Journal of the Royal Statistical Society Series B: Statistical Methodology, 2025, **00**, 1–26
https://doi.org/10.1093/jrsssb/qkaf077

ROYAL STATISTICAL SOCIETY
DATA | EVIDENCE | DECISIONS
Journal of the Statistics Society
Series B
Statistical Methodology
B

Original Article

Inference with Mondrian random forests

Matias Damian Cattaneo[1], **Jason Matthew Klusowski**[1] **and William George Underwood**[2]

[1]Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA
[2]Statistical Laboratory, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Rd, Cambridge CB3 0WB, UK

Address for correspondence: Matias D. Cattaneo, Department of Operations Research and Financial Engineering, Princeton University, Sherrerd Hall, Charlton Street, Princeton, NJ 08544, USA. Email: cattaneo@princeton.edu

Abstract

Random forests are popular methods for regression and classification analysis, and many different variants have been proposed in recent years. One interesting example is the Mondrian random forest, in which the underlying constituent trees are constructed via a Mondrian process. We give precise bias and variance characterizations, along with a Berry–Esseen-type central limit theorem, for the Mondrian random forest regression estimator. By combining these results with a carefully crafted debiasing approach and an accurate variance estimator, we present valid statistical inference methods for the unknown regression function. These methods come with explicit error bounds in terms of the sample size, tree complexity parameter, and number of trees in the forest, and include coverage error rates for feasible confidence interval estimators. Our debiasing procedure for the Mondrian random forest also allows it to achieve the minimax-optimal point estimation convergence rate in mean squared error for multivariate β-Hölder regression functions, for all $\beta > 0$, provided that the underlying tuning parameters are chosen appropriately. Efficient and implementable algorithms are devised for both batch and online learning settings, and we study the computational complexity of different Mondrian random forest implementations. Finally, simulations with synthetic data validate our theory and methodology, demonstrating their excellent finite-sample properties.

Keywords: Berry–Esseen theorem, bias correction, minimax estimation, random forests, regression trees, statistical inference

1. Introduction

Random forests, first introduced by Breiman (2001), are a workhorse in modern machine learning for regression and classification tasks. Their desirable traits include computational efficiency in big data settings (via parallelization and greedy heuristics), simplicity of configuration and amenability to tuning parameter selection, ability to adapt to latent structure in high-dimensional data sets, and flexibility in handling mixed data types, among other virtues. Random forests have also achieved great empirical successes in many fields of study, including healthcare, finance, online commerce, causal inference, text analysis, bioinformatics, image classification, and ecology.

Since Breiman introduced random forests over twenty years ago, the study of their statistical properties remains an active area of research. Many fundamental questions about Breiman's random forests remain unanswered, owing in part to the subtle ingredients present in the estimation procedure which make standard analytical tools ineffective. These technical difficulties stem from the way the constituent trees greedily partition the covariate space, utilizing both the covariate and response data. This creates complicated dependencies on the data that are often exceedingly hard to untangle without overly stringent assumptions, thereby hampering theoretical progress.

Received: April 8, 2025. **Revised:** November 7, 2025. **Accepted:** November 17, 2025

To address the aforementioned technical challenges while retaining the phenomenology of Breiman's random forests, a variety of stylized versions of random forest procedures have been proposed and studied in the literature. Early proposals include centered random forests (Arnould et al., 2023; Biau, 2012) and median random forests (Arnould et al., 2023; Duroux & Scornet, 2018). Each tree in a centered random forest is constructed by first choosing a covariate uniformly at random and then splitting the cell at the midpoint along the direction of the chosen covariate. Median random forests operate in a similar way, but involve the covariate data by splitting at the empirical median along the direction of the randomly chosen covariate. Known as purely random forests, these procedures simplify Breiman's original, more data-adaptive version by growing trees that partition the covariate space in a way that is statistically independent of the response data.

Yet another variant of random forests, Mondrian random forests (Lakshminarayanan et al., 2014), have received significant attention from the statistics and machine learning communities in recent years (Baptista et al., 2024; Gao et al., 2022; Ma et al., 2020; Mourtada et al., 2020, 2021; O'Reilly, 2024; O'Reilly & Tran, 2022, 2024; Osborne & O'Reilly, 2025; Scillitoe et al., 2021; Vicuna et al., 2021; Zhan et al., 2024). Like other purely random forest variants, Mondrian random forests offer a simplified modification of Breiman's original proposal in which the partition is generated independently of the data and according to a canonical stochastic process known as the Mondrian process (Roy & Teh, 2008). The Mondrian process takes a single tuning parameter $\lambda > 0$ known as the 'lifetime' and enjoys various mathematical properties. These properties allow Mondrian random forests to be fitted in an online manner (Lakshminarayanan et al., 2014; Mourtada et al., 2021) and permit a rigorous statistical analysis, while also retaining some of the appealing features of other random forest methods. The lifetime parameter λ, in analogy with the number of refinements of a data-adaptive recursive partitioning algorithm, governs the extent to which the response data is smoothed, with a large λ resulting in a more complicated partition and therefore less smoothing. However, unlike a data-adaptive partition which modulates the smoothing intensity based on local data characteristics, the Mondrian process applies a uniform smoothing effect globally across all covariates. Thus, one can draw parallels between (axis-aligned) data-dependent partitioning schemes and more flexible—albeit, more involved—versions of the Mondrian process which employ adaptive directional smoothing, with unique lifetimes λ_j learned for each covariate, permitting tailored smoothing of the responses.

This paper studies the theoretical statistical properties of Mondrian random forests with an emphasis on *inference* techniques specific to this procedure. Although Mondrian random forests are not state-of-the-art for high-dimensional problems in practice, we focus on this purely random forest variant not only because of its importance in the development of random forest theory in general, but also because the Mondrian process (along with generalizations including the oblique Mondrian process and stationary random tessellation processes) is, to date, the only known randomized recursive tree mechanism for which the resulting random forest is minimax-optimal for point estimation over a class of smooth multivariate regression functions, without requiring sample splitting (Mourtada et al., 2020; O'Reilly, 2024; O'Reilly & Tran, 2024). In fact, when the covariate dimension exceeds one, the aforementioned centered and median random forests are both minimax *suboptimal*, due to their large biases, over the class of Lipschitz smooth regression functions (J. Klusowski, 2021). It is therefore natural to focus our study of inference for random forests on versions that at the very least exhibit competitive bias and variance, as this will have important implications for the trade-off between confidence and precision. Moreover, recent studies of generalized random forests and distributional random forests identify similar fundamental challenges, namely those of establishing Gaussian approximations and combining them with strategies for bias reduction (Näf et al., 2023).

Despite their recent popularity, relatively little is known about the formal statistical properties of Mondrian random forests. Focusing on nonparametric regression, Mourtada et al. (2020) recently showed that Mondrian forests containing just a single tree (called a Mondrian tree) can be minimax-optimal in integrated mean squared error whenever the regression function is β-Hölder continuous for some $\beta \in (0, 1)$. The authors also showed that, when appropriately tuned, large Mondrian random forests can be similarly minimax-optimal for $\beta \in (0, 2)$, while the constituent trees cannot. See also O'Reilly and Tran (2022) for analogous results on more general Mondrian tree and forest constructions. These results formally demonstrate the value of

ensembling with random forests from a point estimation perspective. No results are currently available in the literature for statistical inference using Mondrian random forests.

As already mentioned, a different strand of the literature studies the statistical properties of Breiman's random forests which form ensembles of *adaptive* decision trees. In such models, each constituent tree is constructed with a greedy algorithm that recursively optimizes a goodness-of-fit metric (such as mean squared error) using both the covariates and response data; a leading example in practice is the celebrated Classification and Regression Tree (CART) methodology (Breiman, 2001; Breiman et al., 1984). The underlying complexity of the resulting procedures make their formal theoretical analysis quite difficult, and therefore only a more restricted set of results is currently available in the literature. In terms of estimation theory, Scornet et al. (2015) established consistency of adaptive random forests for additive models with a fixed number of covariates, and Chi et al. (2022), J. M. Klusowski and Tian (2024), and Cattaneo, Chandak, et al. (2024) provided rates of convergence for models with a growing number of covariates, under different assumptions on the statistical and algorithmic features of the constituent decision trees. A framework for tuning tree depths via data-adaptive early stopping was developed by Miftachov and Reiß (2025). In contrast, formal inference theory is far less developed, since there are arguably no satisfactory theoretical results for fully adaptive decision tree or random forest methods. For example, Wager and Athey (2018) provide asymptotic estimation and inference results for adaptive random forests, but they employ sample splitting (i.e. the so-called 'honesty' property where the partitioning and, separately, the output in the terminal cells are formed using independent subsamples), and make assumptions that rule out procedures commonly used in practice such as CART and other sum-of-squares based splitting criteria; cf., the so-called 'α-regularity' condition (Cattaneo, Klusowski, et al., 2024; Cattaneo, Klusowski, Yu, 2025). From a broad perspective, our paper connects with this distinct thread in the random forest literature by demonstrating optimal estimation and inference results for nonadaptive Mondrian random forests with explicit probability deviation guarantees. Furthermore, by providing a clean theoretical framework, our results may inform and guide the development of more adaptive Mondrian random forests methods, a direction we return to briefly in the conclusion of this paper.

1.1 Contributions

Our paper contributes to the literature on the foundational statistical properties of Mondrian random forest regression estimation with two main results. Firstly, we give a central limit theorem for the classical Mondrian random forest point estimator under weak conditions, and propose valid large-sample inference procedures employing a consistent standard error estimator. We establish these results by deploying a restricted moments version of the Berry–Esseen theorem for independent but not identically distributed (i.n.i.d.) random variables (Petrov, 1995, Theorem 5.7) because we need to handle delicate probabilistic features of the Mondrian random forest estimator. In particular, we deal with the existence of Mondrian cells which are 'too small' and lead to a reduced effective (local) sample size for some trees in the forest. Such pathological cells are in fact typical in Mondrian random forests and complicate the probability limits of certain sample averages; in fact, small Mondrian random forests (or indeed single Mondrian trees) remain random even in the limit due to the lack of ensembling. The presence of such small cells renders inapplicable prior distributional approximation results for partitioning-based estimators in the literature (Cattaneo et al., 2020; Huang, 2003), since the commonly required quasi-uniformity assumption on the underlying partitioning scheme (cf., α-regularity in the adaptive random forest literature) is violated by partitions generated using the Mondrian process. We circumvent this technical challenge by establishing new theoretical results for Mondrian partitions and their associated Mondrian trees and forests, which may be of independent interest. Our distributional approximation does not rely on sample splitting; unlike approaches based on 'honest' trees, the estimator is fit using the entire sample of covariates and responses simultaneously.

