

On the Consistency of a Random Forest Algorithm in the Presence of Missing Entries

Irving Gómez-Méndez Emilien Joly

May 21, 2021

Abstract

This paper tackles the problem of constructing a non-parametric predictor when the latent variables are given with incomplete information. The convenient predictor for this task is the random forest algorithm in conjunction to the so-called CART criterion. The proposed technique enables a partial imputation of the missing values in the data set in a way that suits both a consistent estimator of the regression function as well as a partial recovery of the missing values. A proof of the consistency of the random forest estimator is given in the case where each latent variable is missing completely at random (MCAR).

1 Introduction

Random forests algorithms are widely used in applied computer science nowadays. They provide very useful tools for extracting information of a possibly large scaled data set throughout two tasks: regression and classification. First introduced by Breiman (2001), this non-parametric technique has various computational benefits over many regression/prediction techniques as it is fast and simple. Therefore it is really common to see the random forest algorithm (or parts of it) as a building block of more intricate machine learning algorithms. The small number of hyper-parameters to be tuned, along with the bootstrap-aggregation technique (Breiman, 1996) -commonly called bagging- allows a straightforward parallelization and has made random forests one of the most popular tools for handling data sets of large size. It has been successfully involved in various practical problems, including chemoinformatics (Svetnik et al., 2003), ecology (Prasad, Iverson, and Liaw, 2006; Cutler et al., 2007), 3D object recognition (Shotton et al., 2011), bioinformatics (Díaz-Uriarte and De Andres, 2006) and econometrics (Varian, 2014).

On the theoretical side, little is known about the mathematical guarantees of the method. This has lead the community to underline a “gap” between theory and practice. Part of this gap can be explained by the bagging mechanism and the splitting criterion. Each of these processes introduces a source of randomness into the construction of the trees which makes the random forest algorithms very challenging to study in their full generality. One way to overcome the theoretical difficulties of random forests is through simplified versions of the original procedure. This is often done by simply ignoring the bagging step and/or by replacing the splitting criterion with a more elementary splitting protocol (e.g. Breiman (2004)). However, in recent years, important theoretical studies have been performed to analyze more elaborated models (e.g. Biau, Devroye, and Lugosi (2008), Ishwaran and Kogalur (2010), Biau (2012), Genuer (2012), and Zhu, Zeng, and Kosorok (2015)). Consistency and asymptotic normality for Breiman’s infinite forests were proven in Wager and Athey (2018) simplifying the splitting step. Mentch and Hooker (2016) proved a similar result for finite forests and Scornet, Biau, and Vert (2015) proved a consistency result in the context of additive

regression models. The review presented by Biau and Scornet (2016) gives a very good introduction to those concepts and proposes solutions for some of the challenges presented above.

The relative simplicity of random forests advocates for its use to handle missing data. As an initial solution, Breiman and Cutler (2003) proposed an algorithm which takes advantage of the so-called proximity-matrix to impute the missing values and Ishioka (2013) modified it in order to obtain a new version more robust to outliers. The algorithm MissForest, proposed by Stekhoven and Bühlmann (2011) is another alternative to iteratively improve the imputation of the missing values using random forests. An interesting alternative comes from algorithms that attempt to handle the missing values without performing an imputation. The so-called surrogate splits (see Breiman et al. (1984), Venables and Ripley (2002), Ripley (2007), Friedman, Hastie, and Tibshirani (2009), and Hapfelmeier, Hothorn, and Ulm (2012)) is an example of this family of algorithms. This technique has spread from the original CART trees and has been introduced in conditional trees (Hothorn, Hornik, and Zeileis, 2006). Missing Incorporated in Attributes (MIA) (Twala, Jones, and Hand, 2008) is another approach that has driven attention in recent years. In a practical point of view, Josse et al. (2019) perform a simulation study to compare different methods to handle missing data using trees and study the consistency of two approaches to estimate the prediction function with missing values in a framework of multiple imputation (Rubin, 2004; Rubin, 1996). An overview of different methodologies to handle missing data can be found in Josse, Reiter, et al. (2018).

In the present paper, we propose a new splitting protocol and prove the consistency of the algorithm in the context of an additive model in presence of missing data accordingly to a Missing Completely at Random (MCAR) paradigm. Moreover, in a simulation study, we compare this new algorithm with other approaches that handle missing values using random forests also built with the CART criterion. We choose the following regression function which has been used in previous simulation studies (Breiman, 1996; Rieger, Hothorn, and Strobl, 2010; Josse et al., 2019; Friedberg et al., 2020),

$$m(\mathbf{x}) = 10 \sin \left(\pi \mathbf{x}^{(1)} \mathbf{x}^{(2)} \right) + 20 \left(\mathbf{x}^{(3)} - 0.5 \right)^2 + 10 \mathbf{x}^{(4)} + 5 \mathbf{x}^{(5)}.$$

Following the schema presented by Rieger, Hothorn, and Strobl (2010), we simulate i.i.d. observations uniformly distributed on $[0, 1]^5$ and introduce missing values in $\mathbf{X}^{(1)}$, $\mathbf{X}^{(3)}$ and $\mathbf{X}^{(4)}$. We let the percentage of the missing values in $\mathbf{X}^{(4)}$ vary, to show the performance of this new algorithm in a possibly highly corrupted sample.

The paper is organized as follows. In Section 2, we introduce the random forests, the concept of data-missing mechanisms and show the problems presented by the original CART criterion when there are missing entries. The main consistency result is presented in Section 3 and its proof is given in Section 4. In Section 5, we describe the simulation study and present the conclusions of this study.

2 Random Forests: Aggregating Decision Trees

The random forests algorithms consist in building a collection of objects called decision trees (which are rooted trees) that, most of the time, are binary trees. For a specific tree, each path starting from the root corresponds to a decision (sometimes called prediction) and the descent mechanism is usually a key part in the decision. Each bifurcation in a single tree is done using a random component and an optimization component, representing a split of the space into two parts. Then, moving along a path in the decision tree represents a decision between one of the two parts of the space. Finally, the random forest is made of a set of decision trees that are later aggregated all together. This aggregation step is also important. In this paper, this takes the form of a simple mean (see Equation (2) for a precise statement). The main strength of random forests is that it

aggregates the information of many different decision trees in a global predictor that ends to be a lot more stable (and then a lot more informative) than each specific tree.

Besides being conceptually simple, decision trees are powerful and attractive in practice for several reasons:

- They can model arbitrarily complex relations between the input and the output space.
- They handle categorical or numerical variables, or a mix of both.
- They can be used in regression or supervised classification problems.
- They are easily interpretable, even for non-statisticians.

To construct the partition of the input space, decision trees work in a recursive way. The root of the tree is the whole input space, \mathcal{X} , which is split into disjoint regions. Then each region is split into more regions, and this process is continued until some stopping rule is applied (see Figure 1). Note that, for each cell considered during the construction of the tree, we can estimate the target variable. Hence, the decision tree is the collection of all the nodes and the estimation of the target variable for each one of those nodes. At each step of the tree construction, the partition performed over a cell (or equivalently its corresponding node) is determined by maximizing some splitting criterion (see Equation (3) for a proper definition).




Figure 1: At each step of the tree construction a partition is performed over a cell (or equivalently its corresponding node) maximizing some split criterion.

In our framework, we assume to have access to a training set $\mathcal{D}_n = (\mathbf{X}_i, Y_i)_{i=1, \dots, n}$ where the response variables Y_i are real-valued and the input variables \mathbf{X}_i belong to some space \mathcal{X} . For simplicity, we assume throughout the discussion that the vector \mathbf{X} is a uniform random variable on the square \mathcal{X} . Although appearing restrictive at first sight, this assumption is not essential for the following. It could be replaced by an assumption of the form: \mathbf{X} has a positive density over the full space \mathcal{X} . In most applications the space \mathcal{X} is a compact portion of a p dimensional space. Hence we assume, with very little loss of generality, that $\mathcal{X} = [0, 1]^p$. The task is to predict the random variable

Y with respect to the input vector $\mathbf{X} = (\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(p)})$ where $\mathbf{X}^{(j)} \in [0, 1]$ (for $j = 1, \dots, p$). For example, one could be interested in finding a function $f : [0, 1]^p \rightarrow \mathbb{R}$ that minimizes the functional loss $\mathbb{E}_{\mathbf{X}, Y} [\mathcal{L}(f(\mathbf{X}), Y)]$ where $\mathcal{L} : \mathbb{R}^2 \rightarrow \mathbb{R}$ is the squared error $\mathcal{L}(f(\mathbf{X}), Y) = (f(\mathbf{X}) - Y)^2$. The solution of this optimization problem

$$m = \arg \min_{f: [0, 1]^p \rightarrow \mathbb{R}} \mathbb{E}_{\mathbf{X}, Y} [(f(\mathbf{X}) - Y)^2]$$

is given by the regression function $m(\mathbf{X}) = \mathbb{E}[Y|\mathbf{X}]$. For practical problems, the distribution of (\mathbf{X}, Y) and hence, the regression function, is unknown. The task is then to use the data \mathcal{D}_n to construct a learning model, also called learner, predictor or estimator, $m_n : [0, 1]^p \rightarrow \mathbb{R}$ that estimates the regression function m , and enables us to predict the outcome for new incoming inputs. At least two distinct strategies can be used to solve the estimation problem: the parametric and the non-parametric estimation. In a nutshell, the non-parametric estimation uses a set of functions \mathcal{F}_n (often called model) that serves the minimization (often referred to Empirical Risk Minimization)

$$m_n = \arg \min_{f \in \mathcal{F}_n} \mathbb{E}_{\mathcal{D}_n} [(f(\mathbf{X}) - Y)^2].$$

The main focus of this work is to construct a non-parametric estimation of m through random forests algorithms. As mentioned above, a random forest is a predictor consisting of $M (> 1)$ randomized trees. The randomization is introduced in two different steps of the tree construction and will be governed by the following parameters.

- An integer $a_n \in \{1, \dots, n\}$ defines the number of observations to be subsampled in \mathcal{D}_n for the construction of each tree. Then, a tree built upon those observations (say $\mathbf{X}_{i_1}, \dots, \mathbf{X}_{i_{a_n}}$) depends on \mathcal{D}_n only through those variables.
- An integer $\text{mtry} \in \{1, \dots, p\}$ that corresponds to the number of directions (features) chosen randomly at each step. Those selected features are denoted by $\mathcal{M}_{\text{try}} \subset \{1, \dots, p\}$ and are redefined at each node of the tree construction.
- Two integers nodesize and q_n such that $2q_n - 1 \leq \text{nodesize} \leq a_n$. The number nodesize (resp. q_n) holds for the maximum (resp. minimum) number of observations in a final cell. In the algorithm, a cell A that contains a number of observation $q_n \leq N(A) \leq \text{nodesize}$ is declared **final** and will not be split again.

This randomization (independent from the original source of randomness in the sample \mathcal{D}_n) is represented in a symbolic random variable Θ . To each tree – randomized with the random variable Θ_k – is associated a predicted value at a query point \mathbf{x} , denoted as $m_n(\mathbf{x}; \Theta_k)$. The different trees are constructed by the same procedure but with independent randomization so the random variables $\Theta_1, \dots, \Theta_M$ are i.i.d. with common law Θ . The nature and dimension of Θ depends on its use in the tree construction. In our choice of construction rules, Θ consists of the observations selected for the tree and the candidate directions to be split at each step. The tree construction stops when all the remaining cells are final cells. At the end of the tree construction, a partition of the space \mathcal{X} is returned in a form of the collection of final cells $(A_{n,i})_{i \geq 1}$. Each of these cells corresponds to a leaf of the tree. Finally, the k -th tree estimate is defined as

$$m_n(\mathbf{x}; \Theta_k) = \sum_{i \in \mathcal{I}_{n, \Theta_k}} \frac{Y_i \mathbb{1}_{\mathbf{x} \in A_{n,i}(\mathbf{x}; \Theta_k)}}{N(A_{n,i}(\mathbf{x}; \Theta_k))}, \quad (1)$$

where \mathcal{I}_{n,Θ_k} is the set of the a_n observations selected prior to the construction of the k th tree, $A_n(\mathbf{x}; \Theta_k)$ is the unique final cell that contains \mathbf{x} , and $N(A_n(\mathbf{x}; \Theta_k))$ is the number of observations which belong to the cell $A_n(\mathbf{x}; \Theta_k)$. The aggregation of the trees forms the finite random forest estimator given by

$$m_{M,n}(\mathbf{x}; \Theta_1, \dots, \Theta_M) = \frac{1}{M} \sum_{k=1}^M m_n(\mathbf{x}; \Theta_k). \quad (2)$$

It is known from the work of Breiman (2001) that the random forest estimator does not overfit when M tends to infinity so the value of M is only restricted by the computational power. Besides being important in practice, we do not make big case of the choice of M since the following results show the consistency of each of the tree estimators, hence the consistency of the random forest for any number M . It remains a challenge to be able to show the consistency of the random forest in the case when some trees are not consistent estimators.

Different split criteria have been proposed depending on the statistical problem and the nature of the input space. For regression purposes, the most commonly used is the CART criterion (Breiman et al., 1984). For supervised classification, the split criterion takes the form of an impurity function, like the misclassification error, the Gini index (Gini, 1912) or the Shannon entropy (Shannon, 1948), considered in the criteria known as ID3 and C4.5 (Quinlan, 1986; Quinlan, 1993).

The CART criterion The present work focuses on the CART split criterion that we introduce now. We begin with some important notations.

