0 & 0 & 2 \end{array} \right)$$

FIG 2. Stretching matrix $\mathbf{A}(\mathcal{E})$ and the “Gram matrix” $\mathbf{A}(\mathcal{E})' \mathbf{A}(\mathcal{E})$ from Figure 1.

DEFINITION 5.3. An ensemble $\mathcal{E} = \{\mathcal{T}^1, \dots, \mathcal{T}^T\}$ is valid if each \mathcal{T}^t is valid according to Definition 3.1. For tree sizes $\mathbf{K} = (K^1, \dots, K^T)'$ and subsets $\mathcal{S} = \{\mathcal{S}^1, \dots, \mathcal{S}^T\}$, we denote the set of all valid ensembles by $\mathcal{V}\mathcal{E}_{\mathcal{S}}^{\mathbf{K}}$.

The representation flexibility of additive trees also pertains to jump sizes. The global step size coefficients under additive trees are intertwined due to the tree overlap. This can be seen from the following, more compact, representation of (5.1):

$$f_{\mathcal{E}, \mathcal{B}}(\mathbf{x}) = \sum_{k=1}^{K(\mathcal{E})} \bar{\beta}_k \mathbb{I}(\mathbf{x} \in \tilde{\Omega}_k) \quad \text{with} \quad \bar{\beta}_k \equiv \sum_{t=1}^T \sum_{l=1}^{K^t} \beta_l^t \mathbb{I}(\Omega_l^t \cap \tilde{\Omega}_k \neq \emptyset),$$

where $\bar{\boldsymbol{\beta}} = (\bar{\beta}_1, \dots, \bar{\beta}_{K(\mathcal{E})})' \in \mathbb{R}^{K(\mathcal{E})}$ are aggregated step sizes. A closer look reveals the following key connection between $\bar{\boldsymbol{\beta}}$ and \mathcal{B}

$$(5.4) \quad \bar{\boldsymbol{\beta}} = \mathbf{A}(\mathcal{E}) \mathcal{B},$$

where $\mathbf{A}(\mathcal{E})$ is the stretching matrix defined in (5.2). This link unfolds the theoretical analysis of tree ensembles using tools that we have already developed for single trees. Note that the condition number $\kappa(\mathcal{E})$ determines how much the relative change in $\bar{\boldsymbol{\beta}}$ influences the relative change in \mathcal{B} .

The mapping (5.4) can be in principle many-to-one in the sense that many tree-structured step functions can sum towards the same target (5.1). Such over-parametrization occurs, for instance, when $K(\mathcal{E}) < T\bar{K}$ or, more generally, when $\mathbf{A}(\mathcal{E})$ has zero eigenvalues. This redundancy is not entirely unwanted and, in fact, it endows sum-of-trees models with a lot of modeling freedom.

We now formally define the space of approximating additive trees. For variable sets $\mathcal{S} = \{\mathcal{S}^1, \dots, \mathcal{S}^T\}$ and a vector of tree sizes $\mathbf{K} = (K^1, \dots, K^T)'$,

we denote by

$$(5.5) \quad \mathcal{F}(\mathcal{V}\mathcal{E}_{\mathcal{S}}^{\mathbf{K}}) = \left\{ f_{\mathcal{E},\mathcal{B}} : [0, 1]^p \rightarrow \mathbb{R} : f_{\mathcal{E},\mathcal{B}}(\mathbf{x}) = \sum_{t=1}^T f_{\mathcal{T}^t, \beta^t}(\mathbf{x}); \mathcal{E} \in \mathcal{V}\mathcal{E}_{\mathcal{S}}^{\mathbf{K}}, \beta^t \in \mathbb{R}^{K^t} \right\}$$

the set of all additive tree step functions supported on *valid* ensembles $\mathcal{V}\mathcal{E}_{\mathcal{S}}^{\mathbf{K}}$. The union of these over the number of trees T , all possible sets \mathcal{S} of sizes $\mathbf{q} = (q^1, \dots, q^T)'$ and tree sizes \mathbf{K} gives rise to

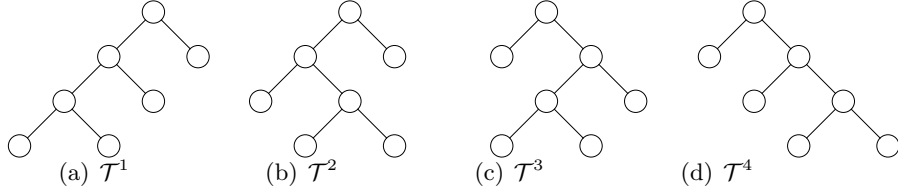
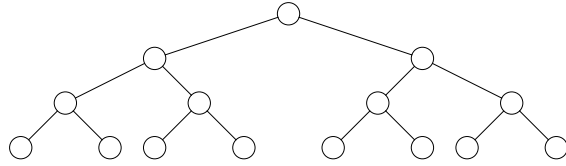
$$(5.6) \quad \mathcal{F}_{\mathcal{E}} = \bigcup_{T=1}^{\infty} \bigcup_{\mathbf{q}} \bigcup_{\mathcal{S}: |\mathcal{S}^t|=q^t} \bigcup_{\mathbf{K}} \mathcal{F}(\mathcal{V}\mathcal{E}_{\mathcal{S}}^{\mathbf{K}}),$$

our approximating space of additive regression tree functions.

5.2. Additive Regression Trees are Adaptive. This section provides an interesting initial perspective on the behavior of Bayesian additive regression trees in Regime 1. We will continue with the more general Regime 2 in the next section. We will focus on a variant of the popular Bayesian Additive Regression Trees (BART) model [14], modified in three ways. First, the tree prior will be according to [18] rather than [15]. The second modification is that the trees are built on the *same set of variables*, endowed with a subset selection prior construction. In the next section, we allow for the fully general case where each tree builds on a potentially different set of variables. Third, rather than fixing the number of trees, we endow T with a prior distribution.

We will see that having a good control of the regression function variation inside each *global* cell together with a good choice of the prior on the total number of leaves $\sum K^t$ will be sufficient to ensure optimal behavior. The approximation ability of tree ensembles hinges on the diameter of the *global* partition. Each tree partition does not need to have a small diameter (i.e. can be a weak learner), as long as the global one does. An important building block in our proof will be the construction of a single tree ensemble that can approximate well. As will be shown in Lemma 10.1, we can construct such ensemble by first finding a single k - d tree from Lemma 3.2 (a strong learner) and then redistributing the cuts among small trees (weak learners) in a way that the global partition is exactly equal to the k - d tree. An example of this deconstruction is depicted in Figure 3, where a full symmetric tree from Figure 4, say $\widehat{\mathcal{T}}$, has been trimmed into many smaller imbalanced trees which add up towards $\widehat{\mathcal{T}}$. More details on this decomposition are in the proof of Lemma 10.1.

The following theorem is an ensemble variant of Theorem 4.1 which will serve as a useful stepping stone towards the full-fledged result presented in

FIG 3. Trees in the approximating ensemble $\hat{\mathcal{E}}$.FIG 4. Approximating k - d tree partition $\hat{\mathcal{T}} = \tilde{\mathcal{T}}(\hat{\mathcal{E}})$

the next section. Instead of approximating $f_0 \in \mathcal{H}_p^\alpha \cap \mathcal{C}(\mathcal{S}_0)$ with one large tree in Regime 1, we build a forest made up of many smaller trees (weak learners). The first two layers of the prior (T1) and (T2) are the same. Next, we assign a prior on the number of trees

$$(T3^*) \quad \pi(T) \propto e^{-C_T T} \quad \text{for } T \in \mathbb{N} \setminus \{0\} \quad \text{with } C_T > \log 2,$$

which is sufficiently diffuse so as to promote ensembles with many trees. Next, conditionally on T , we assign a prior on $\mathbf{K} = (K^1, \dots, K^T)'$ as an independent product of Poisson distributions (T3). An important distinction between the Poisson prior deployed for single trees in Section 4 and the one deployed here for ensembles is that the hyper-parameter depends on T . Namely, our prior on leaves satisfies

$$(T4^*) \quad \pi(\mathbf{K} | T) = \prod_{t=1}^T \frac{(\lambda/T)^{K^t}}{(e^{\lambda/T} - 1)K^t!} \quad \text{for } K^t \in \mathbb{N} \setminus \{0\}.$$

As a prior on \mathcal{E} , given $(\mathcal{S}, \mathbf{K})$, we use the uniform prior over valid ensembles

$$(T5^*) \quad \pi(\mathcal{E} | \mathcal{S}, \mathbf{K}) = \frac{1}{\Delta(\mathcal{V}\mathcal{E}_S^{\mathbf{K}})} \mathbb{I}(\mathcal{E} \in \mathcal{V}\mathcal{E}_S^{\mathbf{K}}),$$

where $\Delta(\mathcal{V}\mathcal{E}_S^{\mathbf{K}})$ is the overall number of ensembles consisting of T valid trees that can be obtained by splitting on data points \mathcal{X} . Recall that throughout this section, all trees are constrained to use split variables in the same set

\mathcal{S} . To mark this difference, we have denoted the partition ensembles with $\mathcal{V}\mathcal{E}_{\mathcal{S}}^{\mathbf{K}}$ instead of $\mathcal{V}\mathcal{E}_{\mathcal{S}}^{\mathbf{K}}$.

Finally, given (T, \mathbf{K}) , we assign an iid Gaussian prior on the step heights $\mathcal{B} \in \mathbb{R}^{\sum_{t=1}^T K^t}$ with a variance $1/T$ (as suggested by [14])

$$(T6^*) \quad \pi(\mathcal{B} | T, \mathbf{K}) = \prod_{t=1}^T \prod_{k=1}^{K^t} \phi(\beta_k^t; 1/T).$$

The following theorem shows that additive regression trees, in combination with a subset selection prior, can nicely adapt to the ambient dimensionality and smoothness, *also* achieving the optimal concentration rate in Regime 1.

THEOREM 5.1. *Assume $f_0 \in \mathcal{H}_p^\alpha \cap \mathcal{C}(\mathcal{S}_0)$ with $0 < \alpha \leq 1$ and $0 < q_0 = |\mathcal{S}_0|$ such that $q_0 \|f_0\|_{\mathcal{H}^\alpha} \lesssim \log n$ and $\|f_0\|_\infty \lesssim \log n$. Moreover, we assume $\log p \lesssim n^{q_0/(2\alpha+q_0)}$ and that \mathcal{X} is (M, \mathcal{S}_0) -regular. We endow $\mathcal{F}_{\mathcal{E}}$ with priors (T1),(T2),(T3*)-(T6*). With $\varepsilon_n = n^{-\alpha/(2\alpha+q_0)} \log n$ we have*

$$\Pi \left(f \in \mathcal{F}_{\mathcal{E}} : \|f_0 - f\|_n > M_n \varepsilon_n \mid \mathbf{Y}^{(n)} \right) \rightarrow 0,$$

for any $M_n \rightarrow \infty$ in P_0^n -probability, as $n, p \rightarrow \infty$.

PROOF. Supplemental Material (Section 2). □

Similarly as for single trees, we obtain the following corollary which states that the posterior concentrates on ensembles whose *overall number of leaves* is not much larger than the optimal value $n^{q_0/(2\alpha+q_0)} \log n$.

COROLLARY 5.1. *Under the assumptions of Theorem 5.1 with $0 < \alpha \leq 1$ we have*

$$\Pi \left((T, \mathbf{K}) : \sum_{t=1}^T K^t > C_k n^{q_0/(2\alpha+q_0)} \log n \mid \mathbf{Y}^{(n)} \right) \rightarrow 0$$

in P_0^n -probability, as $n, p \rightarrow \infty$, for a suitable constant $C_k > 0$.

PROOF. This follows from Lemma 1 of [24] and Section 10.3. □

Corollary 5.1 shows that the posterior distribution rewards either many weak learners or a few strong ones. The compromise between the two is regulated by the prior (T3*), where stronger shrinkage (i.e. larger C_T) will result in fewer trees. This corollary provides an important theoretical justification for why Bayesian additive regression trees have been so resilient to overfitting in practice.

