



On the asymptotics of random forests

Erwan Scornet

► To cite this version:

| Erwan Scornet. On the asymptotics of random forests. 2014. hal-01061506

HAL Id: hal-01061506

<https://hal.science/hal-01061506>

Preprint submitted on 7 Sep 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the asymptotics of random forests

Erwan Scornet

Sorbonne Universités, UPMC Univ Paris 06, F-75005, Paris, France

erwan.scornet@upmc.fr

Abstract

The last decade has witnessed a growing interest in random forest models which are recognized to exhibit good practical performance, especially in high-dimensional settings. On the theoretical side, however, their predictive power remains largely unexplained, thereby creating a gap between theory and practice. The aim of this paper is twofold. Firstly, we provide theoretical guarantees to link finite forests used in practice (with a finite number M of trees) to their asymptotic counterparts (with $M = \infty$). Using empirical process theory, we prove a uniform central limit theorem for a large class of random forest estimates, which holds in particular for Breiman’s original forests. Secondly, we show that infinite forest consistency implies finite forest consistency and thus, we state the consistency of several infinite forests. In particular, we prove that q quantile forests—close in spirit to Breiman’s forests but easier to study—are able to combine inconsistent trees to obtain a final consistent prediction, thus highlighting the benefits of random forests compared to single trees.

Index Terms — Random forests, randomization, consistency, central limit theorem, empirical process, number of trees, q -quantile.

2010 Mathematics Subject Classification: 62G05, 62G20.

1 Introduction

Random forests are a class of algorithms used to solve classification and regression problems. As ensemble methods, they grow several trees as base estimates and aggregate them to make a prediction. In order to obtain many different trees based on a single training set, random forests procedures introduce randomness in the tree construction. For instance, trees can be built by randomizing the set of features (Dietterich and Kong, 1995; Ho, 1998), the data set (Breiman, 1996, 2000), or both at the same time (Breiman, 2001; Cutler and Zhao, 2001).

Among all random forest algorithms, the most popular one is that of Breiman (2001), which relies on CART procedure (Classification and Regression Trees, Breiman et al., 1984) to grow the individual trees. As highlighted by many applied studies (see, e.g., Hamza and Laroque, 2005; Díaz-Uriarte and de Andrés, 2006), Breiman’s (2001) random forests often outperform state-of-the-art methods. They are recognized for their ability to handle high-dimensional data sets, thus being useful in fields such as genomics (Qi, 2012) and pattern recognition (Rogez et al., 2008), just to name a few. On the computational side, Breiman’s (2001) forests are easy to run and robust to changes in the parameters they depend on (Liaw and Wiener, 2002; Genuer et al., 2008). As a proof of their success, many extensions have been developed in ranking problems (Cléménçon

et al., 2013), quantile estimation (Meinshausen, 2006), and survival analysis (Ishwaran et al., 2008). Interesting new developments in the context of massive data sets have been achieved. For instance, Geurts et al. (2006) modified the procedure to reduce calculation time, while other authors extended the procedure to online settings (Denil et al., 2013; Lakshminarayanan et al., 2014, and the reference therein).

While Breiman’s (2001) forests are extensively used in practice, some of their mathematical properties remain under active investigation. In fact, most theoretical studies focus on simplified versions of the algorithm, where the forest construction is independent of the training set. Consistency of such simplified models has been proved (e.g., Biau et al., 2008; Ishwaran and Kogalur, 2010; Denil et al., 2013). However, these results do not extend to Breiman’s original forests whose construction critically depends on the whole training set. Recent attempts to bridge the gap between theoretical forest models and Breiman’s (2001) forests have been made by Wager (2014) and Scornet et al. (2014) who establish consistency of the original algorithm under suitable assumptions.

Apart from the dependence of the forest construction on the data set, there is another fundamental difference between existing forest models and ones implemented. Indeed, in practice, a forest can only be grown with a finite number M of trees although most theoretical works assume, by convenience, that $M = \infty$. Since the predictor with $M = \infty$ does not depend on the specific tree realizations that form the forest, it is therefore more amenable to analysis. However, surprisingly, no study aims at clarifying the link between finite forests (finite M) and infinite forests ($M = \infty$) even if some authors (Mentch and Hooker, 2014; Wager et al., 2014) proved results on finite forest predictions at a fixed point \mathbf{x} .

In the present paper, our goal is to study the connection between infinite forest models and finite forests used in practice in the context of regression. We start by proving a uniform central limit theorem for various random forests estimates, including Breiman’s (2001) ones. In Section 3, we also point out that the \mathbb{L}^2 risk of infinite forests is lower than that of finite forests, which supports the interest of theoretical studies for infinite forests. Besides, this result shows that infinite forest consistency implies finite forest consistency. Finally, in Section 4, we prove the consistency of several infinite random forests. In particular, taking one step toward the understanding of Breiman’s (2001) forests, we prove that q quantile forests, a variety of forests whose construction depends on the positions \mathbf{X}_i ’s of the data, are consistent. As for Breiman’s forests, each leaf of each tree in q quantile forests contains a small number of points that does not grow to infinity with the sample size. Thus, q quantile forests average inconsistent trees estimate to build a consistent prediction.

We start by giving some notation in Section 2. All proofs are postponed to Section 5.

2 Notation

Throughout the paper, we assume to be given a training sample $\mathcal{D}_n = (\mathbf{X}_1, Y_1), \dots, (\mathbf{X}_n, Y_n)$ of $[0, 1]^d \times \mathbb{R}$ -valued independent random variables distributed as

the prototype pair (\mathbf{X}, Y) , where $\mathbb{E}[Y^2] < \infty$. We aim at predicting the response Y , associated with the random variable \mathbf{X} , by estimating the regression function $m(\mathbf{x}) = \mathbb{E}[Y|\mathbf{X} = \mathbf{x}]$. In this context, we use random forests to build an estimate $m_n : [0, 1]^d \rightarrow \mathbb{R}$ of m , based on the data set \mathcal{D}_n .

A random forest is a collection of M randomized regression trees (for an overview on tree construction, see Chapter 20 in [Györfi et al., 2002](#)). For the j -th tree in the family, the predicted value at point \mathbf{x} is denoted by $m_n(\mathbf{x}, \Theta_j, \mathcal{D}_n)$, where $\Theta_1, \dots, \Theta_M$ are independent random variables, distributed as a generic random variable Θ , independent of the sample \mathcal{D}_n . This random variable can be used to sample the training set or to select the candidate directions or positions for splitting. The trees are combined to form the finite forest estimate

$$m_{M,n}(\mathbf{x}, \Theta_1, \dots, \Theta_M) = \frac{1}{M} \sum_{m=1}^M m_n(\mathbf{x}, \Theta_m). \quad (1)$$

By the law of large numbers, for any fixed \mathbf{x} , conditionally on \mathcal{D}_n , the finite forest estimate tends to the infinite forest estimate

$$m_{\infty,n}(\mathbf{x}) = \mathbb{E}_{\Theta} [m_n(\mathbf{x}, \Theta)].$$

The risk of $m_{\infty,n}$ is defined by

$$R(m_{\infty,n}) = \mathbb{E}[m_{\infty,n}(\mathbf{X}) - m(\mathbf{X})]^2, \quad (2)$$

while the risk of $m_{M,n}$ equals

$$R(m_{M,n}) = \mathbb{E}[m_{M,n}(\mathbf{X}, \Theta_1, \dots, \Theta_M) - m(\mathbf{X})]^2. \quad (3)$$

It is stressed that both risks $R(m_{\infty,n})$ and $R(m_{M,n})$ are deterministic since the expectation in (2) is over $\mathbf{X}, \mathcal{D}_n$, and the expectation in (3) is over $\mathbf{X}, \mathcal{D}_n$ and $\Theta_1, \dots, \Theta_M$. Throughout the paper, we say that $m_{\infty,n}$ (resp. $m_{M,n}$) is \mathbb{L}^2 consistent if $R(m_{\infty,n})$ (resp. $R(m_{M,n})$) tends to zero as $n \rightarrow \infty$.

As mentioned earlier, there is a large variety of forests, depending on how trees are grown and how the randomness Θ influences the tree construction. For instance, tree construction can be independent of \mathcal{D}_n ([Biau, 2012](#)), depend only on the \mathbf{X}_i 's ([Biau et al., 2008](#)) or depend on the whole training set ([Cutler and Zhao, 2001](#); [Geurts et al., 2006](#); [Zhu et al., 2012](#)). Throughout the paper, we use Breiman's forests and uniform forests to exemplify our results. In Breiman's original procedure, splits depend on the whole sample and are performed to minimize variance within the two resulting cells. The algorithm stops when each cell contains less than a small pre-specified number of points (typically, 5 in regression and 1 in classification). On the other hand, uniform forests are a simpler procedure since, at each node, a coordinate is uniformly selected among $\{1, \dots, d\}$ and a split position is uniformly chosen in the range of the cell, along the pre-chosen coordinate. The algorithm stops when a full binary tree of level k is built, that is if each cell has been cut exactly k times, where $k \in \mathbb{N}$ is a parameter of the algorithm.

In the rest of the paper, we will repeatedly use the random forest connection function K_n , defined as

$$K_n : \begin{array}{ccc} [0, 1]^d \times [0, 1]^d & \rightarrow & [0, 1] \\ (\mathbf{x}, \mathbf{z}) & \mapsto & \mathbb{P}_\Theta \left[\mathbf{x} \overset{\Theta}{\leftrightarrow} \mathbf{z} \right], \end{array}$$

where $\mathbf{x} \overset{\Theta}{\leftrightarrow} \mathbf{z}$ is the event where \mathbf{x} and \mathbf{z} belong to the same cell in the tree $\mathcal{T}_n(\Theta)$ designed with Θ and \mathcal{D}_n . Moreover, notation \mathbb{P}_Θ denotes the probability with respect to Θ , conditionally on \mathcal{D}_n . The same notational convention holds for the expectation \mathbb{E}_Θ and the variance \mathbb{V}_Θ . Thus, if we fix the training set \mathcal{D}_n , we see that the connection $K_n(\mathbf{x}, \mathbf{z})$ is just the proportion of trees in which \mathbf{x} and \mathbf{z} are connected.

We say that a forest is discrete (resp. continuous) if, keeping \mathcal{D}_n fixed, its connection function $K_n(\bullet, \bullet)$ is piecewise constant (resp. continuous). In fact, most existing forest models fall in one of these two categories. For example, if, at each cell, the number of possible splits is finite, then the forest is discrete. This is the case of Breiman's forests, where splits can only be performed at the middle of two consecutive data points along any coordinate. However, if splits are drawn according to some density along each coordinate, the resulting forest is continuous. For instance, uniform forests are continuous.

3 Finite and infinite random forests

Contrary to finite forests which depend upon the particular Θ_j 's used to design trees, infinite forests do not and are therefore more amenable to mathematical analysis. Besides, finite forests predictions can be difficult to interpret since they depend on the random parameters Θ_j 's. In addition, the Θ_j 's are independent of the data set and thus unrelated to the particular prediction problem.

In this section, we study the link between finite forests and infinite forests. More specifically, assuming that the data set \mathcal{D}_n is fixed, we examine the asymptotic behavior of the finite forest estimate $m_{M,n}(\bullet, \Theta_1, \dots, \Theta_M)$ as M tends to infinity. This setting is consistent with practical problems, where the \mathcal{D}_n is fixed, and one can grow as many trees as possible.

Clearly, by the law of large numbers, we know that conditionally on \mathcal{D}_n , for all $\mathbf{x} \in [0, 1]^d$, almost surely,

$$m_{M,n}(\mathbf{x}, \Theta_1, \dots, \Theta_M) \xrightarrow{M \rightarrow \infty} m_{\infty,n}(\mathbf{x}). \quad (4)$$

The following theorem extend the pointwise convergence in (4) to the convergence of the whole functional estimate $m_{M,n}(\bullet, \Theta_1, \dots, \Theta_M)$, towards the functional estimate $m_{\infty,n}(\bullet)$.

