



HAL
open science

Trees, forests, and impurity-based variable importance

Erwan Scornet

► **To cite this version:**

| Erwan Scornet. Trees, forests, and impurity-based variable importance. 2020. hal-02436169v2

HAL Id: hal-02436169

<https://hal.science/hal-02436169v2>

Preprint submitted on 9 Nov 2020 (v2), last revised 23 Dec 2021 (v3)

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

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

Trees, forests, and impurity-based variable importance

Erwan Scornet

*Centre de Mathématiques Appliquées, Ecole Polytechnique, CNRS,
Institut Polytechnique de Paris, Palaiseau, France*
erwan.scornet@polytechnique.edu

Abstract

Tree ensemble methods such as random forests [Breiman, 2001] are very popular to handle high-dimensional tabular data sets, notably because of their good predictive accuracy. However, when machine learning is used for decision-making problems, settling for the best predictive procedures may not be reasonable since enlightened decisions require an in-depth comprehension of the algorithm prediction process. Unfortunately, random forests are not intrinsically interpretable since their prediction results from averaging several hundreds of decision trees. A classic approach to gain knowledge on this so-called black-box algorithm is to compute variable importances, that are employed to assess the predictive impact of each input variable. Variable importances are then used to rank or select variables and thus play a great role in data analysis. Nevertheless, there is no justification to use random forest variable importances in such way: we do not even know what these quantities estimate. In this paper, we analyze one of the two well-known random forest variable importances, the Mean Decrease Impurity (MDI). We prove that if input variables are independent and in absence of interactions, MDI provides a variance decomposition of the output, where the contribution of each variable is clearly identified. We also study models exhibiting dependence between input variables or interaction, for which the variable importance is intrinsically ill-defined. Our analysis shows that there may exist some benefits to use a forest compared to a single tree.

Index Terms — Random forests, variable importance, Mean Decrease Impurity.

2010 Mathematics Subject Classification: 62G08, 62G20.

1 Introduction

Decision trees [see Loh, 2011, for an overview] can be used to solve pattern recognition tasks in an interpretable way: in order to build a prediction, trees ask to each observation a series of questions, each one being of the form “Is variable $X^{(j)}$ larger than a threshold s ?” where j, s are to be determined by the algorithm. Thus the prediction for a new observation only depends on the sequence of questions/answers for this observation. One of the most popular decision trees is the CART procedure

[Classification And Regression Trees, [Breiman et al., 1984](#)], which unfortunately suffers from low accuracy and intrinsic instability: modifying slightly one observation in the training set can change the entire tree structure.

To overcome this last issue, and also improve their accuracy, ensemble methods, which grow many base learners and aggregate them to predict, have been designed. Random forests, introduced by [Breiman \[2001\]](#), are among the most famous tree ensemble methods. They proceed by growing trees based on CART procedure [[Breiman et al., 1984](#)] and randomizing both the training set and the splitting variables in each node of the tree. Breiman’s [2001] random forests have been under active investigation during the last two decades mainly because of their good practical performance and their ability to handle high-dimensional data sets. They are acknowledged to be state-of-the-art methods in fields such as genomics [[Qi, 2012](#)] and pattern recognition [[Rogez et al., 2008](#)], just to name a few. If empirical performances of random forests are not to be demonstrated anymore [[Fernández-Delgado et al., 2014](#)], their main flaw relies on their lack of interpretability because their predictions result from averaging over a large number of trees (typically several hundreds): the ease of interpretation of each tree cannot be straightforwardly extended to random forests.

Since interpretability is a concept difficult to define precisely, people eager to gain insights about the driving forces at work behind random forests predictions often focus on variable importance, a measure of the influence of each input variable to predict the output. In Breiman’s [2001] original random forests, there exist two importance measures: the Mean Decrease Impurity [MDI, or Gini importance, see [Breiman, 2002](#)], which sums up the gain associated to all splits performed along a given variable; and the Mean Decrease Accuracy [MDA, or permutation importance, see [Breiman, 2001](#)] which shuffles entries of a specific variable in the test data set and computes the difference between the error on the permuted test set and the original test set. Both measures are used in practice even if they possess several major drawbacks.

MDI is known to favor variables with many categories [see, e.g., [Strobl et al., 2007](#), [Nicodemus, 2011](#)]. Even when variables have the same number of categories, MDI exhibits empirical bias towards variables that possess a category having a high frequency [[Nicodemus, 2011](#), [Boulesteix et al., 2011](#)]. MDI is also biased in presence of correlated features [[Nicodemus and Malley, 2009](#)]. A promising way to assess variable importance in random forest consists in designing new tree building process or new feature importances as in the R package `party` [[Strobl et al., 2008, 2009](#)] or in [Li et al. \[2019\]](#), [Zhou and Hooker \[2019\]](#). On the other hand, MDA seems to exhibit less bias but tends to overestimate correlated features [Strobl et al. \[2008\]](#). Besides, its scale version (the default version in the R package `randomForest`) depends on the sample size and on the number of trees, the last being an arbitrary chosen parameter [[Strobl and Zeileis, 2008](#)]. The interested reader may refer to [Genuer et al. \[2008\]](#) and [Genuer et al. \[2010\]](#) for an extensive simulation study about the influence of the number of observations, variables, and trees on MDA together with the impact of correlation on this importance measure. Despite all their shortcomings, one great advantage of MDI and MDA is their ability to take into account interaction between features even if unable to determine the part of marginal/joint effect [[Wright et al., 2016](#)].

From a theoretical perspective, there are only but a few results on variable importance computed via tree procedures. The starting point for theory is by [Ishwaran \[2007\]](#) who studies a modified version of permutation importance and derives its asymptotic positive expression, when the regression function is assumed to be piecewise constant.

A scaled population version of the permutation importance is studied by [Zhu et al. \[2012\]](#) who establish rate of convergence for the probability that a relevant feature is chosen in a modified random forest algorithm. The exact population expression of permutation importance is computed by [Gregorutti et al. \[2017\]](#) for a Gaussian linear model.

Regarding the MDI, there are even fewer results. [Louppe et al. \[2013\]](#) study the population MDI when all features are categorical. In this framework, they propose a decomposition of the total information, depending on the MDI of each variable and of interaction terms. They also prove that the MDI of an irrelevant variable is zero and that adding irrelevant variables does not impact the MDI of relevant ones. Even if these results may appear to be very simple, the proofs are not straightforward at all, which explains the few results on this topic. This work was later extended by [Sutera et al. \[2016\]](#) to the case of context-dependent features.

It is instructive to notice that there is no general result establishing the consistency of these variable importance measures toward population quantities when computed with the original random forest algorithms: all existing results focus on modified random forests version with, in some cases, strong assumptions on the regression model. Therefore, there are no guarantees that using variable importance computed via random forests is suitable to select variables, which is nevertheless often done in practice.

Agenda. In this paper, we focus on the regression framework and derive theoretical properties about the MDI used in Breiman’s [\[2001\]](#) random forests. Section [2](#) is devoted to notations and describes tree and forest construction, together with MDI calculation. The main results are gathered in Section [3](#). We prove that both population MDI and empirical MDI, computed with Breiman’s random forests, can be seen as a percentage of explained variance and are directly related to the quadratic risk of the forest. We also derive the expression of the sum of MDIs for two groups of variables that are independent and do not interact, which is at the core of our analysis. The analysis highlights that fully-developed random forests (whose tree leaves contain exactly one observation) produce inaccurate variable importances and should not be used for that purpose. In Section [4](#), we establish that for additive regression models with independent inputs, the MDI of each variable takes a very simple, interpretable form, which is the variance of each univariate component $m_j(X^{(j)})$. We prove that empirical MDI computed with Breiman’s forest targets this population quantity, which is nothing but a standard way to decompose the variance of the output. Hence, in the absence of input dependence and correlation, MDI computed with random forests provides a good assessment of the variable importance. To the best of our knowledge, it is the first result proving the good properties of the empirical MDI computed with Breiman’s forests. To move beyond the additive models and understand the impact of interactions, we study a multiplicative model in Section [5](#). In this model, MDI is ill-defined and one needs to aggregate many trees to obtain a consistent MDI value. Averaging trees is known to be of great importance to improve accuracy but it turns out to be mandatory to obtain consistent variable importance measure in the presence of interactions. The impact of dependence between input variables is studied in Section [6](#). We stress out that correlated variables are more likely to be chosen in the splitting procedure and, as a consequence, exhibit a larger MDI than independent ones. As for the presence of interactions, this highlights the importance of computing MDI by averaging many trees. Experiments presented in Section [7](#) illustrate the prop-

erties of MDI computed via trees and forests in a finite-sample regime. Additional experiments are gathered in Section 8 while proofs are postponed to Section 9.

2 Notation and general definition

We assume to be given a data set $\mathcal{D}_n = \{(\mathbf{X}_i, Y_i), i = 1, \dots, n\}$ of *i.i.d.* observations (\mathbf{X}_i, Y_i) distributed as the generic pair (\mathbf{X}, Y) where $\mathbf{X} \in [0, 1]^d$ and $Y \in \mathbb{R}$ with $\mathbb{E}[Y^2] < \infty$. Our aim is to estimate the regression function $m : [0, 1]^d \rightarrow \mathbb{R}$, defined as $m(\mathbf{X}) = \mathbb{E}[Y|\mathbf{X}]$, using a random forest procedure and to identify relevant variables, i.e. variables on which m depends.

2.1 CART construction and empirical MDI

CART [Breiman et al., 1984] is the elementary component of random forests [Breiman, 2001]. CART construction works as follows. Each node of a single tree \mathcal{T} is associated with a hyper-rectangular cell included in $[0, 1]^d$. The root of the tree is $[0, 1]^d$ itself and, at each step of the tree construction, a node (or equivalently its corresponding cell) is split in two parts. The terminal nodes (or leaves), taken together, form a partition of $[0, 1]^d$. An example of such tree and partition is depicted in Figure 1, whereas the full procedure is described in Algorithm 1.

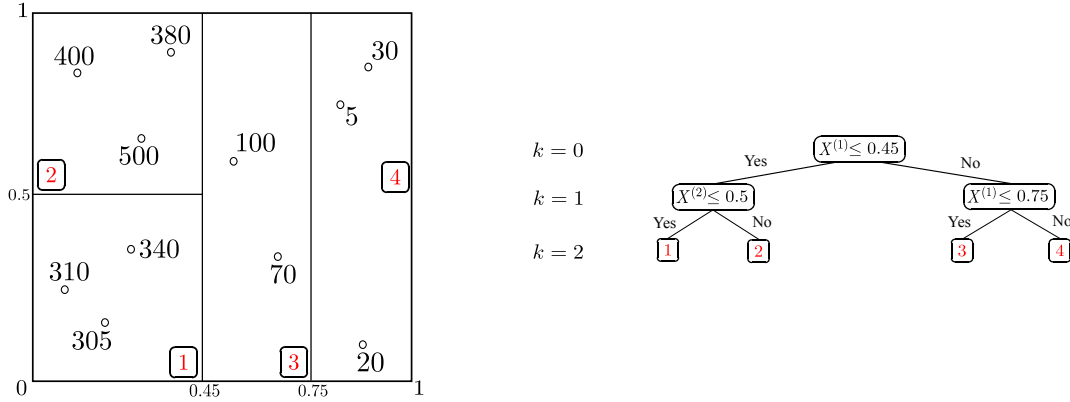


Figure 1: A decision tree of depth $k = 2$ in dimension $d = 2$ (right) and the corresponding partition (left).

We define now the CART-split criterion. To this aim, we let $A \subset [0, 1]^d$ be a generic cell and $N_n(A)$ be the number of data points falling into A . A split in A is a pair (j, z) , where j is a dimension in $\{1, \dots, d\}$ and z is the position of the cut along the j -th coordinate, within the limits of A . We let \mathcal{C}_A be the set of all such possible cuts in A . Then, for any $(j, z) \in \mathcal{C}_A$, the CART-split criterion [Breiman et al., 1984] takes the form

$$L_{n,A}(j, z) = \frac{1}{N_n(A)} \sum_{i=1}^n (Y_i - \bar{Y}_A)^2 \mathbb{1}_{\mathbf{x}_i \in A} - \frac{1}{N_n(A)} \sum_{i=1}^n (Y_i - \bar{Y}_{A_L} \mathbb{1}_{\mathbf{x}_i^{(j)} < z} - \bar{Y}_{A_R} \mathbb{1}_{\mathbf{x}_i^{(j)} \geq z})^2 \mathbb{1}_{\mathbf{x}_i \in A}, \quad (1)$$

where $A_L = \{\mathbf{x} \in A : \mathbf{x}^{(j)} < z\}$, $A_R = \{\mathbf{x} \in A : \mathbf{x}^{(j)} \geq z\}$, and \bar{Y}_A (resp., \bar{Y}_{A_L} , \bar{Y}_{A_R}) is the average of the Y_i 's belonging to A (resp., A_L , A_R), with the convention $0/0 = 0$.

Algorithm 1 CART construction.

- 1: **Inputs:** data set \mathcal{D}_n , query point $\mathbf{x} \in [0, 1]^d$, maximum number of observations in each leaf `nodesize`
 - 2: **Set** $\mathcal{A}_{\text{leaves}} = \emptyset$ and $\mathcal{A}_{\text{inner nodes}} = \{[0, 1]^d\}$.
 - 3: **while** $\mathcal{A}_{\text{inner nodes}} \neq \emptyset$ **do**
 - 4: Select the first element $A \in \mathcal{A}_{\text{inner nodes}}$
 - 5: **if** A contains more than `nodesize` observations **then**
 - 6: Select the best split minimizing the CART-split criterion defined below (see equation (1))
 - 7: Split the cell accordingly.
 - 8: Concatenate the two resulting cell A_L and A_R at the end of $\mathcal{A}_{\text{inner nodes}}$.
 - 9: Remove A from $\mathcal{A}_{\text{inner nodes}}$
 - 10: **else**
 - 11: Concatenate A at the end of $\mathcal{A}_{\text{leaves}}$.
 - 12: Remove A from $\mathcal{A}_{\text{inner nodes}}$
 - 13: **end if**
 - 14: **end while**
 - 15: For a query point \mathbf{x} , the tree outputs the average $m_n(\mathbf{x})$ of the Y_i falling into the same leaf as \mathbf{x} .
-

At each cell A , the best cut $(j_{n,A}, z_{n,A})$ is finally selected by maximizing $L_{n,A}(j, z)$ over \mathcal{C}_A , that is

$$(j_{n,A}, z_{n,A}) \in \arg \max_{(j,z) \in \mathcal{C}_A} L_{n,A}(j, z). \quad (2)$$

To remove ties in the argmax, the best cut is always performed along the best cut direction $j_{n,A}$, at the middle of two consecutive data points.

The Mean Decrease in Impurity (MDI) for the variable $X^{(j)}$ computed via a tree \mathcal{T} is defined by

$$\widehat{\text{MDI}}_{\mathcal{T}}(X^{(j)}) = \sum_{\substack{A \in \mathcal{T} \\ j_{n,A}=j}} p_{n,A} L_{n,A}(j_{n,A}, z_{n,A}), \quad (3)$$

where the sum ranges over all cells A in \mathcal{T} that are split along variable j and $p_{n,A}$ is the fraction of observations falling into A . In other words, the MDI of $X^{(j)}$ computes the weighted decrease in impurity related to splits along the variable $X^{(j)}$.

2.2 Theoretical CART construction and population MDI

In order to study the (empirical) MDI defined in equation (3), we first need to define and analyze the population version of MDI. First, we define a theoretical CART tree, as in Algorithm 1, except for the splitting criterion which is replaced by its population version. Namely, for all cells A and all splits $(j, z) \in \mathcal{C}_A$, the population version of the CART-split criterion is defined as

$$\begin{aligned} L_A^*(j, z) = & \mathbb{V}[Y|\mathbf{X} \in A] - \mathbb{P}[\mathbf{X}^{(j)} < z | \mathbf{X} \in A] \mathbb{V}[Y|\mathbf{X}^{(j)} < z, \mathbf{X} \in A] \\ & - \mathbb{P}[\mathbf{X}^{(j)} \geq z | \mathbf{X} \in A] \mathbb{V}[Y|\mathbf{X}^{(j)} \geq z, \mathbf{X} \in A]. \end{aligned} \quad (4)$$

Therefore in each cell of a theoretical tree \mathcal{T}^* , the best split (j_A^*, z_A^*) is chosen by maximizing the population version of the CART-split criterion that is

$$(j_A^*, z_A^*) \in \arg \max_{(j,z) \in \mathcal{C}_A} L_A^*(j, z).$$

Of course, in practice, we cannot build a theoretical tree \mathcal{T}^* since it relies on the true distribution (\mathbf{X}, Y) which is unknown. A theoretical tree is just an abstract mathematical object, for which we prove properties that will be later extended to the (empirical) CART tree, our true object of interest.

As for the empirical MDI defined above, the population MDI for the variable $X^{(j)}$ computed via the theoretical tree \mathcal{T}^* is defined as

$$\text{MDI}_{\mathcal{T}^*}^*(X^{(j)}) = \sum_{\substack{A \in \mathcal{T}^* \\ j_A^* = j}} p_A^* L_A^*(j_A^*, z_A^*),$$

where all empirical quantities have been replaced by their population version and the theoretical CART tree is used instead of the empirical CART tree defined in Section 2.1. We also let $p_A^* = \mathbb{P}[\mathbf{X} \in A]$.

Random forest A random forest is nothing but a collection of several decision trees whose construction has been randomized. In Breiman's forest, CART procedure serves as base learner and the randomization is performed in two different ways:

1. Prior to the construction of each tree, a sample of the original data set is randomly drawn. Only this sample is used to build the tree. Sampling is typically done using bootstrap, by randomly drawing n observations out of the original n observations with replacement.
2. Additionally, for each cell, the best split is not selected along all possible variables. Instead, for each cell, a subsample of `mtry` variables is randomly selected without replacement. The best split is selected only along these variables. By default, `mtry` = $d/3$ in regression.

Random forest prediction is then simply the average of the predictions of such randomized CART trees. Similarly, the MDI computed via a forest is nothing but the average of the MDI computed via each tree of the forest.

In the sequel, we will use the theoretical random forest (forest that aggregate theoretical CART trees) and the population MDI to derive results on the empirical random forest (forest that aggregate empirical CART trees; the one widely used in practice) and the empirical MDI.

3 Main theoretical result

For any theoretical CART tree \mathcal{T}^* , and for any $k \in \mathbb{N}$, we denote by \mathcal{T}_k^* the truncation of \mathcal{T}^* at level k , that is the subtree of \mathcal{T}^* rooted at $[0, 1]^d$, in which each leaf has been cut at most k times. Let $A_{\mathcal{T}_k^*}(\mathbf{x})$ be the cell of the tree \mathcal{T}_k^* containing \mathbf{x} . For any function $f : [0, 1]^d \rightarrow \mathbb{R}$ and any cell $A \subset [0, 1]^d$, let

$$\Delta(f, A) = \mathbb{V}[f(\mathbf{X}) | \mathbf{X} \in A]$$

be the variance of the function f in the cell A with respect to the distribution of \mathbf{X} . Proposition 1 states that the population MDI computed via theoretical CART trees can be used to decompose the variance of the output, up to some residual term which depends on the risk of the theoretical tree estimate.

Proposition 1. *Assume that $Y = m(\mathbf{X}) + \varepsilon$, where ε is a noise satisfying $\mathbb{E}[\varepsilon|\mathbf{X}] = 0$ and $\mathbb{V}[\varepsilon|\mathbf{X}] = \sigma^2$ almost surely, for some constant σ^2 . Consider a theoretical CART tree \mathcal{T}^* . Then, for all $k \geq 0$,*

$$\mathbb{V}[Y] = \sum_{j=1}^d \text{MDI}_{\mathcal{T}_k^*}^*(X^{(j)}) + \mathbb{E}_{\mathbf{X}'}[\mathbb{V}[Y|\mathbf{X} \in A_{\mathcal{T}_k^*}(\mathbf{X}')]], \quad (5)$$

where \mathbf{X}' is an i.i.d. copy of \mathbf{X} .

Note that few assumptions are made in Proposition 1, which makes it very general. In addition, one can see that the sum of population MDI is close to the R^2 measure used to assess the quality of regression model. The population R^2 is defined as

$$1 - \frac{\mathbb{E}_{\mathbf{X}'}[\mathbb{V}[Y|\mathbf{X} \in A_{\mathcal{T}_k^*}(\mathbf{X}')]]}{\mathbb{V}[Y]} = \frac{\sum_{j=1}^d \text{MDI}_{\mathcal{T}_k^*}^*(X^{(j)})}{\mathbb{V}[Y]}.$$

Hence, the sum of MDI divided by the variance of the output corresponds to the percentage of variance explained by the model, which is exactly the population R^2 . Therefore, Proposition 1 draws a clear connection between MDI and the very classical R^2 measure.

We say that a theoretical tree \mathcal{T}^* is consistent if $\lim_{k \rightarrow \infty} \Delta(m, A_{\mathcal{T}_k^*}(\mathbf{X}')) = 0$. If a theoretical tree is consistent, then its population R^2 tends to $\mathbb{V}[m(\mathbf{X})]/\mathbb{V}[Y]$ as shown in Corollary 1 below.