REMARK 5.1. *The dependence on T in the Poisson prior (T_4^*) works nicely in tandem with the exponential prior (T_3^*) . Removing T from (T_4^*) would have to be balanced with a bit stronger prior $\pi(T)$. Theorem 5.1 also holds with T fixed (with slight modifications of the proof). The prior $\pi(T)$ will be instrumental in the additive case (next section). The dependence on T in (T_6^*) , though recommended in practice [14], is not needed in Theorem 5.1.*

6. Tree Ensembles in Additive Regression. In Section 4.2 we have shown that the posterior distribution under the Bayesian CART prior has optimal properties. However, it is now well known that the practical deployments of Bayesian CART suffer from poor MCMC mixing. Additive aggregations of single small trees [14] have proven to have far more superior mixing properties. One may wonder whether the benefits of additive trees are purely computational or whether there are some aspects that make them more attractive also theoretically. We will address this fundamental question.

For estimating a single smooth function, we were not able to tell apart single trees from tree ensemble in terms of their convergence rate (besides perhaps a small difference in the log factor). They are both optimal in Regime 1. Tree ensembles are inherently additive and, as such, are well-equipped for approximating additive f_0 (Regime 2). Throughout this section, we assume

$$(6.1) \quad f_0(\mathbf{x}) = \sum_{t=1}^{T_0} f_0^t(\mathbf{x}),$$

where $f_0^t \in \mathcal{H}_p^{\alpha^t} \cap \mathcal{C}(\mathcal{S}_0^t)$. Note that each component f_0^t depends only on a potentially very small subset \mathcal{S}_0^t of covariates, where $|\mathcal{S}_0^t| = q_0^t$. However, the additive structure allows f_0 to depend on a larger number of variables, say q_0 , where $\max_{1 \leq t \leq T_0} q_0^t \leq q_0 \leq \sum_{t=1}^{T_0} q_0^t$. The minimax rate r_n^2 for estimating f_0 in Regime 2 satisfies $C_1 \underline{\varepsilon}_n^2 \leq r_n^2 \leq C_2 \bar{\varepsilon}_n^2$ [46], where

$$\begin{aligned} \underline{\varepsilon}_n^2 &= \sum_{t=1}^{T_0} \left[\lambda^{t^2} (\sqrt{n} \lambda^t)^{-4\alpha^t / (2\alpha^t + q_0^t)} + \frac{q_0^t}{n} \log \left(\frac{p}{\sum_t q_0^t} \right) \right], \\ \bar{\varepsilon}_n^2 &= \sum_{t=1}^{T_0} \left[\lambda^{t^2} (\sqrt{n} \lambda^t)^{-4\alpha^t / (2\alpha^t + q_0^t)} + \frac{q_0^t}{n} \log \left(\frac{p}{\min_t q_0^t} \right) \right] \end{aligned}$$

and where λ^t is the Hölder norm of f_0^t . The sparsity constraint in Regime 2 is less strict than in Regime 1, where q_0 can be potentially larger than $\log n$ while still allowing for consistent estimation. For the isotropic case

($\alpha^t = \alpha$ and $q_0^t = \tilde{q}_0$), single trees can achieve the slower rate $n^{-2\alpha/(2\alpha+q_0)}$ where $\tilde{q}_0 \leq q_0 \leq T_0 \tilde{q}_0$. As will be shown below, tree ensembles can achieve a faster rate $\varepsilon_n^2 = \sum_{t=1}^{T_0} (\varepsilon_n^t)^2$ where $\underline{\varepsilon}_n^2 \lesssim \varepsilon_n^2 \lesssim \bar{\varepsilon}_n^2$ and where $(\varepsilon_n^t)^2 = \lambda^{t2} (\sqrt{n}\lambda^t)^{-4\alpha^t/(2\alpha^t+q_0^t)} + \frac{q_0^t}{n} \log\left(\frac{p}{q_0^t}\right)$.

We will approximate f_0 with tree ensembles $f_{\mathcal{E},\mathcal{B}} \in \mathcal{F}_{\mathcal{E}}$. The ensembles here differ from the ones considered in Section 5.2. The crucial difference is that now we allow each of the trees \mathcal{T}^t to depend on a *different set of variables* \mathcal{S}^t . Now we have a vector of subset sizes $\mathbf{q} = (q^1, \dots, q^T)'$ and a set of subsets $\mathcal{S} = \{\mathcal{S}^1, \dots, \mathcal{S}^T\}$, one for each tree. We consider the following independent product variant of the complexity prior (T1)

$$(T1^*) \quad \pi(\mathbf{q} | T) \propto \prod_{t=1}^T c^{-q^t} p^{-aq^t}, \quad q^t = 0, 1, \dots, p, \quad \text{for } a > 2$$

and a product prior variant of (T2), given (T, \mathbf{q}) ,

$$(T2^*) \quad \pi(\mathcal{S} | T, \mathbf{q}) \propto \prod_{t=1}^T 1/\binom{p}{q^t}.$$

The prior on the number of trees, the number of leaves, ensembles and step sizes is the same as in (T3*), (T4*), (T5*).

We are now ready to present our final result showing that the posterior concentration for Bayesian additive regression trees is near-minimax rate optimal when f_0 has an additive structure.

THEOREM 6.1. *Assume that f_0 is as in (6.1) where $f_0^t \in \mathcal{H}_p^{\alpha^t} \cap \mathcal{C}(\mathcal{S}_0^t)$ with $0 < \alpha^t \leq 1$ and $q_0^t = |\mathcal{S}_0^t|$ such that $0 < q_0^t \|f_0^t\|_{\mathcal{H}_{\alpha^t}} \lesssim \log^{1/2} n$ and $\|f_0^t\|_{\infty} \lesssim \log^{1/2} n$. Moreover, assume $\log p \lesssim \min_{1 \leq t \leq T_0} n^{q_0^t/(2\alpha^t+q_0^t)}$ and that \mathcal{X} is (M, \mathcal{S}_0^t) -regular for $1 \leq t \leq T_0$. We endow $\mathcal{F}_{\mathcal{E}}$ with priors (T1*)-(T6*). With $\varepsilon_n^2 = \sum_{t=1}^{T_0} n^{-2\alpha^t/(2\alpha^t+q_0^t)} \log n$ and $T_0 \lesssim n$, we have*

$$\Pi\left(f \in \mathcal{F}_{\mathcal{E}} : \|f_0 - f\|_n > M_n \varepsilon_n \mid \mathbf{Y}^{(n)}\right) \rightarrow 0,$$

for any $M_n \rightarrow \infty$ in P_0^n -probability, as $n, p \rightarrow \infty$.

PROOF. Supplemental Material (Section 1). □

REMARK 6.1. *Under the assumptions $\log p \lesssim n^{q_0^t/(2\alpha^t+q_0^t)}$ and $q_0^t \|f_0^t\|_{\mathcal{H}_{\alpha^t}} \lesssim \log^{1/2} n$, the second term $\frac{q_0^t}{n} \log\left(\frac{p}{q_0^t}\right)$ in $(\varepsilon_n^t)^2$ (defined earlier) is dominated by the first term $n^{-2\alpha^t/(2\alpha^t+q_0^t)} \log n$. This is why the second term does not appear in the rate ε_n^2 in Theorem 6.1.*

Failing to recognize the additive structure in f_0 , single regression trees achieve the slower rate $n^{-\alpha/(2\alpha+q_0)} \log^{1/2} n$, according to Theorem 4.1. Theorem 6.1 thus provides an additional theoretical justification for Bayesian additive tree models suggesting their performance superiority over single trees when f_0 is additive.

6.1. Implementation Considerations. Our priors differ from the widely used BART implementations in three ways: (1) we focus on the uniform prior of Denison et al. [18], (2) we assign a prior distribution on the number of trees and (c) we deploy the spike-and-slab wrapper. Implementations of our priors are feasible with some modifications of the existing software. For the Bayesian CART prior that we analyze, Denison et al. [18] propose a reversible jump MCMC implementation. While this algorithm is different from BART, the acceptance ratios in the Metropolis-Hastings step differ only very slightly. Liu, Ročková and Wang [32] extended their sampler to the spike-and-tree (spike-and-forest) versions in two ways. The first one is a Metropolis-Hastings strategy that consists of joint sampling from variable subsets as well as trees (forests). As a faster alternative, they proposed an approximate ABC sampling strategy based on data splitting (called *ABC Bayesian Forests*).

7. Discussion. In this work, we have laid down foundations for the theoretical study of Bayesian regression trees and their additive variants. We have shown an optimal behavior of Bayesian CART, the first theoretical result on this method. We have developed several useful tools for analyzing additive regression trees (variants of the BART method), showing their optimal performance in both additive and non-additive regression. The smoothness order of studied functions is restricted to values not exceeding one, a main limitation of our approach due to the fact that our approximations are piecewise constants [24, 37, 22]. While in the one-dimensional case, step functions are not appealing estimators of a regression function that is thought to be smooth, methods like CART and BART are attractive and feasible solutions in complex high-dimensional data. The limitation $\alpha \leq 1$ could be overcome by extending our approach to piecewise polynomials or kernels, an elaboration that we leave for future investigation. One such extension was recently proposed in a related paper by Linero and Yang [30]. These authors obtained concentration results for a kernel method that can be regarded as a smooth variant of BART. The results of [30] do not apply for single trees, only aggregates of kernels. In contrast, we study actual posteriors of single trees, as well as forests, and analyze sieves of step functions which are the essence of the actual BART method.

While our priors do not exactly match the BART prior, BART could be adapted to achieve the same optimality properties. The first modification is the splitting probability (as pointed out in Remark 4.3). The second modification is the spike-and-slab wrapper (as pointed out in Section 4.1) or a modification of the prior on the split variables (as in [30]). The prior distribution on the number of trees will only be beneficial in the additive model and is not needed when f_0 has one layer.

The assumption of a known $\sigma_0 = \text{Var}(\varepsilon_i) = 1$ can be relaxed. It has been noted in the literature (e.g. [43, 17]) that the general result of Ghosal and van der Vaart [24] (which we build upon) can be extended to the unknown σ_0 case. Such an extension was formally proved in Jonge and van Zanten (2013), who assume that σ_0 belongs to a compact interval $[a, b]$ and the prior $\pi(\sigma)$ concentrates on $[a, b]$. We could obtain our results under this restriction as well by verifying suitably adapted conditions (2.1)-(2.3). In related work, Yoo and Ghosal [47] show optimal posterior concentration (in both L_2 and L_∞ sense) in non-parametric regression with unknown variance and B-spline tensor product priors. Next, they show that under an inverse-gamma prior, the posterior for σ contracts at σ_0 at the same rate. Moreover, for any prior on σ with positive and continuous density, the posterior of σ is consistent. These results are obtained under the assumption that f_0 is uniformly bounded. We anticipate that similar results will hold also for our priors when f_0 is uniformly bounded.

Acknowledgement We are grateful to the Associate Editor and two referees for helpful comments, and to Johannes Schmidt-Hieber for helpful discussions on the assumption $\alpha \leq 1$.

8. Proof of Theorem 4.1. Our approach consists of establishing conditions (2.1), (2.2) and (2.3) for $\varepsilon_n = n^{-\alpha/(2\alpha+q_0)} \log^\beta n$ for some $\beta \geq 1/2$. Note that Theorem 4.1 is stated for $\beta = 1/2$. We give a proof for the general case $\beta \geq 1/2$ under the assumptions $q_0 \|f_0\|_{\mathcal{H}^\alpha} \lesssim \log^\beta n$ and $\|f_0\|_\infty \lesssim \log^\beta n$ (Remark 4.2). The first step requires constructing the sieve $\mathcal{F}_{\mathcal{T}}^n \subset \mathcal{F}_{\mathcal{T}}$. For a given $n \in \mathbb{N}$ and suitably large integers $q_n < k_n$ (chosen later), we define the sieve $\mathcal{F}_{\mathcal{T}}^n$ as consisting of step functions over small trees that split only on a few variables, i.e.

$$\mathcal{F}_{\mathcal{T}}^n = \bigcup_{q=0}^{q_n} \bigcup_{K=1}^{k_n} \bigcup_{\mathcal{S}:|\mathcal{S}|=q} \mathcal{F}(\mathcal{V}_{\mathcal{S}}^K),$$

where $\mathcal{F}(\mathcal{V}_{\mathcal{S}}^K)$ was defined in (3.3). The optimal choice of k_n and q_n will follow from our considerations below.