Theorem 3.1. *Consider a continuous or discrete random forest. Then, conditionally on \mathcal{D}_n , almost surely, for all $\mathbf{x} \in [0, 1]^d$, we have*

$$m_{M,n}(\mathbf{x}, \Theta_1, \dots, \Theta_M) \xrightarrow{M \rightarrow \infty} m_{\infty,n}(\mathbf{x}).$$

Remark 1. *Since the set $[0, 1]^d$ is not countable, we cannot reverse the “almost sure” and “for all $\mathbf{x} \in [0, 1]^d$ ” statements in (4). Thus, Theorem 3.1 is not a consequence of (4).*

Theorem 3.1 is a first step to prove that infinite forest estimates can be uniformly approximated by finite forest estimates. To pursue the analysis, a natural question is to determine the rate of convergence in Theorem 3.1. The pointwise rate of convergence is provided by the central limit theorem which says that, conditionally on \mathcal{D}_n , for all $\mathbf{x} \in [0, 1]^d$,

$$\sqrt{M}(m_{M,n}(\mathbf{x}, \Theta_1, \dots, \Theta_M) - m_{\infty,n}(\mathbf{x})) \xrightarrow[M \rightarrow \infty]{\mathcal{L}} \mathcal{N}(0, \tilde{\sigma}^2(\mathbf{x})), \quad (5)$$

where

$$\tilde{\sigma}^2(\mathbf{x}) = \mathbb{V}_{\Theta} \left(\frac{1}{N_n(\mathbf{x}, \Theta)} \sum_{i=1}^n Y_i \mathbb{1}_{\mathbf{x} \in \mathcal{C}_{\mathbf{x}_i}} \right) \leq 4 \max_{1 \leq i \leq n} Y_i^2$$

and $N_n(\mathbf{x}, \Theta)$ is the number of data points falling into the cell of the tree $\mathcal{T}_n(\Theta)$ which contains \mathbf{x} .

Equation (5) is not sufficient to determine the asymptotic distribution of the functional estimate $m_{M,n}(\bullet, \Theta_1, \dots, \Theta_M)$. To make it explicit, we need to introduce the empirical process \mathbb{G}_M (see van der Vaart and Wellner, 1996) defined by

$$\mathbb{G}_M = \sqrt{M} \left(\frac{1}{M} \sum_{m=1}^M \delta_{\Theta_m} - \mathbb{P}_{\Theta} \right),$$

where δ_{Θ_m} is the Dirac function at Θ_m . We also let $\mathcal{F}_2 = \{g_{\mathbf{x}} : \theta \mapsto m_n(\mathbf{x}, \theta); \mathbf{x} \in [0, 1]^d\}$ be the collection of all possible tree estimates in the forest. In order to prove that a uniform central limit theorem holds for random forest estimates, we need to show that there exists a Gaussian process \mathbb{G} such that

$$\sup_{g \in \mathcal{F}_2} \left\{ \int_{\Theta} |g(\theta)| d\mathbb{G}_M(\theta) - \int_{\Theta} |g(\theta)| d\mathbb{G}(\theta) \right\} \xrightarrow[M \rightarrow \infty]{} 0, \quad (6)$$

where the first part on the left side can be written as

$$\int_{\Theta} |g(\theta)| d\mathbb{G}_M(\theta) = \sqrt{M} \left(\frac{1}{M} \sum_{m=1}^M |g(\Theta_m)| - \mathbb{E}_{\Theta} [|g(\Theta)|] \right).$$

For more clarity, instead of (6), we will write

$$\sqrt{M} \left(\frac{1}{M} \sum_{m=1}^M m_n(\bullet, \Theta_m) - \mathbb{E}_{\Theta} [m_n(\bullet, \Theta)] \right) \xrightarrow{\mathcal{L}} \mathbb{G}g_{\bullet}. \quad (7)$$

To establish identity (7), we first define, for all $\varepsilon > 0$, the random forest grid step $\delta(\varepsilon)$ by

$$\delta(\varepsilon) = \sup \left\{ \eta \in \mathbb{R} : \sup_{\substack{\mathbf{x}_1, \mathbf{x}_2 \in [0, 1]^d \\ \|\mathbf{x}_1 - \mathbf{x}_2\|_{\infty} \leq \eta}} |1 - K_n(\mathbf{x}_1, \mathbf{x}_2)| \leq \frac{\varepsilon^2}{8} \right\},$$

where K_n is the connection function of the forest. The function δ can be seen as the modulus of continuity of K_n in the sense that it is the distance such that $K_n(\mathbf{x}_1, \mathbf{x}_2)$ does not vary of much that $\varepsilon^2/8$ if $\|\mathbf{x}_1 - \mathbf{x}_2\|_{\infty} \leq \delta(\varepsilon)$. We will also need the following assumption.

(H1) One of the following properties is satisfied:

- The random forest is discrete,
- There exist $C, A > 0, \alpha < 2$ such that, for all $\varepsilon > 0$,

$$\delta(\varepsilon) \geq C \exp(-A/\varepsilon^\alpha).$$

Observe that **(H1)** is mild since most forests are discrete and the only continuous forest we have in mind, the uniform forest, satisfies **(H1)**, as stated in Lemma 1 below.

Lemma 1. *Let $k \in \mathbb{N}$. Then, for all $\varepsilon > 0$, the grid step $\delta(\varepsilon)$ of uniform forests of level k satisfies*

$$\delta(\varepsilon) \geq \exp\left(-\frac{A_{k,d}}{\varepsilon^{2/3}}\right),$$

where $A_{k,d} = (8de(k+2)!)^{1/3}$.

The following theorem states that a uniform central limit theorem is valid over the class of random forest estimates, providing that **(H1)** is satisfied.

Theorem 3.2. *Consider a random forest which satisfies **(H1)**. Then,*

$$\sqrt{M}(m_{M,n}(\bullet) - m_{\infty,n}(\bullet)) \xrightarrow{\mathcal{L}} \mathbb{G}g_{\bullet},$$

where \mathbb{G} is a Gaussian process with mean zero and a covariate function

$$\text{Cov}_{\Theta}(\mathbb{G}g_{\mathbf{x}}, \mathbb{G}g_{\mathbf{z}}) = \text{Cov}_{\Theta}\left(\sum_{i=1}^n Y_i \frac{\mathbb{1}_{\mathbf{x} \leftrightarrow \mathbf{x}_i}}{N_n(\mathbf{x}, \Theta)}, \sum_{i=1}^n Y_i \frac{\mathbb{1}_{\mathbf{z} \leftrightarrow \mathbf{x}_i}}{N_n(\mathbf{z}, \Theta)}\right).$$

According to the discussion above, Theorem 3.2 holds for uniform forests (by Lemma 1) and Breiman's forests (since they are discrete). Moreover, according to this Theorem, the finite forest estimates tend uniformly to the infinite forest estimates, with the standard rate of convergence \sqrt{M} . This result contributes to bridge the gap between finite forests used in practice and infinite theoretical forests.

The proximity between two estimates can also be measured in terms of their \mathbb{L}^2 risk. In this respect, Theorem 3.3 states that the risk of infinite forests is lower than the one of finite forests and provides a bound on the difference between these two risks. We first need an assumption on the regression model.

(H2) One has

$$Y = m(\mathbf{X}) + \varepsilon,$$

where ε is a centered Gaussian noise with finite variance σ^2 , independent of \mathbf{X} , and $\|m\|_{\infty} = \sup_{\mathbf{x} \in [0,1]^d} |m(\mathbf{x})| < \infty$.

Theorem 3.3. Assume that **(H2)** is satisfied. Then, for all $M, n \in \mathbb{N}^*$,

$$R(m_{M,n}) = R(m_{\infty,n}) + \frac{1}{M} \mathbb{E}_{\mathbf{X}, \mathcal{D}_n} \left[\mathbb{V}_{\Theta} [m_n(\mathbf{X}, \Theta)] \right].$$

In particular,

$$0 \leq R(m_{M,n}) - R(m_{\infty,n}) \leq \frac{8}{M} \times (\|m\|_{\infty}^2 + \sigma^2(1 + 4 \log n)).$$

Theorem 3.3 reveals that the prediction accuracy of infinite forests is better than that of finite forests. In practice however, there is no simple way to implement infinite forests and, in fact, finite forests are nothing but Monte Carlo approximations of infinite forests. But, since the difference of risks between both types of forests is bounded (by Theorem 3.3), the prediction accuracy of finite forests is almost as good as that of infinite forests providing the number of trees is large enough. More precisely, under **(H2)**, for all $\varepsilon > 0$, if

$$M \geq \frac{8(\|m\|_{\infty}^2 + \sigma^2)}{\varepsilon} + \frac{32\sigma^2 \log n}{\varepsilon},$$

then $R(m_{M,n}) - R(m_{\infty,n}) \leq \varepsilon$.

Another interesting consequence of Theorem 3.3 is that, assuming that **(H2)** holds and that $M/\log n \rightarrow \infty$ as $n \rightarrow \infty$, finite random forests are consistent as soon as infinite random forests are. This allows to extend all previous consistency results regarding infinite forests (see, e.g., Meinshausen, 2006; Biau et al., 2008) to finite forests. It must be stressed that the “log n ” term comes from the Gaussian noise, since, if $\varepsilon_1, \dots, \varepsilon_n$ are independent and distributed as a Gaussian noise $\varepsilon \sim \mathcal{N}(0, \sigma^2)$, we have,

$$\mathbb{E} \left[\max_{1 \leq i \leq n} \varepsilon_i^2 \right] \leq \sigma^2(1 + 4 \log n),$$

(see, e.g., Chapter 1 in Boucheron et al., 2013). Therefore, the required number of trees depends on the noise in the regression model. For instance, if Y is bounded, then the condition turns into $M \rightarrow \infty$.

4 Consistency of some random forest models

Section 3 was devoted to the connection between finite and infinite forests. In particular, we proved in Theorem 3.3 that the consistency of infinite forests implies that of finite forests, as soon as **(H2)** is satisfied and $M/\log n \rightarrow \infty$. Thus, it is natural to focus on the consistency of infinite forest estimates, which can be written as

$$m_{\infty,n}(\mathbf{X}) = \sum_{i=1}^n W_{ni}^{\infty}(\mathbf{X}) Y_i, \quad (8)$$

where

$$W_{ni}^{\infty}(\mathbf{X}) = \mathbb{E}_{\Theta} \left[\frac{\mathbb{1}_{\mathbf{X} \in \mathcal{R}_i}}{N_n(\mathbf{X}, \Theta)} \right]$$

are the random forest weights.

Proving consistency of infinite random forests is in general a difficult task, mainly because forest construction can depend on both the \mathbf{X}_i 's and the Y_i 's. This feature makes the resulting estimate highly data-dependent, and therefore difficult to analyze (this is particularly the case for Breiman's forests). To simplify the analysis, we investigate hereafter infinite random forest estimates whose weights depends only on $\mathbf{X}, \mathbf{X}_1, \dots, \mathbf{X}_n$ which is called the X -property. The good news is that when infinite forest estimates have the X -property, they fall in the general class of local averaging estimates, whose consistency can be addressed using Stone's (1977) theorem.

Therefore, using Stone's theorem as a starting point, we first prove the consistency of random forests whose construction is independent of \mathcal{D}_n , which is the simplest case of random forests satisfying the X -property. For such forests, the construction is based on the random parameter Θ only. As for now, we say that a forest is totally non adaptive of level k ($k \in \mathbb{N}$, with k possibly depending on n) if each tree of the forest is built independently of the training set and if each cell is cut exactly k times. The resulting cell containing \mathbf{X} , designed with randomness Θ , is denoted by $A_n(\mathbf{X}, \Theta)$.

Theorem 4.1. *Assume that \mathbf{X} is distributed on $[0, 1]^d$ and consider a totally non adaptive forest of level k . In addition, assume that for all $\rho, \varepsilon > 0$, there exists $N > 0$ such that, with probability $1 - \rho$, for all $n > N$,*

$$\text{diam}(A_n(\mathbf{X}, \Theta)) \leq \varepsilon.$$

Then, providing $k \rightarrow \infty$ and $2^k/n \rightarrow 0$, the infinite random forest is \mathbb{L}^2 consistent, that is

$$R(m_{\infty, n}) \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Theorem 4.1 is a generalization of some consistency results in Biau et al. (2008) for the case of totally non adaptive random forest. Together with Theorem 3.3, we see that if **(H2)** is satisfied and $M/\log n \rightarrow \infty$ as $n \rightarrow \infty$, then the finite random forest is \mathbb{L}^2 consistent.

According to Theorem 4.1, a totally non adaptive forest of level k is consistent if the cell diameters tend to zero as $n \rightarrow \infty$ and if the level k is properly tuned. This is in particular true for uniform random forests, as shown in the following corollary.

Corollary 1. *Assume that \mathbf{X} is distributed on $[0, 1]^d$ and consider a uniform forest of level k . Then, providing that $k \rightarrow \infty$ and $2^k/n \rightarrow 0$, the uniform random forest is \mathbb{L}^2 consistent.*

For totally non adaptive forests, the main difficulty that consists in using the data set to build the forest and to predict at the same time, vanishes. However, because of their simplified construction, these forests are far from accurately modelling Breiman's forest. To take one step further into the understanding of Breiman's (2001) forest behavior, we study the q ($q \in [1/2, 1)$) quantile random forest, which satisfies the X -property. Indeed, their construction depends on

the X_i 's which is a good trade off between the complexity of Breiman's forests and the simplicity of totally non adaptive forests. As an example of q quantile trees, the median tree ($q = 1/2$) has already been studied by Devroye et al. (1996), such as the k -spacing tree (Devroye et al., 1996) whose construction is based on quantiles.

