



HAL
open science

Random forest estimation of conditional distribution functions and conditional quantiles

Kévin Elie-Dit-Cosaque, Véronique Maume-Deschamps

► **To cite this version:**

Kévin Elie-Dit-Cosaque, Véronique Maume-Deschamps. Random forest estimation of conditional distribution functions and conditional quantiles. *Electronic Journal of Statistics*, 2022, 16 (2), 10.1214/22-EJS2094 . hal-02733460v5

HAL Id: hal-02733460

<https://hal.science/hal-02733460v5>

Submitted on 4 Feb 2023

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.

Random Forest Estimation of Conditional Distribution Functions and Conditional Quantiles

Kévin Elie-Dit-Cosaque

SCOR, France
e-mail: edckev@gmail.com

Véronique Maume-Deschamps

*Univ Lyon, Université Claude Bernard Lyon 1,
CNRS UMR 5208, Institut Camille Jordan,
F-69622 Villeurbanne, France*
e-mail: veronique.maume@univ-lyon1.fr

Abstract: We propose a theoretical study of two realistic estimators of conditional distribution functions using random forests. The estimation process uses the bootstrap samples generated from the original dataset when constructing the forest. Bootstrap samples are reused to define the first estimator, while the second uses the original sample, once the forest has been built. We prove that both proposed estimators of the conditional distribution functions are consistent uniformly a.s. To the best of our knowledge, it is the first proof of a.s. consistency (previous consistency results are in L^2 norm or in probability) and including the bootstrap part. The consistency result holds for a large class of functions, including additive models and products. The consistency of conditional quantiles estimators follows that of distribution functions estimators using standard arguments.

MSC2020 subject classifications: Random forests, consistency, conditional distribution functions, conditional quantiles.

Contents

1	Introduction	2
2	Breiman's random forest	4
3	Conditional Distribution Forests	6
3.1	Bootstrap samples based estimator	6
3.2	Original sample based estimator	7
4	Consistency results	8
5	On variation in Random Forest leaves	12
5.1	Some comments on Assumption 4.2	12
5.2	Sufficient conditions to control the variation on leaves	13
6	Proofs of Theorem 4.2	19
7	Conclusion	23
A	Technical complements	24

Acknowledgments 30
References 30

1. Introduction

Conditional distribution functions and conditional quantiles estimation is an important task in several domains including environment, insurance or industry. It is also an important tool for Quantile Oriented Sensitivity Analysis (QOSA), see e.g., [16, 23, 9, 14]. In order to estimate conditional quantiles, various methods exist such as kernel based estimation or quantile regression [21] but they present some limitations. Indeed, the performance of kernel methods strongly depends on the bandwidth parameter selection and quickly breaks down as the number of covariates increases. On the other hand, quantile regression as presented in [21] may not be adapted in a non-gaussian setting since the true conditional quantile is not necessarily a linear combination of the input variables [24]. To overcome these issues, we propose to explore the random forest estimation of conditional quantiles [25].

Random forest algorithms allow a flexible modeling of interactions in high dimension by building a large number of regression trees and averaging their predictions. The most famous random forest algorithm is that of [6] whose construction is based on the seminal work of [1, 19, 11]. Breiman’s random forest estimate is a combination of two essential components: Bagging and Classification And Regression Trees (CART)-split criterion [8]. Bagging for *bootstrap-aggregating* was proposed by [5] in order to improve the performance of weak or unstable learners. Two types of randomness are included: a bootstrap sample is used to construct each tree and at each split step, some input variables are chosen randomly in $\{X_1, \dots, X_d\}$.

Random forests are also related to some local averaging algorithms such as nearest neighbors methods [22, 3] or kernel estimators [33]. More precisely, thanks to [22], the random forest method can be seen as an adaptive neighborhood regression procedure and therefore the prediction (estimation of the conditional mean) can be formulated as a weighted average of the observed response variables.

Based on that approach, we develop a Weighted Conditional Empirical Cumulative Distribution Function (W_C_ECDF) approximating the Conditional Cumulative Distribution Function (C_CDF). Then, α -quantile estimates are obtained by using W_C_ECDF instead of C_CDF. [25] defined a W_C_ECDF with weights using the original dataset whereas we allow to construct the weights using the bootstrap samples, as it is done practically in regression random forests. We prove the almost sure consistency of these two estimators. The main hypothesis are:

1. the convergence to 0 of the variation of the conditional distribution function on leaves;
2. a control of the size of the leaves, implying that the trees are not fully developed.

In Section 5, we prove that the first hypothesis is satisfied for a large class of models, including additive models and product functions. Also, this class of functions is dense in the set of non-negative continuous functions on $[0, 1]^d$. Moreover, the first hypothesis is satisfied for some modified version of the CART algorithm, such as in [25]. The second hypothesis may be seen as a stopping rule in the algorithm, so that it should be controlled by practitioners.

An implementation of both algorithms is made available within a Julia package called `ConditionalDistributionForest` [15] as well as a Python package named `qosa-indices`, [13].

The `C_CDF` can be seen as a regression function. Several authors such as [7, 4, 38, 35, 27, 37, 17] have established asymptotic properties of particular variants and simplifications of the original Breiman’s random forest algorithm. For instance, in [4], the tree construction is done by choosing at random the splitting coordinate and once the latter is selected, the split is at the midpoint of the chosen side. Facing some theoretical issues with the bootstrap, most studies replace it by subsampling, assuming that each tree is grown with $s_n < n$ observations randomly chosen without replacement from the original dataset. Most of the time, in order to ensure the convergence of the simplified model, the subsampling rate s_n/n is assumed to tend to zero at some prescribed rate, assumption that excludes the bootstrap mode. Besides, consistency is generally showed by assuming that the number of trees goes to infinity which is not fully relevant in practice. Under some conditions, [31] showed that if the infinite random forest regression estimator is \mathbb{L}^2 consistent then so does the finite random forest regression estimator when the number of trees goes to infinity in a controlled way.

Recent attempts to bridge the gap between theory and practice, provide some results on random forest algorithms at the price of fairly strong conditions. For example, [35] showed the \mathbb{L}^2 consistency of random forests in an additive regression framework by replacing the bootstrap step by subsampling. Their result rests on a fundamental lemma developed in [34] which reviews theoretical random forest, the additive assumption being required here. Furthermore, consistency and asymptotic normality of the whole algorithm were recently proved under strong conditions by [37] replacing bootstrap by subsampling and simplifying the splitting step. One of the strong conditions used in Theorem 3.1. of [37] is that the individual trees satisfy a condition called *honesty*. An example of an honest tree given by the authors is *one where the tree is grown using one subsample, while the predictions at the leaves of the tree are estimated using a different subsample*. Due to this assumption, the authors admit that their theorems are not valid for the practical applications most of the time because almost all implementations of random forests use the training sample twice. In a recent work [20], the \mathbb{L}^2 consistency in a regression framework is also obtained for the expectation over all possible trees of the regression function (infinite forest), the bootstrap process is not taken into account. The result in [20] holds for a class of regression functions broader than the additive model. This class of functions requires some control on the partial derivatives of the model.

Thus, despite an active investigation during the last decade, further consis-

tenacy results are still welcome. Our major contribution is the proof of the almost sure uniform convergence of the estimator `W_C_ECDF` both using the bootstrap samples (Theorem 4.2) or the original one (Theorem 4.3). Remark that [25] gave a proof of the consistency in probability of the `W_C_ECDF` for a simplified model where the weights are considered as constant while they are indeed random variables heavily data-dependent. We provide an a.s. consistency proof under realistic assumptions for a method based on bootstrap samples. This consistency result holds for a large class of functions called **♠-class** which contains, for instance, additive functions, product functions. Our **♠-class** also contains some dense class of functions, as does the class considered in [20]. Contrary to [20], our result holds for finite random forests and takes the bootstrap into account, so that it is closer to practical situations. Also, a sub-product of the consistency proof is an asymptotic proximity result between empirical and theoretical trees (see Proposition 5.4), which is interesting in itself for further theoretical studies on random forests.

The paper is organized as follows. Breiman’s random forest algorithm is detailed in Section 2 and notations are stated. The random forest estimations of `C_CDF` based both on bootstrap samples and the original dataset are introduced in Section 3 as a natural generalization of regression random forests. The main consistency results are presented in Section 4 and a deep discussion on the variation of the conditional distribution function on leaves is proposed in Section 5. The main proofs are gathered in Section 6 and some of them are postponed in Appendix A. A short conclusion is given in Section 7.

2. Breiman’s random forest

The aim of this section is to present Breiman’s random forest algorithm as well as notations used throughout this paper.

Random forest is a generic term to name an aggregation scheme of decision trees allowing to deal with both supervised classification and regression tasks and we focus on the latter in this paper.

The general framework is the nonparametric regression estimation where an input random vector $\mathbf{X} \in \mathcal{X} = \prod_{i=1}^d [u_i, v_i] \subset \mathbb{R}^d$ is observed and a response $Y \in \mathbb{R}$

is predicted by estimating the regression function $m(\mathbf{x}) = \mathbb{E}[Y | \mathbf{X} = \mathbf{x}]$. We assume that we are given a training sample $\mathcal{D}_n = (\mathbf{X}^j, Y^j)_{j=1, \dots, n}$ of independent random variables distributed as the prototype pair (\mathbf{X}, Y) which is a $(d + 1)$ -dimensional random vector. The purpose is to use the dataset \mathcal{D}_n to construct an estimator $m_n : \mathcal{X} \rightarrow \mathbb{R}$ of the function m .

Random forests proposed by [6] build a predictor consisting of a collection of k randomized regression trees grown based on the CART algorithm.

The CART-split criterion of [8] is used in the construction of the individual trees to recursively partition the input space \mathcal{X} in a dyadic manner. More precisely, at each step of the partitioning, a part of the space is divided into two sub-parts according to the best cut perpendicular to the axes. This best cut is selected in

each node of the tree by optimizing the CART-split criterion over the d variables, i.e. minimizing the prediction squared error in the two child nodes. The trees are then grown recursively until reaching a stopping rule. There are several rules, but one generally proposed is that the tree construction continues while leaves contain at least $min_samples_leaf$ elements. This criterion is implemented in the `RandomForestRegressor` class of the python package `Scikit-Learn` [29] or in the `build_forest` function of the Julia [2] package `DecisionTree`.

Building several different trees from a single dataset requires to randomize the tree building process. [6] proposed to inject some randomness both in the dataset and in the tree construction. First of all, prior to the construction of each tree, a resampling step is done by bootstrapping [12] from the original dataset, that is, by randomly sampling n observations with replacement. Only these bootstrap observations are taken into account in the tree building. Accordingly, the $min_samples_leaf$ hyperparameter introduced previously refers, in the random forest method, to the minimum number of bootstrap observations (counted with repetitions) contained in each leaf of a tree. Secondly, at each step of the tree construction, instead of optimizing the CART-split criterion over the d variables, a number of variables called $max_features$ is selected uniformly at random among the d variables, the set of selected variables is denoted \mathcal{M}_{TRY} . In what follows, we shall assume that the $max_features$ co-variables are selected randomly, with positive probability for each co-variate to be selected, which includes of course the uniform selection. Then, the best split is chosen as the one

optimizing the CART-split criterion as follows: for a given node $A = \prod_{i=1}^d [a_i, b_i]$, the CART-split is given by maximizing over $j \in \mathcal{M}_{\text{TRY}}$ and $z \in A^j = [a_j, b_j]$

$$L_A^n(j, z) = \frac{1}{N_n^b(A)} \sum_{i=1}^n (Y^i - \bar{Y}_A)^2 \mathbb{1}_{\{\mathbf{X}^i \in A\}} - \frac{1}{N_n^b(A)} \sum_{i=1}^n \left(Y^i - \bar{Y}_{A_L} \mathbb{1}_{\{X_j^i \leq z\}} - \bar{Y}_{A_R} \mathbb{1}_{\{X_j^i > z\}} \right)^2 \mathbb{1}_{\{\mathbf{X}^i \in A\}}, \quad (2.1)$$

where $A_L = \{\mathbf{x} \in A, x_j \leq z\}$ (resp. A_R) is the left (resp. right) child of A , \bar{Y}_B is the mean of the Y^i 's in the bootstrap sample with $\mathbf{X}^i \in B$ and $N_n^b(B)$ denotes the number of elements in the bootstrap sample belonging to B .

Now, for any query point $\mathbf{x} \in \mathcal{X}$, the ℓ -th tree estimates $m(\mathbf{x})$ as follows

$$m_n^b(\mathbf{x}; \Theta_\ell, \mathcal{D}_n) = \sum_{j \in \mathcal{D}_n^*(\Theta_\ell)} \frac{\mathbb{1}_{\{\mathbf{x}^j \in A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)\}}}{N_n^b(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)} Y^j, \quad (2.2)$$

where:

- $\Theta_\ell, \ell = 1, \dots, k$ are independent random vectors, distributed as a generic random vector $\Theta = (\Theta^1, \Theta^2)$ and independent of \mathcal{D}_n . Θ^1 contains indexes of observations that are used to build each tree, i.e. the bootstrap sample and Θ^2 indexes of splitting candidate variables in each node, we assume that Θ^2 gives a positive probability to each co-variate,

- \mathcal{D}_n^* (Θ_ℓ) is the bootstrap sample selected prior to the tree construction,
- $A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)$ is the tree cell (subspace of \mathcal{X}) containing \mathbf{x} ,
- $N_n^b(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)$ is the number of elements of $\mathcal{D}_n^*(\Theta_\ell)$ that fall into $A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)$.