8.1. *Condition (2.1).* We start with a useful lemma that characterizes a useful upper bound on the covering number of the smaller sets $\mathcal{F}(\mathcal{V}_S^K)$.

LEMMA 8.1. *Let $\mathcal{F}(\mathcal{V}_S^K)$ be the class of step functions (3.3). Then*

$$(8.1) \quad N\left(\frac{\varepsilon}{36}, \left\{f \in \mathcal{F}(\mathcal{V}_S^K) : \|f - f_0\|_n < \varepsilon\right\}, \|\cdot\|_n\right) \leq \Delta(\mathcal{V}_S^K) \left(\frac{108}{\bar{C}} \sqrt{n}\right)^K,$$

where $\Delta(\mathcal{V}_S^K)$ is the partitioning number of \mathcal{V}_S^K .

PROOF. Let $f_{\mathcal{T}, \beta_1}$ and $f_{\mathcal{T}, \beta_2} \in \mathcal{F}(\mathcal{V}_S^K)$ be two step functions supported on a single valid tree partition $\mathcal{T} \in \mathcal{V}_S^K$ with steps $\beta_1 \in \mathbb{R}^K$ and $\beta_2 \in \mathbb{R}^K$. Then, by the minimum leaf size condition, we have

$$\frac{\bar{C}^2}{n} \|\beta_1 - \beta_2\|_2^2 \leq \|f_{\mathcal{T}, \beta_1} - f_{\mathcal{T}, \beta_2}\|_n^2 = \sum_{k=1}^K \mu(\Omega_k) (\beta_{1k} - \beta_{2k})^2 \leq \|\beta_1 - \beta_2\|_2^2.$$

Denote by $f_{\mathcal{T}, \hat{\beta}}$ the $\|\cdot\|_n$ projection of f_0 onto $\mathcal{F}(\mathcal{T}) \subset \mathcal{F}(\mathcal{V}_S^K)$, the set of all step functions that live on a given partition \mathcal{T} . Then $\{\beta : \|f_{\mathcal{T}, \beta} - f_0\|_n \leq \varepsilon\} \subset \{\beta : \|\beta - \hat{\beta}\|_2 \leq \varepsilon \sqrt{n/\bar{C}}\}$ and $\{\beta : \|f_{\mathcal{T}, \beta} - f_0\|_n \leq \varepsilon/36\} \supset \{\beta : \|\beta - \hat{\beta}\|_2 \leq \varepsilon/36\}$. This relationship shows that the $\varepsilon/36$ covering number of an $\|\cdot\|_n$ ball $\{\beta : \|f_{\mathcal{T}, \beta} - f_0\|_n \leq \varepsilon\}$ can be bounded from above by the $\varepsilon/36$ covering number of an Euclidean ball of a radius $\varepsilon \sqrt{n/\bar{C}}$, which is bounded by $(\frac{108}{\bar{C}} \sqrt{n})^K$. We can repeat this argument by projecting f_0 onto any valid tree topology $\mathcal{T} \in \mathcal{V}_S^K$. The number of such valid trees is no larger than $\Delta(\mathcal{V}_S^K)$, which completes the proof. \square

The covering number for the entire sieve $\mathcal{F}_{\mathcal{T}}^n$ is then seen to satisfy

$$\begin{aligned} & N\left(\frac{\varepsilon}{36}, \left\{f \in \mathcal{F}_{\mathcal{T}}^n : \|f - f_0\|_n < \varepsilon\right\}, \|\cdot\|_n\right) \\ & < \sum_{q=0}^{q_n} \sum_{K=1}^{k_n} \sum_{\mathcal{S}: |\mathcal{S}|=q} N\left(\frac{\varepsilon}{36}, \left\{f \in \mathcal{F}(\mathcal{V}_S^K) : \|f - f_0\|_n < \varepsilon\right\}, \|\cdot\|_n\right). \end{aligned}$$

From Lemma 8.1 and Lemma 3.1, we obtain the following upper bound

$$(8.2) \quad \begin{aligned} & \sum_{q=0}^{q_n} \sum_{K=1}^{k_n} \binom{p}{q} (K-1)! q^K n^K \left(\frac{108}{\bar{C}} \sqrt{n}\right)^K \\ & < \sum_{K=0}^{k_n} \left(\frac{108}{\bar{C}} q_n n^{3/2} K\right)^K \sum_{q=0}^{q_n} \binom{p}{q}, \end{aligned}$$

where we used the fact $K! < K^K$. Next, using the regularized incomplete beta function representation of the Binomial *cdf*, we can write

$$(8.3) \quad \begin{aligned} \sum_{q=0}^{q_n} \binom{p}{q} &= 2^p (p - q_n) \binom{p}{q_n} \int_0^{1/2} x^{p-q_n-1} (1-x)^{q_n} dx \\ &\leq 2^{q_n+1} (p - q_n) \binom{p}{q_n} \frac{1 - 1/2^{q_n+1}}{q_n + 1} \leq 2^{q_n+1} \binom{p}{q_n} \frac{p - q_n}{q_n + 1}. \end{aligned}$$

Using $\binom{p}{q_n} \leq (ep/q_n)^{q_n}$, the quantity in (8.2) can be bounded by

$$\left(\frac{2ep}{q_n}\right)^{q_n+1} \sum_{K=0}^{k_n} \left(\frac{108}{C} q_n n^{3/2} k_n\right)^K < \left(\frac{2ep}{q_n}\right)^{q_n+1} \frac{(C q_n n^{3/2} k_n)^{k_n+1} - 1}{C q_n n^{3/2} k_n - 1},$$

where $C = (\frac{108}{C})$. Finally, the entropy condition requires that the log-covering number, now upper bounded by

$$(q_n + 1) \log(2ep/q_n) + (k_n + 1) \log(C q_n n^{3/2} k_n),$$

is no larger than (a constant multiple of) $n \varepsilon_n^2 = n^{q_0/(2\alpha+q_0)} \log^{2\beta} n$. Under our assumption $\log p \lesssim n^{q_0/(2\alpha+q_0)}$, this will be satisfied with $q_n = \lceil C_q \min\{p, n^{q_0/(2\alpha+q_0)} \log^{2\beta} n / \log(p \vee n)\} \rceil$ for some $C_q \geq 1$ and $k_n = \lfloor C_k n \varepsilon_n^2 / \log n \rfloor \asymp n^{q_0/(2\alpha+q_0)} \log^{2\beta-1} n$. The constants C_q and C_k will be determined later.

8.2. *Condition (2.2)*. We wish to show that the prior assigns enough mass around the truth in the sense that

$$(8.4) \quad \Pi(f \in \mathcal{F}_{\mathcal{T}} : \|f - f_0\|_n \leq \varepsilon_n) \geq e^{-d n \varepsilon_n^2}$$

for some large enough $d > 0$. We establish this condition by finding a lower bound on the prior probability in (8.4), using all step functions supported on a single good partition. We denote by \mathcal{S}_0 the true index set of active covariates, where $|\mathcal{S}_0| = q_0$. According to Lemma 3.2, there exists a tree-structured step function $f_{\tilde{\mathcal{T}}, \tilde{\beta}} \in \mathcal{F}(\mathcal{V}_{\mathcal{S}_0}^K)$ for some $K = 2^{q_0 s}$ and $s \in \mathbb{N} \setminus \{0\}$ such that

$$(8.5) \quad \|f_0 - f_{\tilde{\mathcal{T}}, \tilde{\beta}}\|_n \leq \|f_0\|_{\mathcal{H}^\alpha} C_1 q_0 / K^{\alpha/q_0}$$

for some $C_1 > 0$. The proof Lemma 3.2 is in the Supplemental Material (Section 3).

To continue with the proof of Theorem 4.1, We find the smallest $K = 2^s q_0$ such that the function $f_{\tilde{\mathcal{T}}, \tilde{\beta}} \in \mathcal{F}(\mathcal{V}_{\mathcal{S}_0}^K)$ in (8.5) safely approximates f_0 with

an error that is no larger than $\varepsilon_n/2$, a constant multiple of the target rate. Such a K will be denoted by a_n and is defined as the smallest K such that $\varepsilon_n \geq 2C_0q_0K^{-\alpha/q_0}$ for $C_0 = \|f_0\|_{\mathcal{H}^\alpha}C_1$. Then we have

$$(8.6) \quad \left(\frac{2C_0q_0}{\varepsilon_n}\right)^{\frac{q_0}{\alpha}} \leq a_n \leq \left(\frac{2C_0q_0}{\varepsilon_n}\right)^{\frac{q_0}{\alpha}} + 1.$$

We denote by $\widehat{\mathcal{T}} = \{\widehat{\Omega}_k\}_{k=1}^{a_n} \in \mathcal{T}_{\mathcal{S}_0}^{a_n}$ the k - d tree partition from Lemma 3.2 obtained with the choice $K = a_n$. The tree $\widehat{\mathcal{T}}$ is not only valid, but also balanced in the sense that $C_{min}^2/a_n < \mu(\widehat{\Omega}_k) \leq C_{max}^2/a_n$ for some $C_{min} < 1 < C_{max}$. The associated step sizes of the approximating tree will be denoted by $\widehat{\beta} \in \mathbb{R}^{a_n}$. Now, we lower-bound the prior probability of the neighborhood $\{f \in \mathcal{F}_{\mathcal{T}} : \|f - f_0\|_n^2 \leq \varepsilon_n^2\}$ by the prior probability of all regression trees supported on $\widehat{\mathcal{T}}$ inside this neighborhood (denoted by $\mathcal{F}(\widehat{\mathcal{T}})$):

$$(8.7) \quad \Pi(f \in \mathcal{F}_{\mathcal{T}} : \|f - f_0\|_n^2 \leq \varepsilon_n^2) \geq \pi(q_0)\pi(a_n) \frac{\Pi(f \in \mathcal{F}(\widehat{\mathcal{T}}) : \|f - f_0\|_n^2 \leq \varepsilon_n^2)}{\binom{p}{q_0}\Delta(\mathcal{V}_{\mathcal{S}_0}^{a_n})}.$$

For any $\beta \in \mathbb{R}^{a_n}$, we have

$$\|f_{\widehat{\mathcal{T}},\beta} - f_{\widehat{\mathcal{T}},\widehat{\beta}}\|_n^2 = \sum_{k=1}^{a_n} \mu(\widehat{\Omega}_k)(\beta_k - \widehat{\beta}_k)^2 \leq \|\beta - \widehat{\beta}\|_2^2$$

and by the reverse triangle inequality

$$\|\beta - \widehat{\beta}\|_2 \geq \|f_{\widehat{\mathcal{T}},\beta} - f_{\widehat{\mathcal{T}},\widehat{\beta}}\|_n \geq \left| \|f_{\widehat{\mathcal{T}},\beta} - f_0\|_n - \|f_{\widehat{\mathcal{T}},\widehat{\beta}} - f_0\|_n \right|.$$

Then the statement $\|\beta - \widehat{\beta}\|_2 < \varepsilon_n/2$ implies $\|f_0 - f_{\widehat{\mathcal{T}},\beta}\|_n < \|f_0 - f_{\widehat{\mathcal{T}},\widehat{\beta}}\|_n + \varepsilon/2 < \varepsilon$, where the last inequality follows from the definition of a_n . Thus, we have

$$\{\beta : \|\beta - \widehat{\beta}\|_2 \leq \varepsilon_n/2\} \subset \{f \in \mathcal{F}(\widehat{\mathcal{T}}) : \|f_0 - f\|_n < \varepsilon_n\}.$$