In the spirit of Breiman's algorithm, before growing each tree, data are subsampled, that is a_n points ($a_n < n$) are selected without replacement. Then, each split is performed on an empirical q_n -quantile (where $q_n \in [1 - q, q]$ can be pre-specified by the user or randomly chosen) along a coordinate, chosen uniformly at random among the d coordinates. Recall that the q' -quantile ($q' \in [1 - q, q]$) of $\mathbf{X}_1, \dots, \mathbf{X}_n$ is defined as the only $\mathbf{X}_{(\ell)}$ satisfying $F_n(\mathbf{X}_{(\ell-1)}) \leq q_n < F_n(\mathbf{X}_{(\ell)})$, where the $\mathbf{X}_{(i)}$'s are ordered increasingly. Note that data points on which splits are performed are not sent down to the resulting cells. Finally, the algorithm stops when each cell contains exactly one point. The full procedure is described in **Algorithm 2**.

Algorithm 1: q quantile forest predicted value at \mathbf{x} .

Input: Fix $a_n \in \{1, \dots, n\}$, and $\mathbf{x} \in [0, 1]^d$.

Data: A training set \mathcal{D}_n .

```

1 for  $j = 1, \dots, M$  do
2   Select  $a_n$  points, without replacement, uniformly in  $\mathcal{D}_n$ .
3   Set  $\mathcal{P} = \{[0, 1]^p\}$  the partition associated with the root of the tree.
4   while there exists  $A \in \mathcal{P}$  which contains strictly more than two points do
5     Select uniformly one dimension  $j$  within  $\{1, \dots, p\}$ .
6     Let  $N$  be the number of data points in  $A$  and select  $q_n \in [1 - q, q] \cap (1/N, 1 - 1/N)$ .
7     Cut the cell  $A$  at the position given by the  $q_n$  empirical quantile (see definition below) along the  $j$ -th coordinate.
8     Call  $A_L$  and  $A_R$  the two resulting cell.
9     Set  $\mathcal{P} \leftarrow (\mathcal{P} \setminus \{A\}) \cup A_L \cup A_R$ .
10  end
11  for each  $A \in \mathcal{P}$  which contains exactly two points do
12    Select uniformly one dimension  $j$  within  $\{1, \dots, p\}$ .
13    Cut along the  $j$ -th direction, in the middle of the two points.
14    Call  $A_L$  and  $A_R$  the two resulting cell.
15    Set  $\mathcal{P} \leftarrow (\mathcal{P} \setminus \{A\}) \cup A_L \cup A_R$ .
16  end
17  Compute the predicted value  $m_n(\mathbf{x}, \Theta_j)$  at  $\mathbf{x}$  equal to the single  $Y_i$  falling in the cell of  $\mathbf{x}$ , with respect to the partition  $\mathcal{P}$ .
18 end
19 Compute the random forest estimate  $m_{M,n}(\mathbf{x}; \Theta_1, \dots, \Theta_M, \mathcal{D}_n)$  at the query point  $\mathbf{x}$  according to equality (1).

```

Since the construction of q quantile forests depends on the \mathbf{X}_i 's and is based on subsampling, it is a more realistic modeling of Breiman's forests than totally non adaptive forests. It also provides a good understanding on why random forests are still consistent even when there is exactly one data point in each leaf.

Theorem 4.2 states that with a proper subsampling rate of the training set, the q quantile random forests are consistent.

(H3) *One has*

$$Y = m(\mathbf{X}) + \varepsilon,$$

where ε is a centred Gaussian noise with finite variance σ^2 variable, independent of \mathbf{X} . Moreover, \mathbf{X} has a density bounded from below and from above and m is continuous.

Theorem 4.2. *Assume that (H3) is satisfied. Then, providing $a_n \rightarrow \infty$ et $a_n/n \rightarrow \infty$, the infinite q quantile random forest is \mathbb{L}^2 consistent.*

Some remarks are in order. At first, observe that each tree in the q quantile forest is inconsistent (see Problem 4.3 in Györfi et al., 2002), because each leaf contains exactly one data point, a number which does not grow to infinity as $n \rightarrow \infty$. Thus, Theorem 4.2 shows that q quantile forest combines inconsistent trees to form a consistent estimate.

Secondly, many random forests can be seen as quantile forests if they satisfy the X -property and if splits do not separate a small fraction of data points from the rest of the sample. The last assumption is true, for example, if \mathbf{X} has a density on $[0, 1]^d$ bounded from below and from above, and if some splitting rule forces splits to be performed far away from the cell edges. This assumption is explicitly made in the analysis of Meinshausen (2006) and Wager (2014) to ensure that cell diameters tend to zero as $n \rightarrow \infty$, which is a necessary condition to prove the consistency of partitioning estimates (see Chapter 4 in Györfi et al., 2002).

We note finally that Theorem 4.2 does not cover the bootstrap case since in that case, $a_n = n$ data points are selected with replacement. However, the condition on the subsampling rate can be replaced by the following one: for all \mathbf{x} ,

$$\max_i \mathbb{P}_\Theta \left[\mathbf{x} \overset{\Theta}{\leftrightarrow} \mathbf{X}_i \right] \rightarrow 0 \text{ as } n \rightarrow \infty. \quad (9)$$

Condition (9) can be interpreted by saying that a point \mathbf{x} should not be connected too often to the same data point in the forest, thus meaning that trees have to be various enough to ensure the forest consistency. This idea of diversity among trees has already been suggested by Breiman (2001). In bootstrap case, a single data point is selected in about 64% of trees. Thus, the term $\max_i \mathbb{P}_\Theta \left[\mathbf{x} \overset{\Theta}{\leftrightarrow} \mathbf{X}_i \right]$ is roughly upper bounded by 0.64 which is not sufficient to prove (9). It does not mean that random forests based on bootstrap are inconsistent but that a more detailed analysis is required. A possible, but probably difficult, route is an in-depth analysis of the connection function $K_n(\mathbf{x}, \mathbf{X}_i) = \mathbb{P}_\Theta \left[\mathbf{x} \overset{\Theta}{\leftrightarrow} \mathbf{X}_i \right]$.

5 Proofs

5.1 Proof of Theorem 3.1

We assume that \mathcal{D}_n is fixed and prove Theorem 3.1 for $d = 2$. The general case can be treated similarly. Throughout the proof, we write, for all $\theta, \mathbf{x}, \mathbf{z} \in [0, 1]^2$,

$$f_{\mathbf{x}, \mathbf{z}}(\theta) = \frac{\mathbb{1}_{\mathbf{x} \leftrightarrow \mathbf{z}}^\theta}{N_n(\mathbf{x}, \theta)}.$$

Let us first consider a discrete random forest. By definition of such random forests, there exists $p \in \mathbb{N}^*$ and a partition $\{A_i : 1 \leq i \leq p\}$ of $[0, 1]^2$ such that the connection function K_n is constant over the sets $A_i \times A_j$'s ($1 \leq i, j \leq p$). For all $1 \leq i \leq p$, denote by \mathbf{a}_i , the center of the cell A_i . Take $\mathbf{x}, \mathbf{z} \in \mathbb{R}^2$. There exist i, j such that $\mathbf{x} \in A_i, \mathbf{z} \in A_j$. Thus, for all θ ,

$$\begin{aligned} \left| \frac{\mathbb{1}_{\mathbf{x} \leftrightarrow \mathbf{z}}^\theta}{N_n(\mathbf{x}, \theta)} - \frac{\mathbb{1}_{\mathbf{a}_i \leftrightarrow \mathbf{a}_j}^\theta}{N_n(\mathbf{a}_i, \theta)} \right| &\leq \left| \frac{\mathbb{1}_{\mathbf{x} \leftrightarrow \mathbf{z}}^\theta}{N_n(\mathbf{x}, \theta)} - \frac{\mathbb{1}_{\mathbf{a}_i \leftrightarrow \mathbf{z}}^\theta}{N_n(\mathbf{a}_i, \theta)} + \frac{\mathbb{1}_{\mathbf{a}_i \leftrightarrow \mathbf{z}}^\theta}{N_n(\mathbf{a}_i, \theta)} - \frac{\mathbb{1}_{\mathbf{a}_i \leftrightarrow \mathbf{a}_j}^\theta}{N_n(\mathbf{a}_i, \theta)} \right| \\ &\leq \frac{1}{N_n(\mathbf{a}_i, \theta)} \left| \mathbb{1}_{\mathbf{x} \leftrightarrow \mathbf{z}}^\theta - \mathbb{1}_{\mathbf{a}_i \leftrightarrow \mathbf{z}}^\theta \right| \\ &\quad + \frac{1}{N_n(\mathbf{a}_i, \theta)} \left| \mathbb{1}_{\mathbf{a}_i \leftrightarrow \mathbf{z}}^\theta - \mathbb{1}_{\mathbf{a}_i \leftrightarrow \mathbf{a}_j}^\theta \right| \\ &\leq \frac{1}{N_n(\mathbf{a}_i, \theta)} \mathbb{1}_{\mathbf{x} \leftrightarrow \mathbf{a}_i}^\theta + \frac{1}{N_n(\mathbf{a}_i, \theta)} \mathbb{1}_{\mathbf{a}_j \leftrightarrow \mathbf{z}}^\theta \\ &\leq 0. \end{aligned}$$

Thus, the set

$$\mathcal{H} = \{\theta \mapsto f_{\mathbf{x}, \mathbf{z}}(\theta) : \mathbf{x}, \mathbf{z} \in [0, 1]^2\}$$

is finite. Therefore, by the strong law of large numbers, almost surely, for all $f \in \mathcal{H}$,

$$\frac{1}{M} \sum_{m=1}^M f(\Theta_m) \xrightarrow{M \rightarrow \infty} \mathbb{E}_\Theta[f(\Theta)].$$

Noticing that $W_{ni}^M(\mathbf{x}) = \frac{1}{M} \sum_{m=1}^M f_{\mathbf{x}, \mathbf{x}_i}(\Theta_m)$, we obtain that, almost surely, for all $\mathbf{x} \in [0, 1]^2$,

$$W_{ni}^M(\mathbf{x}) \rightarrow W_{ni}^\infty(\mathbf{x}), \quad \text{as } M \rightarrow \infty.$$

Since \mathcal{D}_n is fixed and random forest estimates are linear in the weights, the proof of the discrete case is complete.

Let us now consider a continuous random forest. We define, for all $\mathbf{x}, \mathbf{z} \in [0, 1]^2$,

$$W_n^M(\mathbf{x}, \mathbf{z}) = \frac{1}{M} \sum_{m=1}^M \frac{\mathbb{1}_{\mathbf{x} \leftrightarrow \mathbf{z}}^{\Theta_m}}{N_n(\mathbf{x}, \Theta_m)},$$

and

$$W_n^\infty(\mathbf{x}, \mathbf{z}) = \mathbb{E}_\Theta \left[\frac{\mathbb{1}_{\mathbf{x} \leftrightarrow \mathbf{z}}}{N_n(\mathbf{x}, \Theta)} \right].$$

According to the strong law of large numbers, almost surely, for all $\mathbf{x}, \mathbf{z} \in [0, 1]^2 \cap \mathbb{Q}^2$,

$$\lim_{M \rightarrow \infty} W_n^M(\mathbf{x}, \mathbf{z}) = W_n^\infty(\mathbf{x}, \mathbf{z}).$$

Set $\mathbf{x}, \mathbf{z} \in [0, 1]^2$ where $\mathbf{x} = (x^{(1)}, x^{(2)})$ and $\mathbf{z} = (z^{(1)}, z^{(2)})$. Assume, without loss of generality, that $x^{(1)} < z^{(1)}$ and $x^{(2)} < z^{(2)}$. Let

$$A_{\mathbf{x}} = \{\mathbf{u} \in [0, 1]^2, u^{(1)} \leq x^{(1)} \text{ and } u^{(2)} \leq x^{(2)}\},$$

and $A_{\mathbf{z}} = \{\mathbf{u} \in [0, 1]^2, u^{(1)} \geq z^{(1)} \text{ and } u^{(2)} \geq z^{(2)}\}.$

Choose $\mathbf{x}_1 \in A_{\mathbf{x}} \cap \mathbb{Q}^2$ (resp. $\mathbf{z}_2 \in A_{\mathbf{z}} \cap \mathbb{Q}^2$) and take $\mathbf{x}_2 \in [0, 1]^2 \cap \mathbb{Q}^2$ (resp. $\mathbf{z}_1 \in [0, 1]^2 \cap \mathbb{Q}^2$) such that $\mathbf{x}_1, \mathbf{x}, \mathbf{x}_2$ (resp. $\mathbf{z}_2, \mathbf{z}, \mathbf{z}_1$) are aligned in this order (see Figure 1).