The trees are then combined to form the finite forest estimator

$$m_{k,n}^b(\mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n) = \frac{1}{k} \sum_{\ell=1}^k m_n^b(\mathbf{x}; \Theta_\ell, \mathcal{D}_n) . \quad (2.3)$$

We may now present the conditional distribution function estimators.

3. Conditional Distribution Forests

We aim to estimate $F(y|\mathbf{x}) = \mathbb{P}(Y \leq y | \mathbf{X} = \mathbf{x})$. Two estimators may be defined. One uses the bootstrap samples both in the forest construction and in the estimation. The other uses the original sample in the estimation part. Once the distribution function has been estimated, the conditional quantiles may be estimated straightforwardly.

3.1. Bootstrap samples based estimator

First of all, let us define the random variable $B_j(\Theta_\ell^1, \mathcal{D}_n)$ as the number of times that the observation (\mathbf{X}^j, Y^j) has been drawn from the original dataset in the bootstrap sample for the ℓ -th tree construction. The conditional mean estimator in Equation (2.3) rewrites as

$$\begin{aligned} m_{k,n}^b(\mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n) &= \sum_{j=1}^n \left(\frac{1}{k} \sum_{\ell=1}^k \frac{B_j(\Theta_\ell^1, \mathcal{D}_n) \mathbb{1}_{\{\mathbf{x}^j \in A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)\}}}{N_n^b(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)} \right) Y^j \\ &= \sum_{j=1}^n w_{n,j}^b(\mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n) Y^j , \end{aligned}$$

where the weights are defined by

$$w_{n,j}^b(\mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n) = \frac{1}{k} \sum_{\ell=1}^k \frac{B_j(\Theta_\ell^1, \mathcal{D}_n) \mathbb{1}_{\{\mathbf{x}^j \in A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)\}}}{N_n^b(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)} . \quad (3.1)$$

Note that the weights $w_{n,j}^b(\mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n)$ are non-negative random variables as functions of $\Theta_1, \dots, \Theta_k, \mathcal{D}_n$ and their sum for $j = 1, \dots, n$ equals 1.

The random forest estimator (3.1) can be seen as a local averaging estimate. Indeed, as mentioned by [32], the regression trees make an average of the observations located in a neighborhood of \mathbf{x} , this neighborhood being defined as the leaf of the tree containing \mathbf{x} . The forest, which aggregates several trees, also operates by calculating a weighted average of the observations in a neighborhood of \mathbf{x} . However, in the case of forests, this neighborhood results from the

superposition of the neighborhoods of each tree, and therefore has a more complex shape. Several works have tried to study the random forest algorithm from this point of view (local averaging estimate) such as [22] who was the first to point out the connection between the random forest and the adaptive nearest-neighbors methods, further developed by [3]. Some works such as [33] have also studied random forests through their link with the kernel methods.

We are interested in the Conditional Cumulative Distribution Function (C_CDF) of Y given $\mathbf{X} = \mathbf{x}$ in order to obtain the conditional quantiles. Pairing the following equality

$$F(y|\mathbf{X} = \mathbf{x}) = \mathbb{P}(Y \leq y | \mathbf{X} = \mathbf{x}) = \mathbb{E} [\mathbb{1}_{\{Y \leq y\}} | \mathbf{X} = \mathbf{x}] , \quad (3.2)$$

with the weighted approach described above, we propose to estimate the C_CDF as follows

$$F_{k,n}^b(y|\mathbf{X} = \mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n) = \sum_{j=1}^n w_{n,j}^b(\mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n) \mathbb{1}_{\{Y^j \leq y\}} . \quad (3.3)$$

Hence, given a level $\alpha \in]0, 1[$, the conditional quantile estimator $\hat{q}^\alpha(Y|\mathbf{X} = \mathbf{x})$ is defined as follows

$$\hat{q}^\alpha(Y|\mathbf{X} = \mathbf{x}) = \inf \{Y^p, p = 1, \dots, n : F_{k,n}^b(Y^p|\mathbf{X} = \mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n) \geq \alpha\} .$$

Let us turn now to the estimator using the original sample.

3.2. Original sample based estimator

Trees are still grown with their respective bootstrap sample $\mathcal{D}_n^*(\Theta_\ell)$, $\ell = 1, \dots, k$. But instead of considering them in the estimation, we may use the original sample \mathcal{D}_n . Consider the weights

$$w_{n,j}^o(\mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n) = \frac{1}{k} \sum_{\ell=1}^k \frac{\mathbb{1}_{\{\mathbf{X}^j \in A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)\}}}{N_n^o(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)} , \quad (3.4)$$

where $N_n^o(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)$ is the number of points of \mathcal{D}_n that fall into $A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)$. As previously, the weights $w_{n,j}^o(\mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n)$ are non-negative random variables as functions of $\Theta_1, \dots, \Theta_k, \mathcal{D}_n$ and their sum over $j = 1, \dots, n$ equals 1.

It was proposed in [25] to estimate the C_CDF with

$$F_{k,n}^o(y|\mathbf{X} = \mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n) = \sum_{j=1}^n w_{n,j}^o(\mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n) \mathbb{1}_{\{Y^j \leq y\}} . \quad (3.5)$$

The conditional quantiles are then estimated by plugging

$F_{k,n}^o(y|\mathbf{X} = \mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n)$ instead of $F(y|\mathbf{X} = \mathbf{x})$ as before.

A python library named `qosa-indices` has also been developed to perform the numerical estimations of conditional distributions and quantiles for

both methods. It is available at [13] and uses `Scikit-Learn`, `Numpy`, `Numba`. Both approaches are also implemented in a `Julia` package based on the library `DecisionTree` and that is available at [15].

It has to be noted that a package called `quantregForest` has been made available in `R` [30] and can be found at [26]. The estimation method currently implemented in `quantregForest` is different from the method described in [25]. As a matter of fact, for a new observation \mathbf{x} and the ℓ -th tree, one element of $\mathcal{D}_n = (\mathbf{X}^j, Y^j)_{j=1, \dots, n}$ falling into in the leaf node $A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)$ is chosen at random. This gives, k values of Y and allows to estimate the conditional distribution function with the classical Empirical Cumulative Distribution Function associated with the empirical measure.

4. Consistency results

In this section, we state our main results, which are the uniform a.s. consistency of both estimators $F_{k,n}^b$ and $F_{k,n}^o$ of the conditional distribution function. It constitutes the most interesting result of this paper because it handles the bootstrap component and gives the almost sure uniform convergence. Indeed, most of the studies [35, 37, 17] replaces the bootstrap by subsampling without replacement in order to avoid the mathematical difficulties induced by this one and therefore differs slightly from the procedure used in practice. In [20], the bootstrap procedure is not analyzed and the \mathbb{L}^2 consistency is proved for the expectation over Θ of the estimator of the regression function, which does not lead directly to the result for finite forests.

[25] showed the uniform convergence in probability of a simplified version of the estimator $F_{k,n}^o$. In [25], the weights $w_{n,j}^o(\mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n)$ are in fact considered to be non-random while they are indeed random variables depending on $(\Theta_\ell)_{\ell=1, \dots, k}$ and \mathcal{D}_n .

Overall, proving the consistency of the forest methods whose construction depends both on the \mathbf{X}^j 's and on the Y^j 's is a difficult task. This feature makes the resulting estimate highly data-dependent, and therefore difficult to analyze. A simplification widely used by most authors from a theoretical point of view is to work with random forest estimates whose tree shape depends only on \mathbf{X}^j 's which [10] called the \mathbf{X} -property but the Y^j 's are still used to compute the prediction, either the conditional mean or the conditional distribution function for example. One of the first results dealing with data-dependent random forest estimator of the regression function is [35] who showed the \mathbb{L}^2 consistency in an additive regression framework by replacing the bootstrap by subsampling. Our main result is stated in Theorems 4.2 and 4.3 hereafter but we first rewrite it in a more explicit way in order to emphasize their applicability. We begin by a description of the class of functions on which our result holds, then we state the hypothesis and the result.

In what follows, \mathcal{X} is a compact hyper-rectangle of \mathbb{R}^d :

$$\mathcal{X} = \prod_{i=1}^d [u_i, v_i], \quad -\infty < u_i \leq v_i < \infty \text{ and we denote by } \mathcal{A} \text{ the set of hyper-}$$

rectangles in \mathcal{X} : $A \in \mathcal{A}$ writes $A = \prod_{i=1}^d [a_i, b_i]$ with $u_i \leq a_i \leq b_i \leq v_i$. Also, we denote by $A^{-j} = \prod_{k \neq j} [a_k, b_k]$ and $A^J = \prod_{k \in J} [a_k, b_k]$ for any $J \subset \{1, \dots, d\}$. Given $\mathbf{x} \in \mathbb{R}^d$, \mathbf{x}^{-j} is the vector of \mathbb{R}^{d-1} where the j -th coordinate has been removed and \mathbf{x}^J is the vector of \mathbb{R}^J whose coordinates are x_j , $j \in J$.

Definition 4.1. Let $f : \mathcal{X} \rightarrow \mathbb{R}$, it does NOT belong to the \spadesuit -class if there exists a rectangle $A = \prod_{j=1}^d [a_j, b_j] \subset \mathcal{X}$, with $a_j \leq b_j$ such that for all $j = 1, \dots, d$, $z \mapsto \mathbb{E} [f(z, \mathbf{X}^{-j}) \mathbb{1}_{\{\mathbf{X}^{-j} \in A^{-j}\}}]$ is constant on $[a_j, b_j]$ and f is not constant on A .

Remark 4.1. From Definition 4.1, if f belongs to the \spadesuit -class then either for any rectangle A , there exists $j = 1, \dots, d$ such that $z \mapsto \mathbb{E} [f(z, \mathbf{X}^{-j}) \mathbb{1}_{\{\mathbf{X}^{-j} \in A^{-j}\}}]$ is not constant on $[a_j, b_j]$ or if for some rectangle A , $z \mapsto \mathbb{E} [f(z, \mathbf{X}^{-j}) \mathbb{1}_{\{\mathbf{X}^{-j} \in A^{-j}\}}]$ is constant for all $j = 1, \dots, d$ then f is constant on A .

Let us give some examples of functions of the \spadesuit -class, the proof is more or less straightforward, a sketch of proof is given in the Appendix (see Lemma A.3).

Remark 4.2. The \spadesuit -class contains:

1. additive functions $f(\mathbf{x}) = \sum_{j=1}^d f_j(x_j)$,
2. product functions $f(\mathbf{x}) = \prod_{j=1}^d f_j(x_j)$ provided that for all $j = 1, \dots, d$, $\mathbb{E} \left[\prod_{k \neq j} f_k(X_k) \mathbb{1}_{\{X_k \in [a_k, b_k]\}} \right] \neq 0$,
3. sums of product functions $f(\mathbf{x}) = \sum_{I \in \mathcal{I}} \prod_{j \in I} f_j^I(x_j)$ with \mathcal{I} a partition of $\{1, \dots, d\}$ provided that for all $j = 1, \dots, d$, for $\mathcal{I} \ni I \ni j$, $\mathbb{E} \left[\prod_{k \in I, k \neq j} f_k^I(X_k) \mathbb{1}_{\{X_k \in [a_k, b_k]\}} \right] \neq 0$,
4. sums of product functions $f(\mathbf{x}) = \sum_{I \in \mathcal{I}} \prod_{j \in I} f_j^I(x_j)$ with I a family of subsets of $\{1, \dots, d\}$ provided that the f_j^I are either all increasing or all decreasing and they are either all positive or all negative,
5. other functions, such as - for example - in dimension 2 with $(X_1, X_2) \rightsquigarrow \mathcal{U}([0, 1]^2)$ - $f(x, y) = \ln(x + y)$.

A particular case of item 4) above is the set of linear combination of Gaussian radial basis functions on $[0, 1]^d$, with positive weights:

$$\mathcal{G} = \left\{ \sum_{i=1}^p a_i \exp \left[\sum_{j=1}^d (x_j - \mu_j)^2 \sigma_j^2 \right], a_i \geq 0, \sigma_j \geq 0, \mu_j \in \mathbb{R} \right\}.$$

It is known that the class \mathcal{G} is dense in the set of non-negative continuous functions on $[0, 1]^d$ (see [28] and also [20] where the class \mathcal{G} is also considered).

Assumption 4.1.

- $Y = m(\mathbf{X}) + \varepsilon$;
- $\mathbf{X} = (X_1, \dots, X_d)$ is a continuous random vector with independent coordinates;
- ε and \mathbf{X} are independent, ε is a continuous, centered random variable with increasing distribution function, ε has light tails i.e. there exists $0 < \theta < 1$ such that for any $D > 0$, $\mathbb{P}(|\varepsilon| > D) \leq C\theta^D$.
- \mathbf{X} takes its values in \mathcal{X} which is assumed to be a compact hyper-rectangle of \mathbb{R}^d : $\mathcal{X} = \prod_{i=1}^d [u_i, v_i]$, $-\infty < u_i \leq v_i < \infty$;
- for any y , $\mathbf{x} \mapsto F(y|\mathbf{x})$ is continuous and increasing.

