# When does Subagging Work?

Christos Revelas Otilia Boldea Bas J.M. Werker C.REVELAS@TILBURGUNIVERSITY.EDU O.BOLDEA@TILBURGUNIVERSITY.EDU B.J.M.WERKER@TILBURGUNIVERSITY.EDU

Tilburg University Department of Econometrics and Operations Research Warandelaan 2, 5037 AB Tilburg, Netherlands

# Abstract

We study the effectiveness of subagging, or subsample aggregating, on regression trees, a popular non-parametric method in machine learning. First, we give sufficient conditions for pointwise consistency of trees. We formalize that (i) the bias depends on the diameter of cells, hence trees with few splits tend to be biased, and (ii) the variance depends on the number of observations in cells, hence trees with many splits tend to have large variance. While these statements for bias and variance are known to hold globally in the covariate space, we show that, under some constraints, they are also true locally. Second, we compare the performance of subagging to that of trees across different numbers of splits. We find that (1) for any given number of splits, subagging improves upon a single tree, and (2) this improvement is larger for many splits than it is for few splits. However, (3) a single tree grown at optimal size can outperform subagging if the size of its individual trees is not optimally chosen. This last result goes against common practice of growing large randomized trees to eliminate bias and then averaging to reduce variance.

**Keywords:** CART, regression trees, pointwise consistency, bias-variance trade-off, bagging, performance across sizes, performance at optimal sizes

## 1 Introduction

Decision trees are a popular method for non-parametric regression estimation in statistics, machine learning, economics and data science practice. Given a dataset, a decision tree partitions the covariate space and estimates the regression function locally, i.e., in each cell of the partition, by a simple parametric form, typically a mean. CART<sup>1</sup> is a prominent tree construction methodology. Bagging<sup>2</sup> is a randomization technique that consists of averaging multiple trees grown on bootstrap samples of the dataset. Early experiments suggested that bagging can improve the accuracy of learners, and in particular trees. On the one hand, this led to the development of randomization variants and extensions, such as random forests<sup>3</sup>. On the other hand, it led to many studies trying to *clarify whether and why randomization methods "work*".

© Christos Revelas, Otilia Boldea and Bas J.M. Werker.

License: CC-BY 4.0, see https://creativecommons.org/licenses/by/4.0/.

<sup>1.</sup> Short for classification and regression trees, introduced by Breiman et al., 1984. In the present paper, we refer to CART for *regression* trees constructed based on Breiman et al.'s methodology.

<sup>2.</sup> Bagging, introduced for CART by Breiman, 1996, is a blend word for bootstrap aggregating.

<sup>3.</sup> Random forests, introduced by Breiman, 2001, are averages of trees grown on bootstrap samples with an additional randomization of the covariate selection at each split.

Breiman, 1996 argued that bagging mimics having several datasets available for estimation and heuristically showed that bagging can improve upon unstable<sup>4</sup> learners, such as trees, through *aggregation*. These arguments hold on average over covariates. In his simulations, trees are fully grown and then pruned<sup>5</sup>, both for the single tree and the bagged estimator. Improvement is measured as a reduction in average, over covariates, out-ofsample mean-squared error. Bühlmann and Yu, 2002 were the first to establish rigorous results for subagging in the context of trees. They argued that instability<sup>6</sup> comes from hard decisions, such as indicators, and that randomization helps in reducing the variance through smoothing of hard decisions. They prove that subagging reduces mean-squared error locally around the split point in the case of stumps, i.e., trees with a single split. In their simulations, trees are grown large and without pruning, and improvement is also measured as a reduction in average mean-squared error. What these two studies have in common, is that they compare trees with (su)bagging trees of similar size<sup>7</sup>. Breiman compares pruned trees with bagging of pruned trees. Bühlman and Yu at first compare stumps with subagged stumps, then large trees with subagging of large trees. What neither of these two studies does, is to compare trees with (su)bagging across sizes. For example, how does (su)bagging small trees compare to a single large tree, and vice versa? Moreover, both studies place tree instability at the centre of (su)bagging's performance, but, are trees always unstable? If not, how does (su)bagging perform for stable trees?

There has been a tendency to grow large trees, i.e., with many splits, when considering ensemble methods such as (su)bagging and random forests<sup>8,9</sup>, the idea being that by growing large trees, bias is eliminated, and by averaging, the variance is reduced. This idea goes back to Breiman, 2001, who suggested that in a forest, trees should be fully grown and *not* be pruned, and claimed that random forests do *not* overfit<sup>10</sup>. As Segal, 2004 pointed out, while fully growing trees indeed reduces the bias, there is a trade-off with variance. He argued that the datasets used by Breiman where inherently difficult to overfit using forests and provided a counter-example in which indeed, contradicting Breiman, random forests overfit if grown deeply. In the same direction, Lin and Jeon, 2006 showed, by establishing

<sup>4.</sup> Defined in Breiman, 1996 as learners for which small changes in the data can lead to large changes in predictions: in other words, estimators with high variance.

<sup>5.</sup> Pruning, introduced in Breiman et al., 1984, is a procedure consisting of merging one-by-one the terminal nodes of a fully grown tree and choosing as final tree the one with smallest out-of-sample mean-squared error.

<sup>6.</sup> Defined in Bühlmann and Yu, 2002 as predictors that have non-zero asymptotic variance. Stability is defined locally in the covariates, contrary to Breiman, 1996 where, implicitly, stability refers to global, i.e., on average over covariates, low variance.

<sup>7.</sup> The size of a tree is defined as the number of terminal nodes, also known as leaves, or cells. The size of a tree is therefore equal to one plus the number of splits. Tree size is not to be confused with tree depth, defined as the maximum distance between the root and a terminal node. A tree of depth  $\delta$  can have up to  $2^{\delta}$  cells, and two trees of same depth can have different sizes.

<sup>8.</sup> See e.g. Hastie et al., 2009 and James et al., 2013.

<sup>9.</sup> See also the **randomForest** package in  $\mathbb{R}$ , version 4.7-1.1: the default minimum node size is five observations, and by default there is no bound on the number of nodes allowed.

<sup>10.</sup> Overfitting describes the situation in which a learner has small "training" mean-squared error, i.e., performs well in the sample used for estimation, and large "test" mean-squared error, i.e., performs poorly out-of-sample.

lower bounds on the rate of convergence of the prediction error of nonadaptive<sup>11</sup> random forests and via simulations for the classical, i.e., adaptive, forests, that growing large trees does not always give the best out-of-sample performance. In line with these studies, we ask the following question: how does (su)bagging perform, compared to optimally - in terms of bias and variance - grown trees, if its subtrees<sup>12</sup>, either are, or are not, also optimally grown?

The present paper contributes to the literature on tree-based methods in two ways.

First, we formalize the bias-variance trade-off associated with tree size by establishing *pointwise* consistency of trees under assumptions similar to those of Breiman et al., 1984. To our knowledge, pointwise consistency has not previously been established for CART. Trees are a particular case of local averaging estimators. A general proof of consistency for such estimators dates back to Stone, 1977 in the case where weights only depend on the covariates and not on the target variable. For data-driven partitions<sup>13</sup>, i.e., partitions in which weights depend on both covariates and the target variable. Breiman et al., 1984 give sufficient conditions for  $L^2$  consistency but do not show that CART satisfies those conditions. The first proof of consistency for CART was given in Scornet et al., 2015 in the context of random forests with subsampling instead of bootstrapping: they show  $L^2$ consistency assuming an additive regression function. Mentch and Hooker, 2016 showed pointwise asymptotic normality of subagging and random forests, but their limiting normal distributions are centered at the expected prediction and not necessarily the true regression function. Wager and Athey, 2018 showed pointwise consistency of random forests, but their result is not applicable to trees as they assume that the subsample size is asymptotically negligible compared to the dataset size. Recently Klusowski and Tian, 2023 also proved, among other things,  $L^2$  consistency for an additive regression function. In the present paper, in order to guarantee pointwise consistency, we assume in particular that the number of observations in a cell grows at a certain rate with the dataset size. The classical CART is not guaranteed to satisfy these assumptions and we give an explanation as to why that might be the case: depending on the dataset at hand, CART might not split enough in some regions of the feature space. However, we provide an algorithm<sup>14</sup> that satisfies our theorem's assumptions while still partitioning based on the CART criterion: we simply do not allow for splits to be performed if they were to give cells with fewer observations than a well-chosen lower bound. In some sense, we uniformize the number of observations across *cells* in order to guarantee pointwise consistency. In our proof, we use honesty<sup>15</sup> to explicitly calculate the bias and variance of a tree estimate *locally*, i.e., at a given value of interest. We show that the bias depends on the diameter of the partition's cell containing this value of interest, and that the variance depends on the number of observations inside the cell. This allows us to formalize the bias-variance trade-off associated with tree size. In simulations, we illustrate this trade-off by comparing "small", "large" and "consistent" trees. Moreover,

<sup>11.</sup> Term used in Lin and Jeon, 2006 to describe a random forest the partition of which is independent of the target variable.

<sup>12.</sup> In the present paper, we use the term *subtree* to describe a tree grown on a subsample.

<sup>13.</sup> See Györfi et al., 2002's chapter on data-dependent partitioning for general consistency results.

<sup>14.</sup> See Implementation 1 in Section 2.2.

<sup>15.</sup> Honesty is a concept that allows to get rid of some dependencies to the data when constructing an estimator. See e.g. Athey and Imbens, 2016 for a use of honesty in the context of trees.

we point out that consistency implies stability<sup>16</sup>, which means that trees can be stable, if appropriately grown. This goes against the commonly stated view that "randomization works because trees are unstable"<sup>17</sup>. We show in simulations that subagging can improve consistent trees by variance reduction. In other words, we find that subagging can also improve upon stable learners.

