

# Boulevard: Regularized Stochastic Gradient Boosted Trees and Their Limiting Distribution

**Yichen Zhou** \*

*Department of Statistics and Data Science  
Cornell University  
Ithaca, NY 14853, USA*

YZ793@CORNELL.EDU

**Giles Hooker**

*Department of Statistics  
University of California, Berkeley  
Berkeley, CA 94720, USA*

GHOOKER@BERKELEY.EDU

**Editor:** Mehryar Mohri

## Abstract

This paper examines a novel gradient boosting framework for regression. We regularize gradient boosted trees by introducing subsampling and employ a modified shrinkage algorithm so that at every boosting stage the estimate is given by an average of trees. The resulting algorithm, titled “Boulevard”, is shown to converge as the number of trees grows. This construction allows us to demonstrate a central limit theorem for this limit, providing a characterization of uncertainty for predictions. A simulation study and real world examples provide support for both the predictive accuracy of the model and its limiting behavior.

**Keywords:** gradient boosting, regression tree, regularization, limiting distribution

## 1. Introduction

This paper presents a theoretical study deriving inferential results for regularized gradient boosted decision trees (GBM, GBDT or GBT: Friedman, 2001). Machine learning methods for prediction have generally been thought of as trading off both intelligibility and statistical uncertainty quantification in favor of accuracy. Recent results have started to provide a statistical understanding of methods based on ensembles of decision trees (Breiman et al., 1984). In particular, the consistency of methods related to random forests (RF: Breiman, 2001) has been demonstrated in Biau (2012); Scornet et al. (2015) while Wager et al. (2014); Mentch and Hooker (2016); Wager and Athey (2017) and Athey et al. (2016) prove central limit theorems for RF predictions. These have then been used for tests of variable importance and nonparametric interactions in Mentch and Hooker (2017).

In this paper, we extend this analysis to GBDT. Analyses of RF have relied on a subsampling structure to express the estimator in the form of a U-statistic from which central limit theorems can be derived. By contrast, GBDT produces trees sequentially with the current tree depending on the values in those built previously, requiring a different analytical approach. While the algorithm proposed in Friedman (2001) is intended to be generally

---

\*. Currently working at Google Research: yichenzhou@google.com

applicable to any loss function, in this paper we focus specifically on squared-error regression (Stone, 1977, 1982). Given a sample of  $n$  observations  $(x_1, y_1), \dots, (x_n, y_n) \in [0, 1]^d \times \mathbb{R}$ , assume they follow the relation

$$X \sim \mu, \quad Y = f(X) + \epsilon$$

which satisfies the following:

- (M1) The density,  $\mu$ , is bounded from above and below, i.e.  $\exists 0 < c_1 < c_2$  s.t.  $c_1 \leq \mu \leq c_2$ .
- (M2)  $f$  is bounded Lipschitz, i.e.  $\exists M_f > 0$  s.t.  $|f(x)| \leq M_f < \infty$ , and  $\exists \alpha > 0$  s.t.  $|f(x_1) - f(x_2)| \leq \alpha \|x_1 - x_2\|, \forall x_1, x_2 \in [0, 1]^d$ .
- (M3)  $\epsilon$  is sub-Gaussian error with  $\mathbb{E}[\epsilon] = 0, \mathbb{E}[\epsilon^2] = \sigma_\epsilon^2, \mathbb{E}[\epsilon^4] < \infty$ .

To fit this sample, GBDT builds correlated trees sequentially to perform gradient descent in functional space (Friedman et al., 2000). With  $L^2$  loss, the process resembles the Robbins-Monro algorithm (Robbins and Monro, 1951), applying trees to iteratively fit residuals. The procedure is given as follows.

- Start with the initial estimate  $\hat{f}_0 = 0$ .
- For  $b \geq 1$ , given the current functional estimate  $\hat{f}_b$ , calculate the negative gradient

$$z_i \triangleq -\frac{\partial}{\partial u_i} \sum_{j=1}^n \frac{1}{2} (u_j - y_j)^2 \Big|_{u_j = \hat{f}_b(x_j)} = y_i - \hat{f}_b(x_i).$$

- Construct a tree regressor  $t_b(\cdot)$  on  $(x_1, z_1), \dots, (x_n, z_n)$ .
- Update by a small learning rate  $\lambda > 0$ ,

$$\hat{f}_{b+1} = \hat{f}_b + \lambda t_b.$$

Gradient boosting is initially developed from attempts to understand Adaboost (Freund et al., 1999) in Friedman et al. (2000). Mallat and Zhang (1993) studied the Robbins-Monro algorithm and demonstrated convergence when the regressors are taken from a Hilbert space. Their result was extended by Bühlmann (2002) to show the consistency of decision tree boosting under the settings of  $L^2$  norm and early stopping. From a broader point of view, discussions of consistency and convergence of general  $L^2$  boosting framework can be found in Bühlmann and Yu (2003), Zhang et al. (2005) and Bühlmann and Hothorn (2007).

Besides the original implementation, there are a number of variations on the algorithm presented above. Friedman (2002) incorporated subsampling in each iteration and empirically showed significant improvement in predictive accuracy. Rashmi and Gilad-Bachrach (2015) argued that GBDT is sensitive towards the initial trees, requiring lots of subsequent trees to make an impact. Their solution incorporates the idea of dropout (Wager et al., 2013; Srivastava et al., 2014) to train each new iteration with a subset of the existing ensemble. Further, Rogozhnikov and Likhomanenko (2017) suggested to downscale the learning rate and studied the convergence of the boosting path when the learning rate is small enough

to guarantee contraction. We note that these analyzes differ from the analysis of the convergence of Adaboost (e.g. Bartlett et al., 1998; Rudin et al., 2004; Mukherjee et al., 2013) which focus on examining a classification margin.

In this paper we intend to provide a universal framework combining the factors mentioned above into GBDT to study its asymptotic behavior. Our method is particularly inspired by the recent development of the RF inferential framework (Mentch and Hooker, 2016; Wager and Athey, 2017; Mentch and Hooker, 2017), in which the averaging structure of random forests contributes to asymptotic normality based on theories of U-statistics. Similarly, among classic stochastic gradient methods, Ruppert-Polyak (Polyak and Juditsky, 1992; Ruppert, 1988) averaging implies asymptotic normality for model parameter estimators by averaging the gradient descent history. We therefore regularize the boosting scheme towards this averaging structure, naming this algorithm Boulevard boosting and summarizing it in the following Algorithm 1.

**Algorithm 1 (Boulevard)**

- Start with the initial estimate  $\hat{f}_0 = 0$ .
- Given the current functional estimate  $\hat{f}_b$ , calculate the current residuals<sup>1</sup>

$$z_i = y_i - \hat{f}_b(x_i). \tag{1}$$

- If needed, generate a subsample  $w \subseteq \{1, 2, \dots, n\}$  (otherwise let  $w = \{1, 2, \dots, n\}$ ).
- Construct a tree regressor  $t_b(\cdot)$  on  $\{(x_i, z_i), i \in w\}$ .
- Update by learning rate  $1 > \lambda > 0$ , presumably closer to 1, that

$$\hat{f}_{b+1} = \frac{b-1}{b} \hat{f}_b + \frac{\lambda}{b} t_b = \frac{\lambda}{b} \sum_{i=1}^b t_i. \tag{2}$$

- When boosting ends after the  $b^*$ -th iteration with  $\hat{f}_{b^*}$ , rescale it by  $\frac{1+\lambda}{\lambda}$  to obtain the final prediction  $\frac{1+\lambda}{\lambda} \hat{f}_{b^*}$ .

The name Boulevard re-imagines the random forest stretched out along a row, with the averaging structure diminishing the influence of early trees as the algorithm proceeds. Boulevard is designed to leverage some of the advantages of GBDT while still allowing us to derive the inferential results that enable statistical inference. As indicated by the last two steps, they differ in how they update the ensemble, and Boulevard regularizes the boosting path by de-emphasizing any single tree in the ensemble. Introducing the  $(b-1)/b$  term not only makes the ensemble less sensitive to its early members, but more importantly also guarantees there is always signal in the gradients. As a result, the final form of the predictor sits between an ordinary GBDT and a RF doing averaging.

We will show in the rest of the paper that, following such designs, Boulevard provides us with theoretical guarantees regarding its limiting behavior. Writing  $\hat{f}_{b,n}$  to be the functional estimate after  $b$  iterations for a sample of size  $n$ , we obtain the following results in their simplified forms.

---

1. These can be generalized to the negative gradient of the loss  $z_i \triangleq -\frac{\partial}{\partial u} \sum_{j=1}^n \frac{1}{2} (u - y_j)^2 \Big|_{u=\hat{f}_b(x_j)}$ .

1. Finite sample convergence. Conditioned on any finite sample of size  $n$ , should we keep boosting using Boulevard, the boosting path converges almost surely. In other words,  $\exists \hat{f}_{\infty,n} : \mathbb{R}^d \rightarrow \mathbb{R}$  relying on the sample such that for *a.s.*  $x$ ,

$$\hat{f}_{b,n}(x) \rightarrow \hat{f}_{\infty,n}(x), b \rightarrow \infty.$$

2. Consistency. Under several conditions regarding the construction of regression trees, this limit aligns with the truth after rescaling by  $\frac{1+\lambda}{\lambda}$ . For *a.s.*  $x$ ,

$$\hat{f}_{\infty,n}(x) \xrightarrow{P} \frac{\lambda}{1+\lambda} f(x), n \rightarrow \infty,$$

where the convergence in probability is with respect to sample variability.

3. Asymptotical normality. Under the same conditions as for consistency, we can prove that the prediction is asymptotically normal. For *a.s.*  $x$ ,

$$\frac{\hat{f}_{\infty,n}(x) - \frac{\lambda}{1+\lambda} f(x)}{\text{sd}(\hat{f}_{\infty,n}(x))} \xrightarrow{d} N(0, 1), n \rightarrow \infty.$$

We will demonstrate in detail these results along with their required conditions in later sections. So far as we are aware, these represent the first results on a distributional limit for GBDT and hence the potential for inference using this framework; we hope that they inspire further refinements. It is worth noticing that Bayesian Additive Regression Trees (BART) Chipman et al. (2010) were also motivated by GBDT and allow the development of Bayesian credible intervals. However, the training procedure for BART resembles backfitting a finite number of trees, resulting in a somewhat different model class.

The remainder of the paper is organized as follows: In Section 2, we introduce a counterpart of honesty for boosted trees named structure-value isolation and non-adaptivity to regularizes the boosted tree behavior. In Section 3, we show finite sample Boulevard convergence to a fixed point. In Section 4, we further prove the limiting distribution. In Section 5, we focus on non-adaptivity and discuss how to achieve it either manually or spontaneously. In Section 6, we present our empirical study.

## 2. Non-adaptivity

The sequential process employed in constructing decision trees renders them challenging to describe mathematically and more tractable variants of the original implementation in Breiman et al. (1984) are often employed. In particular, separating the structure of the tree from the observed responses – which may still be used to determine leaf values – has been particularly useful. Early analyses employed completely randomized trees in which the structure was determined independently of observed responses: see Bühlmann et al. (2002); Breiman (2004); Biau et al. (2008); Biau (2012); Arlot and Genuer (2014) that, in particular, generate a connection to kernel methods observed in Davies and Ghahramani (2014) and Scornet (2016). See Section 3.1 of Biau and Scornet (2016) for an overview of these developments.

More recent results in Wager and Athey (2017) employ sample splitting so that the structure of a tree is determined by part of the data while the values in the leaves are decided by an independent set to create a property that is termed honesty – independence between tree structure and leaf values. Wager and Athey (2017) provide some evidence for the practical, as well as theoretical, utility of this construction.

In this paper, we rely on two restrictions on the tree building process: an extension of honesty that we call structure-value isolation to require the structure of trees independent of leaf values for the whole ensemble rather than individual trees, and non-adaptivity that requires the distribution of tree structures (conditional on the sample) to be constant over boosting iterations. Non-adaptivity allows us to demonstrate the convergence of the Boulevard algorithm for a fixed sample and to show that its limit has the form of kernel ridge regression. Structure-value isolation then allows us to provide distributional results by separating the form of the kernel from the response.

We recognize that these place restrictive conditions on the tree building process and note parallels between these and the early developments in the theory of random forests referenced above. Completely randomized trees satisfy both conditions but are not strictly necessary. Non-adaptivity can be replaced by an asymptotic version which we argue may arise spontaneously or can be enforced in various ways. Structure-value isolation can be achieved, for example, by global subsample splitting – replacing the leaf values of trees by a held-out subset as Boulevard progresses, although note that our results do not cover uncertainty in the resulting kernel structure. Further details of these conditions are pursued below; we speculate in Section 4.6 that our distributional results may also be obtained via alternative representations, but developing these is beyond the scope of this paper.

## 2.1 Structure-Value Isolation

A decision tree (Breiman et al., 1984) predicts by iteratively segmenting the covariate space into disjoint subsets (leaves) within each of which a terminal leaf value is chosen for making predictions. Therefore we can separate the construction of a decision tree into two stages: deciding the tree structure, and deciding the values in the terminal nodes. Traditional greedy tree building algorithms use the same sample points for both of these two steps. One drawback of these algorithms is the difficulty of providing mathematical guarantees about isolating sample points with large observation errors, i.e. outliers, thereby de-stabilizing the resulting predicted values. We think of this behavior as “chasing order statistics”. As a result, a plethora of conclusions about trees and tree ensembles rely on randomization, for example, using completely randomized trees or applying honesty as noted above.

In the context of Boulevard, the sequential dependence of trees requires us to employ a stronger condition. We define structure-value isolation as the requirement that leaf values are independent of tree structures across the entire ensemble, rather than on a tree-by-tree basis. To make this precise, here we introduce our decision tree notations. A regression tree  $t_n(\cdot)$  segments the covariate space of our interest  $[0, 1]^d$  into disjoint hyper-rectangular leaves  $[0, 1]^d = \bigsqcup_{j=1}^m A_j$ . Its prediction can be explicitly expressed as

$$t_n(x) = \sum_{i=1}^n s_{n,i}(x)y_i,$$

where, given  $x \in A_j$  for some  $j$ ,

$$s_{n,k}(x) = \frac{I(x_k \in A_j)}{\sum_{i=1}^n I(x_i \in A_j)},$$

describing the influence of  $x_k$  on predicting the value at  $x$  with  $I(\cdot)$  the indicator function; when used without parentheses  $I$  will denote the identity matrix. Slight changes are required when a subsample is used instead of the full sample to calculate the leaf values. For a subsample  $w \subset \{1, \dots, n\}$ , we write

$$t_n(x; w) = \sum_{i=1}^n s_{n,i}(x; w)y_i.$$

In this case, for any  $x \in A_j$ ,

$$s_{n,k}(x) = s_{n,k}(x; w) = \frac{I(x_k \in A_j)I(k \in w)}{\sum_{i=1}^n I(x_i \in A_j)I(i \in w)} = \frac{I(x_k \in A_j)I(k \in w)}{\sum_{x_i \in A_j} I(i \in w)}.$$

We define  $s_n(x) = (s_{n,1}(x), \dots, s_{n,n}(x))$  the (row) structure vector of  $x$ , and

$$S_n = \begin{bmatrix} s_{n,1}(x_1) & \dots & s_{n,n}(x_1) \\ \vdots & \ddots & \\ s_{n,1}(x_n) & \dots & s_{n,n}(x_n) \end{bmatrix} = \begin{bmatrix} s_n(x_1) \\ \vdots \\ s_n(x_n) \end{bmatrix}$$

the structure matrix as the stacked structure vectors of the full sample. With this notation, structure-value isolation can also be viewed as the separation between the tree structure matrix and the response vector. Formally, we have

**Definition 1** *An ensemble of trees*

$$\hat{f}_{b,n}(x) = \sum_{j=1}^b t_n^j(x; w) = \sum_{j=1}^b \sum_{i=1}^n s_{n,i}^j(x; w)y_i$$

*exhibits structure-value isolation if  $s_{n,i}^j(x; w)$  is independent of the vector  $Y = (y_1, \dots, y_n)^T$  for all  $1 \leq j \leq b$  and  $x$ , where we use the superscript  $j$  to represent the  $j$ -th tree.*

For example, random forests utilizing completely randomized trees achieve structure-value isolation, while standard boosting and standard random forests do not.

## 2.2 Non-adaptivity

In addition to structure-value isolation, in order to demonstrate that Boulevard converges for a fixed sample, we also impose a condition which ensures that the distribution of tree structures stabilizes as Boulevard progresses. We describe this in two senses: a strong sense in which the distribution, conditional on the sample, is fixed for the whole sample, and a weak sense in which this occurs eventually.