Theorem 4.1. *Let Assumption 4.1 be satisfied, with m belonging to the \spadesuit -class, assume that for fixed $\beta > \frac{5}{2}$, $C > 0$, each constructed tree is the highest such that $C\sqrt{n}(\ln n)^\beta \leq N_n^b(\Theta_\ell, \mathcal{D}_n)$, let $k = O(n^\alpha)$, with $\alpha > 0$. Then*

$$\forall \mathbf{x} \in \mathcal{X}, \quad \sup_{y \in \mathbb{R}} |F_{k,n}^b(y|\mathbf{X} = \mathbf{x}) - F(y|\mathbf{X} = \mathbf{x})| \xrightarrow[n \rightarrow \infty]{a.s.} 0,$$

and

$$\forall \mathbf{x} \in \mathcal{X}, \quad \sup_{y \in \mathbb{R}} |F_{k,n}^o(y|\mathbf{X} = \mathbf{x}) - F(y|\mathbf{X} = \mathbf{x})| \xrightarrow[n \rightarrow \infty]{a.s.} 0.$$

Remark 4.3. *Let us mention that the number $N_n^b(\Theta_\ell, \mathcal{D}_n)$ of elements from $\mathcal{D}_n^*(\Theta_\ell)$ in each leaf may be controlled during the construction of the forest. Indeed, in most of the implementations, it is a parameter that may be chosen by the user.*

The algorithm used to get the highest tree such that $C\sqrt{n}(\ln n)^\beta \leq N_n^b(\Theta_\ell, \mathcal{D}_n)$ is as follows: choose (j, z) , with $j \in \mathcal{M}_{try}$ maximizing $L_A^n(j, z)$ while the number of elements from $\mathcal{D}_n^(\Theta_\ell)$ in $A \cap \{[a_j, z]\}$ and $A \cap \{[z, b_j]\}$ is greater than $C\sqrt{n}(\ln n)^\beta$.*

In section 5, we shall see that with this procedure the height of the trees goes to infinity with n a.s. (we do not have and do not need further control on the height).

It will be clear from the discussion of Section 5 that Theorem 4.1 is a direct consequence of Theorems 4.2 and 4.3 that we now state. Remark that Theorems 4.2 and 4.3 could apply in other contexts than Theorem 4.1, possibly with some dependencies, the main point being to satisfy Assumption 4.2. For example, for some adhoc tree constructions, Assumption 4.2 is satisfied without independence assumption, see Remark 4.4. In [20], some hypothesis weaker than independence between the X_i 's are given but the specific examples provided there are also with independent X_i 's. The following assumptions allow to prove the a.s. consistency of our estimators in a general framework.

Assumption 4.2. For all $\ell \in \llbracket 1, k \rrbracket$, we assume that the variation of the conditional cumulative distribution function within any cell goes to 0:

$$\forall \mathbf{x} \in \mathcal{X}, \forall y \in \mathbb{R}, \quad \sup_{\mathbf{z} \in A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)} |F(y|\mathbf{z}) - F(y|\mathbf{x})| \xrightarrow[n \rightarrow \infty]{a.s.} 0.$$

We shall discuss further on Assumption 4.2. In particular in Section 5, we show that this assumption is verified if Assumption 4.1 is satisfied with m in the \spadesuit -class.

Remark 4.4. Let us remark that Assumption 4.2 is satisfied, for example, provided that the diameter of each tree cell goes to zero and for all y , $F(y|\cdot)$ is continuous. This is satisfied, for example if the splitting rules in the construction of the trees imply that each direction $j = 1, \dots, d$ is chosen with positive probability at each split and a positive proportion of the sub-sample goes in each child node (see [25] Assumption 3 and Lemma 2). Imposing that each direction is chosen with a positive probability at each split seems unrealistic since even non informative variables will be chosen. This is why we propose an approach where no additional assumptions on the splitting rules are needed, see Remark 4.3 and Section 5.

Assumption 4.3. We shall make the following assumptions on k (number of trees) and $N_n^b(\mathbf{x}; \Theta, \mathcal{D}_n)$ (number of bootstrap observations in a leaf node):

1. $k = \mathcal{O}(n^\alpha)$, with $\alpha > 0$.
 2. $\forall \mathbf{x} \in \mathcal{X}, \quad N_n^b(\mathbf{x}; \Theta, \mathcal{D}_n) = \Omega\left(\sqrt{n}(\ln(n))^\beta\right)$, with ¹ $\beta > 1$, a.s.
 - or
 3. $\forall \mathbf{x} \in \mathcal{X}, \quad \mathbb{E}[N_n^b(\mathbf{x}; \Theta, \mathcal{D}_n)] = \Omega\left(\sqrt{n}(\ln(n))^\beta\right)$, with $\beta > 1$, and
- $$\forall \mathbf{x} \in \mathcal{X}, \quad \text{CV}(N_n^b(\mathbf{x}; \Theta, \mathcal{D}_n)) = \mathcal{O}\left(\frac{1}{n^{(\alpha+1)/2}(\ln(n))^{\gamma/2}}\right), \text{ with } ^2 \gamma > 1.$$

Remark 4.5. In order to prove our main consistency result, either Assumption 4.3 item 2. or item 3. is needed. Item 2. may seem much stronger than item 3. but it has to be noted that the number of bootstrap observations in a tree leaf is a construction parameter of the forest, so that it can be controlled. Using item 2. simplifies the proof but item 3. is sufficient.

Assumption 4.4. For every $\mathbf{x} \in \mathcal{X}$, the conditional cumulative distribution function $y \mapsto F(y|\mathbf{X} = \mathbf{x})$ is continuous and increasing.

The two theorems below give the uniform a.s. consistency of our two estimators.

Theorem 4.2. Consider a random forest which satisfies Assumptions 4.2 to 4.4. Then,

$$\forall \mathbf{x} \in \mathcal{X}, \quad \sup_{y \in \mathbb{R}} |F_{k,n}^b(y|\mathbf{X} = \mathbf{x}) - F(y|\mathbf{X} = \mathbf{x})| \xrightarrow[n \rightarrow \infty]{a.s.} 0.$$

¹ $f(n) = \Omega(g(n)) \iff \exists k > 0, \exists n_0 > 0 \mid \forall n \geq n_0 \quad |f(n)| \geq k \cdot |g(n)|$

² $\text{CV}(X) = \sigma_X / \mathbb{E}[X]$

Theorem 4.3. *Consider a random forest which satisfies Assumptions 4.2 to 4.4. Then,*

$$\forall \mathbf{x} \in \mathcal{X}, \quad \sup_{y \in \mathbb{R}} |F_{k,n}^o(y | \mathbf{X} = \mathbf{x}) - F(y | \mathbf{X} = \mathbf{x})| \xrightarrow[n \rightarrow \infty]{a.s.} 0 .$$

Remark 4.6. *Using standard arguments, the consistency of quantile estimates stems from Assumption 4.4 as well as the uniform convergence of the conditional distribution function estimators obtained above.*

Let us mention that Assumption 4.3 allows to control the estimation error of our estimators and expresses that cells should contain a sufficiently large number of points so that averaging among the observations is effective.

Finally, Assumption 4.4 is used to get uniform convergence of the estimators.

Section 5 is devoted to comments on Assumption 4.2 and to prove that Assumption 4.1 together with a control on the leave size is sufficient to have Assumption 4.2 for functions of the ♠-class.

5. On variation in Random Forest leaves

This section is devoted to some further analysis on the variation of $F(y|\cdot)$ on leaves, as formalized in Assumption 4.2.

5.1. Some comments on Assumption 4.2

Assumption 4.2 ensures a control on the approximation error of the estimators. It is drawn from Proposition 2 of [35] which shows the consistency of Breiman's random forest estimate in an additive regression framework. Their Proposition 2 allows to manage the approximation error of the estimator by showing that the variation of the regression function m within a cell of a random empirical tree is small provided n is large enough. This result is based on Lemma 1 of [35] which states that the variation of the regression function m within a cell of a random theoretical tree goes to zero for an additive regression model. A random theoretical tree is grown following the same rules as a random empirical tree, except that the theoretical equivalent of the empirical CART-split criterion (2.1) on a node A below is used to choose the best split

$$\begin{aligned} L_A^*(i, z) = & \text{Var}(Y | \mathbf{X} \in A) \\ & - \mathbb{P}(X_i < z | \mathbf{X} \in A) \text{Var}(Y | X_i < z, \mathbf{X} \in A) \\ & - \mathbb{P}(X_i \geq z | \mathbf{X} \in A) \text{Var}(Y | X_i \geq z, \mathbf{X} \in A) . \end{aligned} \tag{5.1}$$

Hence, a theoretical tree is obtained thanks to the best consecutive cuts (i^*, z^*) , among $i \in \mathcal{M}_{\text{try}}$, $z \in A^i$ optimizing the previous criterion $L^*(\cdot, \cdot)$. Remark that \mathcal{M}_{try} is selected with respect to Θ^2 , as for the empirical forest.

General results on standard partitioning estimators whose construction is independent of the label in the training set (see Chapter 4 in [18] or Chapter 6 in [10]) state that a sufficient condition to prove the consistency is that the

diameter of the cells tend to zero as $n \rightarrow \infty$. Instead of such a geometrical assumption, Proposition 2 in [35] ensures that the variation of m inside a node is small thanks to their Lemma 1. But the cornerstone of the Lemma 1 is the Technical Lemma 1 of [34] recalled below for completeness.

Lemma (Technical). *Assume that:*

- $Y = m(\mathbf{X}) + \varepsilon$ with $m(\mathbf{X}) = \sum_{i=1}^d m_i(X_i)$, $\mathbf{X} \sim \mathcal{U}([0, 1]^d)$ and $\varepsilon \sim \mathcal{N}(0, \sigma^2)$,
- $L_A^*(i, z) = 0 \quad \forall i = 1, \dots, d, \forall z \in [a_i, b_i] \quad (0 \leq a_i < b_i \leq 1)$,

then the regression function m is constant on A .

This lemma states that if the theoretical split criterion is zero for all cuts in a node, then the regression function m is constant on this node, i.e. the variation of m on the cell is zero. But, examples for which $L_A^*(i, z) = 0 \quad \forall i, \forall z \in [a_i, b_i]$ and the regression function is not constant can be easily constructed. Let us consider a two-dimensional example, let $A = A_1 \times A_2 = [a_1, b_1] \times [a_2, b_2]$ and suppose that the response Y is

$$Y = X_1 X_2 + c_1 X_1 + c_2 X_2 + \varepsilon =: m(X_1, X_2) + \varepsilon, \quad (5.2)$$

with

- $\mathbf{X} = (X_1, X_2)$ independent random inputs,
- $c_1 = -\frac{\mathbb{E}[X_2 \mathbb{1}_{\{X_2 \in A_2\}}]}{\mathbb{P}(X_2 \in A_2)}$ and $c_2 = -\frac{\mathbb{E}[X_1 \mathbb{1}_{\{X_1 \in A_1\}}]}{\mathbb{P}(X_1 \in A_1)}$,
- and ε a centered noise independent of \mathbf{X} .

It can be shown for this model that within the node A , $L^* \equiv 0$ for all $i \in \{1, 2\}$, for all $z \in [a_i, b_i]$ and yet the regression function m is not constant.

Accordingly, the technical lemma above is well-designed for an additive regression framework. But this context is far from reality for many concrete examples.

5.2. Sufficient conditions to control the variation on leaves

Our purpose is now to prove that Assumption 4.1 with a control on the size of the leaves and a regression function in the \spadesuit -class is sufficient to ensure Assumption 4.2 which controls the variation of $F(y|\cdot)$ on leaves. An interesting sub-product of this study is Proposition 5.4 which shows the asymptotic proximity of empirical and theoretical trees.

In what follows, C denotes any positive constant so that we may write $C + C = C$, $uC = C$, with $u > 0$, ...

We are interested in $F(y|\cdot)$, remark that $F(y|\cdot)$ is constant on A if and only if $m(\cdot)$ is constant on A . Indeed,

$$F(y|\mathbf{x}) = \mathbb{P}(Y \leq y | \mathbf{X} = \mathbf{x}) = \mathbb{P}(\varepsilon \leq y - m(\mathbf{x})) = F_\varepsilon(y - m(\mathbf{x})) .$$

We shall use theoretical trees construction as in [35], the theoretical CART-split is done by maximizing $L_A^*(i, z)$ defined in (5.1). If L_A^* admits several maxima, then one is chosen uniformly at random.

Let us denote by $A_n(\Theta_\ell, \mathcal{D}_n)$ any leaf in an empirical tree, and $N_n^b(\Theta_\ell, \mathcal{D}_n)$ be the number of elements of $\mathcal{D}_n^*(\Theta_\ell)$ that fall into $A_n(\Theta_\ell, \mathcal{D}_n)$. We have the following result.

Theorem 5.1. *Let Y satisfy Assumption 4.1, with m belonging to the \spadesuit -class, let $\beta > \frac{5}{2}$, $C > 0$, let the constructed trees be the highest such that $C\sqrt{n}(\ln n)^\beta \leq N_n^b(\Theta_\ell, \mathcal{D}_n)$, then Assumption 4.2 is verified.*

Note that for additive m it is proven in [35] that the variation of m on leaves goes to zero in probability, with an assumption on the height of trees. We extend this result to the \spadesuit -class and obtain an a.s. convergence. The control on the height of trees is replaced by a control on the size of leaves. The strategy of proof is partially inspired by [34] and [35]. The following lemma proves that the variation of $F(y|\cdot)$ on leaves of theoretical trees goes to 0 a.s. as the height of the trees goes to infinity.