Second, we study in simulations how subagging<sup>18</sup> performs for trees of different sizes. We measure performance in terms of mean-squared error. We fix the same number of splits among a single tree and every randomized tree constituting subagging when comparing the performance of the two methods. This can be obtained using, e.g., the **randomForest** package in  $\mathbf{R}$  but we implemented our own versions of CART and subagging<sup>19</sup> in order to extract additional information. Starting with stumps, we look at the effect of subagging on the weight of observations around the value of interest. We illustrate that subagging assigns positive weight to observations that had zero weight in the tree, conditionally on the dataset. To our knowledge, we are the first to show this explicitly for CART. The idea however that bagging stabilizes predictions by *balancing* the importance of observations in the data goes back to Grandvalet, 2004. They argued that bagging can improve or deteriorate the base learner depending on the quality of influential observations in the context of point estimation and regression, but not for trees. In Grandvalet, 2006 classification trees are considered, but the effect of bagging on weights is not illustrated as in the present paper. Then, we find that subagging reduces the variance of a tree around split points, which is in line with Bühlmann and Yu, 2002's "instability regions"<sup>20</sup>. The more we split, the larger the instability region, and improvement is larger for trees with many splits than it is for small trees. However, a single tree grown at optimal size can outperform subagging with large trees. Therefore, subagging large trees is not always a good idea. Nonetheless, subagging can still improve upon a single tree if both methods are optimally grown, and hence to be preferred in practice. Additionally, both for the single tree and subagging, our simulations suggest that there is a *linear relation* between the number of observations in the data and the optimal number of splits. In practice, this means that one can first find the optimal size for a single tree in the usually way, e.g., with cross-validation, and then deduce the right size for the ensemble method at hand.

**Related work.** We are not the first to consider the importance of size in tree-based methods. The closest existing study that we are aware of is Zhou and Mentch, 2023<sup>21</sup>. They argue that tree depth<sup>22</sup> has a regularizing effect on randomized ensembles. They compare the performance of random forests versus bagging as a function of size for different signal-to-noise ratios and find that small trees are advantageous when the ratio is low while larger

<sup>16.</sup> Here stability is defined as in Bühlmann and Yu, 2002, i.e., locally and asymptotically.

<sup>17.</sup> For example, Soloff et al., 2023 have established "algorithmic" stability of bagging, a finite sample definition that formalizes the heuristic definition originally given by Breiman. They show in simulations that subagging has a highly stabilizing effect on regression trees, concluding that trees are very unstable, but their trees are grown large in the first place: they use a depth of 50 for a dataset of 500 observations.

Subagging has previously been used instead of bagging in the context of trees in e.g. Scornet et al., 2015, Wager and Athey, 2018 and Klusowski and Tian, 2023.

<sup>19.</sup> See Implementation 2 in Section 2.2, and Section 7.2 for replication details.

<sup>20.</sup> Defined as a neighbourhoods around the split points. See Section 5.2.

<sup>21.</sup> Our respective starting problematics partially overlap, as both build on the observations made by Segal, 2004 and Lin and Jeon, 2006.

<sup>22.</sup> Note that, in their simulations, size, rather than depth, is considered, "as a proxy for depth".

trees should be used when the ratio is high. Duroux and Scornet, 2018 show that the global mean-squared error of fully grown quantile forests<sup>23</sup> with subsampling and that of small randomized trees without subsampling have the same bounds. Both studies always keep the randomization of feature selection in their simulations, while we compare subagging with deterministic, i.e., non-random given the data, trees. Very recently, Curth et al., 2024 were able to quantify the degree of smoothing of forests, by looking at forests through the lens of nearest neighbours (Lin and Jeon, 2006). Moreover they bridge a gap between the notions of bias in statistics and in machine learning (Dietterich and Kong, 1995) and disentangle multiple distinct mechanisms through which forests improve on trees. In relation to tree size, in their Figure 4, the smoothing effect of ensembling is compared for small and large trees, but trees are also randomized through feature selection.

**Organization.** Section 2 sets the statistical framework and notation, and discusses the considered methods. In particular we give some intuition as to why CART "avoids" splitting near the edges of the covariate space. Moreover, we heuristically establish a relation between the bias and variance of subagging to that of a single tree grown on the full sample<sup>24</sup>. Our theorem for pointwise consistency of trees, and an illustration of the bias-variance trade-off associated with tree size, are given in Section 3. The effect of subagging on consistent trees is illustrated in Section 4. Section 5 examines small trees. We illustrate the effect of subagging on stumps conditionally on the data as well as the effect of subagging if the size of its subtrees is not well chosen. The optimal number of splits as a function of the dataset size is also shown. The effect of subagping size and bootstrapping, as well as the replication of our results with readily-available implementations, are finally presented in Section 7. Section 8 concludes and proofs are gathered in appendix.

#### 2 Decision-Tree Methods for Regression

We consider a regression model of the form

$$Y = f(X) + \varepsilon \tag{1}$$

with  $(X,Y) \in [0,1]^p \times \mathbb{R}$  and  $\varepsilon \in \mathbb{R}$  such that  $\mathbb{E}[\varepsilon|X] = 0$  and  $\mathbb{E}[\varepsilon^2|X] = \sigma^2$ . Given an random sample  $D_n = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$  drawn from (1), we want to construct an estimator of  $f(x) = \mathbb{E}[Y|X = x]$  based on  $D_n$ .

A decision tree partitions  $[0, 1]^p$  into rectangular cells and, for any x, gives as estimator for f(x) the average, noted  $T_n(x)$ , of all the  $Y_i$ 's that fall in the same cell as x:

$$T_n(x) = \sum_{i=1}^n W_{n,i}(x) Y_i \quad \text{with} \quad W_{n,i}(x) = \frac{\mathbbm{1}_{X_i \in C_n(x)}}{\sum_{j=1}^n \mathbbm{1}_{X_j \in C_n(x)}}$$
(2)

where  $C_n(x)$  is the cell that contains x. A decision tree, given  $D_n$ , is fully deterministic. The partition depends on  $D_n$  and is obtained by recursive binary splitting of  $[0,1]^p$  based

<sup>23.</sup> A special type of forests in which splits only depend on the covariates but not on the target.

<sup>24.</sup> We do not expect this relation to be difficult to formalize, but leave it for future research.

on a splitting criterion and a stopping rule. In the case of CART, the splitting criterion is the maximization over splits of the *sample analogue* of

$$\mathbb{V}[Y|X \in C] - \frac{\mathbb{P}(X \in C_L)}{\mathbb{P}(X \in C)} \mathbb{V}[Y|X \in C_L] - \frac{\mathbb{P}(X \in C_R)}{\mathbb{P}(X \in C_R)} \mathbb{V}[Y|X \in C_R]$$
(3)

where  $\mathbb{V}$  denotes the variance, C is a cell to be split and  $C_L$  and  $C_R$  are the two cells that we obtain after splitting C. CART splits aim at reducing the variation in  $Y_i$ 's. The stopping rule can be the number of splits, a minimum number of observations in each cell, or a minimum gain in variance reduction.

Bagging is an ensemble method which gives as estimator for f(x) the average, noted  $\overline{T}_n^*(x)$ , of multiple trees, which are i.i.d. conditionally on  $D_n$ :

$$\bar{T}_{n}^{*}(x) = \frac{1}{B} \sum_{b=1}^{B} \sum_{i=1}^{n} W_{n,i,b}^{*}(x) Y_{i}$$
(4)

where B is the number of trees and  $W_{n,i,b}^*(x)$  is the random weight for  $Y_i$  in the  $b^{\text{th}}$  tree. Randomization comes from bootstrapping: for every b, we generate a bootstrap sample  $D_n^*(b)$  of  $D_n$ , i.e., a sample of n observations drawn independently and with replacement, and grow a tree on  $D_n^*(b)$ . Defining  $W_{n,i}^*(x) = \frac{1}{B} \sum_{b=1}^{B} W_{n,i,b}^*(x)$ , then (4) can be re-written as

$$\bar{T}_{n}^{*}(x) = \sum_{i=1}^{n} W_{n,i}^{*}(x) Y_{i}.$$
(5)

Note that all weights sum up to one:  $\sum_{i=1}^{n} W_{n,i}(x) = 1$ ,  $\sum_{i=1}^{n} W_{n,i,b}^*(x) = 1$  for all b and  $\sum_{i=1}^{n} W_{n,i,i}^*(x) = 1$ . Defining  $T_{n,b}^*(x) = \sum_{i=1}^{n} W_{n,i,b}^*(x)Y_i$ , then (4) can also be re-written as

$$\bar{T}_{n}^{*}(x) = \frac{1}{B} \sum_{b=1}^{B} T_{n,b}^{*}(x).$$
(6)

In this paper we consider subagging, where bootstrap sampling is replaced by subsampling of size k: each replicate  $D_n^*(b)$  of  $D_n$  on which a subtree is grown is obtained by drawing, without replacement,  $k \leq n$  observations from  $D_n$ . Several studies, e.g., Bühlmann and Yu, 2002, Grandvalet, 2004, Buja and Stuetzle, 2006, Friedman and Hall, 2007, have suggested that subagging with half the observations has a similar performance to bagging. Buja and Stuetzle, 2006 proved more than that in the case where the base learner is a U-statistic (Hoeffding, 1992), and supported with simulations that the same may hold for trees: bagging with  $\alpha_w n$  observations is equivalent to subagging with  $\alpha_{w/o} n$  observations if  $\alpha_w = \frac{\alpha_{w/o}}{1-\alpha_{w/o}}$ . In such case, subagging has the advantage of being computationally cheaper than bagging (because trees are grown on samples of smaller size). In the present paper we use subagging with half the observations, i.e., we take k = 0.5n, and show, in Section 7.1, that our simulation results are still valid when we vary the subsample size k as well as when we use bagging instead of subagging.

#### 2.1 CART Criterion and the Location of Splits

CART, as opposed to other partitioning estimators such as kernel regression (Nadaraya, 1964; Watson, 1964), do *not* split the covariate space uniformly. In order to better understand how CART splits behave, we show the following.