The second main contribution of our paper is to propose a debiasing approach for the Mondrian random forest point estimator. We accomplish this by first precisely characterizing the probability limit of the large sample conditional bias, and then applying a debiasing procedure based on the generalized jackknife (Schucany & Sommers, 1977). We thus exhibit a Mondrian random forest variant which is minimax-optimal in pointwise mean squared error when the regression function is β-Hölder for any $\beta > 0$. Our method works by generating an ensemble of Mondrian random

forests carefully chosen to have smaller misspecification bias when extra smoothness is available, resulting in minimax optimality even for $\beta > 2$. This result complements (Mourtada et al., 2020) by demonstrating the existence of a class of Mondrian random forests that can efficiently exploit the additional smoothness of the unknown regression function for minimax-optimal point estimation. Our proposed debiasing procedure is also useful when conducting statistical inference because it provides a principled method for ensuring that the bias is negligible relative to the standard deviation of the estimator. More specifically, we use our debiasing approach to construct valid confidence intervals based on robust bias correction (Calonico et al., 2018, 2022), and include an explicit bound on their coverage error probability.

For the purposes of implementation, we propose techniques for tuning parameter selection and demonstrate the practical applicability and accuracy of our methodology through empirical studies with simulated data. We also discuss applications to batch and online learning settings, presenting computationally efficient algorithms along with bounds for their average case time complexity.

1.2 Organization

Section 2 gives the assumptions on the data generating process, using a Hölder smoothness condition on the regression function to control the bias of various estimators. We also introduce the Mondrian process and use it to define the Mondrian random forest estimator, stating the assumptions on its lifetime parameter and the number of trees.

Section 3 presents our first set of main results. We begin by precisely characterizing the bias of the Mondrian random forest estimator in Lemma 1, with the aim of subsequently applying a debiasing procedure. We similarly analyze the variance of this estimator (Lemma 2), and deduce its rate of convergence in Theorem 1. Next, we present our Berry–Esseen-type central limit theorem for the centered Mondrian random forest estimator under weak conditions as Theorem 2, and discuss implications for lifetime parameter selection. To enable valid feasible statistical inference, we provide a consistent variance estimator in Lemma 3, and use it to construct confidence intervals in Theorem 3.

Section 4 introduces our proposed debiased Mondrian random forests, a family of estimators based on linear combinations of Mondrian random forests with varying lifetime parameters. These parameters are carefully chosen to annihilate leading terms in our bias characterization, yielding an estimator with superior bias properties (Lemma 4). We also study the variance of this debiased estimator (Lemma 5), and derive its rate of convergence in Theorem 4. The resulting rate is shown to be minimax-optimal in mean squared error for each Hölder parameter $\beta > 0$, under regularity conditions. Furthermore, Theorem 5 verifies that a Berry–Esseen theorem holds for the debiased Mondrian random forest. We again discuss the implications for the lifetime parameter, and provide a consistent variance estimator (Lemma 6) for constructing confidence intervals (Theorem 6).

Section 5 discusses implementation details and empirical results, beginning by presenting a data-driven approach to selecting the crucial lifetime parameter using polynomial estimation. We also give advice on choosing the number of trees, as well as other parameters associated with the debiasing procedure. Empirical simulation results are presented using synthetic data, demonstrating the practical value of our methods for optimal point estimation and feasible robust bias-corrected inference.

Section 6 considers the computational aspects of our methodology, presenting algorithmic procedures with precisely characterized average case time complexity bounds for the batch setting (Algorithm 1, Lemma 7) and for online learning regimes (Algorithm 2, Lemma 8).

Concluding remarks are given in Section 7, while the supplementary material (Cattaneo, Klusowski, Underwood, 2025) contains all the mathematical proofs of our theoretical results, alongside additional empirical studies.

1.3 Notation

We write $\|\cdot\|_2$ for the usual Euclidean ℓ^2 norm on $\mathbb{R}^d$. The natural numbers are $\mathbb{N} = \{0, 1, 2, \ldots\}$. We use $a \wedge b$ for the minimum and $a \vee b$ for the maximum of two real numbers. For nonnegative sequences a_n and b_n, we write $a_n \lesssim b_n$ to indicate that a_n/b_n is bounded for $n \geq 1$. If $a_n \lesssim b_n \lesssim a_n$, we write $a_n \asymp b_n$. For random nonnegative sequences A_n and B_n, similarly we write $A_n \lesssim_{\mathbb{P}} B_n$ if A_n/B_n is bounded in probability. Let $\Phi : \mathbb{R} \to \mathbb{R}$ be the cumulative distribution function of the standard normal distribution, and for $\alpha \in (0, 1)$, let q_α be the normal quantile satisfying $\Phi(q_\alpha) = \alpha$.

Algorithm 1 Batch learning with Mondrian random forests

Input: Data (X_i, Y_i) for $1 \le i \le n$, forest size $B \ge 1$, debiasing order $J \ge 0$.

1 Select λ using one of the methods from Section 5.2.

2 Construct the union cell $U(x)$ and active indices $I(x)$ as in (12).

3 Calculate $N_{br}(x)$, $S_{br}(x)$ and $V_{br}(x)$ for each $1 \le b \le B$ and $0 \le r \le J$ as in (13).

4 Compute $\hat{\mu}_{\mathrm{d}}(x)$ and $\hat{\sigma}^2(x)$ with (14).

5 Calculate $\hat{\Sigma}_{\mathrm{d}}(x)$ and $\mathrm{CI}_{\mathrm{d}}(x)$ using (15) and (11) respectively.

Algorithm 2 Online learning with Mondrian random forests

Input: Data (X_i, Y_i) for $1 \le i \le n$, forest size $B \ge 1$, debiasing order $J \ge 0$, lifetime λ, forest exponent $\xi \in (0, 1)$, lifetime exponent $\zeta \in (0, 1/d)$, active region $U(x)$, active indices $I(x)$, trees $T_{br}(x)$ and $N_{br}(x)$, $S_{br}(x)$, $V_{br}(x)$ for $1 \le b \le B$ and $0 \le r \le J$, new data (X_i, Y_i) for $n+1 \le i \le n+k$, recalculation gap $K \ge 1$.

1 Get the updated number of trees $B^* = \lfloor (n+k)^{\xi} B / n^{\xi} \rfloor$.

2 With probability $1 \wedge (k/K)$, select λ^* as in Section 5.2; otherwise, set $\lambda^* = (n+k)^{\zeta}\lambda / n^{\zeta}$.

3 Generate the incrementally updated forest $T^*_{br}(x)$ as in (16) and (17).

4 Construct the updated union cell $U^*(x)$ and active indices $I^*(x)$ as in (18).

5 Calculate $N^*_{br}(x)$, $S^*_{br}(x)$, and $V^*_{br}(x)$ as in (19) and derive $\hat{\mu}^*_{\mathrm{d}}(x)$ and $\sigma^{2*}(x)$ from (13).

6 With probability $1 \wedge (k/K)$, recalculate $\hat{\Sigma}^*_{\mathrm{d}}(x)$ using (15); otherwise, set $\hat{\Sigma}^*_{\mathrm{d}}(x) = \hat{\Sigma}_{\mathrm{d}}(x)$.

7 Compute $\mathrm{CI}^*_{\mathrm{d}}(x)$ using (11) with $\hat{\mu}^*_{\mathrm{d}}(x)$ and $\hat{\Sigma}^*_{\mathrm{d}}(x)$.

2. Setup

When using a Mondrian random forest, there are two sources of randomness. The first is the data, and here we consider the nonparametric regression setting with d-dimensional covariates. The second source is injected purposely from a collection of independent trees drawn from a Mondrian process using a specified lifetime parameter.

2.1 Data generation

We begin with a definition of Hölder continuity, which is used to determine a target class of regression functions, and which participates in controlling the bias of various estimators.

Definition 1 (Hölder continuity) Take $\beta > 0$ and define $\underline{\beta} = \lceil \beta - 1 \rceil$ as the largest integer strictly less than β. We say a function $g : [0, 1]^d \to \mathbb{R}$ is β-Hölder continuous and we write $g \in \mathcal{H}^{\beta}$ if g is $\underline{\beta}$ times differentiable and $\max_{|\nu|=\underline{\beta}} |\partial^{\nu} g(x) - \partial^{\nu} g(x')| \le C \|x - x'\|_2^{\beta - \underline{\beta}}$ for some constant $C > 0$ and all $x, x' \in [0, 1]^d$. Here, $\nu \in \mathbb{N}^d$ is a multi-index with $|\nu| = \sum_{j=1}^d \nu_j$ and $\partial^{\nu} g(x) = \partial^{|\nu|} g(x) \big/ \prod_{j=1}^d \partial x_j^{\nu_j}$.

Throughout this paper, we assume that the data satisfies the following assumption.

Assumption 1 (Data generation). Fix $d \ge 1$ and let (X_i, Y_i) be independent and identically distributed (i.i.d.) samples from a distribution on $\mathbb{R}^d \times \mathbb{R}$, writing $\mathbf{X} = (X_1, \ldots, X_n)$ and $\mathbf{Y} = (Y_1, \ldots, Y_n)$. Suppose X_i has Lebesgue density function $f(x)$ on $[0, 1]^d$ which is bounded away from zero and satisfies $f \in \mathcal{H}^{\beta_f}$ for some $\beta_f > 0$. Suppose $\mathbb{E}[Y_i^2 \mid X_i]$ is bounded, let $\mu(X_i) = \mathbb{E}[Y_i \mid X_i]$, and assume $\mu \in \mathcal{H}^{\beta_\mu}$ where $\beta_\mu > 0$. Let $\varepsilon_i = Y_i - \mu(X_i)$ and assume $\sigma^2(X_i) = \mathbb{E}[\varepsilon_i^2 \mid X_i]$ is bounded away from zero with $\sigma^2 \in \mathcal{H}^{\beta_\sigma}$ for some $\beta_\sigma > 0$. Set $\beta = \beta_\mu \wedge (\beta_f + 1)$.

Some comments are in order surrounding Assumption 1. The requirement that the covariate density $f(x)$ should be strictly positive on all of $[0, 1]^d$ may seem restrictive, particularly when d is moderately large. However, since our theory is presented pointwise in the design point x, it is sufficient for this condition to hold only on some neighborhood of x. To see this, note that continuity implies the density is positive on some hypercube containing x. Upon rescaling the covariates, this hypercube can be mapped onto $[0, 1]^d$. The same argument holds for the Hölder smoothness assumptions, and for the upper and lower bounds on the conditional variance function.