Lemma 5.2. *Assume that Assumption 4.1 is satisfied with the regression function m in the \spadesuit -class, let $S^\infty = (s_j, j = 1, \dots)$ with $s_j \in \{L, R\}$, it describes an infinite path in a binary tree, let $S^h = (s_j, j = 1, \dots, h)$, it describes a path in a binary tree of height h . Let $A_h(S^h, \Theta)$ be the corresponding leaf in a theoretical tree. Then the variation of $F(y|\cdot)$ on $A_h(S^h, \Theta)$ goes to 0 a.s. as h goes to infinity.*

Proof. Since $\mathbf{x} \mapsto F(y|\mathbf{x})$ is assumed to be continuous, then the result holds if the diameter of $A_h(S^h, \Theta)$ goes to zero. Let $A_\infty(S^\infty, \Theta) = \bigcap_{h \geq 1} A_h(S^h, \Theta)$,

it is a decreasing intersection of rectangles, if its diameter is non zero then it writes: for a non empty subset $J \subset \{1, \dots, d\}$, let J^c be its complementary, $(x_j, j \in J^c, a_j, j \in J); (x_j, j \in J^c, b_j, j \in J) \in \mathcal{X}$

$$A_\infty(S^\infty, \Theta) = \{x_j, j \in J^c\} \times \prod_{j \in J} [a_j, b_j] = \{x_j, j \in J^c\} \times A^J, \text{ with } a_j < b_j.$$

Remark that for any rectangle $A = \prod_{k=1}^d [\alpha_k, \beta_k]$,

$$\begin{aligned} L_A^*(i, z) &= \mathbb{P}(X_i < z | \mathbf{X} \in A) (\mathbb{E}[Y | X_i < z, \mathbf{X} \in A] - \mathbb{E}[Y | \mathbf{X} \in A])^2 \\ &\quad + \mathbb{P}(X_i \geq z | \mathbf{X} \in A) (\mathbb{E}[Y | X_i \geq z, \mathbf{X} \in A] - \mathbb{E}[Y | \mathbf{X} \in A])^2 \end{aligned}$$

can be seen as a continuous function of $(\alpha_k, \beta_k)_{k=1, \dots, d}, i, z$.

For $j \in J$, $L_{A_\infty(S^\infty, \Theta)}^*(j, z)$ rewrites

$$\begin{aligned} &L_{A_\infty(S^\infty, \Theta)}^*(j, z) \\ &= \mathbb{P}(X_j < z | \mathbf{X}^J \in A^J) (\mathbb{E}[Y | X_j < z, \mathbf{X}^J \in A^J] - \mathbb{E}[Y | \mathbf{X}^J \in A^J])^2 \\ &\quad + \mathbb{P}(X_j \geq z | \mathbf{X}^J \in A^J) (\mathbb{E}[Y | X_j \geq z, \mathbf{X}^J \in A^J] - \mathbb{E}[Y | \mathbf{X}^J \in A^J])^2. \end{aligned}$$

Also, for $j \in J^c$, $L_{A_\infty(S^\infty)}^*(j, z) = 0$.

Write $A_h(S^h, \Theta) = \prod_{j=1}^d [a_j^h, b_j^h]$, then for $j \in J$, $a_j^h \rightarrow a_j$ and $b_j^h \rightarrow b_j$ as $h \rightarrow \infty$, also $L_{A_h(S^h)}^*(j, z) \rightarrow L_{A_\infty(S^\infty)}^*(j, z)$.

Let $(j^*, z^*) \in \operatorname{argmax} L_{A_\infty(S^\infty)}^*(j, z)$, if $j^* \in J^c$ then $L_{A_\infty(S^\infty)}^*(j^*, z^*) = 0$ and thus $L_{A_\infty(S^\infty)}^*(j, z) = 0$ for any j, z . Assume $j^* \in J$ and consider a subsequence $(h_p)_{p \in \mathbb{N}}$ such that $j^* \in \mathcal{M}_{\text{try}}$ at each level h_p , it exists a.s. because Θ gives a positive probability to each $j \in \{1, \dots, d\}$ of belonging to \mathcal{M}_{try} at each level. Then, consider $(j^p, z^p) \in \operatorname{argmax} L_{A_{h_p}(S^{h_p})}^*(j, z)$ and (j^∞, z^∞) any limit point of the sequence (j^p, z^p) , which means that for a subsequence $p_q \rightarrow \infty$, $(j^{p_q}, z^{p_q}) \rightarrow (j^\infty, z^\infty)$, thus for q large enough, $j^{p_q} = j^\infty$. It easily follows that $L_{A_\infty(S^\infty)}^*(j^*, z^*) \leq L_{A_\infty(S^\infty)}^*(j^\infty, z^\infty)$ and thus $(j^\infty, z^\infty) \in \operatorname{argmax} L_{A_\infty(S^\infty)}^*(j, z)$. Now, if $j^\infty \in J$, each z^h is either $a_{j^h}^h$ or $b_{j^h}^h$; z^∞ is either a_{j^∞} or b_{j^∞} . So that either $\{X_{j^\infty} < z, \mathbf{X}^J \in A^J\} = \{X^J \in A^J\}$ and $\{X_{j^\infty} \geq z, \mathbf{X}^J \in A^J\} = \emptyset$ or $\{X_{j^\infty} < z, \mathbf{X}^J \in A^J\} = \emptyset$ and $\{X_{j^\infty} \geq z, \mathbf{X}^J \in A^J\} = \{X^J \in A^J\}$, thus

$$\begin{aligned} & L_{A_\infty(S^\infty)}^*(j^\infty, z^\infty) \\ &= \mathbb{P}(X_{j^\infty} < z | \mathbf{X}^J \in A^J) (\mathbb{E}[Y | X_{j^\infty} < z, \mathbf{X}^J \in A^J] - \mathbb{E}[Y | \mathbf{X}^J \in A^J])^2 \\ & \quad + \mathbb{P}(X_{j^\infty} \geq z | \mathbf{X}^J \in A^J) (\mathbb{E}[Y | X_{j^\infty} \geq z, \mathbf{X}^J \in A^J] - \mathbb{E}[Y | \mathbf{X}^J \in A^J])^2 \\ &= 0. \end{aligned}$$

If $j^\infty \in J^c$, then $L_{A_\infty(S^\infty)}^*(j^\infty, z^\infty) = 0$. Finally, since $L_{A_\infty(S^\infty)}^*(j, z) \leq L_{A_\infty(S^\infty)}^*(j^\infty, z^\infty) = 0$, we conclude that $L_{A_\infty(S^\infty)}^*(j, z) = 0$ for all (j, z) . This is equivalent to: for all $i = 1, \dots, d$, $z \in [a_i, b_i]$,

$$\begin{aligned} & \mathbb{E}[Y | X_i \leq z, \mathbf{X} \in A_\infty] - \mathbb{E}[Y | \mathbf{X} \in A_\infty] = 0 \Leftrightarrow \\ & \mathbb{P}(\mathbf{X} \in A_\infty) \mathbb{E}[Y \mathbb{1}_{\{X_i \leq z, \mathbf{x} \in A_\infty\}}] = \mathbb{P}(X_i \leq z, \mathbf{X} \in A_\infty) \mathbb{E}[Y \mathbb{1}_{\{\mathbf{x} \in A_\infty\}}]. \end{aligned}$$

By derivating with respect to z , we may see that it is equivalent to $z \mapsto \mathbb{E}\left[m(z, \mathbf{X}^{-i}) \mathbb{1}_{\{\mathbf{X}^{-i} \in A_\infty^{-i}\}}\right]$ is constant for all $i = 1, \dots, d$. Since we assumed that m belongs to the \spadesuit -class, either m is constant on $A_\infty(S^\infty)$ or the diameter of $A_\infty(S^\infty)$ is zero. In both cases, the variation of $F(y|\cdot)$ on $A_h(S^h, \Theta)$ goes to 0 as h goes to infinity. \square

We now show that empirical and theoretical trees are close. The main tool for this step is the following result which has an intrinsic interest since it proves that the theoretical and empirical cost functions used for the tree constructions are uniformly close, provided that the leaf sizes are not too small. Recall that $L_A^n(j, z)$ is the empirical cost on the rectangle A on direction j and at level z .

Proposition 5.3. *Let Assumption 4.1 be satisfied. Let $\beta > \frac{5}{2}$, let A be a rectangle in \mathcal{X} , we shall say that $(A, j, z) \in \mathcal{A}^n$ if the numbers $N_{A_L}^b, N_{A_R}^b$ of elements of \mathcal{D}_n^* belonging to $A_L := A \cap \{x_j \leq z\}$ and $A_R := A \cap \{x_j > z\}$ are greater*

than $C\sqrt{n}(\ln n)^\beta$. We have

$$\sup_{(A,j,z) \in \mathcal{A}^n} |L_A^*(j,z) - L_A^n(j,z)| \xrightarrow{a.s.} 0 \text{ as } n \rightarrow \infty.$$

Proof. Let $(A, j, z) \in \mathcal{A}^n$ be fixed. Rewrite the difference $|L_A^*(j, z) - L_A^n(j, z)| = |T_L + T_R|$ with

$$\begin{aligned} T_L &= T_{L,1} + T_{L,2} \\ &=: \frac{N_{A_L}^b}{N_A^b} \left((\overline{Y_A} - \overline{Y_{A_L}})^2 - (\mathbb{E}[Y|\mathbf{X} \in A] - \mathbb{E}[Y|\mathbf{X} \in A_L])^2 \right) \\ &\quad + (\mathbb{E}[Y|\mathbf{X} \in A] - \mathbb{E}[Y|\mathbf{X} \in A_L])^2 \left(\frac{N_{A_L}^b}{N_A^b} - \mathbb{P}(\mathbf{X} \in A_L|\mathbf{X} \in A) \right). \end{aligned}$$

The term T_R is defined in the same way by using A_R instead of A_L . In order to prove the proposition, we shall prove that $\sup_{A,j,z} T_{L,1}$ and $\sup_{A,j,z} T_{L,2}$ go to 0 a.s. The same holds in the same way for $T_{R,1}$ and $T_{R,2}$. Using Vapnik-Chervonenkis theory on rectangles in \mathcal{A} (see Lemma A.1 in Appendix A) we have:

$$\mathbb{P} \left(\sup_{B \in \mathcal{A}} \left| \frac{N_B^b}{n} - \mathbb{P}(\mathbf{X} \in B) \right| > \kappa \right) \leq 16(n+1)^{2d} e^{-n\kappa^2/128}. \quad (5.3)$$

$T_{L,2}$ decomposes into:

$$\begin{aligned} |T_{L,2}| &\leq (\mathbb{E}(Y|\mathbf{X} \in A) - \mathbb{E}(Y|\mathbf{X} \in A_L))^2 \\ &\quad \times \left[\frac{n}{N_A^b} \left| \frac{N_{A_L}^b}{n} - \mathbb{P}(\mathbf{X} \in A_L) \right| + \left| \frac{1}{\mathbb{P}(\mathbf{X} \in A)} - \frac{n}{N_A^b} \right| \mathbb{P}(\mathbf{X} \in A_L) \right]. \end{aligned}$$

Remark that for $B \in \mathcal{A}$, if

$$\left| \frac{N_B^b}{n} - \mathbb{P}(\mathbf{X} \in B) \right| \leq \frac{C}{2} \frac{(\ln n)^\beta}{\sqrt{n}} \quad (5.4)$$

and $N_B^b \geq C\sqrt{n}(\ln n)^\beta$, then $\mathbb{P}(\mathbf{X} \in B) \geq \frac{C}{2} \frac{(\ln n)^\beta}{\sqrt{n}}$. So that, for $(A, j, z) \in \mathcal{A}^n$, we have, provided that (5.4) holds for A ,

$$\begin{aligned} \mathbb{E}(Y|\mathbf{X} \in A) &= \frac{1}{\mathbb{P}(\mathbf{X} \in A)} \left[\mathbb{E}(Y \mathbf{1}_{\{\mathbf{X} \in A\}} \mathbf{1}_{\{Y \leq D\}}) + \mathbb{E}(Y \mathbf{1}_{\{\mathbf{X} \in A\}} \mathbf{1}_{\{Y > D\}}) \right] \\ &\leq D + \mathbb{E}(Y^p)^{\frac{1}{p}} \frac{\mathbb{P}(Y > D)^{\frac{1}{q}}}{\mathbb{P}(X \in A)^{1-\frac{1}{r}}} \end{aligned}$$

the second term is obtained using Hölder inequality

$$\begin{aligned} &\leq (\ln n)^\gamma + C e^{(\ln n)^\gamma \frac{\ln \theta}{q}} \frac{n^{\frac{1}{2}(1-\frac{1}{r})}}{(\ln n)^{\beta(1-\frac{1}{r})}} \\ &\leq C(\ln n)^\gamma, \end{aligned}$$

by taking $D = (\ln n)^\gamma$, $\gamma > 1$, $p, q, r > 0$ with $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} = 1$. Now, Equation (5.3) gives that (5.4) is satisfied for A and A_L with probability greater than $1 - 16(n+1)^{2d} e^{-C \frac{(\ln n)^{2\beta}}{512}}$ and the condition on N_{A_L} for $(A, j, z) \in \mathcal{A}^n$ gives that

$$\mathbb{P}(T_{L,2} > \kappa) \leq 16(n+1)^{2d} e^{-C \frac{(\ln n)^{2\beta}}{512}} + C(n+1)^{2d} e^{-C \kappa^2 \frac{(\ln n)^{2\beta} (\ln n)^{-4\gamma}}{128}}.$$

Then, Borel-Cantelli Lemma gives that $\sup_{(A,j,z) \in \mathcal{A}^n} T_{L,2}$ goes to 0 a.s. provided that $2\beta - 4\gamma > 1$.

The term $T_{L,1}$ is treated in the same way, instead of using (5.3), we use Lemma A.2 in Appendix A to get that for any $D > 0$, for any $\kappa > 0$, $A \in \mathcal{A}$ and $\frac{1}{p} + \frac{1}{q} = 1$

$$\begin{aligned} & \mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n Y_i^* \mathbb{1}_{\{\mathbf{x}_i^* \in A\}} - \mathbb{E}(Y \mathbb{1}_{\{\mathbf{x} \in A\}}) \right| > \kappa \right) \\ & \leq 6 \left(\frac{24eD}{\kappa} \ln \left(\frac{48eD}{\kappa} \right) \right)^{2d} e^{-n\kappa^2/(128D^2)} + C \frac{\mathbb{E}(Y^p)^{\frac{1}{p}} \mathbb{P}(Y > D)^{\frac{1}{q}}}{\kappa}, \end{aligned} \quad (5.5)$$

where the (\mathbf{X}_i^*, Y_i^*) 's form a bootstrap sample from $(\mathbf{X}_1, Y_1), \dots, (\mathbf{X}_n, Y_n)$ and we take $D = (\ln n)^\gamma$ as before. \square

The last stone for the proof of Theorem 5.1 is to prove that at each level h , each node of the empirical tree $A^n(S)$ is close to a level h node of a theoretical tree.