Corollary 1. *Grant assumptions of Proposition 1. Additionally, assume that $\|m\|_\infty < \infty$ and, almost surely, $\lim_{k \rightarrow \infty} \Delta(m, A_{\mathcal{T}_k^*}(\mathbf{X}')) = 0$. Then, for all $\gamma > 0$, there exists K such that, for all $k > K$,*

$$\left| \sum_{j=1}^d \text{MDI}_{\mathcal{T}_k^*}^*(X^{(j)}) - \mathbb{V}[m(\mathbf{X})] \right| \leq \gamma$$

Consistency (and, in this case, tree consistency) is a prerequisite when dealing with algorithm interpretability. Indeed, it is hopeless to think that we can leverage information about the true data distribution from an inconsistent algorithm, intrinsically incapable of modelling these data. Therefore, we assume in this section that theoretical trees are consistent and study afterwards variable importances produced by such algorithm. In the next sections, we will be able to prove tree consistency for specific regression models, allowing us to remove the generic consistency assumption in our results, as the one in Corollary 1.

Corollary 1 states that if the theoretical tree \mathcal{T}_k^* is consistent, then the sum of the MDI of all variables tends to the total amount of information of the model, that is $\mathbb{V}[m(\mathbf{X})]$. This gives a nice interpretation of MDI as a way to decompose the total variance $\mathbb{V}[m(\mathbf{X})]$ across variables. Louppe et al. [2013] prove a similar result when all features are categorical. This result was later extended by Suter et al. [2016] in the case of context-dependent features. In their case, trees are finite since there exists only

a finite number of variables with a finite number of categories. The major difference with our analysis is that trees we study can be grown to arbitrary depth. We thus need to control the error of a tree stopped at some level k , which is exactly the right-hand term in equation (5).

According to Corollary 1, the sum of the MDI is the same for all consistent theoretical trees: even if there may exist many theoretical trees, the sum of MDI computed for each theoretical tree tends to the same value, regardless of the considered tree. Therefore, all consistent theoretical tree produce the same asymptotic value for the sum of population MDI.

In practice, one cannot build a theoretical CART tree or compute the population MDI. Proposition 2 below is the extension of Proposition 1 for the empirical CART procedure. Let \mathcal{T}_n be the (empirical) CART tree built with data set \mathcal{D}_n and let, for all k , $\mathcal{T}_{n,k}$ be the truncation of \mathcal{T}_n at level k . Denote by $\widehat{\mathbb{V}}[Y]$ the empirical variance of Y computed on the data set \mathcal{D}_n . Define, for any function $f : [0, 1]^d \mapsto \mathbb{R}$, the empirical quadratic risk via

$$R_n(f) = \frac{1}{n} \sum_{i=1}^n (Y_i - f(X_i))^2.$$

Proposition 2 (Empirical version of Proposition 1). *Let \mathcal{T}_n be the CART tree, based on the data set \mathcal{D}_n . Then,*

$$\widehat{\mathbb{V}}[Y] = \sum_{j=1}^d \widehat{\text{MDI}}_{\mathcal{T}_n}(X^{(j)}) + R_n(f_{\mathcal{T}_n}), \quad (6)$$

where $f_{\mathcal{T}_n}$ is the estimate associated to \mathcal{T}_n , as defined in Algorithm 1.

Note that equations (5) and (6) in Proposition 1 and Proposition 2 are valid for general tree construction. The main argument of the proof is simply a telescopic sum of the MDI which links the variance of Y in the root node of the tree to the variance of Y in each terminal node. These equalities hold for general impurity measures by providing a relation between root impurity and leaves impurity. As in Proposition 1, Proposition 2 proves that the R^2 of a tree can be written as

$$\frac{\sum_{j=1}^d \widehat{\text{MDI}}_{\mathcal{T}_n}(X^{(j)})}{\widehat{\mathbb{V}}[Y]}, \quad (7)$$

which allows us to see the sum of MDI as the percentage of variance explained by the model. This is particularly informative about the quality of the tree modelling when the depth k of the tree is fixed.

Trees are likely to overfit when they are fully grown. Indeed, the risk of trees which contain only one observation per leaf is exactly zero. Hence, according to equation (6), we have

$$\widehat{\mathbb{V}}[Y] = \sum_{j=1}^d \widehat{\text{MDI}}_{\mathcal{T}_n}(X^{(j)})$$

and the R^2 of such tree is equal to one. Such R^2 does not mean that trees have good generalization error but that they have enough approximation capacity to overfit the

data. Additionally, for a fully grown tree \mathcal{T}_n , we have

$$\lim_{n \rightarrow \infty} \sum_{j=1}^d \widehat{\text{MDI}}_{\mathcal{T}_n}(X^{(j)}) = \mathbb{V}[m(\mathbf{X})] + \sigma^2.$$

For fully grown trees, the sum of MDI contains not only all available information $\mathbb{V}[m(\mathbf{X})]$ but also the noise in the data. This implies that the MDI of some variables are higher than expected due to the noise in the data. The noise, when having a larger variance than the regression function, can induce a very important bias in the MDI by overestimating the importance of some variables. The bias of MDI has already been empirically studied. To circumvent the overfitting problem of MDI computations (using a single data set to build the trees and estimate the variable importances), [Zhou and Hooker \[2019\]](#) use a test set to compute MDI that are less biased. A quantification of the MDI bias has been proposed recently in [Li et al. \[2019\]](#) in which they show that deep trees tend to exhibit larger MDI bias for irrelevant variables. Their analysis provides finite-sample bounds for MDI of irrelevant variables whereas our analysis is asymptotic for both relevant and irrelevant variables.

The above discussion stresses out that, when interested by interpreting a decision tree, one must never use the MDI measures output by a fully grown tree. However, if the depth of the tree is fixed and large enough, the MDI output by the tree provides a good estimation of the population MDI as shown in [Theorem 1](#).

Model 1. *The regression model satisfies $Y = m(\mathbf{X}) + \varepsilon$ where m is continuous, \mathbf{X} admits a density bounded from above and below by some positive constants and ε is an independent Gaussian noise of variance σ^2 .*

Theorem 1. *Assume [Model 1](#) holds and that for all theoretical CART tree \mathcal{T}^* , almost surely,*

$$\lim_{k \rightarrow \infty} \Delta(m, A_{\mathcal{T}_k^*}(\mathbf{X})) = 0.$$

Let \mathcal{T}_n be the CART tree based on the data set \mathcal{D}_n . Then, for all $\gamma > 0, \rho \in (0, 1]$, there exists $K \in \mathbb{N}^$ such that, for all $k > K$, for all n large enough, with probability at least $1 - \rho$,*

$$\left| \sum_{j=1}^d \widehat{\text{MDI}}_{\mathcal{T}_{n,k}}(X^{(j)}) - \mathbb{V}[m(\mathbf{X})] \right| \leq \gamma.$$

As for [Corollary 1](#), [Theorem 1](#) relies on the consistency of theoretical trees, which is essential to obtain positive results on tree interpretability. It is worth noticing that [Theorem 1](#) is not a straightforward extension of [Corollary 1](#). The proof is based on the uniform consistency of theoretical trees and combines several recent results on tree partitions to establish the consistency of empirical CART tree.

It is not possible in general to make explicit the contribution of each MDI to the sum. In other words, it is not easy to find an explicit expression of the MDI of each variable. This is due to interactions and correlation between inputs, which make difficult to distinguish the impact of each variable. Nevertheless, when the regression model can be decomposed into a sum of two independent terms, we can obtain more precise information on MDI. To this aim, let, for any set $\mathcal{J} = \{j_1, \dots, j_J\} \subset \{1, \dots, d\}$, $\mathbf{X}^{(\mathcal{J})} = (X^{(j_1)}, \dots, X^{(j_J)})$.

Model 2. There exists $\mathcal{J} \subset \{1, \dots, d\}$ such that $\mathbf{X}^{(\mathcal{J})}$ and $\mathbf{X}^{(\mathcal{J}^c)}$ are independent and

$$Y = m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) + m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)}) + \varepsilon, \quad (8)$$

where $m_{\mathcal{J}}$ and $m_{\mathcal{J}^c}$ are continuous functions, ε is a Gaussian noise $\mathcal{N}(0, \sigma^2)$ independent of \mathbf{X} and \mathbf{X} admits a density bounded from above and below by some positive constants.

Lemma 1. Assume Model 2 holds. Then, for all $j \in \mathcal{J}$, the criterion $L_A^*(j, s)$ does not depend on $m_{\mathcal{J}^c}$ and is equal to $L_{A^{(\mathcal{J})}}^*(j, s)$. Besides, any split $j \in \mathcal{J}$ leaves the variance of $m_{\mathcal{J}^c}$ unchanged.

Lemma 1 will be used to prove Proposition 3 below but is informative on its own. It states that if the regression model can be decomposed as in Model 2, the best split along coordinates in \mathcal{J} does not depend on $m_{\mathcal{J}^c}$ or on the ranges of coordinates in \mathcal{J}^c .

Proposition 3. Assume Model 2 holds. Let $\mathcal{T}_{\mathcal{J}}^*$ be a tree built on distribution $(\mathbf{X}^{(\mathcal{J})}, m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}))$ and let $\mathcal{T}_{\mathcal{J},k}^*$ be the truncation of the tree at level k . Similarly, let $\mathcal{T}_{\mathcal{J}^c}^*$ be a tree built on distribution $(\mathbf{X}^{(\mathcal{J}^c)}, m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)}))$ and let $\mathcal{T}_{\mathcal{J}^c,k}^*$ be the truncation of the tree at level k . If for any tree $\mathcal{T}_{\mathcal{J}}^*$ and $\mathcal{T}_{\mathcal{J}^c}^*$, almost surely,

$$\lim_{k \rightarrow \infty} \Delta(m_{\mathcal{J}}, A_{\mathcal{T}_{\mathcal{J},k}^*}(\mathbf{X}^{(\mathcal{J})})) = 0, \quad \text{and} \quad \lim_{k \rightarrow \infty} \Delta(m_{\mathcal{J}^c}, A_{\mathcal{T}_{\mathcal{J}^c,k}^*}(\mathbf{X}^{(\mathcal{J}^c)})) = 0$$

then, for any tree \mathcal{T}^* built on the original distribution (\mathbf{X}, Y) , almost surely,

$$\lim_{k \rightarrow \infty} \Delta(m, A_{\mathcal{T}_k^*}(\mathbf{X})) = 0.$$

Proposition 3 states that the consistency of any theoretical trees results from the consistency of all theoretical trees built on subsets (\mathcal{J} and \mathcal{J}^c) of the original set of variables, under Model 2. This is of particular significance for proving the consistency of theoretical tree for a new model which can be rewritten as a sum of independent regression functions: we can then deduce the tree consistency from existing consistency results for each regression function. Note that there are no structural assumptions on the functions $m_{\mathcal{J}}$ and $m_{\mathcal{J}^c}$ apart from the continuity. In particular, the class of functions in Model 2 is wider than additive functions.

Theorem 2 below proves that the sum of the MDI inside groups of variables that (i) do not interact between each other and (ii) are independent of each other is well-defined and equals the variance of the corresponding component of the regression function.

Theorem 2. Under the assumptions of Proposition 3, for all $\gamma > 0$, there exists K such that, for all $k > K$,

$$\left| \sum_{j \in \mathcal{J}} \text{MDI}_{\mathcal{T}_k^*}^*(X^{(j)}) - \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})})] \right| \leq \gamma,$$

and

$$\left| \sum_{j \in \mathcal{J}^c} \text{MDI}_{\mathcal{T}_k^*}^*(X^{(j)}) - \mathbb{V}[m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)})] \right| \leq \gamma.$$

Besides, for all $\gamma > 0, \rho \in (0, 1]$, there exists K such that, for all $k > K$, for all n large enough, with probability at least $1 - \rho$,

$$\left| \sum_{j \in \mathcal{J}} \widehat{MDI}_{\mathcal{T}_{n,K}}(X^{(j)}) - \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})})] \right| \leq \gamma,$$

and

$$\left| \sum_{j \in \mathcal{J}^c} \widehat{MDI}_{\mathcal{T}_{n,K}}(X^{(j)}) - \mathbb{V}[m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)})] \right| \leq \gamma.$$

Theorem 2 establishes the limiting value of the sum of MDI in subgroups \mathcal{J} and \mathcal{J}^c , both for the population version and the empirical version of the MDI, as soon as all theoretical trees are consistent. We will use Theorem 2 in the Section 4 to derive explicit expressions for the MDI of each variable, assuming a specific form of the regression function.

4 Additive Models

One class of models particularly easy to interpret is additive models (sum of univariate functions) with independent inputs defined as below.

Model 3 (Additive model). *The regression model writes*

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

where each m_j is continuous; ε is a Gaussian noise $\mathcal{N}(0, \sigma^2)$, independent of \mathbf{X} ; and $\mathbf{X} \sim \mathcal{U}([0, 1]^d)$.

Due to the model additivity, there is no interaction between variables: the contribution of each variable for predicting the output is reduced to its marginal contribution. Since we also assume that features are independent, the information contained in one variable cannot be inferred by the values of other variables.

By considering a model with no interaction and with independent inputs, we know that the contribution of a variable has an intrinsic definition which does not depend on other variables that are used to build the model. In Model 3, the explained variance of the model $\mathbb{V}[m(\mathbf{X})]$ takes the form

$$\mathbb{V}[m(\mathbf{X})] = \sum_{j=1}^d \mathbb{V}[m_j(X^{(j)})],$$

which holds only because of independent inputs *and* the absence of interactions. The variance explained by the j th variable is defined unambiguously, and independently of which variables are considered in the model, as $\mathbb{V}[m_j(X^{(j)})]$. Therefore any importance measure for $X^{(j)}$ which aims at decomposing the explained variance must output $\mathbb{V}[m_j(X^{(j)})]$. It turns out that the MDI computed via CART trees converges to this quantity.

Theorem 3 (Additive model). *Assume that Model 3 holds. Let \mathcal{T}^* be any theoretical CART tree and \mathcal{T}_n be the empirical CART tree. Then, for all $\gamma > 0$, there exists K such that, for all $k > K$, for all j ,*

$$\left| \text{MDI}_{\mathcal{T}_k^*}^*(X^{(j)}) - \mathbb{V}[m_j(X^{(j)})] \right| \leq \gamma,$$

Moreover, for all $\gamma > 0, \rho \in (0, 1]$, there exists K such that, for all $k > K$, for all n large enough, with probability at least $1 - \rho$, for all j ,

$$\left| \widehat{\text{MDI}}_{\mathcal{T}_{n,k}}(X^{(j)}) - \mathbb{V}[m_j(X^{(j)})] \right| \leq \gamma.$$

Since the MDI computed via random forests is nothing but the average of MDI computed by trees, Theorem 3 also holds for MDI computed via random forests. To the best of our knowledge, Theorem 3 is the first result highlighting that empirical MDI computed with CART procedure converge to reasonable values that can be used for variable selection, in the framework of Model 3.

Contrary to Theorem 2, Theorem 3 does not assume the consistency of the theoretical tree. Indeed, in the context of additive models, one can directly take advantage of Scornet et al. [2015] to prove the consistency of theoretical CART trees.

Surprisingly, the population version of MDI has the same expression as the population version of MDA (up to factor 2), as stated in Gregorutti et al. [2017]. Thus, in the particular context of additive models with independent features, both MDI and MDA target the same quantity, which is the natural way to decompose the total variance. In this context, MDI and MDA can then be employed to rank features and select variables based on the importance values they produce.

Note that MDI is computed with truncated trees, i.e. trees that contain a large number of observations in each node. This is mandatory to ensure that the variance of the output in the resulting cell is correctly estimated. As mentioned before, considering fully grown trees would result in positively-biased MDI, which can lead to unintended consequences for variable selection. We therefore stress that MDI must not be computed using fully grown trees.

Theorem 3 proves that MDI computed with any CART tree are asymptotically identical: the MDI are consistent across trees. The only interest to use many trees to compute the MDI instead of a single one relies on the variance reduction property of ensemble methods, which allows us to reduce the variance of the MDI estimate when using many trees.

Remark: invariance under monotonous transformations Tree-based methods are known to be invariant by strictly monotonous transformations of each input. Letting f_j being any monotonous transformation applied to variable $X^{(j)}$, the new regression model writes

$$Y = \sum_{j=1}^d m_j \circ f_j^{-1}(f_j(X^{(j)})) + \varepsilon,$$

where $f_j(X^{(j)})$ are the new input variables. Assuming that Theorem 3 can be applied in the more general setting where $X^{(j)} \sim f_j(\mathcal{U}([0, 1]))$, the MDI for this new problem satisfies

$$\lim_{k \rightarrow \infty} \text{MDI}_{\mathcal{T}_k^*}^*(X^{(j)}) = \mathbb{V}[m_j \circ f_j^{-1}(f_j(X^{(j)}))] = \mathbb{V}[m_j(X^{(j)})].$$

Therefore the asymptotic MDI is invariant by monotonic transformation, assuming that Theorem 3 holds for any marginal distribution of \mathbf{X} (with independent components).

One particularly famous instance of additive models are the well-studied and extensively used linear models.

Model 4 (Linear model). *The regression model writes*

$$Y = \sum_{j=1}^d \alpha_j X^{(j)} + \varepsilon,$$

where, for all j , $\alpha_j \in \mathbb{R}$; ε is a Gaussian noise $\mathcal{N}(0, \sigma^2)$, independent of \mathbf{X} ; and $\mathbf{X} \sim \mathcal{U}([0, 1]^d)$.

As for Model 3, there is only one way to decompose the explained variance for linear models with independent input, as stated by Nathans et al. [2012] and Grömping [2015] (and the references therein), which is $\alpha_j^2 \mathbb{V}[X^{(j)}]$. Theorem 4 below proves that MDI converge to these quantities and provides information about the theoretical tree structure in Model 4.

Theorem 4 (Linear model). *Assume that Model 4 holds. For any cell $A = \prod_{\ell=1}^d [a_\ell, b_\ell] \subset [0, 1]$, the coordinate which maximizes the population CART-split criterion (4) satisfies*

$$j^* \in \arg \max_{\ell \in \{1, \dots, d\}} \alpha_\ell^2 (b_\ell - a_\ell)^2$$

and the splitting position is the center of the cell; the associated variance reduction is

$$\frac{\alpha_{j^*}^2}{4} \left[\frac{b_{j^*} - a_{j^*}}{2} \right]^2.$$

Let \mathcal{T}^* be any theoretical CART tree and \mathcal{T}_n be the empirical CART tree. Then, for all $\gamma > 0$, there exists K such that, for all $k > K$ and for all j ,

$$\left| \text{MDI}_{\mathcal{T}_k^*}^*(X^{(j)}) - \frac{\alpha_j^2}{12} \right| \leq \gamma.$$

Moreover, for all $\gamma > 0, \rho \in (0, 1]$, there exists K such that, for all $k > K$, for all n large enough, with probability at least $1 - \rho$, for all j ,

$$\left| \widehat{\text{MDI}}_{\mathcal{T}_{n,k}}(X^{(j)}) - \frac{\alpha_j^2}{12} \right| \leq \gamma.$$

Theorem 4 establishes that, for linear models with independent inputs, the MDI computed with CART trees whose depth is limited is exactly the quantity of interest, which depends only on the magnitude of α_j^2 , since all variables have the same variance. Theorem 4 is a direct consequence of Theorem 3. However, for linear models, one can compute explicitly the best splitting location and the associated variance decreasing, which is stated in Theorem 4. Being able to compute the variance decreasing for each split allows us to compute the MDI directly, without using Theorem 3. Note that the splitting location and variance reduction has been first stated by Biau [2012] and then by Ishwaran [2013].

5 A model with interactions

In the absence of interaction and dependence between inputs, the MDI is a measure of variable importance which is independent of the set of variables included in the model, as stated in Theorem 3. However, this is not true as soon as we consider classes of models that present interactions or dependence. In this section, we study the influence of interactions via the following model.

Model 5 (Multiplicative model). *Let $\alpha \in \mathbb{R}$. The regression model writes*

$$Y = 2^d \alpha \prod_{j=1}^d X^{(j)} + \varepsilon.$$

where $\alpha \in \mathbb{R}$; ε is a Gaussian noise $\mathcal{N}(0, \sigma^2)$, independent of \mathbf{X} ; and $\mathbf{X} \sim \mathcal{U}([0, 1]^d)$.

Theorem 5 (Multiplicative model). *Assume that Model 5 holds. For any cell $A = \prod_{\ell=1}^d [a_\ell, b_\ell] \subset [0, 1]$, the coordinate which maximizes the population CART-split criterion (4) satisfies*

$$j^* \in \arg \max_{\ell \in \{1, \dots, d\}} \frac{b_\ell - a_\ell}{a_\ell + b_\ell},$$

and the splitting position is the center of the cell; the associated variance reduction is

$$\frac{\alpha^2}{4} (b_{j^*} - a_{j^*})^2 \prod_{\ell \neq j^*} (a_\ell + b_\ell)^2$$

Let \mathcal{T}^* be any theoretical CART tree and \mathcal{T}_n be the empirical CART tree. Then, for all $\gamma > 0$, there exists K such that, for all $k > K$,

$$\left| \sum_{j=1}^d \text{MDI}_{\mathcal{T}_k^*}^*(X^{(j)}) - \alpha^2 \left(\left(\frac{4}{3} \right)^d - 1 \right) \right| \leq \gamma.$$

Moreover, for all $\gamma > 0, \rho \in (0, 1]$, there exists K such that, for all $k > K$, for all n large enough, with probability at least $1 - \rho$,

$$\left| \sum_{j=1}^d \widehat{\text{MDI}}_{\mathcal{T}_{n,k}}(X^{(j)}) - \alpha^2 \left(\left(\frac{4}{3} \right)^d - 1 \right) \right| \leq \gamma.$$

Theorem 5 gives the explicit splitting position for a model with interactions. Deriving splitting positions allows us to prove that the theoretical tree is consistent which, in turns, proves that the sum of MDI converges to the explained variance $\mathbb{V}[m(\mathbf{X})]$, according to Theorem 1.