**Proposition 1** The criterion (3) can be re-written as

$$\frac{\mathbb{P}(X \in C_L)\mathbb{P}(X \in C_R)}{\mathbb{P}(X \in C)} \{\mathbb{E}[Y|X \in C_L] - \mathbb{E}[Y|X \in C_R]\}^2.$$
(7)

A proof is given in the appendix. Expression (7) is the product of two factors. To understand them, consider the one-dimensional case (p = 1), assume X is uniformly distributed in [0, 1], and let c and  $c_l$  be the lengths of intervals C and  $C_L$  respectively. The first factor in (7) is then  $\frac{c_l(c-c_l)}{c}$ . The length c is fixed when we search for a split inside C, and maximizing  $c_l(c-c_l)$  over  $c_l \in [0,c]$  gives  $c_l = \frac{c}{2}$ . In other words, the first factor in (7) is maximized in the middle of C, and this is true irrespectively of the distribution of Y. The second factor is more complex as it depends on f. Note that  $\mathbb{E}[Y|X \in C_L] = \mathbb{E}[f(X)|X \in C_L]$ and similarly for  $C_R$ . If f is constant, then the second factor in (7) is zero for any split, therefore the entire expression equals zero for any split. If f is linear, then the second factor is constant<sup>25</sup> for any split and therefore (7) is maximized in the middle of C (because the first factor is maximized in the middle of C). Suppose that  $f(x) = x^2$ . Then the second factor is maximized where f is steeper, i.e., at the extremity x = 1. The first factor being maximized in the middle, i.e., at  $x = \frac{1}{2}$ , this brings the overall argmax at x = 0.64, which is away from the extremity x = 1. In other words, CART tends to split away from the boundaries of a cell, and, by extension, CART tends to split away from the boundaries of the feature space. This supports previous knowledge (Wager and Athey, 2018) that CART can be inconsistent at the boundaries of the feature space. Intuitively, not splitting close to the boundaries implies that some cells (precisely those at the boundaries) remain large and hence tree estimates inside those cells will tend to be biased. In order to guarantee pointwise consistency we therefore need somehow to force trees to sometimes split close to the boundary, in order to guarantee that cells become smaller throughout the feature space. *How* we enforce it is given in Section 3.

CART partitions the feature space not based on (7) but instead based on its empirical analogue, which incorporates the noise variable  $\varepsilon$ . On the one hand, the larger the noise variance  $\sigma^2$ , the more a CART split can deviate from its theoretical counter-part obtained by maximizing (7). On the other hand, the more the observations in a given cell, the closer a CART split will be to its theoretical counter-part.

#### 2.2 Stopping Rules and Tree Size

A decision tree is obtained by recursive binary splitting of the data. A CART tree cannot be grown further if any of its cells either contains a single observation or yields the same CART criterion for every possible split. A *stopping rule* is a constraint that makes the tree

<sup>25.</sup> To see that the terms in s cancel each other out, use the identity  $(u-v)^2 = (u+v)(u-v)$ .

stop growing earlier. When there is no stopping rule, we say that the tree is *unconstrained*. A tree is *fully grown* when each cell contains a single observation. In Proposition 2 we show that, *if unconstrained*, *CART trees are necessarily fully grown*. More precisely, let C be a cell and let  $\nu$  be the number of observations from  $D_n$  that are in C. Denote L(s) the CART criterion at a split value s.

### **Proposition 2** If $\nu \geq 3$ , then there exists $s_1$ and $s_2$ such that $L(s_1) \neq L(s_2)$ .

In other words, the CART criterion cannot be constant in any cell with at least three observations. A proof is given in the appendix. Because fully grown trees have large variance, stopping rules are needed to guarantee consistency. In this paper we consider two stopping rules. First, referred to as *Implementation 1* hereafter, we use bounds on the number of observations in a cell. Second, referred to as Implementation 2 hereafter, we control for the number of splits. Bounds on the number of observations indirectly also control for the number of splits: if we start with a hundred observations and constrain cells to have between ten and twenty observations, then there can only be between five and ten splits. Bounds on the number of observations to some extent also control the location of splits. Indeed, consider the one-dimensional case (p = 1) and the extreme scenario where we impose a lower bound of fifty observations per cell. Then there is only one admissible split, leaving exactly fifty observations on both sides of the split, precisely between the fiftieth and the fifty-first observations ordered with respect to their X value. Conversely, controlling the number of splits controls the *average*, across cells, number of observations. For example, four splits of a hundred observations will give five cells. Noting the number of observations in each cell by  $c_1, ..., c_5$ , then  $\sum_{i=1}^{5} c_i = 100$ , i.e.,  $(\sum_{i=1}^{5} c_i)/5 = 20$ . The greater the number of splits, the smaller the average number of observations. However, controlling the number of splits alone, gives no control over the location of splits: they are chosen based on the CART criterion.

Implementation 1: control the number of observations. We implemented a recursive function that takes as input a dataset  $D_n$  and a minimum cell size  $h_n$  and returns a partition in which each cell contains at least  $h_n$  and at most  $2h_n-1$  observations. To do so, at each splitting step, we define admissible splits as those that leave at least  $h_n$  observations on each side and we choose among the admissible splits the one that maximizes the CART criterion. On the one hand, this implementation plays an important role in our results on consistency of Section 3: we show that consistency is guaranteed if  $h_n$  is of the form  $n^{\alpha}$  for some  $\alpha > \frac{1}{2}$ . By controlling for both the minimum and maximum number of observations allowed in a cell, we are able to guarantee pointwise consistency: we uniformize number of observations the best split in terms of CART criterion, which can happen if such split would yield a cell with fewer than  $h_n$  observations.

Implementation 2: control the number of splits. Also recursive, takes as input a dataset  $D_n$  and a number of splits N and returns a partition with *exactly* N splits. This implementation is used in later sections where we study the performance of trees and subagging as a function of the number of splits. It always finds the best split in terms of CART criterion, but it does not provide exact control over the size of cells obtained from such splits. This implementation replicates existing implementations in  $\mathbb{R}^{26}$ , but we used

<sup>26.</sup> See Section 7.2 for details and an illustration.

it in order to be able to extract and illustrate the balancing effect of subagging on weights (see Section 5.1).

**Tree size.** We say that a tree is *small*, if either  $h_n$  is large or N is small. The smallest possible tree is called a *stump*, for which  $h_n = 0.5n$  and N = 1. Conversely, we say that a tree is *large*, if either  $h_n$  is small relative to n or if N is large relative to n. The largest possible tree is the fully grown tree for which  $h_n = 1$  and N = n - 1. In Section 3 we look at a particular range for  $h_n$  that guarantees consistency of trees and in Section 4 we look at the effect of subagging on consistent trees. The case N = 1 is detailed in Section 5. Section 6 shows the performance of subagging compared to trees as a function of N.

### 2.3 Bias and Variance of Subagging

Here we heuristically establish a relation between the statistical bias and variance of subagging and that of a single tree grown on the *full* sample. Fix  $x_0$  a feature value of interest and let  $T_n := T_n(x_0)$  and  $\bar{T}^*_{n,k} := \bar{T}^*_{n,k}(x_0)$  denote the tree and subagged estimates for  $f(x_0)$ as in (2) and (6) respectively, simplifying in notation the dependency in  $x_0$  and making explicit the dependency in the subsample size k. Also note  $T^*_{n,k,1}, \ldots, T^*_{n,k,B}$  the subtrees constituting  $\bar{T}^*_{n,k}$ .

### **Proposition 3** We have

$$\mathbb{E}[\bar{T}_{n,k}^*] = \mathbb{E}[T_{n,k,1}^*] \tag{8}$$

and

$$\mathbb{V}[\bar{T}_{n,k}^*] = \mathbb{V}\mathbb{E}[T_{n,k,1}^*|D_n] + \frac{1}{B}\mathbb{E}\mathbb{V}[T_{n,k,1}^*|D_n].$$
(9)

The proof follows immediately<sup>27</sup> from the fact that  $T_{n,k,1}^*, \ldots, T_{n,k,B}^*$  are i.i.d. conditionally on  $D_n$ .

**Bias.** From (8) we deduce that the bias of subagging with subsample size k is the same as the bias of a single of its subtrees, i.e., a tree grown on a subsample of size k. This in turn implies that, for large n, we expect subagging with subsample size k to have similar bias to a single tree grown on the **full** sample, as long as we use the same number of splits for both methods. To our knowledge, we are the first to point this out. In later sections we support this statement with simulations. Here we give an informal explanation. Assume that k grows proportionally with n, for example, k = 0.5n. If n is large, then k is large as well. In this case, a tree grown on the entire sample (n observations) and a tree grown on a subsample (k observations) will give splits that are close to each other. Therefore, the cell containing  $x_0$  will be of similar diameter for both trees. In Section 3, we show that the bias of a tree at  $x_0$  depends on the diameter of the cell that contains  $x_0$ . Therefore, since a tree grown on n observations and a tree grown on k observations give cells of similar diameter for large n, we expect indeed both trees to have similar bias.