- A denotes a general node (or cell).
- $N(A)$ holds for the number of points in A .
- The notation $d = (h, z)$ denotes a cut in A , where the integer $h \in \{1, \dots, p\}$ is called a direction (or feature) and $z \in [0, 1]$ is the position of the cut in the h -th direction.
- \mathcal{C}_A is the set of all possible cuts in the node A . It means that h do belong to the set \mathcal{M}_{try} and that for any chosen h , z lies between the bounds of the cell A in that specific direction (see for example Figure 1).
- A cell A is split accordingly to a cut $d = (h, z)$ into two cells denoted as $A_L = \{\mathbf{x} \in A : \mathbf{x}^{(h)} < z\}$ and $A_R = \{\mathbf{x} \in A : \mathbf{x}^{(h)} \geq z\}$.
- \bar{Y}_A (resp. $\bar{Y}_{A_L}, \bar{Y}_{A_R}$) is the empirical mean of the response variable Y_i for the indexes such that \mathbf{X}_i belongs to the cell A (resp. A_L, A_R).

Then, the empirical version of the CART split criterion for a generic cell A is defined as

$$L_n(A, d) = \frac{1}{N(A)} \sum_{i=1}^n (Y_i - \bar{Y}_A)^2 \mathbb{1}_{\mathbf{x}_i \in A} - \frac{1}{N(A)} \sum_{i=1}^n \left(Y_i - \bar{Y}_{A_L} \mathbb{1}_{\mathbf{x}_i^{(h)} < z} - \bar{Y}_{A_R} \mathbb{1}_{\mathbf{x}_i^{(h)} \geq z} \right)^2 \mathbb{1}_{\mathbf{x}_i \in A}, \quad (3)$$

with the convention $0/0 = 0$.

Intuitively, the CART criterion compares (in a weighted manner) the empirical variance inside the cell A with the sum of the two empirical variances of the sub-cells A_R and A_L . At each step of

the creation of the random tree, a current cell A is selected and cut by choosing the best empirical cut $\hat{d} = (\hat{h}, \hat{z})$ so that $L_n(A, d)$ is maximal over the set \mathcal{C}_A , that is

$$\hat{d} = (\hat{h}, \hat{z}) \in \arg \max_{d \in \mathcal{C}_A} L_n(A, d).$$

In the presence of missing values Let us emphasize, through an example, the issues that one faces with the original CART criterion in a context of missing values. Consider the space $\mathcal{X} = [0, 1]^2$ and two observations, \mathbf{X}_1 and \mathbf{X}_2 , that belong to a cell $A \subset \mathcal{X}$. We assume that a direction and location for a cut of A have been chosen and we denote by A_L and A_R the two resulting cells (see Table 1, Figure 2 for an illustration). Since its value is missing, $\mathbf{X}_1^{(1)}$ is represented as a dashed line over the interval $[0, 1]$.

	$\mathbf{x}^{(1)}$	$\mathbf{x}^{(2)}$
\mathbf{X}_1	NA	0.5
\mathbf{X}_2	0.75	0.25
A	[0.3, 0.9]	[0.2, 0.7]
A_L	[0.3, 0.6]	[0.2, 0.7]
A_R	[0.6, 0.9]	[0.2, 0.7]

Table 1: Toy example data. We consider a cell A that has been split into A_L and A_R . The feature $\mathbf{X}^{(1)}$ is missing in the first observation (denoted as NA).

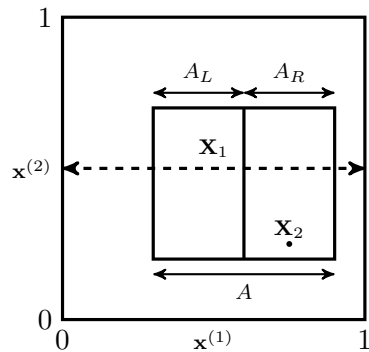


Figure 2: Illustration of the example data, \mathbf{X}_1 is represented as a dashed line since $\mathbf{X}^{(1)}$ is missing.

It is clear that $\mathbf{X}_2 \in A_R$, however it is not possible to decide, without any further operation, if $\mathbf{X}_1 \in A_R$ or $\mathbf{X}_1 \in A_L$, or even if $\mathbf{X}_1 \in A$. The CART criterion is then intractable since the quantities $N(A)$, $N(A_L)$, $N(A_R)$, \bar{Y}_A , \bar{Y}_{A_L} , \bar{Y}_{A_R} , $\mathbb{1}_{\mathbf{x}_i \in A}$, $\mathbb{1}_{\mathbf{x}_i^{(h)} < z}$ and $\mathbb{1}_{\mathbf{x}_i^{(h)} \geq z}$ can not be computed. We propose a new partial imputation algorithm in Section 3.1 to handle these problems.

Definition of Missingness The concept of data-missing mechanisms (introduced by Rubin (1976)) establishes the relationship between missingness and the data. It is common to define the data-missing mechanisms through the data matrix. However, to have a useful definition from a theoretical point of view, we formally define them using the random variables \mathbf{X} and Y . First, let us define a new variable, called the indicator of missing value,

$$\mathbf{M}^{(h)} = \begin{cases} 1 & \text{if } \mathbf{X}^{(h)} \text{ is missing} \\ 0 & \text{otherwise} \end{cases}, \quad 1 \leq h \leq p.$$

A data point (\mathbf{X}, Y) such that $\mathbf{M}^{(h)} = 1$ is said to be *missing in the direction h* . We are assuming throughout this work that the response Y has no missing values, so it is not necessary to define an indicator of missing value for Y . The missingness mechanisms are fully characterized by the information of the conditional distribution of $\mathbf{M}^{(h)}$ given (\mathbf{X}, Y) . One of the data-missing mechanisms identified by Rubin (1976) is the so-called Missing Completely at Random (MCAR). We say that

the variable $\mathbf{X}^{(h)}$ is MCAR if $\mathbf{M}^{(h)}$ is independent of (\mathbf{X}, Y) . In other words, under the MCAR assumption, a coordinate $\mathbf{X}^{(h)}$ has some probability to be missing in the sample and this probability does not depend on the value of \mathbf{X} nor the response variable Y . Note that it is allowed to get missing rates that differ from one direction to another (i.e. $(\mathbf{M}^{(h)})_h$ are not identically distributed) as seen in the hypothesis of Theorem 1.

In this work, we allow the probability of missingness to depend on n . This hypothesis is relevant when the corrupted data in a sample represents a proportion of the entire sample; for example when measuring or sending data of fixed size in a network, a certain number of bits is lost, in average. Note that this does not imply that the distribution of the data is depending on the sample size n but the fact that one specific entry is lost do depend on n . We define the expected proportion of missing values by $p_n^{(h)} = \mathbb{P}[\mathbf{M}^{(h)} = 1 | \mathcal{D}_n]$. See Theorem 1 for the condition on the values of $p_n^{(h)}$.

3 Main Result

3.1 Introduction of our Algorithm

The approach proposed in this paper overpasses the CART computation problems through iterative partial imputations of the missing values that are refined at each step. Unlike most of the so-called imputation techniques, the imputation step is *not* performed independently of the evaluation of the CART criterion. On the contrary, the imputation is itself part of the optimization. The split of a cell (in the usual framework) is then replaced by the couple (split, imputation), where the additional imputation part refers to assign deterministically each data point missing in the direction of the split to one of the two resulting cells. We repeat this step on the resulting cells until a stopping rule is achieved. In the present article we consider splits that keep at least a determined number of points q_n to each child node, and we stop when there is no split with such property.

This algorithm looks similar to other approaches that perform imputation, except for the fact that the imputations take the form of closed intervals rather than punctual values. Each iteration of the algorithm consists of the (split, imputation) task throughout the optimization of the CART criterion. As a result, each iteration takes as an input the current cell A and the current imputations of the missing variables belonging to A , and respond a pair of cells (A_L, A_R) (for left and right) and new imputations. The initialization of the imputation of each missing variable is given by $\widehat{\mathbf{X}}_i^{(h)} = [0, 1]$ which is interpreted as giving to the missing value all the possible values in $[0, 1]$. To avoid confusion in the sequel, we make the subtle difference between the imputation of the missing values at the start (referred to as *in*) of an iteration and the imputation at the end (referred to *out*) of the iteration.

For one cut $d = (h, z)$, we give the following formalism for the assignments of the missing random variables in the direction h . An assignation is a binary vector w of size equal to the number of missing values in the direction h that belong to A , and such that $w_i = 1$ if the i th missing observation is chosen to belong to the left cell A_L and $w_i = 0$ if it is chosen to belong to the right cell A_R . The result of this operation is a vector $\widehat{\mathbf{X}}_w$, such that

$$\widehat{\mathbf{X}}_{i,w}^{(h)} = \begin{cases} [a^{(h)}, z] & \text{if } w_i = 1 \\ [z, b^{(h)}] & \text{if } w_i = 0 \end{cases}$$

if the value of \mathbf{X}_i is missing in the direction h and $\widehat{\mathbf{X}}_{i,w}^{(h)} = \mathbf{X}_i^{(h)}$ otherwise. The values $a^{(h)}$ and $b^{(h)}$ hold for the extreme point of the cell A in the direction h . See Figure 3 for an illustration with $p = 2$.




Figure 3: We perform a cut and assignation of points were the variable is missing, maximizing the CART criterion $L_n(A, d, w)$.

Finally, denoting by $I(A)$ the indexes of the observations belonging the cell A (by previous assignation or being completely observed), the CART criterion in the context of missing values is defined on a cell A by the formula

$$\begin{aligned}
 L_n(A, d, w) &= \frac{1}{\widehat{N}(A)} \sum_{i \in I(A)} \left(Y_i - \widehat{Y}_A \right)^2 \\
 &\quad - \frac{1}{\widehat{N}(A)} \sum_{i \in I(A)} \left(Y_i - \widehat{Y}_{A_L} \right)^2 \mathbb{1}_{a^{(h)} \leq \widehat{\mathbf{X}}_{i,w}^{(h)} \leq z} \\
 &\quad - \frac{1}{\widehat{N}(A)} \sum_{i \in I(A)} \left(Y_i - \widehat{Y}_{A_R} \right)^2 \mathbb{1}_{z < \widehat{\mathbf{X}}_{i,w}^{(h)} \leq b^{(h)}}
 \end{aligned}$$

where \widehat{Y}_{A_L} (resp. \widehat{Y}_{A_R}) is the empirical mean of the Y_i such that $\widehat{\mathbf{X}}_{i,w}$ belongs to the cell A_L (resp. A_R) and $\widehat{N}(A)$ is the number of points of the cell A .

Let $\widehat{N}_{miss}^{(h)}(A) = \sum_{i=1}^n \mathbb{1}_{\mathbf{M}_i^{(h)}=1}$ be the number of observations assigned to the cell A that are missing in the direction h , note that $\sum_i w_i$ is the number of these observations which are assigned to the left node. Thus, $\widehat{p}_L^{(h)} = \sum_i w_i / \widehat{N}_{miss}^{(h)}(A)$ estimates the probability of an observation with the direction h missing to belong to the left node, given that it belongs to the cell A . Similarly $\widehat{p}_R^{(h)} = 1 - \widehat{p}_L^{(h)}$ denotes the probability of an observation with the direction h missing to belong to the right node, given that it belongs to the cell A . We can use these probabilities to predict the target variable for a new observation whose direction h is missing. These can be done stochastically, assigning the new observation to the left node with probability $\widehat{p}_L^{(h)}$ and to the right node with probability $\widehat{p}_R^{(h)}$. If the new observation has a direction missing that was always observed during the training phase, then we might not be able to assign this observation to the left or right nodes when that particular direction was chosen in the construction of the tree. In such case, we can still estimate the target value using the estimation of the current cell. The prediction of the random forest would still be the average of the predictions of the trees. Note that this average would stabilize the stochastic behavior of the individual predictors, in a “law of large numbers” fashion way.

We give the precise algorithm for the CART optimization with assignation in Algorithm 1 and the construction of the random forest in Algorithm 2.

One could ask if the condition for a possible (cut, assignation) in line 5 in Algorithm 1 could never be fulfilled. This case happens if and only if the number of data points in A is less than $2q_n$. This condition is check in Algorithm 2 anytime before Algorithm 1 is called. It basically checks

Algorithm 1 CART optimization with assignation.

Input: A cell A , the corresponding data $(\widehat{\mathbf{X}}_{i,in}, Y_i)_i$ that belongs to A

Output: Two cells A_L, A_R and assignations $(\widehat{\mathbf{X}}_{i,out}, Y_i)_i$ of the data.