Proposition 5.4. *Let Assumption 4.1 be satisfied. Assume that for $\beta > \frac{5}{2}$, $N_n^b(\Theta_\ell, \mathcal{D}_n) \geq C\sqrt{n}(\ln n)^\beta$. For $h \in \mathbb{N}$, let $S \in \{L, R\}^h$ describe a path of length h in a binary tree, let $A^n(S)$ and $A(S)$ be corresponding nodes in empirical and theoretical trees. Denote*

$$A(S) = \prod_{j=1}^d [a_j, b_j] \text{ and } A^n(S) = \prod_{j=1}^d [a_j^n, b_j^n].$$

Denote \mathcal{T}_h the set of theoretical trees of height h , then

$$\inf_{\mathcal{T}_h} \max_{j=1, \dots, d} \max (|a_j - a_j^n|, |b_j - b_j^n|) \longrightarrow 0 \text{ a.s. as } n \rightarrow \infty \quad (5.6)$$

Remark 5.1. *Proposition 5.4 implies that for any $h \in \mathbb{N}$ fixed, for any empirical tree with node sizes greater than $C\sqrt{n}(\ln n)^\beta$, we may find a theoretical one as close as we wish, until height h . This is an a.s. and more precise version of Lemma 3 in [35].*

Proof. We proceed by induction on h . If $h = 0$ then $A(S) = A^n(S) = \mathcal{X}$ and the assertion holds. Let us assume that the result holds for $h \in \mathbb{N}$, let $S' \in \{L, R\}^{h+1}$

describe a path of length $h + 1$ in a binary tree, $S' = (S, u)$ with $S \in \{L, R\}^h$ and $u \in \{L, R\}$. Let $(j^n, z^n) \in \operatorname{argmax} L_{A^n(S)}^n(j, z)$, then

$$A^n(S') = A^n(S) \cap \{x_{j^n} \leq z^n\} \text{ or } A^n(S') = A^n(S) \cap \{x_{j^n} > z^n\}.$$

Let us denote $A^n(S) = \prod_{j=1}^d [a_j^n, b_j^n]$. Our hypothesis on the construction of the empirical tree implies that for all $n \in \mathbb{N}$, $(A^n(S), j^n, z^n) \in \mathcal{A}^n$ so that by Proposition 5.3:

$$|L_{A^n(S)}^n(j^n, z^n) - L_{A^n(S)}^*(j^n, z^n)| \longrightarrow 0 \text{ a.s. as } n \rightarrow \infty. \quad (5.7)$$

The induction hypothesis implies that, taking if necessary a subsequence, we may find $A(S) = \prod_{j=1}^d [a_j, b_j]$ a node in a theoretical tree at height h such that $|a_j^n - a_j| \rightarrow 0$ and $|b_j^n - b_j| \rightarrow 0$ a.s. Let (j^∞, z^∞) be any limit point of (j^n, z^n) , the continuity of L^* as a function of (a_j, b_j, j, z) and (5.7), lead to $|L_{A^n(S)}^n(j^n, z^n) - L_{A(S)}^*(j^\infty, z^\infty)| \rightarrow 0$, a.s. as n goes to infinity, taking if needed a subsequence.

Let us consider any (j, z) , if $(A^n(S), j, z) \notin \mathcal{A}^n$ for n large enough, then $L_{A(S)}^*(j, z) = 0$ because, in that case, either:

$\mathbb{P}(X_j < z, \mathbf{X} \in A(S)) = 0$ and $\mathbb{P}(X_j \geq z, \mathbf{X} \in A(S)) = \mathbb{P}(\mathbf{X} \in A(S))$ or
 $\mathbb{P}(X_j \geq z, \mathbf{X} \in A(S)) = 0$ and $\mathbb{P}(X_j < z, \mathbf{X} \in A(S)) = \mathbb{P}(\mathbf{X} \in A(S))$ (using (5.3)).

Otherwise, $(A^n, j, z) \in \mathcal{A}^n$ for infinitely many n and using again Proposition 5.3, we have $L_{A(S)}^*(j^\infty, z^\infty) \geq L_{A(S)}^*(j, z)$. We conclude that $(j^\infty, z^\infty) \in \operatorname{argmax} L_{A(S)}^*(j, z)$ and thus

$$A(S) \cap \{x^{j^\infty} \leq z^\infty\} \text{ and } A(S) \cap \{x^{j^\infty} > z^\infty\}$$

are level $h + 1$ nodes of a theoretical tree. In other words, z^n and z^∞ are respectively new end points of $A^n(S')$, $A(S')$ and since z^∞ is a limit point of $(z^n)_{n \in \mathbb{N}}$, this concludes the induction step and the proof. \square

End of proof of Theorem 5.1. In order to conclude the proof of Theorem 5.1, we need to prove that the variation of $F(y|\cdot)$ on any leaf in an empirical tree goes to 0 a.s.

Consider $h \in \mathbb{N}$, a leaf is described by a finite sequence $S \in \{L, R\}^h$ of some $S^\infty \in \{L, R\}^{\mathbb{N}}$. If $A^n(S)$ is a node of the empirical tree, it means that the number of elements from \mathcal{D}_n^* in $A^n(S)$ is greater than $C\sqrt{n}(\ln n)^\beta$.

Let $(j^n, z^n) \in \operatorname{argmax} L_{A^n(S)}^n(j, z)$. On one side, if $(A^n(S), j^n, z^n) \notin \mathcal{A}^n$ for infinitely many n , let $A(S)$ be a node in a theoretical tree, given by Proposition 5.4. As in the proof of Proposition 5.4, we have that $L_{A(S)}^*(j, z) = 0$ for all (j, z) , since m belongs to the \spadesuit -class, this implies that the variation of $F(y|\cdot)$ on $A(S)$ is zero. Using again Proposition 5.4, we conclude that for any $\kappa > 0$, for n large

enough, the variation of $F(y|\cdot)$ on $A^n(S)$ is less than κ . In this case, $A^n(S)$ is a leaf of the empirical tree (since the condition on the number of elements in the child nodes is not satisfied).

Otherwise, $(A^n(S), j^n, z^n) \in \mathcal{A}^n$ for n large enough, which means that the construction of the empirical tree continues.

By induction on h , we get that if $A^n(S)$ is a leaf in an empirical tree, with $|S| = h_n$ then either the variation of $F(y|\cdot)$ on $A^n(S)$ goes to 0 or h_n goes to infinity. Consider the second case. Fix $\kappa > 0$ and using Lemma 5.2, choose h so that for $|S| = h$ the variation of $F(y|\cdot)$ on a node $A(S)$ in an empirical tree is less than κ . Choose η so that if $|z - z'| < \eta$ then $|F(y|z) - F(y|z')| < \kappa$. Using Proposition 5.4, choose n such that for any $z \in A^n(S)$ there is $z' \in A(S)$ with $|z - z'| < \eta$. Then, the variation of $F(y|\cdot)$ on $A_n(S)$ is less than 3κ . \square

To conclude this section, let us say that Theorem 5.1 gives simple and realistic conditions under which the hypothesis of Theorems 4.2 and 4.3 are verified.

The next section is devoted to the proof of Theorem 4.2. The proof of Theorem 4.3 is similar and left to the reader.

6. Proofs of Theorem 4.2

The proofs of Theorems 4.2 and 4.3 are close. We only provide that of Theorem 4.2 below.

The main ingredient of the proof is to use a second sample \mathcal{D}_n^\diamond in order to deal with the data-dependent aspect. Thus, we first define a dummy estimator based on two samples \mathcal{D}_n and \mathcal{D}_n^\diamond which will be used below. The trees are grown using \mathcal{D}_n , but we consider another sample \mathcal{D}_n^\diamond (independent of \mathcal{D}_n and Θ) which is used to define a dummy estimator

$$\begin{aligned} F_{k,n}^\diamond(y|\mathbf{X} = \mathbf{x}; \Theta_1, \dots, \Theta_k, \mathcal{D}_n^\diamond, \mathcal{D}_n) \\ = \sum_{j=1}^n w_{n,j}^\diamond(\mathbf{x}; \Theta_1, \dots, \Theta_k, \mathbf{X}^{\diamond 1}, \dots, \mathbf{X}^{\diamond n}, \mathcal{D}_n) \mathbb{1}_{\{Y^{\diamond j} \leq y\}}, \end{aligned} \quad (6.1)$$

where the weights are for $j = 1, \dots, n$,

$$w_{n,j}^\diamond(\mathbf{x}; \Theta_1, \dots, \Theta_k, \mathbf{X}^{\diamond 1}, \dots, \mathbf{X}^{\diamond n}, \mathcal{D}_n) = \frac{1}{k} \sum_{\ell=1}^k \frac{\mathbb{1}_{\{\mathbf{X}^{\diamond j} \in A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)\}}}{N_n^\diamond(\mathbf{x}; \Theta_\ell, \mathbf{X}^{\diamond 1}, \dots, \mathbf{X}^{\diamond n}, \mathcal{D}_n)},$$

with $N_n^\diamond(\mathbf{x}; \Theta_\ell, \mathbf{X}^{\diamond 1}, \dots, \mathbf{X}^{\diamond n}, \mathcal{D}_n)$, the number of elements of \mathcal{D}_n^\diamond that fall into $A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)$. Throughout this section, we shall use the convention $\frac{0}{0} = 0$ in case $N_n^\diamond(\mathbf{x}; \Theta_\ell, \mathbf{X}^{\diamond 1}, \dots, \mathbf{X}^{\diamond n}, \mathcal{D}_n) = 0$ and thus $\mathbb{1}_{\{\mathbf{X}^{\diamond j} \in A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)\}} = 0$ for $j = 1, \dots, n$.

The weights $w_{n,j}^\diamond(\mathbf{x}; \Theta_1, \dots, \Theta_k, \mathbf{X}^{\diamond 1}, \dots, \mathbf{X}^{\diamond n}, \mathcal{D}_n)$ are nonnegative random variables, as function of $\Theta_1, \dots, \Theta_k, \mathbf{X}^{\diamond 1}, \dots, \mathbf{X}^{\diamond n}, \mathcal{D}_n$. To lighten the notation in the sequel, we will simply write $F_{k,n}^\diamond(y|\mathbf{X} = \mathbf{x}) = \sum_{j=1}^n w_j^\diamond(\mathbf{x}) \mathbb{1}_{\{Y^{\diamond j} \leq y\}}$ instead of (6.1).

Let $\mathbf{x} \in \mathcal{X}$ and $y \in \mathbb{R}$, we have

$$\begin{aligned} |F_{k,n}^b(y|\mathbf{X}=\mathbf{x}) - F(y|\mathbf{X}=\mathbf{x})| &\leq |F_{k,n}^\circ(y|\mathbf{X}=\mathbf{x}) - F(y|\mathbf{X}=\mathbf{x})| \\ &\quad + |F_{k,n}^\circ(y|\mathbf{X}=\mathbf{x}) - F_{k,n}^b(y|\mathbf{X}=\mathbf{x})| . \end{aligned}$$

The convergence of the two right-hand terms is handled separately into the following Proposition 6.1 and Lemma 6.2.

Proposition 6.1. *Consider a random forest which satisfies Assumptions 4.2 and 4.3. Then,*

$$\forall \mathbf{x} \in \mathcal{X}, \forall y \in \mathbb{R}, \quad F_{k,n}^\circ(y|\mathbf{X}=\mathbf{x}) \xrightarrow[n \rightarrow \infty]{a.s.} F(y|\mathbf{X}=\mathbf{x}) .$$

Hence, Proposition 6.1 establishes the consistency for a random forest estimator based on a second sample \mathcal{D}_n° independent of \mathcal{D}_n and Θ . [37] proved that estimators built from honest forests are asymptotically Gaussian. Remark that in [37], it is also required to control the proportion of chosen observations at each split and in each direction. In our case, going through a kind of honest trees is just a theoretical tool. We go one step further with the following lemma by showing that the estimators built with \mathcal{D}_n^* and \mathcal{D}_n° are close.

Lemma 6.2. *Consider a random forest which satisfies Assumption 4.3. Then,*

$$\forall \mathbf{x} \in \mathcal{X}, \forall y \in \mathbb{R}, \quad |F_{k,n}^\circ(y|\mathbf{X}=\mathbf{x}) - F_{k,n}^b(y|\mathbf{X}=\mathbf{x})| \xrightarrow[n \rightarrow \infty]{a.s.} 0 .$$

Hence, according to Proposition 6.1 and Lemma 6.2, we get

$$\forall \mathbf{x} \in \mathcal{X}, \forall y \in \mathbb{R}, \quad F_{k,n}^b(y|\mathbf{X}=\mathbf{x}) \xrightarrow[n \rightarrow \infty]{a.s.} F(y|\mathbf{X}=\mathbf{x}) . \quad (6.2)$$

Now, thanks to Dini's second theorem, let us sketch how to obtain the almost sure uniform convergence relative to y of the estimator.

