
Bayesian Dyadic Trees and Histograms for Regression

Stéphanie van der Pas and Veronika Ročková

Abstract

Many machine learning tools for regression are based on recursive partitioning of the covariate space into smaller regions, where the regression function can be estimated locally. Among these, regression trees and their ensembles have demonstrated impressive empirical performance. In this work, we shed light on the machinery behind Bayesian variants of these methods. In particular, we study Bayesian regression histograms, such as Bayesian dyadic trees, in the simple regression case with just one predictor. We focus on the reconstruction of regression surfaces that are piecewise constant, where the number of jumps is unknown. We show that with suitably designed priors, posterior distributions concentrate around the true step regression function at the minimax rate (up to a log factor). These results *do not* require the knowledge of the true number of steps, nor the width of the true partitioning cells. Thus, Bayesian dyadic regression trees are fully adaptive and can recover the true piecewise regression function nearly as well as if we knew the exact number and location of jumps. Our results constitute the first step towards understanding why Bayesian trees and their ensembles have worked so well in practice. As an aside, we discuss prior distributions on balanced interval partitions and how they relate to a problem in geometric probability. Namely, we quantify the probability of covering the circumference of a circle with random arcs whose endpoints are confined to a grid, a new variant of the original problem.

1 Introduction

Histogram regression methods, such as regression trees [1] and their ensembles [2], have an impressive record of empirical success in many areas of application [3, 4, 5, 6, 7]. Tree-based machine learning (ML) methods build a piecewise constant reconstruction of the regression surface based on ideas of recursive partitioning. Perhaps the most popular partitioning schemes are the ones based on parallel-axis splits. One recent example is the Mondrian process [8], which was introduced to the ML community as a prior over tree data structures with interesting self-consistency properties. Many efficient algorithms exist that can be deployed to fit regression histograms underpinned by some partitioning scheme. Among these, Bayesian variants, such as Bayesian CART [9, 10] and BART [11], have appealed to umpteen practitioners. There are several reasons why. Bayesian tree-based regression tools (a) can adapt to regression surfaces without any need for pruning, (b) are reluctant to overfit, (c) provide an avenue for uncertainty statements via posterior distributions. While practical success stories abound [3, 4, 5, 6, 7], the theoretical understanding of Bayesian regression tree methods has largely been lacking. In this work, we study the quality of posterior distributions with regard to the three properties mentioned above. We provide first theoretical results that contribute to the understanding of Bayesian Gaussian regression methods based on partitioning.

Our performance metric will be the speed of posterior concentration/contraction around the true regression function. This is ultimately a frequentist assessment, describing the average behavior of the posterior under the true generative model [12]. Posterior concentration results are now slowly entering the machine learning community as a tool for obtaining more insight into Bayesian methods [13, 14, 15, 16, 17]. Typically, our estimator will be some measure of center of the posterior distribution, such as the mean, median or mode. Posterior concentration results quantify not only the average distance between our estimator and the truth, but also the average spread of the posterior

around the truth. Ideally, most posterior mass should be concentrated in a ball centered around the true value with a radius proportional to the square root of the minimax rate [12, 18]. Being inherently a performance measure of both location and spread, optimal posterior concentration provides a necessary certificate for further uncertainty quantification [19, 20, 21]. Beyond uncertainty assessment, theoretical guarantees that describe the average posterior shrinkage behavior have also been a valuable instrument for assessing the suitability of priors. As such, these results can often provide useful guidelines for the choice of tuning parameters, e.g. the latent Dirichlet allocation model [14].

Despite the rapid growth of this Bayesian-frequentist theory field, posterior concentration results for Bayesian regression histograms/trees/forests have, so far, been unavailable. Here, we adopt this theoretical framework to get new insights into why these methods tend to work so well.

Related Work

Bayesian density estimation with step functions is a relatively well-studied problem [22, 23, 24]. The literature on Bayesian histogram regression is a bit less crowded. Perhaps the closest to our conceptual framework is the work by Coram and Lalley [25], who studied Bayesian non-parametric binary regression with uniform mixture priors on step functions. The authors focused on L_1 consistency. Here, we focus on posterior concentration rather than consistency. We are not aware of any related theoretical study of Bayesian histogram methods for Gaussian regression.

Our Contributions

In this work we focus on a canonical regression setting with merely one predictor. We study hierarchical priors on step functions and provide conditions under which the posteriors concentrate optimally around the true regression function. We consider the case when the true regression function itself is a step function, i.e. a tree or a tree ensemble, where the number and location of jumps is unknown.

We start with a very simple space of approximating step functions, supported on equally sized intervals where the number of splits is equipped with a prior. These partitions include dyadic regression trees. We show that for a suitable complexity prior, all relevant information about the true regression function (jump sizes and the number of jumps) is learned from the data automatically. During the course of the proof, we develop a notion of the complexity of a piecewise constant function relative to its approximating class, and find that the prior should pile on functions with at most order $\log n$ distinct pieces.

Next, we take a larger approximating space consisting functions supported on balanced partitions that do not necessarily have to be of equal size. These correspond to more general trees. With a uniform prior over all balanced partitions, we are able to achieve a nearly ideal performance (as if we knew the number and the location of jumps). As an aside, we describe the distribution of interval lengths obtained when the splits are sampled uniformly from a grid. We relate this distribution to the probability of covering the circumference of a circle with random arcs, a problem in geometric probability that dates back to [26, 27]. Our version of this problem assumes that the splits are chosen from a discrete grid rather than from a unit interval.