- 1: **for** each cut $d = (h, z)$ in \mathcal{C}_A **do**
 - 2: Compute $\widehat{N}_{miss}^{(h)}(A) = \sum_{i=1}^n \mathbb{1}_{\mathbf{M}_i^{(h)}=1}$.
 - 3: Let $\mathcal{W}_A^{(h)} = \{0, 1\}^{\widehat{N}_{miss}^{(h)}(A)}$ be the possible assignations in the direction h .
 - 4: Let $\overline{\mathcal{W}}_A^{(h)} \subset \mathcal{W}_A^{(h)}$ such that A_L and A_R have at least q_n elements.
 - 5: **if** $\overline{\mathcal{W}}_A^{(h)} = \emptyset$ **then**
 - 6: Remove d from \mathcal{C}_A .
 - 7: **else**
 - 8: Compute $\widehat{w}_d \in \arg \max_{w \in \overline{\mathcal{W}}_A^{(h)}} L_n(A, d, w)$.
 - 9: **end if**
 - 10: **end for**
 - 11: Compute $\widehat{d} \in \arg \max_{d \in \mathcal{C}_A} L_n(A, d, \widehat{w}_d)$
 - 12: Define $\widehat{\mathbf{X}}_{i,out}^{(h)} = \widehat{\mathbf{X}}_{i,\widehat{w}_{\widehat{d}}}^{(h)}$
 - 13: Define $A_L = \{\mathbf{x} \in A : \mathbf{x}^{(\widehat{h})} \leq \widehat{z}\}$ and $A_R = \{\mathbf{x} \in A : \mathbf{x}^{(\widehat{h})} > \widehat{z}\}$
-

if the resulting two cells in a split have enough number of points. Such condition is important to ensure the convergence of our algorithm and is restated in Theorem 1. The task of Algorithm 1 is to compute the optimal couple (split, assignation) given by

$$(\widehat{d}, \widehat{w}) \in \arg \max_{\substack{d \in \mathcal{C}_A \\ w \in \overline{\mathcal{W}}_A^{(h)}}} L_n(A, d, w).$$

On admissible assignations. At first sight, one could think that the number of possible assignations defined by $\mathcal{W}_A^{(h)}$ in Algorithm 1 contains all the binary vectors of size $\widehat{N}_{miss}^{(h)}(A)$ (that we denote abusively N in this paragraph). This would lead to exponential complexity in the number of calculations of the CART criterion and then to an intractable algorithm. The following remark shows that only a linear (in N) number of assignations has to be considered. Fix a cell A and a cut (h, z) in A , to keep a simple notation we denote by $Y_{(1)} \leq \dots \leq Y_{(N)}$ the ordered values of the response variable such that $\mathbf{M}^{(h)} = 0$. Moreover, denote $\bar{Y}_{L,obs}$ (resp. $\bar{Y}_{R,obs}$) the empirical mean of the non-missing variables belonging to the cell A_L (resp. A_R). Suppose, without loss of generality, that $\bar{Y}_{L,obs} \leq \bar{Y}_{R,obs}$. Since $\bar{Y}_{L,obs} \leq \bar{Y}_{R,obs}$, the only candidates to the maximization of the CART criterion are the one such that there exists a $0 \leq k \leq N$ such that $Y_{(1)}, \dots, Y_{(k)} \in A_L$ and $Y_{(k+1)}, \dots, Y_{(N)} \in A_R$. The case $k = N$ (resp. $k = 0$) holds for an assignation of all the variables to the left (resp. to the right). Indeed, if the assignation w does not belong to those candidates, there exists a couple such that $Y_{(i)} < Y_{(j)}$ and $Y_{(i)} \in A_R, Y_{(j)} \in A_L$. The assignation obtained by the simple transposition of those two variables strictly increases the CART criterion so that those assignations cannot be optimal. Finally, this shows that only $N + 1$ assignation are admissible. The cases where no observed variables are in A_L or in A_R can be treated in the same way.

In particular, note that if we have extra information on the missingness mechanism, the admissible assignations w might be even reduced. For example, assume that the only missing values are in $\mathbf{X}^{(1)}$ and that those values are missing exactly when they are bigger than some threshold τ , that

Algorithm 2 Random forest with assignation of missing entries.

Input: Training sample \mathcal{D}_n , number of trees $M > 1$, $\text{mtry} \in \{1, \dots, p\}$, $a_n \in \{1, \dots, n\}$, $\text{nodesize} \in \{2q_n - 1, \dots, a_n\}$, $q_n \in \{1, \dots, \lfloor a_n/2 \rfloor\}$.

Output: Random forest $m_{M,n}$.

- 1: **for** $i = 1, \dots, M$ **do**
- 2: Select a_n points uniformly in \mathcal{D}_n , denoted (abusively) $(\mathbf{X}_1, Y_1), \dots, (\mathbf{X}_{a_n}, Y_{a_n})$.
- 3: Set $\widehat{\mathbf{X}}_i^{(h)} = [0, 1]$ for each missing variable.
- 4: Let $\mathcal{A} = \{[0, 1]^p\}$, the initialization of the set of **active** cells.
- 5: Let $m_n(\cdot, \Theta_i) = \emptyset$, the initialization of the decision tree.
- 6: **while** $\mathcal{A} \neq \emptyset$ **do**
- 7: Let A be the first element of \mathcal{A} .
- 8: **if** $q_n \leq N(A) \leq \text{nodesize}$ **then**
- 9: Remove A from \mathcal{A} , add mark it as a **final** cell.
- 10: Compute the mean of the points in A , denoted as \widehat{Y}_A , and add $\{A, \widehat{Y}_A\}$ to $m_n(\cdot, \Theta_i)$.
- 11: **else**
- 12: Select uniformly, without replacement, a subset $\mathcal{M}_{\text{try}} \subset \{1, \dots, h\}$ of cardinality mtry .
- 13: Apply Algorithm 1 over the directions in \mathcal{M}_{try} and \mathcal{C}_A .
- 14: Let be $\widehat{d} = (\widehat{h}, \widehat{z})$ the optimal cut selected and \widehat{w} the optimal assignation.
- 15: Compute the mean of the points in A , denoted as \widehat{Y}_A .
- 16: Compute the estimated probabilities $\widehat{p}_L^{(\widehat{h})}$ and $\widehat{p}_R^{(\widehat{h})}$
- 17: Add $\{A, \widehat{Y}_A, \widehat{d}, \widehat{p}_L^{(\widehat{h})}, \widehat{p}_R^{(\widehat{h})}\}$ to $m_n(\cdot, \Theta_i)$.
- 18: Remove A from \mathcal{A} , and add A_L and A_R to \mathcal{A} .
- 19: **end if**
- 20: **end while**
- 21: **end for**
- 22: Set $m_{M,n} = \{m_n(\cdot, \Theta_i)\}_i$.

is, when we have right censoring. Also consider that we know the value of τ , so the missingness mechanism is completely known. So, if we perform a cut at τ in $\mathbf{X}^{(1)}$ then the only admissible assignation for $\widehat{\mathbf{X}}_i$ is to the right child node.

3.2 Hypothesis and Main Theorem

We consider an additive regression model satisfying the following properties,

Hypothesis 1 (H1) *The response variable Y is of the form*

$$Y = \sum_{j=1}^p m_j(\mathbf{X}^{(j)}) + \varepsilon,$$

where \mathbf{X} is uniformly distributed over $[0, 1]^p$, ε is an independent Gaussian centered noise with finite variance $\sigma^2 > 0$ and each component m_j is continuous.

Hypothesis 2 (H2) *The random variables $\mathbf{X}^{(h)}$ are missing by following an MCAR mechanism. The probability of missingness $p_n^{(h)} = \mathbb{P}[\mathbf{M}^{(h)} = 1]$ only depends on the size n of the sample \mathcal{D}_n and $\lim_{n \rightarrow \infty} p_n^{(h)} = c^{(h)}$ where $0 < c^{(h)} < 1$ is constant for all $h \in \{1, \dots, p\}$.*

Theorem 1 *Assume that H1 and H2 hold. Then, under the condition $q_n \rightarrow \infty$, the random forest estimator with missing values is consistent in probability, i.e., for all $\xi, \rho > 0$ there exists $N \in \mathbb{N}^*$, such that for all $n > N$*

$$\mathbb{P}\left[|m_n(\mathbf{X}) - m(\mathbf{X})| \leq \xi\right] \geq 1 - \rho$$

Additive regression models, which decompose the regression function as a sum of univariate functions, are flexible and easy to interpret, providing a good trade-off between model complexity, calculation time and interpretation. We have taken the previous work of Scornet, Biau, and Vert (2015) as inspiration, they proved the consistency of the original random forest algorithm (without missing values) considering the same structure for the regression function and the same assumptions for the input variables and the errors. However, the proof of the results differs from theirs in several parts. Considering other assumptions for the input variables and the errors, like sub-Gaussian errors or distributions with support in $[0, 1]^p$ for \mathbf{X} , would require to adapt the technical parts of the proof with little relevance.

The fact that $q_n \rightarrow \infty$ implies that the number of points selected in each tree a_n , tends to infinity. Hence, in the sequel, we assume that $a_n \rightarrow \infty$ and omit the dependence of the trees over a_n . The condition that $q_n \rightarrow \infty$ is sub-optimal, and we think that with some extra effort the results presented in this work can be adapted to a more classical (and more powerful) condition like $a_n/t_n \rightarrow \infty$, where t_n holds for the number of final cells in the trees. In this case it is possible that not every tree would be consistent but the random forest would still converge to the regression function.

4 Proof of Theorem 1

First of all, we define the theoretical counterpart of the empirical CART criterion by

$$\begin{aligned} L^*(A, d) = & \mathbb{V}[Y|\mathbf{X} \in A] - \mathbb{V}[Y|\mathbf{X}^{(h)} < z, \mathbf{X} \in A]\mathbb{P}[\mathbf{X}^{(h)} < z|\mathbf{X} \in A] \\ & - \mathbb{V}[Y|\mathbf{X}^{(h)} \geq z, \mathbf{X} \in A]\mathbb{P}[\mathbf{X}^{(h)} \geq z|\mathbf{X} \in A]. \end{aligned}$$

Analogously to the empirical case, we define the best theoretical cut $d^* = (h^*, z^*)$ in A as

$$d^* = (h^*, z^*) \in \arg \max_{d \in \mathcal{C}_A} L^*(A, d).$$

By the classical strong law of large numbers, $L_n(A, d)$ converges almost surely to $L^*(A, d)$ as n tends to infinity, for all cuts $d \in \mathcal{C}_A$. This fact leads to the interpretation that the chosen cut at each step of the tree construction tends to decrease the variability of the sets of data points corresponding to the resulting nodes. This implies that the empirical mean in each cell tends to stabilize around a value that correspond to the conditional mean for the cell. Applying basic algebra, it is easy to get an equivalent expression for the empirical CART criterion, given by

$$L_n(A, d) = \frac{N(A_L)N(A_R)}{N(A)N(A)} (\bar{Y}_{A_L} - \bar{Y}_{A_R})^2. \quad (4)$$

From Equation (4), we can get an alternative expression to the theoretical CART criterion, given by

$$\begin{aligned} L^*(A, d) = & \mathbb{P}[\mathbf{X}^{(h)} < z|\mathbf{X} \in A]\mathbb{P}[\mathbf{X}^{(h)} \geq z|\mathbf{X} \in A] \\ & \times \left(\mathbb{E}[Y|\mathbf{X} < z, \mathbf{X} \in A] - \mathbb{E}[Y|\mathbf{X} \geq z, \mathbf{X} \in A] \right)^2. \end{aligned}$$

Remember that, in our procedure, the CART criterion takes “imputed” intervals which are updated after a cut is selected. From a theoretical point of view we do not longer have observations. Hence, we introduce the notions of the input assigned random variable $\widehat{\mathbf{X}}_{in}$ and output assigned random variable $\widehat{\mathbf{X}}_{out}$. The random variable $\widehat{\mathbf{X}}_{in}$ corresponds to a prior assignation whereas $\widehat{\mathbf{X}}_{out}$ corresponds to the assignations after the theoretical cut is performed. Furthermore, the binary assignations w are translated into probabilities.

Formally, let \mathcal{W} be the collection of functions from \mathbb{R} to $[0, 1]$, and $\widehat{\mathbf{X}}_{in} = \left(\widehat{\mathbf{X}}_{in}^{(1)}, \dots, \widehat{\mathbf{X}}_{in}^{(p)} \right)$ be the input distribution of the imputation, then $\widehat{\mathbf{X}}_{out}^{(h)} | \widehat{\mathbf{X}}_{in}$ is defined as

$$\widehat{\mathbf{X}}_{out}^{(h)} | \widehat{\mathbf{X}}_{in} \in A = \begin{cases} \mathbf{X}^{(h)} | \mathbf{X} \in A & \text{if } \mathbf{M}^{(h)} = 0 \\ \mathbf{B}^{(h)} & \text{if } \mathbf{M}^{(h)} = 1 \end{cases},$$

where

$$\mathbf{B}^{(h)} = \begin{cases} (a^{(h)}, z) & \text{if } \text{Ber}(w(Y)) = 1 \\ (z, b^{(h)}) & \text{if } \text{Ber}(w(Y)) = 0 \end{cases}, \quad w \in \mathcal{W}. \quad (5)$$

The imputation variable $w(Y)$ is the (random) probability that $\widehat{\mathbf{X}}_{out}^{(h)} < z$ conditionally to $\mathbf{M}^{(h)} = 1$ and Y . Note that $\widehat{\mathbf{X}}_{in}$ always belongs to a cell A , so the above definition of $\widehat{\mathbf{X}}_{out}^{(h)} | \widehat{\mathbf{X}}_{in} \in A$ is well defined. We define the theoretical CART over a cut $d = (h, z)$ and a function $w \in \mathcal{W}$ as

$$\begin{aligned} L^*(A, d, w) = & \mathbb{V}[Y | \widehat{\mathbf{X}}_{in} \in A] - \mathbb{V}[Y | \widehat{\mathbf{X}}_{out}^{(h)} < z, \widehat{\mathbf{X}}_{in} \in A] \mathbb{P}[\widehat{\mathbf{X}}_{out}^{(h)} < z | \widehat{\mathbf{X}}_{in} \in A] \\ & - \mathbb{V}[Y | \widehat{\mathbf{X}}_{out}^{(h)} \geq z, \widehat{\mathbf{X}}_{in} \in A] \mathbb{P}[\widehat{\mathbf{X}}_{out}^{(h)} \geq z | \widehat{\mathbf{X}}_{in} \in A,] \end{aligned}$$

and the best cut and assignation (d^*, w^*) is selected by maximizing $L^*(A, d, w)$ over $\mathcal{C}_A \times \mathcal{W}$, that is

$$(d^*, w^*) \in \arg \max_{\substack{d \in \mathcal{C}_A \\ w \in \mathcal{W}}} L^*(A, d, w).$$

We define, for any subset $A \subset \mathcal{X}$, the variation of m within A as

$$\Delta(m, A) = \sup_{\mathbf{x}, \mathbf{x}' \in A} |m(\mathbf{x}) - m(\mathbf{x}')|.$$

Furthermore, we denote by $A_{s(n)}(\mathbf{X}, \Theta)$ the final cell of the tree built with the random variable Θ that contains \mathbf{X} , where $s(n)$ is the number of cuts necessary to construct the cell in the tree. After the last step of the tree construction, we end with a collection of imputed values that corresponds to the last imputed sample $(\widehat{\mathbf{X}}_{1,out}, \dots, \widehat{\mathbf{X}}_{n,out})$. Each of these vectors are non ambiguously assigned to a specific final cell of the tree. In all that follows, when we write an imputation $\widehat{\mathbf{X}}$ without any specification of *in* or *out*, we refer to the final imputation. The proof of Theorem 1 relies on Proposition 1 below, which states that the variation of the regression function within a cell of a random forest is small for n large enough and helps us to control the error of our predictor. In the sequel, we will abuse of the notation and will not write the dependence of the cells over \mathbf{X} and Θ .