**Variance.** Equation (9) is more complex. Again, we proceed with an informal treatment. If the number of subtrees, or subsamples, B, is large, then the second term in (9) is

<sup>27.</sup> Use that  $\mathbb{E}[A] = \mathbb{E}[\mathbb{E}[A|B]]$  and  $\mathbb{V}[A] = \mathbb{V}[\mathbb{E}[A|B]] + \mathbb{E}[\mathbb{V}[A|B]]$  for random variables A and B.

negligible. The first term, i.e.,  $\mathbb{VE}[T_{n,k,1}^*|D_n]$ , is the variance of a U-statistic. Indeed,

$$\mathbb{E}[T_{n,k,1}^*|D_n] = \frac{1}{\binom{n}{k}} \sum_{1 \le i_1 < \dots < i_k \le n} T_{i_k}$$
(10)

where  $T_{i_k} := T_{i_k}(X_{i_1}, Y_{i_1}, \ldots, X_{i_k}, Y_{i_k})$  is the prediction for  $f(x_0)$  based on a single, deterministic given  $D_n$ , tree grown on  $(X_{i_1}, Y_{i_1}), \ldots, (X_{i_k}, Y_{i_k})^{28}$ . Because (10) is a projection, its variance is smaller than the variance of a single tree based on  $(X_1, Y_1), \ldots, (X_k, Y_k)^{29}$ . Following the same reasoning as for the bias, for large n, the variance of the subagged estimator with subsample size k cannot exceed the variance of a single tree estimator grown on the **full** dataset by more than a factor of  $\frac{n}{k}$ , provided that the same number of splits is used for both methods. To our knowledge, we are the first to state this. Indeed, in Section 3 we show that the variance of a single tree is inversely related to the number of observations used to estimate  $f(x_0)$ . For large n, a tree on the full sample and a tree grown on a subsample will give approximately the same splits, and hence, if  $\nu$  is the number of observations in the cell of interest for the full tree, then the subtree will give cells of roughly  $\frac{k}{n}\nu$  observations. In particular, if a single tree has close to zero variance for some  $x_0$ , something that happens away from split points, as we will see Section 5, then subagging will also have close to zero variance around the same  $x_0$ .

# 3 Pointwise Consistency of Trees and the Bias-Variance Trade-Off Associated with Tree Size

In line with other non-parametric methods, such as kernel regression, the main intuition behind consistency still holds for trees: the finer the partition of the feature space, the smaller the bias<sup>30</sup> of the estimator, while its variance decreases when the number of observations in each cell increases. Therefore, we want trees to be grown deeper when the dataset size n increases, but slowly enough so that the number of observations in each cell increases when n increases. In this section, we consolidate this intuition by showing in Theorem 1 that trees are pointwise consistent when grown under some constraints. In particular, our first implementation described in Section 2.2 satisfies these constraints. We start with some assumptions.

Assumption 1 (DGP)  $X \sim \mathcal{U}([0,1]^p)$ , f continuous,  $\mathbb{E}[\varepsilon|X] = 0$ ,  $\mathbb{E}[\varepsilon^2|X] = \sigma^2$ .

The continuity of f together with the compactness of the feature space allow us to calculate limits of integrals.

Assumption 2 (honesty) For all n, for all i,  $Y_i \perp \!\!\!\perp W_{n,i} | X_1, \ldots, X_n$ .

In other words, for all i,  $Y_i$  is independent of  $W_{n,i}$  conditionally on  $X_1, \ldots, X_n$ . Honesty allows us to compute the bias and variance of the tree estimator locally.

<sup>28.</sup> Note that if we would define subagging as the average over **all** possible subsamples of size k, then its variance would be precisely the variance of (10). In fact subagging would be a U-statistic. See Mentch and Hooker, 2016 for an in-depth analysis of subagging as U-statistics.

<sup>29.</sup> See for example Section 12.1 in Van der Vaart, 2000.

<sup>30.</sup> To be precise, this holds locally where the regression function being estimated is monotone.

Assumption 3 (number of observations) There exists  $\frac{1}{2} < \alpha < 1$  such that for all n and all  $x_0$ , noting  $h_n = n^{\alpha}$ , almost surely

$$h_n \le \sum_{i=1}^n \mathbb{1}_{X_i \in C_n(x_0)} \le 2h_n - 1 \tag{11}$$

where  $C_n(x_0)$  again denotes the cell containing  $x_0$ . This assumption guarantees that the empirical measure of cells tends to zero, and hence, by the Glivenko-Cantelli theorem, the Lebesgue measure of cells will also tend to zero.

Assumption 4 (diameter) For all  $x_0$ , almost surely  $\bigcap_{n>1} C_n(x_0) = \{x_0\}$ .

This additional assumption is imposed to avoid the scenario in which from a certain point onwards, all splits are performed along the same direction. For example, if p = 2, this would mean obtaining asymptotically a segment in the two-dimensional space, in which case our estimate for  $f(x_0)$  would be the average of f over the segment, and hence not necessarily equal to  $f(x_0)$ . Although Assumption 4 is useful for our proof, the above scenario does not occur in simulations.

# **Theorem 1** Under Assumptions 1-4, for all $x_0$ , $T_n(x_0) \to f(x_0)$ in probability.

Here we give a sketch of the proof of Theorem 1 in order to illustrate the bias-variance tradeoff between cell size and number of observations. A complete proof is given in the appendix. Using honesty, we calculate the squared bias and the variance of the tree estimator at a given  $x_0$ . We obtain three terms:

• the squared bias of  $T_n(x_0)$ , of the form

$$\left(\mathbb{E}\left[\frac{\int_{C_n(x_0)} f(x)dx}{\int_{C_n(x_0)} dx}\right] - f(x_0)\right)^2,\tag{12}$$

• the variance of the error term  $\sum_{i=1}^{n} W_{n,i}(x_0)\varepsilon_i$ , of the form

$$\mathbb{E}\left[\frac{\sigma^2}{\sum_{i=1}^n \mathbbm{1}_{X_i \in C_n(x_0)}}\right],\tag{13}$$

• and the variance of the regression term  $\sum_{i=1}^{n} W_{n,i}(x_0) f(X_i)$ , of the form

$$\mathbb{E}\left[\frac{\int_{C_n(x_0)} f^2(x) dx}{\int_{C_n(x_0)} dx}\right] - \mathbb{E}\left[\frac{\int_{C_n(x_0)} f(x) dx}{\int_{C_n(x_0)} dx}\right]^2.$$
(14)

Consistency is obtained when all three terms converge to zero. On the one hand, the diameter of  $C_n(x_0)$  needs to tend to zero as *n* increases, so that the terms (12) and (14) tend to zero. Indeed, if that's the case, we use the uniform continuity of *f* to guarantee for example that  $\frac{\int_{C_n(x_0)} f(x)dx}{\int_{C_n(x_0)} dx}$  converges to  $f(x_0)$ . On the other hand, the number of observations in



Figure 1: (In)consistency of trees: on the x-axis is n; the solid (resp. dotted) line represents the sample mean (resp. mean  $\pm$  one standard deviation) of the tree estimate for  $f(x_0)$  (grey line) in each scenario a), b) and c).

 $C_n(x_0)$  needs to tend to infinity, so that the term (13) converges to zero. Therefore, the diameter of  $C_n(x_0)$  should not tend to zero too quickly. This summarizes the bias-variance trade-off when choosing how much to grow a regression tree.

Small trees, as defined in Section 2.2, generate partitions of large cells, hence tend to have small variance while remaining biased. Large trees generate finer partitions, and will tend to have small bias but large variance. Trees that satisfy Assumption 3 are somewhere in between small and large trees: they tend to have a smaller bias compared to small trees and a smaller variance compared to large trees. We illustrate this with a simulation based on our first implementation of Section 2.2. We consider three scenarios:

- a) small tree:  $h_n = \frac{n}{3}$  which gives a small number of splits for any n,
  - b) consistent tree:  $h_n = n^{0.65}$ , i.e., satisfying Assumption 3, and
- c) large tree:  $h_n = 4$ , i.e., cells have between 4 and 7 observations for any n.

Simulation I. We take  $f(x) = x^2$ ,  $X \sim \mathcal{U}(0, 1)$  and  $\varepsilon \sim \mathcal{N}(0, 0.2^2)$  in (1). We generate 200 datasets of size 2000, all of which have the same realization of  $X_i$ 's. For each dataset, and for  $n = 50, 100, 150, \ldots, 2000$ , we grow a tree on the first n of the 2000 observations following each of the three scenarios a,b,c). Figure 1 shows the empirical mean and the mean  $\pm$  one standard deviation conditionally on X of the tree estimate for  $f(0.5) = 0.5^2$  across replicates for each n and each scenario. The left plot shows scenario a): small trees. The middle plot shows scenario b): consistent trees. The right plot shows scenario c): large trees. These graphs illustrate the bias-variance trade-off: small trees have low variance but are biased while large trees are unbiased but have high variance. Trees satisfying Assumption 3 are in-between: they are unbiased *and* have low variance.

Figure 2 shows the corresponding biases, variances, and mean-squared errors. The mean squared error at  $x_0 = 0.5$  converges to zero for the consistently grown tree while it does not converge to zero for the small and large trees. In the case of the small tree, the bias does not vanish. For the large tree, the variance does not vanish.



Figure 2: Bias-variance trade-off associated with tree size.

### 4 Subagging Consistent Trees

We show via simulation that subagging consistent - and hence stable - trees does not affect the bias while it can improve the variance. We saw in Theorem 1 that a tree grown on the full sample is consistent when the minimum cell size  $h_n$  grows appropriately with n, i.e., is of the form  $n^{\alpha}$  for some  $\alpha > \frac{1}{2}$ . To guarantee consistency of each subtree, we similarly choose  $h_k$  of the form  $k^{\alpha}$ , again for some  $\alpha > \frac{1}{2}$ . Then the subagged estimator, which is an average of such subtrees, will also be consistent.

Keeping the same  $\alpha$  for the tree and the subagged estimator implies that cells will be of similar diameter for both estimators and hence, based on (12), we expect the two estimators to have similar bias. Each subtree will have cells of fewer observations than the original tree and therefore, based on (13), each subtree is expected to have a larger variance than the original tree. However the subagged estimator is an average of subtrees, which means that ultimately the average is taken over more observations compared to the original tree, and hence we expect a reduction in variance. The effect on variance is further examined in Section 5.

Figure 3 shows the effect of subagging on consistent trees in the same simulation as in Figure 1 (bottom plot) for scenario b), i.e.,  $h_n = n^{0.65}$ . Here the subagged estimator is defined as the average of B = 50 trees<sup>31</sup>, each of which is grown on a subsample of size k = 0.5n and with a minimum cell size of  $h_k = k^{0.65}$ . We observe indeed that, in terms of bias, the two estimators behave similarly, while subagging reduces variance.