The parameter β represents the effective degree of smoothness which is captured by the (suitably debiased) Mondrian random forest. Its definition is motivated as follows: firstly, we take $\beta \leq \beta_\mu$ in order to compare our rates of convergence with classical results for β-Hölder regression functions. Secondly, due to the presence of design bias, we require in our analysis that the density function $f(x)$ should also be smooth (though not necessarily as smooth as μ), imposing $\beta \leq \beta_f + 1$. Our proofs characterize the roles of each of the smoothness parameters β_μ, β_f and β_σ precisely, though our main results depend only on β. By allowing for $\beta_f < 1$, we strictly generalize the Lipschitz density assumption of Mourtada et al. (2020, Theorem 3).

2.2 The Mondrian process

The Mondrian process was introduced by Roy and Teh (2008) and offers a canonical method for generating random rectangular partitions, which can be used as the trees for a random forest (Lakshminarayanan et al., 2014). For the reader's convenience, we give a brief description of this process here; see Mourtada et al. (2020, Section 3) for a more complete construction.

For a fixed dimension d and lifetime parameter $\lambda > 0$, the Mondrian process is a stochastic process taking values in the set of finite rectangular partitions of $[0, 1]^d$. For a rectangle $D = \prod_{j=1}^d [a_j, b_j] \subseteq [0, 1]^d$, we denote the side aligned with dimension j by $D_j = [a_j, b_j]$, write $D_j^- = a_j$ and $D_j^+ = b_j$ for its left and right endpoints respectively, and use $|D_j| = D_j^+ - D_j^-$ for its length. The volume of D is $|D| = \prod_{j=1}^d |D_j|$ and its linear dimension is $|D|_1 = \sum_{j=1}^d |D_j|$.

To sample a partition T from the Mondrian process $\mathcal{M}([0, 1]^d, \lambda)$, start at time $t = 0$ with the trivial partition of $[0, 1]^d$ which has no splits. Then repeatedly apply the following procedure to each cell D in the partition. Let t_D be the time at which the cell was formed, and sample $E_D \sim \text{Exp}(|D|_1)$, where $\text{Exp}(a)$ is the exponential distribution on $[0, \infty)$ with Lebesgue density ae^{-ax}. If $t_D + E_D \leq \lambda$, then split D. This is done by first selecting a split dimension J with $\mathbb{P}(J = j) = |D_j|/|D|_1$, and then sampling a split location $S_J \sim \text{Unif}[D_J^-, D_J^+]$. The cell D splits into the two new cells $\{x \in D : x_J \leq S_J\}$ and $\{x \in D : x_J > S_J\}$, each with formation time $t_D + E_D$. The output of this sampling procedure is the partition T consisting of the cells D which were not split because $t_D + E_D > \lambda$. The cell in T containing a point $x \in [0, 1]^d$ is written $T(x)$. Figure 1 shows typical realizations of $T \sim \mathcal{M}([0, 1]^d, \lambda)$ for $d = 2$ and with different lifetime parameters λ.

2.3 Mondrian random forests

We define the Mondrian random forest estimator (1) as in Lakshminarayanan et al. (2014) and Mourtada et al. (2020), and will later extend it to a debiased version in Section 4. For a lifetime parameter $\lambda > 0$ and forest size $B \geq 1$, let $\mathbf{T} = (T_1, \ldots, T_B)$ be a Mondrian forest where $T_b \sim \mathcal{M}([0, 1]^d, \lambda)$ are mutually independent Mondrian trees which are independent of the data. For $x \in [0, 1]^d$, we write $N_b(x) = \sum_{i=1}^n \mathbb{I}\{X_i \in T_b(x)\}$ for the number of samples in $T_b(x)$, with $\mathbb{I}$ denoting an indicator function. Then the Mondrian random forest estimator of $\mu(x)$ is

$$\hat{\mu}(x) = \frac{1}{B} \sum_{b=1}^{B} \frac{\sum_{i=1}^n Y_i \mathbb{I}\{X_i \in T_b(x)\}}{N_b(x)}. \tag{1}$$

If there are no samples X_i in $T_b(x)$ then $N_b(x) = 0$, so we define $0/0 = 0$ (see the supplementary material Cattaneo, Klusowski, Underwood, 2025 for details). To ensure the bias and variance of the Mondrian random forest estimator converge to zero (see Section 3), and to avoid boundary issues,

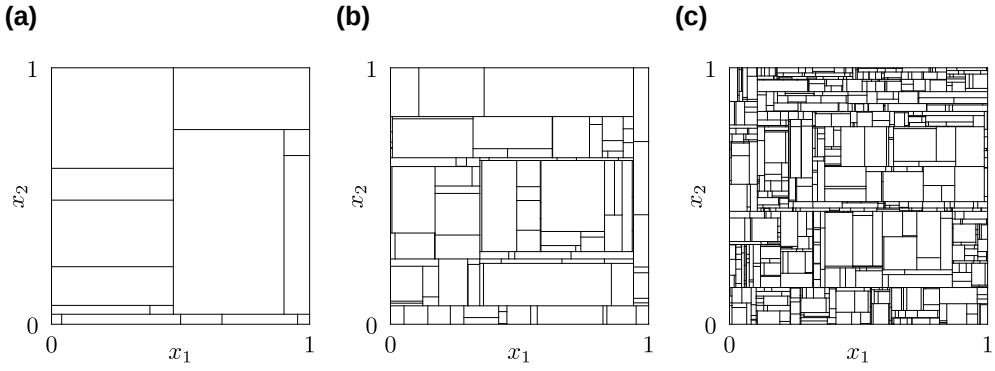

Figure 1. The Mondrian process $T \sim \mathcal{M}([0, 1]^d, \lambda)$ with $d = 2$ and lifetime parameters λ. (a) $\lambda = 3$. (b) $\lambda = 10$ and (c) $\lambda = 30$.

we impose some basic conditions on x, λ, and B in Assumption 2. We emphasize that our results are for the low-dimensional nonparametric regime where d does not depend on the sample size n.

Assumption 2 (Mondrian random forest estimator). Suppose $x \in (0, 1)^d$ is a fixed interior point of the support of X_i, and also that $\lambda \gtrsim (\log n)^3$, $n/\lambda^d \to \infty$, and $B \gtrsim (\log n)^d$.

The requirement that $n/\lambda^d \to \infty$ ensures that the number of data points $N_b(x)$ falling inside a typical Mondrian cell $T_b(x)$, and hence the effective sample size of the Mondrian random forest, diverges in large samples. Assumption 2 also implies that the size of the forest B should grow with n. In our theory, the forest size B plays three roles. Firstly, we assume $B \gtrsim (\log n)^d$ to ensure sufficient averaging across independent trees; this 'large forest' condition is used in the central limit theorem proof (Theorems 2 and 5) to overcome heavy tail issues from very small Mondrian cells, and not to control bias. Secondly, the Berry–Esseen bounds for our centered estimators (Theorems 2 and 5) include a $1/B$ term, so increasing B improves the Gaussian approximation. Thirdly, for feasible inference and mean squared error optimality, B must be large enough that the finite-forest squared bias terms (e.g. $1/(\lambda^{2(1\wedge\beta)}B)$; see Lemma 1 and Theorems 1, 3, and 4) are negligible relative to the leading 'variance' term of the infinite-forest estimator, which scales as λ^d/n (Lemma 2), necessitating that B grows as a positive fractional power of n. Thus, for the purpose of both meeting our statistical assumptions and mitigating the computational burden, we recommend choosing $B \asymp \sqrt{n}$ for Mondrian random forests; selecting B for our debiased estimator requires a different set of conditions (see Sections 4 and 5). Large forests usually do not present significant computational challenges in practice as the ensemble estimator is easily parallelizable over the trees (see Section 6 for more discussion). We will emphasize explicitly where the 'large forest' condition is important for our theory.

3. Inference with Mondrian random forests

Our analysis begins with a standard conditional bias–variance decomposition for the Mondrian random forest estimator:

$$\begin{aligned}\hat{\mu}(x) - \mu(x) &= \Big(\hat{\mu}(x) - \mathbb{E}\big[\hat{\mu}(x) \mid \mathbf{X}, \mathbf{T}\big]\Big) + \Big(\mathbb{E}\big[\hat{\mu}(x) \mid \mathbf{X}, \mathbf{T}\big] - \mu(x)\Big) \\ &= \frac{1}{B}\sum_{b=1}^{B} \frac{\sum_{i=1}^{n} \big(\mu(X_i) - \mu(x)\big)\mathbb{I}\big\{X_i \in T_b(x)\big\}}{N_b(x)} \qquad (2)\end{aligned}$$

$$+ \frac{1}{B}\sum_{b=1}^{B} \frac{\sum_{i=1}^{n} \varepsilon_i \mathbb{I}\big\{X_i \in T_b(x)\big\}}{N_b(x)}. \qquad (3)$$

Our approach to estimation and inference is as follows. Firstly, we precisely characterize the probability limit of the 'bias' term (2), and compute the second conditional moment of the 'variance' term (3). This allows us to understand the bias–variance trade-off, and to derive upper bounds on the rate of convergence for the Mondrian random forest point estimator.

Secondly, we provide a central limit theorem for the 'variance' term (3). By ensuring that the standard deviation dominates the conditional bias, we may conclude that a corresponding central limit theorem holds for the Mondrian random forest (1). With an appropriate estimator for the variance, we then establish procedures for valid and feasible statistical inference on the unknown regression function $\mu(x)$.

3.1 Bias and variance characterizations

We begin with (2), which captures the bias of the Mondrian random forest estimator conditional on the covariates $\mathbf{X}$ and the forest $\mathbf{T}$. The next lemma demonstrates that this conditional bias converges in L^2 at a certain rate, and provides a precise characterization of the resulting nonrandom limiting bias.