Proposition 1 *Assume that H1 and H2 hold. Then,*

$$\Delta(m, A_{s(n)}) \rightarrow 0 \quad \text{almost surely.}$$

For all $\mathbf{x} \in [0, 1]^p$, we denote by $L^*(A_{s(n)}, d, w)$ the theoretical CART criterion over the cell $A_{s(n)}$ evaluated at a cut $d \in \mathcal{C}_{A_{s(n)}}$ and assignation $w \in \mathcal{W}$. Let (d^*, w^*) be the optimal couple (cut, assignation) of the cell $A_{s(n)}$ for the theoretical criterion and let $(\widehat{d}, \widehat{w})$ be the optimal couple (cut, assignation) of the cell $A_{s(n)}$ for the empirical criterion.

Lemma 1 Assume that *H1* and *H2* are satisfied and fix $\mathbf{x} \in [0, 1]^p$. Then for all $\rho, \xi > 0$, there exists $N \in \mathbb{N}^*$ such that, for all $n \geq N$ if $\mathbb{P} [L^* (A_{s(n)}, d^*, w^*) \leq \xi] \geq 1 - \rho$, then

$$\Delta(m, A_{s(n)}(\mathbf{x})) \rightarrow 0 \quad \text{almost surely.}$$

Lemma 2 Assume that *H1* and *H2* are satisfied and fix $\mathbf{x} \in [0, 1]^p$. Then for all $\rho, \xi > 0$, there exists $N \in \mathbb{N}^*$ such that, for all $n \geq N$

$$\mathbb{P} \left[L_n \left(A_{s(n)}, \hat{d}, \hat{w} \right) \leq \xi \right] \geq 1 - \rho$$

4.1 Proof of Lemma 1

Technical Lemma 1 Assume that *H1* and *H2* are satisfied and for a cell A , $L^*(A, d, \tilde{w}_{A,d}) \equiv 0$ for all cuts $d = (h, z) \in \mathcal{C}_A$, where

$$\tilde{w}_{A,d} = \mathbb{P} \left[a^{(h)} \leq \mathbf{X}^{(h)} < z \mid \mathbf{X} \in A, \widehat{\mathbf{M}}^{(h)} = 1 \right].$$

Then, m is constant on the cell A .

Proof

Without loss of generality, we will assume that the cut $d = (1, z)$ is performed in the first direction so that $h = 1$ and that the bounds of the cell A are a on the left and b on the right. We omit the direction h and simply note X instead of $\mathbf{X}^{(1)}$. Note that if $\widehat{\mathbf{X}}_{out}$ is assigned to the left child node using $\tilde{w}_{A,d}$, which does not depend on Y , then $\widehat{\mathbf{X}}_{out}$ follows the same distribution than \mathbf{X} and then our theoretical CART criterion is similar to the usual one for every cut $d = (h, z) \in \mathcal{C}_A$, that is,

$$L^*(A, d, \tilde{w}_{A,d}) = \mathbb{P} [a \leq X < z \mid \mathbf{X} \in A] \mathbb{P} [z \leq X \leq b \mid \mathbf{X} \in A] (\mu_{A_L} - \mu_{A_R})^2$$

where the notation μ_{A_L} (resp. μ_{A_R}) holds for the conditional expected value of Y given $\widehat{\mathbf{X}}_{in} \in A$ and $a \leq X < z$ (resp. $z \leq X \leq b$). Because \mathbf{X} is uniformly distributed over $[0, 1]^p$,

$$\mathbb{P} [a \leq X < z \mid \mathbf{X} \in A] = \frac{z - a}{b - a}$$

and

$$\mathbb{P} [z \leq X \leq b \mid \mathbf{X} \in A] = \frac{b - z}{b - a}.$$

Next, we need to understand the distribution of Y conditionally to $\widehat{\mathbf{X}}_{in} \in A$. This random variable is a mixture of the values of Y such that $\mathbf{X} \in A$ and the ones that were assigned to the cell A through the notion of the variable $\widehat{\mathbf{X}}_{in}$. Since the vector $(\widehat{\mathbf{X}}_{in}, Y)$ has a density, we can give a precise meaning to the quantity

$$\tilde{m}(\mathbf{x}) = \mathbb{E} \left[Y \mid \widehat{\mathbf{X}}_{in} = \mathbf{x} \right], \quad (6)$$

for all $\mathbf{x} \in A$. By assumption *H1*, we introduce the notation $Y^{(j)} = m_j(X^{(j)}) + \epsilon^{(j)}$, where $\epsilon^{(j)} \sim \mathcal{N}(0, \sigma^2/p^2)$ so that we have $Y \sim \sum_{j=1}^p Y^{(j)}$. Since under the condition $\mathbf{X} \in A$, the random variable

\mathbf{X} is also uniformly distributed on the cell A , we have that

$$\begin{aligned}\mu_{A_L} &= \mathbb{E} \left[\sum_{j=1}^p Y^{(j)} | a \leq X < z, \widehat{\mathbf{X}}_{in} \in A \right] \\ &= \sum_{j \geq 2} \mathbb{E} \left[Y^{(j)} | \widehat{\mathbf{X}}_{in} \in A \right] + \mathbb{E} \left[Y^{(1)} | a \leq X < z, \widehat{\mathbf{X}}_{in} \in A \right] \\ &= \tilde{K} + \frac{1}{z-a} C_a^z,\end{aligned}$$

where \tilde{K} does not depend on z , $C_x^y = \int_x^y \tilde{m}_1(t) dt$ and

$\tilde{m}_1(t) = \mathbb{E} \left[Y^{(1)} | \widehat{\mathbf{X}}_{in} \in A, \mathbf{X}^{(1)} = t \right]$. Analogously, we have that

$$\mu_{A_R} = \tilde{K} - \frac{1}{b-z} C_a^z + \frac{1}{b-z} C_a^b.$$

Therefore,

$$\begin{aligned}L^*(A, d, \tilde{w}_{A,d}) &= \left(\frac{z-a}{b-a} \right) \left(\frac{b-z}{b-a} \right) \left(\frac{1}{z-a} C_a^z + \frac{1}{b-z} C_a^z - \frac{1}{b-z} C_a^b \right)^2 \\ &= \frac{1}{(z-a)(b-z)} \left(C_a^z - \frac{z-a}{b-a} C_a^b \right)^2.\end{aligned}$$

Since $L^*(A, d, \tilde{w}_{A,d}) = 0$ by assumption, we obtain that for any $a \leq z \leq b$,

$$C_a^z = \frac{z-a}{b-a} C_a^b.$$

This proves that $z \mapsto C_a^z$ is linear in z and thus, \tilde{m}_1 is constant on $[a, b]$. Since the law of Y_1 is a mixture of distribution such that $\mathbf{M}^{(1)} = 0$ and $\mathbf{M}^{(1)} = 1$, we have that

$$\tilde{m}_1(t) = (1 - p^{(1)}) m_1(t) + p^{(1)} \mathbb{E} \left[Y^{(1)} | \widehat{\mathbf{X}}_{in} \in A, \mathbf{X}^{(1)} = t, \mathbf{M}^{(1)} = 1 \right]$$

where the second term does not depend on t (since the value is missing). This forces the function m_1 to be constant on the interval $[a, b]$. And, by additivity, the function m is constant on A .

We need a second technical lemma that states that any cell within a cell for which each split gives a small value of L^* has also a uniformly small value of the CART criterion L^* .

Technical Lemma 2 *Assume constructed the imputation $\widehat{\mathbf{X}}_{in}$. Let $\epsilon > 0$, then there exists $\delta > 0$ such that for all cell A and cut $d \in \mathcal{C}_A$ satisfying $L^*(A, d, \tilde{w}_{A,d}) \leq \delta$, we have*

$$L^*(B, d, \tilde{w}_{A,d}) \leq \epsilon$$

for all cell $B \subset A$, where both notions of CART criteria are constructed with the same vector of imputation $\widehat{\mathbf{X}}_{in}$.

Proof

We first prove the unidimensional case. Without loss of generality, we can assume that the cell A is the interval $[0, 1]$ and we first consider the cell B of the form $[0, a]$ with $0 \leq a \leq 1$. Following the lines of Technical Lemma 1, we have that

$$L^*(A, z, \tilde{w}_{A,d}) = \frac{1}{z(1-z)} (C_0^z - C_0^1 z)^2,$$

where for all x and y , $C_x^y = \int_x^y \tilde{m}(t) dt$ and \tilde{m} is defined as in Equation (6). Evaluating the CART criterion at the value $z = a$, we have

$$\frac{1}{a(1-a)} (C_0^a - C_0^1 a)^2 \leq \delta.$$

For the cell B , at a cut level $0 < z < a$, we have that

$$\begin{aligned} L^*(B, z, \tilde{w}_{A,d}) &= \frac{1}{z(a-z)} \left(C_0^z - C_0^a \frac{z}{a} \right)^2 \\ &= \frac{1}{z(a-z)} \left(C_0^z - C_0^1 z + C_0^1 z - C_0^a \frac{z}{a} \right)^2 \\ &\leq \frac{2}{z(a-z)} (C_0^z - C_0^1 z)^2 + \frac{2z}{a-z} \left(C_0^1 - \frac{C_0^a}{a} \right)^2 \\ &\leq 2 \frac{1-z}{a-z} L^*(A, z, \tilde{w}) + \frac{2z(1-a)}{a(a-z)} \delta \\ &\leq 2 \frac{a(1-z) + z(1-a)}{a(a-z)} \delta \leq \frac{2}{a(a-z)} \delta. \end{aligned}$$

But since the function $z \mapsto C_0^z$ is differentiable, we have that $C_0^z = C_0^a - (a-z)\tilde{m}(a) + o(a-z)$ when $z \rightarrow a$ by a Taylor expansion. This shows that $L^*(B, z, \tilde{w}_{A,d}) \rightarrow 0$ when $z \rightarrow a$ and then there exists a $\delta_0 > 0$ such that $L^*(B, z, \tilde{w}_{A,d}) \leq \epsilon$ when $(a-z) \leq \delta_0$. Then, using the previous inequalities for $0 < z < a - \delta_0$, we have that for $\delta = \epsilon \delta_0 a / 2$, $L^*(B, z, \tilde{w}_{A,d}) \leq \epsilon$. In the same way, we generalize the previous arguments for $A = [0, 1]^p$ and for a specific type of cell

$$B = \{\mathbf{x} \in A : \mathbf{x}^{(1)} \leq a\}$$

with $0 < a < 1$. The values C_y^z are then replaced by $\int_y^z \int_0^1 \dots \int_0^1 \tilde{m}(t) dt_1 dt_2 \dots dt_p$. The result can be repeated for the case $B = \{\mathbf{x} \in A : \mathbf{x}^{(1)} \geq a\}$.

For the general case of $B \subset A$ we see that any B can be obtained by a finite sequence of $B = B_k \subset B_{k-1} \subset \dots \subset B_1 \subset A$ of subset constructed by the scheme described above. This finishes the proof.