### 5 Subagging Small Trees

In this section we look at stumps, i.e., single-split trees, as proxy for small trees, because this allows us to explain the effect of subagging on weights. To do so, we start with an analysis conditionally on  $D_n$ . We show that conditionally on  $D_n$ , subagging increases the number of distinct observations used to estimate  $f(x_0)$  compared to a single tree, and these observations cover a wider part of the feature space. Additionally, the closer  $x_0$  is to the split

<sup>31.</sup> Extensive empirical evidence exists in the literature that the number of trees need not be large in general. In our simulations, increasing B to 100 or 150 did not bring any change to the results.



Figure 3: Subagging consistent trees: on the x-axis is n; the solid (respectively dashed) line represents the squared bias (left plot), variance (middle) and mean squared error (right) of the tree (respectively subagged tree) estimates for  $f(x_0)$  when consistently grown ( $\alpha = 0.65$ ).

point, the more subagging adds weight to observations that had zero weight in the single tree. Then we look at the tree and subagged estimates conditionally on X. We support existing knowledge, e.g., Bühlmann and Yu, 2002, that aggregating improves upon single trees by reducing the variance around the split point, region in which a stump has high variance.

#### **5.1** Analysis Conditionally on $D_n$

Fix  $D_n$  and  $x_0$  and take p = 1. Given  $D_n$ , the first split of a decision tree is fixed and partitions [0, 1] in two regions, one of which contains  $x_0$ . Let  $s \in [0, 1]$  be the split point and assume, without loss of generality, that  $x_0 < s$ . Then the tree estimate for  $f(x_0)$ , noted  $T_n(x_0)$ , is the average of all observations in  $D_n$  such that  $X_i \leq s$ . Now let  $D_n^*(b)$  be a subsample of  $D_n$  and let  $s^*$  be the first split point of a tree grown on  $D_n^*(b)$ . There are three possible scenarios:

A) 
$$x_0 < s^* < s$$
,  
B)  $x_0 < s < s^*$ ,  
C)  $s^* < x_0 < s$ .

First, assume that  $x_0 < s^*$ . Then the subtree estimate for  $f(x_0)$  based on  $D_n^*(b)$ , noted  $T_{n,b}^*(x_0)$ , is the average of all observations in  $D_n$  such that  $X_i \leq s^*$  and  $X_i \in D_n^*(b)$ . In case  $x_0 < s^* < s$ , then  $T_{n,b}^*(x_0)$  is also an average of observations that are in  $T_n(x_0)$  but fewer of them. When  $x_0 < s < s^*$ , then  $T_{n,b}^*(x_0)$  is an average of observations that are in  $T_n(x_0)$  but fewer of them. When  $x_0 < s < s^*$ , then  $T_{n,b}^*(x_0)$  is an average of observations that are in  $T_n(x_0)$  and some observations that are not in  $T_n(x_0)$  since they are such that  $s < X_i$ . The same holds for the case where  $s^* < x_0 < s$ . The subagged estimator being an average of several subtrees, if we take enough subsamples,  $\overline{T}_n^*(x_0)$  will be an average of observations such that  $X_i < s$  and observations such that  $X_i > s$ . In other words, conditionally on  $D_n$ , subagging increases the number of distinct observations used to estimate  $f(x_0)$  compared to a single tree, and these observations cover a wider part of [0, 1]. In order to illustrate the effect of subagging on weights, we consider the following simulation with n fixed.

Simulation II. We use our second implementation as described in Section 2.2. We generate one realization of n = 100 observations from the same DGP (1) as in Simulation I, i.e., with  $f(x) = x^2$ ,  $X_i$  i.i.d.  $\mathcal{U}[0,1]$  and  $\varepsilon_i$  i.i.d.  $\mathcal{N}(0,0.2^2)$ . The first (left) plot in Figure 4 shows the estimate obtained from a stump (in blue) and a subagged stump (in red). We use subsamples of size k = 0.5n and average over B = 50 randomized trees to get the subagged estimates. The optimal CART split<sup>32</sup> is at s = 0.63.<sup>33</sup> Given  $D_n$  and  $x_0$ , the tree weights  $\{W_{n,i}(x_0)\}_{i=1,...,n}$  are deterministic. They are constant and equal to  $\frac{1}{\sum_{j=1}^{n} \mathbb{I}(X_j < s)}$  for all  $Y_i$  such that  $X_i < s$  and equal to zero elsewhere. For a subsample b, the weights  $\{W_{n,i,b}(x_0)\}_{i=1,...,n}$  are random conditionally on  $D_n$ . They are constant for some (because of subsampling)  $Y_i$ 's such that  $X_i < s^*$ , and zero elsewhere. The average over subsamples, i.e. the weights  $W_{n,i}^*(x_0)$ , are also random conditionally on  $D_n$ , and are as follows.

The case where  $x_0$  is far from s. The second plot in Figure 4 shows the weights associated with  $x_0 = 0.1$ , which is represented by the first (from left to right) perpendicular line in the first plot.

- (i) For some subsamples, we will have  $s^* < s$  (scenario A) and hence  $W^*_{n,i,b}(x_0)$  will be zero for those observations that satisfy  $s^* < X_i < s$ , while for those same observations,  $W_{n,i}(x_0) > 0$ . Thus, we expect  $W^*_{n,i}(x_0)$  to be smaller than  $W_{n,i}(x_0)$  close to and to the left of s. We indeed observe this in the second plot for observations in the region [0.5, s].
- (ii) For other subsamples, we will have  $s < s^*$  (scenario B) and hence  $W_{n,i,b}^*(x_0)$  will be non-zero for those observations that satisfy  $s < X_i < s^*$  and  $X_i \in D_n^*(b)$ , while for those same observations,  $W_{n,i}(x_0) = 0$ . Thus, we expect  $W_{n,i}^*(x_0)$  to be larger than  $W_{n,i}(x_0)$  (and in particular non-zero) close to and to the right of s. This is observed in the second plot in the region [s, 0.6].
- (iii)  $s^*$  will, in general, for subsamples of enough observations, not fall very far from s. Thus observations for which  $X_i$  far to the left of s, will also satisfy  $X_i < s^*$ , and hence  $W_{n,i}^*(x_0)$  will be close to  $W_{n,i}(x_0)$  (close and not exact, because of subsampling). This is observed in the second plot in the region [0, 0.5].
- (iv) For the same reason, observations for which  $X_i$  is far to the right of s, will also satisfy  $s^* < X_i$  and hence will have zero weight in the subagged estimate  $(W_{n,i}^*(x_0) = 0)$ . This is observed in the second plot in the region [0.6, 1].

The case where  $x_0$  gets closer to s. The closer  $x_0$  is to s, the more likely it is that for some subsamples, we will have  $s^* < x_0$ . For such subsamples, every observation to the left of  $s^*$  will be excluded from the estimate (therefore, zero weight). Consequently,  $W_{ni}^*(x_0)$  will tend to be smaller than  $W_{ni}(x_0)$ . At the same time, observations to the right of  $s^*$  will be included, which gives the non-zero weights  $W_{ni}^*(x_0)$  for all observations to the right of s. This is shown in the third and fourth plots of Figure 4 for  $x_0 = 0.5$  and  $x_0 = 0.6$ , corresponding to the second and third respectively perpendicular lines in the first plot. In other words, the closer  $x_0$  is to the split point, the more subagging adds weight

<sup>32.</sup> I.e., obtained by maximization of the *empirical* analogue of (3), or equivalently, of (7).

<sup>33.</sup> Note the small difference between 0.63, which is obtained based on the data, and the value of 0.64 previously obtained in Section 2.1 based on (7).



Figure 4: First plot (left): stump (blue) and subagged stump (red) estimates as a function of  $x_0$  for one realization of  $D_n$  (gray points). In black is the true regression function  $f(x_0)$ . Second plot: weights  $W_{n,i}(x_0)$  (stump, blue) and  $W_{n,i}^*(x_0)$  (subagged stump, red) for  $x_0 = 0.1$ . Third plot: weights for  $x_0 = 0.5$ . Fourth plot: weights for  $x_0 = 0.6$ .

to observations that had zero weight in the single tree. To our knowledge, we are the first to make this observation. This also helps understand the variance reduction of subagging observed around the split points, illustrated next in Section 5.2, as *estimates obtained from subagging are based on more distinct observations than in a single tree*.

#### 5.2 Statistical Bias and Variance

In Section 5.1 we looked at a single realization of  $D_n$ . Here we are interested in the effect of subagging on the statistical bias and variance of stumps, hence we look at several realizations of  $D_n$ , keeping the  $X_i$ 's fixed. Even though the  $X_i$ 's are fixed, different realizations of the  $Y_i$ 's will perturb the split point. We show that subagging has a **smoothing** effect which reduces the bias compared to a stump but this effect **disappears** when the number of splits increases. Additionally, subagging has a **stabilizing** effect which reduces the variance compared to a tree and this effect **persists** when the number of splits increases.

Figure 5 shows the squared bias and variance conditionally on X of a stump versus a subagged stump obtained over 200 replicates of  $D_n$  generated from Simulation II, as in Figure 4. Figure 6 shows trees with N = 3 splits. As a reference, the first two theoretically



Figure 5: Squared bias, variance and mean squared error of a stump (blue) and subagged stump (red) as a function of  $x_0$ . The vertical line shows the theoretical split (x = 0.64).



Figure 6: Three-split tree (blue) versus subagging of three-split trees (red).

optimal CART splits, i.e., maximizing (3), are at x = 0.64 and x = 0.83 respectively. The smoothing effect is visible around x = 0.64 for a single stump but no longer in the case N = 3. The variance reduction is visible in both the stump and in the case N = 3. In particular, Bühlmann and Yu, 2002's "instability region" is well illustrated: they argued that this region is a  $n^{-\frac{1}{3}}$ -neighborhood of the best split. Here n = 100, i.e.,  $n^{-\frac{1}{3}} \approx 0.22$ . While it is easy to see that subagging improves upon a tree in terms of global mean-squared error in both the case of a stump and the case N = 3, it is less clear in which case (between N = 1 and N = 3) subagging improves, overall in  $x_0$ , the most. This is further discussed in Section 6.