Lemma 1 (Bias). *Suppose Assumptions 1 and 2 hold. For each* $1 \le r \le \lfloor \beta/2 \rfloor$ *there exists* $B_r(x) \in \mathbb{R}$, *which is a function of the derivatives of* f *and* μ *at* $\bar{\mathrm{x}}$ *up to order* $2r$, *with*

$$\mathbb{E}\left[\left(\mathbb{E}\big[\hat{\mu}(x) \mid \mathbf{X}, \mathbf{T}\big] - \mu(x) - \sum_{r=1}^{\lfloor \beta/2 \rfloor} \frac{B_r(x)}{\lambda^{2r}}\right)^2\right] \lesssim \frac{1}{\lambda^{2\beta}} + \frac{1}{\lambda^{2(1\wedge\beta)} B} + \frac{1}{\lambda^{2(1\wedge\beta)}} \frac{\lambda^d}{n}. \quad (4)$$

Whenever $\beta > 2$, *the leading bias is the quadratic term*

$$\frac{B_1(x)}{\lambda^2} = \frac{1}{2\lambda^2} \sum_{j=1}^{d} \frac{\partial^2 \mu(x)}{\partial x_j^2} + \frac{1}{2\lambda^2} \frac{1}{f(x)} \sum_{j=1}^{d} \frac{\partial \mu(x)}{\partial x_j} \frac{\partial f(x)}{\partial x_j}.$$

If $X_i \sim \mathrm{Unif}([0,1]^d)$ *then* $f(x) = 1$, *and using multi-index notation we have*

$$\frac{B_r(x)}{\lambda^{2r}} = \frac{1}{\lambda^{2r}} \sum_{|v|=r} \partial^{2v} \mu(x) \prod_{j=1}^{d} \frac{1}{v_j + 1}.$$

The bias characterization in Lemma 1 incorporates some high-degree polynomial terms in the lifetime parameter λ which for now may seem ignorable. The magnitude of the bias is determined by the leading term in (4), typically of order $1/\lambda^2$ whenever $\beta \ge 2$. This suffices for ensuring a negligible contribution from the bias with an appropriate choice of lifetime. However, the advantage of specifying higher-order terms will become apparent in Section 4, where we construct a debiased Mondrian random forest estimator, directly targeting and annihilating the higher-order terms in order to furnish superior estimation and inference properties. We also demonstrate numerically the detrimental role of bias in estimation and inference in Section 5.

In Lemma 1 we give some explicit examples of calculating the limiting bias when $\beta > 2$ or X_i are uniformly distributed. The general form of $B_r(x)$ is provided in the supplementary material (Cattaneo, Klusowski, Underwood, 2025) but is somewhat unwieldy except in specific situations. Nonetheless, the most important properties are that $B_r(x)$ are nonrandom and do not depend on the lifetime λ; these are crucial features for our debiasing procedure given in Section 4. If the forest size B does not diverge to infinity then we suffer the first-order conditional bias term $1/(\lambda^{1\wedge\beta}\sqrt{B})$. This phenomenon was explained by Mourtada et al. (2020), who noted that it allows individual Mondrian trees ($B = 1$) to achieve minimax optimality in integrated mean squared error only when

$\beta \in (0, 1)$. In contrast, large forests remove this first-order bias through ensemble averaging and as such are optimal for all $\beta \in (0, 2)$.

We now turn to (3), which captures the stochastic part of the Mondrian random forest. Lemma 2 determines the probability limit of the scaled conditional variance of this term, alongside its L^2 convergence rate. First, define

$$\tilde{\Sigma}(x) = \frac{n}{\lambda^d}\mathrm{Var}\big[\hat{\mu}(x) \mid \mathbf{X}, \mathbf{T}\big] \quad \text{and} \quad \Sigma(x) = \frac{\sigma^2(x)}{f(x)}\left(\frac{4 - 4\log 2}{3}\right)^d.$$

Lemma 2 (Variance). Suppose Assumptions 1 and 2 hold. Then

$$\mathbb{E}\Big[\big(\tilde{\Sigma}(x) - \Sigma(x)\big)^2\Big] \lesssim \frac{\lambda^d}{n} + \frac{1}{B} + \frac{1}{\lambda^{2(1\wedge\beta_f\wedge\beta_\sigma)}}.$$

As $n/\lambda^d \to \infty$, $B \to \infty$ and $\lambda \to \infty$ by Assumption 2, it follows from Lemma 2 that

$$\mathbb{E}\Big[\mathrm{Var}\big[\hat{\mu}(x) \mid \mathbf{X}, \mathbf{T}\big]\Big] = \frac{\lambda^d}{n}\mathbb{E}\big[\tilde{\Sigma}(x)\big] \lesssim \frac{\lambda^d}{n}\left(\Sigma(x) + \sqrt{\frac{\lambda^d}{n}} + \frac{1}{\sqrt{B}} + \frac{1}{\lambda^{1\wedge\beta_f\wedge\beta_\sigma}}\right) \lesssim \frac{\lambda^d}{n}. \tag{5}$$

An upper bound on the L^2 rate of convergence of the Mondrian random forest estimator can therefore be deduced from the bias–variance decomposition, Lemma 1, and (5). This rate of convergence depends on the sequence of lifetime parameters λ; for optimal point estimation, we may balance the contributions from the bias and from the standard deviation by ensuring that $1/\lambda^{2\wedge\beta} + 1/(\lambda^{1\wedge\beta}\sqrt{B}) \asymp \sqrt{\lambda^d/n}$, or equivalently if $\lambda \asymp n^{\frac{1}{d+2(2\wedge\beta)}}$ and $B \gtrsim n^{\frac{2(2\wedge\beta)-2(1\wedge\beta)}{d+2(2\wedge\beta)}}$. We formalize these deductions in Theorem 1 and note that they imply that the Mondrian random forest is rate-minimax-optimal (Stone, 1982) in pointwise mean squared error for β-Hölder functions with $\beta \in (0, 2)$; a corresponding result for integrated mean squared error was provided by Mourtada et al. (2020, Theorem 2).

Theorem 1 (Mean squared error). Suppose Assumptions 1 and 2 hold. Then

$$\mathbb{E}\Big[\big(\hat{\mu}(x) - \mu(x)\big)^2\Big] \lesssim \frac{\lambda^d}{n} + \frac{1}{\lambda^{2(2\wedge\beta)}} + \frac{1}{\lambda^{2(1\wedge\beta)}B}.$$

If further $\lambda \asymp n^{\frac{1}{d+2(2\wedge\beta)}}$ and $B \gtrsim n^{\frac{2(2\wedge\beta)-2(1\wedge\beta)}{d+2(2\wedge\beta)}}$, then

$$\mathbb{E}\Big[\big(\hat{\mu}(x) - \mu(x)\big)^2\Big] \lesssim n^{-\frac{2(2\wedge\beta)}{d+2(2\wedge\beta)}}.$$

We take this opportunity to contrast Mondrian random forests with classical nonparametric local smoothing methods. For example, the lifetime λ plays a similar role to the inverse bandwidth for kernel smoothing as it determines both the effective sample size n/λ^d and the scale of localization $1/\lambda$, and thus also the associated rate of convergence. Likewise, $1/\lambda$ controls the diameter of a typical cell in Mondrian partition-based smoothing. However, due to the Mondrian process construction, some cells are typically 'too small' (equivalent to an insufficiently large bandwidth) to give an appropriate effective sample size. In the same manner, classical methods based on nonrandom partitioning such as spline estimators typically impose a quasi-uniformity assumption to ensure all the cells are of comparable size (Cattaneo et al., 2020; Huang, 2003), a property which does not hold for the Mondrian process (not even with high probability).

3.2 Central limit theorem

Having discussed the point estimation properties of the Mondrian random forest estimator, we present a central limit theorem which forms the core of our methodology for performing statistical inference. As well as establishing asymptotic normality of the appropriately centered and scaled estimator, we also provide a rate of convergence in terms of a Berry–Esseen-style bound on the Kolmogorov–Smirnov distance from the normal distribution. In addition to precisely quantifying the quality of the Gaussian distributional approximation, this allows us to obtain explicit bounds on the coverage error of feasible confidence intervals.

Before stating the theorem, we highlight some of the challenges involved in establishing such a result. At first glance, the summands in (3) appear independent over $1 \le i \le n$, conditional on the forest $\mathbf{T}$, depending only on X_i and ε_i. However, the $N_b(x)$ appearing in the denominator depends on all X_i simultaneously, violating this independence assumption and rendering classical central limit theorems inapplicable. A natural preliminary attempt to resolve this issue is to observe that

$$N_b(x) = \sum_{i=1}^{n} \mathbb{I}\{X_i \in T_b(x)\} \approx n\mathbb{P}(X_i \in T_b(x) \mid T_b) \approx nf(x)|T_b(x)|$$

with high probability. One could attempt to use this by approximating the estimator with an average of i.i.d. random variables, or by employing a central limit theorem conditional on $\mathbf{X}$ and $\mathbf{T}$. However, such an approach fails because $\mathbb{E}[1/|T_b(x)|^2] = \infty$; the possible existence of small cells causes the law of the inverse cell volume to have heavy tails. For similar reasons, attempts to directly establish a central limit theorem based on $2+\delta$ moments, such as the classical Lyapunov central limit theorem, are ineffective.

We circumvent these problems by directly analyzing $\mathbb{I}\{N_b(x) \ge 1\}/N_b(x)$. We establish concentration properties for this non-linear function of X_i via the Efron–Stein inequality (Boucheron et al., 2013, Section 3.1) along with a sequence of delicate preliminary lemmas regarding inverse moments of truncated (conditional) binomial random variables. In particular, we show that $\mathbb{E}[\mathbb{I}\{N_b(x) \ge 1\}/N_b(x)] \lesssim \lambda^d/n$ and $\mathbb{E}[\mathbb{I}\{N_b(x) \ge 1\}/N_b(x)^2] \lesssim \lambda^{2d}(\log n)^d/n^2$. Asymptotic normality is then established by a careful application of a Berry–Esseen theorem (Petrov, 1995) conditional on $(\mathbf{X}, \mathbf{T})$. The supplementary material (Cattaneo, Klusowski, Underwood, 2025) provides all the technical details.

The following theorem gives our Berry–Esseen-type central limit theorem for the centered (zero mean conditional on the covariates and the trees) 'variance' term from (3), scaled and standardized by its conditional variance $\tilde{\Sigma}(x)$. Note that on the event $\tilde{\Sigma}(x) = 0$, we also have $\hat{\mu}(x) = 0$ and $\mathbb{E}[\hat{\mu}(x) \mid \mathbf{X}, \mathbf{T}] = 0$, so continue to define $0/0 = 0$.

Theorem 2 (Central limit theorem). If Assumptions 1 and 2 hold, and $\mathbb{E}[|Y_i|^{2+\delta} \mid X_i]$ is bounded almost surely with $\delta > 0$, then

$$\sup_{t\in\mathbb{R}} \left| \mathbb{P}\left(\sqrt{\frac{n}{\lambda^d}} \frac{\hat{\mu}(x) - \mathbb{E}[\hat{\mu}(x) \mid \mathbf{X}, \mathbf{T}]}{\sqrt{[b]\tilde{\Sigma}(x)}} \le t \right) - \Phi(t) \right| \lesssim \left(\frac{\lambda^d}{n}\right)^{\frac{1\wedge\delta}{2}} + \frac{1}{B}. \tag{6}$$

We make some remarks on Theorem 2. Firstly, since n/λ^d is the effective sample size and Y_i has only $2+\delta$ finite moments, the first term in (6) is likely to be unimprovable (Ibragimov & Linnik, 1975, Theorem 3.4.1). In particular, we attain the classical Berry–Esseen rate when $\mathbb{E}[|Y_i|^3 \mid \mathbf{X}]$ is bounded and $B \gtrsim \sqrt{n/\lambda^d}$.