Now we can lower-bound (8.7) with

$$(8.8) \quad L(q_0, a_n, \mathcal{S}_0, \widehat{\beta}, \varepsilon_n) \equiv \pi(q_0)\pi(a_n) \frac{\Pi(\beta \in \mathbb{R}^{a_n} : \|\beta - \widehat{\beta}\|_2 \leq \varepsilon_n/2)}{\binom{p}{q_0}\Delta(\mathcal{V}_{\mathcal{S}_0}^{a_n})}.$$

In order to bound $\Pi(\beta \in \mathbb{R}^{a_n} : \|\beta - \widehat{\beta}\|_2 \leq \varepsilon_n/2)$, we follow the computations of [24], Theorem 12, to obtain

$$(8.9) \quad \Pi(\beta \in \mathbb{R}^{a_n} : \|\beta - \widehat{\beta}\|_2 \leq \varepsilon_n/2) \geq \frac{2^{-a_n} e^{-\|\widehat{\beta}\|_2^2 - \varepsilon_n^2/8}}{\Gamma(\frac{a_n}{2}) \frac{a_n}{2}} \left(\frac{\varepsilon_n^2}{4}\right)^{\frac{a_n}{2}}.$$

From the triangle inequality (and because $\widehat{\mathcal{T}}$ is balanced) we have

$$\|\widehat{\beta}\|_2 \leq \frac{\sqrt{a_n}}{C_{\min}} \|f_{\widehat{\mathcal{T}}, \widehat{\beta}}\|_n \leq \frac{\sqrt{a_n}}{C_{\min}} (\|f_{\widehat{\mathcal{T}}, \widehat{\beta}} - f_0\|_n + \|f_0\|_\infty) \leq \frac{\sqrt{a_n}}{C_{\min}} \left(\frac{\varepsilon_n}{2} + \|f_0\|_\infty \right).$$

Thereby we can write $\|\widehat{\beta}\|_2^2 \leq C_2 \|f_0\|_\infty^2 a_n$ for some constant $C_2 > 0$. Now we continue with a lower bound to $L(q_0, a_n, \mathcal{S}_0, \widehat{\beta}, \varepsilon_n)$ defined in (8.8). Using the following facts $\Gamma(x) \leq x^x$, $\binom{p}{q_0} \leq (ep/q_0)^{q_0}$ and $\Delta(\mathcal{V}_{\mathcal{S}_0}^{a_n}) < (a_n q_0 n)^{a_n}$ (Lemma 3.1) and using (8.9), we arrive at the following lower bound:

$$(8.10) \quad \frac{\pi(a_n) c^{-q_0} p^{-aq_0}}{\left(\frac{ep}{q_0}\right)^{q_0} (a_n q_0 n)^{a_n}} e^{-\varepsilon_n^2/8 - a_n(C_2 \|f_0\|_\infty^2 + \log 2)} \left(\frac{\varepsilon_n^2}{4}\right)^{\frac{a_n}{2}} \left(\frac{2}{a_n}\right)^{a_n/2+1}.$$

Condition 2.2 will be satisfied if this quantity is at least as large as $e^{-dn\varepsilon_n^2}$ for some large $d > 0$. We denote $c(p, q_0, a_n) = c^{-q_0} p^{-aq_0} (q_0/ep)^{q_0}$. Then we can rewrite (8.10) as

$$\pi(a_n) c(p, q_0, a_n) \left(\sqrt{2} e^{C_2 \|f_0\|_\infty^2 + \log 2} a_n^{3/2} n \right)^{-a_n} \frac{2}{a_n} \left(\frac{\varepsilon_n}{q_0} \right)^{a_n} e^{-\varepsilon_n^2/8}.$$

Taking minus the log of this quantity, Condition (2.2) will be met when (8.11)

$$-\log c(p, q_0, a_n) - \log \pi(a_n) + a_n \log \left(\sqrt{2} e^{C_2 \|f_0\|_\infty^2 + \log 2} a_n^{3/2} n \right) + a_n \log \left(\frac{q_0}{\varepsilon_n} \right)$$

is smaller than a constant multiple of $n\varepsilon_n^2$. Above, we omitted the small terms $\varepsilon_n^2/8$ (since $\varepsilon_n \rightarrow 0$) and $\log(a_n/2)$. First, we note that

$$-\log c(p, q_0, a_n) \leq q_0 \log(cp^{a+1}e/q_0).$$

With $q_0 \lesssim \log^\beta n$ and $\log p \lesssim n^{q_0/(2\alpha+q_0)}$, we obtain $-\log c(p, q_0, a_n) \lesssim n\varepsilon_n^2$. Next, focusing on the last term in (8.11), we obtain (from the left inequality in (8.6)) the following bound

$$q_0/\varepsilon_n \leq a_n^{\alpha/q_0} / (2C_0)$$

and hence $q_0/\varepsilon_n \lesssim a_n$ for $\alpha \leq 1$ and $q_0 \geq 1$. Moreover, from the right inequality in (8.6) we obtain for $\varepsilon_n = n^{-\alpha/(2\alpha+q_0)} \log^\beta n$ and $2C_0 q_0 \lesssim \log^\beta n$

$$(8.12) \quad a_n \lesssim n^{q_0/(2\alpha+q_0)}.$$

Under our assumption $\|f_0\|_\infty \lesssim \log^\beta n$, (8.12) immediately yields $a_n \|f_0\|_\infty^2 \lesssim n\varepsilon_n^2$. All of these considerations, combined with the fact $-\log \pi(a_n) \lesssim a_n \log a_n$, yield the following leading term behind the last three summands in (8.11): $a_n \log(a_n^{3/2} n)$. Using (8.12), we obtain $a_n \log(a_n^{3/2} n) \lesssim a_n \log n \lesssim n\varepsilon_n^2$ for $\beta \geq 1/2$. Altogether, there exists $d > 0$ such that (8.4) is satisfied.

8.3. *Condition (2.3)*. In order to establish Condition 2.3, we begin by noting $\Pi(\mathcal{F}_{\mathcal{T}} \setminus \mathcal{F}_{\mathcal{T}}^n) < \Pi(q > q_n) + \Pi(K > k_n)$. Thus, the condition will be met when both $\Pi(q > q_n) = o(e^{-(d+2)n\varepsilon_n^2})$ and $\Pi(K > k_n) = o(e^{-(d+2)n\varepsilon_n^2})$, where d is the constant deployed in Section 8.2. First, we find that

$$\Pi(q > q_n) \lesssim \sum_{k=q_n+1}^p (cp^a)^{-k} = (cp^a)^{-(q_n+1)} \frac{1 - (cp^a)^{-(p-q_n)}}{1 - 1/(cp^a)} < (cp^a)^{-q_n}.$$

With our choice $q_n = \lceil C_q \min\{p, n^{q_0/(2\alpha+q_0)} \log^{2\beta} n / \log(p \vee n)\} \rceil$, it turns out that

$$(8.13) \quad \Pi(q > q_n) e^{(d+2)n\varepsilon_n^2} < e^{-q_n[\log c + a \log p] + (d+2)n\varepsilon_n^2} \rightarrow 0$$

for a large enough constant $C_q > 0$. Indeed, for $p > n$ we have $q_n \log p \asymp n\varepsilon_n^2$. For $p < n\varepsilon_n^2 / \log n$ we have $q_n = p$ and $\Pi(q > q_n) = 0$. Finally, for $n\varepsilon_n^2 / \log n \leq p \leq n$, we have $q_n \log n \asymp n\varepsilon_n^2$ and $\log p \geq \frac{q_0}{2\alpha+q_0} \log n + (2\beta - 1) \log \log n$. For $c > 1$ and $\beta \geq 1/2$, we can write $q_n[\log c + a \log p] \leq C_q \frac{aq_0}{2\alpha+q_0} n^{q_0/(2\alpha+q_0)} \log^{2\beta} n$ and (8.13) holds for C_q large enough. Next, we apply the Chernoff bound for $\Pi(K > k_n)$. Namely, for any $t > 0$ we can write

$$(8.14) \quad \Pi(K > k_n) < e^{-t(k_n+1)} \mathbb{E} e^{tK} \propto e^{-t(k_n+1)} \sum_{k=1}^{\infty} \frac{(e^t \lambda)^k}{k!} \propto e^{-t(k_n+1)} (e^{e^t \lambda} - 1).$$

With our choice $k_n = \lfloor C_k n \varepsilon_n^2 / \log n \rfloor \asymp C_k n^{q_0/(2\alpha+q_0)} \log^{2\beta-1} n$ (Section 8.1) and with $t = \log k_n$ we obtain

$$\Pi(K > k_n) e^{(d+2)n\varepsilon_n^2} \lesssim e^{-(k_n+1) \log k_n + \lambda k_n + (d+2)n\varepsilon_n^2} \rightarrow 0$$

for a large enough constant C_k .

9. Proof of Theorem 6.1. We aim to establish conditions (2.1), (2.2) and (2.3) for $\varepsilon_n^2 = \sum_{t=1}^{T_0} (\varepsilon_n^t)^2$, where $\varepsilon_n^t = n^{-\alpha^t/(2\alpha^t+q_0^t)} \log^{\beta^t} n$. Our sieve consists of valid forests with either (a) many trees that are small (weak learners), or (b) a few large trees (strong learners). We impose a joint requirement on $\sum_{t=1}^T K^t$ so that the overall number of leaves in the ensemble is small. At the same time, we require that $\sum_{t=1}^T q^t$ (the upper bound on the number of active variables in the ensemble) is small as well. The sieve is constructed as follows:

$$(9.1) \quad \mathcal{F}_{\mathcal{E}}^n = \bigcup_{T=1}^{\infty} \bigcup_{\mathbf{q}: \sum_{t=1}^T q^t \leq s_n} \bigcup_{\mathcal{S}: |\mathcal{S}^t|=q^t} \bigcup_{\mathbf{K}: \sum_{t=1}^T K^t \leq z_n} \mathcal{F}(\mathcal{V} \mathcal{E}_{\mathcal{S}}^{\mathbf{K}})$$

for some integer values s_n and z_n . Throughout this section we denote $\bar{K} = \frac{1}{T} \sum_{t=1}^T K^t$.