We are also interested in obtaining the exact MDI expression for each variable, as established for additive models. However, Theorem 2 no longer applies since the regression model cannot be decomposed into independent additive components. It turns out to be impossible to derive explicit expression for each MDI in this model. To see this, note that there exist many theoretical trees in Model 5. Two of them are displayed in Figure 2. They result in the same partition but the first split can be made either along $X^{(1)}$ (Figure 2, left) or $X^{(2)}$ (Figure 2, right). Surprisingly, the variable importance computed with these trees is larger for the variable that is split in

the second step. This exemplifies that splits occurring in the first levels of trees do not lead to the largest decrease in variance. Nonetheless, this holds only when interactions are present: in the case of the additive models defined in Section 4, the decrease in variance is stronger in the first tree levels. A direct consequence of this fact is that two different theoretical tree can output two different variable importance, as shown by Lemma 2 below.

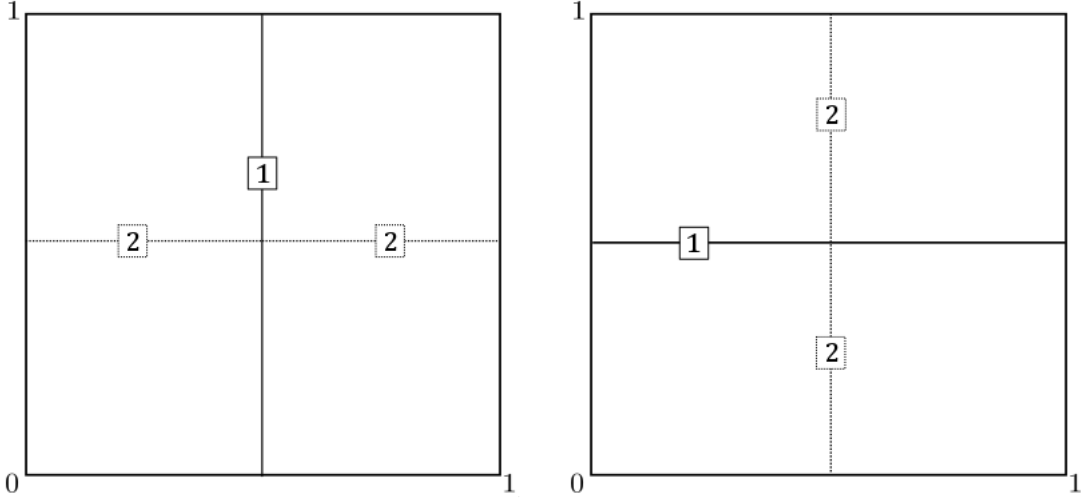


Figure 2: Two theoretical tree partitions of level $k = 2$: the first split is performed along $X^{(1)}$ (resp. $X^{(2)}$) for the tree on the left (resp. on the right).

Lemma 2. *Assume that Model 5 holds. Then, there exists two theoretical trees \mathcal{T}_1 and \mathcal{T}_2 such that*

$$\lim_{k \rightarrow \infty} \left(MDI_{\mathcal{T}_2, k}^*(X^{(1)}) - MDI_{\mathcal{T}_1, k}^*(X^{(1)}) \right) = \alpha^2/16.$$

According to Lemma 2, there exist many different theoretical trees for Model 5 and the MDI may vary when computed with different theoretical trees. This is a major difference with additive models for which each theoretical tree asymptotically output the same MDI. Since MDI are usually used to rank and select variables, the fact that each tree can output a different MDI, and therefore a different variable ranking, is a major drawback for employing this measure in presence of interactions. One way to circumvent this issue could be to compute the MDI via a random forest: the randomization of splitting variables yields different trees and the aggregation of MDI across trees provides a mean effect of the variable importance. For example, in Model 5, one can hope to obtain importance measure that are equal since Model 5 is symmetric in all variables. This is impossible with only one tree but is a reachable goal by computing the mean of MDI across many trees as done in random forests. More generally, for simple regression problems, a single tree may be sufficient to predict accurately but many shallow trees are needed to obtain accurate MDI values.

6 A model with dependent inputs

Up until now, we have focused on independent input variables. This is often an unrealistic assumption. To gain insights on the impact of input dependence on MDI,

we study a very specific model in which input features are not independent. We will prove that even in this very simple model, MDI has some severe drawbacks that need to be addressed carefully. Accordingly, we consider the simple case where the input vector $\mathbf{X} = (X^{(1)}, X^{(2)})$ is of dimension two and has a distribution, parametrized by $\beta \in \mathbb{N}$, defined as

$$\mathbf{X} \sim \mathcal{U} \left(\bigcup_{j=0}^{2^\beta-1} \left[\frac{j}{2^\beta}, \frac{j+1}{2^\beta} \right]^2 \right) = \mathcal{U}^{\otimes 2^\beta}. \quad (9)$$

The distribution of \mathbf{X} is uniform on 2^β squares on the diagonal. Examples of such distribution are displayed in Figure 3. For such distribution, the correlation between

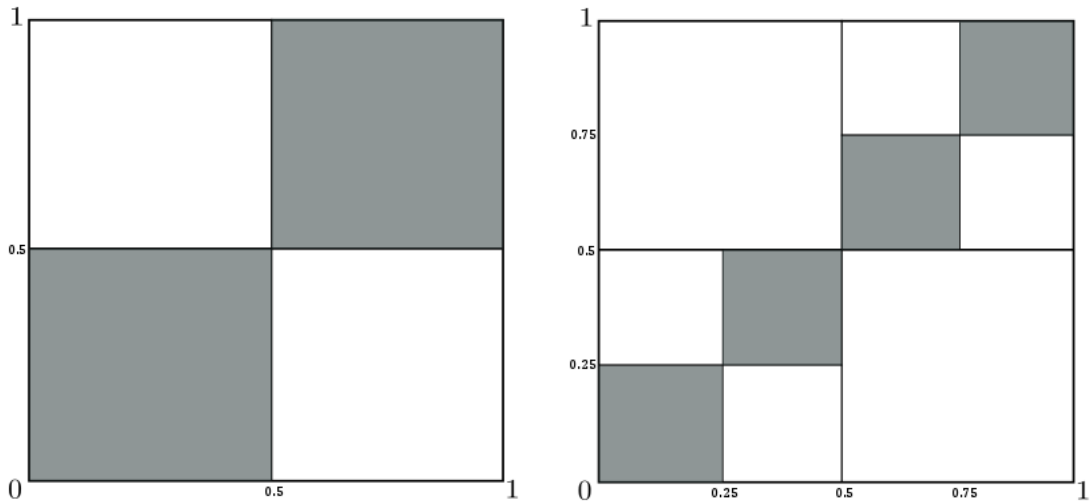


Figure 3: Illustration of distributions for \mathbf{X} , with parameter $\beta = 1$ (left) and $\beta = 2$ (right).

$X^{(1)}$ and $X^{(2)}$ is given by Lemma 3.

Lemma 3. *Let $\mathbf{X} = (X^{(1)}, X^{(2)}) \sim \mathcal{U}^{\otimes 2^\beta}$ as defined in (9). Then*

$$\text{Corr}(X^{(1)}, X^{(2)}) = 1 - \frac{1}{4^\beta}.$$

For $\beta = 0$, the distribution of \mathbf{X} is uniform over $[0, 1]^2$ and accordingly, the correlation is null between the two components of \mathbf{X} , as stated in Lemma 3. When β (*i.e.*, the number of squares) increases, that is the size of each square decreases, the distribution concentrates on the line $X^{(1)} = X^{(2)}$. Therefore, when β tends to infinity, the correlation should tend to 1, which is proved in Lemma 3.

The distribution defined in (9) is rather unusual, and one can wonder why not considering Gaussian distributions instead. This choice is related to the thresholding nature of decision trees. Since we want to compute explicitly the splitting criterion along each coordinate, we must have a closed expression for the truncated marginals (restriction to any rectangle of $[0, 1]^2$), which is not possible with Gaussian distribution. The distribution defined in (9) allows us to compute easily both truncated marginals and the joint distribution.

Model 6. *Let $\beta \in \mathbb{N}$. Assume that*

$$Y = X^{(1)} + X^{(2)} + \alpha X^{(3)} + \varepsilon,$$

where $(X^{(1)}, X^{(2)}) \sim \mathcal{U}^{\otimes 2^\beta}$, $X^{(3)} \sim \mathcal{U}([0, 1])$ is independent of $(X^{(1)}, X^{(2)})$, and ε is an independent noise distributed as $\mathcal{N}(0, \sigma^2)$.

In the framework of Model 6, Proposition 4 below states that the CART-split criterion has an explicit expression and highlights that splits along positively correlated variables ($X^{(1)}$ and $X^{(2)}$) are more likely to occur compared to splits along independent ones ($X^{(3)}$). Even if the model is very simple, it is the first theoretical proof that CART splitting procedure tends to favor positively correlated variables. Note however that considering negatively correlated variables will result in an opposite effect, *i.e.*, a tendency to favour independent variables compared to negatively correlated ones.

Proposition 4. *Assume that Model 6 holds. Then, the following statements hold:*

- (i) *The splitting criterion has an explicit expression. For $\beta = 0, \dots, 5$, each split is performed at the center of the support of the chosen variable.*
- (ii) *For $\beta = 0, \dots, 5$, the split in the root node $[0, 1]^3$ is performed along $X^{(1)}$ or $X^{(2)}$ if, and only if, $|\alpha| \leq 2$. In that case, the variance reduction is equal to $1/4$.*

Statement (ii) in Proposition 4 proves that a positive correlation between two variables increases the probability to split along one of these two variables. Indeed, in our particular model, the marginal effect of $X^{(3)}$ on the output must be at least twice the effect of $X^{(1)}$ or $X^{(2)}$ in order to split along $X^{(3)}$. It is likely that the way the correlation impact the variable selection in trees depends both on the sign of the correlation and on the signs of coefficients in the linear model. Nevertheless, Proposition 4 proves that input dependence has an influence on variable selection in CART procedures. Theorem 6 provides the limiting MDI values given by theoretical and empirical CART trees, in Model 6.

Theorem 6. *Let $\beta \in \{0, \dots, 5\}$. Assume that Model 6 holds. Let \mathcal{T}^* be any theoretical CART tree and \mathcal{T}_n be the empirical CART tree. Then, for all $\gamma > 0$, there exists K such that, for all $k > K$,*

$$\left| \text{MDI}_{\mathcal{T}_k^*}^*(X^{(1)}) + \text{MDI}_{\mathcal{T}_k^*}^*(X^{(2)}) - \mathbb{V}[X^{(1)} + X^{(2)}] \right| \leq \gamma.$$

and

$$\left| \text{MDI}_{\mathcal{T}_k^*}^*(X^{(3)}) - \alpha^2 \mathbb{V}[X^{(3)}] \right| \leq \gamma.$$

Additionally, for all $\gamma > 0, \rho \in (0, 1]$, there exists K such that, for all $k > K$, for all n large enough, with probability at least $1 - \rho$,

$$\left| \widehat{\text{MDI}}_{\mathcal{T}_{n,k}}(X^{(1)}) + \widehat{\text{MDI}}_{\mathcal{T}_{n,k}}(X^{(2)}) - \mathbb{V}[X^{(1)} + X^{(2)}] \right| \leq \gamma.$$

and

$$\left| \widehat{\text{MDI}}_{\mathcal{T}_{n,k}}(X^{(3)}) - \alpha^2 \mathbb{V}[X^{(3)}] \right| \leq \gamma.$$

Theorem 6 gives the limiting values for the MDI. Unfortunately, since $X^{(1)}$ and $X^{(2)}$ are correlated, it is only possible to have information on the sum of the two MDI of $X^{(1)}$ and $X^{(2)}$ rather than on each one of them. According to Lemma 3, a simple calculation shows that the sum of importances of $X^{(1)}$ and $X^{(2)}$ satisfies,

$$\begin{aligned} \mathbb{V}[X^{(1)} + X^{(2)}] &= 2\mathbb{V}[X^{(1)}] + 2\mathbb{V}[X^{(1)}] \left(1 - \frac{1}{2^{2\beta}}\right) \\ &> \mathbb{V}[X^{(1)}] + \mathbb{V}[X^{(2)}]. \end{aligned}$$

Therefore, the group of variables $X^{(1)}$ and $X^{(2)}$ has a larger importance because of their positive correlation. In this case, the amount of information provided by the two variables is larger than that provided by each one of them. This is a very important difference compared to the independent additive case, in which the sum of contributions of each variable is equal to the contributions of all variables. Here, even without interaction effects, this property does not hold, which leads to larger MDI for variables that are positively correlated. As mentioned before, the opposite effect would hold if the variables had a negative correlation or if they had coefficients of opposite signs in the linear model. The impact of correlation has been thoroughly investigated empirically for MDI Nicodemus and Malley [2009] and MDA Strobl et al. [2008] but we have theoretically shown above under Model 6 that correlation changes the MDI.

Even if the limiting value $\text{MDI}_{\mathcal{T}_K}^*(X^{(1)}) + \text{MDI}_{\mathcal{T}_K}^*(X^{(2)})$ is known, we have no information on how this quantity is split between variable $X^{(1)}$ and $X^{(2)}$ inside the tree \mathcal{T}_K^* . Indeed, we can find two theoretical trees which produce very different MDI as stated in Lemma 4.

Lemma 4. *Let $\beta \in \{0, \dots, 5\}$. Assume that Model 6 holds. Then, there exists two theoretical trees \mathcal{T}_1 and \mathcal{T}_2 such that*

$$\lim_{k \rightarrow \infty} \left(\text{MDI}_{\mathcal{T}_{2,k}}^*(X^{(1)}) - \text{MDI}_{\mathcal{T}_{1,k}}^*(X^{(1)}) \right) = \frac{1}{3} - \frac{1}{3} \left(\frac{1}{4} \right)^\beta.$$

Even in the case of a simple linear model with correlated input, it is not wise to use the importance of a single tree to measure the impact of this variable on the output, since this measure can vary drastically between two different trees. Instead, one should rather use an average of this measure across many shallow trees, hoping that randomizing the eligible variables for splitting ensure enough tree diversity in the forest, which in turn, leads to more balanced variable importances.

The proof of Lemma 4 relies on exhibiting a tree whose early splits are made along $X^{(1)}$ only, until \mathbf{X} is uniformly distributed in each cell. For this tree, the importance of $X^{(1)}$ is larger than the importance of $X^{(2)}$. Since $X^{(1)}$ and $X^{(2)}$ have symmetric roles, one can also build a tree whose early splits are made along $X^{(2)}$ only. The difference of importance for $X^{(1)}$ between these two trees can be computed exactly, thanks to statement (i) and (ii) in Proposition 4.

7 Experiments

Now, we illustrate numerically the theoretical results of the previous sections. We start by defining the general framework of the experiments. We then compare the sum of MDI and the individual MDI computed with random forests and single trees respectively. Since variable importances are often use to rank variables, we also compare the rankings provided by MDI computed with random forests and individual trees.

7.1 Experimental setting

In all experiments, we use the R package `randomForests` to train random forests and individual trees. For random forests, we use the default parameter values, except for the number of trees which is arbitrary set to 100 (instead of 500) to reduce computations. To train a tree, we also use the package `randomForests`, for which we set the number of trees to one, the number of eligible features for splitting to d , with no subsampling. For both random forests and trees, we use the default value for the parameter `nodesize`, which stands for the minimal number of observations for a cell to be split. The parameter `maxnodes`, the maximum number of leaves, is used to stop the tree construction (we call the resulting estimates shallow trees/forests) or to grow the tree at their maximum (the resulting estimates are named fully-grown trees/forests).

Remark 1. *Note that fully-grown trees/forests are composed of leaves containing between 1 and 5 observations, and thus the error of the corresponding tree on the training set may be different from zero. This departs slightly from the framework discussed after Proposition 2 in which the leaves of a tree contain exactly one observation. However, considering such estimates leads to similar results as the one shown below. We thus prefer to consider the default parameter of random forests.*

In all experiments, we consider the following simple linear model because additive models with independent inputs are among the few models for which the importance of a variable can be defined unambiguously.

Simulated Model 1 We let $Y = \alpha_1 X^{(1)} + \alpha_2 X^{(2)} + \alpha_3 X^{(3)} + \varepsilon$, where ε is an independent noise, distributed as $\mathcal{N}(0, \sigma^2)$ and $(X^{(1)}, X^{(2)}, X^{(3)})$ is distributed as $\mathcal{N}(0, \Sigma)$ where

$$\Sigma = \begin{pmatrix} 1 & \rho & 0 \\ \rho & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

For all j , we let $\alpha_j = \sqrt{j}$. Hence, the variable importance of the j -th component is j (for $\rho = 0$) and the difference between two consecutive variable importance is constant, and equal to one. Studying the impact of the noise σ^2 is easier in this setting.

We consider different sample sizes in the experiments. For each sample size, we generate 10 data sets distributed as above. On each data set, we build one tree and one forest (using the parameters described above). For both trees and forests respectively, we compute the mean of the quantity of interest (sum of MDI, MDI of the first variable) and its standard deviation over the 10 data sets. These means and standard deviations are displayed in the following Figures. When available, the theoretical counterpart of the quantity of interest is displayed with a black horizontal line.

7.2 Estimating the strength of the signal via the sum of MDI

Here, we use both individual trees and random forests to compute the sum of MDI in Simulated Model 1. The sum of MDI can be a good estimate of the signal strength $\mathbb{V}[m(\mathbf{x})]$, if computed with trees and forests whose construction is stopped early, as stated in Theorem 1.

Results are displayed in Figure 4 and 5. We observe that the variance of the sum of MDI computed with random forests is of the same order as the one computed with individual trees. We do not observe any variance reduction phenomenon, probably due to the fact that all theoretical trees target the same MDI values (for each variable respectively, as established in Theorem 3), and consequently there may not exist enough diversity in trees for the averaging effect of random forest to be noticeable.

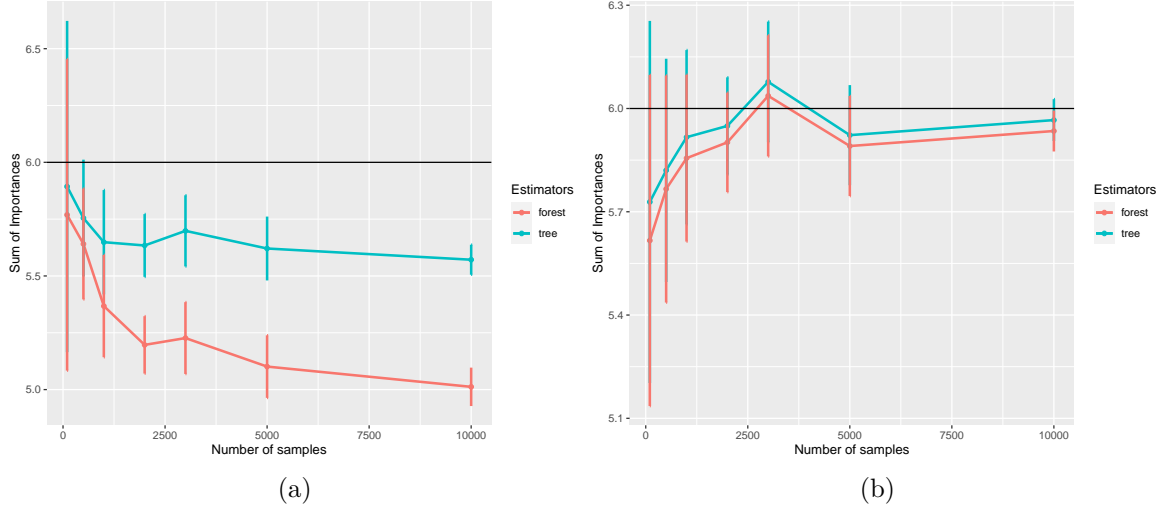


Figure 4: Sum of importances in Simulated Model 1, with $\sigma^2 = 0, \rho = 0$ and, from left to right $\text{maxnodes} = \lfloor n^{0.6} \rfloor, n$

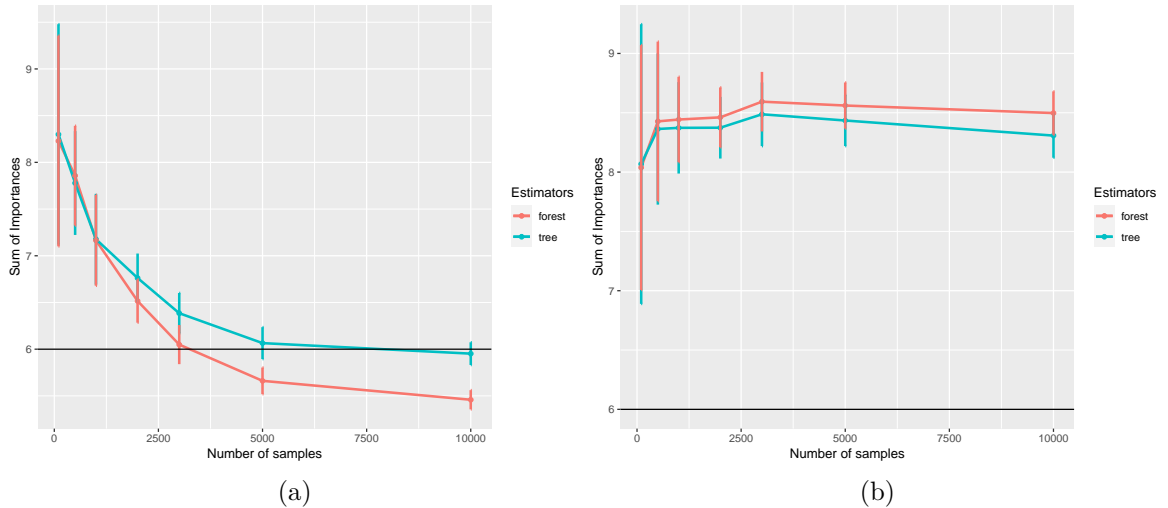


Figure 5: Sum of importances in Simulated Model 1, with $\sigma^2 = 3, \rho = 0$ and, from left to right $\text{maxnodes} = \lfloor n^{0.6} \rfloor, n$

Results for fully grown forests are consistent with Theorem 1. When the noise variance is null, the sum of MDI computed with fully-grown forests is a consistent estimate of the signal strength (horizontal black line, Figure 4b). However, when the noise variance is important, the sum of MDI exhibits a positive bias: it estimates a quantity close to $\mathbb{V}[Y] = \mathbb{V}[m(\mathbf{X})] + \sigma^2$ instead of $\mathbb{V}[m(\mathbf{X})]$ (Figure 5b). Since the

parameter `nodesize` is set to 5, the sum of MDI does not account for all the noise variance, yielding a difference between the estimated values and the true value (see Figure 5b) smaller than 3.