The condition of $B \gtrsim (\log n)^d$ is central to our proof of Theorem 2, ensuring sufficient 'mixing' of different Mondrian cells to escape the heavy-tailed phenomenon detailed in the preceding discussion. For concreteness, the large forest condition allows us to deal with expressions such as $\mathbb{E}[1/(|T_b(x)||T_{b'}(x)|)] = \mathbb{E}[1/|T_b(x)|]\mathbb{E}[1/|T_{b'}(x)|] \approx \lambda^{2d} < \infty$ where $b \ne b'$, by independence of the trees, rather than the 'no ensembling' single tree analog $\mathbb{E}[1/|T_b(x)|^2] = \infty$.

Nonetheless, it is not clear whether the $1/B$ term is strictly necessary in (6) or if it is an artifact of the proof. When B is bounded, $\tilde{\Sigma}(x)$ remains random in the limit, and in fact it is not difficult to show that in this regime we have that $\mathbb{E}[\{\tilde{\Sigma}(x)\}^2] \geq (\log n)^d$, which diverges (cf., Lemma 2). While these mildly pathological properties may not necessarily render the central limit theorem invalid, they certainly highlight some issues associated with inference based on a single tree or a small forest.

Theorem 2 applies only to the centered Mondrian random forest estimator; in order for it to be useful in a feasible inference setting, we must combine it with methods for controlling the conditional bias (see Lemma 1). In Section 4 we will show how the estimator can be debiased, giving weaker lifetime conditions for inference, improved rates of convergence, and superior coverage guarantees, whenever additional smoothness is available.

3.3 Confidence intervals

We demonstrate how to use our previous results to construct valid confidence intervals for the regression function $\mu(x)$. To do this, there are two preliminary issues which must be resolved. Firstly, the Berry–Esseen central limit theorem presented in Theorem 2 is stated for the Mondrian random forest estimator $\hat{\mu}(x)$ centered at its conditional expectation $\mathbb{E}[\hat{\mu}(x) \mid \mathbf{X}, \mathbf{T}]$, rather than at the true value $\mu(x)$. As such, we use Lemma 1 to ensure that the bias $\mathbb{E}[\hat{\mu}(x) \mid \mathbf{X}, \mathbf{T}] - \mu(x)$ is taken into account when establishing procedures for inference. Specifically, the bias should shrink faster than the standard deviation; this requires $1/\lambda^{2\wedge\beta} + 1/(\lambda^{1\wedge\beta}\sqrt{B}) \ll \sqrt{\lambda^d/n}$, which is satisfied by imposing the restrictions $\lambda \gg n^{\frac{1}{d+2(2\wedge\beta)}}$ and $B \gg n^{\frac{2(2\wedge\beta)-2(1\wedge\beta)}{d+2(2\wedge\beta)}}$ on the lifetime λ and forest size B.

The second issue is that the variances $\tilde{\Sigma}(x)$ and $\Sigma(x)$ depend on the unknown quantities $\sigma^2(x)$ and $f(x)$. To conduct feasible inference, we must therefore provide a consistent variance estimator. To this end, define

$$\begin{aligned} \hat{\sigma}^2(x) &= \tfrac{1}{B}\sum_{b=1}^{B}\sum_{i=1}^{n}\frac{\big(Y_i - \hat{\mu}(x)\big)^2 \mathbb{I}\{X_i \in T_b(x)\}}{N_b(x)}, \\ \hat{\Sigma}(x) &= \hat{\sigma}^2(x)\,\tfrac{n}{\lambda^d}\sum_{i=1}^{n}\left(\frac{1}{B}\sum_{b=1}^{B}\frac{\mathbb{I}\{X_i \in T_b(x)\}}{N_b(x)}\right)^2. \end{aligned} \tag{7}$$

Lemma 3 (Variance estimation). If Assumptions 1 and 2 hold, and $\mathbb{E}[|Y_i|^{2+\delta} \mid X_i]$ is bounded almost surely with $\delta > 0$, then

$$\left(\mathbb{E}\left[\left|\hat{\Sigma}(x) - \Sigma(x)\right|^{\frac{2-\mathbb{I}\{\delta<2\}}{2}}\right]\right)^{\frac{2}{2-\mathbb{I}\{\delta<2\}}} \lesssim \left(\frac{\lambda^d}{n}\right)^{\frac{1}{2}\frac{\mathbb{I}\{\delta<2\}}{2+\delta}} + \frac{1}{\sqrt{B}} + \frac{1}{\lambda^{1\wedge\beta_\mu\wedge\beta_f\wedge\beta_\sigma}}.$$

For a confidence level $\alpha \in (0, 1)$, Theorem 3 shows how to construct an asymptotically valid $100(1-\alpha)\%$ confidence interval for the regression function $\mu(x)$. The restrictions on the lifetime λ and forest size B are the same as those previously discussed, and an explicit upper bound on the coverage error rate is provided. Define the interval estimator

$$\mathrm{CI}(x) = \left[\hat{\mu}(x) - \sqrt{\frac{\lambda^d}{n}}\hat{\Sigma}(x)^{1/2} q_{1-\alpha/2},\ \ \hat{\mu}(x) - \sqrt{\frac{\lambda^d}{n}}\hat{\Sigma}(x)^{1/2} q_{\alpha/2}\right].$$

Theorem 3 (Confidence intervals). If Assumptions 1 and 2 hold, and $\mathbb{E}[|Y_i|^{2+\delta} \mid X_i]$ is bounded almost surely with $\delta > 0$, then

$$\left|\mathbb{P}\big(\mu(x) \in \mathrm{CI}(x)\big) - (1-\alpha)\right| \lesssim \frac{n}{\lambda^d}\frac{1}{\lambda^{2(2\wedge\beta)}} + \left(\left(\frac{\lambda^d}{n}\right)^{1-\frac{2\mathbb{1}[\delta<2]}{2+\delta}} + \frac{1}{B} + \frac{1}{\lambda^{2(1\wedge\beta_\mu\wedge\beta_f\wedge\beta_\sigma)}} + \frac{n}{\lambda^d}\frac{1}{\lambda^{2(1\wedge\beta)}B}\right)^{\frac{1}{3+2\mathbb{1}[\delta<2]}}.$$

When coupled with an appropriate lifetime selection method (see Section 5), Theorem 3 gives a feasible procedure for uncertainty quantification in Mondrian random forests. Our procedure requires no adjustment of the original Mondrian random forest estimator beyond ensuring that the bias is negligible, and in particular does not rely on sample splitting. The construction of confidence intervals is just one corollary of the result given in Theorem 2; other applications include hypothesis testing based on the value of $\mu(x)$ at a design point x by inversion of the confidence interval, as well as specification testing by comparison with a $\sqrt{n}$-consistent parametric regression estimator. The construction of simultaneous confidence intervals for finitely many points $x_1, \ldots, x_D$ can be accomplished either using standard multiple testing corrections or by first establishing a multivariate central limit theorem using the Cramér–Wold device and formulating a consistent variance matrix estimator.

4. Debiased Mondrian random forests

We give our next main contribution: a novel variant of the Mondrian random forest estimator that corrects for higher-order bias with an approach based on generalized jackknifing (Schucany & Sommers, 1977). This estimator retains the basic form of a Mondrian random forest in the sense that it is a linear combination of Mondrian tree estimators, but in this section we allow for non-identical linear coefficients, some of which may be negative, and for differing lifetime parameters across the trees. Since the basic Mondrian random forest estimator is a special case of this more general debiased version, we will discuss only the latter throughout the rest of the paper.

We use the explicit form of the bias given in Lemma 1 to construct the debiased Mondrian forest estimator as follows, letting $J \geq 0$ be the bias correction order. With $J = 0$ we preserve the original Mondrian random forest, with $J = 1$ we remove second-order bias, and with $J = \lfloor \underline{\beta}/2 \rfloor$ we remove bias terms up to and including order $2\lfloor \underline{\beta}/2 \rfloor$, giving the maximum possible bias reduction achievable in the Hölder class $\mathcal{H}^{\underline{\beta}}$ (Stone, 1982). As such, only bias terms of order $1/\lambda^{\underline{\beta}}$ will remain.

For $0 \leq r \leq J$, let $\hat{\mu}_r(x)$ be a Mondrian forest estimator based on the trees $T_{br} \sim \mathcal{M}([0,1]^d, \lambda_r)$ for $1 \leq b \leq B$, where $\lambda_r = a_r\lambda$ for some $a_r > 0$ and $\lambda > 0$. We write $\mathbf{T}$ to denote the collection of all the trees, and suppose they are mutually independent. We find values of a_r along with coefficients $\omega_r \in \mathbb{R}$ which annihilate the leading J bias terms of the debiased Mondrian random forest estimator

$$\hat{\mu}_{\mathrm{d}}(x) = \sum_{r=0}^{J} \omega_r \hat{\mu}_r(x) = \sum_{r=0}^{J} \omega_r \frac{1}{B}\sum_{b=1}^{B} \frac{\sum_{i=1}^{n} Y_i \mathbb{1}\{X_i \in T_{br}(x)\}}{N_{br}(x)}. \tag{8}$$

This ensemble estimator retains the 'forest' structure of the original estimators, but with varying lifetime parameters λ_r and coefficients ω_r. Thus, referring to (4), we desire

$$\sum_{r=0}^{J} \omega_r \left(\mu(x) + \sum_{s=1}^{J} \frac{B_s(x)}{a_r^{2s}\lambda^{2s}}\right) = \mu(x)$$

for all λ, or equivalently the system of linear equations $\sum_{r=0}^{J} \omega_r = 1$ and $\sum_{r=0}^{J} \omega_r a_r^{-2s} = 0$ for each $1 \leq s \leq J$. We solve these as follows. Define the $(J+1) \times (J+1)$ Vandermonde matrix $A_{rs} = a_{r-1}^{2-2s}$,

let $\omega = (\omega_0, \ldots, \omega_J)^T \in \mathbb{R}^{J+1}$ and take $e_0 = (1, 0, \ldots, 0)^T \in \mathbb{R}^{J+1}$. Then a solution for the debiasing coefficients is given by $\omega = A^{-1}e_0$ whenever A is nonsingular. In practice we can take a_r to be a fixed geometric or arithmetic sequence to ensure this is the case, appealing to the Vandermonde determinant formula: $\det A = \prod_{0\le r<s\le J}(a_r^{-2} - a_s^{-2}) \neq 0$ whenever a_r are distinct. For example, one could set $a_r = (1+\gamma)^r$ or $a_r = 1 + \gamma r$ for some $\gamma > 0$. Because we assume β, and therefore the choice of J, do not depend on n, there is no need to quantify the invertibility of A by, for example, bounding its eigenvalues away from zero as a function of J and the choice of a_r.