9.1. *Condition 2.1.* We first obtain the following upper bound on the log-covering number

$$(9.2) \quad \log N \left(\frac{\varepsilon}{36}, \{f_{\mathcal{E}, \mathcal{B}} \in \mathcal{F}(\mathcal{E}) : \|f_{\mathcal{E}, \mathcal{B}} - f_0\|_n < \varepsilon\}, \|\cdot\|_n \right) \lesssim (T \times \bar{K}) \log(108\sqrt{n})$$

where $\mathcal{F}(\mathcal{E}) = \left\{ f_{\mathcal{E}, \mathcal{B}} : [0, 1]^p \rightarrow \mathbb{R} : f_{\mathcal{E}, \mathcal{B}}(\mathbf{x}) = \sum_{t=1}^T f_{\mathcal{T}^t, \beta^t}(\mathbf{x}); \beta^t \in \mathbb{R}^{K^t} \right\}$ is the set of all additive step functions supported on a single partition ensemble $\mathcal{E} \in \mathcal{V}\mathcal{E}_{\mathcal{S}}^K$. We denote by $\tilde{\mathcal{T}}(\mathcal{E}) = \{\tilde{\Omega}_k\}_{k=1}^{K(\mathcal{E})}$ the global partition associated with \mathcal{E} , consisting of $K(\mathcal{E})$ global cells. For $\mathcal{B}_1, \mathcal{B}_2 \in \mathbb{R}^{T \times \bar{K}}$, we denote by $f_{\mathcal{E}, \mathcal{B}_1}, f_{\mathcal{E}, \mathcal{B}_2} \in \mathcal{F}(\mathcal{E})$ two additive regression trees that sit on the same partition ensemble \mathcal{E} . Let $\bar{\beta}_1 = \mathbf{A}(\mathcal{E})\mathcal{B}_1$ and $\bar{\beta}_2 = \mathbf{A}(\mathcal{E})\mathcal{B}_2$ be the aggregated step sizes, as defined in (5.4), where $\mathbf{A}(\mathcal{E})$ is the stretching matrix. Then we can write

$$\frac{1}{n} \|\bar{\beta}_1 - \bar{\beta}_2\|_2^2 \leq \|f_{\mathcal{E}, \mathcal{B}_1} - f_{\mathcal{E}, \mathcal{B}_2}\|_n^2 = \sum_{k=1}^{K(\mathcal{E})} \mu(\tilde{\Omega}_k) (\bar{\beta}_{1k} - \bar{\beta}_{2k})^2 \leq \|\bar{\beta}_1 - \bar{\beta}_2\|_2^2.$$

Deploying the singular value decomposition $\mathbf{A}(\mathcal{E}) = \mathbf{U}\mathbf{D}\mathbf{V}^T$ we write $\tilde{\mathcal{B}}_1 = \mathbf{V}^T \mathcal{B}_1 \in \mathbb{R}^{\tilde{K}}$ and $\tilde{\mathcal{B}}_2 = \mathbf{V}^T \mathcal{B}_2 \in \mathbb{R}^{\tilde{K}}$, where $\tilde{K} \leq \min\{K(\mathcal{E}), T\bar{K}\}$. Using the fact that \mathbf{U} is unitary, we have

$$\frac{1}{n} \|\mathbf{D}(\tilde{\mathcal{B}}_1 - \tilde{\mathcal{B}}_2)\|_2^2 \leq \|f_{\mathcal{E}, \mathcal{B}_1} - f_{\mathcal{E}, \mathcal{B}_2}\|_n^2 \leq \|\mathbf{D}(\tilde{\mathcal{B}}_1 - \tilde{\mathcal{B}}_2)\|_2^2.$$

We write $f_{\mathcal{E}, \mathcal{B}_2}$ to be the projection of f_0 onto $\mathcal{F}(\mathcal{E})$ and note that $\{\mathcal{B} \in \mathbb{R}^{T\bar{K}} : \|f_{\mathcal{E}, \mathcal{B}} - f_{\mathcal{E}, \mathcal{B}_2}\|_n \leq \varepsilon\} \subset \{\tilde{\mathcal{B}} : \|\mathbf{D}(\tilde{\mathcal{B}} - \tilde{\mathcal{B}}_2)\|_2 \leq \sqrt{n}\varepsilon\}$ and $\{\mathcal{B} \in \mathbb{R}^{T\bar{K}} : \|f_{\mathcal{E}, \mathcal{B}} - f_{\mathcal{E}, \mathcal{B}_2}\|_n \leq \varepsilon/36\} \supset \{\tilde{\mathcal{B}} : \|\mathbf{D}(\tilde{\mathcal{B}} - \tilde{\mathcal{B}}_2)\|_2 \leq \varepsilon/36\}$. The covering number of $\{f_{\mathcal{E}, \mathcal{B}} \in \mathcal{F}(\mathcal{E}) : \|f_{\mathcal{E}, \mathcal{B}} - f_0\|_n \leq \varepsilon\}$ can be thus bounded from above by the minimal number of smaller ellipsoids $\{\tilde{\mathcal{B}}_1 \in \mathbb{R}^{\tilde{K}} : \|\mathbf{D}(\tilde{\mathcal{B}}_1 - \tilde{\mathcal{B}}_2)\|_2 \leq \varepsilon/36\}$ needed to cover a larger ellipsoid $\{\tilde{\mathcal{B}}_1 \in \mathbb{R}^{\tilde{K}} : \|\mathbf{D}(\tilde{\mathcal{B}}_1 - \tilde{\mathcal{B}}_2)\|_2 \leq \sqrt{n}\varepsilon\}$. Since these ellipsoids have *the same scaling factors* \mathbf{D} , this number is *the same* as the minimal number of little balls $\{\tilde{\mathcal{B}}_1 \in \mathbb{R}^{\tilde{K}} : \|\tilde{\mathcal{B}}_1 - \tilde{\mathcal{B}}_2\|_2 \leq \varepsilon/36\}$ needed to cover $\{\tilde{\mathcal{B}}_1 \in \mathbb{R}^{\tilde{K}} : \|\tilde{\mathcal{B}}_1 - \tilde{\mathcal{B}}_2\|_2 \leq \sqrt{n}\varepsilon\}$. The links between coverings of ellipsoids and balls can be found, for instance, in Dumer [?]. This altogether implies that the covering number is bounded by $(108\sqrt{n})^{T\bar{K}}$.

Now we find an upper bound on the number of valid ensembles $\mathcal{E} \in \mathcal{V}\mathcal{E}_{\mathcal{S}}^K$ inside the sieve $\mathcal{F}_{\mathcal{E}}^n$. To start, we note that given $(T, \mathbf{q}, \mathbf{K}, \mathcal{S})$, there are at

most $\prod_{t=1}^T (K^t q^t n)^{K^t}$ valid ensembles $\mathcal{V}\mathcal{E}_{\mathcal{S}}^{\mathbf{K}}$. This bound is obtained from Lemma 3.1 by combining all possible T -tuples of trees.¹ Given (T, \mathbf{q}) , there are $\prod_{t=1}^T \binom{p}{q^t}$ sets of subsets $\mathcal{S} = \{\mathcal{S}^1, \dots, \mathcal{S}^T\}$ satisfying the constraint $|\mathcal{S}^t| = q^t$. This leads to an overall upper bound

$$\begin{aligned} & \sum_{T=1}^{\min\{s_n, z_n\}} \sum_{\mathbf{K}: \sum_{t=1}^T K^t \leq z_n} \sum_{\mathbf{q}: \sum_{t=1}^T q^t \leq s_n} \prod_{t=1}^T \binom{p}{q^t} (K^t q^t n)^{K^t} \\ & < \sum_{T=1}^{\min\{s_n, z_n\}} \sum_{\mathbf{K}: \sum_{t=1}^T K^t \leq z_n} (z_n s_n n)^{z_n} \sum_{\mathbf{q}: \sum_{t=1}^T q^t \leq s_n} \prod_{t=1}^T \left(\frac{pe}{q^t}\right)^{q^t} \\ & < s_n^{s_n+1} z_n^{z_n+1} (z_n s_n n)^{z_n} (pe)^{s_n}. \end{aligned}$$

Combining this bound with (9.2), we obtain the following bound

$$\begin{aligned} \log N\left(\frac{\varepsilon}{36}, \left\{f \in \mathcal{F}_{\mathcal{E}}^n : \|f - f_0\|_n < \varepsilon\right\}, \|\cdot\|_n\right) & < (s_n + 1) \log s_n + (z_n + 1) \log z_n \\ (9.3) \quad & + z_n \log(z_n s_n n) + s_n \log(pe) + z_n \log(108\sqrt{n}). \end{aligned}$$

Condition 2.1 will be met when (9.3) is smaller than (a constant multiple of) $n\varepsilon_n^2 = \sum_{t=1}^{T_0} n(\varepsilon_n^t)^2$. With the choice $z_n = \lfloor C_z n \varepsilon_n^2 / \log n \rfloor$ and $s_n = \lceil C_s n^{q_0 / (2\alpha + q_0)} \log^{2\beta} n / \log(p \vee n) \rceil$, where C_s and C_z are large enough constants to be determined later, this condition is satisfied.

9.2. *Condition (2.2).* To establish Condition (2.2) for tree ensembles, we begin by finding a single additive tree that approximates well. We will heavily leverage our findings from Section 8.2, noting that the problem of approximating an additive function f_0 with a sum of trees can be decomposed into smaller problems of approximating each layer f_0^t separately.

Denote by a_n^t the smallest leaf size of a k - d tree (defined in Remark 3.1) needed to approximate f_0^t with an error smaller than $\varepsilon_n^t/2$, where $\varepsilon_n^t = n^{-\alpha^t / (2\alpha^t + q_0^t)} \log^{\beta^t} n$. Such a tree step function approximation exists according to Lemma 3.2 when \mathcal{X} is (M, \mathcal{S}_0^t) -regular. We will denote this approximation with $f_{\hat{\mathcal{T}}^t, \hat{\beta}^t}(\mathbf{x})$. Moreover, with $\mathbf{a}_n = (a_n^1, \dots, a_n^{T_0})'$ we denote the vector of such minimal tree sizes, where each a_n^t satisfies (8.6) with q_0^t, α^t and ε_n^t . Next, we will denote by $\hat{\mathcal{E}} = \{\hat{\mathcal{T}}^1, \dots, \hat{\mathcal{T}}^{T_0}\}$ the approximating partition ensemble with step heights $\hat{\beta} = (\hat{\beta}^1, \dots, \hat{\beta}^{T_0})'$. The individual tree approximations $f_{\hat{\mathcal{T}}^t, \hat{\beta}^t}(\mathbf{x})$ are woven into an approximating forest $f_{\hat{\mathcal{E}}, \hat{\beta}}(\mathbf{x}) = \sum_{t=1}^{T_0} f_{\hat{\mathcal{T}}^t, \hat{\beta}^t}(\mathbf{x}) \in \mathcal{F}(\mathcal{V}\mathcal{E}_{\mathcal{S}_0}^{\mathbf{a}_n})$, where $\mathcal{S}_0 = \{\mathcal{S}_0^1, \dots, \mathcal{S}_0^{T_0}\}$.

¹The order of trees in \mathcal{E} matters.

Arguing as in Section 8.2, the statement $\|\widehat{\beta}^t - \beta^t\|_2 < \frac{\varepsilon_n^t}{2}$ for all $1 \leq t \leq T_0$ implies $\|f_0^t - f_{\widehat{\mathcal{T}}^t, \beta^t}\|_n < \varepsilon_n^t$ for all $1 \leq t \leq T_0$ and $\beta^t \in \mathbb{R}^{a_n^t}$. This further implies

$$\|f_0 - f_{\widehat{\mathcal{E}}, \mathcal{B}}\|_n \leq \sum_{t=1}^{T_0} \|f_0^t - f_{\widehat{\mathcal{T}}^t, \beta^t}\|_n \leq \sum_{t=1}^{T_0} \varepsilon_n^t \leq \sqrt{T_0} \varepsilon_n,$$

for any $\mathcal{B} = (\beta^{1'}, \dots, \beta^{T_0'})' \in \mathbb{R}^{T_0 \times \bar{a}_n}$, where the final inequality is due to Cauchy-Schwarz and where $\bar{a}_n = \frac{1}{T_0} \sum_{t=1}^{T_0} a_n^t$. Denote by $\mathcal{F}(\widehat{\mathcal{E}})$ the set of all additive trees supported on the partition ensemble $\widehat{\mathcal{E}}$. Then we can write

$$\begin{aligned} \Pi(f \in \mathcal{F}(\widehat{\mathcal{E}}) : \|f_0 - f\|_n \leq \varepsilon_n) &\geq \Pi\left(\beta \in \mathbb{R}^{\sum_{t=1}^{T_0} a_n^t} : \|\widehat{\beta}^t - \beta^t\|_2 \leq \frac{\varepsilon_n^t}{2\sqrt{T_0}} \text{ for each } t = 1, \dots, T_0\right) \\ &= \prod_{t=1}^{T_0} \Pi\left(\beta^t \in \mathbb{R}^{a_n^t} : \|\widehat{\beta}^t - \beta^t\|_2 \leq \frac{\varepsilon_n^t}{2\sqrt{T_0}}\right). \end{aligned}$$

Because we assumed $\beta_j^t \sim \mathcal{N}(0, 1/T)$, given T , we can directly use (8.9) to lower-bound the above with

$$(9.4) \quad \prod_{t=1}^{T_0} \frac{2^{-a_n^t} e^{-\|\widehat{\beta}^t\|_2^2 - (\varepsilon_n^t)^2/8}}{\Gamma(\frac{a_n^t}{2})^{\frac{a_n^t}{2}}} \left(\frac{(\varepsilon_n^t)^2}{4}\right)^{\frac{a_n^t}{2}}.$$