Notation

With \asymp and \preceq will denote an equality and inequality, up to a constant. 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 the set Ω . We denote by $\phi(\cdot)$ the standard normal density and by $P_f^n = \otimes P_{f,i}$ the n -fold product measure of the n independent observations under (1) with a regression function $f(\cdot)$. By $\mathbb{P}_n^x = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ we denote the empirical distribution of the observed covariates, by $\|\cdot\|_n$ the norm on $L_2(\mathbb{P}_n^x)$ and by $\|\cdot\|_2$ the standard Euclidean norm.

2 Bayesian Histogram Regression

We consider a classical nonparametric regression model, where response variables $\mathbf{Y}^{(n)} = (Y_1, \dots, Y_n)'$ are related to input variables $\mathbf{x}^{(n)} = (x_1, \dots, x_n)'$ through the function f_0 as follows

$$Y_i = f_0(x_i) + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, 1), \quad i = 1, \dots, n. \quad (1)$$

We assume that the covariate values x_i are one-dimensional and have been rescaled so that $x_i \in (0, 1]$. The underlying regression function $f_0 : [0, 1] \rightarrow \mathbb{R}$ is assumed to be a step function, i.e.

$$f_0(x) = \sum_{k=1}^{K_0} \beta_k^0 \mathbb{I}_{\Omega_k^0}(x),$$

where $\{\Omega_k^0\}_{k=1}^{K_0}$ is a partition of $[0, 1]$ into K_0 non-overlapping intervals. We assume that $\{\Omega_k^0\}_{k=1}^{K_0}$ is minimal, meaning that f_0 cannot be represented with a smaller partition (with less than K_0 pieces). Each partitioning cell Ω_k^0 is associated with a step size β_k^0 , determining the level of the function f_0 on Ω_k^0 . The entire vector of K_0 step sizes will be denoted by $\beta^0 = (\beta_1^0, \dots, \beta_{K_0}^0)'$.

One might like to think of f_0 as a regression tree with K_0 bottom leaves. Indeed, every step function can be associated with an equivalence class of trees that live on the same partition but differ in their tree topology. The number of bottom leaves K_0 will be treated as unknown throughout this paper. Our goal will be designing a suitable class of priors on step functions so that the posterior concentrates tightly around f_0 . Our analysis with a single predictor has served as a precursor to a full-blown analysis for high-dimensional regression trees (our followup manuscript).

We consider an approximating space of all step functions (with $K = 1, 2, \dots$ bottom leaves)

$$\mathcal{F} = \cup_{K=1}^{\infty} \mathcal{F}_K, \quad (2)$$

which consists of smaller spaces (or shells) of all K -step functions

$$\mathcal{F}_K = \left\{ f_{\beta} : [0, 1] \rightarrow \mathbb{R}; f_{\beta}(x) = \sum_{k=1}^K \beta_k \mathbb{I}_{\Omega_k}(x) \right\},$$

each indexed by a partition $\{\Omega_k\}_{k=1}^K$ and a vector of K step heights β . The fundamental aspect of our theoretical analysis is how to build a prior on \mathcal{F} . This prior distribution has three main ingredients, described in detail below, (a) a prior on the number of steps K , (b) a prior on the partitions $\{\Omega_k\}_{k=1}^K$ of size K , and (c) a prior on step sizes $\beta = (\beta_1, \dots, \beta_K)'$.

2.1 Prior $\pi_K(\cdot)$ on the Number of Steps K

To avoid overfitting, we assign an exponentially decaying prior distribution that penalizes partitions with too many jumps.

Definition 2.1. (Prior on K) *The prior on the number of partitioning cells K satisfies*

$$\pi_K(k) \equiv \Pi(K = k) \propto \exp(-c_K k \log k). \quad (3)$$

This prior is no stranger to non-parametric problems. It was deployed for stepwise reconstructions of densities [24, 23] and regression surfaces [25]. When c_K is large, this prior is concentrated on models with small complexity where overfitting should not occur. Decreasing c_K leads to the smearing of the prior mass over partitions with more jumps. This is illustrated in Figure 1, which depicts the prior for various choices of c_K . We provide recommendations for the choice of c_K in Section 3.3.

2.2 Prior $\pi_{\Omega}(\cdot | K)$ on Interval Partitions $\{\Omega_k\}_{k=1}^K$

After selecting the number of steps K from $\pi_K(k)$, we assign a prior over interval partitions $\pi_{\Omega}(\cdot | K)$. We will consider two important special cases.

2.2.1 Equivalent Blocks

Perhaps the simplest partition is based on statistically equivalent blocks [28], where all the cells are required to have the same number of points. This is also known as the K -spacing rule that partitions the unit interval using the order statistics of the observations.

Definition 2.2. (Equivalent Blocks) *Let $x_{(i)}$ denote the i^{th} order statistic of $\mathbf{x} = (x_1, \dots, x_n)'$, where $x_{(n)} \equiv 1$ and $n = Kc$ for some $c \in \mathbb{N} \setminus \{0\}$. Denote by $x_{(0)} \equiv 0$. A partition $\{\Omega_k\}_{k=1}^K$ consists of K equivalent blocks, when $\Omega_k = (x_{(j_k)}, x_{(j_{k+1})}]$, where $j_k = (k-1)c$.*

A variant of this definition can be obtained in terms of interval lengths rather than numbers of observations.