The debiased Mondrian random forest estimator defined in (8) is a linear combination of standard Mondrian random forests, and as such contains both a sum over $0 \le r \le J$, representing the debiasing procedure, and a sum over $1 \le b \le B$, representing the forest averaging. We have been interpreting this estimator as a debiased version of the standard Mondrian random forest given in (1), but it is equally valid to swap the order of these sums. This gives rise to an alternative point of view: we replace each Mondrian random tree with a 'debiased' version, and then take a forest of such modified trees. This perspective is perhaps more in line with existing techniques for constructing randomized ensembles, where the outermost operation represents a B-fold average of randomized base learners (not necessarily locally constant decision trees), each of which has a small bias component (Caruana et al., 2004; Zhou & Feng, 2019).

4.1 Bias and variance characterizations

In Lemma 4 we verify that this debiasing procedure does indeed annihilate the desired bias terms; it is a direct consequence of Lemma 1 and the construction of the debiased Mondrian random forest estimator $\hat{\mu}_{\mathrm{d}}(x)$.

Lemma 4 (Bias of the debiased estimator). Suppose Assumptions 1 and 2 hold. Then in the notation of Lemma 1 and with $\bar{\omega} = \sum_{r=0}^{J} \omega_r a_r^{-2J-2}$,

$$\mathbb{E}\left[\left(\mathbb{E}\big[\hat{\mu}_{\mathrm{d}}(x) \mid \mathbf{X}, \mathbf{T}\big] - \mu(x) - \mathbb{I}\{2J+2<\beta\}\frac{\bar{\omega}B_{J+1}(x)}{\lambda^{2J+2}}\right)^2\right] \lesssim \frac{1}{\lambda^{2((2J+4)\wedge\beta)}} + \frac{1}{\lambda^{2(1\wedge\beta)}B} + \frac{1}{\lambda^{2(1\wedge\beta)}}\frac{\lambda^d}{n}.$$

Lemma 4 has the following consequence: the leading bias term is characterized in terms of $B_{J+1}(x)$ whenever $J < \beta/2 - 1$, or equivalently $J < \lfloor \underline{\beta}/2 \rfloor$, that is, the debiasing order J does not exhaust the Hölder smoothness β. If this condition does not hold, then the estimator is fully debiased; the resulting leading bias term is bounded above by $1/\lambda^{\beta}$ up to constants but its form is left unspecified.

The following lemma controls the variance of the debiased Mondrian random forest estimator. With $\ell_{rr'} = 2a_r(1 - a_r \log(1 + a_{r'}/a_r)/a_{r'})/3$, define

$$\tilde{\Sigma}_{\mathrm{d}}(x) = \frac{n}{\lambda^d}\mathrm{Var}\big[\hat{\mu}_{\mathrm{d}}(x) \mid \mathbf{X}, \mathbf{T}\big] \quad \text{and} \quad \Sigma_{\mathrm{d}}(x) = \frac{\sigma^2(x)}{f(x)}\sum_{r=0}^{J}\sum_{r'=0}^{J}\omega_r\omega_{r'}(\ell_{rr'} + \ell_{r'r})^d.$$

Lemma 5 (Variance of the debiased estimator). Supposing Assumptions 1 and 2 hold,

$$\mathbb{E}\Big[\big(\tilde{\Sigma}_{\mathrm{d}}(x) - \Sigma_{\mathrm{d}}(x)\big)^2\Big] \lesssim \frac{\lambda^d}{n} + \frac{1}{B} + \frac{1}{\lambda^{2(1\wedge\beta_f\wedge\beta_\sigma)}}.$$

As in (5), it follows from Lemma 5 that

$$\mathbb{E}\Big[\mathrm{Var}\big[\hat{\mu}_{\mathrm{d}}(x) \mid \mathbf{X}, \mathbf{T}\big]\Big] \lesssim \frac{\lambda^d}{n}. \tag{9}$$

4.2 Minimax optimality

Our next main result, Theorem 4, shows that when using an appropriate sequence of lifetime parameters λ, the debiased Mondrian random forest estimator achieves, up to constants, the minimax-optimal rate of convergence for pointwise mean squared error estimation of a d-dimensional regression function $\mu \in \mathcal{H}^\beta$ (Stone, 1982). This result holds for all $d \geq 1$ and all $\beta > 0$, complementing a previous result (see Theorem 1) established only for $\beta \in (0, 2)$ and in integrated mean squared error by Mourtada et al. (2020).

Theorem 4 (Mean squared error of the debiased estimator). Grant Assumptions 1 and 2. Then

$$\mathbb{E}\Big[\big(\hat{\mu}_{\mathrm{d}}(x) - \mu(x)\big)^2\Big] \lesssim \frac{\lambda^d}{n} + \frac{1}{\lambda^{2((2J+2)\wedge\beta)}} + \frac{1}{\lambda^{2(1\wedge\beta)}B}.$$

Thus with $J \geq \lfloor \underline{\beta}/2 \rfloor$, $\lambda \asymp n^{\frac{1}{d+2\beta}}$ and $B \gtrsim n^{\frac{2\beta-2(1\wedge\beta)}{d+2\beta}}$, we have

$$\mathbb{E}\Big[\big(\hat{\mu}_{\mathrm{d}}(x) - \mu(x)\big)^2\Big] \lesssim n^{-\frac{2\beta}{d+2\beta}}.$$

The sequence of lifetime parameters λ required in Theorem 4 is chosen to balance the bias and standard deviation bounds implied by Lemma 4 and (9) respectively, in order to minimize the pointwise mean squared error. While selecting an optimal debiasing order J needs only knowledge of an upper bound on the smoothness β, choosing an optimal sequence of λ values does assume that β is known a priori. The problem of adapting to β from data is beyond the scope of this paper; we provide some practical advice for tuning parameter selection in Section 5.

Theorem 4 complements the minimaxity results proven by Mourtada et al. (2020) for Mondrian trees (with $\beta \leq 1$) and for Mondrian random forests (with $\beta \leq 2$), with one modification: our version is stated in pointwise rather than integrated mean squared error. This is because our debiasing procedure is designed to handle interior smoothing bias and as such does not provide any correction for boundary bias. We leave the development of such boundary corrections to future work, but constructions similar to higher-order boundary-correcting kernels should be possible. If the region of integration is a compact set in the interior of $[0, 1]^d$ then we do obtain an optimal integrated mean squared error bound: if $a \in (0, 1/2)$ is fixed then under the same conditions as Theorem 4, with appropriate tuning of λ and B,

$$\mathbb{E}\Big[\int_{[a,1-a]^d} \big(\hat{\mu}_{\mathrm{d}}(x) - \mu(x)\big)^2 \,\mathrm{d}x\Big] \lesssim \frac{\lambda^d}{n} + \frac{1}{\lambda^{2\beta}} + \frac{1}{\lambda^{2(1\wedge\beta)}B} \lesssim n^{-\frac{2\beta}{d+2\beta}}.$$

4.3 Central limit theorem

In Theorem 5, we verify that a central limit theorem holds for the debiased random forest estimator $\hat{\mu}_{\mathrm{d}}(x)$. The strategy and challenges associated with proving Theorem 5 are identical to those discussed earlier surrounding Theorem 2. In fact in the supplementary material (Cattaneo, Klusowski, Underwood, 2025) we provide a direct proof only for Theorem 5 and deduce Theorem 2 as a special case. Again on the event $\tilde{\Sigma}_{\mathrm{d}}(x) = 0$, we also have $\hat{\mu}_{\mathrm{d}}(x) = 0$ and $\mathbb{E}[\hat{\mu}_{\mathrm{d}}(x) \mid \mathbf{X}, \mathbf{T}] = 0$, so we take $0/0 = 0$.

Theorem 5 (Central limit theorem with debiasing). Suppose Assumptions 1 and 2 hold, and $\mathbb{E}[|Y_i|^{2+\delta} \mid X_i]$ is bounded almost surely with $\delta > 0$. Then

$$\sup_{t\in\mathbb{R}}\left|\mathbb{P}\left(\sqrt{\frac{n}{\lambda^d}}\frac{\hat{\mu}_{\mathrm{d}}(x) - \mathbb{E}\big[\hat{\mu}_{\mathrm{d}}(x) \mid \mathbf{X}, \mathbf{T}\big]}{\sqrt{\hat{\Sigma}_{\mathrm{d}}(x)}} \le t\right) - \Phi(t)\right| \lesssim \left(\frac{\lambda^d}{n}\right)^{\frac{1\wedge\delta}{2}} + \frac{1}{B}.$$

4.4 Confidence intervals

As before, to conduct valid feasible inference we must ensure that the bias (now significantly reduced due to our debiasing procedure) is negligible when compared to the standard deviation (which is of the same order as before). We treat here the general 'partial debiasing' setting where either the debiasing order J or the Hölder smoothness β may determine the magnitude of the bias, which is $1/\lambda^{(2J+2)\wedge\beta}$. For optimal results, one should take $J \ge \lfloor\beta/2\rfloor$ to ensure total debiasing, as in Theorem 4. We thus require $1/\lambda^{(2J+2)\wedge\beta} + 1/(\lambda^{1\wedge\beta}\sqrt{B}) \ll \sqrt{\lambda^d/n}$, satisfied by imposing $\lambda \gg n^{\frac{1}{d+2((2J+2)\wedge\beta)}}$ and $B \gg n^{\frac{2((2J+2)\wedge\beta)-2(1\wedge\beta)}{d+2((2J+2)\wedge\beta)}}$ on the lifetime parameter λ and forest size B.