## 6 Subagging Large Trees

Now we show that subagging large trees is not always a good idea. As seen in Section 5.2, for stumps, subagging has a small effect on bias around the split point and such effect disappears with a few more splits: we expect that subagging brings no improvement on bias when the number of splits increases further, and in particular if it is large. Moreover, subagging reduces the variance around the split points, and this effect persists when trees are grown deeper: the more we split, the more instability we create, and hence the larger the region over which subagging helps by reducing the variance<sup>34</sup>. Figure 7 shows the global performance, i.e., on average over  $x_0$ , of trees versus subagging as a function of the number of splits in the same simulation framework as in Figures 5 and 6.

Large trees become very unstable, and even though subagging significantly reduces the variance of such trees, subagging performs worse than a single tree appropriately grown. In Figure 7, obtained again from Simulation II, the optimal (in terms of mean-squared error) number of splits for a single tree is N = 4, while for subagging it is N = 3. These are represented by the two dotted vertical lines in the right plot of Figure 7. The corresponding mean-squared errors are 0.87% (horizontal grey line) and 0.44% respectively, meaning an improvement of factor 2. In comparison, the mean-squared error of a tree with N = 49 splits is 3.74% and for its subagged analogue 1.48%, i.e., almost twice the mean-squared error of a single tree optimally grown.

<sup>34.</sup> This is in line with Bühlmann and Yu, 2002.



Figure 7: Global performance of trees (blue) and subagged trees (red) as a function of the number of splits.

#### 6.1 Optimal Number of Splits as a Function of the Dataset Size

We showed that for n = 100 observations, the optimal number of splits in terms of meansquared error is much lower than the maximum possible, both for trees and subagging (N = 4 and N = 3 respectively). This seems to still hold for larger dataset sizes. Figure 8 shows the optimal number of splits and corresponding global mean-squared error for trees and subagging as a function of n. We observe a roughly linear relation. This is of practical importance, as when the sample size is large, subagging trees of different sizes can become computationally costly. Instead, we can grow a single tree at different sizes, find the optimal number of splits with e.g. cross-validation, then use subagging with that same size for prediction, or, alternatively, perform cross-validation to choose the best number of splits around N = 3, e.g., to choose between 2,3 or 4 splits. Moreover, if interpretability of the model at hand is important, and if the difference in performance between a tree at its best and subagging at its best is not very large, then the single tree could be preferred.



Figure 8: Optimal number of splits (left) as a function of n and corresponding global meansquared error (right) for trees (blue) versus subagging (red).

# 7 Robustness

Here we test the robustness of our results and implementations.

#### 7.1 Subsample Size and Replacement

So far we have considered subagging instead of bagging, and fixed subsample size to k = 0.5n. In Figure 9, we show that our main empirical finding, namely that of Figure 7, still holds for different subsample sizes k as well as if we do bagging instead of subagging. In all cases, the conclusion is the same: the optimal number of splits is around 3. A decrease in k shows a decrease in variance but a bias is introduced, and vice-versa; increasing k brings the subagged estimator closer to the tree<sup>35</sup>. Moreover, bagging indeed behaves similarly to subagging with  $k = 0.5n^{36}$ .



Figure 9: Global performance of trees (blue), subagged trees with k = 50% n (red), k = 20% n (dashed), k = 95% n (dotted), and bagging (solid black).

#### 7.2 Readily Available Implementations

So far we have used our own implementations of recursive partitioning. Here we show that our main empirical finding, namely that of Figure 7, still holds if, instead of using our second implementation, we use the readily-available **R** package **randomForest**<sup>37</sup>, which incorporates both trees and subagging as special cases. We set the number of trees equal to 1 (**ntree = 1**) and sample without replacement n = 100 observations (**replace = FALSE** and **sampsize = 100**) in order to obtain a tree. Then we set the number of trees to B = 50(**ntree = 50**) and the subsample size to k = 0.5n (**sampsize = 50**) to obtain subagging. We can control the *maximum* number of splits by controlling the maximum number of nodes allowed, implemented with the option **maxnodes**: setting e.g. **maxnodes = 5** guarantees that we will grow trees with a maximum of 4 splits. We also set the minimum number of observations allowed in a terminal cell to one (**nodesize = 1**) and kept the default choices for all other options. Figure 10 shows the performance of a tree versus subagging

<sup>35.</sup> These observations were already made by Bühlmann and Yu, 2002, who pointed out that the subsample size can also be interpreted as a "smoothing" parameter - see their Section 3.2.

<sup>36.</sup> Also in line with previous literature, e.g., Bühlmann and Yu, 2002; Buja and Stuetzle, 2006.

<sup>37.</sup> Liaw and Wiener, 2002, version 4.7-1.1.

as a function of the maximum number of splits allowed. We obtain best performance at a maximum of 3 splits for both methods (vertical dotted line in the right plot), resulting in mean-squared errors of 0.87% (horizontal grey line) for the tree and 0.41% for subagging, in line with our results: subagging large trees performs worse than a single tree at optimal size.



Figure 10: Performance of trees versus subagging using the **randomForest** package.

Finally, in Figure 11 we show that the results of Figure 1 based on our first implementation can also be obtained using the **randomForest** package. Again, we set the number of trees equal to 1 (**ntree = 1**) and sample without replacement n = 2000 observations (**replace = FALSE and sampsize = 2000**) in order to obtain a single tree. Then, we used **maxnodes =**  $\frac{n}{h_n}$  as a *proxy* to control for the number of observations in cells in scenario b) (so **maxnodes =**  $\frac{n}{n^{0.65}}$ ); for small trees, we set **maxnodes = 2**; for large trees, we set **maxnodes =**  $\frac{n}{5}$ . In all cases we kept the minimum **nodesize = 1** and we used the default values for all other parameters. The bias-variance trade-off is again illustrated.



Figure 11: (In)consistency of trees using the randomForest package.

# 8 Conclusion

We studied CART and subagging in the context of regression estimation and have shown that tree size plays an important role in the performance of these methods. We established sufficient conditions for point-wise consistency of trees and provided an algorithm that satisfies them. Based on this, we formalized and illustrated the bias-variance trade-off associated with tree size. Then, we studied the effect of subagging on trees of varying sizes. We have illustrated the effect on weights and showed in simulations how subagging performs, compared to a single tree, when the number of splits is increased, both for the single tree and the trees constituting the subagged estimator. We found that a single tree optimally grown can outperform subagging if its subtrees are large. For practical applications, our findings suggest that large trees should not be the default choice in ensemble methods, and computational time can be saved upon by first finding the optimal size for a single tree in the problem at hand before building the ensemble.

# Appendix A. Proofs

# A.1 Proof of Proposition 1

Consider

$$\mathcal{C} = \mathbb{V}[Y|X \in C] - \frac{\mathbb{P}(X \in C_L)}{\mathbb{P}(X \in C)} \mathbb{V}[Y|X \in C_L] - \frac{\mathbb{P}(X \in C_R)}{\mathbb{P}(X \in R)} \mathbb{V}[Y|X \in C_R]$$
(15)

as in (3). We have  $\mathbb{V}(Y|X \in C) = \mathbb{E}[Y^2|X \in C] - \mathbb{E}[Y|X \in C]^2$  and similarly for  $C_L$  and  $C_R$ . Using that  $\mathbb{1}(X \in C) = \mathbb{1}(X \in C_L) + \mathbb{1}(X \in C_R)$  and the linearity of the expectation we get

$$\mathbb{P}(X \in C)\mathbb{V}(Y|X \in C) = \mathbb{E}[Y^2\mathbb{1}(X \in C_L)] + \mathbb{E}[Y^2\mathbb{1}(X \in C_R)] - \frac{1}{\mathbb{P}(X \in C)} \Big\{\mathbb{E}[Y\mathbb{1}(X \in C_L)] + \mathbb{E}[Y\mathbb{1}(X \in C_R)]\Big\}^2.$$
(16)

We also have

$$\mathbb{P}(X \in C_L)\mathbb{V}(Y|X \in C_L) = \mathbb{E}[Y^2\mathbb{1}(X \in C_L)] - \frac{\mathbb{E}[Y\mathbb{1}(X \in C_L)]^2}{\mathbb{P}(X \in C_L)}$$
(17)

and similarly for  $C_R$ . Plugging everything into (15), the terms containing  $Y^2$  cancel out and we are left with

$$\mathcal{C} = \mathbb{E}[Y\mathbbm{1}(X \in C_L)]^2 \left(\frac{1}{\mathbb{P}(X \in C_L)} - \frac{1}{\mathbb{P}(X \in C)}\right) - 2\mathbb{E}[Y\mathbbm{1}(X \in C_L)]\mathbb{E}[Y\mathbbm{1}(X \in C_R)]\frac{1}{\mathbb{P}(X \in C)} + \mathbb{E}[Y\mathbbm{1}(X \in C_R)]^2 \left(\frac{1}{\mathbb{P}(X \in C_R)} - \frac{1}{\mathbb{P}(X \in C)}\right) = \frac{\mathbb{P}(X \in C_L)\mathbb{P}(X \in C_R)}{\mathbb{P}(X \in C)} \left\{\mathbb{E}[Y|X \in C_L] - \mathbb{E}[Y|X \in C_R]\right\}^2$$
(18)

which concludes the proof of Proposition 1.