Proof of Lemma 1

We will show that $\Delta(m, A_{s(n)}(\mathbf{x})) \rightarrow 0$ a.s. by contradiction. We assume that with positive probability, there exists a positive constant $c > 0$ and a sub-sequence $\phi(n)$ of cells $A_{s(\phi(n))}$ such that $\Delta(m, A_{s(\phi(n))}(\mathbf{x})) > c$. This means that in each set $A_{s(\phi(n))}(\mathbf{x})$ one can find a pair of elements (x_n, y_n) such that

$$|m(x_n) - m(y_n)| > c.$$

The sequences $(x_n)_n$ and $(y_n)_n$ belong to the compact set $[0, 1]^p$ so one can extract a sub-sequence $\psi(n)$ such that $(x_n)_n$ and $(y_n)_n$ converge respectively to two points x and y . By continuity of m , we have that

$$|m(x) - m(y)| > c.$$

At the cost of taking an x' and y' close to x and y , satisfying

$$|m(x') - m(y')| > \frac{c}{2},$$

and such that, for n large enough, all the cells $A_{s(\phi(n))}$ contain the pair of points x' and y' . By hypothesis, we know that

$$\sup_{d \in \mathcal{C}_{A_{s(\phi \circ \psi(n))}}} L^*(A_{s(\phi \circ \psi(n))}, d, \tilde{w}) \rightarrow 0 \quad \text{in probability,}$$

where we just wrote \tilde{w} for the choice of the assignation given in Technical Lemma 1. So one can extract a subsequence $\chi(n)$ such that the $\sup_d L^*(A_{s(\phi \circ \psi \circ \chi(n))}, d, \tilde{w})$ converges to 0 almost surely. For simplicity, we keep denoting n for the sub-sequence $\phi \circ \psi \circ \chi(n)$ in the following part of the proof. Lastly, we define the sequence of cells $(C_i)_{i \geq 1}$ such that

$$C_i = \bigcap_{k=1}^i A_{s(k)}.$$

These cells form a non-increasing sequence for the inclusion order and for each i , $C_i \subset A_{s(i)}$. The cells C_i inherit the same vector of imputation as for $A_{s(i)}$. We can use Technical Lemma 2 to get that $\sup_{d \in \mathcal{C}_{A_{s(n)}}} L^*(C_n, d, \tilde{w}) \rightarrow 0$ a.s. Since the sequence of cells $(C_i)_i$ is a non-increasing sequence, there exists a cell, denoted C_∞ that is the limit of the cells C_i when $i \rightarrow \infty$ in the sense

$$C_\infty = \bigcap_{i \geq 1} C_i.$$

To see that, one can write $C_i = \prod_h [a_i^{(h)}, b_i^{(h)}]$ and take the limits of the sequences $(a_i^{(h)})_i$ and $(b_i^{(h)})_i$. The objective is to show that there is no other possibility than having $L^*(C_\infty, d, \tilde{w}) = 0$ for every cut d . For a cell A , the CART criterion $L^*(A, d, \tilde{w}_{A,d})$ has a continuous behavior with respect to A . Indeed, we see that

$$L^*(A, d, \tilde{w}_{A,d}) = \mathbb{P} \left[a^{(h)} \leq \mathbf{X}^{(h)} < z | \mathbf{X} \in A \right] \mathbb{P} \left[z \leq \mathbf{X}^{(h)} \leq b^{(h)} | \mathbf{X} \in A \right] (\hat{\mu}_{AL} - \hat{\mu}_{AR})^2,$$

which is a product of three terms that are uniformly continuous in A (since m is a continuous function) with respect to the distance between cells given by $\Delta(A, B) = \max_h |x_A^{(h)} - x_B^{(h)}| + \max_h |y_A^{(h)} - y_B^{(h)}|$ where $x_A^{(h)}$ and $y_A^{(h)}$ are defined as $A = \bigcap_h [x_A^{(h)}, y_A^{(h)}]$. Since $C_i \rightarrow C_\infty$ for the distance Δ , we have that for all $\epsilon > 0$, for all i large enough, for all cut d of the final cell C_∞ ,

$$|L^*(C_i, d, \tilde{w}) - L^*(C_\infty, d, \tilde{w})| \leq \epsilon.$$

But since the $L^*(C_i, d, \tilde{w})$ converges almost surely uniformly in d to 0, and ϵ is arbitrary, for every cut d inside the valid cuts of C_∞ , we have that $L^*(C_\infty, d, \tilde{w}) = 0$. But then by Technical Lemma 1, the function m has to be constant in the cell C_∞ . But the points x' and y' do belong to the cell C_∞ since they belong to each of the cells in the intersection. So $m(x') = m(y')$ which contradicts the fact that $|m(x') - m(y')| > c/2$. This proves that

$$\Delta(m, A_{s(n)}(\mathbf{x})) \rightarrow 0 \quad \text{almost surely.}$$

4.2 Proof of Lemma 2

Remember that $A_{s(n)}$ denotes the final cell of the tree built with the random variable Θ that contains \mathbf{X} , where $s(n)$ is the number of cuts necessary to construct the cell. Similarly, A_k is the same cell but where only the first k cuts have been performed.

Technical Lemma 3 *Assume that H1 and H2 hold and fix $\mathbf{x} \in [0, 1]^p$. For all $\rho, \xi > 0$ there exists $N \in \mathbb{N}^*$ such that for all $n \geq N$ there exists $k_0(n) \in \mathbb{N}^*$ such that*

$$\mathbb{P} \left[\left| L_n \left(A_{s(n)}, \hat{d}, \hat{w} \right) - L_n \left(A_k, \hat{d}, \hat{w} \right) \right| \leq \xi \right] \geq 1 - \rho,$$

for all $k \geq k_0(n)$.

Proof

Fix $\alpha, \rho > 0$ and consider the following standard inequality on a Gaussian tail

$$\mathbb{P}[\varepsilon_1 \geq \alpha] \leq \frac{\sigma^2}{t\sqrt{2\pi}} \exp\left(-\frac{t^2}{2\sigma^2}\right)$$

then, simple calculations show that, for all $n \in \mathbb{N}^*$

$$\mathbb{P} \left\{ \left| \sum_{i=1}^n \varepsilon_i \right| \geq n\alpha \right\} \leq \frac{\sigma}{\alpha\sqrt{n}} \exp\left\{-\frac{\alpha^2 n}{2\sigma^2}\right\}. \quad (7)$$

Note that there are at most $n(n+1)/2$ sets of the form $\{i : \mathbf{X}_i^{(h)} \in [a_n, b_n], \mathbf{M}_i^{(h)} = 0\}$ for $0 \leq a_n < b_n \leq 1$. On the other hand, let $(Y_{(1)}, \dots, Y_{(n)})$ be the ordered vector of Y , since missing observations are assigned to the cell using Y and maximizing the CART criterion, this implies that close values of Y must be assigned to the same cell, thus once again note that there are at most $n(n+1)/2$ sets of the form $\{i : Y_{(i)} \in [a_n, b_n], \mathbf{M}_i^{(h)} = 1\}$ for $0 \leq a_n < b_n \leq 1$.

We deduce from Equation (7) and the union bound, that there exists $N_1 \in \mathbb{N}^*$ such that, with probability at least $1 - \rho$, for all $n \geq N_1$ and all $0 \leq a_n < b_n \leq 1$ satisfying $\hat{N} \left(\prod_{h=1}^p [a_n^{(h)}, b_n^{(h)}] \right) \geq q_n$,

$$\left| \frac{1}{\hat{N} \left(\prod_{h=1}^p [a_n^{(h)}, b_n^{(h)}] \right)} \sum_{i=1}^n \varepsilon_i \mathbb{1}_{\hat{\mathbf{x}}_i \in A} \right| \leq \frac{\sigma n^{4p}}{\alpha\sqrt{n}} \exp\left\{-\frac{\alpha^2 n}{2\sigma^2}\right\} \leq \alpha. \quad (8)$$

Furthermore, making use of the same ideas and applying the inequality $\mathbb{P}[\chi^2(n) \geq 5n] \leq e^{-n}$ (for interested readers, see Laurent and Massart (2000)), there exists $N_2 \in \mathbb{N}^*$ such that, with probability at least $1 - \rho$ for all $n \geq N_2$ and all $0 \leq a_n < b_n \leq 1$ satisfying $\hat{N} \left(\prod_{h=1}^p [a_n^{(h)}, b_n^{(h)}] \right) \geq q_n$

$$\frac{1}{\hat{N} \left(\prod_{h=1}^p [a_n^{(h)}, b_n^{(h)}] \right)} \sum_{i=1}^n \varepsilon_i^2 \mathbb{1}_{\hat{\mathbf{x}}_i \in A} \leq \tilde{\sigma}^2, \quad (9)$$

where $\tilde{\sigma}^2$ is a positive constant, depending only on ρ . Since $(A_k)_k$ is a decreasing sequence of compact sets, for every $\xi > 0$ there exists k_0 such that, for all $k \geq k_0$

$$\max \left(\|\mathbf{a}_k - \mathbf{a}_{s(n)}\|_\infty, \|\mathbf{b}_k - \mathbf{b}_{s(n)}\|_\infty \right) \leq \xi, \quad (10)$$

where $\mathbf{a}_k = (a_k^{(1)}, \dots, a_k^{(p)}) \in [0, 1]^p$ and $\mathbf{b}_k = (b_k^{(1)}, \dots, b_k^{(p)}) \in [0, 1]^p$ such that

$$A_k = \prod_{j=1}^p [a_k^{(j)}, b_k^{(j)}].$$

We define $\mathbf{a}_{s(n)}$ and $\mathbf{b}_{s(n)}$ analogously for the cell $A_{s(n)}$. Since the variable \mathbf{X} is uniformly distributed in the hypercube $[0, 1]^p$ and the missing entries follow the MCAR mechanism we have that, for any cell A , $\widehat{N}_{obs}(A)/\widehat{N}_{obs}([0, 1]^p) \rightarrow \text{Vol}(A)$ almost surely as $n \rightarrow \infty$. Furthermore, for any $k \geq k_0(n)$, Equation (10) implies that $\text{Vol}(A_k) \leq (1 + 2\xi)^p \text{Vol}(A_{s(n)})$. Then there exists $N_3 \in \mathbb{N}^*$ such that for all $n \geq N_3$, $\widehat{N}_{obs}(A_k) \leq (1 + 3\xi)^p \widehat{N}_{obs}(A_{s(n)})$, and

$$\widehat{N}_{obs}(A_k \setminus A_{s(n)}) \leq \xi' \widehat{N}_{obs}(A_{s(n)}) \leq \xi' \widehat{N}_{obs}(A_k), \quad (11)$$

where $\xi' = (1 + 3\xi)^p - 1$. On the other hand, for any fixed n , the quantity $u_k = \widehat{N}_{miss}(A_k) - \widehat{N}_{miss}(A_{s(n)})$ converges to 0 almost surely as $k \rightarrow \infty$. Now taking $n \geq N_3$ fixed, we have that there exists a $k_1(n)$ such that for all $k \geq k_1(n)$, with probability at least $1 - \rho$, we have

$$\begin{aligned} \widehat{N}(A_k \setminus A_{s(n)}) &= (\widehat{N}_{obs}(A_k) + \widehat{N}_{miss}(A_k)) - (\widehat{N}_{obs}(A_{s(n)}) + \widehat{N}_{miss}(A_{s(n)})) \\ &= \widehat{N}_{obs}(A_k \setminus A_{s(n)}) + u_k \\ &\leq 2\xi' \widehat{N}_{obs}(A_{s(n)}), \end{aligned} \quad (12)$$

where we used Equation (11) and the fact that convergence almost sure implies convergence in probability. For the rest of the proof, we take $k \geq \max\{k_0(n), k_1(n)\}$ and assume that Equations (8),(9) and (12) are satisfied, which occurs with probability at least $1 - 3\rho$ for every $n > N$ with $N = \max\{N_1, N_2, N_3\}$. Note that $A_k \setminus A_{s(n)}$ is either a final node, contains at least one final node or is empty (in which case the result holds trivially). Since each final node contains at least q_n points then for q_n sufficiently large, using equation (8) we conclude that $|\widehat{Y}_{A_k}| \leq \|m\|_\infty + \alpha$, $|\widehat{Y}_{A_{s(n)}}| \leq \|m\|_\infty + \alpha$ and $|\widehat{Y}_{A_k \setminus A_{s(n)}}| \leq \|m\|_\infty + \alpha$. We use the following decomposition

$$|L_n(A_k, \widehat{d}, \widehat{w}) - L_n(A_{s(n)}, \widehat{d}, \widehat{w})| \leq K_0 + K_L + K_R,$$

where the three terms K_0 , K_L and K_R are given by

$$K_0 = \left| \frac{1}{\widehat{N}(A_k)} \sum_{i=1}^n (Y_i - \widehat{Y}_{A_k})^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_k} - \frac{1}{\widehat{N}(A_{s(n)})} \sum_{i=1}^n (Y_i - \widehat{Y}_{A_{s(n)}})^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}} \right|,$$

$$\begin{aligned} K_L &= \left| \frac{1}{\widehat{N}(A_k)} \sum_{i=1}^n (Y_i - \widehat{Y}_{A_{L,k}})^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_k, \widehat{\mathbf{x}}_i^{(\widehat{h})} < \widehat{z}} \right. \\ &\quad \left. - \frac{1}{\widehat{N}(A_{s(n)})} \sum_{i=1}^n (Y_i - \widehat{Y}_{A_{L,s(n)}})^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}, \widehat{\mathbf{x}}_i^{(\widehat{h})} < \widehat{z}} \right| \end{aligned}$$

and

$$\begin{aligned} K_R &= \left| \frac{1}{\widehat{N}(A_k)} \sum_{i=1}^n (Y_i - \widehat{Y}_{A_{R,k}})^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_k, \widehat{\mathbf{x}}_i^{(\widehat{h})} \geq \widehat{z}} \right. \\ &\quad \left. - \frac{1}{\widehat{N}(A_{s(n)})} \sum_{i=1}^n (Y_i - \widehat{Y}_{A_{R,s(n)}})^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}, \widehat{\mathbf{x}}_i^{(\widehat{h})} \geq \widehat{z}} \right|. \end{aligned}$$

We first consider the term K_0 that can be upper bounded once again by using a similar split in $K_0 \leq K_{0,1} + K_{0,2} + K_{0,3}$ where

$$K_{0,1} = \left| \frac{1}{\widehat{N}(A_k)} \sum_{i=1}^n (Y_i - \widehat{Y}_{A_k})^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}} - \frac{1}{\widehat{N}(A_k)} \sum_{i=1}^n (Y_i - \widehat{Y}_{A_{s(n)}})^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}} \right|,$$

$$K_{0,2} = \left| \frac{1}{\widehat{N}(A_k)} \sum_{i=1}^n (Y_i - \widehat{Y}_{A_{s(n)}})^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}} - \frac{1}{\widehat{N}(A_{s(n)})} \sum_{i=1}^n (Y_i - \widehat{Y}_{A_{s(n)}})^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}} \right|$$

and

$$K_{0,3} = \left| \frac{1}{\widehat{N}(A_k)} \sum_{i=1}^n (Y_i - \widehat{Y}_{A_k})^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_k \setminus A_{s(n)}} \right|.$$