Figure 1: Respective positions of $\mathbf{x}, \mathbf{x}_1, \mathbf{x}_2$ and $\mathbf{z}, \mathbf{z}_1, \mathbf{z}_2$

Thus,

$$\begin{aligned} |W_n^M(\mathbf{x}, \mathbf{z}) - W_n^\infty(\mathbf{x}, \mathbf{z})| &\leq |W_n^M(\mathbf{x}, \mathbf{z}) - W_n^M(\mathbf{x}_1, \mathbf{z}_2)| \\ &\quad + |W_n^M(\mathbf{x}_1, \mathbf{z}_2) - W_n^\infty(\mathbf{x}_1, \mathbf{z}_2)| \\ &\quad + |W_n^\infty(\mathbf{x}_1, \mathbf{z}_2) - W_n^\infty(\mathbf{x}, \mathbf{z})|. \end{aligned} \quad (10)$$

Set $\varepsilon > 0$. Because of the continuity of K_n , we can choose $\mathbf{x}_1, \mathbf{x}_2$ close enough to \mathbf{x} and $\mathbf{z}_2, \mathbf{z}_1$ close enough to \mathbf{z} such that,

$$\begin{aligned} |K_n(\mathbf{x}_2, \mathbf{x}_1) - 1| &\leq \varepsilon, \\ |K_n(\mathbf{z}_1, \mathbf{z}_2) - 1| &\leq \varepsilon, \\ |1 - K_n(\mathbf{x}_1, \mathbf{x})| &\leq \varepsilon, \\ |1 - K_n(\mathbf{z}_2, \mathbf{z})| &\leq \varepsilon. \end{aligned}$$

Let us consider the second term in equation (10). Since $\mathbf{x}_1, \mathbf{z}_2$ belong to $[0, 1]^2 \cap \mathbb{Q}^2$, almost surely, there exists $M_1 > 0$ such that, if $M > M_1$,

$$|W_n^M(\mathbf{x}_1, \mathbf{z}_2) - W_n^\infty(\mathbf{x}_1, \mathbf{z}_2)| \leq \varepsilon.$$

Considering the first term in (10), we have

$$|W_n^M(\mathbf{x}, \mathbf{z}) - W_n^\infty(\mathbf{x}_1, \mathbf{z}_2)| \leq \frac{1}{M} \sum_{m=1}^M \left| \frac{\mathbb{1}_{\mathbf{x} \overset{\Theta_m}{\leftrightarrow} \mathbf{z}}}{N_n(\mathbf{x}, \Theta_m)} - \frac{\mathbb{1}_{\mathbf{x}_1 \overset{\Theta_m}{\leftrightarrow} \mathbf{z}_2}}{N_n(\mathbf{x}, \Theta_m)} \right|.$$

Observe that, given the positions of $\mathbf{x}, \mathbf{x}_1, \mathbf{z}, \mathbf{z}_2$, the only case where

$$\left| \frac{\mathbb{1}_{\mathbf{x} \overset{\Theta_m}{\leftrightarrow} \mathbf{z}}}{N_n(\mathbf{x}, \Theta_m)} - \frac{\mathbb{1}_{\mathbf{x}_1 \overset{\Theta_m}{\leftrightarrow} \mathbf{z}_2}}{N_n(\mathbf{x}, \Theta_m)} \right| \neq 0$$

occurs when $\mathbf{x}_1 \overset{\Theta_m}{\leftrightarrow} \mathbf{z}_2$ and $\mathbf{x} \not\overset{\Theta_m}{\leftrightarrow} \mathbf{z}$. Thus,

$$\begin{aligned} & \frac{1}{M} \sum_{m=1}^M \left| \frac{\mathbb{1}_{\mathbf{x} \overset{\Theta_m}{\leftrightarrow} \mathbf{z}}}{N_n(\mathbf{x}, \Theta_m)} - \frac{\mathbb{1}_{\mathbf{x}_1 \overset{\Theta_m}{\leftrightarrow} \mathbf{z}_2}}{N_n(\mathbf{x}, \Theta_m)} \right| \\ &= \frac{1}{M} \sum_{m=1}^M \left| \frac{\mathbb{1}_{\mathbf{x} \overset{\Theta_m}{\leftrightarrow} \mathbf{z}}}{N_n(\mathbf{x}, \Theta_m)} - \frac{\mathbb{1}_{\mathbf{x}_1 \overset{\Theta_m}{\leftrightarrow} \mathbf{z}_2}}{N_n(\mathbf{x}, \Theta_m)} \right| \mathbb{1}_{\mathbf{x}_1 \overset{\Theta_m}{\leftrightarrow} \mathbf{z}_2} \mathbb{1}_{\mathbf{x} \not\overset{\Theta_m}{\leftrightarrow} \mathbf{z}} \\ &\leq \frac{1}{M} \sum_{m=1}^M \mathbb{1}_{\mathbf{x} \overset{\Theta_m}{\leftrightarrow} \mathbf{z}} \mathbb{1}_{\mathbf{x}_1 \overset{\Theta_m}{\leftrightarrow} \mathbf{z}_2}. \end{aligned}$$

Again, given the relative positions of $\mathbf{x}, \mathbf{x}_1, \mathbf{x}_2, \mathbf{z}, \mathbf{z}_2, \mathbf{z}_1$, we obtain

$$\begin{aligned} \frac{1}{M} \sum_{m=1}^M \mathbb{1}_{\mathbf{x} \overset{\Theta_m}{\leftrightarrow} \mathbf{z}} \mathbb{1}_{\mathbf{x}_1 \overset{\Theta_m}{\leftrightarrow} \mathbf{z}_2} &\leq \frac{1}{M} \sum_{m=1}^M \left(\mathbb{1}_{\mathbf{x}_1 \overset{\Theta_m}{\leftrightarrow} \mathbf{x}} + \mathbb{1}_{\mathbf{z}_2 \overset{\Theta_m}{\leftrightarrow} \mathbf{z}} \right) \\ &\leq \frac{1}{M} \sum_{m=1}^M \left(\mathbb{1}_{\mathbf{x}_1 \overset{\Theta_m}{\leftrightarrow} \mathbf{x}_2} + \mathbb{1}_{\mathbf{z}_2 \overset{\Theta_m}{\leftrightarrow} \mathbf{z}_1} \right) \\ &\leq \frac{1}{M} \sum_{m=1}^M \mathbb{1}_{\mathbf{x}_1 \overset{\Theta_m}{\leftrightarrow} \mathbf{x}_2} + \frac{1}{M} \sum_{m=1}^M \mathbb{1}_{\mathbf{z}_2 \overset{\Theta_m}{\leftrightarrow} \mathbf{z}_1}. \end{aligned}$$

Collecting the previous inequalities, we have

$$\begin{aligned} |W_n^M(\mathbf{x}, \mathbf{z}) - W_n^\infty(\mathbf{x}_1, \mathbf{z}_2)| &\leq \frac{1}{M} \sum_{m=1}^M \mathbb{1}_{\mathbf{x}_1 \overset{\Theta_m}{\leftrightarrow} \mathbf{x}_2} + \frac{1}{M} \sum_{m=1}^M \mathbb{1}_{\mathbf{z}_2 \overset{\Theta_m}{\leftrightarrow} \mathbf{z}_1} \\ &\leq 2 - \frac{1}{M} \sum_{m=1}^M \mathbb{1}_{\mathbf{x}_1 \overset{\Theta_m}{\leftrightarrow} \mathbf{x}_2} - \frac{1}{M} \sum_{m=1}^M \mathbb{1}_{\mathbf{z}_2 \overset{\Theta_m}{\leftrightarrow} \mathbf{z}_1}. \end{aligned}$$

Since $\mathbf{x}_2, \mathbf{z}_1, \mathbf{x}_1, \mathbf{z}_2 \in [0, 1]^2 \cap \mathbb{Q}^2$, we deduce that there exists M_2 such that, for all $M > M_2$,

$$|W_n^M(\mathbf{x}, \mathbf{z}) - W_n^\infty(\mathbf{x}_1, \mathbf{z}_2)| \leq 2 - K_\infty(\mathbf{x}_2, \mathbf{x}_1) - K_\infty(\mathbf{z}_1, \mathbf{z}_2) + 2\varepsilon. \quad (11)$$

Considering the third term in (10), using the same arguments as above, we see that

$$\begin{aligned}
|W_n^\infty(\mathbf{x}_1, \mathbf{z}_2) - W_n^\infty(\mathbf{x}, \mathbf{z})| &\leq \mathbb{E}_\Theta \left| \frac{\mathbb{1}_{\mathbf{x}_1 \overset{\Theta}{\leftrightarrow} \mathbf{z}_2}}{N_n(\mathbf{x}_1, \Theta)} - \frac{\mathbb{1}_{\mathbf{x} \overset{\Theta}{\leftrightarrow} \mathbf{z}}}{N_n(\mathbf{x}, \Theta)} \right| \\
&\leq \mathbb{E}_\Theta \left| \frac{\mathbb{1}_{\mathbf{x}_1 \overset{\Theta}{\leftrightarrow} \mathbf{z}_2}}{N_n(\mathbf{x}_1, \Theta)} - \frac{\mathbb{1}_{\mathbf{x} \overset{\Theta}{\leftrightarrow} \mathbf{z}}}{N_n(\mathbf{x}, \Theta)} \mathbb{1}_{\mathbf{x}_1 \overset{\Theta}{\leftrightarrow} \mathbf{z}_2} \mathbb{1}_{\mathbf{x} \overset{\Theta}{\leftrightarrow} \mathbf{z}} \right| \\
&\leq \mathbb{E}_\Theta \left[\mathbb{1}_{\mathbf{x}_1 \overset{\Theta}{\leftrightarrow} \mathbf{z}_2} \mathbb{1}_{\mathbf{x} \overset{\Theta}{\leftrightarrow} \mathbf{z}} \right] \\
&\leq \mathbb{E}_\Theta \left[\mathbb{1}_{\mathbf{x}_1 \overset{\Theta}{\leftrightarrow} \mathbf{x}_2} + \mathbb{1}_{\mathbf{z}_2 \overset{\Theta}{\leftrightarrow} \mathbf{z}_1} \right] \\
&\leq 2 - K_n(\mathbf{x}_1, \mathbf{x}_2) - K_n(\mathbf{z}_2, \mathbf{z}_1). \tag{12}
\end{aligned}$$

Using inequalities (11) and (12) in (10), we finally conclude that, for all $M > \max(M_1, M_2)$,

$$\begin{aligned}
|W_n^M(\mathbf{x}, \mathbf{z}) - W_n^\infty(\mathbf{x}, \mathbf{z})| &\leq 4 - 2K_\infty(\mathbf{x}_2, \mathbf{x}_1) - 2K_\infty(\mathbf{z}_1, \mathbf{z}_2) + 3\varepsilon \\
&\leq 7\varepsilon.
\end{aligned}$$

This completes the proof of Theorem 3.1.

5.2 Proof of Lemma 1 and Theorem 3.2

Proof of Lemma 1. Set $k \in \mathbb{N}$ and $\varepsilon > 0$. We start by considering the case where $d = 1$. Take $x, z \in [0, 1]$ and let $w = -\log(|x - z|)$. The probability that x and z are not connected in the uniform forest after k cuts is given by

$$\begin{aligned}
1 - K_k(x, z) &\leq 1 - K_k(0, |z - x|) \\
&\quad \text{(according to Technical Lemma 1, see the end of the section)} \\
&\leq e^{-w} \mathbb{1}_{k>0} \sum_{i=0}^{k-1} \frac{w^i}{i!} \\
&\quad \text{(according to Technical Lemma 2, see the end of the section)} \\
&\leq \frac{(k+2)!e}{w^3},
\end{aligned}$$

for all $w > 1$. Now, consider the multivariate case, and let $\mathbf{x}, \mathbf{z} \in [0, 1]^d$. Set, for all $1 \leq j \leq d$, $w_j = -\log(|x_j - z_j|)$. By union bound, recalling that $1 - K_k(\mathbf{x}, \mathbf{z}) = \mathbb{P}_\Theta(\mathbf{x} \not\overset{\Theta}{\leftrightarrow} \mathbf{z})$, we have

$$\begin{aligned}
1 - K_k(\mathbf{x}, \mathbf{z}) &\leq \sum_{j=1}^d (1 - K_k(x_j, z_j)) \\
&\leq \frac{d(k+2)!e}{\min_{1 \leq j \leq d} w_j^3}.
\end{aligned}$$

Thus, if, for all $1 \leq j \leq d$,

$$|x_j - z_j| \leq \exp\left(-\frac{(A_{k,d})^{1/3}}{\varepsilon^{2/3}}\right),$$

then

$$1 - K_k(\mathbf{x}, \mathbf{z}) \leq \frac{\varepsilon^2}{8},$$

where $A_{k,d} = (8de(k+2)!)^{1/3}$. Consequently,

$$\delta(\varepsilon) \geq \exp\left(-\frac{(A_{k,d})^{1/3}}{\varepsilon^{2/3}}\right).$$

□

Proof of Theorem 3.2. We start the proof by proving that the class

$$\mathcal{H} = \{\theta \mapsto f_{\mathbf{x},\mathbf{z}}(\theta) : \mathbf{x}, \mathbf{z} \in \mathbb{R}^2\}$$

is \mathbb{P}_Θ -Donsker, that is, there exists a Gaussian process \mathbb{G} such that

$$\sup_{f \in \mathcal{H}} \{\mathbb{E}|f| (d\mathbb{G}_M - d\mathbb{G})\} \xrightarrow{M \rightarrow \infty} 0.$$

At first, let us consider a finite random forest. As noticed in the proof of Theorem 3.1, the set \mathcal{H} is finite. Consequently, by the central limit theorem, the set \mathcal{H} is \mathbb{P}_Θ -Donsker.