**Definition 2** Given the sample  $(x_1, y_1), \dots, (x_n, y_n)$  of size  $n$ , denote  $(Q_{n,b}, \mathcal{Q}_{n,b})$  the probability space of  $b$ -th tree's structures after  $b$  trees have been built. Here  $Q_{n,b}$  consists of all possible tree structures and  $\mathcal{Q}_{n,b}$  the probability measure on  $Q_{n,b}$ . A tree ensemble is non-adaptive if  $(Q_{n,b}, \mathcal{Q}_{n,b}) = (Q_n, \mathcal{Q}_n)$  identical for all  $b$ . A tree ensemble is eventually non-adaptive if  $(Q_{n,b}, \mathcal{Q}_{n,b}) = (Q_n, \mathcal{Q}_n)$  identical for sufficiently large  $b$ .

Following this definition, both standard random forests and random forests utilizing completely randomized trees are non-adaptive. In terms of boosting, eventual non-adaptivity is a desirable condition should a GBDT ensemble become stationary after enough iterations, while the following algorithm provides a straightforward example to guarantee non-adaptivity.

**Algorithm 2 (Trees for Non-adaptive Boosting)**

- Start with  $(x_1, z_1), \dots, (x_n, z_n)$ , where  $z_1, \dots, z_n$  are current residuals.
- Obtain the tree structure  $q = \{A_j\}_{j=1}^m$  independently of  $z_1, \dots, z_n$ .
- If needed, uniformly subsample an index set  $w \subseteq \{1, \dots, n\}$  of size  $\theta n$ ,  $0 < \theta < 1$  (otherwise let  $w = \{1, 2, \dots, n\}$ ).
- Decide the leaf values, hence  $t_n(\cdot)$ , merely with respect to  $w$  as for  $x \in A_j$ ,

$$t_n(x) = \sum_{x_i \in A_j} \frac{I(i \in w)}{\sum_{x_l \in A_j} I(l \in w)} \cdot z_i,$$

with  $0/0$  defined to be 0.

Algorithm 2 is not specific about how to decide the tree structures. One can use completely randomized trees in which the gradients only influence the leaf values. Another strategy is to acquire another independent sample  $(x'_1, y'_1), \dots, (x'_n, y'_n)$  and implement a random forest on this sample solely for extracting tree structures.

Eventual non-adaptivity is helpful when the early trees are adaptive but have their influence diminishing gradually, ending up with no impact on the boosting limiting behavior. To fulfill eventual non-adaptivity, we may consider a strategy which uses a parallel adaptive boosting procedure on another independent sample to produce tree structures. Eventual non-adaptivity allows the use of any tree building strategy for a finite number of trees. This assumption turns out to be sufficient for our theoretical development as long as the ensemble does not stop too early. Meanwhile it is compatible with any pragmatic tree building algorithms for the beginning part where most signals get explored, which potentially accelerates the learning process. We will delay the discussion of fulfilling eventual non-adaptivity either spontaneously or manually to Section 5, and will focus on the Boulevard algorithm equipped with non-adaptivity mechanism as non-adaptive Boulevard unless otherwise stated.

### 2.3 Trees as Kernel Method

Another motivation for isolating tree structures and leaf values comes from the fact that it relates decision trees to a kernel method. To clarify, recall the definition of the tree structure matrix and denote by  $(Q_n, \mathcal{Q}_n)$  the probability space of all possible tree structures given sample  $(x_1, y_1), \dots, (x_n, y_n)$  of size  $n$  and, when necessary, a subsampling framework  $w$ , where  $q = \{A_i\}_{i=1}^{m_q} \in Q_n$  is the structure of a single possible tree. When we have structure-value isolation, we can write the expected decision tree prediction on this sample as

$$\hat{Y} = \mathbb{E}_q \mathbb{E}_w [S_n] \cdot Y = \mathbb{E}_{q,w} [S_n] \cdot Y, \tag{3}$$

where  $Y = (y_1, \dots, y_n)^T$ ,  $\mathbb{E}_w$  the expectation with respect to subsample indices and  $\mathbb{E}_q$  the expectation with respect to the probability measure  $\mathcal{Q}_n$ .

**Theorem 3** *Denote by  $\mathbb{E}_{q,w}$  the expectation over all possible tree structures and subsample index sets, then*

- (i)  $\mathbb{E}_{q,w}[S_n]$  is symmetric, element-wise nonnegative.
- (ii)  $\mathbb{E}_{q,w}[S_n]$  is positive semidefinite.
- (iii)  $\|\mathbb{E}_{q,w}[S_n]\|_1 \leq 1$ ,  $\|\mathbb{E}_{q,w}[S_n]\|_\infty \leq 1$ ,  $\|\mathbb{E}_{q,w}[S_n]\| \leq 1$ . Here the last  $\|\cdot\|$  stands for the spectral norm.

This  $\mathbb{E}_{q,w}[S_n]$  in (3) is similar to the random forest kernel (Scornet et al., 2015) defined by the corresponding tree structure space, subsampling strategy and tree structure randomization approach. In terms of a single tree, honesty contributes to its symmetry and positive semidefiniteness, while subsampling decides the concentration level of the kernel. While for an ensemble, non-adaptivity guarantees that  $\mathbb{E}_{q,w}[S_n]$  applies to all its tree components, making it possible to use  $\hat{Y}_b = \mathbb{E}_{q,w}[S_n] \cdot Y_b$ , to describe the prediction from any tree, where  $Y_b$  is the target negative gradients and  $\hat{Y}_b$  the fitted values after  $b$ -th iteration.

It is also worth noticing that in part (iii) of Theorem 3 we only guarantee inequality rather than setting the norm exactly to 1. This is a side effect of structure-value isolation: for instance, when applying subsample splitting, it is possible that a leaf node decided by the first subsample contains no points of the second subsample. For such cases we can predict 0 for expediency: we demonstrate in Section 4.2 that this choice has an asymptotically negligible effect below.

## 3. Convergence

As stated in Zhang et al. (2005), a first theoretical challenge in analyzing boosting methods is the difficulty of attaining convergence. As a starting point we will show that Boulevard guarantees point-wise convergence under finite sample settings.

### 3.1 Stochastic Contraction and Boulevard Convergence

To study Boulevard convergence, we start by informally postulating a convergent point  $Y^*$  for any boosting algorithm, such that the boosting path should remain ideally “stationary” at and after  $Y^*$ . For standard boosting with  $L^2$  loss, supposing  $Y^*$  happens at the  $(b-1)$ -th

iteration, the next update is  $Y^* + \lambda_b S_b \cdot (Y - Y^*)$  with  $S_b$  the tree structure matrix decided by the gradient vector.  $Y^*$  being the convergent point implies that  $\lambda_b S_b \cdot (Y - Y^*) = 0$ . Either we can implement a diminishing sequence of learning rates  $\lambda_b \rightarrow 0$ , or we must guarantee that we stop boosting with  $S_b \cdot (Y - Y^*) \rightarrow 0$ . The latter condition depends on both the properties of the gradient vector  $Y - Y^*$  and the tree building algorithm deciding  $S_b$ , making it challenging to study either the existence of such point  $Y^*$  or the convergence of the boosting path to  $Y^*$ .

For Boulevard, the update will be  $\frac{b-1}{b}Y^* + \frac{\lambda}{b}S_b(Y - Y^*)$ , implying convergence at  $Y^* = \lambda \mathbb{E}[S_b](Y - Y^*)$  when  $S_b$  is random. For trees with non-adaptivity and structure-value isolation, we can solve this relation for  $Y^*$  to prove its existence. Furthermore, Theorem 14 can be used to study the convergence of the boosting path to such  $Y^*$ . We formalize our discussion as the following theorem.

**Theorem 4** *Given sample  $(x_1, y_1), \dots, (x_n, y_n)$ . If we construct gradient boosted trees with structure-value isolation and non-adaptivity using tree structure space  $(Q_n, \mathcal{Q}_n)$ , by choosing  $M \gg \max\{M_f, y_1, \dots, y_n\}$  and defining  $\Gamma_M(x) = \text{sign}(x)(|x| \wedge M)$  as a truncation function, let the Boulevard iteration take the form*

$$\hat{f}_b(x) = \frac{b-1}{b} \hat{f}_{b-1}(x) + \frac{\lambda}{b} s_b(x)(Y - \Gamma_M(\hat{Y}_{b-1})),$$

where  $Y = (y_1, \dots, y_n)^T$  is the observed response vector,  $\hat{Y}_b = (\hat{f}_b(x_1), \dots, \hat{f}_b(x_n))^T$  is the predicted response vector by the first  $b$  trees, and  $s_b$  is the random tree structure vector. For a.s.  $x$ ,

$$\hat{Y}_b \longrightarrow \left[ \frac{1}{\lambda} I + \mathbb{E}[S_b] \right]^{-1} \mathbb{E}[S_b] Y, b \rightarrow \infty,$$

where  $\mathbb{E}[\cdot] = \mathbb{E}_{q,w}[\cdot]$ ,  $S_b$  the random tree structure matrix defined above.

The proof of this theorem relies on a stochastic contraction theorem introduced by Almudevar (2022) regarding a special class of stochastic processes. We refer the readers to the original manuscript, but the theorem and its key points of the proof are briefly reproduced in Appendix A.2 and extended to cover a Kolmogorov inequality. Theorem 4 guarantees the convergence of Boulevard paths under finite sample setting once we threshold by a large  $M$ . As a corollary we can express the prediction at any point of interest  $x$ , which coincides with a regularized kernel regression with ridge penalty using the random forest kernel.

**Corollary 5** *By defining  $\hat{f} = \lim_{b \rightarrow \infty} \hat{f}_b$ ,*

$$\hat{f}(x) = \mathbb{E}[s_n(x)] \left[ \frac{1}{\lambda} I + \mathbb{E}[S_n] \right]^{-1} Y. \quad (4)$$

The shrinkage form of Corollary 5 indicates that Boulevard does not recover the full signal. We will see in later proofs that Boulevard can be made to converge to  $\frac{\lambda}{1+\lambda} f(x)$  under proper rate assumptions. In fact, Boulevard down-weights the boosting history, thereby ensuring that each tree in the finite ensemble must predict some partial signal during training. It thus avoids being dominated by the first few trees then repeatedly fitting on noise. In

practice, since we will show that the prediction from Boulevard is consistent with respect to  $\frac{\lambda}{1+\lambda}f(x)$ , we simply rescale it by  $\frac{1+\lambda}{\lambda}$  to retrieve the full signal if needed. This convergence can be further bounded loosely when we choose  $\lambda < \sqrt{2} - 1$ .

**Theorem 6** For given  $0 < \lambda < \sqrt{2} - 1$ ,

$$\mathbb{E} \left[ \left\| \hat{Y}_b - Y^* \right\|^2 \right] = O \left( \frac{1}{b + \lambda} \right). \quad (5)$$

The convergence rate proportional to the square root of the ensemble size comes at the cost that we cannot credit a single decision tree beyond the given threshold.

### 3.2 Beyond $L^2$ Loss

Besides regression, other tasks may require alternative loss functions for boosting. For instance, the exponential loss  $L(w, y) = \exp(-wy)$  in adaboost (Freund and Schapire, 1995). Analogous to the proof for  $L^2$  loss, we can write the counterparts for any general loss  $L(u) = \sum_i L(u_i, y_i)$  whose non-adaptive Boulevard iteration takes the form of

$$\hat{Y}_b = \frac{b-1}{b} \hat{Y}_{b-1} - \frac{\lambda}{b} S_b \nabla_w L(w) \Big|_{w=\hat{Y}_{b-1}}.$$

Suppose the existence of the fix point  $\hat{Y}^* = -\lambda \mathbb{E}[S_b] \nabla_w L(w) \Big|_{w=\hat{Y}^*}$ , then

$$\mathbb{E}[\hat{Y}_b - \hat{Y}^* | \mathcal{F}_{b-1}] = \frac{b-1}{b} (\hat{Y}_{b-1} - \hat{Y}^*) - \frac{\lambda}{b} \mathbb{E}[S_b] \left( \nabla_w L(w) \Big|_{w=\hat{Y}_{b-1}} - \nabla_w L(w) \Big|_{w=\hat{Y}^*} \right).$$

If the gradient term is bounded and Lipschitz (which could be enforced by truncation), i.e.

$$\left\| \nabla_w L(w) \Big|_{w=w_1} - \nabla_w L(w) \Big|_{w=w_2} \right\| \leq M \|w_1 - w_2\|,$$

we can similarly show such Boulevard iteration converges by choosing  $\lambda \leq M^{-1}$ . However, the closed form of  $\hat{Y}^*$  can be intractable to obtain and potentially non-unique. For example for AdaBoost,  $\hat{Y}^*$  is the solution to  $\hat{Y}^* = -\lambda \mathbb{E}[S_n] (\exp(-\hat{Y}_1^* y_1), \dots, \exp(-\hat{Y}_n^* y_n))^T$ .

## 4. Limiting Distribution

Inspired by recent results demonstrating the asymptotic normality of random forest predictions, in this section we prove the asymptotic normality of predictions from Boulevard. Before detailing these results, we need some prerequisite discussion on the rates used for decision tree construction in order to ensure asymptotic local behavior. In general, the variability of model predictions comes from three sources: the sample variability, response errors, and the systematic variability of boosting caused by its subsampling and stopping rule. As shown in the last section, Boulevard convergence relieves the need for considering the stopping rule and the systematic variability. Therefore the strategy for our proof is as follows: we first consider the boosting process conditioned on the sample  $X$  so only the response errors contribute to the variability. We then establish the uniformity over almost all random samples to extend the limiting distribution to the unconditional cases, showing that it is still the response errors that dominate the prediction variability.

#### 4.1 Building Deeper Trees

Decision trees can be thought as k-nearest-neighbor (k-NN: Altman, 1992) models where  $k$  is the leaf size and the distance metric is given by whether two points are in the same leaf. This adapts the metric to the local geometry of the response function. As the conclusions on k-NN predictions require growing-in-size and shrinking-in-radius neighborhoods (Gordon and Olshen, 1984), so are the counterparts of building deeper trees. Assuming non-adaptivity, the following assumptions are sufficient for our analysis.

(L1) Asymptotic locality. We introduce a sequence  $\{d_n\}, d_n > 0, d_n \rightarrow 0$  by which the diameter of any leaf is bounded from above. Writing  $diam(A) = \sup_{x,y \in A} |x - y|$ , we require,

$$\sup_{A \in \mathcal{q} \in Q_n} diam(A) = O(d_n).$$

(L2) Minimal leaf size. We introduce another sequence  $\{v_n\}, v_n > 0, v_n \rightarrow 0$  by which the volume of any leaf is bounded from below. Writing  $V(\cdot)$  as the volume function by Lebesgue measure, we require that

$$\inf_{A \in \mathcal{q} \in Q_n} V(A) = \Omega(v_n).$$

Recall the notation that  $f = \Omega(g) \iff g = O(f)$ . These place requirements on all the leaves of a tree; we specify the rates we require for  $d_n$  and  $v_n$  below.

These assumptions together bound the space occupied by any leaf of any possible tree from being either too extensive or too small. Along with the Lipschitz condition, this geometrically shrinking neighborhood not only indicates that each tree is asymptotically guaranteed to gather points with close response values together, but also helps to control the impact of one point on another across multiple consecutive iterations. To be specific, because boosting fits gradients instead of responses, a pair of distant points with divergent values might still influence each other's fitted values through a bridge created by other points linked by being in the same leaf nodes across multiple iterations. Here (L1) provides a theoretical upper bound for such influence.

Notice that randomized honest tree building strategies can be compatible with (L1) if we force it, whereas the standard CART algorithm does not imply (L1). However, since the intuition of (L1) is to put points with similar responses together, we anticipate that CART, tuned to a proper depth, will satisfy (L1), at least in practice. If so, we can still exercise Boulevard boosting using trees built with greedy algorithms without compromising its theoretical guarantees in practice.

#### 4.2 Conditioned on the Sample

We first consider a sample triangular array, i.e. for each  $n$ , the sample  $(x_{n,1}, y_{n,1}), \dots, (x_{n,n}, y_{n,n})$  is given. The first subscript  $n$  will be dropped when there is no ambiguity. We specify the rates for the size of leaf nodes as:

(R1) For some  $\epsilon_1 > 0$ ,

$$d_n = O\left(n^{-\frac{1}{d+2}-\epsilon_1}\right).$$

(R2) For some  $\epsilon_2 > \epsilon_1 > 0$ ,

$$\inf_{A \in \mathcal{q} \in \mathcal{Q}_n} \sum_{i=1}^n I(x_i \in A) = \Omega\left(n^{\frac{2}{d+2}-d\epsilon_2}\right).$$

One compatible realization is

$$d_n = O\left(n^{-\frac{1}{d+1}}\right), \quad \inf_{A \in \mathcal{q} \in \mathcal{Q}_n} \sum_{i=1}^n I(x_i \in A) = \Omega\left(n^{\frac{1}{d+2}}\right).$$

For simplicity all our proofs are under this setting. However, any other rates satisfying these conditions are also sufficient.