Because each tree $\widehat{\mathcal{T}}^t$ is a k - d tree and is by definition balanced, we have $\|\widehat{\beta}^t\|_n^2 \lesssim a_n^t \|f_0^t\|_\infty^2$. Now we can directly apply all our calculations from Section 8.2. In particular, using (9.4) and noting that $\Delta(\mathcal{V}\mathcal{E}_{\mathbf{S}_0}^{a_n}) < \prod_{t=1}^{T_0} \Delta(\mathcal{V}_{\mathbf{S}_0^t}^{a_n^t})$, we obtain

$$\begin{aligned} \Pi(f \in \mathcal{F}_{\mathcal{E}} : \|f_0 - f\|_n \leq \varepsilon_n) &\geq \pi(T_0) \pi(\mathbf{q}_0 | T_0) \pi(\mathbf{a}_n | T_0) \pi(\mathbf{S}_0 | T, \mathbf{q}_0) \times \\ &\quad \times \pi(\widehat{\mathcal{E}} | \mathbf{S}_0, \mathbf{a}_n) \Pi(f \in \mathcal{F}(\widehat{\mathcal{E}}) : \|f_0 - f_{\widehat{\mathcal{E}}, \mathcal{B}}\|_n \leq \varepsilon_n) \\ &> \pi(T_0) \prod_{t=1}^{T_0} L(q_0^t, a_n^t, \mathbf{S}_0^t, \widehat{\beta}^t, \varepsilon_n^t), \end{aligned}$$

where $L(\cdot)$ was defined in (8.8). It follows from Section (8.2) that

$$-\log L(q_0^t, a_n^t, \mathbf{S}_0^t, \widehat{\beta}^t, \varepsilon_n^t) \lesssim n(\varepsilon_n^t)^2$$

for each $1 \leq t \leq T_0$ when $q_0^t \lesssim \log^{\beta^t} n$, $\log p \lesssim \min_{1 \leq t \leq T_0} n^{q^t/(2\alpha^t + q_0^t)}$ and $T_0 \lesssim n$.

The last condition $T_0 \lesssim n$ is needed under our prior $K^t \sim \text{Poisson}(\lambda/T)$ so

that $-\log \pi(a_n) \lesssim a_n \log a_n + a_n \log T_0 \lesssim n(\varepsilon_n^t)^2 = nq_0^t/(2\alpha^t + q_0^t) \log^{2\beta^t} n$ for $\beta^t \geq 1/2$. Putting all the pieces together, we obtain the following lower bound

$$\Pi(f \in \mathcal{F}_\mathcal{E} : \|f_0 - f\|_n \leq \varepsilon_n) \geq \pi(T_0) e^{-dn \sum_{t=1}^{T_0} (\varepsilon_n^t)^2}$$

for some suitably large $d > 0$. The last requirement needed for Condition (2.2) to be satisfied is that $\pi(T_0) \geq e^{-dn\varepsilon_n^2}$. Our prior $\pi(T) \propto e^{-C_T T}$ safely satisfies this requirement.

9.3. *Condition (2.3)*. The condition entails showing that $\Pi(\mathcal{F}_\mathcal{E} \setminus \mathcal{F}_\mathcal{E}^n) = o(e^{-(d+2)n\varepsilon_n^2})$ for d deployed in the previous section. It suffices to show that

$$\left[\Pi \left((T, \mathbf{K}) : \sum_{t=1}^T K^t > z_n \right) + \Pi \left((T, \mathbf{q}) : \sum_{t=1}^T q^t > s_n \right) \right] e^{(d+2)n\varepsilon_n^2} \rightarrow 0.$$

Because we assume $K^t \mid T \stackrel{iid}{\sim} \text{Poisson}(\lambda/T)$ for some $\lambda \in \mathbb{R}$ (according to our definition in (T4*)), we can apply a similar Chernoff bound as in (8.14). Namely, we have for any $\gamma > 0$

$$\begin{aligned} \Pi \left(\sum_{t=1}^T K^t > z_n \right) &= \sum_{T=1}^{\infty} \pi(T) \Pi \left(\sum_{t=1}^T K^t > z_n \mid T \right) \\ &\lesssim e^{-\gamma(z_n+1)} \sum_{T=1}^{\infty} \pi(T) \left(e^{e^\gamma \lambda/T} - 1 \right)^T \lesssim e^{-\gamma(z_n+1) + e^\gamma \lambda}. \end{aligned}$$

With $z_n = \lfloor C_z n \varepsilon_n^2 / \log n \rfloor \sim \sum_{t=1}^{T_0} n q_0^t / (2\alpha^t + q_0^t) \log^{2\beta^t - 1} n$ and $\gamma = \log z_n$, we have $\Pi \left(\sum_{t=1}^T K^t > z_n \right) e^{dn\varepsilon_n^2} \rightarrow 0$ for a large enough constant $C_z > 0$. Next, with the independent product prior (T1*), the Chernoff bound gives

$$\Pi \left(\sum_{t=1}^T q^t > s_n \mid T \right) \leq e^{-\gamma(s_n+1)} \prod_{t=1}^T \mathbb{E} \left[e^{\gamma q^t} \right] \lesssim e^{-\gamma(s_n+1)} e^{-T \log[1 - e^\gamma / (cp^a)]}$$

for any $\gamma > 0$, where we used the fact

$$\mathbb{E} \left[e^{\gamma q^t} \right] = \sum_{q=0}^p [e^\gamma / (cp^a)]^q < \frac{1}{1 - e^\gamma / (cp^a)}.$$

With $\gamma = \log(p \vee n)$ and $a > 2$, we can write

$$\begin{aligned} \Pi \left(\sum_{t=1}^T q^t > s_n \mid T \right) &\lesssim e^{-(s_n+1) \log(p \vee n)} e^{-T \log[1 - 1/(cp^{a-1})]} \\ &< e^{-(s_n+1) \log(p \vee n) - T \log[1 - 1/(cp)]}. \end{aligned}$$

Next, we have

$$\Pi \left(\sum_{t=1}^T q^t > s_n \right) \lesssim e^{-(s_n+1) \log(p \vee n)} \sum_{T=1}^{\infty} \pi(T) e^{T \log[1+1/(cp-1)]}.$$

With $\pi(T) \propto e^{-C_T T}$, where $C_T > \log 2$, we have

$$\Pi \left(\sum_{t=1}^T q^t > s_n \right) \lesssim e^{-(s_n+1) \log(p \vee n)} \sum_{T=1}^{\infty} e^{-T(C_T - \log 2)} \lesssim e^{-(s_n+1) \log(p \vee n)}.$$

With $s_n = \lfloor C_s n \varepsilon_n^2 / \log(p \vee n) \rfloor$ we have $\Pi \left(\sum_{t=1}^T q^t > s_n \right) e^{(d+2)n\varepsilon_n^2} \rightarrow 0$ for a large enough constant C_q .

10. Proof of Theorem 5.1. The sieve will be very similar to (9.1). The only difference is that each tree in the ensemble is now constrained to depend on the same set of active variables \mathcal{S} . To mark this difference, we have denoted the partition ensembles with $\mathcal{V}\mathcal{E}_{\mathcal{S}}^{\mathbf{K}}$ instead of $\mathcal{V}\mathcal{E}_{\mathcal{S}}^{\mathbf{K}}$. Throughout this section, we use the following sieve:

$$(10.1) \quad \mathcal{F}_{\mathcal{E}}^n = \bigcup_{T=1}^{\infty} \bigcup_{q=0}^{q_n} \bigcup_{\mathcal{S}: |\mathcal{S}|=q} \bigcup_{\mathbf{K}: \sum_{t=1}^T K^t \leq z_n} \mathcal{F}(\mathcal{V}\mathcal{E}_{\mathcal{S}}^{\mathbf{K}}).$$

10.1. *Condition 2.1.* Our sieve (10.1) is embedded in (9.1), where the number of ensembles \mathcal{E} inside $\mathcal{F}_{\mathcal{E}}^n$ is now upper-bounded by

$$\sum_{T=1}^{z_n} \sum_{q=0}^{q_n} \sum_{\mathbf{K}: \sum_{t=1}^T K^t \leq z_n} \binom{p}{q} \prod_{t=1}^T (K^t q n)^{K^t} < z_n^{z_n+1} (q_n + 1) (z_n q_n n)^{z_n} (p e)^{q_n}.$$

Using the same arguments as in Section 9.1, we choose $z_n = \lfloor C_z n \varepsilon_n^2 / \log n \rfloor$ and $q_n = \lceil C_q \min\{p, n^{q_0/(2\alpha+q_0)} \log^{2\beta} n / \log(p \vee n)\} \rceil$, for which the condition holds.

10.2. *Condition (2.2).* The key ingredient for establishing Condition (2.2) is the following lemma on the existence of a tree ensemble that approximates f_0 well.

LEMMA 10.1. *Assume $f \in \mathcal{H}_p^{\alpha} \cap \mathcal{C}(\mathcal{S})$, where $|\mathcal{S}| = q$, and that \mathcal{X} is (M, \mathcal{S}) -regular. Then for any $s \in \mathbb{N} \setminus \{0\}$, there exists an additive tree function $f_{\hat{\mathcal{E}}, \hat{\mathcal{B}}} \in \mathcal{F}(\mathcal{V}\mathcal{E}_{\mathcal{S}}^{\mathbf{K}})$ consisting of $T = 2^{s q - 1}$ trees, each with $K^t = s q + 1$ leaves, such that*

$$\|f - f_{\hat{\mathcal{E}}, \hat{\mathcal{B}}}\|_n \leq \|f\|_{\mathcal{H}^{\alpha}} C M^{\alpha} q / K(\hat{\mathcal{E}})^{\alpha/q}$$

for some $C > 0$, where $K(\widehat{\mathcal{E}}) = 2^{s^q}$.

PROOF. From Lemma 3.2, we know that there exists a single tree function $f_{\widehat{\mathcal{T}}, \widehat{\beta}} \in \mathcal{F}(\mathcal{V}_S^K)$ with $\widehat{K} = 2^{s^q}$ leaves which approximates well. We regard the full symmetric tree $\widehat{\mathcal{T}}$ as the global partition of the approximating partition ensemble $\widehat{\mathcal{E}}$, i.e. $\widetilde{\mathcal{T}}(\widehat{\mathcal{E}}) = \widehat{\mathcal{T}}$. Moreover, $\widehat{\beta}$ is regarded as the vector of aggregated steps, i.e. $\widetilde{\beta} = \widehat{\beta}$ (the definition of the aggregated steps is in (5.4)). The actual ensemble $\widehat{\mathcal{E}}$ is obtained from $\widehat{\mathcal{T}}$ by redistributing the cuts among T trees, each with $K^t = K$ leaves, in the following way. We take completely imbalanced trees that keep refining one cell until the resolution reaches the tree depth $\log_2 \widehat{K}$. These trees have $K^t = \log_2 \widehat{K} + 1$ leaves and we need $T = \widehat{K}/2$ of those to sum up towards $\widehat{\mathcal{T}}$. This decomposition is illustrated in Figure 3, where a full symmetric tree $\widehat{\mathcal{T}}$ (Figure 4) with $\widehat{K} = 8$ leaves has been trimmed into $T = \widehat{K}/2 = 4$ smaller trees with $K^t = \log_2 \widehat{K} + 1 = 4$ leaves. The decomposition yields a tree ensemble $\widehat{\mathcal{E}} = \{\mathcal{T}^1, \dots, \mathcal{T}^T\}$ with a stretching matrix $\mathbf{A}(\widehat{\mathcal{E}}) = [\mathbf{I}_{\widehat{K}}, \mathbf{A}_1]$ (after a suitable permutation of columns), where \mathbf{A}_1 is some binary matrix. It follows from Lemma 1(g) of Govaerts and Pryce [?] that $\lambda_{min}^2(\widehat{\mathcal{E}}) = 1 + \sigma_{min}^2(\mathbf{A}_1) \geq 1$, where $\sigma_{min}(\mathbf{A}_1)$ denotes the smallest singular value of \mathbf{A}_1 . Moreover, we have $K(\widehat{\mathcal{E}}) = \widehat{K}$. Finally, we use (5.4) to obtain the individual tree steps via $\widehat{\beta} = (\mathbf{A}(\widehat{\mathcal{E}})' \mathbf{A}(\widehat{\mathcal{E}}))^\dagger \mathbf{A}(\widehat{\mathcal{E}})' \widetilde{\beta}$, where \mathbf{A}^\dagger denotes the Moore pseudoinverse of \mathbf{A} . The rest follows from Lemma 3.2. \square