Now, consider a random forest which satisfies the second statement in **Assumption 1**. Set $\varepsilon > 0$. Consider a regular grid of $[0, 1]^d$ with a step δ and let \mathcal{G}_δ be the set of nodes of this grid. We start by finding a condition on δ such that the set

$$\tilde{\mathcal{G}}_\delta = \{[f_{\mathbf{x}_1, \mathbf{z}_1}, f_{\mathbf{x}_2, \mathbf{z}_2}] : \mathbf{x}_1, \mathbf{x}_2, \mathbf{z}_1, \mathbf{z}_2 \in \mathcal{G}_\delta\}$$

is a covering of ε -bracket of the set \mathcal{H} , that is, for all $f \in \mathcal{H}$, there exists $\mathbf{x}_1, \mathbf{z}_1, \mathbf{x}_2, \mathbf{z}_2 \in \mathcal{G}_\delta$ such that

$$f_{\mathbf{x}_1, \mathbf{z}_1} \leq f \leq f_{\mathbf{x}_2, \mathbf{z}_2} \text{ and } \mathbb{E}^{1/2}[f_{\mathbf{x}_2, \mathbf{z}_2}(\Theta) - f_{\mathbf{x}_1, \mathbf{z}_1}(\Theta)]^2 \leq \varepsilon. \quad (13)$$

To this aim, set $\mathbf{x}, \mathbf{z} \in [0, 1]^d$ and choose $\mathbf{x}_1, \mathbf{x}_2, \mathbf{z}_1, \mathbf{z}_2 \in \mathcal{G}_\delta$ (see Figure 2). Note that, for all θ ,

$$\frac{\mathbb{1}_{\mathbf{x}_1 \overset{\theta}{\leftrightarrow} \mathbf{z}_2}}{N_n(\mathbf{x}_1, \theta)} \leq \frac{\mathbb{1}_{\mathbf{x} \overset{\theta}{\leftrightarrow} \mathbf{z}}}{N_n(\mathbf{x}, \theta)} \leq \frac{\mathbb{1}_{\mathbf{x}_2 \overset{\theta}{\leftrightarrow} \mathbf{z}_1}}{N_n(\mathbf{x}_2, \theta)},$$

that is, $f_{\mathbf{x}_1, \mathbf{z}_2} \leq f_{\mathbf{x}, \mathbf{z}} \leq f_{\mathbf{x}_2, \mathbf{z}_1}$. To prove the second statement in (13), observe that

$$\begin{aligned} \mathbb{E}^{1/2}[f_{\mathbf{x}_2, \mathbf{z}_2}(\Theta) - f_{\mathbf{x}_1, \mathbf{z}_1}(\Theta)]^2 &= \mathbb{E}_\Theta^{1/2} \left[\frac{\mathbb{1}_{\mathbf{x}_1 \overset{\Theta}{\leftrightarrow} \mathbf{z}_2}}{N_n(\mathbf{x}_1, \Theta)} - \frac{\mathbb{1}_{\mathbf{x}_2 \overset{\Theta}{\leftrightarrow} \mathbf{z}_1}}{N_n(\mathbf{x}_2, \Theta)} \right]^2 \\ &= \mathbb{E}_\Theta^{1/2} \left[\left(\frac{\mathbb{1}_{\mathbf{x}_1 \overset{\Theta}{\leftrightarrow} \mathbf{z}_2}}{N_n(\mathbf{x}_1, \Theta)} - \frac{\mathbb{1}_{\mathbf{x}_2 \overset{\Theta}{\leftrightarrow} \mathbf{z}_1}}{N_n(\mathbf{x}_2, \Theta)} \right) \right. \\ &\quad \left. \times \mathbb{1}_{\mathbf{x}_1 \overset{\Theta}{\leftrightarrow} \mathbf{z}_2} \mathbb{1}_{\mathbf{x}_2 \overset{\Theta}{\leftrightarrow} \mathbf{z}_1} \right]^2 \\ &\leq \mathbb{E}_\Theta^{1/2} \left[\mathbb{1}_{\mathbf{x}_1 \overset{\Theta}{\leftrightarrow} \mathbf{x}_2} + \mathbb{1}_{\mathbf{z}_1 \overset{\Theta}{\leftrightarrow} \mathbf{z}_2} \right]^2 \\ &\leq 2\sqrt{1 - K_n(\mathbf{x}_1, \mathbf{x}_2) + 1 - K_n(\mathbf{z}_1, \mathbf{z}_2)}. \end{aligned}$$




Figure 2: Respective positions of $\mathbf{x}, \mathbf{x}_1, \mathbf{x}_2$ and $\mathbf{z}, \mathbf{z}_1, \mathbf{z}_2$ with $d = 2$.

Thus, we have to choose the grid step δ such that

$$\sup_{\substack{\mathbf{x}_1, \mathbf{x}_2 \in [0,1]^d \\ \|\mathbf{x}_1 - \mathbf{x}_2\|_\infty \leq \delta}} |1 - K_n(\mathbf{x}_1, \mathbf{x}_2)| \leq \frac{\varepsilon^2}{8}. \quad (14)$$

By **Assumption 1** and the definition of the random forest grid step, there exist constants $C, A > 0$ and $0 < \alpha < 2$ such that, for all $\varepsilon > 0$, if

$$\delta \geq C \exp(-A/\varepsilon^\alpha), \quad (15)$$

then (14) is satisfied. Hence, if δ satisfies (15), then $\tilde{\mathcal{G}}_\delta$ is a covering of ε -bracket of \mathcal{H} . In that case, the number $N_{[\cdot]}(\varepsilon, \mathcal{F}, L_2(P))$ of ε -bracket needed to cover \mathcal{H} satisfies

$$N_{[\cdot]}(\varepsilon, \mathcal{F}, L_2(P)) \leq \text{Card}(\tilde{\mathcal{G}}_\delta) \leq \text{Card}(\mathcal{G}_\delta)^4 \leq \left(\frac{1}{\delta}\right)^{4d}.$$

Consequently,

$$\sqrt{\log N_{[\cdot]}(\varepsilon, \mathcal{F}, L_2(P))} \leq \sqrt{\frac{2Ad}{\varepsilon^\alpha} - 2d \log C}$$

where the last term is integrable near zero since $\alpha < 2$. Thus, according to Theorem 2.5.6 in [van der Vaart and Wellner \(1996\)](#) (and the remark at the beginning of Section 2.5.2 in [van der Vaart and Wellner, 1996](#)), the class \mathcal{H} is \mathbb{P}_Θ -Donsker.

To conclude the proof, consider a random forest satisfying **(H1)**. From above, we see that the class \mathcal{H} is \mathbb{P}_Θ -Donsker. Recall that $\mathcal{F}_2 = \{g_{\mathbf{x}} : \theta \mapsto m_n(\mathbf{x}, \theta) : \mathbf{x} \in [0, 1]^d\}$, where

$$m_n(\mathbf{x}, \Theta) = \sum_{i=1}^n Y_i f_{\mathbf{x}, \mathbf{x}_i}(\Theta).$$

Since the training set \mathcal{D}_n is fixed, we have

$$\begin{aligned}
& \sup_{g_{\mathbf{x}} \in \mathcal{F}_2} \left\{ \mathbb{E} |g_{\mathbf{x}}| (d\mathbb{G}_M - d\mathbb{G}) \right\} \\
&= \sup_{\mathbf{x} \in [0,1]^d} \left\{ \mathbb{E} \left| \sum_{i=1}^n Y_i f_{\mathbf{x}, \mathbf{X}_i} \right| (d\mathbb{G}_M - d\mathbb{G}) \right\} \\
&\leq \sum_{i=1}^n |Y_i| \sup_{\mathbf{x} \in [0,1]^d} \left\{ \mathbb{E} |f_{\mathbf{x}, \mathbf{X}_i}| (d\mathbb{G}_M - d\mathbb{G}) \right\} \\
&\leq \left(\sum_{i=1}^n |Y_i| \right) \sup_{\mathbf{x}, \mathbf{z} \in [0,1]^d} \left\{ \mathbb{E} |f_{\mathbf{x}, \mathbf{z}}| (d\mathbb{G}_M - d\mathbb{G}) \right\},
\end{aligned}$$

which tends to zero as M tends to infinity, since the class \mathcal{H} is \mathbb{P}_{Θ} -Donsker.

Finally, note that Breiman's random forests are discrete, thus satisfying **(H1)**. Uniform forests are continuous and satisfy **(H1)** according to Lemma 1.

□

5.3 Proof of Theorem 3.3

Observe that,

$$\begin{aligned}
& \left(m_{M,n}(\mathbf{X}, \Theta_1, \dots, \Theta_m) - m(\mathbf{X}) \right)^2 \\
&= \left(m_{M,n}(\mathbf{X}, \Theta_1, \dots, \Theta_m) - \mathbb{E}_{\Theta} [m_n(\mathbf{X}, \Theta)] \right)^2 + \left(\mathbb{E}_{\Theta} [m_n(\mathbf{X}, \Theta)] - m(\mathbf{X}) \right)^2 \\
&\quad + 2 \left(\mathbb{E}_{\Theta} [m_n(\mathbf{X}, \Theta)] - m(\mathbf{X}) \right) \left(m_{M,n}(\mathbf{X}, \Theta_1, \dots, \Theta_m) - \mathbb{E}_{\Theta} [m_n(\mathbf{X}, \Theta)] \right).
\end{aligned}$$

Taking the expectation on both sides, we obtain

$$R(m_{M,n}, m) = R(m_{\infty,n}, m) + \mathbb{E} \left[m_{M,n}(\mathbf{X}, \Theta_1, \dots, \Theta_m) - \mathbb{E}_{\Theta} [m_n(\mathbf{X}, \Theta)] \right]^2,$$

by noticing that

$$\begin{aligned}
& \mathbb{E} \left[\left(m_{M,n}(\mathbf{X}, \Theta_1, \dots, \Theta_m) - \mathbb{E}_{\Theta} [m_n(\mathbf{X}, \Theta)] \right) \left(\mathbb{E}_{\Theta} [m_n(\mathbf{X}, \Theta)] - m(\mathbf{X}) \right) \right] \\
&= \mathbb{E}_{\mathbf{X}, \mathcal{D}_n} \left[\left(\mathbb{E}_{\Theta} [m_n(\mathbf{X}, \Theta)] - m(\mathbf{X}) \right) \right. \\
&\quad \left. \times \mathbb{E}_{\Theta_1, \dots, \Theta_m} \left[m_{M,n}(\mathbf{X}, \Theta_1, \dots, \Theta_m) - \mathbb{E}_{\Theta} [m_n(\mathbf{X}, \Theta)] \right] \right] \\
&= 0,
\end{aligned}$$

according to the definition of $m_{M,n}$. Fixing \mathbf{X} and \mathcal{D}_n , note that random variables $m_n(\mathbf{X}, \Theta_1), \dots, m_n(\mathbf{X}, \Theta_m)$ are independent and identically distributed.