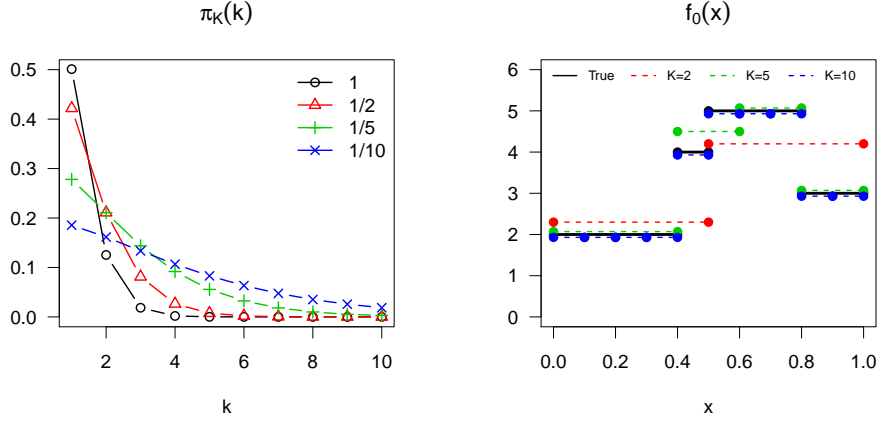


Figure 1: (Left) Prior on the tree size for several values of c_K , (Right) Best approximations of f_0 (in the ℓ_2 sense) by step functions supported on equispaced blocks of size $K \in \{2, 5, 10\}$.

Definition 2.3. (*Equispaced Blocks*) A partition $\{\Omega_k\}_{k=1}^K$ consists of K equispaced blocks Ω_k , when $\Omega_k = (\frac{k-1}{K}, \frac{k}{K}]$ for $k = 1, \dots, K$.

When K is dyadic, the equispaced partition is a full complete binary tree with splits at dyadic rationals. If the observations x_i lie on a regular grid (Assumption 1 to follow), then Definition 2.2 and 2.3 are basically equivalent. We come back to this point in the Section 3. We will focus on equivalent blocks (EB). Because there is only one such partition for each K , the prior $\pi_{\Omega}(\cdot|K)$ has a single point mass, located at Ω_K^{EB} . We will denote by $\Omega^{EB} = \cup_{K=1}^{\infty} \Omega_K^{EB}$ the set of all EB partitions for $K = 1, 2, \dots$. We will use these partitioning schemes as a jump-off point.

2.2.2 Balanced Intervals

Equivalent (equispaced) blocks are deterministic and, as such, do not provide much room for learning about the actual location of jumps in f_0 . Balanced intervals, introduced below, are a richer class of partitions that tolerate a bit more imbalance. First, we introduce the notion of cell counts $\mu(\Omega_k)$. For each interval Ω_k , we write

$$\mu(\Omega_k) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}(x_i \in \Omega_k), \quad (4)$$

the proportion of observations falling inside Ω_k . Note that for equivalent blocks, we can write $\mu(\Omega_1) = \dots = \mu(\Omega_K) = c$.

Definition 2.4. (*Balanced Intervals*) A partition $\{\Omega_k\}_{k=1}^K$ is balanced if

$$\frac{C_{min}^2}{K} \leq \mu(\Omega_k) \leq \frac{C_{max}^2}{K} \quad \text{for all } k = 1, \dots, K \quad (5)$$

for some universal constants $C_{min} \leq 1 \leq C_{max}$ not depending on K .

The following variant of this condition uses interval widths rather than cell counts, where the balance is satisfied if $\tilde{C}_{min}^2/K \leq |\Omega_k| \leq \tilde{C}_{max}^2/K$. Again, under Assumption 1, these two definitions are equivalent. In the sequel, we will denote by Ω_K^{BI} the set of all balanced partitions consisting of K intervals and by $\Omega^{BI} = \cup_{K=1}^{\infty} \Omega_K^{BI}$ the set of all balanced intervals of sizes $K = 1, 2, \dots$.

With balanced partitions, the K^{th} shell \mathcal{F}_K of the approximating space \mathcal{F} in (2) consists of all step functions that are supported on partitions Ω_K^{BI} and have $K - 1$ points of discontinuity $u_k \in I_n \equiv \{x_i : i = 1, \dots, n - 1\}$ for $k = 1, \dots, K - 1$. For equal blocks in Definition 2.2, we assumed that the points of subdivision were *deterministic*, i.e. $u_k = k/K$. For balanced partitions, we assume that u_k are *random* and chosen amongst the observed values x_i . The order statistics of the vector of splits

$\mathbf{u} = (u_1, \dots, u_{K-1})'$ uniquely define a segmentation of $[0, 1]$ into K intervals $\Omega_k = (u_{(k-1)}, u_{(k)}]$, where $u_{(k)}$ designates the k^{th} smallest value in \mathbf{u} and $u_{(0)} \equiv 0, u_{(K)} = x_{(n)} \equiv 1$.

Our prior over balanced intervals $\pi_{\Omega}(\cdot | K)$ will be defined implicitly through a *uniform* prior over the split vectors \mathbf{u} . Namely, the prior over balanced partitions Ω_K^{BI} satisfies

$$\pi_{\Omega}(\{\Omega_k\}_{k=1}^K | K) = \frac{1}{\text{card}(\Omega_K^{BI})} \mathbb{I}(\{\Omega_k\}_{k=1}^K \in \Omega_K^{BI}). \quad (6)$$

In the following Lemma, we obtain upper bounds on $\text{card}(\Omega_K^{BI})$ and discuss how they relate to an old problem in geometric probability.

Lemma 2.1. *Assume that $\mathbf{u} = (u_1, \dots, u_{K-1})'$ is a vector of independent random variables obtained by uniform sampling (without replacement) from I_n . Then we have for $1/n < C < 1/K$*

$$\Pi \left(\min_{1 \leq k \leq K} |\Omega_k| \geq C \right) = \frac{\binom{\lfloor n(1-KC) \rfloor + K - 1}{K-1}}{\binom{n-1}{K-1}} \quad (7)$$

and

$$\Pi \left(\max_{1 \leq k \leq K} |\Omega_k| \leq C \right) = 1 - \sum_{k=1}^{\tilde{n}} (-1)^k \binom{n-1}{k} \frac{\binom{\lfloor n(1-kC) \rfloor + K - 1}{K-1}}{\binom{n-1}{K-1}}, \quad (8)$$

where $\tilde{n} = \min\{n-1, \lfloor 1/C \rfloor\}$.