Starting here we use the abbreviations

$$k_n^T = \mathbb{E}[s_n(x)], \quad K_n = \mathbb{E}[S_n], \quad r_n^T = k_n^T \left[ \frac{1}{\lambda} I + K_n \right]^{-1}.$$

to cover the components of (4). The rates of  $k_n$  and  $r_n$  determine Boulevard asymptotics.

Recall that the structure-value isolation paired with subsampling may lead to the consequence of predicting 0 in some leaves, so that we only guarantee  $\|k_n\|_1 \leq 1$ . Working with the tree construction rates as above, the subsample rate  $\theta$  determines how far away  $\|k_n\|_1$  is from 1. We show in Appendix A.7 for any  $x$ ,  $1 - \|k_n\|_1 = O\left(\frac{1}{n}\right)$  if we use subsample size at least of  $n^{\frac{d+1}{d+2}} \log n$ . This requires the subsample to be relatively large, which is compatible with both  $\theta$  being constant, or  $\theta = (\log n)^{-1}$ .

### 4.3 Exponential Decay of Influence and Asymptotic Normality

The prediction that Boulevard makes at a point is a linear combination of responses  $y_1, \dots, y_n$  whose coefficients are given by  $r_n$ . Therefore the asymptotical normality of Boulevard predictions will rely on whether  $r_n$  can comply with the Lindeberg-Feller condition. Since  $r_n$  only differs from  $k_n$  by a ridge regression style matrix multiplication, we can expect  $r_n$  to have analogous properties to  $k_n$ . Further, by examining  $r_n$  we can show that distant points are less influential on the prediction, and this decay of influence is exponential in our case.

Given any  $n$ -vector  $v$  and an index set  $D$ , denote

$$v|_D = \begin{bmatrix} v_1 \cdot I(1 \in D) \\ \vdots \\ v_n \cdot I(n \in D) \end{bmatrix}.$$

This notation implies the decomposition that  $v = v|_D + v|_{D^c}$ .

**Lemma 7** Given sample  $(x_1, y_1), \dots, (x_n, y_n)$ , a point of interest  $x$ , set  $l_n = \frac{\log n}{-\log \lambda} = c_1 \log n$ , and define index set  $D_n = \{i : |x_i - x| \leq l_n \cdot d_n\}$ , then

(i) Globally,

$$\left| \sum_{i=1}^n r_{n,i} - \frac{\lambda}{1 + \lambda} \right| = O\left(\frac{1}{n}\right).$$

(ii) If we only look nearby,

$$\|r_n|_{D_n^c}\|_1 = O\left(\frac{1}{n}\right).$$

Lemma 7 indicates that Boulevard trees will asymptotically rely on a  $\log n$  shrinking neighborhood around the point of interest. Given sample size  $n$  and a point of interest  $x$ , we can therefore define two neighborhoods of different radii  $B_n = \{i : |x_i - x| \leq d_n\}$  and  $D_n = \{i : |x_i - x| \leq l_n \cdot d_n\}$ .  $B_n$  contains all points that have direct influence on  $x$  in a single tree, and  $D_n$  contains the points that dominate the prediction at  $x$ . Their cardinalities  $|B_n|$  and  $|D_n|$  follow Binomial distributions with parameters depending on the local covariate density.

To show the limiting distribution of the sequence of predictions  $\hat{f}_n(x)$ , the key point is to verify that, by writing the prediction at any point  $x$  as a linear combination of sample responses, all these coefficients are diminishing at a rate allowed by the Lindeberg-Feller condition. Therefore we take a look at the coefficient vectors  $k_n$  and  $r_n$ .

**Lemma 8** Given the triangular array with sample  $(x_{n,1}, y_{n,1}), \dots, (x_{n,n}, y_{n,n})$  at size  $n$ , assume

- The smaller neighborhood is growing fast enough:  $|B_n| = \Omega(n \cdot d_n^d)$ .
- We have enough points in each leaf:

$$\inf_{A \in \mathcal{q} \in Q_n} \sum_{i=1}^n I(x_{n,i} \in A) = \Omega\left(n^{\frac{1}{d+2}}\right).$$

Then

$$\Omega\left(n^{-\frac{1}{2} \frac{1}{d+1}}\right) = \|k_n\|, \|r_n\| = O\left(n^{-\frac{1}{2} \frac{1}{d+2}}\right).$$

We note that these rates correspond to very slowly-growing trees with practical implementations of boosting generally fixing the depth of individual trees between 2 and 4. Combining all our discussions above, the following theorem stating the asymptotic normality conditioned on a given sample sequence follows. Lemma 8 is used here to estimate the order of the variance term.

**Theorem 9 (Conditional Asymptotic Normality for Boulevard Predictions)** For any  $x \in [0, 1]^d$ , suppose the conditions in Lemma 8 hold. Write  $f(X_n) = (f(x_1), \dots, f(x_n))^T$ , then under non-adaptivity and structure-value isolation,

$$\frac{\hat{f}_n(x) - r_n^T f(X_n)}{\|r_n\|} \xrightarrow{d} N(0, \sigma_\epsilon^2).$$

#### 4.4 Sample Variability

To extend the scope of the conditional limiting distribution to consider the sample variability, we require that all the assumptions we made leading to Theorem 9 also hold true for almost surely all random sample sequence. In order to do so, we would like to review the sample variability in a well-defined probability space regarding the triangular array constructed from a random sample sequence. This can be done by the Kolmogorov's extension theorem; see details in Appendix A.11. We also have to make (R2) slightly more strict.

(R2') We increase the minimal leaf geometric volume by any small  $\nu > 0$  s.t.  $v_n$  follows

$$v_n = \frac{n^{\frac{1}{d+2}+\nu}}{n} = n^{-\frac{d+1}{d+2}+\nu} < n^{-\frac{d}{d+1}} = O\left(d_n^d\right).$$

With this assumption, the following lemma proves that Theorem 9 holds in probability for possible random samples.

**Lemma 10** *For given  $x \in [0, 1]^d$ , suppose we have random sample  $(x_1, y_1), \dots, (x_n, y_n)$  for each  $n$ . If we restrict the cardinality of tree space  $Q_n$  using any small  $\alpha > 0$  s.t.*

$$|Q_n| = O\left(\frac{1}{n} \exp\left(\frac{1}{2}n^{\frac{1}{d+2}-\nu} - n^\alpha\right)\right),$$

then

$$\frac{\hat{f}_n(x) - r_n^T f(X_n)}{\|r_n\|} \xrightarrow{d} N(0, \sigma_\epsilon^2).$$

By further showing that the difference between  $r_n^T f(X_n)$  and  $\frac{\lambda}{1+\lambda}f(x)$  is negligible, we obtain the main theorem of this paper that the limiting distribution of the random design non-adaptive Boulevard predictions is normal.

**Theorem 11 (Asymptotic Normality for Boulevard Predictions)** *Under the conditions of Lemmas 8 and 10 as well as non-adaptivity and structure-value isolation, for given  $x \in [0, 1]^d$ ,*

$$\frac{\hat{f}_n(x) - \frac{\lambda}{1+\lambda}f(x)}{\|r_n\|} \xrightarrow{d} N(0, \sigma_\epsilon^2).$$

Theorem 11 employs a deterministic mean term to center the estimate but scales by a stochastic variance term. From results on kernel ridge regression, we would expect that this stochastic variance converges in probability if the random forest kernel behaves as a generic kernel with a shrinking bandwidth. From a theoretical perspective, the optimal rate of  $\|r_n\|$  is bounded from below by  $n^{-\frac{1}{2}\frac{1}{d+1}}$ , which corresponds to the optimal nonparametric regression rate using  $\frac{1}{2}$ -Hölder continuous functions as base learners (Stone, 1982). In practice,  $\|r_n\|$  relies on the specific method of growing the boosted trees, and therefore may vary from case to case.

Furthermore, Theorem 11 demonstrates that with carefully structured trees Boulevard predictions are consistent after scaling and the response errors dominate the prediction variability. It acts as an undersmoothed local smoother whose bias term shrinks faster than the variance term.

#### 4.5 Subsampling and Tree Space Capacity

We have a strict requirement that the tree terminal node size grows at a rate between  $n^{\frac{1}{d+1}}$  and  $n^{\frac{1}{d+2}}$  to guarantee undersmoothing. Any log term is allowed to be added to the existing polynomial result without changing the behavior. We notice that different subsample rates (i.e.  $\theta = (\log n)^{-d}$  in Wager et al. (2014),  $\theta = n^{-1/2}$  in Mentch and Hooker (2016), although see relaxations in Peng et al. (2019)) have been applied for measuring uncertainty. In comparison, the Boulevard algorithm requires a relatively restricted rate between these. In addition, although Boulevard training implements subsampling at each iteration, this does not influence the asymptotic distribution. The impact of subsampling is on the possible variance from the mean process therefore the convergence speed especially if we assume non-adaptivity.

We have required the size of tree space to scale at a rate close to  $O\left(\frac{1}{n} \exp\left(\frac{1}{2}n^{\frac{1}{d+2}}\right)\right)$ . In comparison, Wager and Walther (2015) have shown that, in fixed dimension, any tree can be well approximated by a collection of  $O(\exp((\log n)^2))$  hyper-rectangles. One feasible way to compare these two rates is to consider a tree space satisfying Boulevard rates in which each tree has several leaves that constitute one hyper-rectangle in the collection for approximation. Therefore a tree space of cardinality  $O(\exp((\log n)^2))$  is enough for approximating any tree ensemble. In this sense, the capacity of our designated tree space is large enough from a practical perspective.

#### 4.6 Alternative Analyses

The distributional results above are based on an analysis of a kernel representation of the Boulevard process in which structure-value isolation plays a key role in both consistency and asymptotic normality results. By contrast, Mentch and Hooker (2016) and Wager and Athey (2017) employed results derived from U-statistics to demonstrate asymptotic normality when random forest trees are obtained using subsamples. This analysis uses the subsample structure to bypass the kernel representation and allows the trees to be built adaptively. In a similar manner, when Boulevard employs subsampling we may represent its fixed point as a  $U$  statistic. However, as in our discussion in Section 3.2,  $U$ -statistic kernel is then defined implicitly with respect to a non-smooth tree-building process and further analysis is beyond the scope of this paper.

Wager and Athey (2017) also uses regularity to constrain leaf sizes in order to demonstrate consistency outside of the  $U$ -statistic framework. We expect that a similar weakening of these conditions is possible in Boulevard.

### 5. Eventual Non-adaptivity

All the results discussed above have assumed the non-adaptivity of the boosting procedure of Boulevard, which is unconventional in contrast to how boosted trees pragmatically decide both their structures and values using the same batch of negative gradients. Such procedures are known for their tendency to overfit, behavior which can be relieved by subsampling. However, when seeking to extend our results to this case we lose the easy identifiability of a Boulevard convergence point as the tree structure distribution changes at each iteration.

To extend our convergence results to a more practical Boulevard algorithm that allows the current gradient to determine tree structure, a first approach is to relax non-adaptivity to eventual non-adaptivity, both of which have the same limiting distributional behavior. We postulate a convergent sequence of predictions, indicating that the underlying tree spaces will be stabilized after boosting for a sufficiently long time. If this is the case, eventual non-adaptivity is spontaneous. Here we introduce the notation  $\mathbb{E}[S_n(Y, \hat{Y})]$  where  $Y = (y_1, \dots, y_n)^T$  and  $\hat{Y} = (\hat{f}(x_1), \dots, \hat{f}(x_n))^T$  indicating the expected tree structure given the gradient of the loss between observed responses and current predictions. In regression this is  $Y - \hat{Y}$ , and we will take this form into the following discussion instead of a generic gradient expression.

### 5.1 Local Homogeneity and Contraction Regions

We start with trees whose splits are based on the optimal  $L^2$  gain (Breiman et al., 1984). For  $(x_1, z_1), \dots, (x_n, z_n)$ , the chosen split minimizes the impurity in the form of

$$\inf_{L,R} \sum_{i \in L} (z_i - \bar{z}_L)^2 + \sum_{i \in R} (z_i - \bar{z}_R)^2, \quad (6)$$

where  $L \subset \{1, \dots, n\}, R = L^c$ . Once the optimal split is unique, i.e. the optimum has a positive margin over the rest, we could allow a small change of all  $y$ 's values without changing the split decision. This also holds true if the split is decided by a subsample. In terms of adaptive boosting, this observation demonstrates local homogeneity that, except on a set  $\Omega_0 \subset \mathbb{R}^n$  with Lebesgue measure 0 where  $(z_1, \dots, z_n)^T = Y - \hat{Y} \in \Omega_0$  has multiple optima for (6), we can segment  $\mathbb{R}^n$ , the space of possible vectors  $Y - \hat{Y}$ , into subsets  $\bigsqcup_{i=1}^\alpha C_i = \mathbb{R}^n \setminus \Omega_0$  s.t.  $\mathbb{E}[S_n(Y, \hat{Y})] = \mathbb{E}[S_n(Y, \hat{Y}')] for  $Y - \hat{Y}, Y - \hat{Y}' \in C_i$  the same subset.$

Notice that Gini gain is insensitive to multiplying  $(y_1, \dots, y_n)$  by a nonzero factor. Therefore all  $C_i$ 's are open double cones in  $\mathbb{R}^n$ .

**Definition 12 (Contraction Region)** *Given the sample  $(x_1, y_1), \dots, (x_n, y_n)$ . Write  $Y = (y_1, \dots, y_n)$  and current prediction  $\hat{Y} = (\hat{y}_1, \dots, \hat{y}_n)$ . Following the above segmentation  $\bigsqcup_{i=1}^\alpha C_i = \mathbb{R}^n \setminus \Omega_0$ . We call  $C_i$  a contraction region if  $Y^* \in C_i$  for the following  $Y^*$*

$$Y^* = \lambda \mathbb{E}[S_n(Y, \hat{Y})](Y - Y^*), \text{ i.e. } Y^* = \left[ \frac{1}{\lambda} I + \mathbb{E}[S_n(Y, \hat{Y})] \right]^{-1} \mathbb{E}[S_n(Y, \hat{Y})]Y,$$

for any  $Y - \hat{Y} \in C_i$ , where  $\mathbb{E}[S_n(Y, \hat{Y})]$  is the unique structural matrix in this region.

The intuition behind this definition is that, as long as a Boulevard process stays inside a contraction region, the subsequent tree structures will be conditionally independent of the predicted values. Therefore the process becomes non-adaptive, collapsing to  $Y^*$ . To achieve this eventual non-adaptivity, we would like to know when a Boulevard path is permanently contained in a contraction region.

We should point out here that we have not shown the existence and the uniqueness of such contraction regions. Such an analysis would rely on the method of choosing splits, the training sample and the choice of  $\lambda$ .

## 5.2 Escaping the Contraction Region

In this section we explore possible approaches to restrict a Boulevard process inside a contraction region. Assuming the existence of contraction regions, we recall Theorem 14 which indicates that the Boulevard process has positive probability of not moving far from the fixed point.

**Theorem 13** *Denote  $B(x, r)$  the open ball of radius  $r$  centered at  $x$  in  $\mathbb{R}^n$ . Suppose  $C \subset \mathbb{R}^n$  a contraction region,  $Y^* \in C$  the contraction point and  $B(Y^*, 2r) \subset C$  for some  $r > 0$ . Write  $(\hat{Y}_b)_{b \geq 1}$  the Boulevard process. For sufficiently large  $t$ ,*

$$P\left(\hat{Y}_b \in C, \forall b \geq t \mid \hat{Y}_t \in B(Y^*, r)\right) \rightarrow 1, t \rightarrow \infty.$$

Theorem 13 states that the longer we boost, the more likely the boosting path is trapped in a contraction region forever, resulting in spontaneous eventual non-adaptivity and justifying all our theory assuming non-adaptivity. However, it guarantees neither the existence or the uniqueness of the contraction region.

## 5.3 Forcing Eventual Non-Adaptivity

If we alternatively consider manually forcing eventual non-adaptivity, a possible *ad hoc* solution to the existence is to apply a tail snapshot which uses the tree space of certain iteration  $b^*$  for the rest of the boosting run. This  $b^*$  can be either pre-specified or decided on the fly when the Boulevard path becomes stationary. Ideally, as long as  $b^*$  is in a contraction region, its forced non-adaptivity is indistinguishable from spontaneous eventual non-adaptivity. An example of Boulevard regression implementing the tail snapshot strategy is detailed in Algorithm 3, in which we decide  $b^*$  by examining the training loss reduction.