For $K_{0,1}$, observe that

$$\begin{aligned} |\widehat{Y}_{A_k} - \widehat{Y}_{A_{s(n)}}| &= \left| \frac{1}{\widehat{N}(A_k)} \sum_{i=1}^n Y_i \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_k \setminus A_{s(n)}} + \frac{1}{\widehat{N}(A_k)} \sum_{i=1}^n Y_i \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}} - \widehat{Y}_{A_{s(n)}} \right| \\ &\leq \frac{\widehat{N}(A_k \setminus A_{s(n)})}{\widehat{N}(A_k)} |\widehat{Y}_{A_k \setminus A_{s(n)}} - \widehat{Y}_{A_{s(n)}}| \\ &\leq 4\xi'(\|m\|_\infty + \alpha). \end{aligned}$$

Hence,

$$\begin{aligned} K_{0,1} &\leq \frac{2}{\widehat{N}(A_k)} |\widehat{Y}_{A_{s(n)}} - \widehat{Y}_{A_k}| \left| \sum_{i=1}^n \left(Y_i + \frac{\widehat{Y}_{A_{s(n)}} + \widehat{Y}_{A_k}}{2} \right) \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}} \right| \\ &\leq \frac{8\xi'(\|m\|_\infty + \alpha)}{\widehat{N}(A_k)} \left[\left| \widehat{N}(A_{s(n)}) \widehat{Y}_{A_{s(n)}} \right| + \left| \frac{\widehat{Y}_{A_{s(n)}} + \widehat{Y}_{A_k}}{2} \widehat{N}(A_{s(n)}) \right| \right] \\ &\leq 16\xi'(\|m\|_\infty + \alpha)^2. \end{aligned}$$

For the term $K_{0,2}$, with the help of Equation (9) observe that

$$\begin{aligned} K_{0,2} &\leq \frac{\widehat{N}(A_k \setminus A_{s(n)})}{\widehat{N}(A_k)} \left| \frac{1}{\widehat{N}(A_{s(n)})} \sum_{i=1}^n (Y_i - \widehat{Y}_{A_{s(n)}})^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}} \right| \\ &\leq 2\xi' \left| \frac{1}{\widehat{N}(A_{s(n)})} \sum_{i=1}^n Y_i^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}} + \widehat{Y}_{A_{s(n)}}^2 \right| \\ &\leq 2\xi' \left[(\|m\|_\infty + \alpha)^2 + \frac{1}{\widehat{N}(A_{s(n)})} \sum_{i=1}^n m^2(\mathbf{X}_i) \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}} \right. \\ &\quad \left. + \left| \frac{2}{\widehat{N}(A_{s(n)})} \sum_{i=1}^n m(\mathbf{X}_i) \varepsilon_i \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}} \right| + \frac{1}{\widehat{N}(A_{s(n)})} \sum_{i=1}^n \varepsilon_i^2 \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}} \right] \\ &\leq 2\xi' [(\|m\|_\infty + \alpha)^2 + \|m\|_\infty^2 + 2\|m\|_\infty \alpha + \tilde{\sigma}^2]. \end{aligned}$$

Regarding $K_{0,3}$, observe that

$$\begin{aligned}
K_{0,3} &\leq 2\xi' \left| \frac{1}{\widehat{N}(A_k \setminus A_{s(n)})} \sum_{i=1}^n (Y_i^2 + 2Y_i \widehat{Y}_{A_k} + \widehat{Y}_{A_k}^2) \mathbb{1}_{\widehat{\mathbf{X}}_i \in A_k \setminus A_{s(n)}} \right| \\
&\leq 2\xi' \left[\left| \frac{1}{\widehat{N}(A_k \setminus A_{s(n)})} \sum_{i=1}^n (m(\mathbf{X}_i) + \varepsilon_i)^2 \mathbb{1}_{\widehat{\mathbf{X}}_i \in A_k \setminus A_{s(n)}} \right| \right. \\
&\quad + \frac{2}{\widehat{N}(A_k \setminus A_{s(n)})} (\|m\|_\infty + \alpha) \left| \sum_{i=1}^n (m(\mathbf{X}_i) + \varepsilon_i) \mathbb{1}_{\widehat{\mathbf{X}}_i \in A_k \setminus A_{s(n)}} \right| \\
&\quad \left. + (\|m\|_\infty + \alpha)^2 \right] \\
&\leq 2\xi' \left[\|m\|_\infty^2 + 2\|m\|_\infty \left| \frac{1}{\widehat{N}(A_k \setminus A_{s(n)})} \sum_{i=1}^n \varepsilon_i \mathbb{1}_{\widehat{\mathbf{X}}_i \in A_k \setminus A_{s(n)}} \right| \right. \\
&\quad \left. + \frac{1}{\widehat{N}(A_k \setminus A_{s(n)})} \sum_{i=1}^n \varepsilon_i^2 \mathbb{1}_{\widehat{\mathbf{X}}_i \in A_k \setminus A_{s(n)}} + 2(\|m\|_\infty + \alpha)^2 + (\|m\|_\infty + \alpha)^2 \right] \\
&\leq 2\xi' [3(\|m\|_\infty + \alpha)^2 + \|m\|_\infty^2 + 2\|m\|_\infty \alpha + \tilde{\sigma}^2].
\end{aligned}$$

Therefore, there exists a universal constant $C > 0$ such that $K_0 \leq C\xi$ and with similar arguments we can show that $K_L \leq C\xi$ and $K_R \leq C\xi$. Which concludes that with probability at least $1 - 3\rho$,

$$|L_n(A_k, \widehat{d}, \widehat{w}) - L_n(A_{s(n)}, \widehat{d}, \widehat{w})| \leq 3C\xi.$$

Proof of Lemma 2

Assume that H1 and H2 are satisfied, fix $\mathbf{x} \in [0, 1]^p$ and Θ . Then, let us show by contradiction that for all $\xi > 0$, there exists $N \in \mathbb{N}^*$ such that, with probability at least $1 - \rho$ for all $n \geq N$

$$L_n(A_{s(n)}, \widehat{d}, \widehat{w}) \leq \xi.$$

So, assume that there exists $c > 0$, $0 < p_0 < 1$ and a sub-sequence $\phi(n)$ such that

$$L_{\phi(n)}(A_{s(\phi(n))}, \widehat{d}_{s(\phi(n))}, \widehat{w}_{s(\phi(n))}) > c,$$

with probability at least p_0 . To keep the notation simple, we omit to write $\phi(n)$ and still write n for the indexes of the sub-sequence. Additionally, assume that k is sufficiently large so the conclusion of Technical Lemma 3 is satisfied, hence

$$|L_n(A_k, \widehat{d}, \widehat{w}) - L_n(A_{s(n)}, \widehat{d}, \widehat{w})| \leq \xi \tag{13}$$

and Equation (10) is satisfied. Note that all the feasible cuts d in A_k must be performed in $A_k \setminus A_{s(n)}$, otherwise d would split $A_{s(n)}$ and $(A_k)_k$ would not converge to $A_{s(n)}$ (see Figure 4 for an illustration in $p = 2$). From here we conclude that

$$L_n(A_k, \widehat{d}, \widehat{w}) \leq \sup_{\substack{d \in \mathcal{C}_{A_k} \cap \mathcal{C}_{A_{s(n)}} \\ w \in \mathcal{W}_{A_k}}} L_n(A_k, d, w) \leq \sup_{\substack{d \in \mathcal{C}_{A_k} \\ w \in \mathcal{W}_{A_k}}} L_n(A_k, d, w) \leq \xi.$$

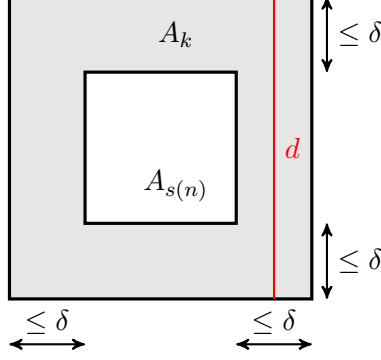


Figure 4: All feasible cuts in A_k must be performed in the $A_k \setminus A_{s(n)}$, like the cut d in the figure, otherwise the cut would split $A_{s(n)}$.

On the other hand, from Equation (13), with probability at least p_0 , we have

$$c - \xi \leq L_n \left(A_{s(n)}, \hat{d}, \hat{w} \right) - \xi^2 \leq L_n \left(A_k, \hat{d}, \hat{w} \right).$$

Hence, we have, with probability at least p_0 ,

$$c - C\xi \leq L_n \left(A_k, \hat{d}, \hat{w} \right) \leq \xi$$

which is absurd, since $c > 0$ is fixed and ξ is arbitrarily small. Thus the result follows.

4.3 Proof of Proposition 1

For a cell A , fix a cut $d \in \mathcal{C}_A$ and consider a function $w \in \mathcal{W}$, we need to define $L_n(A, d, w)$. This is done according to the following procedure, first create a random vector W of dimension $\hat{N}_{miss}^{(h)}(A) = \text{Card}(\mathbf{i}_{A,miss}^{(h)})$, where $W_k = \text{Ber}(w(Y_{j_k}))$ for $j_k \in \mathbf{i}_{A,miss}^{(h)}$, then assign the observations $\hat{\mathbf{X}}_{j_k}$ to the child nodes according to the random vector W . Once we have assigned the observations to the child nodes, we evaluate the empirical CART criterion L_n considering these assignments. Note that in this case L_n is a random variable and the assignments are independent to each other so $L_n(A, d, w)$ is a sum of independent random variables with the same distribution. Hence, by the strong law of large numbers $L_n(A, d, w) \rightarrow L^*(A, d, w)$ almost surely as $n \rightarrow \infty$ for all cuts $d \in \mathcal{C}_A$ and all functions $w \in \mathcal{W}$.

We prove the almost sure convergence of $\Delta(m, A_{s(n)})$ towards 0 by showing that the theoretical CART criterion of the sequence $(A_{s(n)})_n$ tends to 0 and use of Lemmas 1 and 2. Note that

$$\begin{aligned} & L^* \left(A_{s(n)}, d^*, w^* \right) - L_n \left(A_{s(n)}, \hat{d}, \hat{w} \right) \\ &= L^* \left(A_{s(n)}, d^*, w^* \right) - L_n \left(A_{s(n)}, d^*, w^* \right) \\ &\quad + L_n \left(A_{s(n)}, d^*, w^* \right) - L_n \left(A_{s(n)}, \hat{d}, \hat{w} \right) \\ &\leq L^* \left(A_{s(n)}, d^*, w^* \right) - L_n \left(A_{s(n)}, d^*, w^* \right). \end{aligned}$$

Where the last inequality comes from noting that $L_n \left(A_{s(n)}, \hat{d}, \hat{w} \right) \geq L_n \left(A_{s(n)}, d, w \right)$ for all cut $d \in \mathcal{C}_{A_{s(n)}}$ and assignment $w \in \mathcal{W}_{A_{s(n)}}^{(\hat{h})}$, where $\hat{d} = (\hat{h}, \hat{z})$. As discussed above, by strong law of large

numbers

$$L^* (A_{s(n)}, d^*, w^*) - L_n (A_{s(n)}, d^*, w^*) \rightarrow 0,$$

almost surely. Fix $\xi, \rho > 0$, for n sufficiently large, we have

$$L^* (A_{s(n)}, d^*, w^*) - L_n (A_{s(n)}, \widehat{d}, \widehat{w}) \leq \xi \quad \text{almost surely.}$$

On the other hand, by Lemma 2, there exists N_1 such that for all $n \geq N_1$, with probability at least $1 - \rho$

$$L_n (A_{s(n)}, \widehat{d}, \widehat{w}) \leq C\xi.$$

Hence, with the same probability,

$$L^* (A_{s(n)}, d^*, w^*) \leq \xi.$$

And by Lemma 1, we conclude that

$$\Delta(m, A_{s(n)}) \xrightarrow{a.s.} 0.$$

4.4 Proof of Theorem 1

Let

$$m_n(\mathbf{X}) = \frac{1}{\widehat{N}(A_{s(n)}(\mathbf{X}))} \sum_{i=1}^n Y_i \mathbb{1}_{\widehat{\mathbf{x}}_i \in A_{s(n)}(\mathbf{X})}$$

be our tree estimator. Moreover, let us define two other estimators, the first one takes our partition of the cells $A_{s(n)}$ built up using the imputed variables but considers the local means using the complete (unseen) observations \mathbf{X}_i , denoted as

$$m'_n(\mathbf{X}) = \frac{1}{N(A_{s(n)}(\mathbf{X}))} \sum_{i=1}^n Y_i \mathbb{1}_{\mathbf{x}_i \in A_{s(n)}(\mathbf{X})},$$

while the second one takes $m(\mathbf{X}_i)$ for the prediction and the complete observations \mathbf{X}_i , denoted as

$$m''_n(\mathbf{X}) = \frac{1}{N(A_{s(n)}(\mathbf{X}))} \sum_{i=1}^n m(\mathbf{X}_i) \mathbb{1}_{\mathbf{x}_i \in A_{s(n)}(\mathbf{X})}.$$