Proof. In Supplement. □

Remark 2.1. *Flatto and Konheim [29] showed that the probability of covering a circle with random arcs of length C is equal to the probability that all segments of the unit interval, obtained with iid random uniform splits, are smaller than C . Similarly, the probability (8) corresponds to the probability of covering the circle with random arcs whose endpoints lie on a grid of $K-1$ equidistant points on the circumference.*

There are $\binom{n-1}{K-1}$ partitions of size K , of which $\binom{\lfloor n(1-\tilde{C}_{min}^2) \rfloor + K - 1}{K-1}$ satisfy the minimal cell width balancing condition (where $\tilde{C}_{min}^2 > K/n$). This number gives an upper bound on the combinatorial complexity of balanced partitions $\text{card}(\Omega_K^{BI})$.

2.3 Prior $\pi(\beta | K)$ on Step Heights β

To complete the prior on \mathcal{F}^K , we take independent normal priors on each of the coefficients. Namely

$$\pi(\beta | K) = \prod_{k=1}^K \phi(\beta_k), \quad (9)$$

where $\phi(\cdot)$ is a standard normal density.

3 Main Results

Throughout the paper, we assume that the observations x_i are fixed. Partitioning-based regression methods are often invariant to monotone transformations of observations. In particular, when f_0 is a step function, standardizing the distance between the observations, and thereby the split points, has no effect on the nature of the estimation problem. Without loss of generality, we will thereby assume that the observations are aligned on an equispaced grid.

Assumption 1. (*Equispaced Grid*) *We assume that the predictor values satisfy $x_i = \frac{1}{n}(i - \frac{1}{2})$ for each $i = 1, \dots, n$.*

This assumption implies that partitions that are balanced in terms of the Lebesgue measure will be balanced also in terms of the number of observations. A similar assumption was imposed by Donoho [30] in his study of Dyadic CART.

Another crucial ingredient of our proof will be understanding how well f_0 can be approximated by other step functions (supported on partitions Ω , which are either equal/equivalent blocks Ω^{EB} or balanced partitions Ω^{BI}). We will describe the approximation error in terms of the overlap between the true partition $\{\Omega_k^0\}_{k=1}^{K_0}$ and the approximating partitions $\{\Omega_k\}_{k=1}^K \in \Omega$. More formally, we define the *restricted cell count* (according to Nobel [31]) as

$$m\left(V; \{\Omega_k^0\}_{k=1}^{K_0}\right) = |\Omega_k^0 : \Omega_k^0 \cap V \neq \emptyset|,$$

which is the number of cells in $\{\Omega_k^0\}_{k=1}^{K_0}$ that overlap with an interval $V \subset [0, 1]$. Next, we define the *complexity* of f_0 as the smallest size of a partition in Ω needed to completely cover f_0 without any overlap.

Definition 3.1. (*Complexity of f_0 w.r.t. Ω*) We define $K(f_0, \Omega)$ as the smallest K such that there exists a K -partition $\{\Omega_k\}_{k=1}^K$ in the class of partitions Ω for which

$$m\left(\Omega_k; \{\Omega_k^0\}_{k=1}^{K_0}\right) = 1 \quad \text{for all } k = 1, \dots, K.$$

The number $K(f_0, \Omega)$ will be referred to as the *complexity* of f_0 w.r.t. Ω .

The complexity number $K(f_0, \Omega)$ indicates the optimal number of steps needed to approximate f_0 with a step function (supported on partitions in Ω) without any error. It depends on the true number of jumps K_0 as well as the true interval lengths $|\Omega_k^0|$. If the minimal partition $\{\Omega_k^0\}_{k=1}^{K_0}$ resided in the approximating class, i.e. $\{\Omega_k^0\}_{k=1}^{K_0} \in \Omega$, then we would obtain $K(f_0, \Omega) = K_0$, which is the true number of steps. On the other hand, when $\{\Omega_k^0\}_{k=1}^{K_0} \notin \Omega$, the complexity number $K(f_0, \Omega)$ can be much larger. This is illustrated in Figure 1 (right), where the true partition $\{\Omega_k^0\}_{k=1}^{K_0}$ consists of $K_0 = 4$ unequal pieces and we are trying to approximate it with equispaced blocks with $K = 2, 5, 10$ steps. Because the intervals Ω_k^0 are not equal and the smallest one has a length $1/10$, we need $K(f_0, \Omega^{EB}) = 10$ equispaced blocks to perfectly approximate f_0 . For our analysis, we do not need to assume that $\{\Omega_k^0\}_{k=1}^{K_0} \in \Omega$. We only have to assume that $K(f_0, \Omega)$ is finite, i.e. it can be approximated without any error with a member of Ω with finite number of discontinuities. In essence, we assume that f_0 can be represented as an element of one of the approximating classes.

Assumption 2. Let Ω be a class of approximating partitions. We assume that f_0 is Ω -valid, i.e. $K(f_0, \Omega) < \infty$.¹

Another way to view $K(f_0, \Omega)$ is as the ideal partition size on which the posterior should concentrate. If this number were known, we could achieve the minimax posterior concentration rate (Theorem 5.1 in Supplemental materials). Recall that the minimax rate for estimating a univariate step function is $n^{-1/2}$ [32]. This minimax performance is attainable when K_0 is known. For our main results, however, we will adapt to the unknown K_0 sacrificing a log-factor in the rate.