# A.2 Proof of Proposition 2

Let  $(X_{\iota_1}, Y_{\iota_1}), \ldots, (X_{\iota_{\nu}}, Y_{\iota_{\nu}})$  be  $\nu \geq 3$  re-ordered observations constituting a cell C. Suppose that the criterion is constant for any split in C. Using Proposition 1, this means that there exists a constant  $K_0 \geq 0$  such that for all  $\kappa \in \{1, \ldots, \nu - 1\}$ , we have

$$\frac{\kappa(\nu-\kappa)}{\nu} \left( \frac{1}{\kappa} \sum_{\lambda=1}^{\kappa} Y_{\iota_{\lambda}} - \frac{1}{\nu-\kappa} \sum_{\lambda=\kappa+1}^{\nu} Y_{\iota_{\lambda}} \right)^2 = nK_0.$$
(19)

In particular (19) is true for  $\kappa = 1$  and  $\kappa = \nu - 1$ , which, combined, give

$$(\nu - 1)Y_{\iota_1} - \sum_{\lambda=2}^{\nu} Y_{\iota_{\lambda}} = \pm \left(\sum_{\lambda=1}^{\nu-1} Y_{\iota_{\lambda}} - (\nu - 1)Y_{\nu}\right).$$
(20)

Suppose that in (20) the "+" holds. This is equivalent (recall  $\nu \geq 3$ ) to

$$\frac{Y_{\iota_1} + Y_{\iota_\nu}}{2} = \frac{1}{\nu - 2} \sum_{\lambda=2}^{\nu - 1} Y_{\iota_\lambda}.$$
(21)

Plug this into (19) with  $\kappa = 1$ . After simplifying, we obtain

$$Y_{\iota_1} = Y_{\iota_\nu} \pm 2\sqrt{\frac{\nu - 1}{\nu}nK_0}$$
(22)

which is almost surely impossible. Next suppose that in (20) the "-" holds. This is equivalent to  $Y_{\iota_1} = Y_{\iota_{\nu}}$  which is also almost surely impossible as long as Y is a continuous random variable. Therefore, by contradiction, almost surely the criterion is not constant in any cell containing at least three observations. This concludes the proof of Proposition 2.

#### A.3 Proof of Theorem 1

We start by re-formulating the honesty assumption. We assume given two datasets:

- 1.  $D_n = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$ , a dataset used to partition  $[0, 1]^p$ , and
- 2.  $D'_n = \{(X'_1, Y'_1), \dots, (X'_n, Y'_n)\}$ , a dataset used for estimation, independent of  $D_n$ , of same size as  $D_n$ .

Fix  $x_0 \in [0,1]^p$ . We note  $C_n = C_n(x_0; D_n)$  the terminal cell containing  $x_0$  in the partition obtained using the first dataset. Throughout the proof, we also get rid of the dependence of weights on  $x_0$ . The tree estimator is noted

$$T_n(x_0) = \sum_{i=1}^n W_{n,i} Y'_i \quad \text{with} \quad W_{n,i} = \frac{\mathbb{1}_{X'_i \in C_n}}{\sum_{j=1}^n \mathbb{1}_{X'_j \in C_n}}$$
(23)

which can be re-written as

$$T_n(x_0) = \sum_{i=1}^n W_{n,i} f(X'_i) + \sum_{i=1}^n W_{n,i} \varepsilon'_i.$$
 (24)

We treat the two terms separately.

# A.3.1 Bias of the error term

Lemma 1 For all n,

$$\mathbb{E}\left[\sum_{i=1}^{n} W_{n,i}\varepsilon'_{i} \middle| \mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n})\right] = 0.$$
(25)

**Proof** From Assumptions 1 and 2,

$$\forall i, \ \mathbb{E}[\varepsilon_i'|\mathbb{1}(X_1' \in C_n), \dots, \mathbb{1}(X_n' \in C_n)] = \mathbb{E}[\varepsilon_i'] = 0.$$

# A.3.2 Variance of the error term

Lemma 2 For all n,

$$\mathbb{V}\left[\sum_{i=1}^{n} W_{n,i}\varepsilon'_{i} \middle| \mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n})\right] = \frac{\sigma^{2}}{\sum_{i=1}^{n} \mathbb{1}(X'_{i} \in C_{n})}.$$
(26)

 $\mathbf{Proof}$ 

$$\mathbb{V}\left[\sum_{i=1}^{n} W_{n,i}\varepsilon'_{i} \middle| \mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n})\right] \\
= \mathbb{E}\left[\left(\sum_{i=1}^{n} W_{n,i}\varepsilon'_{i}\right)^{2} \middle| \mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n})\right] - 0^{2} \text{ from (25)} \\
= \sum_{i=1}^{n} W_{n,i}^{2}\mathbb{E}[\varepsilon'_{i}^{2} \middle| \mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n})] \\
+ 2\sum_{j\neq i} W_{n,i}W_{n,j}\mathbb{E}[\varepsilon'_{i}\varepsilon'_{j} \middle| \mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n})] \\
= \sigma^{2}\sum_{i=1}^{n} W_{n,i}^{2} + 0 \\
= \frac{\sigma^{2}}{\sum_{i=1}^{n} \mathbb{1}(X'_{i} \in C_{n})}.$$

Lemma 3

$$\frac{1}{\sum_{i=1}^{n} \mathbb{1}(X'_i \in C_n)} \xrightarrow[n \to \infty]{\mathbb{P}} 0.$$
(27)

**Proof** Conditionally on  $C_n$ ,  $\sum_{i=1}^n \mathbb{1}(X'_i \in C_n)$  follows a Binomial distribution of parameters  $(n, \lambda(C_n))$ . Therefore it is sufficient to show

$$n \ \lambda(C_n) \xrightarrow[n \to \infty]{\mathbb{P}} \infty.$$
 (28)

We know from Assumption 3

$$n \ \lambda_n(C_n) \ge h_n \xrightarrow[n \to \infty]{} \infty \tag{29}$$

where

$$\lambda_n(C_n) := \frac{1}{n} \sum_{i=1}^n \mathbb{1}(X_i \in C_n).$$
(30)

From empirical process theory, the set of hyper-rectangles is a Glivenko-Cantelli class (of VC dimension  $2p < \infty$ ), note it  $\mathcal{R}$ , and we have

$$\sqrt{n} \sup_{C \in \mathcal{R}} |\lambda(C) - \lambda_n(C)| = \mathcal{O}_{\mathbb{P}}(1)$$
(31)

and for any realization of  $C_n$ ,

$$\sqrt{n}|\lambda(C_n) - \lambda_n(C_n)| \le \sqrt{n} \sup_{C \in \mathcal{R}} |\lambda(C) - \lambda_n(C)|$$
(32)

therefore, using the decomposition

$$\sqrt{n\lambda(C_n)} = \sqrt{n\lambda_n(C_n)} + \sqrt{n}(\lambda(C_n) - \lambda_n(C_n))$$
(33)

we get that if  $\sqrt{n\lambda_n(C_n)}$  tends to infinity, then so does  $\sqrt{n\lambda(C_n)}$  for any realization of  $C_n$ . By Assumption 3,

$$n\lambda_n(C_n) \ge h_n = n^\alpha \tag{34}$$

therefore

$$\sqrt{n\lambda_n(C_n)} \ge n^{\alpha + \frac{1}{2} - 1} \tag{35}$$

with  $\alpha + \frac{1}{2} - 1 > 0$  because  $\alpha > \frac{1}{2}$ . Therefore  $\sqrt{n\lambda_n(C_n)} \xrightarrow[n \to \infty]{\mathbb{P}} \infty$ , hence  $\sqrt{n\lambda(C_n)} \xrightarrow[n \to \infty]{\mathbb{P}} \infty$ and hence also  $n\lambda(C_n) \xrightarrow[n \to \infty]{\mathbb{P}} \infty$  for any realization of  $C_n$ .

### A.3.3 BIAS OF THE REGRESSION TERM

Lemma 4 For all n,

$$\mathbb{E}\left[\sum_{i=1}^{n} W_{n,i}f(X'_{i})|\mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n}), C_{n}\right] = \mathbb{E}[f(X)|X \in C_{n}].$$
(36)

**Proof** Conditionally on  $\{\mathbb{1}(X'_i \in C_n), C_n\}$ ,  $f(X'_i)$  is independent of  $\mathbb{1}(X'_j \in C_n)$  for all  $j \neq i$ , i.e.,

$$\mathbb{E}[f(X'_i)|\mathbb{1}(X'_1 \in C_n), \dots, \mathbb{1}(X'_n \in C_n), C_n] = \mathbb{E}[f(X'_i)|\mathbb{1}(X'_i \in C_n), C_n].$$
(37)

Moreover

$$\mathbb{E}[f(X_i')|\mathbb{1}(X_i' \in C_n), C_n]$$
  
=  $\mathbb{1}(X_i' \in C_n)\mathbb{E}[f(X)|X \in C_n] + \mathbb{1}(X_i' \notin C_n)\mathbb{E}[f(X)|X \notin C_n].$ 

# A.3.4 Variance of the regression term

Lemma 5 For all n,

$$\mathbb{V}\left[\sum_{i=1}^{n} W_{n,i}f(X'_{i})|\mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n}), C_{n}\right] = \mathbb{V}[f(X)|X \in C_{n}].$$
(38)

**Proof** From Lemma 4, we have

$$\mathbb{V}\left[\sum_{i=1}^{n} W_{n,i}f(X_{i})|\mathbb{1}(X_{1} \in C_{n}), \dots, \mathbb{1}(X_{n} \in C_{n}), C_{n}\right] = \mathbb{E}\left[\left(\sum_{i=1}^{n} W_{n,i}f(X_{i}')\right)^{2}|\mathbb{1}(X_{1}' \in C_{n}), \dots, \mathbb{1}(X_{n}' \in C_{n}), C_{n}\right] - \mathbb{E}[f(X)|X \in C_{n}]^{2} \tag{39}$$

where

$$\mathbb{E}\left[\left(\sum_{i=1}^{n} W_{n,i}f(X'_{i})\right)^{2} |\mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n}), C_{n}\right] \\
= \sum_{i=1}^{n} W_{n,i}^{2} \mathbb{E}[f(X'_{i})^{2} |\mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n}), C_{n}] \\
+ 2\sum_{i \neq j} W_{n,i} W_{n,j} \mathbb{E}[f(X'_{i})f(X'_{j}) |\mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n}), C_{n}].$$
(40)