Note that $\{Y^j \leq y\} = \{U_j \leq F_{Y|\mathbf{X}=\mathbf{x}}(y)\}$ under Assumption 4.4 with $U_j = F_{Y|\mathbf{X}=\mathbf{x}}(Y^j)$, $j = 1, \dots, n$ which are i.i.d random variables. Then, (6.2) is equivalent to

$$\forall \mathbf{x} \in \mathcal{X}, \forall s \in [0, 1], \quad \sum_{j=1}^n w_j^b(\mathbf{x}) \mathbb{1}_{\{U_j \leq s\}} \xrightarrow[n \rightarrow \infty]{a.s.} s .$$

As in the proof of Glivenko–Cantelli's Theorem, using

that $s \mapsto \sum_{j=1}^n w_j^b(\mathbf{x}) \mathbb{1}_{\{U_j(\omega) \leq s\}}$ is increasing and Dini's second theorem, we get the uniform convergence almost everywhere, which concludes the proof of the theorem. \square

We now turn to the proof of Proposition 6.1 while that of Lemma 6.2 is postponed to the Appendix. To that aim, the following lemma, based on Vapnik-Chervonenkis classes [36] is a key tool. The proof is postponed to the Appendix.

Lemma 6.3. *Consider \mathcal{D}_n and \mathcal{D}_n° , two independent datasets of independent n samples of (\mathbf{X}, Y) . Build a tree using \mathcal{D}_n with bootstrap and bagging procedure*

driven by Θ . As before, $N^b(A_n(\Theta)) = N_n^b(\mathbf{x}; \Theta, \mathcal{D}_n)$ is the number of bootstrap observations of \mathcal{D}_n that fall into $A_n(\Theta) = A_n(\mathbf{x}; \Theta, \mathcal{D}_n)$ and $N^\diamond(A_n(\Theta)) = N_n^\diamond(\mathbf{x}; \Theta, \mathbf{X}^{\diamond 1}, \dots, \mathbf{X}^{\diamond n}, \mathcal{D}_n)$, the number of observations of \mathcal{D}_n^\diamond that fall into $A_n(\Theta)$. Then,

$$\forall \varepsilon > 0, \quad \mathbb{P}(|N^b(A_n(\Theta)) - N^\diamond(A_n(\Theta))| > \varepsilon) \leq 24(n+1)^{2d} e^{-\varepsilon^2/288n}.$$

Lemma 6.3 is the main ingredient of the proof of Proposition 6.1.

Proof of Proposition 6.1.

We aim to prove

$$\forall \mathbf{x} \in \mathcal{X}, \forall y \in \mathbb{R}, \quad \mathbb{P}\left(F_{k,n}^\diamond(y|\mathbf{X}=\mathbf{x}) \xrightarrow{n \rightarrow \infty} F(y|\mathbf{X}=\mathbf{x})\right) = 1.$$

Let $\mathbf{x} \in \mathcal{X}$ and $y \in \mathbb{R}$, we have

$$\begin{aligned} |F_{k,n}^\diamond(y|\mathbf{x}) - F(y|\mathbf{x})| &\leq \left| \sum_{j=1}^n w_j^\diamond(\mathbf{x}) (\mathbb{1}_{\{Y^{\diamond j} \leq y\}} - F(y|\mathbf{X}^{\diamond j})) \right| \\ &\quad + \left| \sum_{j=1}^n w_j^\diamond(\mathbf{x}) (F(y|\mathbf{X}^{\diamond j}) - F(y|\mathbf{x})) \right|. \end{aligned}$$

Define

$$W_n = \sum_{j=1}^n w_j^\diamond(\mathbf{x}) (\mathbb{1}_{\{Y^{\diamond j} \leq y\}} - F(y|\mathbf{X}^{\diamond j})) = \sum_{j=1}^n w_j^\diamond(\mathbf{x}) Z_j^\diamond$$

with $Z_j^\diamond = \mathbb{1}_{\{Y^{\diamond j} \leq y\}} - F(y|\mathbf{X}^{\diamond j})$, n i.i.d random variables and

$$V_n = \sum_{j=1}^n w_j^\diamond(\mathbf{x}) (F(y|\mathbf{X}^{\diamond j}) - F(y|\mathbf{x})).$$

Remark that $\mathbb{E}[Z_j^\diamond | \mathbf{X}^{\diamond j}] = 0$.

We first show that $(W_n)_{n \geq 1}$ goes to 0 a.s. in the case of Assumption 4.3 item 2. This is achieved by adapting Hoeffding inequality's proof to our random weighted sum context. For any $\varepsilon > 0$, $t \in \mathbb{R}_+$, we have

$$\mathbb{P}(W_n > \varepsilon) \leq \mathbb{E}[e^{tW_n}] \cdot e^{-t\varepsilon}.$$

We shall make use of the folklore lemma below.

Lemma. *Let X be a centred random variable, a.s. bounded by 1. Then, for any $t \in \mathbb{R}$, $\mathbb{E}[e^{tX}] \leq e^{\frac{t^2}{2}}$.*

Let $t > 0$, we have

$$\begin{aligned}\mathbb{E} [e^{tW_n}] &= \mathbb{E} \left[\prod_{j=1}^n e^{tw_j^\diamond(\mathbf{x})Z_j^\diamond} \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\prod_{j=1}^n e^{tw_j^\diamond(\mathbf{x})Z_j^\diamond} \middle| \mathcal{D}_n, \Theta_1, \dots, \Theta_k, \mathbf{X}^{\diamond 1}, \dots, \mathbf{X}^{\diamond n} \right] \right]\end{aligned}$$

Conditionally to $\mathcal{D}_n, \Theta_1, \dots, \Theta_k, \mathbf{X}^{\diamond 1}, \dots, \mathbf{X}^{\diamond n}$, the w_j^\diamond are constant and the Z_j^\diamond are centred, independent and bounded by 1. Thus, using the folklore lemma,

$$\begin{aligned}\mathbb{E} [e^{tW_n}] &= \mathbb{E} \left[\prod_{j=1}^n \mathbb{E} \left[e^{tw_j^\diamond(\mathbf{x})Z_j^\diamond} \middle| \mathcal{D}_n, \Theta_1, \dots, \Theta_k, \mathbf{X}^{\diamond 1}, \dots, \mathbf{X}^{\diamond n} \right] \right] \\ &\leq \mathbb{E} \left[\prod_{j=1}^n e^{t^2 w_j^\diamond(\mathbf{x})^2 / 2} \right].\end{aligned}$$

Let $K > 0$ be such that for all $\ell = 1, \dots, k$, $N_n^b(A_n(\ell)) = N_n^b(\mathbf{x}; \Theta_\ell, \mathcal{D}_n) \geq K\sqrt{n}(\ln(n))^\beta$ a.s. by using Assumption 4.3 item 2. Denote $\Gamma(\ell)$ the event $\left\{ N_n^\diamond(A_n(\ell)) < \frac{K\sqrt{n}(\ln(n))^\beta}{2} \right\}$. Remark that

$$\Gamma(\ell) \subset \left\{ \left| N_n^\diamond(A_n(\ell)) - N_n^b(A_n(\ell)) \right| > \frac{K\sqrt{n}(\ln(n))^\beta}{2} \right\}.$$

Thus, using Lemma 6.3, we have that $\mathbb{P}(\Gamma(\ell)) \leq 24(n+1)^{2d} \exp\left[-\frac{K^2(\ln(n))^{2\beta}}{1152}\right]$.

We have

$$\begin{aligned}\sum_{j=1}^n w_j^\diamond(\mathbf{x})^2 &= \sum_{j=1}^n \frac{w_j^\diamond(\mathbf{x})}{k} \left(\sum_{\ell=1}^k \frac{\mathbb{1}_{\{\mathbf{X}^{\diamond j} \in A_n(\ell)\}}}{N_n^\diamond(A_n(\ell))} (\mathbb{1}_{\{\Gamma(\ell)^c\}} + \mathbb{1}_{\{\Gamma(\ell)\}}) \right) \\ &\leq \sum_{j=1}^n w_j^\diamond(\mathbf{x}) \left(\frac{2}{K\sqrt{n}(\ln n)^\beta} + \frac{1}{k} \sum_{\ell=1}^k \mathbb{1}_{\{\mathbf{X}^{\diamond j} \in A_n(\ell)\}} \mathbb{1}_{\{\Gamma(\ell)\}} \right).\end{aligned}$$

So that,

$$\begin{aligned}\mathbb{E} \left[\prod_{j=1}^n e^{t^2 w_j^\diamond(\mathbf{x})^2 / 2} \right] &\leq \exp \left[t^2 / \left(K\sqrt{n}(\ln(n))^\beta \right) \right] \times \mathbb{E} \left[\exp \left(\frac{t^2}{2} \cdot \mathbb{1}_{\{\cup_{\ell=1}^k \Gamma(\ell)\}} \right) \right] \\ &\leq \exp \left[t^2 / \left(K\sqrt{n}(\ln(n))^\beta \right) \right] \times \left(1 + e^{t^2/2} \sum_{\ell=1}^k \mathbb{P}(\Gamma(\ell)) \right) \\ &\leq \exp \left[t^2 / \left(K\sqrt{n}(\ln(n))^\beta \right) \right] \\ &\quad \times \left(1 + 24k(n+1)^{2d} \exp \left[\frac{t^2}{2} - \frac{K^2(\ln(n))^{2\beta}}{1152} \right] \right).\end{aligned}$$

Taking $t^2 = \frac{K^2(\ln(n))^{2\beta}}{576}$ leads to

$$\mathbb{P}(W_n > \varepsilon) \leq \left(1 + 24k(n+1)^{2d}\right) \exp \left[\frac{K(\ln(n))^\beta}{576\sqrt{n}} - \frac{\varepsilon K(\ln(n))^\beta}{24} \right].$$

The same upper bound is obtained for $\mathbb{P}(W_n < -\varepsilon)$ by using that $\mathbb{P}(W_n < -\varepsilon) = \mathbb{P}(-W_n > \varepsilon)$.

Thus, by using Assumption 4.3, item 1., $k = O(n^\alpha)$ so that the right hand side is summable, we conclude that W_n goes to 0 almost surely.

In the case where Assumption 4.3 item 3 is satisfied, the proof that $(W_n)_{n \geq 1}$ goes to 0 a.s. is done in a similar spirit, the sketch of proof under Assumption 4.3 item 3, is proposed in the Appendix.

Finally, we show that $(V_n)_{n \geq 1}$ goes to 0 a.s. which easily follows from Assumption 4.2. This allows us to conclude that

$$\forall \mathbf{x} \in \mathcal{X}, \forall y \in \mathbb{R}, \quad F_{k,n}^\diamond(y | \mathbf{X} = \mathbf{x}) \xrightarrow[n \rightarrow \infty]{a.s.} F(y | \mathbf{X} = \mathbf{x}).$$

□

Thus, we have proved Theorem 4.2, the proof of Theorem 4.3 which is a bit simpler is left to the reader.

7. Conclusion

This article proposes two conditional distribution functions and conditional quantiles approximations based on random forests. The former is a natural generalisation of the random forest estimator of the regression function making use of the bootstrap samples, while the latter is based on a variant using only the original dataset.

The consistency of the bootstrap samples based estimator is shown under realistic assumptions and constitutes the major contribution of this paper. Indeed, this is the first consistency result handling the bootstrap component in a random forest method whereas it is usually replaced by subsampling. As for the second estimator, the consistency proof established in [25] for a simplified random forest model is extended to a realistic one by taking into account all the randomness used in the procedure. The two estimators have close performances on our toy example. A specific interest of the bootstrap estimation is that the Out-Of-Bag samples could be used for cross-validation and / or back-testing procedures.

The estimators developed in this paper rest on trees grown with the CART-split criterion. But the assumptions providing the consistency results are detached from the split procedure used. Thus, the theoretical tools developed here could be useful for a large class of methods by just changing the splitting scheme. An ambitious additional work would be to develop a theoretical analysis for obtaining convergence rates and also to construct confidence intervals.

Appendix A: Technical complements

We provide some technical Lemmas used in Section 5 and Section 6.