### Algorithm 3 (Tail Snapshot Boulevard)

- Start with  $\hat{f}_0 = 0$  and  $b^*$  to be chosen.
- For  $b = 0, \dots$ , given  $\hat{f}_b$ , calculate the gradient

$$z_i \triangleq -\frac{\partial}{\partial u_i} \sum_{i=1}^n \frac{1}{2} (u_i - y_i)^2 \Big|_{u_i = \Gamma_M(\hat{f}_b(x_i))} = y_i - \Gamma_M(\hat{f}_b(x_i));$$

- If  $b^*$  has not been chosen, update by  $1 > \lambda > 0$  and the tree structure space  $Q_b$  decided by current gradients,

$$\hat{f}_{b+1}(x) = \frac{b}{b+1} \hat{f}_b(x) + \frac{\lambda}{b+1} s_b(x; Q_b)(z_1, \dots, z_n)^T,$$

where  $s_b(x; Q)$  denotes the random tree structure vector based on tree space  $Q$ . If  $b^*$  has been chosen, update by  $Q_{b^*}$  instead, i.e.

$$\hat{f}_{b+1}(x) = \frac{b}{b+1} \hat{f}_b(x) + \frac{\lambda}{b+1} s_b(x; Q_{b^*})(z_1, \dots, z_n)^T.$$

- When  $b^*$  has not been chosen, check the empirical training loss as a measure of the distance to the fixed point.

$$L_{b+1} = \frac{1}{2n} \sum_{i=1}^n \left( \frac{\lambda}{1+\lambda} y_i - \hat{f}_{b+1}(x_i) \right)^2.$$

If  $L_{b+1} < L^*$  a preset threshold, we claim Boulevard is close enough to a fixed point and choose  $b^*$  to be current  $b+1$ . This step is then skipped for the rest of the boosting run.

Another plausible solution is to collectively build the tree space up based on the following intuition. For the first trees we prefer to construct them using the gradients to capture the signal. However when the boosting process keeps going on, we may consider reusing the tree structures or the tree spaces of previous iterations as to avoid overfitting the data or fitting the errors. Therefore the tree space  $Q_{n,b}$  for a large  $b$  should be the tree structures created by current gradients, plus all tree spaces  $Q_{n,i}$  for  $i < b$ . This accumulation scheme changes as we boost, however asymptotically it leads to a common tree space and thus eventual non-adaptivity.

## 6. Empirical Study

We have conducted an empirical study<sup>2</sup> to demonstrate the performance of Boulevard. Despite the fact that our purpose in developing Boulevard lies in statistical inference, we require its accuracy to be on a par with other state-of-the-art tree ensembles, which is assessed on both simulated and real world data. In addition, we inspect the empirical limiting behavior of non-adaptive Boulevard to show its agreement with our theory. We summarize the result of the empirical study in this section, while additional details can be found in Appendix B.

### 6.1 Predictive Accuracy

We first compare Boulevard predictive accuracy with the following tree ensembles: Random Forest (RF), gradient boosted trees without subsampling (GBT), stochastic gradient boosted trees (SGBT), non-adaptive Boulevard achieved by completely randomized trees (rBLV), adaptive Boulevard whose tree structures are influenced by gradient values (BLV). All the tree ensembles build same depth of trees (see Appendix B.1 for details) throughout the experiment.

Results on simulated data are shown in Figure 1 and 2. We choose sample size of 5000 and use the following two settings as underlying response functions: (1)  $y = x_1 + 3x_2 + x_3x_4$  (top), and (2)  $y = x_1 + 3x_2 + (1 - x_3)^2 + x_4x_5 + (1 - x_6)^6 + x_7$  (bottom). Error terms are Unif[-1,1] (left) and equal point mass on  $\{-1, 1\}$  (right). Training errors are evaluated on the training set with noisy responses, while test errors are evaluated by error from the underlying signal on a separate test set. For each setup we have two plots covering the behavior of the first 50 trees and the following 250 trees respectively. BLV and rBLV are comparable with RF, while all the three equal-weight tree ensembles are slightly inferior to

---

2. The empirical study code is provided at: <https://github.com/siriuz42/boulevard.git>

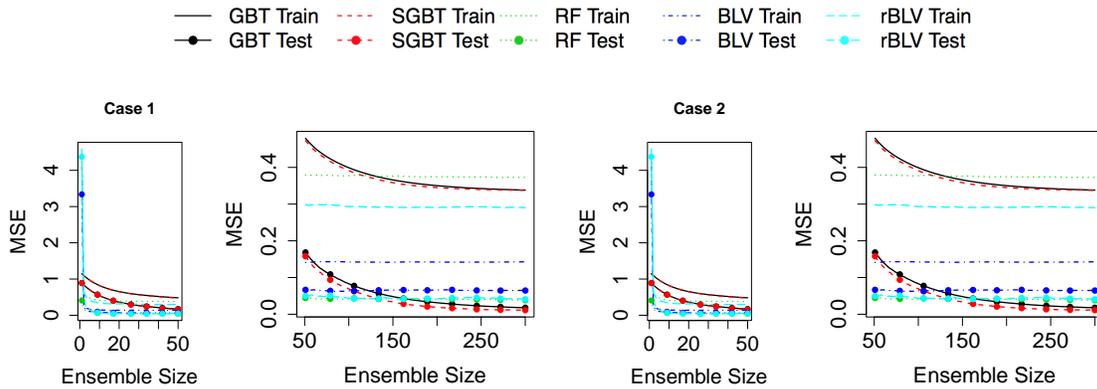


Figure 1: Training and testing error curves of tree ensembles on simulated data generated by function (1). Training errors: left Unif[-1,1], right Unif{-1,1}. Testing errors are smaller because they are calculated on test sets with no errors.

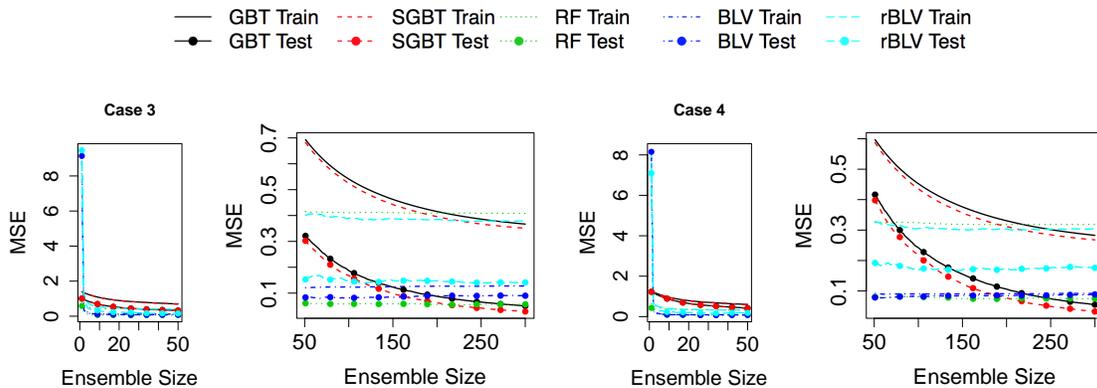


Figure 2: Training and testing error curves of tree ensembles on simulated data generated by function (2). Training errors are: Unif[-1,1] on the left, Unif{-1,1} on the right. Testing errors are smaller because they are calculated on test sets with no errors.

GBT and SGBT. In particular, test errors are smaller as they are calculated on test sets with no errors.

In addition, we have used the same settings and applied different  $\lambda$  values 0.2, 0.5 and 0.8 on two simulated data sets to investigate the impact of  $\lambda$  on Boulevard accuracy for finite sample cases. Results are summarized in Figure 3. BLV outperforms rBLV while  $\lambda = 0.8$  appears to be the best choice. Though this observation aligns with our intuition to choose a large  $\lambda$ , we consider the discrepancy between different  $\lambda$  values to be small.

Results on four real world data sets selected from UCI Machine Learning Repository (Dheeru and Karra Taniskidou, 2017; Tüfekci, 2014; Kaya et al., 2012) are shown in Figure

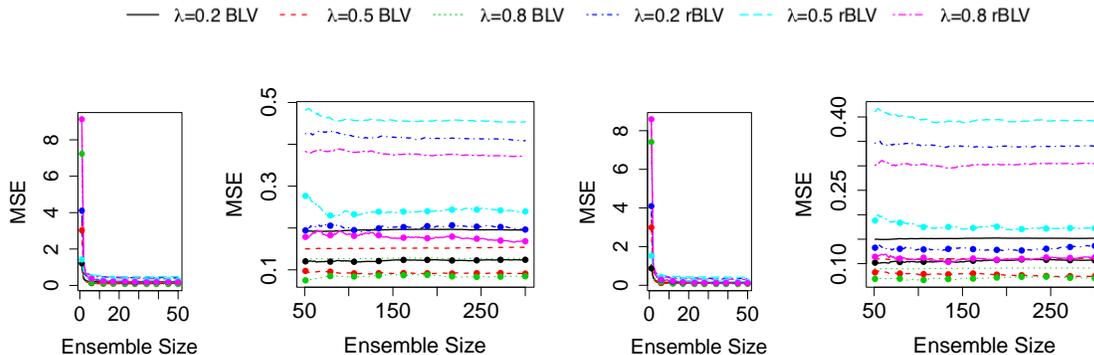


Figure 3: Boulevard accuracy with different choices of  $\lambda$  values. Training errors are: Unif[-1,1] on the left, Unif{-1,1} on the right. Test errors are marked by dots and are smaller as they are calculated on test sets with no errors.

4. All curves are averages after 5-fold cross validation. Different parameters are used for different data sets, see Appendix B.1. Rankings of the five methods in comparison do not show consistency, nevertheless rBLV and BLV manage to achieve similar performance to the other methods on test sets.

### 6.2 Convergence and Kernel Ridge Regression

Theorem 4 and Corollary 5 claim that any Boulevard prediction should converge to a kernel ridge regression (KRR) form. Although the matrix inverse involved in the expression is too time consuming to compute for large  $n$ , it is possible to estimate it using Monte Carlo when  $n$  and  $d$  are small. We take advantage of this setup to compare Boulevard and KRR in order to empirically support the theorem.

We choose  $d = 5$  and  $n = 200$ . We use the following model

$$y = x_1 + 3x_2 + x_3^2 + 2x_4x_5 + \epsilon, \quad \epsilon \sim \text{Unif}[-1, 1]. \tag{7}$$

to generate training data. For 4 test points (see Appendix B.1), we predict their values using non-adaptive Boulevard with 100 trees, subsampling rate 0.8 and  $\lambda = 0.8$ . We also make KRR predictions based on the expected tree structure matrix  $\mathbb{E}[S]$  estimated by Monte Carlo using the same 100 trees. This procedure is repeated 20 times with a new training sample each time. Our results are summarized in Figure 5.

The plot on the left shows the interim predictions made by one Boulevard run. The shaped dots are Boulevard predictions, while the horizontal lines are the predictions made by KRR. We see that Boulevard iterations converge, while the convergent points agree with KRR predictions. On the right we plot the predictions made by Boulevard against the ones made by KRR for each of the 20 repetitions, with the horizontal and vertical lines marking the true responses. All dots are close to the diagonal  $x = y$ , indicating these two methods always produce almost identical results. It also shows that the sample variability is influencing the two methods in the same way. Their MSEs, as labeled in the plots, also appear to be indistinguishable.

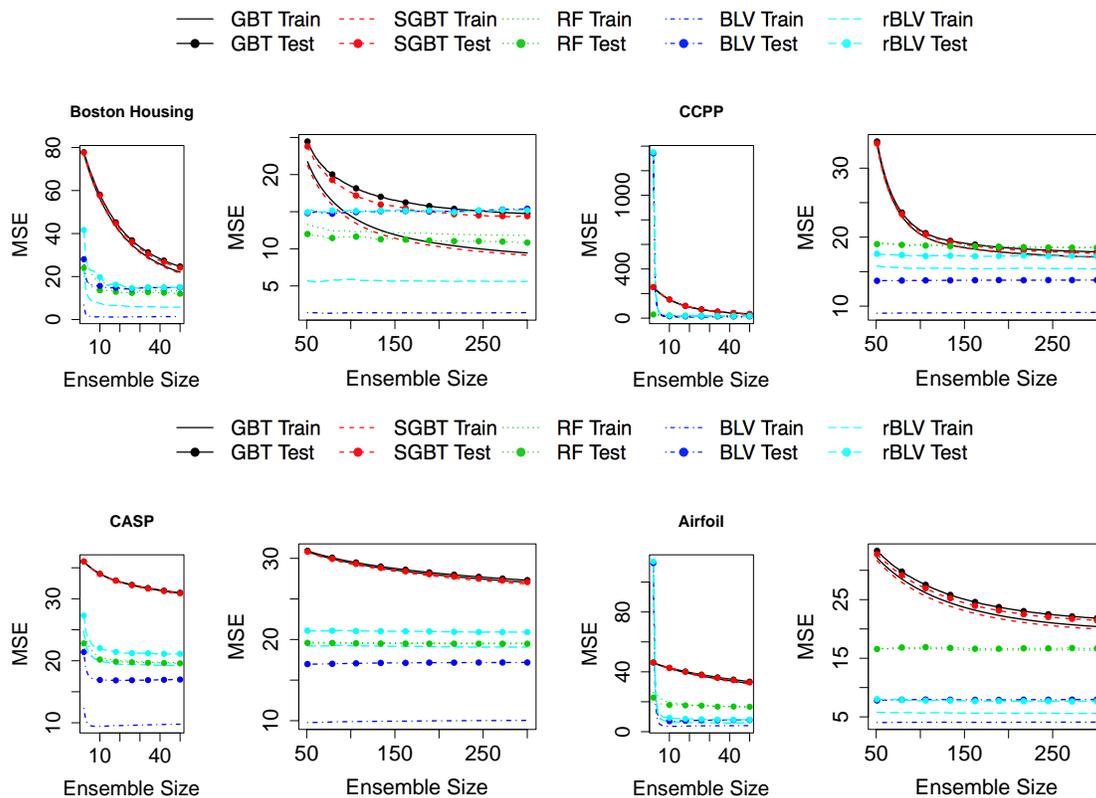


Figure 4: Training and testing error curves of tree ensembles on real world data sets. Data sets are: Boston Housing (top left), Combined Cycle Power Plant (top right), Protein Tertiary Structure (bottom left) and Airfoil (bottom right).

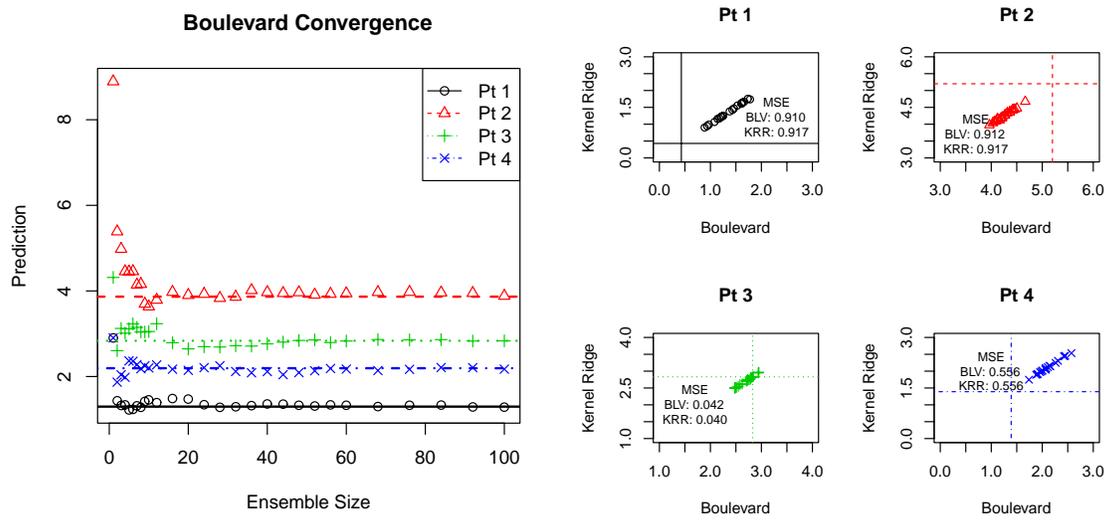


Figure 5: **Left:** Boulevard iterations on 4 test points converge to the theoretical convergent points given by the KRR form. The shaped dots are interim Boulevard predictions at different iterations, while horizontal lines are KRR results. **Right:** Boulevard predictions v.s. KRR predictions after 20 repetitions. Two methods generate indistinguishable predictions along the diagonal with similar MSEs. True responses are marked by the horizontal and vertical lines.