For all i, we have

$$\mathbb{E}[f(X_i')^2 | \mathbb{1}(X_1' \in C_n), \dots, \mathbb{1}(X_n' \in C_n), C_n] = \mathbb{E}[f(X_i')^2 | \mathbb{1}(X_i' \in C_n), C_n]$$
(41)

and for all  $i \neq j$ , we have

$$\mathbb{E}[f(X'_{i})f(X'_{j})|\mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n}), C_{n}] = \mathbb{E}[f(X'_{i})f(X'_{j})|\mathbb{1}(X'_{i} \in C_{n}), \mathbb{1}(X'_{j} \in C_{n}), C_{n}].$$
(42)

Moreover,

$$\mathbb{E}[f(X_i')^2|\mathbb{1}(X_i' \in C_n)] = \mathbb{1}(X_i' \in C_n)\mathbb{E}[f(X)^2|X \in C_n)] + \mathbb{1}(X_i' \notin C_n)\mathbb{E}[f(X)^2|X \notin C_n)]$$

$$(43)$$

and

$$\mathbb{E}[f(X'_{i})f(X'_{j})|\mathbb{1}(X'_{i} \in C_{n}), \mathbb{1}(X'_{j} \in C_{n})] = \\
\mathbb{1}(X'_{i} \in C_{n})\mathbb{1}(X'_{j} \in C_{n})(\mathbb{E}[f(X)|X \in C_{n})])^{2} \\
+\mathbb{1}(X'_{i} \in C_{n})\mathbb{1}(X'_{j} \notin C_{n})\mathbb{E}[f(X)|X \in C_{n})]\mathbb{E}[f(X)|X \notin C_{n})] \\
+\mathbb{1}(X'_{i} \notin C_{n})\mathbb{1}(X'_{j} \in C_{n})\mathbb{E}[f(X)|X \notin C_{n})]\mathbb{E}[f(X)|X \in C_{n})] \\
+\mathbb{1}(X'_{i} \notin C_{n})\mathbb{1}(X'_{j} \notin C_{n})(\mathbb{E}[f(X)|X \notin C_{n})])^{2}.$$
(44)

Note that

- $W_{n,i}\mathbb{1}(X'_i \in C_n) = 1$  for all i,
- $W_{n,i}\mathbb{1}(X'_j \in C_n) = 0$  whenever  $i \neq j$ ,
- $\sum_{i=1}^{n} W_{n,i}^2 = \frac{1}{\sum_{i=1}^{n} \mathbb{1}(X'_i \in C_n)}$ , and

• 
$$2\sum_{i\neq j} W_{n,i}W_{n,j} = \left(\sum_{i=1}^{n} W_{n,i}\right)^2 - \sum_{i=1}^{n} W_{n,i}^2 = 1 - \frac{1}{\sum_{i=1}^{n} \mathbb{I}(X'_i \in C_n)}.$$

Therefore

$$\sum_{i=1}^{n} W_{n,i}^{2} \mathbb{E}[f(X_{i}')^{2} | \mathbb{1}(X_{1}' \in C_{n}), \dots, \mathbb{1}(X_{n}' \in C_{n}), C_{n}]$$

$$= \mathbb{E}[f(X)^{2} | X \in C_{n}] \frac{1}{\sum_{i=1}^{n} \mathbb{1}(X_{i}' \in C_{n})}$$
(45)

and

$$2\sum_{i\neq j} W_{n,i}W_{n,j}\mathbb{E}[f(X'_i)f(X'_j)|\mathbb{1}(X'_1 \in C_n), \dots, \mathbb{1}(X'_n \in C_n), C_n]$$
  
=  $\mathbb{E}[f(X)^2|X \in C_n] \left(1 - \frac{1}{\sum_{i=1}^n \mathbb{1}(X'_i \in C_n)}\right).$  (46)

Putting everything together gives

$$\mathbb{V}\left[\sum_{i=1}^{n} W_{n,i}f(X'_{i})|\mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n}), C_{n}\right] \\
= \mathbb{E}[f(X)^{2}|X \in C_{n}]\left(\frac{1}{\sum_{i=1}^{n}\mathbb{1}(X'_{i} \in C_{n})} + 1 - \frac{1}{\sum_{i=1}^{n}\mathbb{1}(X'_{i} \in C_{n})}\right) \quad (47) \\
- \mathbb{E}[f(X)|X \in C_{n}]^{2} \\
= \mathbb{V}[f(X)|X \in C_{n}].$$

# A.3.5 Putting everything together

From Lemmas 2 and 3 we get

$$\mathbb{V}\left[\sum_{i=1}^{n} W_{n,i}\varepsilon_{i}^{\prime}|\mathbb{1}(X_{1}^{\prime}\in C_{n}),\ldots,\mathbb{1}(X_{n}^{\prime}\in C_{n})\right] \xrightarrow[n\to\infty]{\mathbb{P}} 0$$

$$(48)$$

and with Lemma 1 we obtain

$$\sum_{i=1}^{n} W_{n,i} \varepsilon_i' \xrightarrow[n \to \infty]{\mathbb{P}} 0.$$
(49)

From Lemma 5 and Assumption 4 we get

$$\mathbb{V}\left[\sum_{i=1}^{n} W_{n,i}f(X'_{i})|\mathbb{1}(X'_{1} \in C_{n}), \dots, \mathbb{1}(X'_{n} \in C_{n}), C_{n}\right] \xrightarrow[n \to \infty]{\mathbb{P}} 0$$
(50)

and with Lemma 4 we obtain

$$\sum_{i=1}^{n} W_{n,i}f(X'_i) \xrightarrow[n \to \infty]{\mathbb{P}} f(x_0).$$
(51)

Therefore using Slutsky's theorem

$$\hat{T}_n(x_0) \xrightarrow[n \to \infty]{\mathbb{P}} f(x_0)$$
 (52)

which concludes the proof of Theorem 1.

### References

- Susan Athey and Guido Imbens. Recursive partitioning for heterogeneous causal effects. Proceedings of the National Academy of Sciences, 113(27):7353–7360, 2016.
- Leo Breiman. Bagging predictors. Machine Learning, 24:123–140, 1996.
- Leo Breiman. Random forests. Machine Learning, 45:5–32, 2001.
- Leo Breiman, Jerome Friedman, Richard Olshen, and Charles Stone. Classification and regression trees. *CRC Press*, 1984.
- Peter Bühlmann and Bin Yu. Analyzing bagging. Annals of Statistics, 30(4):927–961, 2002.
- Andreas Buja and Werner Stuetzle. Observations on bagging. *Statistica Sinica*, 16(2): 323–351, 2006.
- Alicia Curth, Alan Jeffares, and Mihaela van der Schaar. Why do random forests work? understanding tree ensembles as self-regularizing adaptive smoothers. doi:10.48550/arXiv.2402.01502, 2024.
- Thomas G Dietterich and Eun Bae Kong. Machine learning bias, statistical bias, and statistical variance of decision tree algorithms. *Working Paper*, 1995.

- Roxane Duroux and Erwan Scornet. Impact of subsampling and tree depth on random forests. *ESAIM: Probability and Statistics*, 22:96–128, 2018.
- Jerome H Friedman and Peter Hall. On bagging and nonlinear estimation. *Journal of* statistical planning and inference, 137(3):669–683, 2007.
- Yves Grandvalet. Bagging equalizes influence. Machine Learning, 55:251–270, 2004.
- Yves Grandvalet. Stability of bagged decision trees. In Proceedings of the XLIII Scientific Meeting of the Italian Statistical Society, pages 221–230. CLEUP, 2006.
- László Györfi, Michael Kohler, Adam Krzyzak, Harro Walk, et al. A distribution-free theory of nonparametric regression, volume 1. Springer, 2002.
- Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer, 2009.
- Wassily Hoeffding. A class of statistics with asymptotically normal distribution. Breakthroughs in Statistics: Foundations and Basic Theory, pages 308–334, 1992.
- Gareth James, Daniela Witten, Trevor Hastie, Robert Tibshirani, et al. An introduction to statistical learning, volume 112. Springer, 2013.
- Jason M Klusowski and Peter M Tian. Large scale prediction with decision trees. *Journal* of the American Statistical Association, 119(545):1–27, 2023.
- Andy Liaw and Matthew Wiener. Classification and regression by randomforest. *R News*, 2(3):18–22, 2002.
- Yi Lin and Yongho Jeon. Random forests and adaptive nearest neighbors. Journal of the American Statistical Association, 101(474):578–590, 2006.
- Lucas Mentch and Giles Hooker. Quantifying uncertainty in random forests via confidence intervals and hypothesis tests. *Journal of Machine Learning Research*, 17(26):1–41, 2016.
- Elizbar A Nadaraya. On estimating regression. Theory of Probability & Its Applications, 9 (1):141–142, 1964.
- Erwan Scornet, Gérard Biau, and Jean-Philippe Vert. Consistency of random forests. Annals of Statistics, 43(4):1716 1741, 2015.
- Mark R Segal. Machine learning benchmarks and random forest regression. *Working Paper*, 2004.
- Jake A Soloff, Rina Foygel Barber, and Rebecca Willett. Bagging provides assumption-free stability. doi:10.48550/arXiv.2301.12600, 2023.
- Charles J Stone. Consistent nonparametric regression. Annals of Statistics, 5(4):595–620, 1977.

Aad W Van der Vaart. Asymptotic statistics, volume 3. Cambridge University Press, 2000.

- Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. Journal of the American Statistical Association, 113(523):1228– 1242, 2018.
- Geoffrey S Watson. Smooth regression analysis. Sankhyā: Indian Journal of Statistics, Series A, 26(4):359–372, 1964.
- Siyu Zhou and Lucas Mentch. Trees, forests, chickens, and eggs: when and why to prune trees in a random forest. Statistical Analysis and Data Mining: The ASA Data Science Journal, 16(1):45–64, 2023.