Thus, we have

$$\begin{aligned}
& \mathbb{E} [m_{M,n}(\mathbf{X}, \Theta_1, \dots, \Theta_M) - \mathbb{E}_\Theta [m_n(\mathbf{X}, \Theta)]]^2 \\
&= \mathbb{E}_{\mathbf{X}, \mathcal{D}_n} \mathbb{E}_{\Theta_1, \dots, \Theta_M} \left[\frac{1}{M} \sum_{m=1}^M m_n(\mathbf{X}, \Theta_m) - \mathbb{E}_\Theta [m_n(\mathbf{X}, \Theta)] \right]^2 \\
&= \frac{1}{M} \times \mathbb{E} [\mathbb{V}_\Theta [m_n(\mathbf{X}, \Theta)]] ,
\end{aligned}$$

which concludes the first part of the proof. Now, note that,

$$\begin{aligned}
R(m_{M,n}) - R(m_{\infty,n}) &= \frac{1}{M} \times \mathbb{E} [\mathbb{V}_\Theta [m_n(\mathbf{X}, \Theta)]] \\
&= \frac{1}{M} \times \mathbb{E} \left[\mathbb{V}_\Theta \left[\sum_{i=1}^n W_{ni}(\mathbf{X}, \Theta) (m(\mathbf{X}_i) + \varepsilon_i) \right] \right] \\
&\leq \frac{1}{M} \times \left[8\|m\|_\infty^2 + 2\mathbb{E} \left[\mathbb{V}_\Theta \left[\sum_{i=1}^n W_{ni}(\mathbf{X}, \Theta) \varepsilon_i \right] \right] \right] \\
&\leq \frac{1}{M} \times \left[8\|m\|_\infty^2 + 2\mathbb{E} \left[\max_{1 \leq i \leq n} \varepsilon_i - \min_{1 \leq j \leq n} \varepsilon_j \right]^2 \right] \\
&\leq \frac{1}{M} \times \left[8\|m\|_\infty^2 + 8\sigma^2 \mathbb{E} \left[\max_{1 \leq i \leq n} \frac{\varepsilon_i}{\sigma} \right]^2 \right].
\end{aligned}$$

The term inside the brackets is the maximum of n χ^2 -squared distributed random variables. Thus, for all $n \in \mathbb{N}^*$,

$$\mathbb{E} \left[\max_{1 \leq i \leq n} \varepsilon_i^2 \right] \leq 1 + 4 \log n,$$

(see, e.g., Chapter 1 in [Boucheron et al., 2013](#)). Therefore,

$$R(m_{M,n}) - R(m_{\infty,n}) \leq \frac{8}{M} \times (\|m\|_\infty^2 + \sigma^2(1 + 4 \log n)).$$

5.4 Proof of Theorem 4.1 and Proposition 1

The proof of Theorem 4.1 is based on Stone's theorem which is recalled here.

Stone's theorem (1977). Assume that the following conditions are satisfied for every distribution of \mathbf{X} :

- (i) There is a constant c such that for every non negative measurable function f satisfying $\mathbb{E}f(\mathbf{X}) < \infty$ and any n ,

$$\mathbb{E} \left(\sum_{i=1}^n W_{ni}(\mathbf{X}) f(\mathbf{X}_i) \right) \leq c \mathbb{E} (f(\mathbf{X})).$$

- (ii) There is a $D > 1$ such that, for all n ,

$$\mathbb{P} \left(\sum_{i=1}^n W_{ni}(\mathbf{X}) < D \right) = 1.$$

(iii) For all $a > 0$,

$$\lim_{n \rightarrow \infty} \mathbb{E} \left(\sum_{i=1}^n W_{ni}(\mathbf{X}) \mathbb{1}_{\|\mathbf{X} - \mathbf{X}_i\| > a} \right) = 0.$$

(iv) The sum of weights satisfies

$$\sum_{i=1}^n W_{ni}(\mathbf{X}) \xrightarrow[n \rightarrow \infty]{} 1 \quad \text{in probability.}$$

(v)

$$\lim_{n \rightarrow \infty} \mathbb{E} \left(\max_{1 \leq i \leq n} W_{ni}(\mathbf{X}) \right) = 0.$$

Then the corresponding regression function estimate m_n is universally \mathbb{L}^2 consistent, that is,

$$\lim_{n \rightarrow \infty} \mathbb{E} [m_{\infty,n}(\mathbf{X}) - m(\mathbf{X})]^2 = 0,$$

for all distributions of (\mathbf{X}, Y) with $\mathbb{E}Y^2 < \infty$.

Proof of Theorem 4.1. We check the assumptions of Stone's theorem. For every non negative measurable function f satisfying $\mathbb{E}f(\mathbf{X}) < \infty$ and for any n , almost surely,

$$\mathbb{E}_{\mathbf{X}, \mathcal{D}_n} \left(\sum_{i=1}^n W_{ni}(\mathbf{X}, \Theta) f(\mathbf{X}_i) \right) \leq \mathbb{E}_{\mathbf{X}} (f(\mathbf{X})),$$

where

$$W_{ni}(\mathbf{X}, \Theta) = \frac{\mathbb{1}_{\mathbf{X}_i \in A_n(\mathbf{X}, \Theta)}}{N_n(\mathbf{X}, \Theta)}$$

are the weights of the random tree $\mathcal{T}_n(\Theta)$ (see the proof of Theorem 4.2 in [Györfi et al., 2002](#)). Taking expectation with respect to Θ from both sides, we have

$$\mathbb{E}_{\mathbf{X}, \mathcal{D}_n} \left(\sum_{i=1}^n W_{ni}^\infty(\mathbf{X}) f(\mathbf{X}_i) \right) \leq \mathbb{E}_{\mathbf{X}} (f(\mathbf{X})),$$

which proves the first condition of Stone's theorem.

According to the definition of random forest weights W_{ni}^∞ , since $\sum_{i=1}^n W_{ni}(\mathbf{X}, \Theta) \leq 1$ almost surely, we have

$$\sum_{i=1}^n W_{ni}^\infty(\mathbf{X}) = \mathbb{E}_\Theta \left[\sum_{i=1}^n W_{ni}(\mathbf{X}, \Theta) \right] \leq 1.$$

To check condition (iii), note that, for all $a > 0$,

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^n W_{ni}^\infty(\mathbf{X}) \mathbb{1}_{\|\mathbf{X} - \mathbf{X}_i\|_\infty > a} \right] &= \mathbb{E} \left[\sum_{i=1}^n \mathbb{1}_{\mathbf{X} \leftrightarrow \mathbf{X}_i} \mathbb{1}_{\|\mathbf{X} - \mathbf{X}_i\|_\infty > a} \right] \\ &= \mathbb{E} \left[\sum_{i=1}^n \mathbb{1}_{\mathbf{X} \leftrightarrow \mathbf{X}_i} \mathbb{1}_{\|\mathbf{X} - \mathbf{X}_i\|_\infty > a} \right. \\ &\quad \left. \times \mathbb{1}_{\text{diam}(A_n(\mathbf{X}, \Theta)) \geq a/2} \right], \end{aligned}$$

because $\mathbb{1}_{\|\mathbf{X} - \mathbf{X}_i\|_\infty > a} \mathbb{1}_{\text{diam}(A_n(\mathbf{X}, \Theta)) < a/2} = 0$. Thus,

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^n W_{ni}^\infty(\mathbf{X}) \mathbb{1}_{\|\mathbf{X} - \mathbf{X}_i\|_\infty > a} \right] &\leq \mathbb{E} \left[\mathbb{1}_{\text{diam}(A_n(\mathbf{X}, \Theta)) \geq a/2} \right. \\ &\quad \left. \times \sum_{i=1}^n \mathbb{1}_{\mathbf{X} \leftrightarrow \mathbf{X}_i} \mathbb{1}_{\|\mathbf{X} - \mathbf{X}_i\|_\infty > a} \right] \\ &\leq \mathbb{P} \left[\text{diam}(A_n(\mathbf{X}, \Theta)) \geq a/2 \right], \end{aligned}$$

which tends to zero, as $n \rightarrow \infty$, by assumption.

To prove assumption (iv), we follow the arguments developed by [Biau et al. \(2008\)](#). For completeness, these arguments are recalled here. Let us consider the partition associated with the random tree $\mathcal{T}_n(\Theta)$. By definition, this partition has 2^k cells, denoted by A_1, \dots, A_{2^k} . For $1 \leq i \leq 2^k$, let N_i be the number of points among $\mathbf{X}, \mathbf{X}_1, \dots, \mathbf{X}_n$ falling into A_i . Finally, set $\mathcal{S} = \{\mathbf{X}, \mathbf{X}_1, \dots, \mathbf{X}_n\}$. Since these points are independent and identically distributed, fixing the set \mathcal{S} (but not the order of the points) and Θ , the probability that \mathbf{X} falls in the i -th cell is $N_i/(n+1)$. Thus, for every fixed $t > 0$,

$$\begin{aligned} \mathbb{P}[N_n(\mathbf{X}, \Theta) < t] &= \mathbb{E} \left[\mathbb{P}[N_n(\mathbf{X}, \Theta) < t \mid \mathcal{S}, \Theta] \right] \\ &= \mathbb{E} \left[\sum_{i: N_i < t+1} \frac{N_i}{n+1} \right] \\ &\leq \frac{2^k}{n+1} t. \end{aligned}$$

Thus, by assumption, $N_n(\mathbf{X}, \Theta) \rightarrow \infty$ in probability, as $n \rightarrow \infty$. Consequently, observe that

$$\begin{aligned} \sum_{i=1}^n W_{ni}^\infty(\mathbf{X}) &= \mathbb{E}_\Theta \left[\sum_{i=1}^n W_{ni}(\mathbf{X}, \Theta) \right] \\ &= \mathbb{E}_\Theta \left[\mathbb{1}_{N_n(\mathbf{X}, \Theta) \neq 0} \right] \\ &= \mathbb{P}_\Theta [N_n(\mathbf{X}, \Theta) \neq 0] \\ &\rightarrow 1 \quad \text{as } n \rightarrow \infty. \end{aligned}$$

At last, to prove (v), note that,

$$\begin{aligned}\mathbb{E} \left[\max_{1 \leq i \leq n} W_{ni}^\infty(\mathbf{X}) \right] &\leq \mathbb{E} \left[\max_{1 \leq i \leq n} \frac{\mathbb{1}_{\mathbf{X}_i \in A_n(\mathbf{X}, \Theta)}}{N_n(\mathbf{X}, \Theta)} \right] \\ &\leq \mathbb{E} \left[\frac{1}{N_n(\mathbf{X}, \Theta)} \right] \\ &\rightarrow 0 \quad \text{as } n \rightarrow \infty,\end{aligned}$$

since $N_n(\mathbf{X}, \Theta) \rightarrow \infty$ in probability, as $n \rightarrow \infty$. \square

Proof of Proposition 1. We check conditions of Theorem 4.1. Let us denote by $V_{nj}(\mathbf{X}, \Theta)$ the length of the j -th side of the cell containing \mathbf{X} and $K_{nj}(\mathbf{X}, \Theta)$ the number of times the cell containing \mathbf{X} is cut along the j -coordinate. Note that, if U_1, \dots, U_n are independent uniform on $[0, 1]$,

$$\begin{aligned}\mathbb{E}[V_{nj}(\mathbf{X}, \Theta)] &\leq \mathbb{E} \left[\mathbb{E} \left[\prod_{l=1}^{K_{nj}(\mathbf{X}, \Theta)} \max(U_l, 1 - U_l) \middle| K_{nj}(\mathbf{X}, \Theta) \right] \right] \\ &= \mathbb{E} \left[\left[\mathbb{E}[\max(U_1, 1 - U_1)] \right]^{K_{nj}(\mathbf{X}, \Theta)} \right] \\ &= \mathbb{E} \left[\left(\frac{3}{4} \right)^{K_{nj}(\mathbf{X}, \Theta)} \right].\end{aligned}$$

Since $K_{nj}(\mathbf{X}, \Theta)$ is distributed as a binomial $\mathcal{B}(k_n, 1/d)$, $K_{nj}(\mathbf{X}, \Theta) \rightarrow +\infty$ in probability, as n tends to infinity. Thus $\mathbb{E}[V_{nj}(\mathbf{X}, \Theta)] \rightarrow 0$ as $n \rightarrow \infty$. \square

5.5 Proof of Theorem 4.2

To prove Theorem 4.2, we need the following lemma which states that the cell diameter of a quantile tree tends to zero.