3.1 Posterior Concentration for Equivalent Blocks

Our first result shows that the minimax rate is achieved, up to a $\sqrt{\log n}$ factor, without any assumptions on the number of pieces of f_0 or the sizes of the pieces.

Theorem 3.1. (*Equivalent Blocks*) Let $f_0 : [0, 1] \rightarrow \mathbb{R}$ be a step function with K_0 steps, where K_0 is unknown. Denote by \mathcal{F} the set of all step functions supported on equivalent blocks, equipped with priors $\pi_K(\cdot)$ and $\pi(\beta | K)$ as in (3) and (9). Assume that f_0 is Ω^{EB} -valid. Then, under Assumption 1 we have

$$\mathbb{P}\left(f \in \mathcal{F} : \|f - f_0\|_n \geq M_n n^{-1/2} \sqrt{\log n} \mid \mathbf{Y}^{(n)}\right) \rightarrow 0 \quad (10)$$

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

Before we proceed with the proof, a few remarks ought to be made. First, it is worthwhile to emphasize that the statement in Theorem 3.1 is a frequentist one, as it relates to an aggregated behavior of the posterior distributions obtained under the true generative model $P_{f_0}^n$.

¹An example of functions f_0 that are *not* Ω -valid are those f_0 supported on a number of intervals that increases with n .

Second, the theorem shows that the Bayesian procedure performs an automatic adaptation to $K(f_0, \Omega^{EB})$. The posterior will concentrate on EB partitions that are fine enough to approximate f_0 well. Thus, we are able to recover the true function almost as well as if we knew K_0 . The extra log-factor is the price we pay for learning $K(f_0, \Omega^{EB})$ from the data.

Third, it is worth mentioning that Theorem 3.1 holds for equivalent as well as equisized blocks, because of Assumption 1. It describes posterior concentration for *dyadic regression trees*. Indeed, with dyadic $K = 2^s$ for some $s \in \mathbb{N}$, the equisized partition corresponds to a full binary tree with splits at dyadic rationals.

Another interesting insight is that the Gaussian prior (9), while selected for mathematical convenience, turns out to be sufficient for recovery. In other words, despite the relatively large amount of mass of near zero, the Gaussian prior does not rule out optimal posterior concentration. Our standard normal prior is a simpler version of the Bayesian CART prior, which determines the variance from the data [9].

The Theorem 3.1 is proved by verifying the three conditions of Theorem 4 of [18], for $\varepsilon_n = n^{-1/2} \sqrt{\log n}$ and $\mathcal{F}_n = \bigcup_{K=0}^{k_n} \mathcal{F}_K$, with k_n of the order $\log n$. The approximating subspace $\mathcal{F}_n \subset \mathcal{F}$ should be rich enough to approximate f_0 well and it should receive most of the prior mass. The conditions for posterior contraction at the rate ε_n are:

$$\begin{aligned} \text{(C1)} \quad & \sup_{\varepsilon > \varepsilon_n} \log N\left(\frac{\varepsilon}{36}, \{f \in \mathcal{F}_n : \|f - f_0\|_n < \varepsilon\}, \|\cdot\|_n\right) \leq n\varepsilon_n^2, \\ \text{(C2)} \quad & \frac{\Pi(\mathcal{F} \setminus \mathcal{F}_n)}{\Pi(f \in \mathcal{F} : \|f - f_0\|_n^2 \leq \varepsilon_n^2)} = o(e^{-2n\varepsilon_n^2}), \\ \text{(C3)} \quad & \frac{\Pi(f \in \mathcal{F}_n : j\varepsilon_n < \|f - f_0\|_n \leq 2j\varepsilon_n)}{\Pi(f \in \mathcal{F} : \|f - f_0\|_n^2 \leq \varepsilon_n^2)} \leq e^{\frac{j^2}{4} n\varepsilon_n^2}, \text{ for all sufficiently large } j. \end{aligned}$$

The entropy condition (C1) restricts attention to EB partitions with small K . As will be seen from the proof, the largest allowed partitions have at most (a constant multiple of) $\log n$ pieces. For large enough n , such partitions will perfectly approximate f_0 because $K(f_0, \Omega^{EB}) < \infty$. Condition (C2) requires that the prior does not promote partitions with more than $\log n$ pieces. This property is guaranteed by the exponentially decaying prior $\pi_K(\cdot)$, which penalizes large partitions. While the result holds regardless of the choice of the tuning constant $c_K > 0$ in (3), we conjecture that values $c_K \geq 1/K(f_0, \Omega^{EB})$ are optimal in terms of the concentration speed, and we recommend a default choice of $c_K = 1$.

The final condition, (C3), requires that the prior charges a $\|\cdot\|_n$ neighborhood of the true function. In our proof, we verify this condition by showing that the prior mass on step functions of the optimal size $K(f_0, \Omega^{EB})$ is sufficiently large.

Proof. We verify the three conditions (C1), (C2) and (C3).

(C1) Let $\varepsilon > \varepsilon_n$ and $K \in \mathbb{N}$. For $f_\alpha, f_\beta \in \mathcal{F}_K$, we have $K^{-1} \|\alpha - \beta\|_2^2 = \|f_\alpha - f_\beta\|_n^2$ because $\mu(\Omega_k) = 1/K$ for each k . We now argue as in the proof of Theorem 12 of [18] to show that $N\left(\frac{\varepsilon}{36}, \{f \in \mathcal{F}_K : \|f - f_0\|_n < \varepsilon\}, \|\cdot\|_n\right)$ can be covered by the number of $\sqrt{K}\varepsilon/36$ -balls required to cover a $\sqrt{K}\varepsilon$ -ball in \mathbb{R}^K . This number is bounded above by 108^K . Summing over K , we recognize a geometric series. Taking the logarithm of the result, we find that (C1) is satisfied if $\log(108)(k_n + 1) \leq n\varepsilon_n^2$.