Once again, we propose a variance estimator and show that it is consistent. With $\hat{\sigma}^2(x)$ as in (7) in Section 3, define

$$\hat{\Sigma}_{\mathrm{d}}(x) = \hat{\sigma}^2(x)\frac{n}{\lambda^d}\sum_{i=1}^{n}\left(\sum_{r=0}^{J}\omega_r\frac{1}{B}\sum_{b=1}^{B}\frac{\mathbb{I}\{X_i \in T_{br}(x)\}}{N_{br}(x)}\right)^2. \tag{10}$$

Lemma 6 (Variance estimation for the debiased estimator). Suppose Assumptions 1 and 2 hold, and $\mathbb{E}[|Y_i|^{2+\delta} \mid X_i]$ is bounded almost surely with $\delta > 0$. Then

$$\left(\mathbb{E}\left[\left|\hat{\Sigma}_{\mathrm{d}}(x) - \Sigma_{\mathrm{d}}(x)\right|^{\frac{2-\mathbb{I}\{\delta<2\}}{2}}\right]\right)^{\frac{2}{2-\mathbb{I}\{\delta<2\}}} \lesssim \left(\frac{\lambda^d}{n}\right)^{\frac{1}{2}-\frac{\mathbb{I}\{\delta<2\}}{2+\delta}} + \frac{1}{\sqrt{B}} + \frac{1}{\lambda^{1\wedge\beta_\mu\wedge\beta_f\wedge\beta_\sigma}}.$$

In analogy to Section 3, we now demonstrate the construction of feasible valid confidence intervals using the debiased Mondrian random forest estimator in Theorem 6. Consider the debiased $100(1-\alpha)\%$ confidence interval estimator

$$\mathrm{CI}_{\mathrm{d}}(x) = \left[\hat{\mu}_{\mathrm{d}}(x) - \sqrt{\frac{\lambda^d}{n}}\hat{\Sigma}_{\mathrm{d}}(x)^{1/2}q_{1-\alpha/2},\ \ \hat{\mu}_{\mathrm{d}}(x) - \sqrt{\frac{\lambda^d}{n}}\hat{\Sigma}_{\mathrm{d}}(x)^{1/2}q_{\alpha/2}\right]. \tag{11}$$

Theorem 6 (Confidence intervals with debiasing). Suppose Assumptions 1 and 2 hold, and $\mathbb{E}[|Y_i|^{2+\delta} \mid X_i]$ is bounded almost surely with $\delta > 0$. Then

$$\begin{aligned}&\left|\mathbb{P}\big(\mu(x) \in \mathrm{CI}_{\mathrm{d}}(x)\big) - (1-\alpha)\right| \\ &\lesssim \frac{n}{\lambda^d}\frac{1}{\lambda^{2((2J+2)\wedge\beta)}} + \left(\left(\frac{\lambda^d}{n}\right)^{1-\frac{2\mathbb{I}\{\delta<2\}}{2+\delta}} + \frac{1}{B} + \frac{1}{\lambda^{2(1\wedge\beta_\mu\wedge\beta_f\wedge\beta_\sigma)}} + \frac{n}{\lambda^d}\frac{1}{\lambda^{2(1\wedge\beta)}B}\right)^{\frac{1}{5+2\mathbb{I}\{\delta<2\}}}.\end{aligned}$$

One important benefit of our debiasing technique is made clear in Theorem 6: the restrictions imposed on the lifetime parameter λ are substantially relaxed, especially in smooth classes with

large β. As well as the high-level of benefit of relaxed conditions, this is also useful for practical selection of appropriate lifetimes for estimation and inference respectively; see Section 5 for more details. Nonetheless, such improvements do not come without concession. The limiting variance of the debiased estimator is typically larger than that of the unbiased version in small samples (the extent of this increase depends on the choice of the debiasing parameters a_r), leading to wider confidence intervals and larger estimation error, despite the theoretical asymptotic improvements. Nonetheless, the empirical results in Section 5 demonstrate that the debiasing effect can overcome the increased variance with moderate sample sizes. Because we employ symmetric confidence intervals, the coverage error depends on the squared bias $1/\lambda^{2((2J+2)\wedge\beta)}$, whereas the corresponding Berry–Esseen rate would depend on the (larger) linear bias $1/\lambda^{(2J+2)\wedge\beta}$.

5. Implementation and empirical results

We discuss procedures for selecting the parameters involved in fitting a debiased Mondrian random forest; namely the base lifetime parameter λ, the number of trees in each forest B, the order of the bias correction J, and the debiasing scale parameters a_r for $0 \leq r \leq J$. We then provide empirical results with simulated data to demonstrate the effectiveness of our methods.

5.1 Data availability and reproducibility

All code used to generate results in this section (e.g. model training and parameter tuning) is available at https://github.com/WGUNDERWOOD/MondrianForests.jl.

5.2 Tuning parameter selection

The most important parameter is the base Mondrian lifetime λ, which plays the role of a complexity parameter and thus governs the overall bias–variance trade-off of the estimator. Correct tuning of λ is especially important in two main respects: firstly, in order to use the central limit theorem established in Theorem 5, we must have that the bias converges to zero, requiring $\lambda \gg n^{\frac{1}{d+2((2J+2)\wedge\beta)}}$. Secondly, the minimax optimality result of Theorem 4 is valid only in the regime $\lambda \asymp n^{\frac{1}{d+2\beta}}$, and so λ requires careful determination in practice. For clarity, in this section we use the notation $\hat{\mu}_{\mathrm{d}}(x;\lambda,J)$ for the debiased Mondrian random forest estimator implemented with lifetime λ and debiasing order J, as in (8). Similarly, we write $\hat{\Sigma}_{\mathrm{d}}(x;\lambda,J)$ for the associated variance estimator (10).

For minimax-optimal point estimation when β is known (for example, when the data come from noisy measurements of a smooth physical system), choose any sequence $\lambda \asymp n^{\frac{1}{d+2\beta}}$ and use $\hat{\mu}_{\mathrm{d}}(x;\lambda,J)$ with $J = \lfloor \underline{\beta}/2 \rfloor$, following the theory given in Theorem 4. For an explicit example of how to choose the lifetime, one can instead use $\hat{\mu}_{\mathrm{d}}(x;\hat{\lambda}_{J-1},J-1)$ so that the leading bias is explicitly characterized by Lemma 4, and with $\hat{\lambda}_{J-1}$ as defined below. This estimator is however not minimax-optimal as the debiasing order of $J-1<J$ does not satisfy the conditions of Theorem 4.

For performing inference, a more careful procedure is required; we suggest the following, when $\beta > 2$ is known. Set $J = \lfloor \underline{\beta}/2 \rfloor$ as before, and use $\hat{\mu}_{\mathrm{d}}(x;\hat{\lambda}_{J-1},J)$ and $\hat{\Sigma}_{\mathrm{d}}(x;\hat{\lambda}_{J-1},J)$ to construct a confidence interval (11), so that one selects a lifetime tailored for a more biased estimator than that which is actually used. This results in an inflated lifetime estimate, guaranteeing the resulting bias is negligible when it is plugged into the fully debiased estimator. This approach to tuning parameter selection and debiasing for nonparametric inference corresponds to an application of robust bias correction (Calonico et al., 2018, 2022), where the point estimator is bias-corrected and the robust standard error estimator incorporates the additional variability introduced by the correction. This gives a refined distributional approximation but may not exhaust the underlying smoothness of the regression function. An alternative approach based on Lepski's method (Birgé, 2001; Lepski, 1992) could be developed with the latter goal in mind, when β is unknown.

It remains to propose a concrete method for computing $\hat{\lambda}_J$ in finite samples; we suggest a procedure based on minimizing the asymptotic mean squared error (AMSE) using plug-in selection with polynomial estimation, building on classical ideas from the nonparametric smoothing literature. Expressions for the AMSE are available as direct consequences of Lemmas 4 and 5, provided that $J < \lfloor \underline{\beta}/2 \rfloor$ so the Hölder smoothness is not fully exhausted.

Selecting the lifetime parameter λ with polynomial estimation

For implementation, we propose a simple rule-of-thumb approach. Suppose that $X_i \sim \text{Unif}([0,1]^d)$ and that the leading bias of $\hat{\mu}_{\text{d}}(x)$ is well approximated by an additively separable function so that, writing $\partial_j^{2J+2}\mu(x)$ for $\partial_j^{2J+2}\mu(x)/\partial x_j^{2J+2}$, the asymptotic bias is

$$\text{ABias}(x;\lambda,J) = \frac{\bar{\omega}B_{J+1}(x)}{\lambda^{2J+2}} = \frac{1}{\lambda^{2J+2}}\frac{\bar{\omega}}{J+2}\sum_{j=1}^{d}\partial_j^{2J+2}\mu(x).$$

Suppose that the model is homoscedastic so $\sigma^2(x)=\sigma^2$ and the asymptotic variance of $\hat{\mu}_{\text{d}}$ is

$$\text{AVar}(x;\lambda,J) = \frac{\lambda^d}{n}\Sigma_{\text{d}}(x) = \frac{\lambda^d\sigma^2}{n}\sum_{r=0}^{J}\sum_{r'=0}^{J}\omega_r\omega_{r'}(\ell_{rr'}+\ell_{r'r})^d.$$

The asymptotic mean squared error is therefore

$$\text{AMSE}(x;\lambda,J) = \frac{1}{\lambda^{4J+4}}\frac{\bar{\omega}^2}{(J+2)^2}\left(\sum_{j=1}^{d}\partial_j^{2J+2}\mu(x)\right)^2 + \frac{\lambda^d\sigma^2}{n}\sum_{r=0}^{J}\sum_{r'=0}^{J}\omega_r\omega_{r'}(\ell_{rr'}+\ell_{r'r})^d.$$

Minimizing over $\lambda > 0$ yields the AMSE-optimal lifetime parameter

$$\lambda_J = \left(\frac{\frac{(4J+4)\bar{\omega}^2}{(J+2)^2}n\left(\sum_{j=1}^{d}\partial_j^{2J+2}\mu(x)\right)^2}{d\sigma^2\sum_{r=0}^{J}\sum_{r'=0}^{J}\omega_r\omega_{r'}(\ell_{rr'}+\ell_{r'r})^d}\right)^{\frac{1}{4J+4+d}}.$$

An estimator of λ_J is given by the plug-in procedure

$$\hat{\lambda}_J = \left(\frac{\frac{(4J+4)\bar{\omega}^2}{(J+2)^2}n\left(\sum_{j=1}^{d}\partial_j^{2J+2}\hat{\mu}(x)\right)^2}{d\hat{\sigma}^2\sum_{r=0}^{J}\sum_{r'=0}^{J}\omega_r\omega_{r'}(\ell_{rr'}+\ell_{r'r})^d}\right)^{\frac{1}{4J+4+d}}$$

for some preliminary estimators $\partial_j^{2J+2}\hat{\mu}(x)$ and $\hat{\sigma}^2$. These can be obtained by fitting a global polynomial regression to the data $(\mathbf{X},\mathbf{Y})$ of order $2J+4$ without interaction terms. To do this, define the $n\times((2J+4)d+1)$ design matrix $\mathbf{P}$ with rows given by

$$P(X_i) = \left(1, X_{i1}, X_{i1}^2, \ldots, X_{i1}^{2J+4}, X_{i2}, X_{i2}^2, \ldots, X_{i2}^{2J+4}, \ldots, X_{id}, X_{id}^2, \ldots, X_{id}^{2J+4}\right).$$