By Equation (8), we know that for all $\alpha, \xi > 0$ there exists $N \in \mathbb{N}^*$, such that for all $n \geq N$,

$$\begin{aligned} & \mathbb{P} [|m'_n(\mathbf{X}) - m''_n(\mathbf{X})| \geq \alpha] \\ &= \mathbb{P} \left[\left| \frac{1}{N(A_{s(n)}(\mathbf{X}))} \sum_{i=1}^n (Y_i - m(\mathbf{X}_i)) \mathbb{1}_{\mathbf{x}_i \in A_{s(n)}(\mathbf{X})} \right| \geq \alpha \right] \\ &= \mathbb{P} \left[\left| \frac{1}{N(A_{s(n)}(\mathbf{X}))} \sum_{i=1}^n \varepsilon_i \mathbb{1}_{\mathbf{x}_i \in A_{s(n)}(\mathbf{X})} \right| \right] \\ &\leq \xi. \end{aligned}$$

On the other hand, note that $m''_n(\mathbf{X}) = \sum_{i=1}^n W_{n,i}(\mathbf{X}) m(\mathbf{X}_i)$, where

$$W_{n,i}(\mathbf{X}) = \frac{1}{N(A_{s(n)}(\mathbf{X}))} \mathbb{1}_{\mathbf{x}_i \in A_{s(n)}(\mathbf{X})}.$$

Then,

$$\begin{aligned}
\mathbb{E}[m_n''(\mathbf{X}) - m(\mathbf{X})]^2 &= \mathbb{E} \left[\sum_{i=1}^n W_{n,i}(\mathbf{X}) m(\mathbf{X}_i) - m(\mathbf{X}) \right]^2 \\
&= \mathbb{E} \left[\sum_{i=1}^n \sqrt{W_{n,i}(\mathbf{X})} \sqrt{W_{n,i}(\mathbf{X})} (m(\mathbf{X}_i) - m(\mathbf{X})) \right]^2 \\
&\quad (\text{Applying Cauchy-Schwartz's inequality}) \\
&\leq \mathbb{E} \left[\sum_{i=1}^n W_{n,i}(\mathbf{X}) \sum_{i=1}^n W_{n,i}(\mathbf{X}) (m(\mathbf{X}_i) - m(\mathbf{X}))^2 \right] \\
&= \mathbb{E} \left[\sum_{i=1}^n W_{n,i}(\mathbf{X}) (m(\mathbf{X}_i) - m(\mathbf{X}))^2 \mathbb{1}_{\mathbf{X}_i \in A_{s(n)}(\mathbf{X})} \right].
\end{aligned}$$

Note that $(m(\mathbf{X}_i) - m(\mathbf{X}))^2 \mathbb{1}_{\mathbf{X}_i \in A_{s(n)}(\mathbf{X})} \leq \Delta(m, A_{s(n)}(\mathbf{X}))^2$, hence

$$\mathbb{E}[m_n''(\mathbf{X}) - m(\mathbf{X})]^2 \leq \mathbb{E} [\Delta(m, A_{s(n)}(\mathbf{X}))^2].$$

Since $\Delta(m, A_{s(n)}(\mathbf{X})) \leq \Delta(m, [0, 1]^p) < \infty$, we can use the dominated convergence theorem, and conclude that, using Proposition 1,

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

Hence, we have the consistency $m_n'(\mathbf{X}) \xrightarrow{\mathbb{P}} m(\mathbf{X})$. This means that a (fictive) estimator built upon the empirical partition but where all the values are observed to compute the empirical mean step is consistent. We will use this fact to show $m_n(\mathbf{X}) \xrightarrow{\mathbb{P}} m(\mathbf{X})$.

First, consider the case in dimension 1. We use the specific case where a cut d and an assignment w (of the cell $A_{s(n)}$) leaves all the observed points to the left and assigns all the missing observations to the right. By Lemma 2, we know that for n sufficiently large $L_n(A_{s(n)}, d, w) \leq \xi$. As already seen in Equation (4),

$$L_n(A_{s(n)}, d, w) = \frac{\widehat{N}_{obs}(A_{s(n)}) \widehat{N}_{miss}(A_{s(n)})}{\widehat{N}(A_{s(n)}) \widehat{N}(A_{s(n)})} \left(\widehat{Y}_{obs} - \widehat{Y}_{miss} \right)^2.$$

By the convergence of $m_n'(\mathbf{X})$, we have, in particular that $\widehat{Y}_{obs} \xrightarrow{\mathbb{P}} m(\mathbf{X})$.

Using Hypothesis 2 and the same ideas as in Lemma 2, since $p_n^{(h)} \rightarrow c^{(h)} < 1$, there exists $N \in \mathbb{N}^*$ such that, with probability at least $1 - \rho$ for all $n \geq N$, $\widehat{N}_{obs}(A_{s(n)}) / \widehat{N}(A_{s(n)}) \geq c$, where $c > 0$ is a constant. On the other hand, if $\widehat{N}_{miss}(A_{s(n)}) / \widehat{N}(A_{s(n)}) \xrightarrow{\mathbb{P}} 0$ then, trivially $m_n(\mathbf{X}) \xrightarrow{\mathbb{P}} m_n'(\mathbf{X})$ so let us consider the case $\widehat{N}_{miss}(A_{s(n)}) / \widehat{N}(A_{s(n)}) \geq c'$, hence with probability at least $1 - \rho$,

$$\left(\widehat{Y}_{obs} - \widehat{Y}_{miss} \right)^2 \leq \frac{\xi}{cc'}.$$

This shows that the random variable $\widehat{Y}_{obs} - \widehat{Y}_{miss}$ converges to 0 in probability. Since $\widehat{Y}_{obs} \xrightarrow{\mathbb{P}} m(\mathbf{X})$, we obtain that $\widehat{Y}_{miss} \xrightarrow{\mathbb{P}} m(\mathbf{X})$. Finally, using the following formula for $m_n(\mathbf{X})$

$$m_n(\mathbf{X}) = \frac{\widehat{N}_{obs}(A_{s(n)}(\mathbf{X}))}{\widehat{N}(A_{s(n)}(\mathbf{X}))} \widehat{Y}_{obs} + \frac{\widehat{N}_{miss}(A_{s(n)}(\mathbf{X}))}{\widehat{N}(A_{s(n)}(\mathbf{X}))} \widehat{Y}_{miss},$$

we conclude that $m_n(\mathbf{X}) \xrightarrow{\mathbb{P}} m(\mathbf{X})$. For dimension bigger than 1, denote again $Y^{(j)} = m_j(X^{(j)}) + \epsilon^{(j)}$, where $\epsilon^{(j)} \sim \mathcal{N}(0, \sigma^2/p^2)$ so that we have $Y \sim \sum_{j=1}^p Y^{(j)}$ and define

$$m_n^{(j)}(\mathbf{X}) = \frac{1}{\widehat{N}(A_{s(n)}(\mathbf{X}))} \sum_{i=1}^n Y_i^{(j)} \mathbb{1}_{\widehat{\mathbf{X}}_i \in A_{s(n)}(\mathbf{X})}.$$

Consider the cut in the direction 1 which leaves all the observations where $\mathbf{M}^{(1)} = 0$ to the left and assigns the observations where $\mathbf{M}^{(1)} = 1$ to the right. We denote $\widehat{Y}_{obs(1)}$ and $\widehat{Y}_{miss(1)}$ the respective (on the left and on the right) empirical means. By the arguments in dimension 1, we have that $\widehat{Y}_{obs(1)} - \widehat{Y}_{miss(1)} \rightarrow 0$ in probability. By definition,

$$\widehat{Y}_{obs(1)} - \widehat{Y}_{miss(1)} = \widehat{Y}_{obs}^{(1)} - \widehat{Y}_{miss}^{(1)} + \left(\sum_{j=2}^p \widehat{Y}_{obs(1)}^{(j)} - \sum_{j=2}^p \widehat{Y}_{miss(1)}^{(j)} \right).$$

Since the random variable $Y^{(j)}$ ($j \neq 1$) is independent of the random variable $\widehat{\mathbf{X}}^{(1)}$ conditionally to $\widehat{\mathbf{X}} \in A_{s(n)}$, the distributions of the two sums on the right hand side are equal. Since each random variable $\widehat{Y}^{(j)}$ converges (see Technical Lemma 1) we conclude that the difference of the two sums converges in probability to 0. Hence, $\widehat{Y}_{obs}^{(1)} - \widehat{Y}_{miss}^{(1)} \rightarrow 0$ in probability. This finally shows that $m_n^{(1)}(\mathbf{X}) \xrightarrow{\mathbb{P}} m_1(\mathbf{X})$. Similarly, we show that for all $j \geq 1$, $m_n^{(j)}(\mathbf{X}) \xrightarrow{\mathbb{P}} m_j(\mathbf{X})$ and then $m_n(\mathbf{X}) \xrightarrow{\mathbb{P}} m(\mathbf{X})$ by summation of the p previous convergences which concludes the proof of the Theorem.

5 Simulation Study

We present a brief simulation study for our proposed method to handle missing values in a regression task. The study compares this approach with other 5 methods to handle missing values with random forests. One of them, corresponding to median imputation before the construction of the random forest is taken as a simple baseline. Three methods correspond to more elaborated imputation algorithms corresponding to the Breiman's method presented in Breiman and Cutler (2003), the improvement suggested by Ishioka (2013) (here presented simply as Ishioka's method) and missForest (Stekhoven and Bühlmann, 2011). Finally, the last method considered in this study corresponds to Missing Incorporated in Attributes (MIA) (Twala, Jones, and Hand, 2008) which uses the observations with missing values directly in the construction of the recursive trees. Two of these approaches, corresponding to missForest and MIA could be considered as state-of-the-art algorithms to handle missing values through random forests¹.

5.1 Approaches considered in the simulation

To properly introduce the algorithms considered in this study, we need to define the connectivity between two points in a tree and the proximity matrix of the forest. Let $K_{\Theta, n}(\mathbf{X}, \mathbf{X}')$ be the indicator that \mathbf{X} is in the same final cell that \mathbf{X}' in the tree designed with \mathcal{D}_n and the parameter Θ . If $K_{\Theta, n}(\mathbf{X}, \mathbf{X}') = 1$ we say that \mathbf{X} and \mathbf{X}' are connected in the tree $m_n(\cdot; \Theta)$. The proximity of two points is the average of times in which they are connected in the forest. Formally, let us define

¹Codes to reproduce our results can be found in: <https://github.com/IrvingGomez/RandomForestsSimulations>

the proximity between \mathbf{X} and \mathbf{X}' in the finite forest, $m_{M,n}(\cdot; \Theta_1, \dots, \Theta_M)$, as

$$K_{M,n}(\mathbf{X}, \mathbf{X}') = \frac{1}{M} \sum_{k=1}^M K_{\Theta_k,n}(\mathbf{X}, \mathbf{X}')$$

Breiman's and Ishioka's approach operate through imputation of missing values in a recursive way. First, they use the original training data set, \mathcal{D}_n , to fill the blank spaces in a roughly way. We denote this new data set as \mathcal{D}_{n,t_1} . The imputed data set \mathcal{D}_{n,t_1} is used to build a random forest. Then the proximity matrix is used to improve the imputation, resulting in a new data set \mathcal{D}_{n,t_2} . The procedure follows iteratively until some stopping rule is applied. Let $K_{M,t_\ell}(i, j)$ be the proximity between \mathbf{X}_i and \mathbf{X}_j at time t_ℓ . And denote by $\widehat{\mathbf{X}}_{i,t_\ell}$ the imputation of the observation \mathbf{X}_i at time t_ℓ . We also define $\mathbf{i}_{miss}^{(h)} \subseteq \{1, \dots, n\}$ as the indexes where $\mathbf{X}^{(h)}$ is missing, and $\mathbf{i}_{obs}^{(h)} = \{1, \dots, n\} \setminus \mathbf{i}_{miss}^{(h)}$ as the indexes where $\mathbf{X}^{(h)}$ is observed.

Breiman's Approach. If $\mathbf{X}^{(h)}$ is a continuous variable, $\widehat{\mathbf{X}}_{j,t_{\ell+1}}^{(h)}$ is the weighted mean of the observed values in $\mathbf{X}^{(h)}$, where the weights are defined by the proximity matrix of the current random forest, that is

$$\widehat{\mathbf{X}}_{j,t_{\ell+1}}^{(h)} = \frac{\sum_{i \in \mathbf{i}_{obs}^{(h)}} K_{M,t_\ell}(i, j) \mathbf{X}_i^{(h)}}{\sum_{i \in \mathbf{i}_{obs}^{(h)}} K_{M,t_\ell}(i, j)}, \quad \begin{array}{l} \ell \geq 1 \\ j \in \mathbf{i}_{miss}^{(h)} \end{array}$$

On the other hand, if $\mathbf{X}^{(h)}$ is a categorical variable, $\widehat{\mathbf{X}}_{j,t_{\ell+1}}^{(h)}$ is given by

$$\widehat{\mathbf{X}}_{j,t_{\ell+1}}^{(h)} = \arg \max_{\mathbf{x} \in \mathcal{X}^{(h)}} \sum_{i \in \mathbf{i}_{obs}^{(h)}} K_{M,t_\ell}(i, j) \mathbb{1}_{\mathbf{X}_i^{(h)} = \mathbf{x}}, \quad \begin{array}{l} \ell \geq 1 \\ j \in \mathbf{i}_{miss}^{(h)} \end{array}$$

That is, $\widehat{\mathbf{X}}_{j,t_{\ell+1}}^{(h)}$ is the class that maximizes the sum of the proximity considering the observed values in the class.