(C2) Let $K(f_0) \equiv K(f_0, \Omega^{EB})$ be as in Definition 3.1. We bound the denominator by:

$$\Pi(f \in \mathcal{F} : \|f - f_0\|_n^2 \leq \varepsilon^2) \geq \pi_K(K_0) \Pi\left(\beta \in \mathbb{R}^{K(f_0)} : \|\beta - \beta_0^{\text{ext}}\|_2^2 \leq \varepsilon^2 K(f_0)\right),$$

where $\beta_0^{\text{ext}} \in \mathbb{R}^{K(f_0)}$ is an extended version of $\beta_0 \in \mathbb{R}^{K_0}$, containing the coefficients for f_0 expressed as a step function on the partition $\{\Omega_k^0\}_{k=1}^{K(f_0)}$. Arguing as in the proof of Theorem 12 of [18], this can be bounded from below by

$$\frac{\pi_K(K(f_0))}{2^{K(f_0)} e^{\|\beta_0^{\text{ext}}\|_2^2} \Gamma(K(f_0)/2)} \int_0^{\varepsilon^2 K(f_0)/2} x^{K(f_0)/2-1} e^{-x/2} dx.$$

We bound this from below by bounding the exponential at the upper integration limit, yielding:

$$\frac{\pi_K(K(f_0))K(f_0)^{K(f_0)/2-1}}{2^{3K(f_0)/2-1}e^{\|\beta_0^{\text{ext}}\|_2^2}\Gamma(K(f_0)/2)}e^{-\varepsilon^2 K(f_0)/4}e^{K(f_0)}. \quad (11)$$

For $\varepsilon = \varepsilon_n \rightarrow 0$, we thus find that the denominator in (C2) is of order $\varepsilon^{K(f_0)}$. We bound the numerator:

$$\Pi(\mathcal{F} \setminus \mathcal{F}_n) = \Pi\left(\bigcup_{k=k_n+1}^{\infty} \mathcal{F}_k\right) \propto \sum_{k=k_n+1}^{\infty} e^{-c_K k \log k} \leq e^{-c_K(k_n+1) \log(k_n+1)} + \int_{k_n+1}^{\infty} e^{-c_K x \log x},$$

which is of order $e^{c_K(k_n+1) \log(k_n+1)}$. Combining this bound with (11), we find that (C2) is met if:

$$e^{-K(f_0) \log \varepsilon_n - c_K(k_n+1) \log(k_n+1) + 2n\varepsilon_n^2} \rightarrow 0 \text{ as } n \rightarrow \infty.$$

(C3) We bound the numerator by one, and use the bound (11) for the denominator. As K_0 is constant and $\varepsilon_n \rightarrow 0$, we obtain the condition $e^{-K(f_0) \log \varepsilon_n} \leq e^{\frac{j^2}{4}n\varepsilon_n^2}$ for sufficiently large j .

Conclusion (C1) holds if $\log(108)(k_n + 1) \leq n\varepsilon_n^2$, (C2) holds if $K(f_0) \log \varepsilon_n + c_K k_n \log k_n - 2n\varepsilon_n^2 \rightarrow \infty$, and (C3) holds if there exists a $j \in \mathbb{N}$ such that $-K(f_0) \log \varepsilon_n \leq \frac{j^2}{4}n\varepsilon_n^2$.

With $\varepsilon_n = n^{-1/2}\sqrt{\log n}$, let k_n be the largest integer value such that $\log(108)(k_n + 1) \leq n\varepsilon_n^2$, thus meeting (C1). Then k_n is of order $\log n$ and thus (C2) holds as well. (C3) is met for $j > \sqrt{2K(f_0)}$. \square

3.2 Posterior Concentration for Balanced Intervals

An analogue of Theorem 3.1 can be obtained for balanced partitions from Section 2.2.2 that correspond to regression trees with splits at actual observations. Now, we assume that f_0 is Ω^{BI} -valid and carry out the proof with $K(f_0, \Omega^{BI})$ instead of $K(f_0, \Omega^{EB})$. We replace the prior (3) by a bit sharper prior, because balanced intervals have more flexibility to capture f_0 with fewer pieces. The posterior concentration rate is only slightly worse.

Theorem 3.2. (Balanced Intervals) Let $f_0 : [0, 1] \rightarrow \mathbb{R}$ be a step function with K_0 steps, where K_0 is unknown. Denote by \mathcal{F} the set of all step functions supported on balanced intervals, equipped with priors $\pi_K(k) \propto e^{-c_K k \log n}$ and $\pi_{\Omega}(\cdot | K), \pi(\beta | K)$ as in (6) and (9). Assume that f_0 is Ω^{BI} -valid. Then, under Assumption 1 we have

$$\Pi\left(f \in \mathcal{F} : \|f - f_0\|_n \geq M_n n^{-1/2} (\log n)^\beta | \mathbf{Y}^{(n)}\right) \rightarrow 0 \quad (12)$$

in $P_{f_0}^n$ -probability, for every $M_n \rightarrow \infty$ as $n \rightarrow \infty$, where $\beta > 1/2$.