Consequently, fully-grown forests should not be used to estimate the signal strength in presence of noise, which is a current setting in many applications. Building a consistent estimate of the signal strength can be done by stopping the tree construction early (here when the number of terminal nodes is equal to $\lfloor n^{0.6} \rfloor$). Results are depicted in Figure 4a and 5a. In Figure 4a, the sum of MDI computed with trees exhibits a negative bias (due to the early-stopped construction) which tends very slowly to zero with the sample size (not visible in Figure 4a). The convergence is more obvious in Figure 5a. Thus, in presence of noise, one may prefer early-stopped trees/forests instead of fully-grown forests to estimate the sum of MDI.

Quite surprisingly, we notice a strong difference between early-stopped trees and forests. This difference is due to the random choice of eligible features for splitting in each cell in random forests. We run another set of experiments setting the parameter `mtry` = d in random forests and we observe that the two curves are closer to each other, which may indicate that this randomization is the most important effect justifying the differences between the two curves (see Figure 12 and 13 in Section 8). One possible explanation is as follows. The randomness introduced in the splitting strategy may create splits that do not reduce much the variance. The overall random tree may thus not be able to create homogeneous terminal nodes, given the limited number of cuts available. Indeed, since the random tree is not fully-grown, it may not be possible to split the cells enough times to decrease enough the variance in terminal nodes. Then the variance may be large in leaves, which leads to underestimating the sum of MDI.

7.3 Convergence of individual MDI

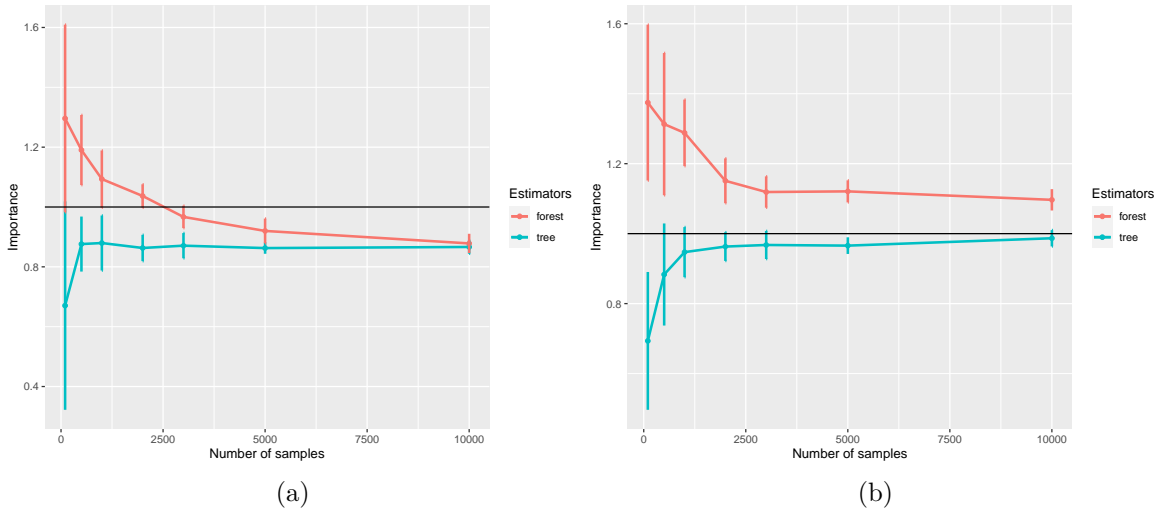


Figure 6: Importance of the first variable in Simulated Model 1, with $\sigma^2 = 0, \rho = 0$ and, from left to right `maxnodes` = $\lfloor n^{0.6} \rfloor, n$

Now, we turn to the problem of estimating the first variable importance via MDI. Fully-grown forests provide good estimation of the variable importance of the first variable in absence of noise (Figure 6b). However, in presence of noise, the MDI

output by a fully grown forest is larger than the theoretical value (Figure 7b). The difference between the estimation and the true value is approximately $\sigma^2/3 = 1$, which seems to suggest that the additional noise variance is split equally between the three variables.

According to Theorem 3, each individual tree outputs asymptotically the same MDI. Thus, the estimation of MDI given by early-stopped trees and forests should tend to the black line with a variance that tends to zero. This is illustrated in Figures 6a and 7a where we notice that the convergence toward the theoretical importance (black line) is very slow in absence of noise. We observe an important difference between trees and forests which is related to the randomization splitting process. Indeed, we run the same experiments setting `mtry = d` and we obtain similar curves for trees and forests (see Figures 14 and 15). Taking into account results with and without noise, we conclude that using early-stopped trees and forests should be preferred to using fully-grown trees and forests. Indeed, in a noise setting, the difference between the estimation given by fully-grown trees/forests and the true target is proportional to the noise variance which can be very large.

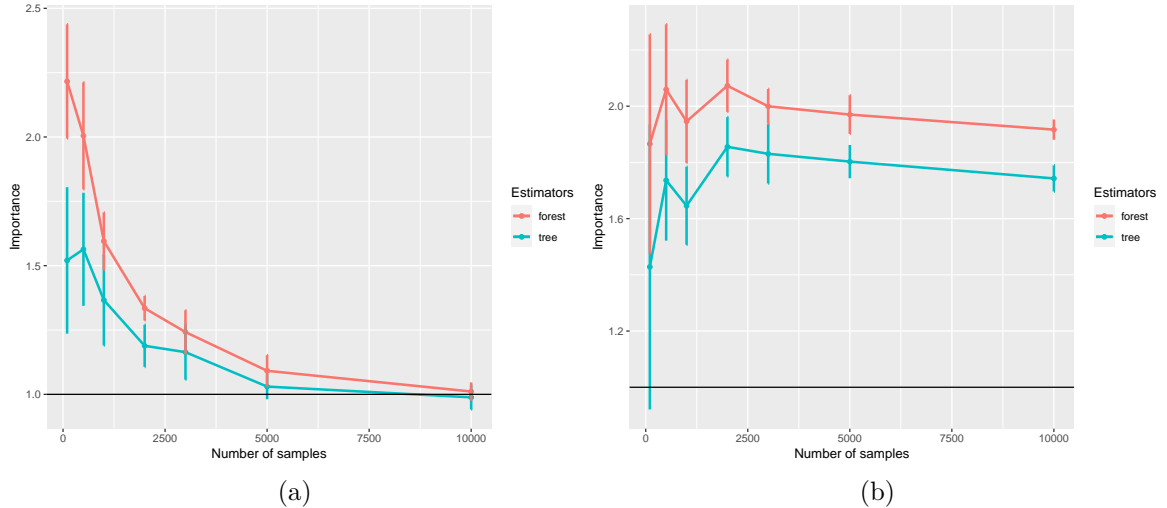


Figure 7: Importance of the first variable in Simulated Model 1, with $\sigma^2 = 3, \rho = 0$ and, from left to right `maxnodes = $\lfloor n^{0.6} \rfloor, n$`

7.4 Impact of correlation on MDI

Now, we study the impact of correlated inputs on the MDI. First, note that the theoretical value for variable importance is not clearly defined in presence of correlation. We choose to display the value $\alpha_1^2 + \alpha_1\alpha_2\rho$ in the horizontal black line. This value corresponds to the marginal importance of the first variable plus half of the interaction term between the first and the second variable (i.e. the interactions are split equally between first and second variables). This theoretical value depends on the correlation ρ .

In Figures 8 and 9, we observe that the standard deviation of MDI computed with individual trees increases compared to the setting with no correlation. Indeed, in presence of correlation, splitting along one variable gives information on another

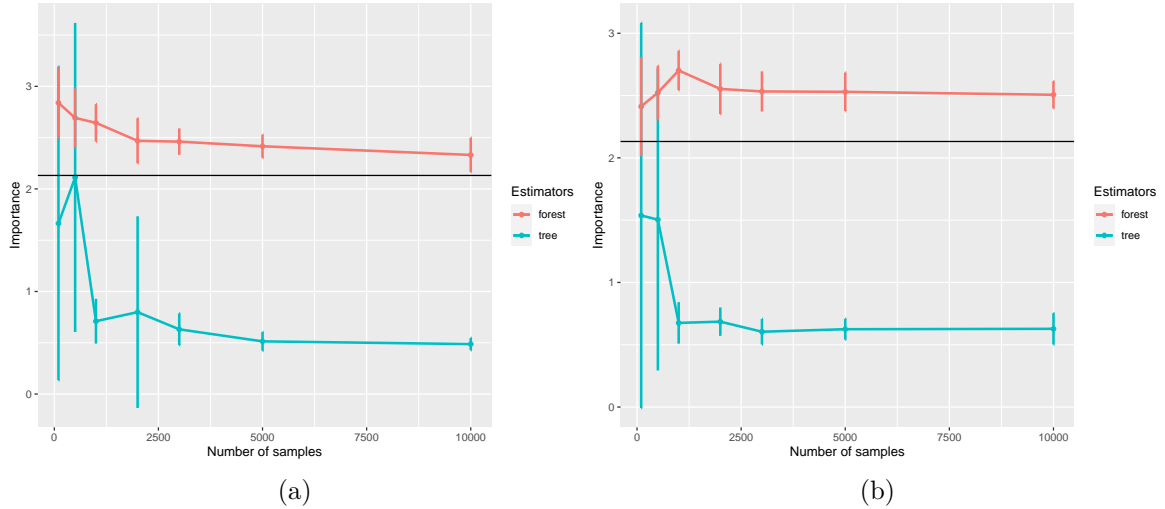


Figure 8: Importance of the first variable in Simulated Model 1, with $\sigma^2 = 0$, $\rho = 0.8$ and, from left to right $\text{maxnodes} = \lfloor n^{0.6} \rfloor, n$

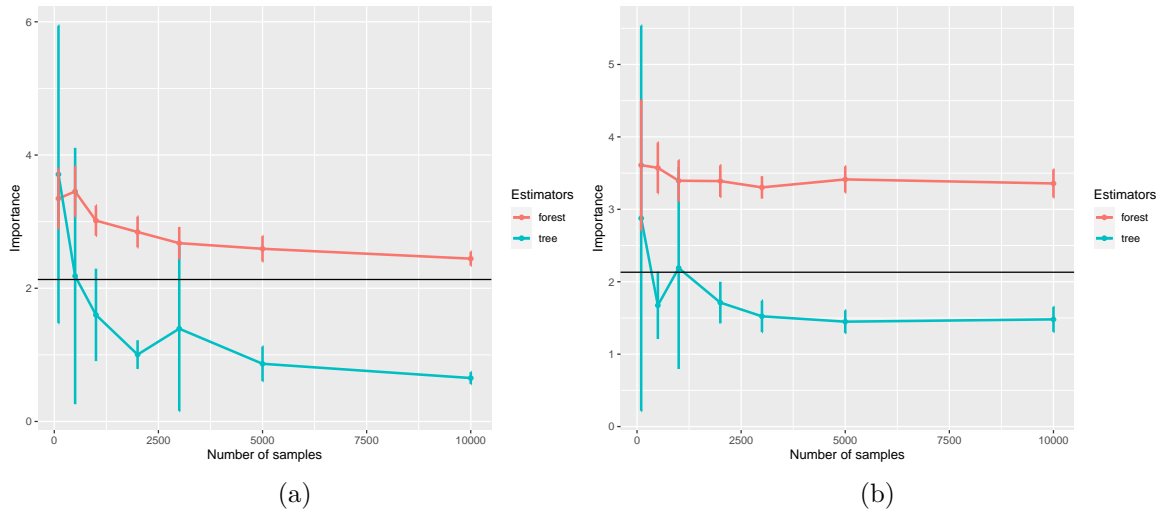


Figure 9: Importance of the first variable in Simulated Model 1, with $\sigma^2 = 3$, $\rho = 0.8$ and, from left to right $\text{maxnodes} = \lfloor n^{0.6} \rfloor, n$

variable so there may exist many sequences of splits (i.e. many trees) that are optimal, and each one may produce a different MDI value.

In this correlated setting, the MDI produced by single trees are smaller than 1 (1 is the marginal importance of the first variable). One possible explanation is that the non-randomized splitting procedure encourages the splits to be performed along the second variable, whose regression coefficient is larger. Therefore, a large part of the importance of the first variable may be attributed to the second, thus automatically reducing the importance of the first variable. Random forests produce larger MDI possibly because of the randomized splitting mechanism which forces some of the splits to be performed along the first variable. Here again, we see that using early-stopped forests is preferable to the use of fully-grown structures.

7.5 Order given by variable importance

Variable importances are often used to rank variables. In this last experiment, we study the ranking performance based on MDI measures computed with trees and forests. For each sample size, we generate 100 data sets and for each data set, we compare the ranking output by MDI with the true ranking, which is given by the following ordering of the input variables (3, 2, 1) (since the regression coefficients are respectively $\sqrt{3}, \sqrt{2}, 1$). The percentage of correct rankings is given in Figure 10 and 11.

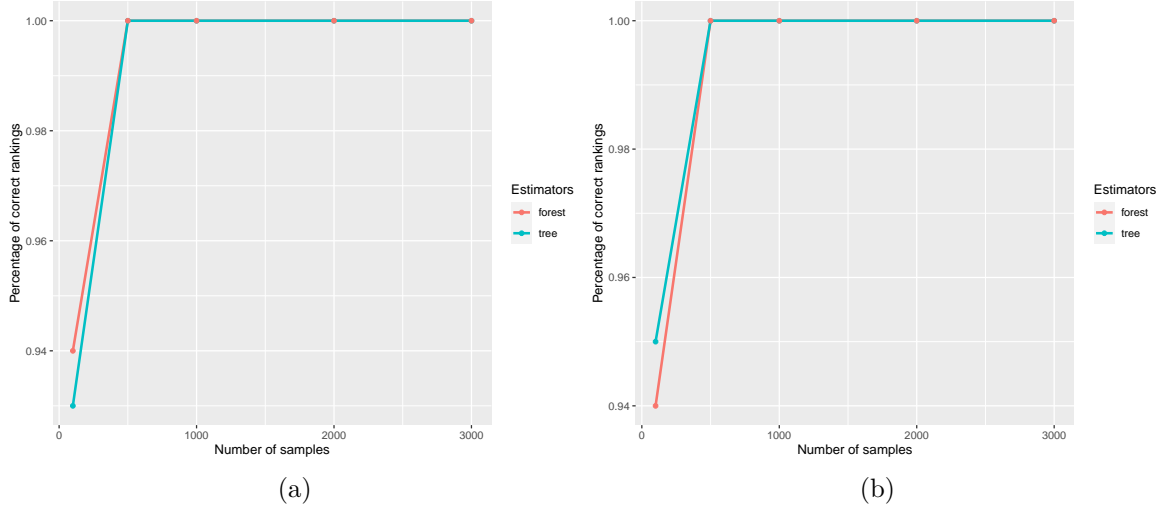


Figure 10: Percent of correct ranking in Simulated Model 1, with $\sigma^2 = 0$ and, from left to right $\text{maxnodes} = \lfloor n^{0.6} \rfloor, n$

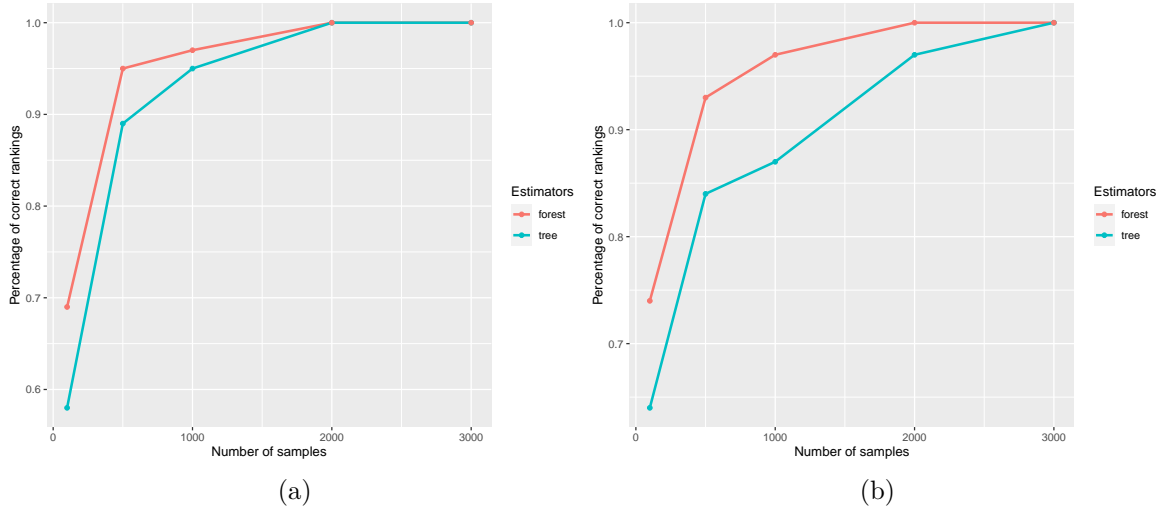


Figure 11: Percent of correct ranking in Simulated Model 1, with $\sigma^2 = 12$ and, from left to right $\text{maxnodes} = \lfloor n^{0.6} \rfloor, n$

One striking comment is that even if the sum of MDI or the MDI itself does not estimate accurately the targeted theoretical counterparts, the overall rankings produced by either trees or forests are surprisingly accurate. Increasing the noise

variance allows us to notice the more accurate rankings produced by random forests compared to single trees (see Figure 11). Note also that, quite logically, increasing the noise variance, leads to decreasing ranking accuracy for both trees and forests.