Now we proceed with Condition 2.2. Denote by $\widehat{\mathcal{E}}$ the approximating ensemble from Lemma 10.1. Recall that the global partition $\widetilde{\mathcal{T}}(\widehat{\mathcal{E}}) = \{\widetilde{\Omega}_k\}_{k=1}^{K(\widehat{\mathcal{E}})}$ is a k - d tree, which is balanced in the sense that $C_{min}^2/K(\widehat{\mathcal{E}}) \leq \mu(\widetilde{\Omega}_k) \leq C_{max}^2/K(\widehat{\mathcal{E}})$ for some constants $C_{min} < 1 < C_{max}$ and $k = 1, \dots, K(\widehat{\mathcal{E}})$. Next, we find the smallest $K(\widehat{\mathcal{E}})$ such that $\|f\|_{\mathcal{H}^\alpha} C M^\alpha q_0 / K(\widehat{\mathcal{E}})^{\alpha/q_0} < \varepsilon_n/2$. This value will be denoted by a_n and it satisfies (8.6). Next, we denote by $\widehat{T} = a_n/2$ the number of approximating trees and by $\widehat{\mathbf{K}} = (\widehat{K}^1, \dots, \widehat{K}^T)'$ the vector of leaves, where $\widehat{K}^t = \log_2 a_n + 1$ (again we are using the construction from Lemma 10.1). Then, using similar arguments as in Section 9.2 we can lower-bound $\Pi(f \in \mathcal{F}_{\mathcal{E}} : \|f - f_0\|_n \leq \varepsilon_n)$ with

$$(10.2) \quad \frac{\pi(\widehat{T})\pi(\widehat{\mathbf{K}}|\widehat{T})\pi(q_0)}{\left(\frac{ep}{q_0}\right)^{q_0} \prod_{t=1}^{\widehat{T}} (\widehat{K}^t q_0 n)^{\widehat{K}^t}} \Pi(f \in \mathcal{F}(\widehat{\mathcal{E}}) : \|f - f_0\|_n \leq \varepsilon_n),$$

where $\mathcal{F}(\widehat{\mathcal{E}})$ consists of all additive tree functions supported on $\widehat{\mathcal{E}}$. We denote by $\widetilde{a}_n = \sum_{t=1}^{\widehat{T}} \widehat{K}^t = a_n/2(\log_2 a_n + 1)$ and by $\widehat{\beta} \in \mathbb{R}^{\widetilde{a}_n}$ the steps of the approximating additive trees from Lemma 10.1. Because $\mu(\widetilde{\Omega}_k) \leq C_{max}^2/K(\widehat{\mathcal{E}})$

we obtain for any arbitrary vector $\mathcal{B} \in \mathbb{R}^{\tilde{a}_n}$ (similarly as in Section 9.1)

$$\|f_{\hat{\mathcal{E}}, \mathcal{B}} - f_{\hat{\mathcal{E}}, \hat{\mathcal{B}}}\|_n \leq C_{\max} \lambda_{\max}(\hat{\mathcal{E}}) / \sqrt{K(\hat{\mathcal{E}})} \|\mathcal{B} - \hat{\mathcal{B}}\|_2$$

and thereby

$$\|\mathcal{B} - \hat{\mathcal{B}}\|_2 \geq \sqrt{K(\hat{\mathcal{E}})} / (C_{\max} \lambda_{\max}(\hat{\mathcal{E}})) \left| \|f_0 - f_{\hat{\mathcal{E}}, \mathcal{B}}\|_n - \|f_0 - f_{\hat{\mathcal{E}}, \hat{\mathcal{B}}}\|_n \right|.$$

Combined with the fact $\lambda_{\max}^2(\mathcal{E}) \leq K(\mathcal{E}) \tilde{a}_n$ (as shown in the proof of Lemma 12.1), the statement $\|\mathcal{B} - \hat{\mathcal{B}}\|_2 < \frac{\varepsilon_n}{2} \frac{1}{C_{\max} \sqrt{\tilde{a}_n}}$ implies $\|f_0 - f_{\hat{\mathcal{E}}, \mathcal{B}}\|_n < \varepsilon_n$. Therefore we have

$$\Pi(f \in \mathcal{F}(\hat{\mathcal{E}}) : \|f - f_0\|_n \leq \varepsilon_n) > \Pi\left(\mathcal{B} \in \mathbb{R}^{\tilde{a}_n} : \|\mathcal{B} - \hat{\mathcal{B}}\|_2 < \frac{\varepsilon_n}{2} \frac{1}{C_{\max} \sqrt{\tilde{a}_n}}\right).$$

Moreover, because $\mu(\tilde{\Omega}_k) \geq C_{\min}^2 / K(\hat{\mathcal{E}})$ for some $C_{\min} < 1$, we have

$$\|\hat{\mathcal{B}}\|_2 \leq \frac{\sqrt{a_n}}{C_{\min} \lambda_{\min}(\hat{\mathcal{E}})} \|f_{\hat{\mathcal{E}}, \hat{\mathcal{B}}}\|_n \leq \frac{\sqrt{a_n}}{C_{\min}} \left(\frac{\varepsilon_n}{2} + \|f_0\|_\infty \right),$$

where we used the fact $\lambda_{\min}^2(\hat{\mathcal{E}}) \geq 1$ (proof of Lemma 10.1). Therefore we have $\|\hat{\mathcal{B}}\|_2^2 \leq C_2 a_n \|f_0\|_\infty^2$ for some $C_2 > 0$. Following the calculations from Section 8.2 (namely (8.10)), we continue to lower-bound (10.2) with (10.3)

$$\frac{\pi(\hat{T}) \pi(\hat{\mathbf{K}} | \hat{T}) \pi(q_0) e^{-\frac{\varepsilon_n^2}{8C_{\max}^2 \tilde{a}_n} - a_n (C_2 \|f_0\|_\infty^2 + \log 2)}}{\left(\frac{ep}{q_0}\right)^{q_0} [(\log_2 a_n + 1) q_0 n]^{\tilde{a}_n}} \left(\frac{\varepsilon_n^2}{4C_{\max}^2 \tilde{a}_n}\right)^{\frac{\tilde{a}_n}{2}} \left(\frac{2}{\tilde{a}_n}\right)^{\tilde{a}_n/2+1}.$$

This quantity should be at least $e^{-dn\varepsilon_n^2}$ for some suitably large $d > 0$. Now, with our prior (T4*) we can write

$$\pi(\hat{\mathbf{K}} | \hat{T}) \gtrsim \prod_{t=1}^{\hat{T}} e^{\hat{K}^t \log(\lambda/\hat{T}) - \hat{K}^t \log \hat{K}^t} \gtrsim e^{-\tilde{a}_n \log(\log_2 a_n + 1) + \tilde{a}_n \log(2\lambda/a_n)}.$$

This quantity can be lower-bounded by $e^{-D\tilde{a}_n \log a_n}$ for some $D > 0$. Then we can write

$$\frac{\pi(\hat{T}) \pi(\hat{\mathbf{K}} | \hat{T}) \pi(q_0)}{\left(\frac{ep}{q_0}\right)^{q_0} [(\log_2 a_n + 1) q_0 n]^{\tilde{a}_n}} > e^{-C_T a_n/2} e^{-D\tilde{a}_n \log a_n - q_0 \log(c e p^{a+1}/q_0)} e^{-\tilde{a}_n \log(q_0 n (\log_2 a_n + 1))}.$$

By our assumptions $q_0 \lesssim \log^\beta n$ and $p \lesssim n^{q_0/(2\alpha+q_0)}$, the term $e^{-q_0 \log(c e p^{a+1}/q_0)}$ will safely be larger than $e^{-d_1 n \varepsilon_n^2}$ for some $d_1 > 0$. We take all the remaining important terms in (10.3), aiming to show that (i) $\tilde{a}_n \log(\tilde{a}_n/\varepsilon_n^2)$, (ii) $a_n \|f_0\|_\infty^2$, (iii) $\tilde{a}_n \log(q_0 n \log_2 a_n)$ and (iv) $\tilde{a}_n \log \tilde{a}_n$ are bounded by a constant multiple of $n \varepsilon_n^2$. From (8.6), we obtain $a_n \lesssim n^{q_0/(2\alpha+q_0)}$ under our assumption $2C_0 q_0 \lesssim \log^\beta n$. Then we can write

$$(10.4) \quad \tilde{a}_n = a_n/2(\log_2 a_n + 1) \lesssim n^{q_0/(2\alpha+q_0)} \log n.$$

Using this bound, we verify that (i)-(iv) are bounded by a constant multiple of $n \varepsilon_n^2 = n^{\frac{q_0}{2\alpha+q_0}} \log^{2\beta} n$. First, note that

$$\tilde{a}_n \log(\tilde{a}_n/\varepsilon_n^2) \lesssim n^{\frac{q_0}{2\alpha+q_0}} \log n \log \left(n^{\frac{q_0}{2\alpha+q_0}} n^{\frac{2\alpha}{2\alpha+q_0}} \log^{1-2\beta} n \right).$$

This quantity is bounded by a multiple of $n^{\frac{q_0}{2\alpha+q_0}} \log^{2\beta} n$ when $\beta \geq 1$. Next, we can write $a_n \|f_0\|_\infty^2 \lesssim n^{\frac{q_0}{2\alpha+q_0}} \log^{2\beta} n$. Lastly, it follows from (10.4) that $\tilde{a}_n \log \tilde{a}_n \lesssim n^{q_0/(2\alpha+q_0)} \log^2 n$ and $\tilde{a}_n \log(q_0 n \log_2 a_n) \lesssim n^{q_0/(2\alpha+q_0)} \log^2 n$. To sum up, there exists $d > 0$ such that $\Pi(f \in \mathcal{F}_\varepsilon : \|f - f_0\|_n \leq \varepsilon_n) > e^{-d n \varepsilon_n^2}$ for $\beta \geq 1$.

10.3. *Condition (2.3)*. The condition $\Pi(\mathcal{F}_\varepsilon \setminus \mathcal{F}_\varepsilon^n) e^{(d+2)n\varepsilon_n^2} \rightarrow 0$ is verified similarly as in Section 9.3 and Section 8.3. For $\Pi(q > q_n)$, we use the bound from Section 8.3 with $q_n = \lceil C_q \min\{p, n^{q_0/(2\alpha+q_0)} \log^{2\beta} n / \log(p \vee n)\} \rceil$ and a large enough constant C_q . For $\Pi\left((T, \mathbf{K}) : \sum_{t=1}^T K^t > z_n\right)$ we use the bound from Section 9.3 with $z_n = \lfloor C_z n \varepsilon_n^2 / \log n \rfloor$ and a large enough constant C_z .

11. Proof of Lemma 3.2. We start with an auxiliary statement showing that when f is α -Hölder continuous, we can grow a step function on any given (tree) partition so that the approximation error will be governed by cell diameters.