Proof. All three conditions (C1), (C2) and (C3) hold if we choose $k_n \sim (\log n)^{2\beta-1}$. The entropy condition will be satisfied when $\sum_{k=1}^{k_n} C^k \text{card}(\Omega_k^{BI}) \preceq n\varepsilon_n^2$ for some $C > 0$, where $\varepsilon_n = n^{-1/2}(\log n)^\beta$. Using the upper bound $\text{card}(\Omega_k^{BI}) < \binom{n-1}{k-1}$ and because $k_n < \frac{n-1}{2}$ for large enough n , the entropy condition (C1) will be satisfied. The condition (C2) will be satisfied when

$$e^{-K(f_0) \log \varepsilon_n + \log \text{card}(\Omega_{K(f_0)}^{BI}) - c_K(k_n+1) \log n + 2n\varepsilon_n^2} \rightarrow 0 \text{ as } n \rightarrow \infty. \quad (13)$$

Because $K(f_0) < \infty$ we have $\log \text{card}(\Omega_{K(f_0)}^{BI}) < \log \text{card}(\Omega_{k_n}^{BI}) \preceq k_n \log(n/k_n)$. Because $k_n \log n$ is the leading term, (13) holds for $c_K > 3$. Condition (C3) is also satisfied for $k_n \sim (\log n)^{2\beta-1}$. \square

4 Discussion

We provided the first posterior concentration rate results for Bayesian non-parametric regression with step functions. We showed that under suitable complexity priors, the Bayesian procedure adapts to the unknown aspects of the target step function. Our approach can be naturally extended in three ways: (a) to smooth f_0 functions, (b) to dimension reduction with many predictors, (c) to more general partitioning schemes that correspond to methods like Bayesian CART and BART. These three extensions are discussed in our followup manuscript.




Figure 2: Two sets E_K of possible stick lengths that satisfy the minimal cell-size condition $|\Omega_k| \geq C$ with $n = 10$, $C = 2/n$ and $K = 2, 3$.

5 Appendix

5.1 Proof of Lemma 2.1

The denominator of (7) follows from the fact that there are $n - 1$ possible splits for the $K - 1$ points of discontinuity u_k . The numerator is obtained after adapting the proof of Lemma 2 of Flatto and Konheim [29], who showed a related probability statement for uniformly distributed jumps on the unit interval. Here, the split points are supported on a grid and thereby a modification of their lemma is needed.

Without loss of generality, we will assume that $C = a/n$ for some $a = 1, \dots, \lfloor n/K \rfloor$ so that $n(1 - KC)$ is an integer. Recall that $|\Omega_k|$ are the lengths of the segments defined through the split points \mathbf{u} . Because the jumps u_k can only occur on the grid I_n , we have $|\Omega_k| = j/n$ for some $j = 1, \dots, n - 1$. It follows from Lemma 1 of Flatto and Konheim [29] that the set $E_K = \{|\Omega_k| : \sum_{k=1}^K |\Omega_k| = 1 \text{ and } |\Omega_k| \geq C \text{ for } k = 1, \dots, K\}$ lies in the interior of a convex hull of K points $v_r = (1 - KC)e_r + C \sum_{k=1}^K e_k$ for $r = 1, \dots, K$, where $e_r = (e_{r1}, \dots, e_{rK})'$ are unit base vectors, i.e. $e_{rj} = \mathbb{I}(r = j)$. Two examples of the set E_K (for $K = 2$ and $K = 3$) are depicted in Figure 2. In both figures, $n = 10$ (i.e. 9 candidate split points) and $a = 2$. With $K = 2$ (Figure 2(a)), there are only $7 = \binom{n(1-KC)+K-1}{K-1}$ pairs of interval lengths $(|\Omega_1|, |\Omega_2|)'$ that satisfy the minimal cell condition. These points lie on a grid between the two vertices $v_1 = (1 - C, C)$ and $v_2 = (C, 1 - C)$. With $K = 3$, the convex hull of points $v_1 = (1 - 2C, C, C)'$, $v_2 = (C, 1 - 2C, C)'$ and $v_3 = (C, C, 1 - 2C)'$ corresponds to a diagonal dissection of a cube of a side length $(1 - 3C)$ (Figure 2(b), again with $a = 2$ and $n = 10$). The number of lattice points in the interior (and on the boundary) of such tetrahedron corresponds to an arithmetic sum $\frac{1}{2}(n - 3a + 2)(n - 3a + 1) = \binom{n-3a+2}{2}$.

So far, we showed (7) for $K = 2$ and $K = 3$. To complete the induction argument, suppose that the formula holds for some arbitrary $K - 1 > 0$. Then the size of the lattice inside (and on the boundary) of a K -tetrahedron of a side length $[1 - (K + 1)C]$ can be obtained by summing lattice sizes inside $(K - 1)$ -tetrahedrons of increasing side lengths $1, 2, \dots, [1 - (K + 1)C]$, i.e.

$$\sum_{j=K-1}^{n[1-(K+1)C]+K-1} \binom{j}{K-1} = \binom{n[1-(K+1)C]+K}{K},$$

where we used the fact $\sum_{j=K}^N \binom{j}{K} = \binom{N+1}{K+1}$. The second statement (8) is obtained by writing the event as a complement of the union of events and applying the method of inclusion-exclusion.

5.2 The case of known K_0

Theorem 5.1. *Let $f_0 : [0, 1] \rightarrow \mathbb{R}$ be a step function satisfying Assumption 2 with known $K(f_0)$. Denote by $\mathcal{F}_{K(f_0)}$ the set of all step functions supported on the equispaced partition consisting of $K(f_0)$ pieces, equipped with prior $\pi(\beta \mid K(f_0))$ as in (9). Then*

$$\Pi(f \in \mathcal{F}_{K(f_0)} : \|f - f_0\|_n \geq M_n n^{-1/2} \mid Y^n) \rightarrow 0$$