Lemma 2. *Assume that \mathbf{X} has a density f over $[0, 1]^d$, with respect to the Lebesgue measure and that there exist two constants $c, C > 0$ such that, for all $\mathbf{x} \in [0, 1]^d$,*

$$c \leq f(\mathbf{x}) \leq C.$$

*Thus, for all $q \in [1/2, 1)$, the q quantile tree defined in **Algorithm 2** satisfies, for all γ ,*

$$\mathbb{P}_{\mathbf{X}, \Theta}[\text{diam}(A_n(\mathbf{X}, \Theta)) > \gamma] \xrightarrow{n \rightarrow \infty} 0.$$

Proof of Lemma 2. Set $q \in [1/2, 1)$. At first, consider a theoretical q quantile tree where cuts are made similarly as in the q quantile tree but by selecting $q_n \in [1 - q, q]$ and by performing the cut at the q_n theoretical quantile (instead of empirical one). The tree is then stopped at level k , where $k \in \mathbb{N}$ is a parameter to be chosen later. Then, consider a cell $A = \prod_{j=1}^d [a_j, b_j]$ of the theoretical q quantile tree. Assume that this cell is cut along the first coordinate and let z be

the split position. Thus, by definition of the theoretical q quantile tree, there exists $q' \in [1 - q, q]$ such that

$$\int_{a_1}^z \int_{a_2}^{b_2} \dots \int_{a_d}^{b_d} f(\mathbf{x}) dx_1 \dots dx_d = q' \int_{a_1}^{b_1} \dots \int_{a_d}^{b_d} f(\mathbf{x}) dx_1 \dots dx_d,$$

that is

$$(1 - q') \int_{a_1}^z g(x_1) dx_1 = q' \int_z^{b_1} g(x_1) dx_1, \quad (16)$$

where $g(x_1) = \int_{a_2}^{b_2} \dots \int_{a_d}^{b_d} f(\mathbf{x}) dx_2 \dots dx_d$. Letting $\mu_{d-1}(A) = \prod_{j=2}^d (b_j - a_j)$, by assumption, we have, for all $x_1 \in [a_1, b_1]$,

$$c\mu_{d-1}(A) \leq g(x_1) \leq C\mu_{d-1}(A).$$

Hence, using (16), we obtain

$$(1 - q')(z - a_1)c\mu_{d-1}(A) \leq q' \int_z^{b_1} g(x_1) dx_1 \leq q'(b_1 - z)C\mu_{d-1}(A),$$

which leads to

$$\frac{z - a_1}{b_1 - a_1} \leq \frac{q'C}{q'C + (1 - q')c} \in]0, 1[. \quad (17)$$

Similarly,

$$q'(b_1 - z)c\mu_{d-1}(A) \leq (1 - q') \int_{a_1}^z g(x_1) dx_1 \leq (1 - q')(z - a_1)C\mu_{d-1}(A),$$

which yields

$$\frac{b_1 - z}{b_1 - a_1} \leq \frac{C(1 - q')}{q'c + (1 - q')C} \in]0, 1[. \quad (18)$$

Combining (17) and (18), we deduce that

$$\begin{aligned} \max\left(\frac{z - a_1}{b_1 - a_1}, 1 - \frac{z - a_1}{b_1 - a_1}\right) &\leq \max\left(\frac{q'C}{q'C + (1 - q')c}, \frac{C(1 - q')}{q'c + (1 - q')C}\right) \\ &\leq \frac{Cq}{(1 - q)(c + C)}. \end{aligned}$$

Consequently, letting

$$\alpha = \frac{Cq}{(1 - q)(c + C)} \in (0, 1),$$

the first dimension of A is reduced at most by a factor $1 - \alpha > 0$ and at least by a factor $\alpha < 1$.

Denote by $V_{ik}(\mathbf{X}, \Theta)$ the length of the i -th side of the cell containing \mathbf{X} at level k in the theoretical q quantile tree, and let $K_{ik}(\mathbf{X}, \Theta)$ the number of times this cell has been cut along the i -th coordinate. Hence, for all $i \in \{1, \dots, d\}$,

$$\mathbb{E}[V_{ik}(\mathbf{X}, \Theta)] \leq \mathbb{E}[\alpha^{K_{ik}(\mathbf{X}, \Theta)}] \xrightarrow{k \rightarrow \infty} 0, \quad (19)$$

which proves that the cell diameter of the theoretical q quantile tree tends to zero, as the level k tends to infinity.

Now, consider the empirical q quantile tree as defined in **Algorithm 2** but stopped at level k . Thus, for n large enough, at each step of the algorithm, q_n is selected in $[1 - q, q]$. Set $\varepsilon, \eta > 0$ and let $G_k(\mathbf{X}, \Theta)$ be the event where all k cuts used to build the cell $A_k(\mathbf{X}, \Theta)$ are distant of less than η from cuts used to build the cell $A_k^*(\mathbf{X}, \Theta)$ of the theoretical q quantile tree. Thus,

$$\begin{aligned} \mathbb{E}[\text{diam}(A_k(\mathbf{X}, \Theta))] &= \mathbb{E}[\text{diam}(A_k(\mathbf{X}, \Theta)\mathbb{1}_{G_k(\mathbf{X}, \Theta)})] \\ &\quad + \mathbb{E}[\text{diam}(A_k(\mathbf{X}, \Theta)\mathbb{1}_{G_k(\mathbf{X}, \Theta)^c})] \\ &\leq \mathbb{E}[\text{diam}(A_k(\mathbf{X}, \Theta)\mathbb{1}_{G_k(\mathbf{X}, \Theta)})] + \mathbb{P}[G_k(\mathbf{X}, \Theta)^c]. \end{aligned} \quad (20)$$

From equation (19), there exists $k_0 \in \mathbb{N}^*$ such that

$$\mathbb{E}[\text{diam}(A_{k_0}^*(\mathbf{X}, \Theta))] < \varepsilon.$$

Since, on the event $G_k(\mathbf{X}, \Theta)$, the k consecutive cuts used to build $A_k^*(\mathbf{X}, \Theta)$ are distant of less than η from the k cuts used to design $A_k(\mathbf{X}, \Theta)$, we have

$$\mathbb{E}[\text{diam}(A_{k_0}(\mathbf{X}, \Theta))] < k_0\eta + \varepsilon. \quad (21)$$

With respect to the second term in equation (20), consider a cell A of the empirical q quantile tree. Without loss of generality, we assume that the next split is performed along the first coordinate. Let F^A (resp. F_n^A) be the one dimensional conditional distribution function (resp. empirical distribution function) of \mathbf{X} given that $\mathbf{X} \in A$. Denote by z_n^A the position of the empirical split performed in A . Since \mathbf{X} is uniformly distributed over $[0, 1]^d$, $(F^A)' \geq c/\mu(A)$. Thus, since F^A is an increasing function, if $\sup_{x \in A} |F^A(x) - F_n^A(x)| \leq c\eta/\mu(A)$ then

$$\inf_{z \in \mathcal{Z}^A} |z_n^A - z| \leq \eta,$$

where $\mathcal{Z}^A = \{z, F^A(z) \in [1 - q, q]\}$. Recall that $A_1(\mathbf{X}, \Theta), \dots, A_k(\mathbf{X}, \Theta)$ are the consecutive cells containing \mathbf{X} designed with Θ . Observe that, conditionally on the position of the split, data on the left side of the split are still independent and identically distributed according to Proposition 2.1 in [Biau et al. \(2012\)](#). Thus, we have

$$\begin{aligned} \mathbb{P}[G_k(\mathbf{X}, \Theta)^c] &\leq \sum_{\ell=1}^k \mathbb{P}\left[\inf_{z \in \mathcal{Z}^{A_\ell(\mathbf{X}, \Theta)}} |z_n^{A_\ell(\mathbf{X}, \Theta)} - z| > \eta\right] \\ &\leq \sum_{\ell=1}^k \mathbb{E}\left[\mathbb{P}\left[\inf_{z \in \mathcal{Z}^{A_\ell(\mathbf{X}, \Theta)}} |z_n^{A_\ell(\mathbf{X}, \Theta)} - z| > \eta \middle| N(A_\ell(\mathbf{X}, \Theta))\right]\right] \\ &\leq \sum_{\ell=1}^k \mathbb{E}\left[\mathbb{P}\left[\sup_x |F(x) - F_n(x)| \geq \frac{c\eta}{\mu(A)} \middle| N(A_\ell(\mathbf{X}, \Theta)), \right. \right. \\ &\quad \left. \left. A_\ell(\mathbf{X}, \Theta)\right]\right]. \end{aligned}$$

Consequently,

$$\begin{aligned}
\mathbb{P}[G_k(\mathbf{X}, \Theta)^c] &\leq 2 \sum_{\ell=1}^k \mathbb{E} \left[\exp \left(- \frac{2N(A_\ell(\mathbf{X}, \Theta))c^2\eta^2}{\mu(A_\ell(\mathbf{X}, \Theta))^2} \right) \right] \\
&\quad (\text{see } \text{Massart, 1990}) \\
&\leq 2k \exp \left(- 2c^2\eta^2 \left((1-q)^k n - \frac{1}{1-q} \right) \right) \\
&\quad \left(\text{since } \min_{\ell} N(A_\ell(\mathbf{X}, \Theta)) \geq (1-q)^k n - \frac{1}{1-q} \right).
\end{aligned}$$

Thus, for all k and for all $n > \log k / (2c^2\eta^2(1-q)^k)$, we obtain

$$\mathbb{P}[G_k(\mathbf{X}, \Theta)^c] \leq \eta. \quad (22)$$

Gathering (20), (21) and (22), we conclude that, for all $n > \log k_0 / (2c^2\eta^2(1-q)^{k_0})$,

$$\mathbb{E}[\text{diam}(A_k(\mathbf{X}, \Theta))] \leq (k_0 + 1)\eta + \varepsilon.$$

Since the diameter is a non increasing function of the level of the tree, the cell diameter of the fully developed tree (which contain exactly one point in each leaf) is lower than that of the tree stopped at k_0 . Letting $A(\mathbf{X}, \Theta)$ the cell of the (fully developed) empirical q quantile tree (defined in **Algorithm 2**), we have

$$\begin{aligned}
\mathbb{E}[\text{diam}(A_n(\mathbf{X}, \Theta))] &\leq \mathbb{E}[\text{diam}(A_k(\mathbf{X}, \Theta))] \\
&\leq (k_0 + 1)\eta + \varepsilon,
\end{aligned}$$

which concludes the proof, since ε and η can be made arbitrarily small. \square

Proof of Theorem 4.2. We check the conditions of Stone's theorem. Condition (i) is satisfied since the regression function is uniformly continuous and $\text{Var}[Y|\mathbf{X}] = \sigma^2$ (see remark after Stone theorem in Györfi et al., 2002).

Condition (ii) is always satisfied for random trees. Condition (iii) is verified since

$$\mathbb{P}_{\mathbf{X}, \Theta} [\text{diam}(A_n(\mathbf{X}, \Theta)) > \gamma] \xrightarrow{n \rightarrow \infty} 0,$$

according to Lemma 2.

Since each cell contains exactly one data point,

$$\begin{aligned}
\sum_{i=1}^n W_{ni}(x) &= \sum_{i=1}^n \mathbb{E}_{\Theta} \left[\frac{\mathbf{1}_{\mathbf{X}_i \in A_n(\mathbf{X}, \Theta)}}{N_n(\mathbf{X}, \Theta)} \right] \\
&= \mathbb{E}_{\Theta} \left[\frac{1}{N_n(\mathbf{X}, \Theta)} \sum_{i=1}^n \mathbf{1}_{\mathbf{X}_i \in A_n(\mathbf{X}, \Theta)} \right] \\
&= 1.
\end{aligned}$$

Thus, conditions (iv) of Stone theorem is satisfied.

To check (v), observe that in the subsampling step, there are exactly $\binom{a_n-1}{n-1}$ choices to pick a fixed observation \mathbf{X}_i . Since \mathbf{x} and \mathbf{X}_i belong to the same cell only if \mathbf{X}_i is selected in the subsampling step, we see that

$$\mathbb{P}_\Theta [\mathbf{X} \overset{\Theta}{\leftrightarrow} \mathbf{X}_i] \leq \frac{\binom{a_n-1}{n-1}}{\binom{a_n}{n}} = \frac{a_n}{n}.$$

So,

$$\mathbb{E} \left[\max_{1 \leq i \leq n} W_{ni}(\mathbf{X}) \right] \leq \mathbb{E} \left[\max_{1 \leq i \leq n} \mathbb{P}_\Theta [\mathbf{X} \overset{\Theta}{\leftrightarrow} \mathbf{X}_i] \right] \leq \frac{a_n}{n},$$

which tends to zero by assumption. \square

5.6 Proofs of Technical Lemmas 2 and 1

Technical Lemma 1. *Take $k \in \mathbb{N}$ and consider a uniform random forest where each tree is stopped at level k . For all $\mathbf{x}, \mathbf{z} \in [0, 1]^d$, its connection function satisfies*

$$K_k(0, |\mathbf{x} - \mathbf{z}|) \leq K_k(\mathbf{x}, \mathbf{z}),$$

where $|\mathbf{x} - \mathbf{z}| = (|x_1 - z_1|, \dots, |x_d - z_d|)$.

Proof. Take $x, z \in [0, 1]$. Without loss of generality, one can assume that $x < z$ and let $\mu = z - x$. Consider the following two configurations.




Figure 3: Scheme of configuration 1 (at the top) and 2 (at the bottom).