8 Additional Figures

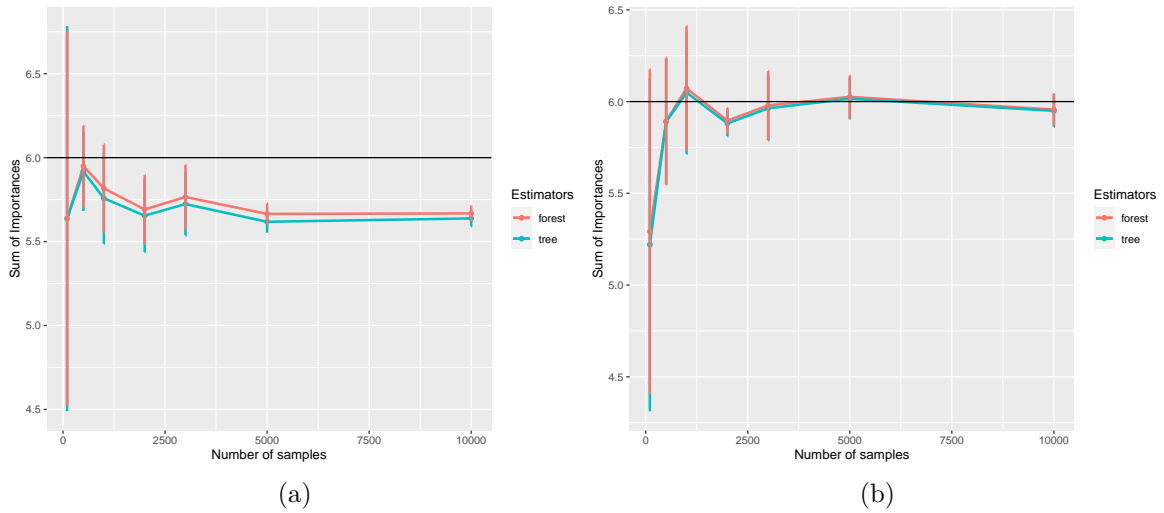


Figure 12: Sum of importances in Simulated Model 1, with $\sigma^2 = 0, \rho = 0$ and, from left to right $\text{maxnodes} = \lfloor n^{0.6} \rfloor, n$, with $\text{mtry} = d$ for trees and forests

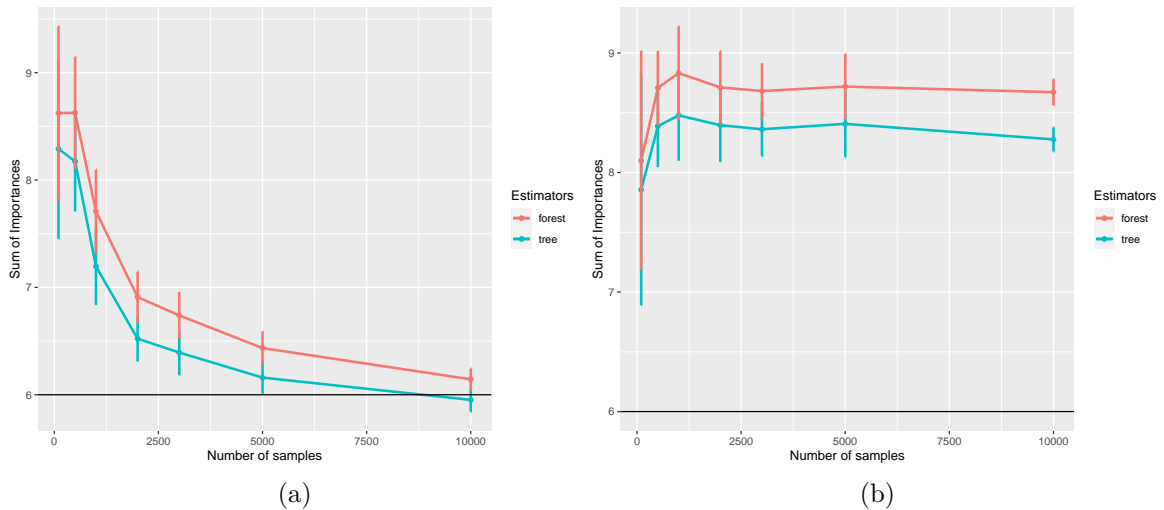


Figure 13: Sum of importances in Simulated Model 1, with $\sigma^2 = 3, \rho = 0$ and, from left to right $\text{maxnodes} = \lfloor n^{0.6} \rfloor, n$, with $\text{mtry} = d$ for trees and forests

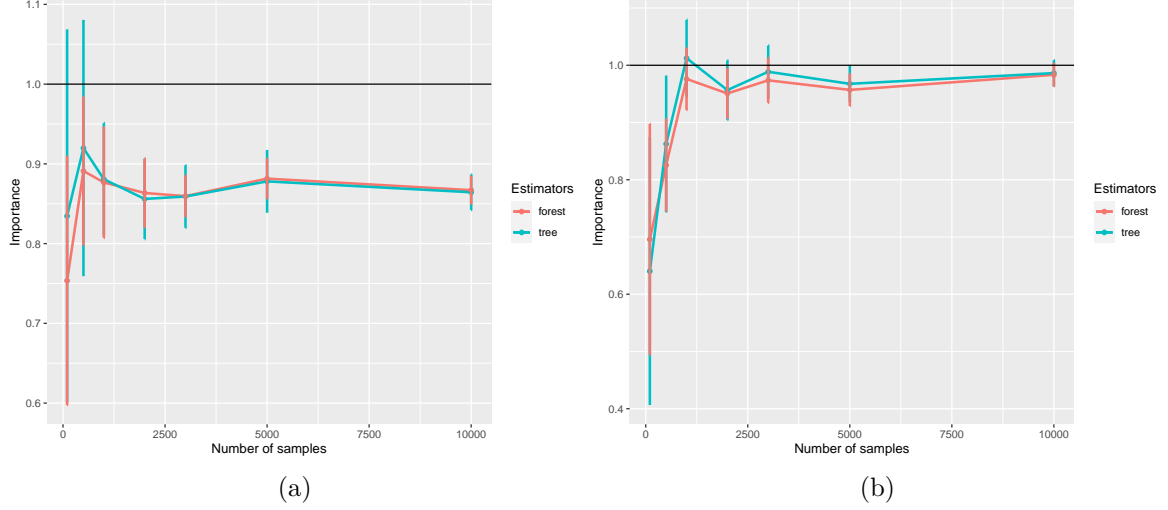


Figure 14: Importance of the first variable in Simulated Model 1, with $\sigma^2 = 0, \rho = 0$ and, from left to right $\text{maxnodes} = \lfloor n^{0.6} \rfloor, n$, with $\text{mtry} = d$ for trees and forests

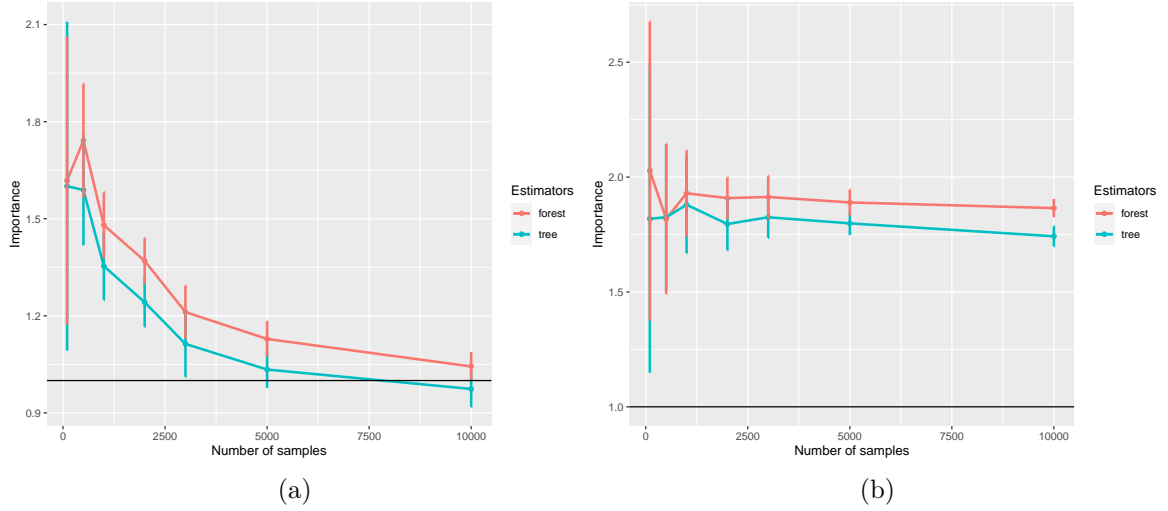


Figure 15: Importance of the first variable in Simulated Model 1, with $\sigma^2 = 3, \rho = 0$ and, from left to right $\text{maxnodes} = \lfloor n^{0.6} \rfloor, n$, with $\text{mtry} = d$ for trees and forests

9 Proofs

Proof of Proposition 1. Let \mathcal{T}^* be a theoretical CART tree and define, for all $k \in \mathbb{N}$, the variance u_k of the tree \mathcal{T}_k^* (the truncation of \mathcal{T} at level k) as

$$u_k = \sum_{\ell=1}^{n_{\text{cell},k}} \mathbb{P}[\mathbf{X} \in A_{\ell,k}] \mathbb{V}[Y|\mathbf{X} \in A_{\ell,k}],$$

where $n_{\text{cell},k}$ stands for the number of terminal nodes in \mathcal{T}_k^* and $\{A_{\ell,k}, \ell = 1, \dots, n_{\text{cell},k}\}$ is the set of terminal nodes in \mathcal{T}_k^* . Let $f_k(\mathbf{x}) = \mathbb{V}[Y|\mathbf{X}' \in A_k(\mathbf{x})]$. By definition of the impurity measure, note that, for all k , the quantity $u_k - u_{k+1}$ corresponds to the reduction of impurity between level k and level $k+1$. Therefore, the impurity measure

rewrites, for any $K \in \mathbb{N}$ as

$$\begin{aligned} \sum_{j=1}^d \text{MDI}_{\mathcal{T}_K^*}^*(X^{(j)}) &= \sum_{k=0}^{K-1} (u_k - u_{k+1}) \\ &= u_0 - u_K, \end{aligned} \tag{10}$$

where $u_0 = \mathbb{V}[Y]$ by definition. Additionally,

$$\begin{aligned} \mathbb{E}[f_k(\mathbf{X})] &= \mathbb{E} \left[\sum_{\ell=1}^{n_{\text{cell},k}} \mathbb{V}[Y|\mathbf{X}' \in A_{\ell,k}] \mathbb{1}_{\mathbf{X} \in A_{\ell,k}} \right] \\ &= \sum_{\ell=1}^{n_{\text{cell},k}} \mathbb{P}[\mathbf{X} \in A_{\ell,k}] \mathbb{V}[Y|\mathbf{X}' \in A_{\ell,k}] \\ &= u_k, \end{aligned}$$

yielding the desire conclusion

$$\sum_{j=1}^d \text{MDI}_{\mathcal{T}_K^*}^*(X^{(j)}) = \mathbb{V}[Y] - \mathbb{E}[f_K(\mathbf{X})]. \tag{11}$$

□

Proof of Corollary 1. Using notations of the proof of Proposition 1, for all $\mathbf{x} \in [0, 1]^d$, $f_k(\mathbf{x}) \leq 4\|m\|_\infty^2 + \sigma^2$ and, by assumption, almost surely, $\lim_{k \rightarrow \infty} f_k(\mathbf{X}) = \sigma^2$. According to the dominated convergence theorem,

$$\lim_{k \rightarrow \infty} u_k = \lim_{k \rightarrow \infty} \mathbb{E}[f_k(\mathbf{X})] = \sigma^2.$$

Therefore,

$$\lim_{K \rightarrow \infty} \sum_{j=1}^d \text{MDI}_{\mathcal{T}_K^*}^*(X^{(j)}) = \mathbb{V}[Y] - \sigma^2. \tag{12}$$

□

Proof of Proposition 2. Equation (11) in the proof of Proposition 1, written with empirical quantities, leads directly to the first statement. □

Proof of Theorem 1.

Uniform convergence of theoretical trees. Note that a tree is nothing but a sequence of splits. Therefore, any tree \mathcal{T} can be seen as a sequence (v_n) such that

- (v_1, \dots, v_d) represents the first split of the tree defined, for a split occurring along variable j at position s , as

$$\begin{cases} v_\ell = 0 & \text{for } \ell \neq j \\ v_j = \frac{s}{d^4} \end{cases}$$

- Each block $(v_{md+1}, \dots, v_{(m+1)d})$ represents the m th split of the tree occurring along variable j at position s defined as

$$\begin{cases} v_\ell = 0 & \text{for } \ell \neq j \\ v_j = \frac{s}{(d(m+1))^4} \end{cases}$$

- Splits are ordered in some arbitrary order.
- In order to obtain a complete tree, we associate with all nodes whose construction is stopped, the split $(0, \dots, 0)$.

Since the split position belongs to $[0, 1]$, we have $\sum_{\ell=1}^{\infty} v_{\ell}^2 < \infty$. Besides,

$$\sum_{\ell=1}^{\infty} \ell^2 v_{\ell}^2 \leq \sum_{\ell=1}^{\infty} \frac{1}{\ell^2} = \frac{\pi^2}{6}.$$

For any tree \mathcal{T} represented by the sequence (v_n) , we have

$$(v_n) \in B(0, \pi^2/6) = \left\{ v, \sum_{\ell=1}^{\infty} \ell^2 v_{\ell}^2 \leq \pi^2/6 \right\},$$

with $B(0, \pi^2/6)$ being a compact of $\ell^2(\mathbb{N})$. Now, let $\gamma > 0$ and

$$\mathcal{A}_{k,\gamma} = \left\{ \mathcal{T}^*, \mathbb{E}[\Delta(m, A_{\mathcal{T}_k^*}(\mathbf{X}))] \geq \gamma \right\}$$

Set $\mathcal{A}_{\infty,\gamma} = \bigcap_{k=0}^{\infty} \mathcal{A}_{k,\gamma}$. Since the population CART-split criterion L^* , defined in (4) is non negative, we have

$$\mathbb{E}[\Delta(m, A_{\mathcal{T}_{k+1}}(\mathbf{X}))] \leq \mathbb{E}[\Delta(m, A_{\mathcal{T}_k}(\mathbf{X}))].$$

Therefore, for all k , $\mathcal{A}_{k+1,\gamma} \subset \mathcal{A}_{k,\gamma}$. Because of the uniform continuity of the splitting criterion (continuous on a compact), for all $k \in \mathbb{N}$, the sets $\mathcal{A}_{k,\gamma}$ are closed.

Assume that for all k , $\mathcal{A}_{k,\gamma} \neq \emptyset$. We know that $B(0, \pi^2/6)$ is compact and that, for all k , $\mathcal{A}_{k,\gamma}$ is a closed subset of $B(0, \pi^2/6)$. Then, according to Cantor's intersection theorem $\mathcal{A}_{\infty,\gamma} \neq \emptyset$. Consequently, there exists $\mathcal{T}^* \in \mathcal{A}_{\infty,\gamma}$. By definition of $\mathcal{A}_{\infty,\gamma}$, for all k ,

$$\mathbb{E}[\Delta(m, A_{\mathcal{T}_k^*}(\mathbf{X}))] \geq \gamma. \quad (13)$$

Again, by definition of $\mathcal{A}_{\infty,\gamma}$, \mathcal{T}^* is a theoretical tree, which is consistent by assumption. Consequently, there exists K such that, for all $k > K$,

$$\mathbb{E}[\Delta(m, A_{\mathcal{T}_k^*}(\mathbf{X}))] \leq \gamma,$$

which is absurd, according to equation (13). Therefore, there exists $K > 0$ such that $\mathcal{A}_{K,\gamma} = \emptyset$. All in all, we proved that, for all $\gamma > 0$, there exists $K > 0$ such that, for all theoretical tree \mathcal{T}^* , for all $k > K$,

$$\mathbb{E}[\Delta(m, A_{\mathcal{T}_k}(\mathbf{X}))] < \gamma. \quad (14)$$

Contribution of cells with small measure. Fix γ and take K such that inequality (14) holds. For any theoretical tree \mathcal{T}_K^* of level K , the MDI associated with this tree is defined as

$$\text{MDI}_{\mathcal{T}_K^*}^* = \sum_{A \in \mathcal{T}_K^*} \mu(A) L_A^*,$$

where the sum ranges for any node A in the tree, $\mu(A)$ is the measure of the cell A defined as

$$\mu(A) = \int_{x \in A} f(\mathbf{x}) d\mathbf{x},$$

where f is the density of \mathbf{X} and L_A^* the value of the population CART-split criterion at the best split of the cell A . Similarly, we let

$$\widehat{\text{MDI}}_{\mathcal{T}_{n,K}} = \sum_{A \in \mathcal{T}_{n,K}} \mu_n(A) L_{n,A},$$

where $\mathcal{T}_{n,K}$ is the empirical CART tree truncated at level K , A a node of the tree, $\mu_n(A)$ the empirical measure of A and $L_{n,A}$ the empirical splitting criterion in cell A . Let $\text{MDI}_{\mathcal{T}_K^*, \gamma}^*$ and $\widehat{\text{MDI}}_{\mathcal{T}_{n,K}, \gamma}$ be computed by removing the variance (empirical or population version) associated to cells A for which $\mu(A) < \gamma$. Define also $\text{MDI}_{\mathcal{T}_{n,K}, \gamma}^*$ as the population version of the MDI but computed with the empirical partition $\mathcal{T}_{n,K}$ where the cells A for which $\mu(A) < \gamma$ have been removed. Now, we want to control

$$\begin{aligned} \inf_{\mathcal{T}^*} |\widehat{\text{MDI}}_{\mathcal{T}_{n,K}} - \text{MDI}_{\mathcal{T}_K^*}^*| &\leq |\widehat{\text{MDI}}_{\mathcal{T}_{n,K}} - \widehat{\text{MDI}}_{\mathcal{T}_{n,K}, \gamma}| + |\widehat{\text{MDI}}_{\mathcal{T}_{n,K}, \gamma} - \text{MDI}_{\mathcal{T}_{n,K}, \gamma}^*| \\ &\quad + \inf_{\mathcal{T}^*} |\text{MDI}_{\mathcal{T}_{n,K}, \gamma}^* - \text{MDI}_{\mathcal{T}_K^*, \gamma}^*| + \inf_{\mathcal{T}^*} |\text{MDI}_{\mathcal{T}_K^*, \gamma}^* - \text{MDI}_{\mathcal{T}_K^*}^*|. \end{aligned} \quad (15)$$

Since \mathbf{X} admits a density on $[0, 1]^d$, the population CART-split criterion varies continuously as a function of the cell on which it is computed. Thus, the third term in (15) is controlled as soon as the empirical partition is close to the theoretical partition. Lemma 3 in [Scornet et al. \[2015\]](#) can be easily extended for any continuous regression function m and any density which is bounded from above and below by some positive constants. Therefore, with probability $1 - \rho$, for all n large enough,

$$\inf_{\mathcal{T}^*} |\text{MDI}_{\mathcal{T}_{n,K}, \gamma}^* - \text{MDI}_{\mathcal{T}_K^*, \gamma}^*| \leq \gamma. \quad (16)$$

Last term in (15) can be treated directly by noting that the theoretical criterion is bounded above by $4\|m\|_\infty^2$. Therefore, we have

$$|\text{MDI}_{\mathcal{T}_K^*, \gamma}^* - \text{MDI}_{\mathcal{T}_K^*}^*| \leq 2^{K+2} \|m\|_\infty^2 \gamma. \quad (17)$$

To control the second term, we make use of the results of [Wager and Walther \[2015\]](#), in particular Corollary 11 which can be adapted to a Gaussian noise by noting that with probability $1 - \rho$, $\max_{1 \leq i \leq n} |\varepsilon_i| \leq C_\rho \log n$. Note that our setting is simpler than that of [Wager and Walther \[2015\]](#) where the dimension d grows to infinity with n . Therefore, with probability at least $1 - \rho$, for all n large enough,

$$|\widehat{\text{MDI}}_{\mathcal{T}_{n,K}, \gamma} - \text{MDI}_{\mathcal{T}_{n,K}, \gamma}^*| \leq \gamma. \quad (18)$$

Regarding the first term in (15), for any cell A such that $\mu(A) < \gamma$,

$$\mu_n(A) \mathbb{V}_n[Y|X \in A] \leq 4\mu_n(A) \|m\|_\infty^2 + \mu_n(A) \mathbb{V}_n[\varepsilon|X \in A].$$

With probability $1 - \rho$, for all n large enough, $\max_{1 \leq i \leq n} |\varepsilon_i| \leq C_\rho \log n$ and, if $n\mu_n(A) \geq \sqrt{n}$,

$$\frac{1}{n\mu_n(A)} \sum_{i, X_i \in A} \varepsilon_i^2 \leq 2\sigma^2.$$

Therefore, with probability $1 - \rho$, for all n large enough,

$$\begin{aligned}
4\mu_n(A)\|m\|_\infty^2 + \mu_n(A)\mathbb{V}_n[\varepsilon|\mathbf{X} \in A] &\leq 8\gamma\|m\|_\infty^2 + \frac{N_n(A)}{n} \frac{1}{N_n(A)} \sum_{i, X_i \in A} \varepsilon_i^2 \\
&\leq 8\gamma\|m\|_\infty^2 + 4\gamma\sigma^2 \mathbf{1}_{N_n(A) \geq \sqrt{n}} + \frac{C_\rho \log n}{\sqrt{n}} \mathbf{1}_{N_n(A) < \sqrt{n}} \\
&\leq 8\gamma\|m\|_\infty^2 + 4\gamma\sigma^2.
\end{aligned} \tag{19}$$

Injecting inequalities (16)-(19) into (15), with probability $1 - \rho$, for all n large enough,

$$\inf_{\mathcal{T}^*} |\widehat{\text{MDI}}_{\mathcal{T}_{n,K}} - \text{MDI}_{\mathcal{T}_K^*}^*| \leq 8\gamma\|m\|_\infty^2 + 4\gamma\sigma^2 + 2\gamma + 2^{K+2}\|m\|_\infty^2\gamma. \tag{20}$$

Since,

$$|\widehat{\text{MDI}}_{\mathcal{T}_{n,K}} - \mathbb{V}[m(\mathbf{X})]| \leq \inf_{\mathcal{T}^*} |\widehat{\text{MDI}}_{\mathcal{T}_{n,K}} - \text{MDI}_{\mathcal{T}_K^*}^*| + |\text{MDI}_{\mathcal{T}_K^*}^* - \mathbb{V}[m(\mathbf{X})]|,$$

we conclude, from (14), (20) and Corollary 1, with probability $1 - \rho$, for all n large enough,

$$|\widehat{\text{MDI}}_{\mathcal{T}_{n,K}} - \mathbb{V}[m(\mathbf{X})]| \leq 3\gamma + 8\gamma\|m\|_\infty^2 + 4\gamma\sigma^2 + 2^{K+2}\|m\|_\infty^2\gamma.$$

Since ρ and γ are arbitrary, this concludes the proof. \square

Proof of Lemma 1. Let $A = \prod_{j=1}^d [a_j, b_j] \subset [0, 1]^d$ be a generic cell. Recall that the regression model satisfies

$$Y = m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) + m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)}) + \varepsilon.$$