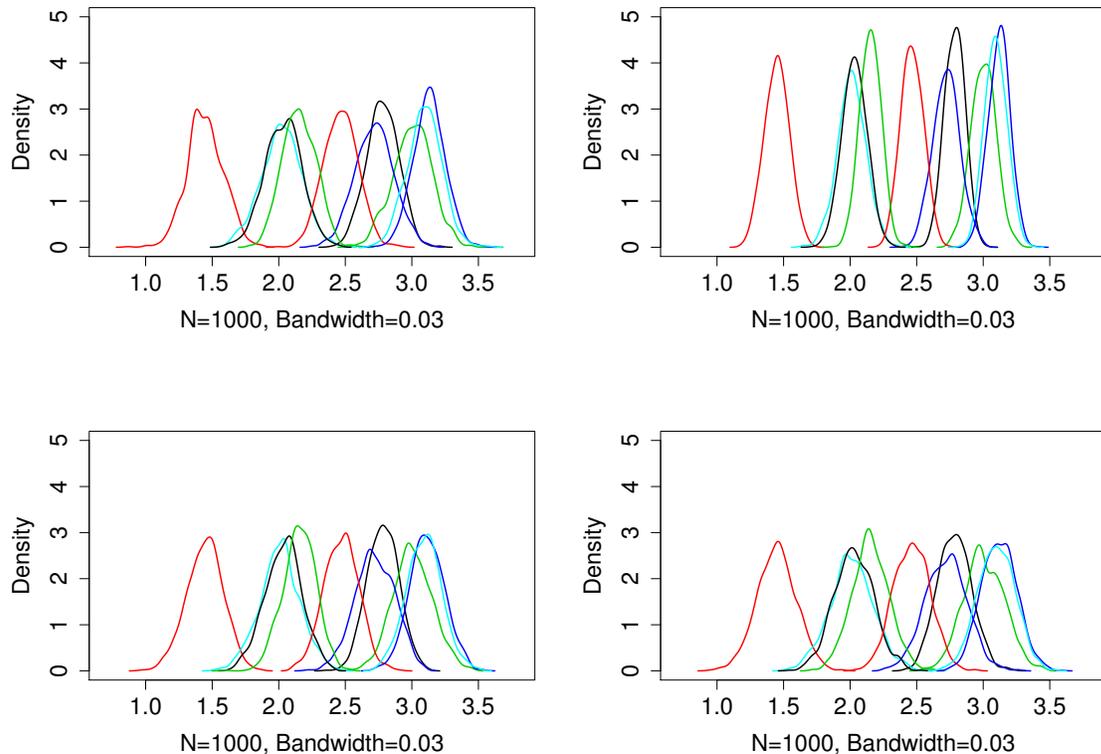


Figure 6: Distributions of predictions of test points with different error terms. The errors are  $N(0,1)$ ,  $\text{Unif}[-1,1]$ ,  $\text{Unif}\{-1,1\}$ , and half chance -1 half chance  $\text{Unif}[0,2]$ , respectively.

### 6.3 Limiting Distribution

To examine the limiting behavior of non-adaptive Boulevard, we again use the generative model

$$y = x_1 + 3x_2 + x_3^2 + 2x_4x_5 + \epsilon.$$

A set of 10 fixed test points (see Appendix B.1) are used along the experiments. We set a sample size of 1000, add different sub-Gaussian error terms to this signal and build non-adaptive Boulevard until ensemble size reaches 2000. This is repeated 1000 times with a new sample each time and we plot the distribution of the predictions in Figure 6. All these curves are undistinguishable from normal distribution by Kolmogorov-Smirnov test.

### 6.4 Reproduction Interval

Similar to prediction intervals which quantify the uncertainty of future predictions, we introduce the reproduction interval as the uncertainty measure for where the prediction

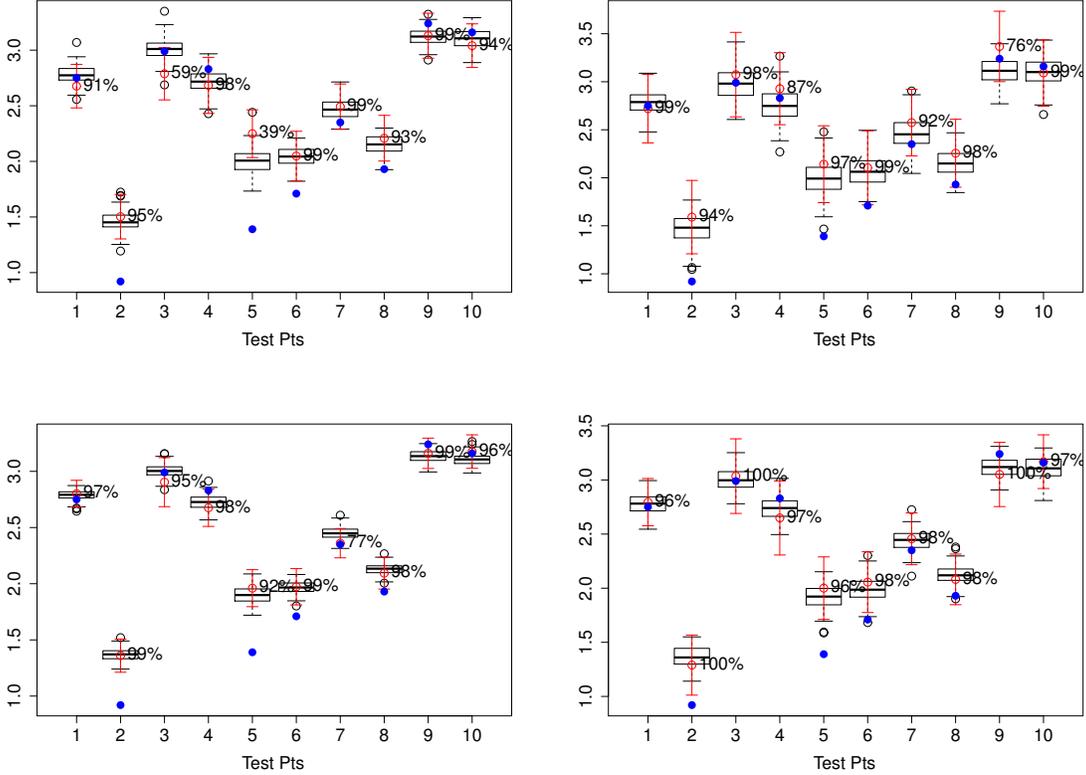


Figure 7: Reproduction intervals. Boxplots show distributions of predictions; red intervals are reproduction intervals; blue dots are truths. Sample sizes are 1000 (top row) and 5000 (bottom row), error terms Unif[-1,1] (left column) and Unif[-2, 2] (right column). Coverage is shown by numbers next to interval centers.

would be if it were made on another independent sample. Theorem 11 is used to create reproduction intervals for Boulevard.  $k_n$  in the stochastic variance is empirically estimated directly using the ensemble, while  $[\frac{1}{\lambda}I + K_n]^{-1}$  is conservatively simplified to its largest possible norm  $\lambda$ . We then scale the variance estimate by 2 to account for having separate independent samples: with two independent samples, we expect the prediction on the same point  $x$  to be  $\hat{f}_1(x) \sim N(f(x), \sigma_x^2)$  and  $\hat{f}_2(x) \sim N(f(x), \sigma_x^2)$  respectively, implying  $\hat{f}_1(x) - \hat{f}_2(x) \sim N(0, 2\sigma_x^2)$ .

We use the training sample to create reproduction intervals for the test points, then repeatedly train and predict each test point for another 100 times with a different sample each time. Figure 7 shows the 95% reproduction intervals we capture under different settings. We anticipate more accurate results with larger sample size.

Furthermore, we notice the uniform pattern of biases in those plots. This bias comes from two known causes. One is that we are using small samples which are far from guaranteeing the consistency. The other is because of the edge effects; the distance of the ten

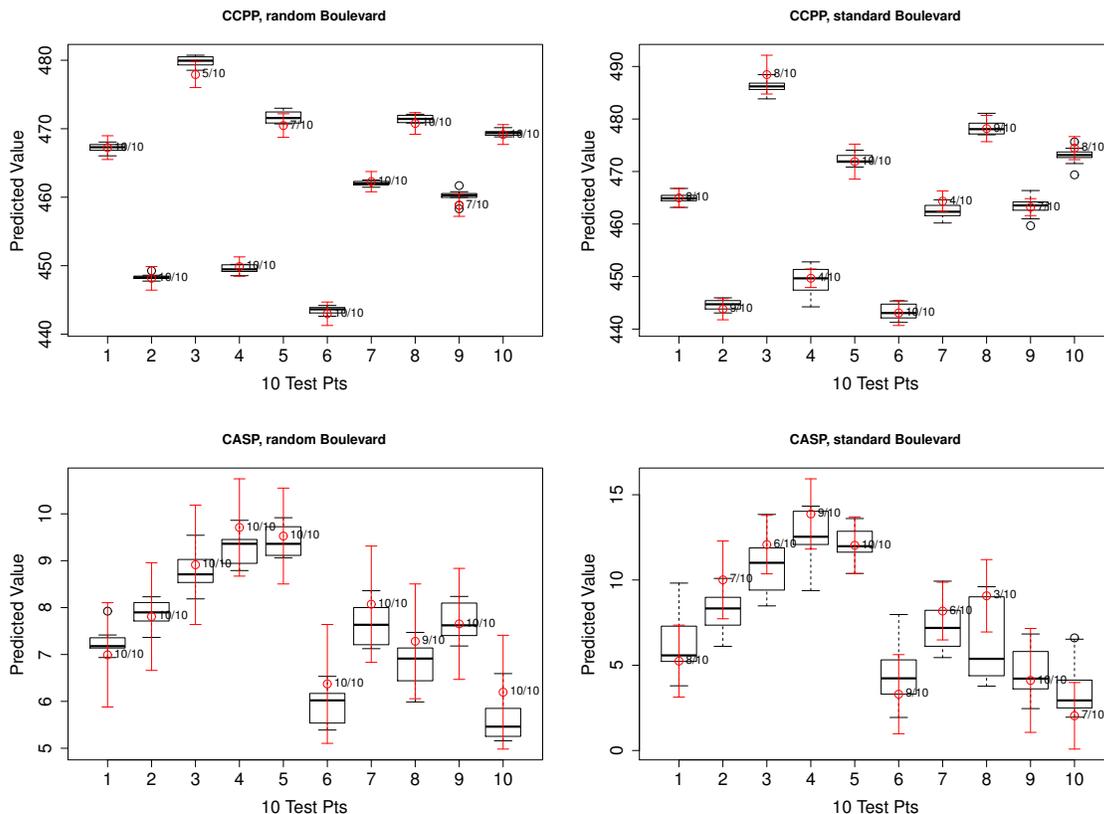


Figure 8: Reproduction intervals for real world datasets. Boxplots show distributions of predictions; red intervals are reproduction intervals. Coverage is shown by numbers next to interval centers.

chosen test points to the center of the hypercube is respectively 0.000, 0.671, 0.894, 0.894, 0.894, 0.693, 0.520, 0.436, 0.510 and 0.469. We in general expect biased prediction when the point is near the boundary.

Figure 8 shows the result when we generate the 90% reproduction intervals for two real world datasets from UCI, namely CCPP and CASP. For each dataset, we take the first 10 examples as test examples, and split the rest of the dataset into 11 folds. The first fold is used to create the reproduction intervals while the rest ten folds are used to create out-sample predictions. That ends up with 850 (CCPP) and 2700 (CASP) examples per fold respectively.

## 7. Discussion

This work is among the earliest attempts we know to have established a limiting distribution for gradient boosted trees. The roadmap consists of the following key components. We implemented the honest tree construction to reduce the chance of chasing order statistics,

applied downweighting towards averaging to achieve convergence, and carefully selected tree construction rate to obtain asymptotic normality. With the uncertainty measure of Boulevard predictions, we are looking forward to exploring the use of regularized gradient boosting within inference, making the model more interpretable and analyzable.

The sequential correlation induced by GBDT is the major issue complicating our analysis, while the resistance of decision trees to mathematical quantification adds further challenges. Much of our effort is spent on seeking reasonable conditions from which to account for these features. Structure-value isolation, non-adaptivity and down-weighting all contribute to making sequential trees less correlated.

As briefly discussed above, there are two remaining questions in our paper. The first question regards the stochastic variance. To allow a richer collection of decision tree construction strategies, it is essential to discover the weakest tree condition under which the variance term converges in probability. The second question is the convergence of adaptive Boulevard which appears to hold in practice. In spite of the *ad hoc* tail snapshot we proposed, the existence and uniqueness of the contraction region of generic Boulevard requires more mathematical formulation of decision trees or some variation of decision trees.

## Acknowledgments

We would like to acknowledge support for this project from the National Science Foundation (NSF DMS-1712554).

## Appendix A. Proofs

### A.1 Proof of Theorem 3

**Proof** We will prove the theorem with respect to  $\mathbb{E}_w[\cdot]$ . The  $\mathbb{E}_{w,q}[\cdot]$  conclusion is an instant corollary as the three properties are additive. To prove (1), element-wise non-negativity is trivial. To show symmetry, consider any given  $i \neq j$  and assume  $x_i \in A$  and  $x_j \in A'$  under the assumption of subsample uniformity,

$$\begin{aligned}\mathbb{E}_w[S_n]_{i,j} &= \mathbb{E}_w[s_{n,j}(x_i)] = \frac{1}{\binom{n}{\theta n}} \sum_w \frac{I(x_j \in A)I(j \in w)}{\sum_{x_l \in A} I(l \in w)} \\ \mathbb{E}_w[S_n]_{j,i} &= \mathbb{E}_w[s_{n,i}(x_j)] = \frac{1}{\binom{n}{\theta n}} \sum_w \frac{I(x_i \in A')I(i \in w)}{\sum_{x_l \in A'} I(l \in w)}.\end{aligned}$$

Therefore  $\mathbb{E}_w[S_n]_{i,j} = \mathbb{E}_w[S_n]_{j,i} = 0$  if  $A \neq A'$ .

In the cases of  $A = A'$ ,  $I(x_j \in A) = I(x_i \in A') = 1$ . We consider the following possibilities of  $w$ .

(a) For  $i \notin w, j \notin w$ ,

$$\frac{I(j \in w)}{\sum_{x_l \in A} I(l \in w)} = \frac{I(i \in w)}{\sum_{x_l \in A} I(l \in w)} = 0.$$

(b) For  $i \in w, j \in w$ ,

$$\frac{I(j \in w)}{\sum_{x_l \in A} I(l \in w)} = \frac{I(i \in w)}{\sum_{x_l \in A} I(l \in w)} = \frac{1}{\sum_{x_l \in A} I(l \in w)}.$$

(c) For  $i \in w, j \notin w$ , consider  $w' = w \setminus \{i\} \cup \{j\}$  s.t.  $\sum_{x_l \in A} I(l \in w) = \sum_{x_l \in A} I(l \in w')$ ,

$$\frac{I(j \in w')}{\sum_{x_l \in A} I(l \in w')} = \frac{I(i \in w)}{\sum_{x_l \in A} I(l \in w)} = \frac{1}{\sum_{x_l \in A} I(l \in w)}.$$

(d) Similarly, for  $i \notin w, j \in w$ , consider  $w' = w \setminus \{j\} \cup \{i\}$ ,

$$\frac{I(j \in w)}{\sum_{x_l \in A} I(l \in w)} = \frac{I(i \in w')}{\sum_{x_l \in A} I(l \in w')} = \frac{1}{\sum_{x_l \in A} I(l \in w)}.$$

Since all  $w$ 's are equally likely, we conclude by symmetry that  $\mathbb{E}_w[S_n]_{i,j} = \mathbb{E}_w[S_n]_{j,i}$ , hence  $\mathbb{E}_w[S_n]$  is symmetric.

To prove (2), notice  $\forall x_i, x_j, x_k \in A$ ,

$$\mathbb{E}_w[S_n]_{k,i} = \frac{1}{\binom{n}{\theta n}} \sum_w \frac{I(i \in w)}{\sum_{x_l \in A} I(l \in w)} = \mathbb{E}_w[S_n]_{j,i}.$$

Therefore  $\mathbb{E}_w[S_n]$ , after proper permutation to gather points in same leaves together, is diagonally blocked with equal non-negative entries in each diagonal block and 0 elsewhere, thus positive semi-definite.

To show (3), notice that  $S_n$  has column sums of  $\leq 1$  (not = 1 due to chances of having no subsample points in the leaf), so does  $\mathbb{E}_w[S_n]$ . Thus  $\|\mathbb{E}_w[S_n]\|_1 \leq 1$ . Similarly,  $\mathbb{E}_w[S_n]$  has rows sums of  $\leq 1$  due to its symmetry therefore  $\|\mathbb{E}_w[S_n]\|_\infty \leq 1$ . Therefore,

$$\begin{aligned} \rho(\mathbb{E}_w[S_n]) = \|\mathbb{E}_w[S_n]\| &\leq \sqrt{\sup_{\|u\|=1} \mathbb{E}_w[S_n]^T \mathbb{E}_w[S_n] u} \\ &\leq \sqrt{\sup_{\|u\|=1} \|\mathbb{E}_w[S_n]^T\|_\infty \|\mathbb{E}_w[S_n]\|_\infty \|u\|_\infty} \\ &\leq \sqrt{\sup_{\|u\|=1} \|\mathbb{E}_w[S_n]\|_1 \|\mathbb{E}_w[S_n]\|_\infty \cdot 1} \\ &\leq 1. \end{aligned}$$

■

## A.2 Stochastic Contraction

We first state the stochastic contraction theorem and its extension to Kolmogorov inequality. The proof follows afterwards.