For any $k \in \mathbb{N}^*$, we let $\mathbf{d}_k = (d_1, \dots, d_k)$ (resp. $\mathbf{d}'_k = (d'_1, \dots, d'_k)$) be k consecutive cuts in configuration 1 (resp. in configuration 2). We denote by \mathcal{A}_k (resp. \mathcal{A}'_k) the set where \mathbf{d}_k (resp. \mathbf{d}'_k) belong.

We show that for all $k \in \mathbb{N}^*$, there exists a coupling between \mathcal{A}_k and \mathcal{A}'_k satisfying the following property: any k -tuple \mathbf{d}_k is associated with a k -tuple \mathbf{d}'_k such that

1. if \mathbf{d}_k separates $[x, z]$ then \mathbf{d}'_k separates $[0, z - x]$,
2. if \mathbf{d}_k does not separate $[x, z]$ and \mathbf{d}'_k does not separate $[0, z - x]$, then the length of the cell built with \mathbf{d}_k is higher than the one built with \mathbf{d}'_k .

We call \mathcal{H}_k this property. We now proceed by induction. For $k = 1$, we use the function g to map \mathcal{A}_1 into \mathcal{A}'_1 such that:

$$g_1(u) = \begin{cases} u & \text{if } u > z \\ z - u & \text{if } u \leq z \end{cases}$$

Thus, for any $d_1 \in \mathcal{A}_1$, if d_1 separates $[x, z]$, then $d'_1 = g_1(d_1)$ separates $[0, z - x]$. Besides, the length of the cell containing $[x, z]$ designed with the cut d_1 is higher than that of the cell containing $[0, z - x]$ designed with the cut d'_1 . Consequently, \mathcal{H}_1 is true.

Now, take $k > 1$ and assume that \mathcal{H}_k is true. Consequently, if \mathbf{d}_k separates $[x, z]$ then $g_k(\mathbf{d}_k)$ separates $[0, z - x]$. In that case, \mathbf{d}_{k+1} separates $[x, z]$ and $g_{k+1}(\mathbf{d}_{k+1})$ separates $[0, z - x]$. Thus, in the rest of the proof, we assume that \mathbf{d}_k does not separate $[x, z]$ and $g_k(\mathbf{d}_k)$ does not separate $[0, z - x]$. Let $[a_k, b_k]$ be the cell containing $[x, z]$ built with cuts \mathbf{d}_k . Since the problem is invariant by translation, we assume, without loss of generality, that $[a_k, b_k] = [0, \delta_k]$, where $\delta_k = b_k - a_k$ and $[x, z] = [x_k, x_k + \mu]$ (see Figure 4).




Figure 4: Configuration 1a (at the top) and 1b (at the bottom).

In addition, according to \mathcal{H}_k , the length of the cell built with \mathbf{d}_k is higher than the one built with \mathbf{d}'_k . Thus, one can find $\lambda \in (0, 1)$ such that $d'_k = \lambda \delta_k$. This is summarized in Figure 5.




Figure 5: Configuration 1b (at the top) and 2b (at the bottom).

Thus, one can map $[0, \delta_k]$ into $[0, \lambda \delta_k]$ with g_{k+1} defined as

$$g_{k+1}(u) = \begin{cases} \lambda u & \text{if } u > x_k + \mu \\ \lambda(x_k + \mu - u) & \text{if } u \leq x_k + \mu \end{cases}$$

Note that, for all d_{k+1} , the length of the cell containing $[x_k, x_k + \mu]$ designed with the cut d_{k+1} (configuration 1b) is bigger than the length of the cell containing $[0, \mu]$ designed with the cut $d'_{k+1} = g_{k+1}(d_{k+1})$ (configuration 2b). Besides, if $d_{k+1} \in [x_k, x_k + \mu]$ then $g_{k+1}(d_{k+1}) \in [0, \mu]$. Consequently, the set of functions g_1, \dots, g_{k+1} induce a mapping of \mathcal{A}_{k+1} into \mathcal{A}'_{k+1} such that \mathcal{H}_{k+1} holds. Thus, Technical Lemma 1 holds for $d = 1$.

To address the case where $d > 1$, note that

$$\begin{aligned}
K_k(\mathbf{x}, \mathbf{z}) &= \sum_{\substack{k_1, \dots, k_d \\ \sum_{j=1}^d k_j = k}} \frac{k!}{k_1! \dots k_d!} \left(\frac{1}{d}\right)^k \prod_{m=1}^d K_{k_m}(x_m, z_m) \\
&\geq \sum_{\substack{k_1, \dots, k_d \\ \sum_{j=1}^d k_j = k}} \frac{k!}{k_1! \dots k_d!} \left(\frac{1}{d}\right)^k \prod_{m=1}^d K_{k_m}(0, |z_m - x_m|) \\
&\geq K_k(0, |\mathbf{z} - \mathbf{x}|),
\end{aligned}$$

which concludes the proof. \square

Technical Lemma 2. Take $k \in \mathbb{N}$ and consider a uniform random forest where each tree is stopped at level k . For all $x \in [0, 1]$, its connection function $K_k(0, x)$ satisfies

$$K_k(0, x) = 1 - x \sum_{j=0}^{k-1} \frac{(-\ln x)^j}{j!},$$

with the notational convention that the last sum is zero if $k = 0$.

Proof of Technical Lemma 2. The result is clear for $k = 0$. Thus, set $k \in \mathbb{N}^*$ and consider a uniform random forest where each tree is stopped at level k . Since the result is clear for $x = 0$, take $x \in]0, 1]$ and let $I = [0, x]$. Thus

$$\begin{aligned}
K_k(0, x) &= \mathbb{P} \left[0 \overset{\Theta}{\underset{k \text{ cuts}}{\longleftrightarrow}} x \right] \\
&= \int_{z_1 \notin I} \int_{z_2 \notin I} \dots \int_{z_k \notin I} p(\mathrm{d}z_k | z_{k-1}) p(\mathrm{d}z_{k-1} | z_{k-2}) \dots p(\mathrm{d}z_2 | z_1) p(\mathrm{d}z_1),
\end{aligned}$$

where z_1, \dots, z_k are the positions of the k cuts (see Figure 6).




Figure 6: Positions of cuts z_1, \dots, z_k and x with $d = 1$

We prove by induction that, for every integer p ,

$$\begin{aligned}
&\int_{z_{k-p} \notin I} \dots \int_{z_k \notin I} p(\mathrm{d}z_k | z_{k-1}) \dots p(\mathrm{d}z_{k-p} | z_{k-p-1}) \\
&= 1 - \frac{x}{z_{k-p-1}} \left(\sum_{j=0}^p \frac{[\ln(z_{k-p-1}/x)]^j}{j!} \right).
\end{aligned}$$

Denote by \mathcal{H}_p this property. Since, given z_{k-1} , z_k is uniformly distributed over $[0, z_{k-1}]$, we have

$$\int_{z_k \notin I} p(dz_k | z_{k-1}) = 1 - \frac{x}{z_{k-1}}.$$

Thus \mathcal{H}_0 is true. Now, fix $p > 0$ and assume that \mathcal{H}_p is true. Let $u = z_{k-p-1}/x$. Thus, integrating both sides of \mathcal{H}_p , we deduce,

$$\begin{aligned} & \int_{z_{k-p-1} \notin I} \int_{z_{k-p} \notin I} \cdots \int_{z_k \notin I} p(dz_k | z_{k-1}) \cdots p(dz_{k-p} | z_{k-p-1}) p(dz_{k-p-1} | z_{k-p-2}) \\ &= \int_{z_{k-p-1} \notin I} \left[1 - \frac{x}{z_{k-p-1}} \left(\sum_{j=0}^p \frac{[\ln(z_{k-p-1}/x)]^j}{j!} \right) \right] p(dz_{k-p-1} | z_{k-p-2}) \\ &= \int_x^{z_{k-p-2}} \left[1 - \frac{x}{z_{k-p-1}} \left(\sum_{j=0}^p \frac{[\ln(z_{k-p-1}/x)]^j}{j!} \right) \right] \frac{dz_{k-p-1}}{z_{k-p-2}} \\ &= \frac{x}{z_{k-p-2}} \int_1^{z_{k-p-2}/x} \left[1 - \frac{1}{u} \left(\sum_{j=0}^p \frac{[\ln(u)]^j}{j!} \right) \right] du. \end{aligned}$$

Using integration by parts on the last term, we conclude that \mathcal{H}_{p+1} is true. Thus, for all $p > 0$, \mathcal{H}_p is verified. Finally, using \mathcal{H}_{k-1} and the fact that $z_0 = 1$, we conclude the proof. \square

References

- G. Biau. Analysis of a random forests model. *Journal of Machine Learning Research*, 13:1063–1095, 2012.
- G. Biau, L. Devroye, and G. Lugosi. Consistency of random forests and other averaging classifiers. *Journal of Machine Learning Research*, 9:2015–2033, 2008.
- G. Biau, F. Cérou, and A. Guyader. New insights into approximate bayesian computation. arXiv:1207.6461, 2012.
- S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford University Press, 2013.
- L. Breiman. Bagging predictors. *Machine Learning*, 24:123–140, 1996.
- L. Breiman. Randomizing outputs to increase prediction accuracy. *Machine Learning*, 40:229–242, 2000.
- L. Breiman. Random forests. *Machine Learning*, 45:5–32, 2001.
- L. Breiman, J. Friedman, R.A. Olshen, and C.J. Stone. *Classification and Regression Trees*. Chapman & Hall, New York, 1984.
- S. Cléménçon, M. Depecker, and N. Vayatis. Ranking forests. *Journal of Machine Learning Research*, 14:39–73, 2013.

- A. Cutler and G. Zhao. Pert - perfect random tree ensembles. *Computing Science and Statistics*, 33:490–497, 2001.
- M. Denil, D. Matheson, and N. de Freitas. *Consistency of online random forests*. arXiv:1302.4853, 2013.
- L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, New York, 1996.
- R. Díaz-Uriarte and S. Alvarez de Andrés. Gene selection and classification of microarray data using random forest. *BMC Bioinformatics*, 7:1–13, 2006.
- T. G. Dietterich and E. B. Kong. Machine learning bias, statistical bias, and statistical variance of decision tree algorithms. Technical report, Department of Computer Science, Oregon State University, 1995.
- R. Genuer, J.-M. Poggi, and C. Tuleau. Random forests: some methodological insights. arXiv:0811.3619, 2008.
- P. Geurts, D. Ernst, and L. Wehenkel. Extremely randomized trees. *Machine Learning*, 63:3–42, 2006.
- L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer, New York, 2002.
- M. Hamza and D. Laroque. An empirical comparison of ensemble methods based on classification trees. *Journal of Statistical Computation and Simulation*, 75: 629–643, 2005.
- T. Ho. The random subspace method for constructing decision forests. *Pattern Analysis and Machine Intelligence*, 20(8), 832–844, 1998.
- H. Ishwaran and U.B. Kogalur. Consistency of random survival forests. *Statistics & Probability Letters*, 80:1056–1064, 2010.
- H. Ishwaran, U.B. Kogalur, E.H. Blackstone, and M.S. Lauer. Random survival forest. *The Annals of Applied Statistics*, 2:841–860, 2008.
- B. Lakshminarayanan, D. M. Roy, and Y. W. Teh. Mondrian forests: Efficient online random forests. arXiv:1406.2673, 2014.
- A. Liaw and M. Wiener. Classification and regression by randomforest. *R news*, 2:18–22, 2002.
- P. Massart. The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *The Annals of Probability*, pages 1269–1283, 1990.
- N. Meinshausen. Quantile regression forests. *Journal of Machine Learning Research*, 7:983–999, 2006.
- L. Mentch and G. Hooker. Ensemble trees and clts: Statistical inference for supervised learning. arXiv:1404.6473, 2014.
- Y. Qi. *Ensemble Machine Learning*, chapter Random forest for bioinformatics, pages 307–323. Springer, 2012.

- G. Rogez, J. Rihan, S. Ramalingam, C. Orrite, and P. H. Torr. Randomized trees for human pose detection. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.
- E. Scornet, G. Biau, and J.-P. Vert. Consistency of random forests. arXiv:1405.2881, 2014.
- C.J. Stone. Consistent nonparametric regression. *The Annals of Statistics*, 5: 595–645, 1977.
- Aad W. van der Vaart and Jon A. Wellner. *Weak Convergence and Empirical Processes : With Applications to Statistics*. Springer, New York, 1996.
- S. Wager. Asymptotic theory for random forests. arXiv:1405.0352, 2014.
- S. Wager, T. Hastie, and B. Efron. Confidence intervals for random forests: The jackknife and the infinitesimal jackknife. *Journal of Machine Learning Research*, 15:1625–1651, 2014.
- R. Zhu, D. Zeng, and M.R. Kosorok. *Reinforcement learning trees*. Technical Report, University of North Carolina, 2012.