For all $j \in \mathcal{J}$, and all splitting position s , define the two resulting cells

$$A_L = A \cap \{\mathbf{x}, x^{(j)} < s\} \quad \text{and} \quad A_R = A \cap \{\mathbf{x}, x^{(j)} \geq s\},$$

where, for notation brevity, the dependence on j and s is omitted. The variance reduction associated to any split (j, s) is defined as

$$\begin{aligned}
L_A^*(j, s) &= \mathbb{V}[m(\mathbf{X})|\mathbf{X} \in A] - \mathbb{P}[\mathbf{X} \in A_L|\mathbf{X} \in A]\mathbb{V}[m(\mathbf{X})|\mathbf{X} \in A_L] \\
&\quad - \mathbb{P}[\mathbf{X} \in A_R|\mathbf{X} \in A]\mathbb{V}[m(\mathbf{X})|\mathbf{X} \in A_R] \\
&= \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})})|\mathbf{X} \in A] - \mathbb{P}[\mathbf{X} \in A_L|\mathbf{X} \in A]\mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})})|\mathbf{X} \in A_L] \\
&\quad - \mathbb{P}[\mathbf{X} \in A_R|\mathbf{X} \in A]\mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})})|\mathbf{X} \in A_R] \\
&\quad + \mathbb{V}[m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)})|\mathbf{X} \in A] - \mathbb{P}[\mathbf{X} \in A_L|\mathbf{X} \in A]\mathbb{V}[m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)})|\mathbf{X} \in A_L] \\
&\quad - \mathbb{P}[\mathbf{X} \in A_R|\mathbf{X} \in A]\mathbb{V}[m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)})|\mathbf{X} \in A_R].
\end{aligned}$$

Since the split is performed along the j th coordinate,

$$\mathbb{V}[m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)})|\mathbf{X} \in A_L] = \mathbb{V}[m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)})|\mathbf{X} \in A_R] = \mathbb{V}[m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)})|\mathbf{X} \in A]. \tag{21}$$

Consequently, the variance reduction satisfies

$$\begin{aligned}
L_A^*(j, s) &= \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})})|\mathbf{X} \in A] - \mathbb{P}[\mathbf{X} \in A_L|\mathbf{X} \in A]\mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})})|\mathbf{X} \in A_L] \\
&\quad - \mathbb{P}[\mathbf{X} \in A_R|\mathbf{X} \in A]\mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})})|\mathbf{X} \in A_R].
\end{aligned} \tag{22}$$

Since $\mathbf{X}^{(\mathcal{J})}$ and $\mathbf{X}^{(\mathcal{J}^c)}$ are independent, letting $A^{(\mathcal{J})} = \prod_{j \in \mathcal{J}} [a_j, b_j]$ be the projection of the cell A along coordinates in \mathcal{J} , we have

$$\begin{aligned} L_A^*(j, s) &= L_{A^{(\mathcal{J})}}^*(j, s) \\ &= \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A^{(\mathcal{J})}] - \mathbb{P}[\mathbf{X} \in A_L^{(\mathcal{J})} | \mathbf{X} \in A^{(\mathcal{J})}] \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_L^{(\mathcal{J})}] \\ &\quad - \mathbb{P}[\mathbf{X} \in A_R^{(\mathcal{J})} | \mathbf{X} \in A^{(\mathcal{J})}] \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_R^{(\mathcal{J})}]. \end{aligned} \quad (23)$$

Therefore, for all $j \in \mathcal{J}$, the criterion $L_A^*(j, s)$ does not depend on $m_{\mathcal{J}^c}$ and is equal to its restriction over the cell $A^{(\mathcal{J})}$. Besides, according to equation (22), any split $j \in \mathcal{J}$ does not change the variance of $m_{\mathcal{J}^c}$. \square

Proof of Proposition 3. Let $\mathbf{x} \in [0, 1]^d$ such that

- (i) any theoretical tree $\mathcal{T}_{k, \mathcal{J}}^*$ built on the distribution $(\mathbf{X}^{(\mathcal{J})}, m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}))$ is consistent, that is

$$\lim_{k \rightarrow \infty} \Delta(m_{\mathcal{J}}, A_{\mathcal{T}_{k, \mathcal{J}}^*}^{(\mathcal{J})}(\mathbf{x}^{(\mathcal{J})})) = 0.$$

- (ii) any theoretical tree $\mathcal{T}_{k, \mathcal{J}^c}^*$ built on the distribution $(\mathbf{X}^{(\mathcal{J}^c)}, m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)}))$ is consistent, that is

$$\lim_{k \rightarrow \infty} \Delta(m_{\mathcal{J}^c}, A_{\mathcal{T}_{k, \mathcal{J}^c}^*}^{(\mathcal{J}^c)}(\mathbf{x}^{(\mathcal{J}^c)})) = 0.$$

Consider a theoretical tree \mathcal{T}_k^* built on the distribution (\mathbf{X}, Y) . For all k , let $A_k(\mathbf{x})$ be the cell in \mathcal{T}_k^* containing \mathbf{x} . For all k , we let

- (i) $A_{\phi(k)}(\mathbf{x})$ be the subsequence of $A_k(\mathbf{x})$ composed of cells split by a coordinate in \mathcal{J} , and $A_{\phi(k)}^{(\mathcal{J})}(\mathbf{x})$ the projection of $A_{\phi(k)}(\mathbf{x})$ over \mathcal{J} ,
- (ii) $A_{\psi(k)}^{(\mathcal{J}^c)}(\mathbf{x})$ be the subsequence of $A_k(\mathbf{x})$ composed of cells split by a coordinate in \mathcal{J}^c , and $A_{\psi(k)}^{(\mathcal{J})}(\mathbf{x})$ the projection of $A_{\psi(k)}(\mathbf{x})$ over \mathcal{J}^c .

According to Lemma 1, the restrictions of the function $m_{\mathcal{J}}$ over $A \in \{A_{\phi(k)+1}^{(\mathcal{J})}(\mathbf{x}), \dots, A_{\phi(k+1)}^{(\mathcal{J})}(\mathbf{x})\}$ are equal. Besides, for any $j \in \mathcal{J}$ and s , the splitting criteria $L_A^*(j, s)$ for $A \in \{A_{\phi(k)+1}^{(\mathcal{J})}(\mathbf{x}), \dots, A_{\phi(k+1)}^{(\mathcal{J})}(\mathbf{x})\}$ are the same. Therefore, the sequence $A_{\phi(k)}^{(\mathcal{J})}(\mathbf{x})$ is a sequence of cells resulting from a theoretical tree built on the distribution $(\mathbf{X}^{(\mathcal{J})}, m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}))$. Similarly, the sequence $A_{\psi(k)}^{(\mathcal{J}^c)}(\mathbf{x})$ is a sequence of cells resulting from a theoretical tree built on the distribution $(\mathbf{X}^{(\mathcal{J}^c)}, m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)}))$.

Let $\mathcal{T}_{\mathcal{J}}^*$ be any theoretical tree whose sequence of cells satisfies $A_{\mathcal{T}_{k, \mathcal{J}}^*}^{(\mathcal{J})}(\mathbf{x}^{(\mathcal{J})}) = A_{\phi(k)}^{(\mathcal{J})}(\mathbf{x})$. Similarly, let $\mathcal{T}_{\mathcal{J}^c}^*$ be any theoretical tree whose sequence of cells satisfies $A_{\mathcal{T}_{k, \mathcal{J}^c}^*}^{(\mathcal{J}^c)}(\mathbf{x}^{(\mathcal{J}^c)}) = A_{\psi(k)}^{(\mathcal{J}^c)}(\mathbf{x})$. Then,

$$\begin{aligned} \Delta(m, A_{\mathcal{T}_k^*}(\mathbf{x})) &= \mathbb{V}[m(\mathbf{X}) | \mathbf{X} \in A_{\mathcal{T}_k^*}(\mathbf{x})] \\ &= \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_{\mathcal{T}_k^*}^{(\mathcal{J})}(\mathbf{x})] + \mathbb{V}[m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)}) | \mathbf{X} \in A_{\mathcal{T}_k^*}^{(\mathcal{J}^c)}(\mathbf{x})] \\ &= \Delta(m_{\mathcal{J}}, A_{\mathcal{T}_{k, \mathcal{J}}^*}^{(\mathcal{J})}(\mathbf{x}^{(\mathcal{J})})) + \Delta(m_{\mathcal{J}^c}, A_{\mathcal{T}_{k, \mathcal{J}^c}^*}^{(\mathcal{J}^c)}(\mathbf{x}^{(\mathcal{J}^c)})). \end{aligned} \quad (24)$$

By assumption, and assuming that $\lim_{\infty} \phi, \psi = \infty$, we have

$$\lim_{k \rightarrow \infty} \Delta(m, A_{\mathcal{T}_k^*}(\mathbf{x})) = 0.$$

Now, consider the case where $\lim_{\infty} \phi, \psi \neq \infty$. Since $\phi(\mathbb{N}) \cup \psi(\mathbb{N}) = \mathbb{N}$, one can assume, without loss of generality, that $\lim_{\infty} \phi = \infty$ and $\lim_{\infty} \psi = p < \infty$. Thus, there exists K such that for all $k > K$, splits in $A_{\mathcal{T}_k^*}(\mathbf{x})$ are performed along coordinates in \mathcal{J} only. Thus, by assumption, the corresponding theoretical tree $\mathcal{T}_{k, \mathcal{J}}^*$ is consistent which implies that the splitting criterion $\mathcal{L}_{A_k(\mathbf{x})}(j, s)$ tends to zero, for all $j \in \mathcal{J}$. This implies that the criterion $\mathcal{L}_{A_k(\mathbf{x})}(j, s)$, for all $j \in \mathcal{J}^c$ is equal to zero (otherwise, it would be larger than the criterion $\mathcal{L}_{A_k(\mathbf{x})}(j, s)$ for $j \in \mathcal{J}$ and consequently, splits would be performed along $j \in \mathcal{J}^c$). This implies that the theoretical tree $\mathcal{T}_{K, \mathcal{J}^c}^*$ is fully grown, and by assumption, we have

$$\Delta(m_{\mathcal{J}^c}, A_{\mathcal{T}_{K, \mathcal{J}^c}^*}(\mathbf{x}^{(\mathcal{J}^c)})) = 0,$$

which concludes the proof, by equality (24). □

Proof of Theorem 2.

First statement. For any theoretical CART tree \mathcal{T}^* , we let, for all $k \in \mathbb{N}$,

$$u_{k, \mathcal{J}} = \sum_{\ell=1}^{n_{cell, k}} \mathbb{P}[\mathbf{X}' \in A_{\ell, k}] \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_{\ell, k}],$$

and

$$f_{k, \mathcal{J}}(\mathbf{x}) = \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_k(\mathbf{x})].$$

We define similarly u_{k, \mathcal{J}^c} and $f_{k, \mathcal{J}^c}(\mathbf{x})$. By definition of the regression model, for all $\mathbf{x} \in [0, 1]^d$,

$$f_{k, \mathcal{J}}(\mathbf{x}) = f_k(\mathbf{x}) - f_{k, \mathcal{J}^c}(\mathbf{x}) \leq f_k(\mathbf{x}),$$

where $f_k(\mathbf{x}) = \mathbb{V}[m(\mathbf{X}) | \mathbf{X} \in A_k(\mathbf{x})]$. By assumption this quantity tends to zero for almost all $\mathbf{X} \in [0, 1]^d$. Besides, it satisfies $f_{k, \mathcal{J}}(\mathbf{x}) \leq 4\|m\|_{\infty}^2$. Therefore, according to the dominated convergence theorem, $\lim_{k \rightarrow \infty} \mathbb{E}[f_{k, \mathcal{J}}(\mathbf{X}')] = 0$. Additionally,

$$\begin{aligned} \mathbb{E}[f_{k, \mathcal{J}}(\mathbf{X}')] &= \mathbb{E}\left[\sum_{\ell=1}^{n_{cell, k}} \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_{\ell, k}] \mathbf{1}_{\mathbf{X}' \in A_{\ell, k}}\right] \\ &= \sum_{\ell=1}^{n_{cell, k}} \mathbb{P}[\mathbf{X}' \in A_{\ell, k}] \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_{\ell, k}] \\ &= u_{k, \mathcal{J}}. \end{aligned}$$

Consequently,

$$\lim_{k \rightarrow \infty} u_{k, \mathcal{J}} = 0 \tag{25}$$

The same result holds for u_{k,\mathcal{J}^c} . To conclude the proof, we just need to show that $u_{k,\mathcal{J}} - u_{k+1,\mathcal{J}}$ corresponds to the variance reduction associated with variables $\mathbf{X}^{(\mathcal{J})}$ between levels k and $k + 1$. Recalling that

$$u_k = \sum_{\ell=1}^{n_{cell,k}} \mathbb{P}[\mathbf{X} \in A_{\ell,k}] \mathbb{V}[m(\mathbf{X}) | \mathbf{X} \in A_{\ell,k}] + \sigma^2,$$

the variance reduction between levels k and $k + 1$ is equal to

$$u_k - u_{k+1} = (u_{k,\mathcal{J}} - u_{k+1,\mathcal{J}}) + (u_{k,\mathcal{J}^c} - u_{k+1,\mathcal{J}^c}),$$

where

$$\begin{aligned} u_{k,\mathcal{J}} - u_{k+1,\mathcal{J}} &= \sum_{\ell=1}^{n_{cell,k}} \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_{\ell,k}] \mathbb{P}[\mathbf{X}' \in A_{\ell,k}] \\ &\quad - \sum_{\ell=1}^{n_{cell,k+1}} \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_{\ell,k+1}] \mathbb{P}[\mathbf{X}' \in A_{\ell,k}]. \end{aligned} \quad (26)$$

Each cell at level $k + 1$ has been created (*i*) either by splitting a cell at level k into two cells, (*ii*) or by copying a cell at level k , if the cell is not split. In that latter case, the contribution of this cell to $u_{k,\mathcal{J}} - u_{k+1,\mathcal{J}}$ is zero. Thus, we can rewrite equation (26) as

$$\begin{aligned} &u_{k,\mathcal{J}} - u_{k+1,\mathcal{J}} \\ &= \sum_{\ell=1}^{n_{\mathcal{J},k}} \left(\mathbb{P}[\mathbf{X}' \in A_{\ell,k}] \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_{\ell,k}] - \mathbb{P}[\mathbf{X}' \in A_{\ell,k,L}] \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_{\ell,k,L}] \right. \\ &\quad \left. - \mathbb{P}[\mathbf{X}' \in A_{\ell,k,R}] \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_{\ell,k,R}] \right) \\ &\quad + \sum_{\ell=n_{\mathcal{J},k+1}}^{n_{split,k}} \left(\mathbb{P}[\mathbf{X}' \in A_{\ell,k}] \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_{\ell,k}] \right. \\ &\quad \left. - \mathbb{P}[\mathbf{X}' \in A_{\ell,k,L}] \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_{\ell,k,L}] \right. \\ &\quad \left. - \mathbb{P}[\mathbf{X}' \in A_{\ell,k,R}] \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})}) | \mathbf{X} \in A_{\ell,k,R}] \right) \end{aligned} \quad (27)$$

where $A_{\ell,k+1,L}$ and $A_{\ell,k+1,R}$ are the children of $A_{\ell,k}$, the first sum (resp. the second) corresponds to the cell split along a coordinate in \mathcal{J} (resp. in \mathcal{J}^c), and $n_{\mathcal{J},k}$ is the number of split performed along a coordinate in \mathcal{J} between level k and $k + 1$.

Note that, according to equation (21), all terms in the second sum in equation (27) are null. Finally, according to equation (22), $u_{k,\mathcal{J}} - u_{k+1,\mathcal{J}}$ is exactly the sum of variance reduction for splits along coordinates in \mathcal{J} that occur between level k and $k + 1$. Consequently, for all K , the population MDI computed with the theoretical tree \mathcal{T}_K^* ,

$$\begin{aligned} \sum_{j \in \mathcal{J}} \text{MDI}_{\mathcal{T}_K^*}^*(X^{(j)}) &= \sum_{k=0}^K (u_{k,\mathcal{J}} - u_{k+1,\mathcal{J}}) \\ &= u_{0,\mathcal{J}} - u_{K,\mathcal{J}}. \end{aligned}$$

Hence, according to equation (25), for all theoretical tree \mathcal{T}^*

$$\lim_{K \rightarrow \infty} \sum_{j \in \mathcal{J}} \text{MDI}_{\mathcal{T}_K^*}^*(X^{(j)}) = \mathbb{V}[m_{\mathcal{J}}(\mathbf{X}^{(\mathcal{J})})], \quad (28)$$

and similarly,

$$\lim_{K \rightarrow \infty} \sum_{j \in \mathcal{J}^c} \text{MDI}_{\mathcal{T}_K^*}^*(X^{(j)}) = \mathbb{V}[m_{\mathcal{J}^c}(\mathbf{X}^{(\mathcal{J}^c)})] \quad (29)$$

This proves the first statement of Theorem 2.

Second statement. To prove the second statement, we need to prove that the convergence of the population MDI in equations (28) and (29) holds uniformly over the set of all possible theoretical trees. Then, we need to extend the convergence to the empirical MDI computed with empirical CART trees. This can be done following the same reasoning as in the proof of Theorem 2. \square

Proof of Theorem 3. Theorem 3 is a direct application of Theorem 2, by noting that the assumption

$$\lim_{k \rightarrow \infty} \Delta(m, A_{\mathcal{T}_k^*}(\mathbf{X}')) = 0$$

in Theorem 2 results from Lemma 1 in Scornet et al. [2015]. \square

To prove Theorem 4, we need the following Technical Lemma, whose proof is postponed to the end of this section.

Technical Lemma 1. *Assume that*

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

where ε is a centred noise of finite variance, independent of \mathbf{X} . Thus, the theoretical splitting criterion in the cell $A \subset [0, 1]^d$ satisfies

$$\begin{aligned} L_A^*(\ell, s) &= -\mathbb{E}^2[m(\mathbf{X})|\mathbf{X} \in A] + \mathbb{P}[\mathbf{X} \in A_L|\mathbf{X} \in A]\mathbb{E}^2[m(\mathbf{X})|\mathbf{X} \in A_L] \\ &\quad + \mathbb{P}[\mathbf{X} \in A_R|\mathbf{X} \in A]\mathbb{E}^2[m(\mathbf{X})|\mathbf{X} \in A_R], \end{aligned}$$

where $A_L = A \cap \{X^{(\ell)} < s\}$ and $A_R = A \cap \{X^{(\ell)} \geq s\}$.

Proof of Theorem 4. Let $A = \prod_{\ell=1}^d [a_\ell, b_\ell] \subset [0, 1]^d$ be a generic cell. According to Lemma 1, the splitting criterion writes

$$\begin{aligned} L_A^*(j, s) &= -\left(\mathbb{E}[\alpha_j X^{(j)}|X \in A]\right)^2 + \frac{s - a_j}{b_j - a_j} \left(\mathbb{E}[\alpha_j X^{(j)}|X \in A, X^{(j)} \leq s]\right)^2 \\ &\quad + \frac{b_j - s}{b_j - a_j} \left(\mathbb{E}[\alpha_j X^{(j)}|X \in A, X^{(j)} > s]\right)^2 \\ &= -\alpha_j^2 \left(\frac{a_j + b_j}{2}\right)^2 + \alpha_j^2 \left(\frac{s - a_j}{b_j - a_j}\right) \left(\frac{a_j + s}{2}\right)^2 + \alpha_j^2 \left(\frac{b_j - s}{b_j - a_j}\right) \left(\frac{b_j + s}{2}\right)^2 \\ &= -\alpha_j^2 \left(\frac{a_j + b_j}{2}\right)^2 + \alpha_j^2 \left(s^2 - a_j^2\right) \left(\frac{a_j + s}{4(b_j - a_j)}\right) + \alpha_j^2 \left(b_j^2 - s^2\right) \left(\frac{b_j + s}{4(b_j - a_j)}\right) \\ &= \left(\frac{\alpha_j^2}{4}\right) s^2 - \alpha_j^2 \left(\frac{a_j + b_j}{4}\right) s + \frac{\alpha_j^2 a_j b_j}{4}. \end{aligned}$$

For any j , the function $s \mapsto L_A^*(j, s)$ reaches its maximum at $s^* = (a_j + b_j)/2$. For this particular value, we have

$$\begin{aligned} L_A^*(j, s^*) &= \left(\frac{\alpha_j^2}{4}\right) \left(\frac{a_j + b_j}{2}\right)^2 - \alpha_j^2 \left(\frac{a_j + b_j}{4}\right) \left(\frac{a_j + b_j}{2}\right) + \frac{\alpha_j^2 a_j b_j}{4} \\ &= \frac{\alpha_j^2}{4} \left[\frac{b_j - a_j}{2}\right]^2. \end{aligned}$$

Thus the coordinate chosen for splitting in cell A must satisfy

$$j^* \in \arg \max_{\ell \in \{1, \dots, d\}} \alpha_\ell^2 (b_\ell - a_\ell)^2. \quad (30)$$

Since the chosen coordinate satisfies equation (30) and the range of the cell along the chosen side is halved after the split, the diameter of any cell tends to zero, which proves that the theoretical tree is consistent. Hence Theorem 2 can be applied and leads to

$$\lim_{k \rightarrow \infty} \text{MDI}_{T_k^*}^*(X^{(j)}) = \alpha_j^2 \mathbb{V}[X^{(j)}] = \frac{\alpha_j^2}{12},$$

since $X^{(j)} \sim \mathcal{U}([0, 1])$. The convergence of the empirical MDI directly results from Theorem 2. \square