Lemma A.1. *Let \mathcal{A} be the class of rectangles in \mathbb{R}^d , let $\kappa > 0$, for any n ,*

$$\mathbb{P} \left(\sup_{B \in \mathcal{A}} \left| \frac{N_B^b}{n} - \mathbb{P}(\mathbf{X} \in B) \right| > \kappa \right) \leq 16(n+1)^{2d} e^{-n\kappa^2/128}.$$

Proof. Denote by N_B the number of elements of the original sample \mathcal{D}_n that are in B . We have:

$$\begin{aligned} & \mathbb{P} \left(\sup_{B \in \mathcal{A}} \left| \frac{N_B^b}{n} - \mathbb{P}(\mathbf{X} \in B) \right| > \kappa \right) \\ & \leq \mathbb{P} \left(\sup_{B \in \mathcal{A}} \left| \frac{N_B}{n} - \mathbb{P}(\mathbf{X} \in B) \right| > \kappa/2 \right) + \mathbb{P} \left(\sup_{B \in \mathcal{A}} \left| \frac{N_B^b}{n} - \frac{N_B}{n} \right| > \kappa/2 \right) \\ & \leq \mathbb{P} \left(\sup_{B \in \mathcal{A}} \left| \frac{N_B}{n} - \mathbb{P}(\mathbf{X} \in B) \right| > \kappa/2 \right) \\ & \quad + \mathbb{E} \left(\mathbb{P} \left(\sup_{B \in \mathcal{A}} \left| \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{\mathbf{X}^{*i} \in B\}} - \mathbb{P}(\mathbf{X}^* \in B | \mathcal{D}_n) \right| > \kappa/2 \mid \mathcal{D}_n \right) \right), \end{aligned}$$

where \mathbf{X}^* is a bootstrap sample, whose distribution conditionally to \mathcal{D}_n is uniform on $\{\mathbf{X}^1, \dots, \mathbf{X}^n\}$. We apply Vapnik-Chervonenkis ([36]) inequality conditionally to \mathcal{D}_n to get that the second term is bounded above by $8(n+1)^{2d} e^{-n\kappa^2/128}$ and Vapnik-Chervonenkis inequality to get that the first term admits the same above bound. \square

Lemma A.2. *Let (\mathbf{X}_i^*, Y_i^*) be a bootstrap sample from $(\mathbf{X}_1, Y_1), \dots, (\mathbf{X}_n, Y_n)$. For any $D > 0$, for any $\kappa > 0$, $A \in \mathcal{A}$ and $\frac{1}{p} + \frac{1}{q} = 1$. Assume that $\mathbb{E}(|Y|^p) < \infty$ then*

$$\begin{aligned} & \mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n Y_i^* \mathbf{1}_{\{\mathbf{X}_i^* \in A\}} - \mathbb{E}(Y \mathbf{1}_{\{\mathbf{X} \in A\}}) \right| > \kappa \right) \\ & \leq 6 \left(\frac{24eD}{\kappa} \ln \left(\frac{48eD}{\kappa} \right) \right)^{2d} e^{-n\kappa^2/(128D^2)} + C \frac{\mathbb{E}(Y^p)^{\frac{1}{p}} \mathbb{P}(Y > D)^{\frac{1}{q}}}{\kappa}. \end{aligned}$$

Proof. Let $Z = \min(Y, D)$, $Z_i = \min(Y_i, D)$, $Z_i^* = \min(Y_i^*, D)$. Write:

$$\begin{aligned} & \left| \frac{1}{n} \sum_{i=1}^n Y_i^* \mathbf{1}_{\{\mathbf{X}_i^* \in A\}} - \mathbb{E}(Y \mathbf{1}_{\{\mathbf{X} \in A\}}) \right| \\ & \leq \mathbb{E}((Y - Z) \mathbf{1}_{\{\mathbf{X} \in A\}}) + \left| \frac{1}{n} \sum_{i=1}^n (Y_i^* - Z_i^*) \mathbf{1}_{\{\mathbf{X}_i^* \in A\}} \right| + \\ & \quad \left| \frac{1}{n} \sum_{i=1}^n Z_i^* \mathbf{1}_{\{\mathbf{X}_i^* \in A\}} - \mathbb{E}(Z \mathbf{1}_{\{\mathbf{X} \in A\}}) \right|. \end{aligned}$$

The probability that the sum of the first two terms is greater than κ is bounded above by $C \frac{\mathbb{E}(Y^p)^{\frac{1}{p}} \mathbb{P}(Y > D)^{\frac{1}{q}}}{\kappa}$. The probability that the last term is greater than κ is bounded by following the proof of Theorem 9.6 in [18] page 155, once conditionally to \mathcal{D}_n to take into account the bootstrap sample and then unconditionally. \square

Let us now turn to the proof of Lemma 6.2 which shows that the dummy estimator $F_{k,n}^\diamond$ is close to the interesting one $F_{k,n}^b$.

Proof of Lemma 6.2.

Let $\mathbf{x} \in \mathcal{X}$ and $y \in \mathbb{R}$, recall that $B_j(\Theta_\ell^1, \mathcal{D}_n)$ has been defined in Section 3.1. We have

$$\begin{aligned} & |F_{k,n}^\diamond(y | \mathbf{X} = \mathbf{x}) - F_{k,n}^b(y | \mathbf{X} = \mathbf{x})| \\ &= \left| \frac{1}{k} \sum_{\ell=1}^k \left(\sum_{j=1}^n \frac{\mathbb{1}_{\{\mathbf{X}^{\diamond j} \in A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)\}} \mathbb{1}_{\{Y^{\diamond j} \leq y\}}}{N_n^\diamond(\mathbf{x}; \Theta_\ell, \mathbf{X}^{\diamond 1}, \dots, \mathbf{X}^{\diamond n}, \mathcal{D}_n)} \right. \right. \\ & \quad \left. \left. - \sum_{j=1}^n \frac{B_j(\Theta_\ell^1, \mathcal{D}_n) \mathbb{1}_{\{\mathbf{X}^j \in A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)\}} \mathbb{1}_{\{Y^j \leq y\}}}{N_n^b(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)} \right) \right| \\ &= \left| \frac{1}{k} \sum_{\ell=1}^k \left(\frac{\#\{j \leq J^\diamond / \mathbf{X}^{\diamond(j)} \in A_n(\Theta_\ell)\}}{N^\diamond(A_n(\Theta_\ell))} - \frac{\sum_{j \in \mathcal{S}} B_j(\Theta_\ell^1, \mathcal{D}_n)}{N^b(A_n(\Theta_\ell))} \right) \right|, \end{aligned}$$

with $\mathcal{S} = \{j \leq J / \mathbf{X}^{(j)} \in A_n(\Theta_\ell)\}$, where we denote $A_n(\Theta_\ell) = A_n(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)$, $N^\diamond(A_n(\Theta_\ell)) = N_n^\diamond(\mathbf{x}; \Theta_\ell, \mathbf{X}^{\diamond 1}, \dots, \mathbf{X}^{\diamond n}, \mathcal{D}_n)$ and $N^b(A_n(\Theta_\ell)) = N_n^b(\mathbf{x}; \Theta_\ell, \mathcal{D}_n)$. J, J^\diamond are such that $Y^{\diamond(J^\diamond)} \leq y < Y^{\diamond(J^\diamond+1)}$ and $Y^{(J)} \leq y < Y^{(J+1)}$, with $Y^{\diamond(j)}$ (resp. $Y^{(j)}$) the order statistics of $(Y^{\diamond 1}, \dots, Y^{\diamond n})$ (resp. (Y^1, \dots, Y^n)) and the $\mathbf{X}^{\diamond(j)}$ (resp. $\mathbf{X}^{(j)}$) the corresponding $\mathbf{X}^{\diamond p}$'s (resp. \mathbf{X}^p 's).

Let us consider for some $\ell \in \llbracket 1, k \rrbracket$,

$$\begin{aligned} G &= \frac{\#\{j \leq J^\diamond / \mathbf{X}^{\diamond(j)} \in A_n(\Theta_\ell)\}}{N^\diamond(A_n(\Theta_\ell))} - \frac{\sum_{j \in \mathcal{S}} B_j(\Theta_\ell^1, \mathcal{D}_n)}{N^b(A_n(\Theta_\ell))} \\ &\stackrel{\text{def}}{=} \frac{N_{J^\diamond}^\diamond(A_n(\Theta_\ell))}{N^\diamond(A_n(\Theta_\ell))} - \frac{N_J(A_n(\Theta_\ell))}{N^b(A_n(\Theta_\ell))}. \end{aligned}$$

We have,

$$\begin{aligned} |G| &\leq \frac{|N^\diamond(A_n(\Theta_\ell)) - N^b(A_n(\Theta_\ell))|}{N^b(A_n(\Theta_\ell))} + \frac{|N_{J^\diamond}^\diamond(A_n(\Theta_\ell)) - N_J(A_n(\Theta_\ell))|}{N^b(A_n(\Theta_\ell))} \\ &\stackrel{\text{def}}{=} |G_1| + |G_2| \end{aligned}$$

We continue the proof below in the case where Assumption 4.3 item 3. is satisfied. The case where item 2. is verified is done easier following the same

lines. Let $\varepsilon > 0$. We are now going to show the almost everywhere convergence to 0 for each term G_1 and G_2 . Let us start with G_1 .

$$\begin{aligned} & \mathbb{P}(|G_1| > \varepsilon) \\ &= \mathbb{P}\left(\frac{|N^\diamond(A_n(\Theta_\ell)) - N^b(A_n(\Theta_\ell))|}{N^b(A_n(\Theta_\ell))} > \varepsilon\right) \\ &= \mathbb{P}\left(|N^\diamond(A_n(\Theta_\ell)) - N^b(A_n(\Theta_\ell))| > \varepsilon N^b(A_n(\Theta_\ell)), N^b(A_n(\Theta_\ell)) > \lambda\right) \\ &\quad + \mathbb{P}\left(|N^\diamond(A_n(\Theta_\ell)) - N^b(A_n(\Theta_\ell))| > \varepsilon N^b(A_n(\Theta_\ell)), N^b(A_n(\Theta_\ell)) \leq \lambda\right) \end{aligned}$$

where $\lambda = \frac{\mathbb{E}[N^b(A_n(\Theta))]}{2}$

$$\leq \mathbb{P}\left(|N^\diamond(A_n(\Theta_\ell)) - N^b(A_n(\Theta_\ell))| > \varepsilon\lambda\right) + \mathbb{P}\left(N^b(A_n(\Theta_\ell)) \leq \lambda\right) .$$

Thanks to Bienaymé-Tchebychev's inequality,

$$\begin{aligned} \mathbb{P}\left(N^b(A_n(\Theta_\ell)) \leq \lambda\right) &\leq 4 \frac{\text{Var}(N^b(A_n(\Theta_\ell)))}{(\mathbb{E}[N^b(A_n(\Theta_\ell))])^2} \\ &\leq 4 (\text{CV}(N^b(A_n(\Theta))))^2 . \end{aligned} \quad (\text{A.1})$$

Now, using Lemma 6.3 and Assumption 4.3, we get

$$\begin{aligned} \mathbb{P}(|G_1| > \varepsilon) &\leq \mathbb{P}\left(|N^\diamond(A_n(\Theta_\ell)) - N^b(A_n(\Theta_\ell))| > \varepsilon\lambda\right) + 4 (\text{CV}(N^b(A_n(\Theta))))^2 \\ &\leq 24(n+1)^{2d} \exp\left[-\frac{\varepsilon^2 K^2 (\ln(n))^{2\beta}}{1152}\right] + \frac{4M^2}{n^{\alpha+1} (\ln(n))^\gamma} . \end{aligned}$$

Then, thanks to Borel–Cantelli Lemma

$$\forall \varepsilon > 0, \quad \mathbb{P}\left(\limsup_{n \rightarrow \infty} \{|G_1| > \varepsilon\}\right) = 0 ,$$

which implies $G_1 \xrightarrow[n \rightarrow \infty]{a.s.} 0$.

The G_2 term is treated by using again the Vapnik-Chervonenkis theory. By considering the class $\mathcal{B} = \left\{ \prod_{i=1}^d [a_i, b_i] \times]-\infty, y] : a_i, b_i \in \overline{\mathbb{R}} \right\}$, it gives (following the lines of proof of Lemma A.1)

$$\begin{aligned} \mathbb{P}(|G_2| > \varepsilon) &\leq \mathbb{P}\left(|N_{J^\diamond}^\diamond(A_n(\Theta_\ell)) - N_J(A_n(\Theta_\ell))| > \varepsilon\lambda\right) + 4 (\text{CV}(N(A_n(\Theta))))^2 \\ &\leq 24(n+1)^{2d} \exp\left[-\frac{\varepsilon^2 K^2 (\ln(n))^{2\beta}}{1152}\right] + \frac{4M^2}{n^{\alpha+1} (\ln(n))^\gamma} . \end{aligned}$$

Thanks to Borel–Cantelli Lemma, we get

$$\forall \varepsilon > 0, \quad \mathbb{P}\left(\limsup_{n \rightarrow \infty} \{|G_2| > \varepsilon\}\right) = 0 ,$$

which implies that $G_2 \xrightarrow[n \rightarrow \infty]{a.s.} 0$.

We conclude that G goes to 0 for all ℓ , thus,

$$\forall \mathbf{x} \in \mathcal{X}, \forall y \in \mathbb{R}, \quad |F_{k,n}^\diamond(y | \mathbf{X} = \mathbf{x}) - F_{k,n}^b(y | \mathbf{X} = \mathbf{x})| \xrightarrow[n \rightarrow \infty]{a.s.} 0.$$

In the case where Assumption 4.3 item 2. is verified, it exists $K > 0$ such that

$$N^b(A_n(\Theta_\ell)) \geq K\sqrt{n}(\ln(n))^\beta \text{ a.s.}$$

So that $\mathbb{P}(|G_1| > \varepsilon)$ and $\mathbb{P}(|G_2| > \varepsilon)$ are bounded above respectively by

- $\mathbb{P}\left(|N^b(A_n(\Theta_\ell)) - N^\diamond(A_n(\Theta_\ell))| > \varepsilon K\sqrt{n}(\ln(n))^\beta\right),$
- and $\mathbb{P}\left(|N_{J^\diamond}^\diamond(A_n(\Theta_\ell)) - N_J(A_n(\Theta_\ell))| > \varepsilon K\sqrt{n}(\ln(n))^\beta\right).$

A simple application of Lemma 6.3 and an adaptation of it to $N_J(A_n(\Theta_\ell))$ show that G_1 and G_2 go to 0 a.s. \square

Now we turn to Lemma 6.3 which is key for the proof of Proposition 6.1 and that of Lemma 6.2.

Proof of Lemma 6.3.