**Theorem 14 (Stochastic Contraction)** *Given  $\mathbb{R}^d$ -valued stochastic process  $\{Z_t\}_{t \in \mathbb{N}}$ , a sequence of  $0 < \lambda_t \leq 1$ , define*

$$\begin{aligned} \mathcal{F}_0 &= \emptyset, \mathcal{F}_t = \sigma(Z_1, \dots, Z_t) \text{ the generated } \sigma \text{ field,} \\ \epsilon_t &= Z_t - \mathbb{E}[Z_t | \mathcal{F}_{t-1}]. \end{aligned}$$

We call  $Z_t$  a stochastic contraction if the following properties hold

(C1) *Vanishing coefficients*

$$\sum_{t=1}^{\infty} (1 - \lambda_t) = \infty, \text{ i.e. } \prod_{t=1}^{\infty} \lambda_t = 0.$$

(C2) *Mean contraction*

$$\|\mathbb{E}[Z_t | \mathcal{F}_{t-1}]\| \leq \lambda_t \|Z_{t-1}\|, \text{ a.s..}$$

(C3) *Bounded deviation*

$$\sup \|\epsilon_t\| \rightarrow 0, \quad \sum_{t=1}^{\infty} \mathbb{E}[\|\epsilon_t\|^2] < \infty.$$

In particular, a multidimensional stochastic contraction exhibits the following behavior

(i) *Contraction*

$$Z_t \xrightarrow{\text{a.s.}} 0.$$

(ii) *Kolmogorov inequality*

$$P \left( \sup_{t \geq T} \|Z_t\| \leq \|Z_T\| + \delta \right) \geq 1 - \frac{4\sqrt{d} \sum_{t=T+1}^{\infty} \mathbb{E}[\epsilon_t^2]}{\min\{\delta^2, \beta^2\}} \quad (8)$$

holds for all  $T, \delta > 0$  s.t.  $\beta = \|Z_T\| + \delta - \sqrt{d} \sup_{t > T} \|\epsilon_t\| > 0$ .

**Definition 15 (Stochastic Contraction)** Given real-valued stochastic process  $\{X_t\}_{t \in \mathbb{N}}$ , a sequence of  $0 < \lambda_t \leq 1$ , define

$$\begin{aligned} \mathcal{F}_0 &= \emptyset, \mathcal{F}_t = \sigma(X_1, \dots, X_t), \\ \epsilon_t &= X_t - \mathbb{E}[X_t | \mathcal{F}_{t-1}]. \end{aligned}$$

We call  $X_t$  a stochastic contraction if the following is satisfied

- *Vanishing coefficients*

$$\sum_{t=1}^{\infty} (1 - \lambda_t) = \infty, \text{ i.e. } \prod_{t=1}^{\infty} \lambda_t = 0.$$

- *Mean contraction*

$$\lambda_t X_{t-1} I(X_{t-1} \leq 0) \leq \mathbb{E}[X_t | \mathcal{F}_{t-1}] \leq \lambda_t X_{t-1} I(X_{t-1} \geq 0), \text{ a.s..}$$

- *Bounded deviation*

$$\sup |\epsilon_t| \rightarrow 0, \quad \sum_{t=1}^{\infty} \mathbb{E}[\epsilon_t^2] \leq \infty.$$

**Lemma 16** If  $\{X_t\}_{t \in \mathbb{N}}$  is a stochastic contraction.

- *Almost sure convergence*

$$X_t \xrightarrow{a.s.} 0.$$

- *Kolmogorov maximal inequality. For any  $T, \delta > 0$  s.t.  $\beta = |X_T| + \delta - \sup_{t>T} |\epsilon_t| > 0$ ,*

$$P\left(\sup_{t \geq T} |X_t| \leq |X_T| + \delta\right) \geq 1 - \frac{4 \sum_{t=T+1}^{\infty} \mathbb{E}[\epsilon_t^2]}{\min\{\delta^2, \beta^2\}}.$$

**Proof** Define the stopping time of sign changes

$$T_0 = 0, T_k = \inf\{t > T_{k-1} | X_{t-1} \leq 0, X_t > 0 \text{ or } X_{t-1} \geq 0, X_t < 0\}.$$

We now look at every realized path and examine the segment of the process holding the same sign. W.o.l.g., suppose  $X_t \geq 0$  for  $T_k < t < T_{k+1}$ . Easy to check

$$X_t = \mathbb{E}[X_t | \mathcal{F}_{T_k}] + \epsilon_t \leq \lambda_t X_{T_k} + \epsilon_t \leq X_{T_k} + \sum_{s=T_k+1}^t \epsilon_s. \quad (9)$$

Therefore  $|X_t| \leq |X_{T_k}| + \left| \sum_{s=T_k+1}^t \epsilon_s \right|$ , same for the negative case. Since  $\epsilon_t$ 's are independent and  $\sum_{t=1}^{\infty} \mathbb{E}[\epsilon_t^2] < \infty$ ,  $\sum_{t=1}^{\infty} \epsilon_t$  exists a.s.. Write  $N = \sup_k \{T_k \leq \infty\}$  the number of sign changes.

If there are infinitely many sign changes, i.e.  $N = \infty$ , by sending  $k \rightarrow \infty$ ,  $|X_{T_k}| \xrightarrow{a.s.} 0$  and  $\left| \sum_{s=T_k+1}^{T_k+n} \epsilon_s \right| \xrightarrow{a.s.} 0, \forall n > 0$ . Hence  $X_t \xrightarrow{a.s.} 0$ .

If there are finitely many sign changes, we assume w.l.o.g. that for some  $k$ ,  $X_t \geq 0, \forall t \geq T_k$ . (9) can be written as  $X_t - \epsilon_t \leq X_{t-1}$  which indicates  $X_t - \sum_{s=T_k+1}^t \epsilon_s$  is decreasing, therefore has a limit  $(-\infty)$ . Since  $\sum_{s=T_k+1}^{\infty} \epsilon_s$  exists a.s.,  $X_t \xrightarrow{a.s.} c \geq 0$ . Assume  $c > 0$ ,

$$\sum_{s=T_k+1}^{\infty} \epsilon_s \geq \sum_{s=T_k+1}^{\infty} X_s - \lambda_s X_{s-1} = -\lambda_{T_k+1} X_{T_k} + \sum_{s=T_k+2}^{\infty} (1 - \lambda_s) X_{s-1} = \infty,$$

which is a contradiction. Therefore  $X_t \xrightarrow{a.s.} 0$ .

To show the maximum inequality, we take the same notations above, and also look at segmentations by sign changes. For any  $t$  in the same segment as  $T$ ,

$$|X_t| \leq |X_T| + \left| \sum_{s=T+1}^t \epsilon_s \right| \leq |X_T| + \sup_{T' > T} \left| \sum_{s=T+1}^{T'} \epsilon_s \right|.$$

For any  $t$  in a different segment starting at  $T'$ ,

$$|X_t| \leq |X_{T'}| + \left| \sum_{s=T'+1}^t \epsilon_s \right| \leq |X_{T'}| + \sup_{S > T'} \left| \sum_{s=T'+1}^S \epsilon_s \right| \leq \sup_{s > T} |\epsilon_s| + \left| \sum_{s=T'+1}^S \epsilon_s \right|.$$

Now we consider any possible sequence of  $\{\epsilon_t, t > T\}$  and allow  $T', S$  to change. Kolmogorov maximal inequality implies

$$P\left(\sup_{i,j>T}\left|\sum_{s=i}^j \epsilon_s\right| \leq x\right) \geq P\left(\sup_{i>T}\left|\sum_{s=T}^i \epsilon_s\right| \leq \frac{x}{2}\right) \geq 1 - \frac{4\sum_{s=T}^{\infty} \mathbb{E}[\epsilon_s^2]}{x^2}.$$

The conclusion is obtained by noticing that  $|X_t| \leq |X_T| + \delta$  for any  $\{\epsilon_t\}_{t>T}$  satisfying

$$\sup_{i,j>T}\left|\sum_{s=i}^j \epsilon_s\right| \leq \min\{\delta, \beta\}.$$

■

### A.3 Proof of Theorem 14

**Proof** The idea is to define a sequence of adaptive orthonormal rotations  $R_t \in \mathcal{F}_{t-1}$  to align the expected update with the previous step so that we can apply the  $\mathbb{R}$  result component-wise. Define  $R_t \mathbb{E}[Z_t | \mathcal{F}_{t-1}] = \gamma_{t-1} Z_{t-1}$ , for some  $\gamma_{t-1} > 0, \gamma_{t-1} \in \mathcal{F}_{t-1}$ . The contraction assumption also implies that  $\gamma_{t-1} \leq \lambda_{t-1}$ . Define a new process  $Z_i^*$  satisfying

1.  $Z_1^* = Z_1, R_1 = I$ ,
2. writing  $R_t^* = \prod_{i=1}^n R_i \in \mathcal{F}_{t-1}$  s.t.  $Z_t^* = R_t^* Z_t = R_t^* \epsilon_t + R_t^* \mathbb{E}[Z_t | \mathcal{F}_{t-1}]$ .

Above implies  $\|Z_t\| = \|Z_t^*\|$ , thus we need to prove the equivalence that  $Z_t^* \xrightarrow{a.s.} 0$ . Notice that  $\sum_{i=1}^n R_i^* \epsilon_i$  is component-wise a martingale with

$$\sum_{i=1}^{\infty} \mathbb{E}[\|R_i^* \epsilon_i\|^2] = \sum_{i=1}^{\infty} \mathbb{E}[\|\epsilon_i\|^2] < \infty,$$

hence  $\sum_{i=1}^n R_i^* \epsilon_i$  exists a.s.. Since the construction aligns  $Z_t^*$  with  $\mathbb{E}[Z_t^* | \mathcal{F}_{t-1}]$  we apply Lemma 16 to obtain almost sure convergence to 0 component-wisely, thus  $\|Z_t^*\| \xrightarrow{a.s.} 0$ . ■

### A.4 Proof of Theorem 4

**Proof** Due to non-adaptivity  $S_n$  is independent of  $\hat{Y}_b$  for any  $b$ . Notice that  $Y^* = \lambda \mathbb{E}[S_n](Y - Y^*)$  for  $Y^* = \left[\frac{1}{\lambda} I + \mathbb{E}[S_n]\right]^{-1} \mathbb{E}[S_n] Y$ . Notice we write  $b$  here in place of  $t$  used by the stochastic contraction. Define the filtration  $\mathcal{F}_b = \sigma(\hat{Y}_0, \dots, \hat{Y}_b)$  and consider the sequence  $Z_b = \hat{Y}_b - Y^*$ . This sequence satisfies the stochastic contraction condition.

First,  $\|Z_0\| = \|Y^*\| \leq \infty$ . Notice

$$\begin{aligned}
\|\mathbb{E}[Z_b|\mathcal{F}_{b-1}]\| &= \left\| \mathbb{E} \left[ \frac{b-1}{b} \hat{Y}_{b-1} + \frac{\lambda}{b} S_n(Y - \Gamma_M(\hat{Y}_{b-1})) - Y^* \middle| \mathcal{F}_{b-1} \right] \right\| \\
&= \left\| \frac{b-1}{b} (\hat{Y}_{b-1} - Y^*) + \frac{\lambda}{b} \mathbb{E}[S_n](Y - \Gamma_M(\hat{Y}_{b-1})) - \frac{1}{b} Y^* \right\| \\
&\leq \frac{b-1}{b} \|\hat{Y}_{b-1} - Y^*\| + \left\| \frac{\lambda}{b} \mathbb{E}[S_n](Y - \Gamma_M(\hat{Y}_{b-1})) - \frac{\lambda}{b} \mathbb{E}[S_n](Y - Y^*) \right\| \\
&\leq \frac{b-1+\lambda}{b} \|\hat{Y}_{b-1} - Y^*\| \triangleq k_b \|Z_{b-1}\|,
\end{aligned}$$

where  $\sum_{b=1}^{\infty} (1 - k_b) = \infty$ . Since entries and row sums of are both  $\leq 1$ ,

$$\|S_n\| \leq \sqrt{\|S_n\|_{\infty} \|S_n\|_1} \leq \sqrt{1 \times n} = \sqrt{n}.$$

Therefore

$$\|\epsilon_b\| = \|Z_b - \mathbb{E}[Z_b|\mathcal{F}_{b-1}]\| = \left\| \frac{\lambda}{b} (\mathbb{E}[S_n] - S_n)(Y - \Gamma_M(\hat{Y}_{b-1})) \right\| \leq \frac{\lambda}{b} (1 + \sqrt{n}) 2\sqrt{n}M.$$

Hence

$$\sum_{b=1}^{\infty} \mathbb{E}[\|\epsilon_b\|^2] \leq \left( \sum_{b=1}^{\infty} \frac{1}{b^2} \right) \cdot \lambda^2 (1 + \sqrt{n})^2 4nM < \infty.$$

We conclude that  $Z_b \xrightarrow{a.s.} 0$ , i.e.  $\hat{Y}_b \xrightarrow{a.s.} Y^*$ . ■

## A.5 Proof of Corollary 5

**Proof** Expanding  $\hat{f}(x)$  gives

$$\begin{aligned}
\hat{f}(x) &= \lim_{B \rightarrow \infty} \frac{1}{B} \sum_{b=1}^B s_b(x)(Y - \hat{Y}_b) \\
&= \lim_{B \rightarrow \infty} \frac{1}{B} \sum_{b=1}^B s_b(x)(Y - Y^* + Y^* - \hat{Y}_b) \\
&= \lim_{B \rightarrow \infty} \frac{1}{B} \sum_{b=1}^B s_b(x)(Y - Y^*) + \lim_{B \rightarrow \infty} \frac{1}{B} \sum_{b=1}^B s_b(x)(Y^* - \hat{Y}_b) \\
&= \mathbb{E}[s_b(x)](Y - Y^*) + 0 \\
&= \mathbb{E}[s_n(x)] \left[ \frac{1}{\lambda} I + \mathbb{E}[S_n] \right]^{-1} Y.
\end{aligned}$$
■

### A.6 Proof of Theorem 6

**Proof** Recall the notations in the proof of Theorem 4 that

$$Z_b = \hat{Y}_b - Y^*, \quad \epsilon_b = Z_b - \mathbb{E}[Z_b | \mathcal{F}_{b-1}], \quad k_b = \frac{b-1+\lambda}{b}.$$

We therefore have, for  $C_1 = (1 + \sqrt{n})2\sqrt{n}M$  that

$$\begin{aligned} \mathbb{E}[\|Z_b\|^2] &= \mathbb{E}[\|\epsilon_b\|^2] + \mathbb{E}[\|\mathbb{E}[Z_b | \mathcal{F}_{b-1}]\|^2] \\ &\leq \frac{C_1 \lambda^2}{b^2} + \left(\frac{b-1+\lambda}{b}\right)^2 \mathbb{E}[\|Z_{b-1}\|^2]. \end{aligned}$$

Proceeding by induction, for any  $C_2 \geq C_1$  s.t.

$$\mathbb{E}[\|Z_{b-1}\|^2] \leq \frac{C_2}{b-1+\lambda},$$

we have

$$\begin{aligned} \mathbb{E}[\|Z_b\|^2] &\leq \frac{C_2 \lambda^2}{b^2} + C_2 \left(\frac{b-1+\lambda}{b}\right)^2 \frac{1}{b-1+\lambda} \\ &= C_2 \left(\frac{b-1+\lambda^2+\lambda}{b^2}\right) \\ &\leq \frac{C_2}{b+\lambda}, \end{aligned}$$

for  $\lambda < \sqrt{2} - 1$  and sufficiently large  $b^* = b^*(\lambda)$ . We take  $C_2 = \max(\mathbb{E}[\|Z_{b^*}\|^2]/\lambda, C_1)$  to complete the proof.  $\blacksquare$

### A.7 Subsample Rate

For Algorithm 2, without loss of generality, suppose we obtain a tree structure where each leaf contains no fewer than  $n^{\frac{1}{d+2}}$  sample points before applying subsampling to decide the fitted values. If the subsample size is  $\theta n = n^{\frac{d+1}{d+2}} \log n$ , i.e.  $\theta = n^{-\frac{1}{d+2}} \log n$ , the chance of missing all sample points in one leaf is

$$\begin{aligned} p(n, \theta) &= \frac{\binom{n-n^{\frac{1}{d+2}}}{\theta n}}{\binom{n}{\theta n}} = \frac{(n-\theta n)(n-\theta n-1)\cdots(n-\theta n-n^{\frac{1}{d+2}}+1)}{n(n-1)\cdots(n-n^{\frac{1}{d+2}}+1)} \\ &\leq \left(\frac{n-\theta n}{n-n^{\frac{1}{d+2}}}\right)^{n^{\frac{1}{d+2}}} = \left(\frac{1-n^{-\frac{1}{d+2}} \log n}{1-n^{-\frac{d+1}{d+2}}}\right)^{n^{\frac{1}{d+2}}} \\ &\leq e \cdot \left(1-n^{-\frac{1}{d+2}} \log n\right)^{n^{\frac{1}{d+2}}} = O\left(\frac{1}{n}\right). \end{aligned}$$

### A.8 Proof of Lemma 7

**Proof** To prove (i), given  $k_n$ , we focus on the matrix multiplication. Consider the expansion

$$\left[ \frac{1}{\lambda} I + K_n \right]^{-1} = \lambda \sum_{i=0}^{\infty} ((\lambda)^{2i} K_n^{2i} - (\lambda)^{2i+1} K_n^{2i+1}).$$

We examine the column sums of each of the matrix powers. Start with  $K_n^2$ ,

$$\sum_{i=1}^n (K_n^2)_{i,1} = \sum_{i=1}^n \sum_{j=1}^n (K_n)_{i,j} (K_n)_{j,1} = \sum_{j=1}^n (K_n)_{j,1} \sum_{i=1}^n (K_n)_{i,j}.$$

Since  $K_n$  consists of structure vectors of sample points, for some  $c > 0$ ,

$$1 - \frac{c}{n} \leq \sum_{j=1}^n (K_n)_{i,j} = \sum_{j=1}^n (K_n)_{j,i} \leq 1, \quad i = 1, \dots, n.$$

Given  $K_n$  is nonnegative,

$$\left(1 - \frac{c}{n}\right)^2 \leq \sum_{i=1}^n (K_n^2)_{i,1} = \sum_{j=1}^n (K_n)_{j,1} \sum_{i=1}^n (K_n)_{i,j} \leq 1.$$

Repeating the same discussion yields

$$\left(1 - \frac{c}{n}\right)^m \leq \sum_{i=1}^n (K_n^m)_{i,1} \leq 1.$$

Therefore,

$$\begin{aligned} \lambda \left( \frac{1}{1 - \lambda^2(1 - \frac{c}{n})^2} - \frac{\lambda}{1 - \lambda^2} \right) &\leq \sum_{j=1}^n \left[ \frac{1}{\lambda} I + K_n \right]_{j,1}^{-1} \\ &= \lambda \left( \sum_{i=0}^{\infty} (\lambda)^{2i} (K_n^{2i})_{j,1} - (\lambda)^{2i+1} (K_n^{2i+1})_{j,1} \right) \\ &\leq \lambda \left( \frac{1}{1 - \lambda^2} - \frac{\lambda}{1 - \lambda^2(1 - \frac{c}{n})^2} \right), \end{aligned}$$

where both the LHS and RHS reduce to  $\frac{\lambda}{1+\lambda} + O(\frac{1}{n})$ . So is true for any column sum of  $[\frac{1}{\lambda} I + K_n]^{-1}$ . Now given  $k_n$  is nonnegative and  $1 - \|k_n\|_1 = O(\frac{1}{n})$  we reach the assertion.