Proof of Theorem 5. Let $A = \prod_{\ell=1}^d [a_\ell, b_\ell] \subset [0, 1]$ be a generic cell. According to Technical Lemma 1, the splitting criterion writes

$$\begin{aligned} L_A^*(\ell, s) &= \mathbb{V}[Y|\mathbf{X} \in A] - \mathbb{P}[\mathbf{X} \in A_L|\mathbf{X} \in A] \mathbb{V}[Y|\mathbf{X} \in A_L] \\ &\quad - \mathbb{P}[\mathbf{X} \in A_R|\mathbf{X} \in A] \mathbb{V}[Y|\mathbf{X} \in A_R] \\ &= -\mathbb{E}^2[m(\mathbf{X})|\mathbf{X} \in A] + \mathbb{P}[\mathbf{X} \in A_L|\mathbf{X} \in A] \mathbb{E}^2[m(\mathbf{X})|\mathbf{X} \in A_L] \\ &\quad + \mathbb{P}[\mathbf{X} \in A_R|\mathbf{X} \in A] \mathbb{E}^2[m(\mathbf{X})|\mathbf{X} \in A_R] \\ &= -\mathbb{E}^2 \left[2^k \alpha \prod_{j=1}^k X_j | \mathbf{X} \in A \right] + \frac{s - a_\ell}{b_\ell - a_\ell} \mathbb{E}^2 \left[2^k \alpha \prod_{j=1}^k X_j | \mathbf{X} \in A, X_\ell < s \right] \\ &\quad + \frac{b_\ell - s}{b_\ell - a_\ell} \mathbb{E}^2 \left[2^k \alpha \prod_{j=1}^k X_j | \mathbf{X} \in A, X_\ell \geq s \right] \\ &= \frac{\alpha^2}{b_\ell - a_\ell} \prod_{j \neq \ell} (a_j + b_j)^2 \left[- (a_\ell + b_\ell)^2 (b_\ell - a_\ell) + (s - a_\ell)(s + a_\ell)^2 \right. \\ &\quad \left. + (b_\ell - s)(s + b_\ell)^2 \right] \\ &= \alpha^2 \left(-s^2 + (a_\ell + b_\ell) - a_\ell b_\ell \right) \prod_{j \neq \ell} (a_j + b_j)^2. \end{aligned}$$

For any ℓ , the function $s \mapsto L_A^*(\ell, s)$ reaches its maximum at $s^* = (a_\ell + b_\ell)/2$. For this particular value, we have

$$L_A^*(\ell, s^*) = \frac{\alpha^2}{4} (b_\ell - a_\ell)^2 \prod_{j \neq \ell} (a_j + b_j)^2. \quad (31)$$

The coordinate ℓ^* is chosen for splitting if $L_A^*(\ell^*, s^*) = \max_{1 \leq k \leq d} L_A^*(k, s^*)$, that is, for all k ,

$$\frac{b_{\ell^*} - a_{\ell^*}}{a_{\ell^*} + b_{\ell^*}} \geq \frac{b_k - a_k}{a_k + b_k}.$$

The variance reduction induced by the split (ℓ^*, s^*) in the cell $A = \prod_{j=1}^d [a_j, b_j]$ is equal to

$$\frac{\alpha_{\ell^*}^2}{4} (b_{\ell^*} - a_{\ell^*})^2 \prod_{j \neq \ell^*} (a_j + b_j)^2. \quad (32)$$

Now, fix $\mathbf{x} \in (0, 1]^d$. We want to prove that $\text{diam}(A_k(\mathbf{x}))$ tends to zero as k tends to infinity. For any cell A , and all $\ell \in \{1, \dots, d\}$, according to equation (31), we have

$$\frac{\alpha^2}{4} (b_\ell - a_\ell)^2 \prod_{j \neq \ell} x_j^2 \leq L_A^*(\ell, s^*) \leq \frac{\alpha^2}{4} (b_\ell - a_\ell)^2. \quad (33)$$

For all k , let $A_k(\mathbf{x}) = \prod_{j=1}^d [a_{j,k}, b_{j,k}]$ be the cell containing \mathbf{x} in the theoretical tree \mathcal{T}_k^* . Let $u_k^{(j)} = b_{j,k} - a_{j,k}$. Since, for all j , $(u_k^{(j)})$ is a positive decreasing sequence, we obtain that $(u_k^{(j)})$ converges toward $\gamma_j \geq 0$. Let $\Gamma = \{j, \gamma_j > 0\}$ and assume that $\Gamma \neq \emptyset$. For all j ,

$$\begin{cases} u_{k+1}^{(j)} = u_k^{(j)} & \text{if the split is not performed along the } j\text{th coordinate} \\ u_{k+1}^{(j)} = \frac{u_k^{(j)}}{2} & \text{otherwise.} \end{cases}.$$

Hence, there exists K such that for all $k > K$, and all $j \in \mathcal{J}$,

$$u_k^{(j)} = \gamma_j. \quad (34)$$

In other words, after level k , all splits are performed along coordinates belonging to \mathcal{J}^c . Now, for all $j \in \mathcal{J}^c$, $\lim_{k \rightarrow \infty} u_k^{(j)} = 0$, and thus, according to equation (33),

$$\lim_{k \rightarrow \infty} \max_{j \in \mathcal{J}^c} L_{A_k(\mathbf{x})}^*(j, s^*) = 0.$$

Therefore, there exists $K_1 > K$, such that, for all $k > K_1$,

$$\max_{j \in \mathcal{J}^c} L_{A_k(\mathbf{x})}^*(j, s^*) < \min_{\ell \in \mathcal{J}} \frac{\alpha^2}{8} \gamma_\ell^2 \prod_{j \neq \ell} x_j^2.$$

Using again Equation (33), we deduce that, for all $k > K_1$,

$$\max_{j \in \mathcal{J}^c} L_{A_k(\mathbf{x})}^*(j, s^*) < \min_{\ell \in \mathcal{J}} L_{A_k(\mathbf{x})}^*(\ell, s^*),$$

thus the next split is performed along $j \in \mathcal{J}$ which is absurd. Hence, either $\mathcal{J} = \emptyset$ or $\mathcal{J}^c = \emptyset$. If $\mathcal{J}^c = \emptyset$, given the expression of the variance reduction in equation (32), an additional split must occur after rank K , with an additional variance decreasing, which is forbidden by equation (34). Finally, $\mathcal{J} = \emptyset$ and consequently, for all $j \in \{1, \dots, d\}$, $\lim_{k \rightarrow \infty} u_k^{(j)} = 0$, which implies, for all $\mathbf{x} \in (0, 1]^d$,

$$\lim_{k \rightarrow \infty} \text{diam}(A_k(\mathbf{x})) = 0.$$

Since \mathbf{X} is uniformly distributed over $[0, 1]^d$, and m is continuous, we have, almost surely,

$$\lim_{k \rightarrow \infty} \Delta(m, A_{\mathcal{T}_k^*}(\mathbf{X})) = 0.$$

A direct application of Theorem 1 yields

$$\sum_{j=1}^d \text{MDI}_{\mathcal{T}_{\infty}^*}^*(X^{(j)}) = \mathbb{V}[m(\mathbf{X})].$$

A direct calculation show that $\mathbb{E}[m(\mathbf{X})] = \alpha$, and $\mathbb{E}[m(\mathbf{X})^2] = \alpha^2 (4/3)^d$, leading directly to the conclusion. The last statement is a direct application of Theorem 2. \square

Proof of Lemma 2. We prove the result in dimension two. The generalization in any dimension will directly follow. According to Theorem 5, in the root cell $[0, 1]^2$, the variable ℓ^* selected for splitting must satisfy

$$\ell^* \in \arg \max_{k \in \{1, 2\}} \left(\frac{a_k + b_k}{2} \right) (b_\ell - a_\ell),$$

where the right-hand term is equal to $1/2$ for $k = 1, 2$. Therefore, the selected variable can be $X^{(1)}$ or $X^{(2)}$. The corresponding diminution in variance is equal to $\alpha^2/4$.

First theoretical tree. Consider the theoretical tree $\mathcal{T}_{2,1}^*$ whose construction is stopped at level two, where the first split is performed along $X^{(1)}$. The same calculation as above shows that splits must occur along the second variable which gives a variance decreasing of $\alpha^2/16$ in the left cell, and $9\alpha^2/16$ in the right cell. The corresponding partition is shown on the left of Figure 2. The variable importances for the tree $\mathcal{T}_{2,1}^*$ are

$$\text{MDI}_{\mathcal{T}_{2,1}^*}^*(X^{(1)}) = \alpha^2/4$$

and

$$\text{MDI}_{\mathcal{T}_{2,1}^*}^*(X^{(2)}) = \frac{1}{2} (\alpha^2/16 + 9\alpha^2/16) = 5\alpha^2/16$$

Second theoretical tree. Consider the theoretical tree $\mathcal{T}_{2,2}^*$ whose construction is stopped at level two, where the first split is made along the second variable. Then the same reasoning as above shows that the second split must be made along the second variable, the resulting partition is displayed on the right of Figure 2. The variable importances for the tree $\mathcal{T}_{2,2}^*$ are thus

$$\text{MDI}_{\mathcal{T}_{2,2}^*}^*(X^{(1)}) = 5\alpha^2/16$$

and

$$\text{MDI}_{\mathcal{T}_{2,2}^*}^*(X^{(2)}) = \alpha^2/4$$

Conclusion The partitions of the two theoretical trees $\mathcal{T}_{2,1}^*$ and $\mathcal{T}_{2,2}^*$ defined above are the same but variable importances are not. Since the partitions are the same, one can choose any splitting strategy for the following splits and use this exact same strategy in both $\mathcal{T}_{2,1}^*$ and $\mathcal{T}_{2,2}^*$. Since the splits occurring at level $k \geq 2$ are set to be the same, we have

$$\text{MDI}_{\mathcal{T}_{k,2}^*}^*(X^{(1)}) - \text{MDI}_{\mathcal{T}_{k,1}^*}^*(X^{(1)}) = 5\alpha^2/16 - \alpha^2/4 = \alpha^2/16.$$

\square

Proof of Lemma 3. We have $\mathbf{X} \sim \mathcal{U}\left(\bigcup_{j=1}^{2^\beta} \left[\frac{j-1}{2^\beta}, \frac{j}{2^\beta}\right]^2\right) = \mathcal{U}^{\otimes 2^\beta}$. Therefore,

$$\begin{aligned}\text{Cov}(X^{(1)}, X^{(2)}) &= \mathbb{E}\left[X^{(1)}X^{(2)} \sum_{j=1}^{2^\beta} \mathbf{1}_{\mathbf{X} \in \left[\frac{j-1}{2^\beta}, \frac{j}{2^\beta}\right]^2}\right] - \frac{1}{4} \\ &= \sum_{j=1}^{2^\beta} \mathbb{E}\left[\mathbf{1}_{\mathbf{X} \in \left[\frac{j-1}{2^\beta}, \frac{j}{2^\beta}\right]^2} \mathbb{E}\left[X^{(1)}X^{(2)} \mid \mathbf{X} \in \left[\frac{j-1}{2^\beta}, \frac{j}{2^\beta}\right]^2\right]\right] - \frac{1}{4},\end{aligned}$$

with

$$\mathbb{E}\left[X^{(1)}X^{(2)} \mid \mathbf{X} \in \left[\frac{j-1}{2^\beta}, \frac{j}{2^\beta}\right]^2\right] = \frac{1}{2^{2\beta}} \left(j - \frac{1}{2}\right)^2,$$

which leads to

$$\begin{aligned}\text{Cov}(X^{(1)}, X^{(2)}) &= \sum_{j=1}^{2^\beta} \frac{1}{2^{3\beta}} \left(j - \frac{1}{2}\right)^2 - \frac{1}{4} \\ &= \frac{1}{2^{3\beta+2}} \sum_{j=1}^{2^\beta} (2j-1)^2 - \frac{1}{4}.\end{aligned}$$

Since we know that

$$\sum_{j=1}^{2^\beta} (2j-1)^2 = \frac{2^\beta(2^{2\beta+2}-1)}{3},$$

the covariance can be expressed as

$$\text{Cov}(X^{(1)}, X^{(2)}) = \frac{1}{12} - \frac{1}{3 \cdot 2^{2\beta+2}}.$$

Since the marginals are uniform on $[0, 1]$, we have $\mathbb{V}[X^{(1)}] = \mathbb{V}[X^{(2)}] = 1/12$, which yields,

$$\text{Corr}(X^{(1)}, X^{(2)}) = 1 - \frac{1}{4^\beta}.$$

□

Proof of Proposition 4. The first statement of Theorem 2 remains valid even when the density of \mathbf{X} is not lower bounded by some positive constant. Besides, using Technical Lemma 1, the splitting criterion for Model 6, computed for a split along the first coordinate writes

$$\begin{aligned}L^*(1, s) &= -\mathbb{E}^2[X^{(1)} + X^{(2)} \mid (X^{(1)}, X^{(2)}) \in [0, 1]^2] \\ &\quad + s\mathbb{E}^2[X^{(1)} + X^{(2)} \mid (X^{(1)}, X^{(2)}) \in [0, s] \times [0, 1]] \\ &\quad + (1-s)\mathbb{E}^2[X^{(1)} + X^{(2)} \mid (X^{(1)}, X^{(2)}) \in [s, 1] \times [0, 1]] \\ &= -1 + s\left(\frac{s}{2} + \mathbb{E}[X^{(2)} \mid (X^{(1)}, X^{(2)}) \in [0, s] \times [0, 1]]\right)^2 \\ &\quad + (1-s)\left(\frac{1+s}{2} + \mathbb{E}[X^{(2)} \mid (X^{(1)}, X^{(2)}) \in [s, 1] \times [0, 1]]\right)^2.\end{aligned}\tag{35}$$

Let us fix a split at position s along variable $X^{(1)}$ and let $\ell = \lfloor 2^\beta s \rfloor$. The conditional density of $X^{(2)}$ given that $(X^{(1)}, X^{(2)}) \in [0, s] \times [0, 1]$ is

$$f_{X^{(2)}|X^{(1)} \in [0, s]}(x_2) = C \left(\frac{1}{2^\beta} \mathbf{1}_{2^\beta x_2 \leq \ell} + \left(s - \frac{\ell}{2^\beta} \right) \mathbf{1}_{\ell < 2^\beta x_2 \leq \ell+1} \right).$$

where, in order for $f_{X_2|X_1 \in [0, s]}$ to be a density, C must satisfy

$$C \left(\frac{\ell}{4^\beta} + \left(s - \frac{\ell}{2^\beta} \right) \frac{1}{2^\beta} \right) = 1,$$

that is $C = 2^\beta/s$. Hence,

$$f_{X^{(2)}|X^{(1)} \in [0, s]}(x_2) = \frac{2^\beta}{s} \left(\frac{1}{2^\beta} \mathbf{1}_{2^\beta x_2 \leq \ell} + \left(s - \frac{\ell}{2^\beta} \right) \mathbf{1}_{\ell < 2^\beta x_2 < \ell+1} \right),$$

and similarly,

$$f_{X^{(2)}|X^{(1)} \in [s, 1]}(x_2) = \frac{2^\beta}{1-s} \left(\left(\frac{\ell+1}{2^\beta} - s \right) \mathbf{1}_{\ell < 2^\beta x_2 < \ell+1} + \frac{1}{2^\beta} \mathbf{1}_{2^\beta x_2 \geq \ell+1} \right).$$

To make explicit the splitting criterion in equation (35), we need to compute, on one hand,

$$\begin{aligned} & \mathbb{E}[X^{(1)} + X^{(2)} | (X^{(1)}, X^{(2)}) \in [0, s] \times [0, 1]] \\ &= \frac{s}{2} + \frac{1}{s} \int_0^{\ell/2^\beta} x_2 dx_2 + \frac{2^\beta}{s} \left(s - \frac{\ell}{2^\beta} \right) \int_{\ell/2^\beta}^{(\ell+1)/2^\beta} x_2 dx_2 \\ &= \frac{s}{2} + \frac{1}{2s} \left(\frac{\ell}{2^\beta} \right)^2 + \frac{2^\beta}{2s} \left(s - \frac{\ell}{2^\beta} \right) \left(\left(\frac{\ell+1}{2^\beta} \right)^2 - \left(\frac{\ell}{2^\beta} \right)^2 \right) \\ &= \frac{s}{2} + \frac{1}{2s} \left(\frac{\ell}{2^\beta} \right)^2 + \frac{2\ell+1}{2^{\beta+1}s} \left(s - \frac{\ell}{2^\beta} \right), \end{aligned} \quad (36)$$

and on the other hand,

$$\begin{aligned} & \mathbb{E}[X^{(1)} + X^{(2)} | (X^{(1)}, X^{(2)}) \in [s, 1] \times [0, 1]] \\ &= \frac{1+s}{2} + \frac{2^\beta}{1-s} \left(\frac{\ell+1}{2^\beta} - s \right) \int_{\ell/2^\beta}^{(\ell+1)/2^\beta} x_2 dx_2 + \frac{2^\beta}{1-s} \frac{1}{2^\beta} \int_{(\ell+1)/2^\beta}^1 x_2 dx_2 \\ &= \frac{1+s}{2} + \frac{2^\beta}{2(1-s)} \left(\frac{\ell+1}{2^\beta} - s \right) \left(\left(\frac{\ell+1}{2^\beta} \right)^2 - \left(\frac{\ell}{2^\beta} \right)^2 \right) + \frac{1}{2(1-s)} \left(1 - \left(\frac{\ell+1}{2^\beta} \right)^2 \right) \\ &= \frac{1+s}{2} - \frac{(2\ell+1)s}{2^{\beta+1}(1-s)} + \frac{1}{2(1-s)} \left(\frac{\ell(\ell+1)}{2^{2\beta}} + 1 \right). \end{aligned} \quad (37)$$

Injecting equations (36) and (37) into (35), we get

$$\begin{aligned} L^*(1, s) &= -1 + s \left(\frac{s}{2} + \frac{2\ell+1}{2^{\beta+1}} - \frac{\ell^2 + \ell}{2^{2\beta+1}s} \right)^2 \\ &\quad + (1-s) \left(\frac{1+s}{2} - \frac{(2\ell+1)s}{2^{\beta+1}(1-s)} + \frac{1}{2(1-s)} \left(\frac{\ell(\ell+1)}{2^{2\beta}} + 1 \right) \right)^2. \end{aligned} \quad (38)$$

The second term in equation (38) can be expressed as

$$\begin{aligned} \left(\frac{s}{2} + \frac{2\ell+1}{2^{\beta+1}} - \frac{\ell^2+\ell}{2^{2\beta+1}s}\right)^2 &= \left(\frac{2\ell+1}{2^{\beta+1}}\right)^2 + \frac{s^2}{4} + \left(\frac{\ell^2+\ell}{2^{2\beta+1}s}\right)^2 - \frac{\ell^2+\ell}{2^{2\beta+1}} \\ &\quad + s\left(\frac{2\ell+1}{2^{\beta+1}}\right) - \left(\frac{2\ell+1}{2^{\beta+1}}\right)\left(\frac{\ell^2+\ell}{2^{2\beta+1}s}\right). \end{aligned}$$

and third term in equation (38) as

$$\begin{aligned} &\left(\frac{1+s}{2} - \frac{(2\ell+1)s}{2^{\beta+1}(1-s)} + \frac{1}{2(1-s)}\left(\frac{\ell(\ell+1)}{2^{2\beta}} + 1\right)\right)^2 \\ &= \frac{1}{4} + \frac{s}{2} + \frac{s^2}{4} + \frac{1}{2}\left(\frac{\ell(\ell+1)}{2^{2\beta}} + 1\right)\frac{1}{1-s} + \frac{s}{1-s}\left(\frac{1}{2}\left(\frac{\ell(\ell+1)}{2^{2\beta}} + 1\right) - \frac{2\ell+1}{2^{\beta+1}}\right) \\ &\quad - \frac{(2\ell+1)s^2}{2^{\beta+1}(1-s)} + \frac{1}{4(1-s)^2}\left(\frac{\ell(\ell+1)}{2^{2\beta}} + 1\right)^2 - \left(\frac{\ell(\ell+1)}{2^{2\beta}} + 1\right)\left(\frac{2\ell+1}{2^{\beta+1}}\right)\frac{s}{(1-s)^2} \\ &\quad + \left(\frac{2\ell+1}{2^{\beta+1}}\right)^2\frac{s^2}{(1-s)^2}, \end{aligned}$$

which, after several simplifications, leads to

$$\begin{aligned} L^*(1, s) &= -\frac{1}{4}s^2 + s\left(\frac{3}{4} + \frac{4\ell^2+4\ell+1}{2^{2\beta+2}} - \left(\frac{2\ell+1}{2^{\beta+1}}\right)\right) - \frac{3}{4} + \frac{1}{2}\left(\frac{\ell(\ell+1)}{2^{2\beta}} + 1\right) \\ &\quad - \left(\frac{2\ell+1}{2^{\beta+1}}\right)\left(\frac{\ell^2+\ell}{2^{2\beta}}\right) + \frac{1}{s}\left(\frac{\ell^2+\ell}{2^{2\beta+1}}\right)^2 + \frac{1}{4(1-s)}\left(\frac{\ell(\ell+1)}{2^{2\beta}} + 1\right)^2 \\ &\quad + \frac{s^2}{1-s}\left(\frac{2\ell+1}{2^{\beta+1}}\right)^2 - \frac{s}{1-s}\left(\frac{2\ell+1}{2^{\beta+1}}\right)\left(\frac{\ell(\ell+1)}{2^{2\beta}} + 1\right). \end{aligned}$$

Thus we have

$$\begin{aligned} \frac{\partial L^*(1, s)}{\partial s} &= -\frac{1}{2}s + \frac{3}{4} + \frac{4\ell^2+4\ell+1}{2^{2\beta+2}} - \left(\frac{2\ell+1}{2^{\beta+1}}\right) - \frac{1}{s^2}\left(\frac{\ell^2+\ell}{2^{2\beta+1}}\right)^2 \\ &\quad + \frac{1}{(1-s)^2}\left(\frac{1}{4}\left(\frac{\ell(\ell+1)}{2^{2\beta}} + 1\right)^2 - \left(\frac{2\ell+1}{2^{\beta+1}}\right)\left(\frac{\ell(\ell+1)}{2^{2\beta}} + 1\right)\right) \\ &\quad + \frac{s(2-s)}{(1-s)^2}\left(\frac{2\ell+1}{2^{\beta+1}}\right)^2, \end{aligned} \tag{39}$$

with

$$\left|\frac{\partial L^*(1, s)}{\partial s}\right| \leq 6 + 12 \times 2^{2\beta}.$$