Namely, for $f \in \mathcal{H}_p^\alpha \cap \mathcal{C}(\mathcal{S})$ and a valid tree partition $\mathcal{T} = \{\Omega_k\}_{k=1}^K \in \mathcal{V}_\mathcal{S}^K$, there exists a step function $f_{\mathcal{T}, \hat{\beta}}(\mathbf{x}) = \sum_{k=1}^K \hat{\beta}_k \mathbb{I}_{\Omega_k}(\mathbf{x})$ such that $\|f - f_{\mathcal{T}, \hat{\beta}}\|_n \leq \|f\|_{\mathcal{H}^\alpha} \sqrt{\sum_{k=1}^K \mu(\Omega_k) \text{diam}^{2\alpha}(\Omega_k; \mathcal{S})}$. Indeed, given \mathcal{T} and design points \mathcal{X} , we take $\hat{\beta}_k = \frac{1}{n_k} \sum_{i=1}^{n_k} f(\mathbf{x}_i) \mathbb{I}_{\Omega_k}(\mathbf{x}_i)$, where $n_k = \mu(\Omega_k)n$. Then for $\mathbf{x}_j \in \Omega_k \cap \mathcal{X}$ we have, from Hölder continuity,

$$\begin{aligned} |f(\mathbf{x}_j) - f_{\mathcal{T}, \hat{\beta}}(\mathbf{x}_j)| &< \frac{1}{n_k} \sum_{\mathbf{x}_i \in \Omega_k} |f(\mathbf{x}_j) - f(\mathbf{x}_i)| \\ &\leq \|f\|_{\mathcal{H}^\alpha} \frac{1}{n_k} \sum_{\mathbf{x}_i \in \Omega_k} \|\mathbf{x}_{j\mathcal{S}} - \mathbf{x}_{i\mathcal{S}}\|_2^\alpha \leq \|f\|_{\mathcal{H}^\alpha} \text{diam}^\alpha(\Omega_k; \mathcal{S}). \end{aligned}$$

Then the approximation error satisfies

$$(11.1) \quad \|f - f_{\mathcal{T}, \hat{\beta}}\|_n \leq \|f\|_{\mathcal{H}^\alpha} \sqrt{\sum_{k=1}^K \mu(\Omega_k) \text{diam}^{2\alpha}(\Omega_k; \mathcal{S})}.$$

To continue with the proof of (3.5), we grow a k - d tree partition $\hat{\mathcal{T}} = \{\hat{\Omega}_k\}_{k=1}^K$ (as explained in Remark 3.1) and construct an approximating step function $f_{\hat{\mathcal{T}}, \hat{\beta}}$, as outlined above.

Using (11.1), the statement (3.5) then follows from

$$\|f - f_{\hat{\mathcal{T}}, \hat{\beta}}\|_n < C_{\max} \|f\|_{\mathcal{H}^\alpha} \max_{1 \leq k \leq K} \text{diam}^\alpha(\hat{\Omega}_k; \mathcal{S}),$$

where we used the fact that $\mu(\hat{\Omega}_k) \leq C_{\max}^2/K$ in k - d trees. A minor modification of the proof of Proposition 6 in [44] yields $\sum_{k=1}^K \mu(\hat{\Omega}_k) \text{diam}(\hat{\Omega}_k; \mathcal{S}) \leq \frac{q}{K^{1/q}}$. The rest follows from Definition 3.3. \square

12. Auxiliary Result.

LEMMA 12.1. *Assume a valid ensemble \mathcal{E} consisting of T trees, each with K^t leaves. Let $\lambda_{\max}^2(\mathcal{E})$ be the largest eigenvalue of $\tilde{\mathbf{A}}(\mathcal{E}) = \mathbf{A}(\mathcal{E})' \mathbf{A}(\mathcal{E})$. Then*

$$(12.1) \quad \lambda_{\max}^2(\mathcal{E}) \leq K(\mathcal{E})(T \times \bar{K}),$$

where $\bar{K} = \frac{1}{T} \sum_{t=1}^T K^t$ and where $K(\mathcal{E})$ denotes the number of rows of $\mathbf{A}(\mathcal{E})$.

PROOF. By the Gershgorin circle theorem, all eigenvalues of $\tilde{\mathbf{A}}(\mathcal{E}) = (\tilde{a}_{ij})$ lie inside the union of intervals $[\tilde{a}_{ii} - \sum_{j \neq i} \tilde{a}_{ij}, \tilde{a}_{ii} + \sum_{j \neq i} \tilde{a}_{ij}]$ for $i = 1, \dots, T \times \bar{K}$. As explained in Section 5.1, the diagonal and off-diagonal entries of $\tilde{\mathbf{A}}(\mathcal{E})$ quantify the persistence and the overlap in terms of the number of intersecting global partitioning cells. The magnitude $|\tilde{a}_{ij}|$ is no larger than $K(\mathcal{E})$ for each $1 \leq i, j \leq T \times \bar{K}$. The upper bound on the maximal eigenvalue $\lambda_{\max}^2(\mathcal{E})$ is thus $K(\mathcal{E})(T \times \bar{K})$. \square

References.

- [1] T. W. Anderson. Some nonparametric multivariate procedures based on statistically equivalent blocks. *Multivariate Analysis*, 1:5–27, 1966.
- [2] J. L. Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, 1975.
- [3] A. Bhattacharya, D. Pati, and D. Dunson. Anisotropic function estimation using multi-bandwidth Gaussian processes. *Annals of Statistics*, 42:352–381, 2014.
- [4] G. Biau. Analysis of a random forests model. *The Journal of Machine Learning Research*, 13(1):1063–1095, 2012.

- [5] G. Biau, L. Devroye, and G. Lugosi. Consistency of random forests and other averaging classifiers. *The Journal of Machine Learning Research*, 9:2015–2033, 2008.
- [6] G. Biau and E. Scornet. A random forest guided tour. *Test*, 25(2):197–227, 2016.
- [7] J. Bleich, A. Kapelner, E. I. George, and S. Jensen. Variable selection for BART: An application to gene regulation. *The Annals of Applied Statistics*, 8(3):1750–1781, 2014.
- [8] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- [9] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification and Regression Trees*. Wadsworth and Brooks, 1984.
- [10] I. Castillo. Pólya tree posterior distributions on densities. *Annales de l’Institut Henri Poincaré (to appear)*, 2016.
- [11] I. Castillo and J. Rousseau. A Bernstein-von Mises theorem for smooth functionals in semiparametric models. *The Annals of Statistics*, 43:2353–2383, 2015.
- [12] I. Castillo, J. Schmidt-Hieber, and A.W. van der Vaart. Bayesian linear regression with sparse priors. *The Annals of Statistics*, 43:1986–2018, 2015.
- [13] I. Castillo and A.W. van der Vaart. Needles and straw in a haystack: posterior concentration for possibly sparse sequences. *The Annals of Statistics*, 40:2069–2101, 2012.
- [14] H. Chipman, E. I. George, and R. McCulloch. BART: Bayesian additive regression trees. *Annals of Applied Statistics*, 4:266–298, 2010.
- [15] H. Chipman, E. I. George, and R. E. McCulloch. Bayesian CART model search. *Journal of the American Statistical Association*, 93:935–960, 1997.
- [16] M. Coram and S. Lalley. Consistency of Bayes estimators of a binary regression function. *Annals of Statistics*, 34:1233–1269, 2006.
- [17] R. de Jonge and H. van Zanten. Semiparametric Bernstein-von Mises for the error standard deviation. *Electron. J. Statist.*, 7:217–243, 2013.
- [18] D. Denison, B. Mallick, and A. Smith. A Bayesian CART algorithm. *Biometrika*, 85:363–377, 1998.
- [19] K. Devroye and L. Györfi. Distribution-free exponential bounds on the l_1 error of partitioning estimates of a regression function. *In proceedings of the fourth Pannonian Symposium of Mathematical Statistics*, pages 67–76, 1985.
- [20] L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer series: Stochastic Modelling and Applied Probability, 1996.
- [21] D. Donoho. CART and best-ortho-basis: a connection. *Annals of Statistics*, 25:1870–1911, 1997.
- [22] J. Engel. A simple wavelet approach to nonparametric regression from recursive partitioning schemes. *Journal of Multivariate Analysis*, 49(2):242–254, 1994.
- [23] S. Ghosal, J. Ghosh, and A. van der Vaart. Convergence rates of posterior distributions. *Annals of Statistics*, 28:500–5311, 2000.
- [24] S. Ghosal and A. van der Vaart. Convergence rates of posterior distributions for noniid observations. *Annals of Statistics*, 35:192–223, 2007.
- [25] L. Gordon and R. Olshen. Consistent nonparametric regression from recursive partitioning schemes. *Journal of Multivariate Analysis*, 10:611–627, 1980.
- [26] L. Gordon and R. Olshen. Almost sure consistent nonparametric regression from recursive partitioning schemes. *Journal of Multivariate Analysis*, 15:147–163, 1984.
- [27] S. Kpotufe. *The curse of dimension in nonparametric regression*. University of California, San Diego, 2010.
- [28] H. Lian. Consistency of Bayesian estimation of a step function. *Statistics & probability letters*, 77(1):19–24, 2007.
- [29] A. R. Linero. Bayesian regression trees for high dimensional prediction and variable

- selection. To appear in the Journal of the American Statistical Association, 2016.
- [30] A. R. Linero and Yang Y. Bayesian regression tree ensembles that adapt to smoothness and sparsity. arXiv:1707.09461, 2018.
 - [31] L. Liu and W. H. Wong. Multivariate density estimation via adaptive partitioning (ii): Posterior concentration. arXiv preprint arXiv:1508.04812, 2015.
 - [32] Y. Liu, V. Ročková, and Y. Wang. ABC variable selection with Bayesian forests. *Submitted Manuscript*, pages 1–50, 2018.
 - [33] L. Lu, H. Jiang, and W. Wong. Multivariate density estimation by Bayesian sequential partitioning. *Journal of the American Statistical Association*, 108:1402–1410, 2013.
 - [34] A. Nobel. Histogram regression estimation using data-dependent partitions. *Annals of Statistics*, 24:1084–1105, 1996.
 - [35] V. Ročková and E. Saha. On theory for BART. *22nd International Conference on Artificial Intelligence and Statistics*, 2019.
 - [36] E. Scornet, G. Biau, and J.P. Vert. Consistency of random forests. *Annals of Statistics*, 43:1716–1741, 2015.
 - [37] C. Scricciolo. On rates of convergence for Bayesian density estimation. *Scandinavian Journal of Statistics*, 34:626–642, 2007.
 - [38] X. Shen and L. Wasserman. Rates of convergence of posterior distributions. *Annals of Statistics*, 29:687–714, 2001.
 - [39] C. Stone. Optimal global rates of convergence for nonparametric regression. *Annals of Statistics*, 10:1040–1053, 1982.
 - [40] C. Stone. An asymptotically optimal histogram selection rule. *In proceedings of the Berkeley conference in honor of Jerzy Neyman and Jack Kiefer*, pages 513–530, 1985.
 - [41] C. Stone. Consistent nonparametric regression. *Annals of Statistics*, 5:595–645, 1985.
 - [42] S. van der Pas and V. Ročková. Bayesian dyadic trees and histograms for regression. *Advances in Neural Information Processing Systems (accepted)*, pages 1–12, 2017.
 - [43] A. van der Vaart and J. van Zanten. Rates of contraction of posterior distributions based on Gaussian process priors. *The Annals of Statistics*, 36:1435–1463, 2008.
 - [44] N. Verma, S. Kpotufe, and S. Dasgupta. Which spatial partition trees are adaptive to intrinsic dimension? In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*, pages 565–574. AUAI Press, 2009.
 - [45] S. Wager and G. Walther. Adaptive concentration of regression trees, with application to random forests. *arXiv preprint arXiv:1503.06388*, 2015.
 - [46] Y. Yang and S. Tokdar. Minimax-optimal nonparametric regression in high dimensions. *Annals of Statistics*, 43:652–674, 2015.
 - [47] W. Yoo and S. Ghosal. Supremum norm posterior contraction and credible sets for nonparametric multivariate regression. *The Annals of Statistics*, 44(3):1069–1102, 06 2016.