To show (ii), under locality,  $k_{nj} = 0$  if  $|x_i - x_j| > d_n$ , while  $[K_n]_{i,j} = 0$  if  $|x_i - x_j| > d_n$ . Recursively, if  $|x_i - x_j| > l_n \cdot d_n$  then  $[K_n^l]_{i,j} = 0$  for  $l \leq l_n$ . As  $k_n$  and  $K_n$  are element-wisely

nonnegative, we again expand the matrix inverse

$$\begin{aligned}
 \left\| r_n|_{D_n^c} \right\|_1 &= \sum_{|x-x_i|>l_n \cdot d_n} |r_{ni}| = \sum_{|x-x_i|>l_n \cdot d_n} \left| \sum_j k_{nj} \left[ \frac{1}{\lambda} I + K_n \right]_{j,i}^{-1} \right| \\
 &= \sum_{|x-x_i|>l_n \cdot d_n} \left| \sum_{|x-x_j| \leq d_n} k_{nj} \left[ \frac{1}{\lambda} I + K_n \right]_{j,i}^{-1} \right| \\
 &\leq \sum_{|x-x_j| \leq d_n} k_{nj} \sum_{|x-x_i|>l_n \cdot d_n} \left| \left[ \frac{1}{\lambda} I + K_n \right]_{j,i}^{-1} \right| \\
 &\leq \sum_{|x-x_j| \leq d_n} k_{nj} \sum_{|x_i-x_j|>(l_n-1) \cdot d_n} \left| \left[ \frac{1}{\lambda} I + K_n \right]_{j,i}^{-1} \right| \\
 &\leq \sum_{|x-x_j| \leq d_n} k_{nj} \sum_{|x_i-x_j|>(l_n-1) \cdot d_n} \lambda \sum_{l=l_n}^{\infty} \lambda^l [K_n^l]_{j,i} \\
 &\leq \sum_{|x-x_j| \leq d_n} k_{nj} \sum_{l=l_n}^{\infty} \lambda^{l+1} \\
 &\leq \sum_{l=l_n}^{\infty} \lambda^{l+1} = \frac{\lambda}{1-\lambda} \frac{1}{n}.
 \end{aligned}$$

■

## A.9 Proof of Lemma 8

**Proof** The idea is to bound  $k_{nj}$  from both above and below. The condition

$$\inf_{A \in \mathcal{q} \in \mathcal{Q}_n} \sum_{i=1}^n I(x_i \in A) = \Omega \left( n^{\frac{1}{d+2}} \right)$$

implies that  $k_{nj} = O \left( n^{-\frac{1}{d+2}} \right)$ . Given  $\|k_n\|_1 \leq 1$ ,

$$\|k_n\| \leq \sqrt{\|k_n\|_1 \|k_n\|_\infty} = O \left( n^{-\frac{1}{2} \frac{1}{d+2}} \right).$$

On the other hand, given  $|B_n| = \Omega(n \cdot d_n^d)$ , there are at most

$$\Omega \left( n \cdot d_n^d \right) = \Omega \left( n^{\frac{1}{d+1}} \right)$$

$k_{nj}$ 's that are positive. Since  $\|k_n\|_1 = 1 - O(n^{-1})$ ,

$$\|k_n\| = \Omega \left( \sqrt{\left( n^{-\frac{1}{d+1}} \right)^2 \cdot n^{\frac{1}{d+1}}} \right) = \Omega \left( n^{-\frac{1}{2} \frac{1}{d+1}} \right).$$

Those bounds also work for  $\|r_n\|$  given

$$\frac{\lambda}{1+\lambda} \leq \text{eigen} \left( \left[ \frac{1}{\lambda} I + K_n \right]^{-1} \right) \leq \lambda.$$

■

### A.10 Proof of Theorem 9

We first introduce a lemma regarding  $|B_n|$  and  $|D_n|$ . They follow Binomial distributions with parameters depending on the local covariate density.

**Lemma 17** *Assume  $X_1, X_2, \dots$  independent binomial random variables s.t.  $X_n \sim \text{Binom}(n, p_n)$  and  $np_n \rightarrow \infty$ .*

$$\frac{X_n - np_n}{\sqrt{np_n(1-p_n)}} \xrightarrow{d} N(0, 1).$$

The proof of Theorem 9 follows.

**Proof** Notice that

$$\hat{f}_n(x) - r_n^T f(X_n) = r_n^T \vec{\epsilon}_n.$$

To obtain a central limit theorem we check the Lindeberg-Feller condition of  $r_n^T \vec{\epsilon}_n$ , i.e. for any  $\delta > 0$ ,

$$\lim_n \frac{1}{\|r_n\|^2 \sigma_\epsilon^2} \sum_{i=1}^n \mathbb{E} \left[ (r_{ni} \epsilon_i)^2 I(|r_{ni} \epsilon_i| > \delta \|r_n\| \sigma_\epsilon) \right] = 0.$$

Since  $\|k_n\|_\infty = O\left(n^{-\frac{1}{d+2}}\right)$  and  $\left[\frac{1}{\lambda} I + K_n\right]^{-1}$  having row sums of  $\frac{\lambda}{1+\lambda} + O(n^{-1})$ , we have

$$\|r_n\|_\infty \leq \|k_n\|_\infty \cdot \left\| \left[ \frac{1}{\lambda} I + K_n \right]^{-1} \right\|_1 = O\left(n^{-\frac{1}{d+2}}\right).$$

Furthermore, since  $\|r_n\| = \Omega\left(n^{-\frac{1}{2} \frac{1}{d+1}}\right)$ , we get

$$\frac{\|r_n\|_\infty}{\|r_n\|} = O\left(n^{-\frac{1}{d+2} + \frac{1}{2} \frac{1}{d+1}}\right),$$

which justifies the Lindeberg-Feller condition when  $\epsilon$  is sub-Gaussian by

$$\begin{aligned} \sum_{i=1}^n \mathbb{E} \left[ (r_{ni} \epsilon_i)^2 I(|r_{ni} \epsilon_i| > \delta \|r_n\| \sigma_\epsilon) \right] &\leq \sum_{i=1}^n r_{ni}^2 \sqrt{\mathbb{E}[\epsilon_i^4] \cdot \mathbb{E}[I(|r_{ni} \epsilon_i| > \delta \|r_n\| \sigma_\epsilon)^2]} \\ &\leq \sum_{i=1}^n r_{ni}^2 \sqrt{\mathbb{E}[\epsilon_i^4]} \cdot \sqrt{P\left(|\epsilon_i| \geq \frac{\delta \|r_n\| \sigma_\epsilon}{r_{ni}}\right)} \\ &\leq \sum_{i=1}^n r_{ni}^2 \sqrt{\mathbb{E}[\epsilon_i^4]} \sqrt{2 \exp\left(-\frac{1}{2\sigma_\epsilon^2} \cdot \left(\frac{\delta \|r_n\| \sigma_\epsilon}{r_{ni}}\right)^2\right)} \\ &\leq \|r_n\|^2 \exp\left(-O\left(n^{\frac{2}{d+2} - \frac{1}{d+1}}\right)\right) \rightarrow 0, \end{aligned}$$

since

$$P(\epsilon > t) \leq \exp\left(-\frac{t^2}{2\sigma_\epsilon^2}\right)$$

for sub-Gaussian  $\epsilon$ . ■

### A.11 Kolmogorov's Extension Theorem

Define  $(x_1, \dots) = \mathbf{X} \in [0, 1]^{d \times \mathbb{N}}$  and  $\epsilon = (\epsilon_1, \dots) \in \mathbb{R}^{\mathbb{N}}$ , where the probability measures on  $[0, 1]^{d \times \mathbb{N}}$  and  $\mathbb{R}^{\mathbb{N}}$  are uniquely decided by the product measures on the cylinder spaces reflecting i.i.d. sampling i.e.  $y_i = f(x_i) + \epsilon_i$  for  $i \in \mathbb{N}$ . Write  $\pi_i$  the cumulative coordinate projection, i.e.  $\pi_i(a_1, \dots, a_n, \dots) = (a_1, \dots, a_i)$ . We can calculate  $k_n$  and  $K_n$  with respect to  $\Pi_n = (\pi_n(\mathbf{X}), \pi_n(\epsilon))$ . Thus

$$\rho_n(\mathbf{X}, \epsilon) = \frac{\hat{f}_n(x; \Pi_n) - k_n^T(x; \Pi_n)[\frac{1}{\lambda}I + K_n(\Pi_n)]^{-1}f(\Pi_n)}{\|k_n(x; \Pi_n)^T[\frac{1}{\lambda}I + K_n(\Pi_n)]^{-1}\|}$$

reflects the prediction after using a random sample of size  $n$ . Using Lemma 18, CLT of  $\rho_n$  requests an almost surely claim of Theorem 9 where the sequence of  $(x_1, y_1), \dots, (x_n, y_n)$  comes from  $(\pi_n(\mathbf{X}), \pi_n(\epsilon))$ .

We also benefit from the following lemma.

**Lemma 18** *Assume  $X : \Omega_1 \rightarrow S$ , independent of  $\epsilon : \Omega_2 \rightarrow S$ ,  $\{f_n : S \times S \rightarrow \mathbb{R}\}$  sequence of measurable functions. Assuming for a.s.  $x \in \Omega_1$ ,*

$$f_n(x, \epsilon) \xrightarrow{d} N(0, 1).$$

Then

$$f_n(X, \epsilon) \xrightarrow{d} N(0, 1).$$

**Proof** Probabilistic DCT guarantees that

$$\begin{aligned} \lim_n P(f_n(X, \epsilon) \leq t) &= \lim_n \int \int \mathbb{1}_{\{f_n(x, \epsilon) \leq t\}} d\mu_x d\mu_\epsilon \\ &= \lim_n \int P(f_n(x, \epsilon) \leq t) d\mu_x \\ &= \int \lim_n P(f_n(x, \epsilon) \leq t) d\mu_x \\ &= \int \Phi(t) d\mu_x = \Phi(t). \end{aligned}$$

■

**A.12 Proof of Lemma 10**

**Proof** In order to prove the lemma, we combine Lemma 8, Theorem 9 and Lemma 18 and show that all assumptions are met from a point-wise perspective on  $[0, 1]^{d \times \mathbb{N}}$ , i.e. fixed sample sequence are given by  $\pi_n \mathbf{X}, n \geq 1$ .

i) We show for a.s.  $\mathbf{X}$ ,  $|B_n| = |B_n(\pi_n \mathbf{X})| = \Omega(n \cdot d_n^d)$ . Consider random  $\mathbf{X}$ . Since  $\mu(x)$  is bounded away from 0, the expected amount of sample points in  $B_n$  should be at least proportional to its volume  $d_n^d$ . Define  $a_n$  s.t.  $na_n = \mathbb{E}[|B_n|] = \Omega(nd_n^d)$ . Considering the multiplicative Chernoff Bound applied to binomial random variables. For fixed  $0 < c < 1$ ,

$$P(|B_n| \leq c \cdot na_n) \leq \exp\left(-\frac{(1-c)^2 na_n}{2}\right).$$

Further, since  $na_n = \Omega(nd_n^d) = \Omega\left(n^{\frac{1}{d+1}}\right)$ , there exists some  $M > 0$  s.t.

$$\begin{aligned} \sum_{n=1}^{\infty} \exp\left(-\frac{(1-c)^2 na_n}{2}\right) &\leq \int_{n=0}^{\infty} \exp\left(-\frac{(1-c)^2 M}{2} \cdot n^{\frac{1}{d+1}}\right) dn \\ &\leq \int_{u=0}^{\infty} \exp\left(-\frac{(1-c)^2 M}{2} \cdot u\right) (d+1)u^d du \\ &< \infty. \end{aligned}$$

As per Borel-Contelli, since

$$\sum_{n=1}^{\infty} P(|B_n(\pi_n \mathbf{X})| \leq c \cdot na_n) < \infty,$$

then for a.s.  $\mathbf{X}$ , events of  $|B_n(\pi_n \mathbf{X})| \leq c \cdot na_n$  happens finite times, which leads to our conclusion.

ii) To show

$$\inf_{A \in \mathcal{Q}_n} \sum_{i=1}^n I(x_i \in A) = \Omega\left(n^{\frac{1}{d+2}}\right)$$

for a.s.  $\mathbf{X}$ , apply the concentration bound again to get:

$$\begin{aligned} P\left(\exists A \in \mathcal{Q}_n \text{ s.t. } \sum_{i=1}^n I(x_i \in A) \leq n^{\frac{1}{d+2}}\right) &\leq |\mathcal{Q}_n| \cdot n^{\frac{d+1}{d+2}} \cdot P\left(\sum_{i=1}^n I(x_i \in A) \leq n^{\frac{1}{d+2}}\right) \\ &\leq |\mathcal{Q}_n| \cdot n^{\frac{d+1}{d+2}} \cdot \exp\left(-\frac{n^{\frac{1}{d+2}+\nu} \cdot n^{-2\nu}}{2}\right) \\ &\leq \frac{1}{n} \exp\left(\frac{1}{2}n^{\frac{1}{d+2}-\nu} - n^\alpha\right) \cdot n \cdot \exp\left(-\frac{1}{2}n^{\frac{1}{d+2}-\nu}\right) \\ &\leq \exp(-n^\alpha). \end{aligned}$$

Therefore, noticing that for  $\alpha > 0$ ,

$$\sum_{n=1}^{\infty} \exp(-n^\alpha) < \infty,$$

the Borel-Cantelli theorem indicates our assertion. Hence, for a.s.  $\mathbf{X}$ , the triangular array  $\pi_n \mathbf{X}$  satisfies the assumptions in Theorem 9.  $\blacksquare$

### A.13 Proof of Theorem 11 (Main Theorem)

**Proof** We first show that for given  $x \in [0, 1]^d$ ,

$$\frac{r_n^T f(X_n) - \frac{\lambda}{1+\lambda} f(x)}{\|r_n\|} \xrightarrow{p} 0.$$