We know that the function $s \mapsto L^*(1, s)$ is symmetric around $1/2$ and we conjecture that it is increasing on $[0, 1/2]$. Note that, we have $L^*(1, 1/2) = 1/4$. For all β , let $m_\beta = L^*(1, 1/2 - 1/2^\beta)$. Setting a regular grid of size

$$\varepsilon = \frac{\frac{1}{4} - m_\beta}{\left\|\frac{\partial L^*(1, s)}{\partial s}\right\|_\infty},$$

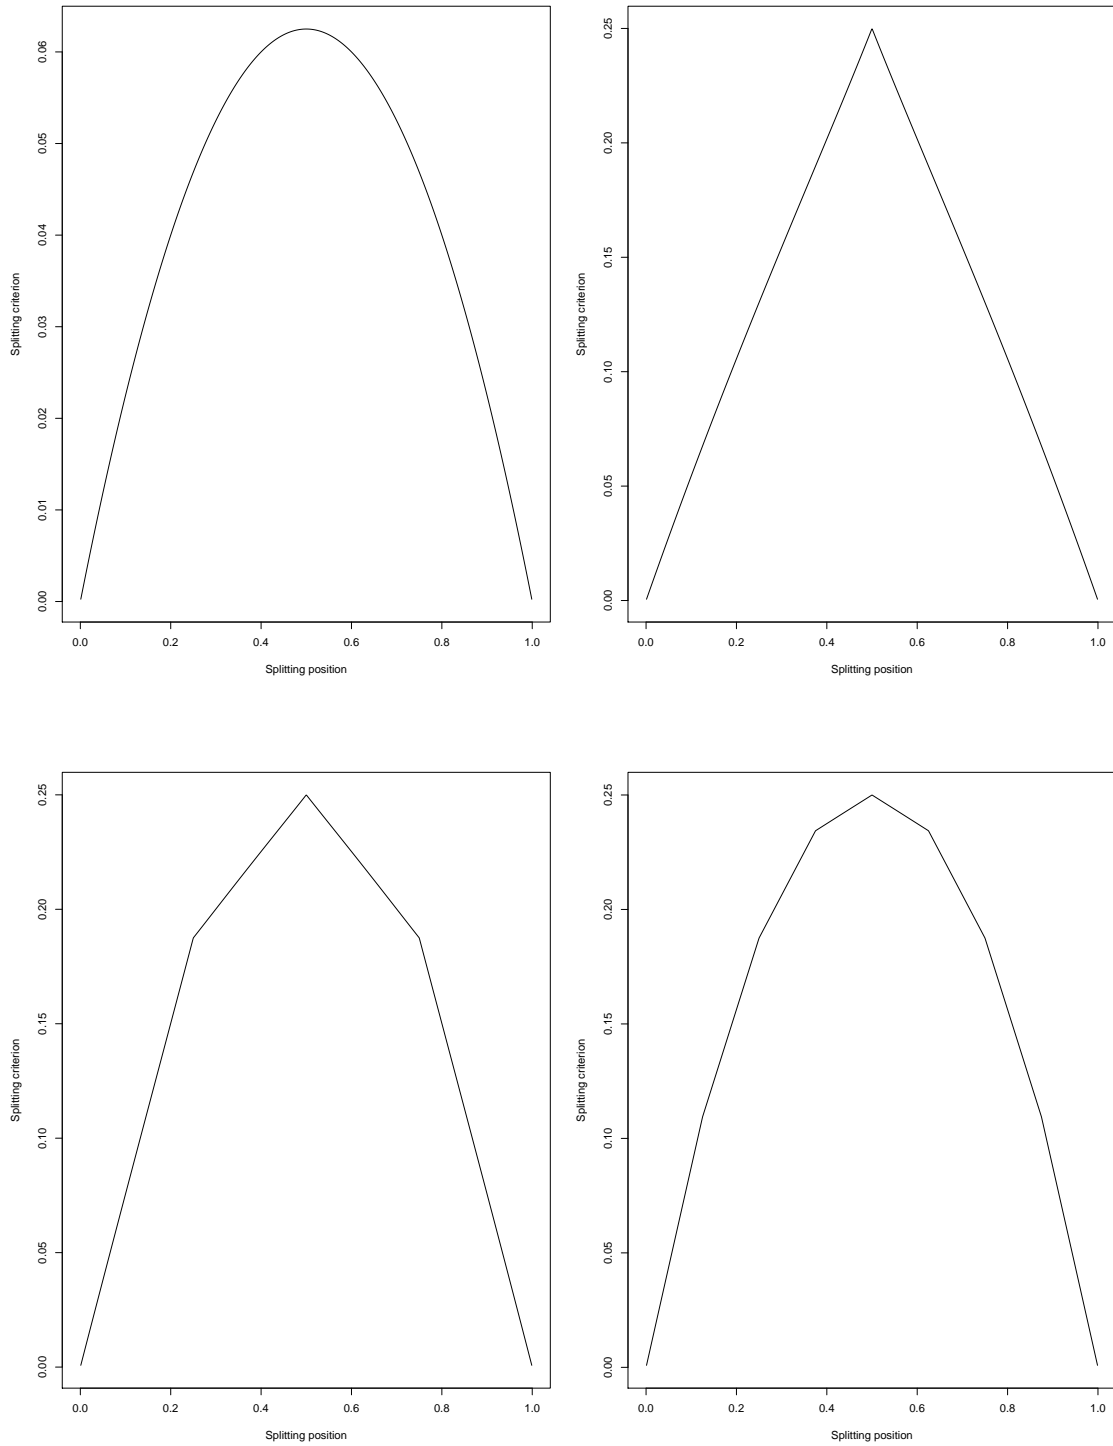


Figure 16: Illustration of the splitting criterion $L^*(1, s)$ for $\beta = 0$ (top left), $\beta = 1$ (top right), $\beta = 2$ (bottom left), $\beta = 3$ (bottom right).

if all evaluations of $L^*(1, s)$ on the previous grid are lower than m_β then the maximum of $s \mapsto L^*(1, s)$ is reached in $[1/2 - 1/2^\beta, 1/2 + 1/2^\beta]$. An numerical optimization with the function `optimize` of `R` shows that it is true for $\beta = 1, \dots, 5$. Figure 16 displays the splitting criterion for $\beta = 0, \dots, 3$. On the interval $[1/2 - 1/2^\beta, 1/2)$, the partial

derivative writes

$$\frac{\partial L^*(1, s)}{\partial s} = -\frac{1}{2}s + \left(\frac{3}{4} + \gamma_\beta\right) - \frac{\gamma_\beta}{s^2},$$

where

$$\gamma_\beta = -\frac{1}{2^{\beta+1}} + \frac{1}{2^{2\beta+2}}.$$

Using equation (39) and the fact that $\frac{\partial L^*(1,1)}{\partial s} = 0$, we have

$$\frac{\partial L^*(1, s)}{\partial s} = -\frac{1}{2}(s-1)\left(s^2 + \left(1 - 2\left(\frac{3}{4} + \gamma_\beta\right)\right)s + 2\left(\frac{1}{4} + \gamma_\beta\right)\right), \quad (40)$$

where the discriminant of the second order expression in equation (40) is, for $\beta \geq 1$,

$$\begin{aligned} \Delta &= 4\gamma_\beta^2 - 6\gamma_\beta - \frac{7}{4} \\ &= -\frac{7}{4} + \frac{3}{2^\beta} - \frac{1}{2^{2\beta+1}} - \frac{1}{2^{3\beta}} + \frac{1}{2^{4\beta+2}} \\ &< 0. \end{aligned}$$

Thus the sign of the expression in the right-hand side of equation (40) does not vary on $[0, 1/2)$ and is positive in zero. Thus, the splitting criterion is increasing on $[1/2 - 1/2^\beta, 1/2)$. Since we know with the optimization procedure that the maximum is reached in this interval, it is reached for $s = 1/2$. This proves the first statement of Proposition 4.

Regarding the second statement of Proposition 4, note that the above calculations hold also for a split along the second variable. According to Theorem 4, the variance reduction associated to a split performed along $X^{(3)}$ equals $\alpha^2/16$. Thus, for $\beta = 0, \dots, 5$, a split will occur along the first or second variable if, and only if

$$\frac{\alpha^2}{16} < \frac{1}{4},$$

that is $|\alpha| < 2$. □

Proof of Theorem 6. The distribution of $X^{(3)}$ in each cell is still uniform. Assume that $|\alpha| < 2$, so that the first split is performed on (without loss of generality) the first variable. Then the resulting left cell is $A = [0, 1/2] \times [0, 1]^2$. For $\beta > 1$, the conditional distribution of $(X^{(1)}, X^{(2)})$ on A is given by

$$(X^{(1)}, X^{(2)}) \sim \mathcal{U}\left(\bigcup_{j=0}^{2^{\beta-1}-1} \left[\frac{j}{2^\beta}, \frac{j+1}{2^\beta}\right]^2\right).$$

On this cell, the regression model still writes

$$Y = \frac{1}{2}(2X^{(1)}) + \frac{1}{2}(2X^{(2)}) + \alpha X^{(3)} + \varepsilon,$$

where

$$(2X^{(1)}, 2X^{(2)}) \sim \mathcal{U}\left(\bigcup_{j=0}^{2^{\beta-1}-1} \left[\frac{j}{2^{\beta-1}}, \frac{j+1}{2^{\beta-1}}\right]^2\right).$$

Therefore, the previous reasoning applies and a split is performed on the first or second variable in this cell if, and only if,

$$\frac{1}{16} > \frac{\alpha^2}{16},$$

that is $|\alpha| < 1$. This holds until each cell has been cut β times along the first or second coordinate. In the resulting cells, the distribution of \mathbf{X} is uniform on its support which is an hyperrectangle, and, according to Theorem 4, the variance of the regression function decreases to zero as a function of the tree level k . The previous reasoning shows that the sequence of splits along $\{X^{(1)} \text{ or } X^{(2)}\}$ and $X^{(3)}$ is deterministic and depends only of the coefficient α . The fourth statement, that is the convergence of the empirical MDI directly results from the convergence of theoretical trees proved above, and from a direct application of Theorem 2. \square

Proof of Lemma 4. According to Theorem 2, splits along $X^{(3)}$ has no impact on the variance of $X^{(1)} + X^{(2)}$. One can thus assume, without loss of generality that the splits of the first β levels are all performed along the first and the second variables. According to the proof of Theorem 4, in each cell of the first β levels, the splitting criterion is the same along variable $X^{(1)}$ and $X^{(2)}$. Consider the tree \mathcal{T}_1^* where all splits in the first β levels are made along the first variable. Then, the total decrease of variance associated to all these splits is

$$\text{MDI}_{\mathcal{T}_{\beta,1}^*}^*(X^{(1)}) = \sum_{j=0}^{\beta-1} \frac{1}{4} \left(\frac{1}{2^j}\right)^2 = \frac{1}{3} - \frac{1}{3} \left(\frac{1}{4}\right)^\beta.$$

In each cell that results from the first β cuts, the distribution of $(X^{(1)}, X^{(2)})$ is uniform. Then the decrease of variance in all these cells is the same for $X^{(1)}$ and $X^{(2)}$. Hence,

$$\lim_{k \rightarrow \infty} \left(\text{MDI}_{\mathcal{T}_{k,1}^*}^*(X^{(1)}) - \text{MDI}_{\mathcal{T}_{k,1}^*}^*(X^{(2)}) \right) = \frac{1}{3} - \frac{1}{3} \left(\frac{1}{4}\right)^\beta.$$

Let \mathcal{T}_2^* be the exact same tree as \mathcal{T}_1^* but where all the first β splits are performed along the second variable. As before, we have,

$$\lim_{k \rightarrow \infty} \left(\text{MDI}_{\mathcal{T}_{k,2}^*}^*(X^{(2)}) - \text{MDI}_{\mathcal{T}_{k,2}^*}^*(X^{(1)}) \right) = \frac{1}{3} - \frac{1}{3} \left(\frac{1}{4}\right)^\beta.$$

According to Theorem 4, we know that

$$\begin{aligned} \lim_{k \rightarrow \infty} \left(\text{MDI}_{\mathcal{T}_{k,1}^*}^*(X^{(1)}) + \text{MDI}_{\mathcal{T}_{k,1}^*}^*(X^{(2)}) \right) &= \lim_{k \rightarrow \infty} \left(\text{MDI}_{\mathcal{T}_{k,2}^*}^*(X^{(1)}) + \text{MDI}_{\mathcal{T}_{k,2}^*}^*(X^{(2)}) \right) \\ &= \mathbb{V}[X^{(1)} + X^{(2)}]. \end{aligned}$$

Finally, we have

$$\begin{aligned} \lim_{k \rightarrow \infty} \left(\text{MDI}_{\mathcal{T}_{k,1}^*}^*(X^{(1)}) - \text{MDI}_{\mathcal{T}_{k,2}^*}^*(X^{(1)}) \right) &= \lim_{k \rightarrow \infty} \left(\text{MDI}_{\mathcal{T}_{k,2}^*}^*(X^{(2)}) - \text{MDI}_{\mathcal{T}_{k,1}^*}^*(X^{(2)}) \right) \\ &= \frac{1}{3} - \frac{1}{3} \left(\frac{1}{4}\right)^\beta. \end{aligned}$$

\square

Proof of Technical Lemma 1. With the notations of Technical Lemma 1, we have

$$\begin{aligned} L_A^*(\ell, s) &= \mathbb{V}[Y|\mathbf{X} \in A] - \mathbb{P}[\mathbf{X} \in A_L | \mathbf{X} \in A] \mathbb{V}[Y|\mathbf{X} \in A_L] \\ &\quad - \mathbb{P}[\mathbf{X} \in A_R | \mathbf{X} \in A] \mathbb{V}[Y|\mathbf{X} \in A_R] \\ &= \mathbb{V}[m(\mathbf{X})|\mathbf{X} \in A] - \mathbb{P}[\mathbf{X} \in A_L | \mathbf{X} \in A] \mathbb{V}[m(\mathbf{X})|\mathbf{X} \in A_L] \\ &\quad - \mathbb{P}[\mathbf{X} \in A_R | \mathbf{X} \in A] \mathbb{V}[m(\mathbf{X})|\mathbf{X} \in A_R], \end{aligned}$$

since the noise ε is independent of \mathbf{X} and the probabilities $\mathbb{P}[\mathbf{X} \in A_L | \mathbf{X} \in A]$ and $\mathbb{P}[\mathbf{X} \in A_R | \mathbf{X} \in A]$ sum to one. Therefore,

$$\begin{aligned} L_A^*(\ell, s) &= \mathbb{E}[m^2(\mathbf{X})|\mathbf{X} \in A] - \mathbb{P}[\mathbf{X} \in A_L | \mathbf{X} \in A] \mathbb{E}[m^2(\mathbf{X})|\mathbf{X} \in A_L] \\ &\quad - \mathbb{P}[\mathbf{X} \in A_R | \mathbf{X} \in A] \mathbb{E}[m^2(\mathbf{X})|\mathbf{X} \in A_R] \\ &\quad - \mathbb{E}^2[m(\mathbf{X})|\mathbf{X} \in A] + \mathbb{P}[\mathbf{X} \in A_L | \mathbf{X} \in A] \mathbb{E}^2[m(\mathbf{X})|\mathbf{X} \in A_L] \\ &\quad + \mathbb{P}[\mathbf{X} \in A_R | \mathbf{X} \in A] \mathbb{E}^2[m(\mathbf{X})|\mathbf{X} \in A_R], \end{aligned} \tag{41}$$

where

$$\begin{aligned} &\mathbb{P}[\mathbf{X} \in A_L | \mathbf{X} \in A] \mathbb{E}[m^2(\mathbf{X})|\mathbf{X} \in A_L] + \mathbb{P}[\mathbf{X} \in A_R | \mathbf{X} \in A] \mathbb{E}[m^2(\mathbf{X})|\mathbf{X} \in A_R] \\ &= \mathbb{E}[m^2(\mathbf{X})\mathbb{1}_{\mathbf{X} \in A_L} + m^2(\mathbf{X})\mathbb{1}_{\mathbf{X} \in A_R} | \mathbf{X} \in A] \\ &= \mathbb{E}[m^2(\mathbf{X}) | \mathbf{X} \in A], \end{aligned} \tag{42}$$

which gives the conclusion by injecting (42) into (41). \square

References

- G. Biau. Analysis of a random forests model. *Journal of Machine Learning Research*, 13:1063–1095, 2012.
- A.-L. Boulesteix, A. Bender, J. Lorenzo Bermejo, and C. Strobl. Random forest gini importance favours snps with large minor allele frequency: impact, sources and recommendations. *Briefings in Bioinformatics*, 13:292–304, 2011.
- L. Breiman. Random forests. *Machine Learning*, 45:5–32, 2001.
- L. Breiman, J.H. Friedman, R.A. Olshen, and C.J. Stone. *Classification and Regression Trees*. Chapman & Hall/CRC, Boca Raton, 1984.
- Leo Breiman. Manual on setting up, using, and understanding random forests v3. 1. *Statistics Department University of California Berkeley, CA, USA*, 1:58, 2002.
- Manuel Fernández-Delgado, Eva Cernadas, Senén Barro, and Dinani Amorim. Do we need hundreds of classifiers to solve real world classification problems? *The Journal of Machine Learning Research*, 15(1):3133–3181, 2014.
- R. Genuer, J.-M. Poggi, and C. Tuleau. Random forests: some methodological insights. *arXiv preprint arXiv:0811.3619*, 2008.
- R. Genuer, J.-M. Poggi, and C. Tuleau-Malot. Variable selection using random forests. *Pattern Recognition Letters*, 31:2225–2236, 2010.

- Baptiste Gregorutti, Bertrand Michel, and Philippe Saint-Pierre. Correlation and variable importance in random forests. *Statistics and Computing*, 27(3):659–678, 2017.
- U. Grömping. Variable importance in regression models. *Wiley Interdisciplinary Reviews: Computational Statistics*, 7:137–152, 2015.
- H. Ishwaran. Variable importance in binary regression trees and forests. *Electronic Journal of Statistics*, 1:519–537, 2007.
- H. Ishwaran. The effect of splitting on random forests. *Machine Learning*, pages 1–44, 2013.
- X. Li, Y. Wang, S. Basu, K. Kumbier, and B. Yu. A debiased mdi feature importance measure for random forests. In *Advances in Neural Information Processing Systems*, pages 8049–8059, 2019.
- Wei-Yin Loh. Classification and regression trees. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 1(1):14–23, 2011.
- G. Louppe, L. Wehenkel, A. Suter, and P. Geurts. Understanding variable importances in forests of randomized trees. In *Advances in Neural Information Processing Systems*, pages 431–439, 2013.
- Laura L Nathans, Frederick L Oswald, and Kim Nimon. Interpreting multiple linear regression: A guidebook of variable importance. *Practical Assessment, Research & Evaluation*, 17(9), 2012.
- K. K. Nicodemus. Letter to the editor: On the stability and ranking of predictors from random forest variable importance measures. *Briefings in bioinformatics*, 12(4):369–373, 2011.
- K. K. Nicodemus and J. D. Malley. Predictor correlation impacts machine learning algorithms: implications for genomic studies. *Bioinformatics*, 25(15):1884–1890, 2009.
- Y. Qi. *Ensemble Machine Learning*, chapter Random forest for bioinformatics, pages 307–323. Springer, 2012.
- G. Rogez, J. Rihan, S. Ramalingam, C. Orrite, and P. H. Torr. Randomized trees for human pose detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- E. Scornet, G. Biau, and J.-P. Vert. Consistency of random forests. *The Annals of Statistics*, 43:1716–1741, 2015.
- C. Strobl and A. Zeileis. Danger: High power!—exploring the statistical properties of a test for random forest variable importance. 2008.
- C. Strobl, A.-L. Boulesteix, A. Zeileis, and T. Hothorn. Bias in random forest variable importance measures: Illustrations, sources and a solution. *BMC bioinformatics*, 8(1):25, 2007.

- C. Strobl, A.-L. Boulesteix, T. Kneib, T. Augustin, and A. Zeileis. Conditional variable importance for random forests. *BMC Bioinformatics*, 9:307, 2008.
- C. Strobl, T. Hothorn, and A. Zeileis. *Party on!* 2009.
- A. Suter, G. Louppe, V. A. Huynh-Thu, L. Wehenkel, and P. Geurts. Context-dependent feature analysis with random forests. *arXiv preprint arXiv:1605.03848*, 2016.
- Stefan Wager and Guenther Walther. Adaptive concentration of regression trees, with application to random forests. *arXiv preprint arXiv:1503.06388*, 2015.
- M. N. Wright, A. Ziegler, and I. R. König. Do little interactions get lost in dark random forests? *BMC bioinformatics*, 17:145, 2016.
- Z. Zhou and G. Hooker. Unbiased measurement of feature importance in tree-based methods. *arXiv preprint arXiv:1903.05179*, 2019.
- R. Zhu, D. Zeng, and M.R. Kosorok. *Reinforcement learning trees*. Technical Report, University of North Carolina, Chapel Hill, 2012.