Ishioka's Approach. If $\mathbf{X}^{(h)}$ is a continuous variable, $\widehat{\mathbf{X}}_{j,t_{\ell+1}}^{(h)}$ is the weighted mean of the k nearest neighbors, according to the proximity matrix, over all the values, both imputed and observed. The k closest values are chosen to make more robust the method and avoid values which are outliers.

$$\widehat{\mathbf{X}}_{j,t_{\ell+1}}^{(h)} = \frac{\sum_{\substack{i \in \text{neigh}_k \\ i \neq j}} K_{M,t_\ell}(i, j) \widehat{\mathbf{X}}_{i,t_\ell}^{(h)}}{\sum_{\substack{i \in \text{neigh}_k \\ i \neq j}} K_{M,t_\ell}(i, j)}, \quad \begin{array}{l} \ell \geq 1 \\ j \in \mathbf{i}_{miss}^{(h)} \end{array}$$

For categorical variables, it is not necessary to see only the k closest values because the outliers of \mathbf{X} will have few attention. Meanwhile the proximity with missing values should have more attention, specially when the missing rate is high. Hence, if $\mathbf{X}^{(h)}$ is a categorical variable, $\widehat{\mathbf{X}}_{j,t_{\ell+1}}^{(h)}$ is given by

$$\widehat{\mathbf{X}}_{j,t_{\ell+1}}^{(h)} = \arg \max_{\mathbf{x} \in \mathcal{X}^{(h)}} \sum_{i \neq j} K_{M,t_\ell}(i, j) \mathbb{1}_{\mathbf{X}_{i,t_\ell}^{(h)} = \mathbf{x}}, \quad \begin{array}{l} \ell \geq 1 \\ j \in \mathbf{i}_{miss}^{(h)} \end{array}$$

MissForest. This algorithm begins with a rough imputation for the missing values, for example the median of each direction. Then, for each direction $\mathbf{X}^{(h)}$ a random forests is built using all the other directions $\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(h-1)}, \mathbf{X}^{(h+1)}, \dots, \mathbf{X}^{(p)}$ and the response, where the random forests is constructed using the observations such that $\mathbf{M}^{(h)} = 0$. Then the missing values of $\mathbf{X}^{(h)}$ are predicted with this random forests, these predictions correspond to the imputation of the missing values. These steps are repeated iteratively until a stopping rule is achieved.

Missing Incorporated in Attributes (MIA). The Missing Incorporated in Attributes (MIA) approach consists in keeping all the missing values together when a split is performed. Thus, the splits with this approach assign the values according to one of the following rules:

- $\{\mathbf{X}^{(h)} < z \text{ and } \mathbf{M}^{(h)} = 1\}$ versus $\{\mathbf{X}^{(h)} \geq z\}$.
- $\{\mathbf{X}^{(h)} < z\}$ versus $\{\mathbf{X}^{(h)} \geq z \text{ and } \mathbf{M}^{(h)} = 1\}$.
- $\{\mathbf{M}^{(h)} = 0\}$ versus $\{\mathbf{M}^{(h)} = 1\}$.

5.2 Parameters considered and results

The regression function is the so-called “friedman1” (Friedman et al., 1991), which has been used in previous simulation studies (Breiman, 1996; Rieger, Hothorn, and Strobl, 2010; Josse et al., 2019; Friedberg et al., 2020), given by

$$m(\mathbf{x}) = 10 \sin\left(\pi \mathbf{x}^{(1)} \mathbf{x}^{(2)}\right) + 20 \left(\mathbf{x}^{(3)} - 0.5\right)^2 + 10 \mathbf{x}^{(4)} + 5 \mathbf{x}^{(5)}.$$

Following the schema presented by Rieger, Hothorn, and Strobl (2010), we simulate \mathbf{X} as a uniformly distributed variable on $[0, 1]^5$ and introduce missing values in $\mathbf{X}^{(1)}, \mathbf{X}^{(3)}$ and $\mathbf{X}^{(4)}$, considering an MCAR data-missing mechanism. We create one testing data set and 100 training data sets. Each training data set contains 200 observations and the testing data set contains 2000 observations. This amount of data is to have an appropriate approximation to the MSE

$$\mathbb{E}_{\mathbf{X}|\mathcal{D}_n} [m_{M,n}(\mathbf{X}) - m(\mathbf{X})]^2.$$

For each training data set $\mathbf{X}^{(1)}$ presents 20% of missing data, $\mathbf{X}^{(3)}$ has 10% of missing data, while for $\mathbf{X}^{(4)}$ we have decided to change the fraction of missingness varying between 5%, 10%, 20%, 40%, 60%, 80%, 90% and 95%. We do not introduce missing values in the testing data set. A random forest is built for each one of the 100 training data sets with missing values and without missing values (which are used as benchmark), using $M = 50$ trees, that has been seen by simulation to be sufficient to stabilize the error in the case of complete training data sets. For the remain parameters we use the default values in the regression mode of the R package `randomForests`, the parameter `mtry` is set to $\lfloor p/3 \rfloor$, we have sampled without replacement, so a_n is set to $\lceil 0.632n \rceil$ and `nodesize` is set to 5.

Figure 5 presents the average MSE over the training data sets for each percentage of missing data considered. Median imputation could be taken as a bound of the minimum expected behavior for a method that attempts to estimate the regression function with missing values. We observe that missForest and the approach here presented generate estimators with the lowest MSE regardless of the percentage of missingness. Important differences between the performance of the methods become clear with the increasing percentage of missing values, especially when this value is over 60%. Moreover, we can see the advantage of searching for the best assignation when there is a high

percentage of missing values with our proposal outperforming all the other algorithms. While all the methods present a similar MSE when this percentage is smaller than 20%.

Another interesting observation corresponds to the relation between our procedure and Missing Incorporated in Attributes (MIA), which assigns all the observations where the split variable is missing to the same node. Our method goes further than MIA, since it does not only consider these assignments, but the best possible partial imputation for the observations with missing values. Furthermore, note that the algorithm proposed in this work has the advantage of not having to calculate extra structures, like the proximity matrix. The work presented in Gómez-Méndez and Joly (2021) is devoted to the computational study of the algorithm here proposed, it presents results of an extensive simulation study considering different data-missing mechanisms beyond MCAR.




Figure 5: Average MSE for the testing data set for each percentage of missingness

References

- Biau, Gérard (2012). “Analysis of a random forests model”. In: *Journal of Machine Learning Research* 13.Apr, pp. 1063–1095 (see p. 1).
- Biau, Gérard, Luc Devroye, and Gábor Lugosi (2008). “Consistency of random forests and other averaging classifiers”. In: *Journal of Machine Learning Research* 9.Sep, pp. 2015–2033 (see p. 1).
- Biau, Gérard and Erwan Scornet (2016). “A random forest guided tour”. In: *Test* 25.2, pp. 197–227 (see p. 2).
- Breiman, Leo (1996). “Bagging predictors”. In: *Machine learning* 24.2, pp. 123–140 (see pp. 1, 2, 26).
- (2001). “Random forests”. In: *Machine learning* 45.1, pp. 5–32 (see pp. 1, 5).
- (2004). “Consistency for a simple model of random forests”. In: (see p. 1).
- Breiman, Leo and Adele Cutler (2003). “Manual—setting up, using, and understanding random forests V4.0.2003”. In: URL: https://www.stat.berkeley.edu/~breiman/Using_random_forests_v4.0.pdf (see pp. 2, 24).

- Breiman, Leo et al. (1984). *Classification and regression trees*. Chapman and Hall/CRC (see pp. 2, 5).
- Cutler, D Richard et al. (2007). “Random forests for classification in ecology”. In: *Ecology* 88.11, pp. 2783–2792 (see p. 1).
- Díaz-Uriarte, Ramón and Sara Alvarez De Andres (2006). “Gene selection and classification of microarray data using random forest”. In: *BMC bioinformatics* 7.1, p. 3 (see p. 1).
- Friedberg, Rina et al. (2020). “Local linear forests”. In: *Journal of Computational and Graphical Statistics*, pp. 1–15 (see pp. 2, 26).
- Friedman, Jerome H et al. (1991). “Multivariate adaptive regression splines”. In: *The annals of statistics* 19.1, pp. 1–67 (see p. 26).
- Friedman, Jerome, Trevor Hastie, and Robert Tibshirani (2009). *The elements of statistical learning*. 2nd ed. Springer series in statistics New York (see p. 2).
- Genuer, Robin (2012). “Variance reduction in purely random forests”. In: *Journal of Nonparametric Statistics* 24.3, pp. 543–562 (see p. 1).
- Gini, Corrado (1912). “Variabilità e mutabilità”. In: *Reprinted in Memorie di metodologica statistica (Ed. Pizetti E, Salvemini, T). Rome: Libreria Eredi Virgilio Veschi* (see p. 5).
- Gómez-Méndez, Irving and Emilien Joly (2021). “Regression with Missing Data, a Comparison Study of Techniques Based on Random Forests”. Preprint (see p. 27).
- Hapfelmeier, Alexander, Torsten Hothorn, and Kurt Ulm (2012). “Recursive partitioning on incomplete data using surrogate decisions and multiple imputation”. In: *Computational Statistics & Data Analysis* 56.6, pp. 1552–1565 (see p. 2).
- Hothorn, Torsten, Kurt Hornik, and Achim Zeileis (2006). “Unbiased recursive partitioning: A conditional inference framework”. In: *Journal of Computational and Graphical statistics* 15.3, pp. 651–674 (see p. 2).
- Ishioka, Tsunenori (2013). “Imputation of missing values for unsupervised data using the proximity in random forests”. In: *International Conference on Mobile, Hybrid, and On-line Learning. Nice*, pp. 30–36 (see pp. 2, 24).
- Ishwaran, Hemant and Udaya B Kogalur (2010). “Consistency of random survival forests”. In: *Statistics & probability letters* 80.13-14, pp. 1056–1064 (see p. 1).
- Josse, Julie, Jerome P Reiter, et al. (2018). “Introduction to the special section on missing data”. In: *Statistical Science* 33.2, pp. 139–141 (see p. 2).
- Josse, Julie et al. (2019). “On the consistency of supervised learning with missing values”. In: *arXiv preprint arXiv:1902.06931* (see pp. 2, 26).
- Laurent, Beatrice and Pascal Massart (2000). “Adaptive estimation of a quadratic functional by model selection”. In: *Annals of Statistics*, pp. 1302–1338 (see p. 17).
- Mentch, Lucas and Giles Hooker (2016). “Quantifying uncertainty in random forests via confidence intervals and hypothesis tests”. In: *The Journal of Machine Learning Research* 17.1, pp. 841–881 (see p. 1).
- Prasad, Anantha M, Louis R Iverson, and Andy Liaw (2006). “Newer classification and regression tree techniques: bagging and random forests for ecological prediction”. In: *Ecosystems* 9.2, pp. 181–199 (see p. 1).
- Quinlan, J. Ross (1986). “Induction of decision trees”. In: *Machine learning* 1.1, pp. 81–106 (see p. 5).
- (1993). *C4.5: Programs for Machine Learning*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc. (see p. 5).
- Rieger, Anna, Torsten Hothorn, and Carolin Strobl (2010). “Random forests with missing values in the covariates”. In: (see pp. 2, 26).

- Ripley, Brian D (2007). *Pattern recognition and neural networks*. Cambridge university press (see p. 2).
- Rubin, Donald B (1976). “Inference and missing data”. In: *Biometrika* 63.3, pp. 581–592 (see p. 6).
- (1996). “Multiple imputation after 18+ years”. In: *Journal of the American statistical Association* 91.434, pp. 473–489 (see p. 2).
- (2004). *Multiple imputation for nonresponse in surveys*. Vol. 81. John Wiley & Sons (see p. 2).
- Scornet, Erwan, Gérard Biau, and Jean-Philippe Vert (2015). “Consistency of random forests”. In: *The Annals of Statistics* 43.4, pp. 1716–1741 (see pp. 1, 11).
- Shannon, Claude Elwood (1948). “A mathematical theory of communication”. In: *Bell system technical journal* 27.3, pp. 379–423 (see p. 5).
- Shotton, Jamie et al. (2011). “Real-time human pose recognition in parts from single depth images”. In: *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*. Ieee, pp. 1297–1304 (see p. 1).
- Stekhoven, Daniel J and Peter Bühlmann (2011). “MissForest—non-parametric missing value imputation for mixed-type data”. In: *Bioinformatics* 28.1, pp. 112–118 (see pp. 2, 24).
- Svetnik, Vladimir et al. (2003). “Random forest: a classification and regression tool for compound classification and QSAR modeling”. In: *Journal of chemical information and computer sciences* 43.6, pp. 1947–1958 (see p. 1).
- Twala, Beth, MC Jones, and David J Hand (2008). “Good methods for coping with missing data in decision trees”. In: *Pattern Recognition Letters* 29.7, pp. 950–956 (see pp. 2, 24).
- Varian, Hal R (2014). “Big data: New tricks for econometrics”. In: *Journal of Economic Perspectives* 28.2, pp. 3–28 (see p. 1).
- Venables and Ripley (2002). “Modern Applied Statistics with S”. In: *Springer, New York* 1228, p. 1229 (see p. 2).
- Wager, Stefan and Susan Athey (2018). “Estimation and inference of heterogeneous treatment effects using random forests”. In: *Journal of the American Statistical Association* 113.523, pp. 1228–1242 (see p. 1).
- Zhu, Ruoqing, Donglin Zeng, and Michael R Kosorok (2015). “Reinforcement learning trees”. In: *Journal of the American Statistical Association* 110.512, pp. 1770–1784 (see p. 1).