Recall the index set  $D_n = \{i : |x_i - x| \leq l_n \cdot d_n\}$ . Denote  $\Delta = \frac{\lambda}{1+\lambda} - \sum_{i=1}^n r_{n,i} = O(n^{-1})$  and  $\tilde{f}(x) = (f(x), \dots, f(x))^T$  an  $n$ -vector. We split

$$\begin{aligned} \frac{r_n^T f(X_n) - \frac{\lambda}{1+\lambda} f(x)}{\|r_n\|} &= \frac{r_n^T [f(X_n) - \tilde{f}(x)]}{\|r_n\|} - \frac{\Delta \cdot f(x)}{\|r_n\|} \\ &= -\frac{\Delta \cdot f(x)}{\|r_n\|} + \frac{\left(r_n|_{D_n}\right)^T [f(X_n) - \tilde{f}(x)]|_{D_n}}{\|r_n\|} + \frac{\left(r_n|_{D_n^c}\right)^T [f(X_n) - \tilde{f}(x)]|_{D_n^c}}{\|r_n\|}. \end{aligned}$$

By replacing  $O(\cdot)$  in the fixed case by  $O_p(\cdot)$  in the random design case, recall that

$$\Omega_p\left(n^{-\frac{1}{2} \frac{1}{d+1}}\right) = \|k_n\|, \|r_n\| = O_p\left(n^{-\frac{1}{2} \frac{1}{d+2}}\right).$$

On one hand, we notice that

$$\left| \left(r_n|_{D_n^c}\right)^T [f(X_n) - \tilde{f}(x)]|_{D_n^c} \right| \leq \left\| r_n|_{D_n^c} \right\|_1 \cdot \left\| [f(X_n) - \tilde{f}(x)]|_{D_n^c} \right\|_\infty = O_p\left(\frac{1}{n} \cdot 2M_f\right) = O_p(n^{-1}).$$

Therefore

$$\frac{\left(r_n|_{D_n^c}\right)^T [f(X_n) - \tilde{f}(x)]|_{D_n^c}}{\|r_n\|} \xrightarrow{p} 0.$$

And similarly since  $|\Delta| = O(n^{-1})$ ,

$$\frac{\Delta \cdot f(x)}{\|r_n\|} \xrightarrow{p} 0.$$

On the other hand, we can show similarly as  $|B_n|$  that  $|D_n| = \Omega_p(n \cdot (l_n \cdot d_n)^d)$  a.s. and therefore

$$\begin{aligned}
\frac{\left| \left( r_n|_{D_n} \right)^T [f(X_n) - \tilde{f}(x)]|_{D_n} \right|}{\|r_n\|} &\leq \frac{\|r_n|_{D_n}\| \cdot \|[f(X_n) - \tilde{f}(x)]|_{D_n}\|}{\|r_n\|} \\
&\leq \|[f(X_n) - \tilde{f}(x)]|_{D_n}\| \\
&= O_p\left(\sqrt{n \cdot (l_n d_n)^d \cdot (l_n d_n \cdot \alpha)^2}\right) \\
&= O_p\left(\sqrt{n \cdot \log_n^{d+2} \cdot d_n^{d+2}}\right) \\
&= O_p\left(\sqrt{n \cdot \log_n^{d+2} \cdot n^{-\frac{d+2}{d+1}}}\right) \\
&= O_p\left((\log n)^{\frac{d+2}{2}} n^{-\frac{1}{2} \frac{1}{d+1}}\right).
\end{aligned}$$

Therefore

$$\frac{\left( r_n|_{D_n} \right)^T [f(X_n) - \tilde{f}(x)]|_{D_n}}{\|r_n\|} \xrightarrow{p} 0.$$

Combining the above calculations gives the result that

$$\frac{r_n^T f(X_n) - \frac{\lambda}{1+\lambda} f(x)}{\|r_n\|} \xrightarrow{p} 0.$$

Therefore by Slutsky's Theorem,

$$\frac{\hat{f}_n(x) - \frac{\lambda}{1+\lambda} f(x)}{\|r_n\|} = \frac{\hat{f}_n(x) - r_n^T f(X_n)}{\|r_n\|} + \frac{r_n^T f(X_n) - \frac{\lambda}{1+\lambda} f(x)}{\|r_n\|} \xrightarrow{d} N(0, \sigma_\epsilon^2).$$

■

#### A.14 Proof of Theorem 13

**Proof** We refer to Theorem 14. Choose  $\delta = r$ , and choose  $T$  s.t.  $\forall t > T$ ,

$$\frac{\lambda}{t}(1 + \sqrt{n})2\sqrt{n}M \leq \frac{r}{\sqrt{d}}, \text{ i.e. } \sup \|\epsilon_t\| \leq \frac{r}{\sqrt{d}},$$

In this case,  $\beta = \|\hat{Y}_t\| + \delta - \sqrt{d} \sup_{t \geq T} \|\epsilon_t\| \geq \delta = r$ . By the conditional independence of  $\hat{Y}_t$  and  $\epsilon_b, b > t$  in the contraction region,

$$\begin{aligned}
P\left(\hat{Y}_b \in C, \forall b \geq t | \hat{Y}_t \in B(Y^*, r)\right) &\geq P\left(\sup_{b>t} \|\hat{Y}_b - Y^*\| \leq \|\hat{Y}_t - Y^*\| + \delta | \hat{Y}_t \in B(Y^*, r)\right) \\
&= P\left(\sup_{b>t} \|\hat{Y}_b - Y^*\| \leq \|\hat{Y}_t - Y^*\| + \delta\right) \\
&\geq 1 - \frac{4\sqrt{d} \sum_{b=t+1}^{\infty} \mathbb{E}[\epsilon_b^2]}{r^2} \rightarrow 1.
\end{aligned}$$



## Appendix B. On Empirical Study

### B.1 Empirical Study Details

- Figure 9 shows the distribution of true responses for our simulated training set.
- The 4 test points used for comparing Boulevard and kernel ridge regression are:  $(0.1, 0.1, 0.1, 0.1, 0.1)$ ,  $(0.6, 0.9, 0.8, 0.9, 0.7)$ ,  $(0.1, 0.1, 0.9, 0.9, 0.9)$ ,  $(0.9, 0.1, 0.1, 0.1, 0.9)$ .
- The 10 test points used for showing limiting distributions and generating reproduction intervals are:  $(0.5, 0.5, 0.5, 0.5, 0.5)$ ,  $(0.2, 0.2, 0.2, 0.2, 0.2)$ ,  $(0.1, 0.9, 0.1, 0.9, 0.1)$ ,  $(0.1, 0.1, 0.9, 0.9, 0.9)$ ,  $(0.9, 0.1, 0.1, 0.1, 0.9)$ ,  $(0.5, 0.1, 0.9, 0.1, 0.5)$ ,  $(0.3, 0.2, 0.7, 0.8, 0.6)$ ,  $(0.4, 0.2, 0.3, 0.6, 0.7)$ ,  $(0.2, 0.7, 0.8, 0.3, 0.5)$ ,  $(0.3, 0.6, 0.4, 0.9, 0.5)$ .
- Table 1 shows the settings we use in studies. The labels are: MSE for Figure 1 and 2, MSE-DatasetName for Figure 4, Limiting for Figure 6, Variance for Table 2 and RI for Figure 7. Abbreviations:  $n$  for sample size,  $\theta$  for subsample rate, ntree for ensemble size,  $k$  for terminal leaf size after subsample (which has to be corrected when no subsample is involved, i.e. GBDT),  $\lambda$  as in Boulevard iterations.

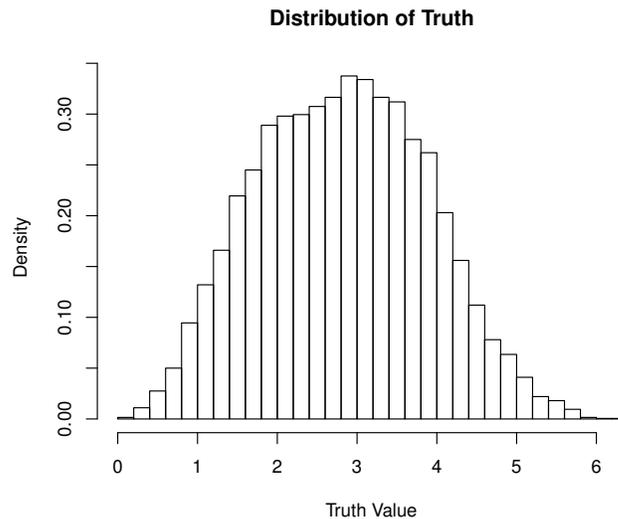


Figure 9: Distribution of truth used when assessing limiting distribution and reproduction interval.

Table 1: Parameters used in empirical study.

label	n	$\theta$	ntree	k	$\lambda$
MSE-(1-4)	5000	0.3	1000	20	0.8
MSE-Boston	506	0.8	1000	5	0.8
MSE-CCPP	9568	0.5	1000	50	0.8
MSE-CASP	20000	0.5	1000	50	0.8
MSE-Airfoil	1503	0.8	1000	40	0.8
Limiting-(1-4)	1000	0.8	2000	10	0.5
Variance-(1-4)	5000	0.8	3000	20	0.5
RI-(1-2)	1000	0.8	2000	10	0.5
RI-(3-4)	5000	0.8	2000	10	0.5

Table 2: Prediction standard deviations scale with error standard deviations.

Error\Fixed Point	1	2	3	4	5	6	7	8	9	10
0	0.030	0.044	0.044	0.049	0.050	0.037	0.038	0.033	0.032	0.040
Unif[-1,1]	0.067	0.089	0.096	0.087	0.096	0.083	0.081	0.074	0.071	0.082
Unif[-2,2]	0.119	0.154	0.172	0.158	0.162	0.152	0.122	0.139	0.137	0.145
Unif[-4,4]	0.243	0.271	0.278	0.278	0.288	0.317	0.284	0.289	0.318	0.254

## B.2 Scaling of Prediction Error with Response Error

Table 2 shows the experiment in which we apply symmetric uniform errors and observe the scaling of prediction standard deviation along with the increase of error standard deviation.

## References

- Anthony Almudevar. A stochastic contraction mapping theorem. [arXiv preprint arXiv:2207.00618](#), 2022.
- Naomi S Altman. An introduction to kernel and nearest-neighbor nonparametric regression. [The American Statistician](#), 46(3):175–185, 1992.
- Sylvain Arlot and Robin Genuer. Analysis of purely random forests bias. [arXiv preprint arXiv:1407.3939](#), 2014.
- S. Athey, J. Tibshirani, and S. Wager. Generalized Random Forests. [ArXiv e-prints](#), October 2016.
- Peter Bartlett, Yoav Freund, Wee Sun Lee, and Robert E Schapire. Boosting the margin: A new explanation for the effectiveness of voting methods. [The annals of statistics](#), 26(5):1651–1686, 1998.
- G erard Biau. Analysis of a random forests model. [Journal of Machine Learning Research](#), 13(Apr):1063–1095, 2012.

- G erard Biau and Erwan Scornet. A random forest guided tour. Test, 25(2):197–227, 2016.
- G erard Biau, Luc Devroye, and G abor Lugosi. Consistency of random forests and other averaging classifiers. Journal of Machine Learning Research, 9(Sep):2015–2033, 2008.
- Leo Breiman. Random forests. Machine learning, 45(1):5–32, 2001.
- Leo Breiman. Consistency for a simple model of random forests. statistical department. University of California at Berkeley. Technical Report,(670), 2004.
- Leo Breiman, Jerome Friedman, Charles J Stone, and Richard A Olshen. Classification and regression trees. CRC press, 1984.
- Peter B uhlmann and Torsten Hothorn. Boosting algorithms: Regularization, prediction and model fitting. Statistical Science, pages 477–505, 2007.
- Peter B uhlmann and Bin Yu. Boosting with the  $l_2$  loss: regression and classification. Journal of the American Statistical Association, 98(462):324–339, 2003.
- Peter B uhlmann, Bin Yu, et al. Analyzing bagging. The Annals of Statistics, 30(4):927–961, 2002.
- Peter Lukas B uhlmann. Consistency for  $l_2$ boosting and matching pursuit with trees and tree-type basis functions. In Research report/Seminar f ur Statistik, Eidgen ossische Technische Hochschule (ETH), volume 109. Seminar f ur Statistik, Eidgen ossische Technische Hochschule (ETH), 2002.
- Hugh A. Chipman, Edward I. George, and Robert E. McCulloch. Bart: Bayesian additive regression trees. Ann. Appl. Stat., 4(1):266–298, 03 2010. doi: 10.1214/09-AOAS285. URL <https://doi.org/10.1214/09-AOAS285>.
- Alex Davies and Zoubin Ghahramani. The random forest kernel and other kernels for big data from random partitions. arXiv preprint arXiv:1402.4293, 2014.
- Dua Dheeru and Efi Karra Taniskidou. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. In European conference on computational learning theory, pages 23–37. Springer, 1995.
- Yoav Freund, Robert Schapire, and Naoki Abe. A short introduction to boosting. Journal-Japanese Society For Artificial Intelligence, 14(771-780):1612, 1999.
- Jerome Friedman, Trevor Hastie, Robert Tibshirani, et al. Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). The annals of statistics, 28(2):337–407, 2000.
- Jerome H Friedman. Greedy function approximation: a gradient boosting machine. Annals of statistics, pages 1189–1232, 2001.

- Jerome H Friedman. Stochastic gradient boosting. Computational Statistics & Data Analysis, 38(4):367–378, 2002.
- Louis Gordon and Richard A Olshen. Almost surely consistent nonparametric regression from recursive partitioning schemes. Journal of Multivariate Analysis, 15(2):147–163, 1984.
- Heysem Kaya, Pmar Tüfekci, and Fikret S Gürgen. Local and global learning methods for predicting power of a combined gas & steam turbine. In Proceedings of the International Conference on Emerging Trends in Computer and Electronics Engineering ICETCEE, pages 13–18, 2012.
- Stéphane G Mallat and Zhifeng Zhang. Matching pursuits with time-frequency dictionaries. IEEE Transactions on signal processing, 41(12):3397–3415, 1993.
- Lucas Mentch and Giles Hooker. Quantifying uncertainty in random forests via confidence intervals and hypothesis tests. The Journal of Machine Learning Research, 17(1):841–881, 2016.
- Lucas Mentch and Giles Hooker. Formal hypothesis tests for additive structure in random forests. Journal of Computational and Graphical Statistics, 26(3):589–597, 2017. doi: 10.1080/10618600.2016.1256817. URL <https://doi.org/10.1080/10618600.2016.1256817>.
- Indraneel Mukherjee, Cynthia Rudin, and Robert E Schapire. The rate of convergence of adaboost. Journal of Machine Learning Research, 14:2315–2347, 2013.
- Wei Peng, Tim Coleman, and Lucas Mentch. Asymptotic distributions and rates of convergence for random forests and other resampled ensemble learners. arXiv preprint arXiv:1905.10651, 2019.
- Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. SIAM Journal on Control and Optimization, 30(4):838–855, 1992.
- KV Rashmi and Ran Gilad-Bachrach. Dart: Dropouts meet multiple additive regression trees. arXiv preprint arXiv:1505.01866, 2015.
- Herbert Robbins and Sutton Monro. A stochastic approximation method. The annals of mathematical statistics, pages 400–407, 1951.
- Alex Rogozhnikov and Tatiana Likhomanenko. Infiniteboost: building infinite ensembles with gradient descent. arXiv preprint arXiv:1706.01109, 2017.
- Cynthia Rudin, Ingrid Daubechies, Robert E Schapire, and Dana Ron. The dynamics of adaboost: cyclic behavior and convergence of margins. Journal of Machine Learning Research, 5(10), 2004.
- David Ruppert. Efficient estimations from a slowly convergent robbins-monro process. Technical report, Cornell University Operations Research and Industrial Engineering, 1988.

- Erwan Scornet. Random forests and kernel methods. IEEE Transactions on Information Theory, 62(3):1485–1500, 2016.
- Erwan Scornet, Gérard Biau, and Jean-Philippe Vert. Consistency of random forests. Ann. Statist., 43(4):1716–1741, 08 2015. doi: 10.1214/15-AOS1321. URL <https://doi.org/10.1214/15-AOS1321>.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. The Journal of Machine Learning Research, 15(1):1929–1958, 2014.
- Charles J Stone. Consistent nonparametric regression. The annals of statistics, pages 595–620, 1977.
- Charles J Stone. Optimal global rates of convergence for nonparametric regression. The annals of statistics, pages 1040–1053, 1982.
- Pınar Tüfekci. Prediction of full load electrical power output of a base load operated combined cycle power plant using machine learning methods. International Journal of Electrical Power & Energy Systems, 60:126–140, 2014.
- Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. Journal of the American Statistical Association, (just-accepted), 2017.
- Stefan Wager and Guenther Walther. Adaptive concentration of regression trees, with application to random forests. arXiv preprint arXiv:1503.06388, 2015.
- Stefan Wager, Sida Wang, and Percy S Liang. Dropout training as adaptive regularization. In Advances in neural information processing systems, pages 351–359, 2013.
- Stefan Wager, Trevor Hastie, and Bradley Efron. Confidence intervals for random forests: The jackknife and the infinitesimal jackknife. The Journal of Machine Learning Research, 15(1):1625–1651, 2014.
- Tong Zhang, Bin Yu, et al. Boosting with early stopping: Convergence and consistency. The Annals of Statistics, 33(4):1538–1579, 2005.