Let $\varepsilon > 0$ and $\mathbf{x} \in \mathcal{X}$, we have

$$\begin{aligned} & \mathbb{P}\left(|N^b(A_n(\Theta)) - N^\diamond(A_n(\Theta))| > \varepsilon\right) \\ \leq & \mathbb{P}\left(\left|\frac{N^b(A_n(\Theta))}{n} - \frac{1}{n} \sum_{j=1}^n \mathbf{1}_{\{\mathbf{X}^j \in A_n(\Theta)\}}\right| > \frac{\varepsilon}{3n}\right) \\ & + \mathbb{P}\left(\left|\frac{1}{n} \sum_{j=1}^n \mathbf{1}_{\{\mathbf{X}^j \in A_n(\Theta)\}} - \mathbb{P}_{\mathbf{X}}(\mathbf{X} \in A_n(\Theta))\right| > \frac{\varepsilon}{3n}\right) \\ & + \mathbb{P}\left(\left|\frac{N^\diamond(A_n(\Theta))}{n} - \mathbb{P}_{\mathbf{X}}(\mathbf{X} \in A_n(\Theta))\right| > \frac{\varepsilon}{3n}\right) \\ \leq & \mathbb{P}\left(\sup_{A \in \mathcal{B}} \left|\frac{1}{n} \sum_{j=1}^n \mathbf{1}_{\{\mathbf{X}^{*j} \in A\}} - \frac{1}{n} \sum_{j=1}^n \mathbf{1}_{\{\mathbf{X}^j \in A\}}\right| > \frac{\varepsilon}{3n}\right) \\ & + \mathbb{P}\left(\sup_{A \in \mathcal{B}} \left|\frac{1}{n} \sum_{j=1}^n \mathbf{1}_{\{\mathbf{X}^j \in A\}} - \mathbb{P}_{\mathbf{X}}(\mathbf{X} \in A)\right| > \frac{\varepsilon}{3n}\right) \\ & + \mathbb{P}\left(\sup_{A \in \mathcal{B}} \left|\frac{1}{n} \sum_{j=1}^n \mathbf{1}_{\{\mathbf{X}^{\diamond j} \in A\}} - \mathbb{P}_{\mathbf{X}}(\mathbf{X} \in A)\right| > \frac{\varepsilon}{3n}\right) \end{aligned}$$

where $\mathbf{X}^{*1}, \dots, \mathbf{X}^{*n}$ denotes a bootstrap sample and $\mathcal{B} = \left\{ \prod_{i=1}^d [a_i, b_i] : a_i, b_i \in \overline{\mathbb{R}} \right\}$. The last two right-hand terms are handled thanks to a direct application of the

Theorem of [36] over the class \mathcal{B} whose Vapnik-Chervonenkis dimension is $2d$. This class is nothing more than an extension of the class \mathcal{R} of rectangles in \mathbb{R}^d . Following the lines of the proof of Theorem 13.8 in [10], one sees that the classes \mathcal{R} and \mathcal{B} have the same Vapnik-Chervonenkis dimension.

The first right hand term is treated by applying Vapnik-Chervonenkis' Theorem under the conditional distribution given \mathcal{D}_n as in the proof of Lemma A.1.

Finally, we get the overall upper bound

$$\mathbb{P}(|N^b(A_n(\Theta)) - N^\diamond(A_n(\Theta))| > \varepsilon) \leq 24(n+1)^{2d} e^{-\varepsilon^2/288n}.$$

□

Let us now sketch the part of proof of Proposition 6.1 which states that $(W_n)_{n \geq 1}$ goes to 0 a.s., under Assumption 4.3 item 3.

Proof. 1. First show that $(W_{n^2})_{n \geq 1}$ goes to 0 a.s. This is achieved by decomposing

$$\begin{aligned} \mathbb{E}[(W_n)^2] &= \mathbb{E}\left[\left(\sum_{j=1}^n w_j^\diamond(\mathbf{x}) Z_j^\diamond\right)^2\right] \\ &= \sum_{j=1}^n \sum_{m=1}^n \mathbb{E}[w_j^\diamond(\mathbf{x}) w_m^\diamond(\mathbf{x}) Z_j^\diamond Z_m^\diamond] \\ &= \sum_{j=1}^n \mathbb{E}[w_j^{\diamond 2}(\mathbf{x}) Z_j^{\diamond 2}] + \sum_{\substack{1 \leq j, m \leq n \\ j \neq m}} \mathbb{E}[w_j^\diamond(\mathbf{x}) w_m^\diamond(\mathbf{x}) Z_j^\diamond Z_m^\diamond] \\ &\stackrel{\text{def}}{=} I_n + J_n \end{aligned}$$

Bienaimé-Tchebychev's inequality, Lemma 6.3 and Assumption 4.3 items 1. and 3., give that there exist C, K and M positive constants such that

$$\begin{aligned} I_n &\leq k\mathbb{P}(|N^b(A_n(\Theta)) - N^\diamond(A_n(\Theta))| > \lambda) \\ &\quad + \frac{4}{\mathbb{E}[N^b(A_n(\Theta))]} + 4k(\text{CV}(N^b(A_n(\Theta))))^2 \\ &\leq 24Cn^\alpha(n+1)^{2d} \exp\left[-\frac{K^2(\ln(n))^{2\beta}}{4608}\right] \\ &\quad + \frac{4}{K\sqrt{n}(\ln(n))^\beta} + \frac{4CM^2}{n(\ln(n))^\gamma}. \end{aligned}$$

Then, the trick of using a second sample \mathcal{D}_n^\diamond independent of the first one and the random variable Θ is really important to handle the J_n term. Indeed, we have $J_n = 0$ because $\mathbb{E}[Z_m^\diamond \mathbf{X}^{\diamond m}] = 0$ while the equivalent

term encountered in the proof of the Theorem 2 developed by [35] is handled using a conjecture regarding the correlation behavior of the CART algorithm that is difficult to verify (cf. assumption (H2) of [35]). Finally,

$$\forall \varepsilon > 0, \quad \mathbb{P}(|W_n| \geq \varepsilon) \leq \frac{\mathbb{E}[(W_n)^2]}{\varepsilon^2} = \frac{I_n}{\varepsilon^2}.$$

Hence, since $\sum_{n \geq 1} I_n < \infty$, Borel–Cantelli Lemma gives

$$\forall \varepsilon > 0, \quad \mathbb{P}\left(\limsup_{n \rightarrow \infty} \{|W_n| \geq \varepsilon\}\right) = 0,$$

which implies that $W_n \xrightarrow[n \rightarrow \infty]{a.s.} 0$.

2. Show that $(W_n)_{n \geq 1}$ converges almost surely to 0. Bienaimé-Tchebychev inequality, Lemma 6.3 and Assumption 4.3 items 1. and 3. as well as Borel-Cantelli Lemma give that

$$\forall \varepsilon > 0, \quad \mathbb{P}\left(\limsup_{n \rightarrow \infty} \{|W_n - W_{p^2}| \geq \varepsilon\}\right) = 0,$$

for $p = p(n) = \lfloor \sqrt{n} \rfloor$. From this, we deduce that $(W_n)_{n \geq 1}$ goes to 0 a.s. \square

We end this Appendix with the proof that sums of products of functions considered in Remark 4.2 belong to the \spadesuit -class. The proof for the other classes of function may be done straightforwardly in the same way.

Lemma A.3. *Let \mathcal{I} be a family of subsets of $\{1, \dots, d\}$ and $f_j^I, j \in \{1, \dots, d\}, I \in \mathcal{I}$, be functions either all increasing or all decreasing and they are either all positive or all negative then sums of product functions*

$$f(\mathbf{x}) = \sum_{I \in \mathcal{I}} \prod_{j \in I} f_j^I(x_j)$$

belong to the \spadesuit -class.

Proof. Let $A = \prod_{k=1}^d [a_k, b_k]$ and $j \in \{1, \dots, d\}$ be fixed, for any $z \in [a_j, b_j]$,

$$\begin{aligned} \mathbb{E}(f(z, \mathbf{X}^{-j}) \mathbb{1}_{\{\mathbf{X}^{-j} \in A^{-j}\}}) &= \sum_{I \in \mathcal{I}, j \in I} f_j^I(z) \mathbb{E} \left(\prod_{k \in I \setminus \{j\}} f_k^I(X_k) \mathbb{1}_{\{X_k \in [a_k, b_k]\}} \right) \\ &+ \sum_{I \in \mathcal{I}, j \notin I} \mathbb{E} \left(\prod_{k \in I} f_k^I(X_k) \mathbb{1}_{\{X_k \in [a_k, b_k]\}} \right). \end{aligned}$$

So that $z \mapsto \mathbb{E}(f(z, \mathbf{X}^{-j})\mathbb{1}_{\{\mathbf{x}^{-j} \in A^{-j}\}})$ is constant on $[a_j, b_j]$ if and only if for any $z, z' \in [a_j, b_j]$,

$$\sum_{I \in \mathcal{I}, j \in I} (f_j^I(z) - f_j^I(z')) \mathbb{E} \left(\prod_{k \in I \setminus \{j\}} f_k^I(X_k) \mathbb{1}_{\{X_k \in [a_k, b_k]\}} \right) = 0,$$

since the functions f_k^I are either all increasing or all decreasing and they are either all positive or all negative, this may happen for all $z, z' \in [a_j, b_j]$, only if f_j^I is constant on $[a_j, b_j]$. \square

Acknowledgments

We are grateful to Andrés Cuberos, Ecaterina Nisipasu, Mathieu Poulin and Przemyslaw Sloma from SCOR for their valuable comments and support. We are also much indebted to Roland Denis and Benoit Fabrèges for intensive support on computational aspects. We are grateful to anonymous reviewers of previous versions of the paper, their comments helped to improve the paper.

References

- [1] Yali Amit and Donald Geman. Shape quantization and recognition with randomized trees. *Neural computation*, 9(7):1545–1588, 1997.
- [2] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017.
- [3] Gérard Biau and Luc Devroye. On the layered nearest neighbour estimate, the bagged nearest neighbour estimate and the random forest method in regression and classification. *Journal of Multivariate Analysis*, 101(10):2499–2518, 2010.
- [4] Gérard Biau. Analysis of a random forests model. *Journal of Machine Learning Research*, 13(Apr):1063–1095, 2012.
- [5] Leo Breiman. Bagging predictors. *Machine learning*, 24(2):123–140, 1996.
- [6] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [7] Leo Breiman. Consistency for a simple model of random forests. 2004.
- [8] Leo Breiman, Jerome H Friedman, Richard A Olshen, and Charles J Stone. Classification and regression trees. *Wadsworth and Brooks/Cole Monterey, CA, USA*, 1984.
- [9] Thomas Browne, Jean-Claude Fort, Bertrand Iooss, and Loïc Le Gratiet. Estimate of quantile-oriented sensitivity indices. 2017.
- [10] Luc Devroye, László Györfi, and Gábor Lugosi. *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media, 2013.
- [11] Thomas G Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2000.
- [12] B. Efron. Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, 7:1–26, 1979.

- [13] Kévin Elie-Dit-Cosaque. qosa-indices, a python package available at: https://gitlab.com/qosa_index/qosa, 2020.
- [14] Kévin Elie-Dit-Cosaque and Véronique Maume-Deschamps. Goal-oriented shapley effects with a special attention to the quantile-oriented case. *SIAM/ASA Journal on Uncertainty Quantification - JUQ*, to appear.
- [15] Benoit Fabrière and Véronique Maume-Deschamps. Conditional distribution forest: a julia package available at <https://github.com/bfabreges/conditionaldistributionforest.jl>, 2020.
- [16] Jean-Claude Fort, Thierry Klein, and Nabil Rachdi. New sensitivity analysis subordinated to a contrast. *Communications in Statistics-Theory and Methods*, 45(15):4349–4364, 2016.
- [17] Benjamin Goehry. Random forests for time-dependent processes. 2019.
- [18] László Györfi, Michael Kohler, Adam Krzyzak, and Harro Walk. *A distribution-free theory of nonparametric regression*. Springer Science & Business Media, 2006.
- [19] Tin Kam Ho. The random subspace method for constructing decision forests. *IEEE transactions on pattern analysis and machine intelligence*, 20(8):832–844, 1998.
- [20] Jason M. Klusowski. Analyzing cart. 2020.
- [21] Roger Koenker and Kevin F Hallock. Quantile regression. *Journal of economic perspectives*, 15(4):143–156, 2001.
- [22] Yi Lin and Yongho Jeon. Random forests and adaptive nearest neighbors. *Journal of the American Statistical Association*, 101(474):578–590, 2006.
- [23] Véronique Maume-Deschamps and Ibrahima Niang. Estimation of quantile oriented sensitivity indices. *Statistics & Probability Letters*, 134:122–127, 2018.
- [24] Véronique Maume-Deschamps, Didier Rullière, and A Usseglio-Carleve. Quantile predictions for elliptical random fields. *Journal of Multivariate Analysis*, 159:1–17, 2017.
- [25] Nicolai Meinshausen. Quantile regression forests. *Journal of Machine Learning Research*, 7(Jun):983–999, 2006.
- [26] Nicolai Meinshausen. Quantile regression forests, a r package available at <https://cran.r-project.org/package=quantregforest>., 2019.
- [27] 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.
- [28] Jooyoung Park and Irwin W Sandberg. Universal approximation using radial-basis-function networks. *Neural computation*, 1991.
- [29] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [30] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2019.

- [31] Erwan Scornet. On the asymptotics of random forests. *Journal of Multivariate Analysis*, 146:72–83, 2016.
- [32] Erwan Scornet. Promenade en forêts aléatoires. *MATAPLI*, 111, 2016.
- [33] Erwan Scornet. Random forests and kernel methods. *IEEE Transactions on Information Theory*, 62(3):1485–1500, 2016.
- [34] Erwan Scornet, Gérard Biau, and Jean-Philippe Vert. Supplementary materials for: Consistency of random forests. *arXiv*, 1510, 2015.
- [35] Erwan Scornet, Gérard Biau, Jean-Philippe Vert, et al. Consistency of random forests. *The Annals of Statistics*, 43(4):1716–1741, 2015.
- [36] V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.
- [37] Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.
- [38] Stefan Wager and Guenther Walther. Adaptive concentration of regression trees, with application to random forests. *arXiv preprint arXiv:1503.06388*, 2015.