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

Proof. The proof is a special case of the proof of Theorem 3.1. □

References

- [1] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification and Regression Trees*. Statistics/Probability Series. Wadsworth Publishing Company, Belmont, California, U.S.A., 1984.
- [2] L. Breiman. Random forests. *Mach. Learn.*, 45:5–32, 2001.
- [3] A. Berchuck, E. S. Iversen, J. M. Lancaster, J. Pittman, J. Luo, P. Lee, S. Murphy, H. K. Dressman, P. G. Febbo, M. West, J. R. Nevins, and J. R. Marks. Patterns of gene expression that characterize long-term survival in advanced stage serous ovarian cancers. *Clin. Cancer Res.*, 11(10):3686–3696, 2005.
- [4] S. Abu-Nimeh, D. Nappa, X. Wang, and S. Nair. A comparison of machine learning techniques for phishing detection. In *Proceedings of the Anti-phishing Working Groups 2nd Annual eCrime Researchers Summit*, eCrime '07, pages 60–69, New York, NY, USA, 2007. ACM.
- [5] M. A. Razi and K. Athappilly. A comparative predictive analysis of neural networks (NNs), nonlinear regression and classification and regression tree (CART) models. *Expert Syst. Appl.*, 29(1):65 – 74, 2005.
- [6] D. P. Green and J. L. Kern. Modeling heterogeneous treatment effects in survey experiments with Bayesian Additive Regression Trees. *Public Opin. Q.*, 76(3):491, 2012.
- [7] E. C. Polly and M. J. van der Laan. Super learner in prediction. Available at: http://works.bepress.com/mark_van_der_laan/200/, 2010.
- [8] D. M. Roy and Y. W. Teh. The Mondrian process. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 1377–1384. Curran Associates, Inc., 2009.
- [9] H. A. Chipman, E. I. George, and R. E. McCulloch. Bayesian CART model search. *JASA*, 93(443):935–948, 1998.
- [10] D. Denison, B. Mallick, and A. Smith. A Bayesian CART algorithm. *Biometrika*, 95(2):363–377, 1998.
- [11] H. A. Chipman, E. I. George, and R. E. McCulloch. BART: Bayesian additive regression trees. *Ann. Appl. Stat.*, 4(1):266–298, 03 2010.
- [12] S. Ghosal, J. K. Ghosh, and A. W. van der Vaart. Convergence rates of posterior distributions. *Ann. Statist.*, 28(2):500–531, 04 2000.
- [13] T. Zhang. Learning bounds for a generalized family of Bayesian posterior distributions. In S. Thrun, L. K. Saul, and P. B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, pages 1149–1156. MIT Press, 2004.
- [14] J. Tang, Z. Meng, X. Nguyen, Q. Mei, and M. Zhang. Understanding the limiting factors of topic modeling via posterior contraction analysis. In T. Jebara and E. P. Xing, editors, *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 190–198. JMLR Workshop and Conference Proceedings, 2014.
- [15] N. Korda, E. Kaufmann, and R. Munos. Thompson sampling for 1-dimensional exponential family bandits. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 1448–1456. Curran Associates, Inc., 2013.

- [16] F.-X. Briol, C. Oates, M. Girolami, and M. A. Osborne. Frank-Wolfe Bayesian quadrature: Probabilistic integration with theoretical guarantees. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 1162–1170. Curran Associates, Inc., 2015.
- [17] M. Chen, C. Gao, and H. Zhao. Posterior contraction rates of the phylogenetic indian buffet processes. *Bayesian Anal.*, 11(2):477–497, 06 2016.
- [18] S. Ghosal and A. van der Vaart. Convergence rates of posterior distributions for noniid observations. *Ann. Statist.*, 35(1):192–223, 02 2007.
- [19] B. Szabó, A. W. van der Vaart, and J. H. van Zanten. Frequentist coverage of adaptive nonparametric Bayesian credible sets. *Ann. Statist.*, 43(4):1391–1428, 08 2015.
- [20] I. Castillo and R. Nickl. On the Bernstein von Mises phenomenon for nonparametric Bayes procedures. *Ann. Statist.*, 42(5):1941–1969, 2014.
- [21] J. Rousseau and B. Szabo. Asymptotic frequentist coverage properties of Bayesian credible sets for sieve priors in general settings. *ArXiv e-prints*, September 2016.
- [22] I. Castillo. Polya tree posterior distributions on densities. *preprint available at <http://www.lpma-paris.fr/pageperso/castillo/polya.pdf>*, 2016.
- [23] L. Liu and W. H. Wong. Multivariate density estimation via adaptive partitioning (ii): posterior concentration. arXiv:1508.04812v1, 2015.
- [24] C. Scricciolo. On rates of convergence for Bayesian density estimation. *Scand. J. Stat.*, 34(3):626–642, 2007.
- [25] M. Coram and S. Lalley. Consistency of Bayes estimators of a binary regression function. *Ann. Statist.*, 34(3):1233–1269, 2006.
- [26] L. Shepp. Covering the circle with random arcs. *Israel J. Math.*, 34(11):328–345, 1972.
- [27] W. Feller. *An Introduction to Probability Theory and Its Applications, Vol. 2, 3rd Edition*. Wiley, 3rd edition, January 1968.
- [28] T. Anderson. Some nonparametric multivariate procedures based on statistically equivalent blocks. In P.R. Krishnaiah, editor, *Multivariate Analysis*, pages 5–27. Academic Press, New York, 1966.
- [29] L. Flatto and A. Konheim. The random division of an interval and the random covering of a circle. *SIAM Rev.*, 4:211–222, 1962.
- [30] D. L. Donoho. CART and best-ortho-basis: a connection. *Ann. Statist.*, 25(5):1870–1911, 10 1997.
- [31] A. Nobel. Histogram regression estimation using data-dependent partitions. *Ann. Statist.*, 24(3):1084–1105, 1996.
- [32] C. J. Stone. Optimal global rates of convergence for nonparametric regression. *Ann. Statist.*, 10(4):1040–1053, 1982.