Then the derivative estimator is

$$\begin{aligned}\partial_j^{2J+2}\hat{\mu}(x) &= \partial_j^{2J+2}P(x)\left(\mathbf{P}^T\mathbf{P}\right)^{-1}\mathbf{P}^T\mathbf{Y}\\ &= (2J+2)!\left(0_{1+(j-1)(2J+4)+(2J+1)}, 1, x_j, x_j^2/2, 0_{(d-j)(2J+4)}\right)\left(\mathbf{P}^T\mathbf{P}\right)^{-1}\mathbf{P}^T\mathbf{Y},\end{aligned}$$

and the variance estimator $\hat{\sigma}^2$ is the based on the residual sum of squared errors of this model:

$$\hat{\sigma}^2 = \frac{1}{n-(2J+4)d-1}\left(\mathbf{Y}^T\mathbf{Y} - \mathbf{Y}^T\mathbf{P}\left(\mathbf{P}^T\mathbf{P}\right)^{-1}\mathbf{P}^T\mathbf{Y}\right).$$

Choosing the number B of trees in each forest

The next parameter to choose is the number of trees in each forest. If no debiasing is applied, we suggest taking $B \asymp \sqrt{n}$ to ensure the coverage error in Theorem 3 converges to zero. If debiasing is used then we recommend setting $B \asymp n^{\frac{2J-1}{2J}}$, consistent with Theorem 4 and Theorem 6.

Setting the debiasing order J

Deciding how many orders of bias to remove requires knowledge of the Hölder smoothness of μ and f, which is in practice very difficult to estimate statistically. As such we recommend removing only the first one or two bias terms, taking $J \in \{0, 1, 2\}$ to avoid inflating the variance of the estimator.

Selecting the debiasing scalars a_r

As in Section 4, take a fixed geometric or arithmetic sequence. For example, $a_r = (1+\gamma)^r$ or $a_r = 1 + \gamma r$ where $\gamma > 0$; we suggest $a_r = (3/2)^r$.

5.3 Empirical results

To demonstrate the empirical properties of our proposed estimation and inference methodology, we present results with simulated data. Throughout this section we use the data generating process given by uniform covariates $X_i \sim \text{Unif}[0,1]^d$ for $d \in \{1, 2\}$, a sinusoidal regression function $\mu(x) = \sum_{j=1}^{d} \sin(\pi x_j)$, and homoscedastic normal errors $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$ with $\sigma = 3/10$. We thus ensure that the regression function is smooth and appropriately bounded but is not a polynomial, so the bias terms given in Lemma 1 do not vanish. The additive structure is not essential; similar performance is expected for more general smooth models with dependence between the covariates, assuming a sufficiently large sample size and densities bounded away from zero. We focus on estimation at the design point $x = (1/2, \ldots, 1/2) \in \mathbb{R}^d$, and use $n = 1{,}000$ samples and $B = 800$ trees in each forest. We demonstrate our procedures both with and without debiasing by setting $J = 1$ and $J = 0$ respectively, and when $J = 1$ we use the debiasing scalars $(a_1, a_2) = (1, 3/2)$ suggested in Section 4, yielding debiasing coefficients of $(\omega_0, \omega_1) = (-4/5, 9/5)$. For lifetime selection (LS), we first show our estimator $\hat{\lambda}_J$ based on polynomial regression (Section 5.2), and then present the infeasible oracle lifetime λ_J which exactly minimizes the asymptotic mean squared error. To illustrate robustness with respect to this tuning parameter, we repeat the same experiments but rescaling λ_J by a lifetime multiplier $\text{LM} \in \{1 \pm \ell/5 : 0 \le \ell \le 2\}$. We further exhibit the robust bias correction (BC) approach discussed in Section 5.2 by using a debiased estimator ($J = 1$) with the AMSE-optimal lifetime parameter λ_0.

For each such estimator, we present the empirical root mean squared error (RMSE), bias, standard deviation (SD), and absolute bias/SD ratio, based on 3,000 repeats. We also show the estimated standard deviation $\widehat{\text{SD}} = \sqrt{\lambda^d/n}\hat{\Sigma}_{\text{d}}(x)$, as well as the estimated variance of the errors $\hat{\sigma}^2(x)$. Since oracle properties available, we give the asymptotic bias (ABias) and asymptotic standard deviation (ASD). Finally, we present the empirical coverage rate (CR) of nominal 95% confidence intervals along with their empirical average widths (CIW).

We now derive the asymptotic oracle properties of our estimators. Firstly, by Lemma 2, the asymptotic variance of the estimator without debiasing is

$$\text{AVar} = \frac{\lambda^d}{n}\Sigma(x) = \frac{\lambda^d}{n}\frac{\sigma^2(x)}{f(x)}\left(\frac{4 - 4\log 2}{3}\right)^d \approx \frac{\lambda^d\sigma^2}{n}0.4091^d.$$

By Lemma 5, $\ell_{00} = \frac{2}{3}(1 - \log 2)$, $\ell_{01} = \frac{2}{3}(1 - \frac{2}{3}\log\frac{5}{2})$, $\ell_{10} = 1 - \frac{3}{2}\log\frac{5}{3}$, and $\ell_{11} = 1 - \log 2$, so the asymptotic variance of its debiased counterpart is

$$\begin{aligned}\text{AVar}_{\text{d}} &= \frac{\lambda^d}{n}\Sigma_{\text{d}}(x) = \frac{\lambda^d}{n}\frac{\sigma^2(x)}{f(x)}\sum_{r=0}^{J}\sum_{r'=0}^{J}\omega_r\omega_{r'}(\ell_{rr'} + \ell_{r'r})^d \\ &\approx \frac{\lambda^d\sigma^2}{n}\left(0.64 \cdot 0.4091^d - 2.88 \cdot 0.4932^d + 3.24 \cdot 0.6137^d\right).\end{aligned}$$

We similarly establish the asymptotic biases. Without debiasing, by Lemma 1,

$$\mathrm{ABias} = \frac{1}{\lambda^2}\sum_{|v|=1}\partial^{2v}\mu(x)\prod_{j=1}^{d}\frac{1}{v_j+1} = \frac{1}{2\lambda^2}\sum_{j=1}^{d}\frac{\partial^2\mu(x)}{\partial x_j^2} = -\frac{\pi^2}{2\lambda^2}\sum_{j=1}^{d}\sin(\pi x_j) = -\frac{\pi^2 d}{2\lambda^2}.$$

For the debiased estimator, with $\bar{\omega} = \omega_0 + \omega_1 a_1^{-4} = -4/9$, we recover

$$\mathrm{ABias}_{\mathrm{d}} = \frac{\bar{\omega}B_2(x)}{\lambda^4} = -\frac{4}{9\lambda^4}\sum_{|v|=2}\partial^{2v}\mu(x)\prod_{j=1}^{d}\frac{1}{v_j+1} = -\frac{4}{27\lambda^4}\sum_{j=1}^{d}\frac{\partial^4\mu(x)}{\partial^4 x_j} = -\frac{4\pi^4 d}{27\lambda^4}.$$

Table 1 gives results in the one-dimensional setting ($d = 1$). Firstly, observe that the polynomial lifetime estimator appears to be moderately accurate, displaying some oversmoothing when fitting a polynomial of order 4 (for $\hat{\lambda}_0$) and some undersmoothing with a polynomial of order 6 (with $\hat{\lambda}_1$). The effects of debiasing on RMSE are clear, with the appropriately tuned debiased Mondrian forest ($J = 1, \lambda_1$) providing the best results (Theorem 4). Likewise, the effect of debiasing is apparent when using an undersmoothed lifetime ($J = 1, \lambda_0$), with the bias being significantly reduced (see Lemma 4) at the expense of a larger standard deviation. The variance estimator performs well, with $\widehat{\mathrm{SD}}$ a good approximation for the finite-sample SD, and $\hat{\sigma}^2$ similarly sits close to $\sigma^2 = 0.09$. The value of robust bias correction ($J = 1, \lambda_0$) for statistical inference is clear, with the coverage rates clustering around the nominal 95% even with perturbed lifetime values (see Theorem 6). In contrast, the no-debiasing estimator ($J = 0, \lambda_0$) fails to attain correct coverage, while its fully debiased counterpart ($J = 1, \lambda_1$) lacks robustness, reaching the nominal level only with larger lifetime values. Accurate coverage is at the expense of wider confidence intervals, but the differences are not large.

Table 2 presents analogous results in the two-dimensional setting ($d = 2$). The debiased estimator ($J = 1, \lambda_1$) again achieves the best RMSE, and the undersmoothed estimator ($J = 1, \lambda_0$) similarly displays the smallest bias/SD ratio. Coverage rates are generally worse than in Table 1, mostly due to the increased difficulty posed by the curse of dimensionality and a reduced effective sample size. Nonetheless, inference based on robust bias correction continues to exhibit a pronounced improvement in coverage when compared to standard nondebiased methods, and again shows a moderate increase in confidence interval widths.

6. Computational complexity and application to online learning

We discuss some computational aspects of (debiased) Mondrian random forests. We firstly consider the batch setting, where all of the data is available simultaneously, and secondly investigate the online regime, where data arrives sequentially and the model must be incrementally updated (Lakshminarayanan et al., 2014). Mondrian random forests have several properties that make them well suited for online learning: (i) in Mourtada et al. (2021) it was shown that some online Mondrian forest variants maintain statistical consistency, achieving the same asymptotic error rates as their batch counterparts under certain conditions; (ii) as we will demonstrate (Lemma 8), online Mondrian forest algorithms exploiting the Markov property of the Mondrian process are computationally efficient, therefore scaling to large streaming datasets; and (iii) the random nature of splits in Mondrian trees allows the forest to naturally adapt to changes in the underlying data distribution over time (concept drift), without requiring explicit drift detection or model reset mechanisms.

Some potential applications of online Mondrian forests with uncertainty quantification include real-time prediction and monitoring in industrial processes (Gomes et al., 2017), adaptive pricing and recommendation systems (Krauss et al., 2017; Li et al., 2018), online anomaly detection with confidence levels (Martindale et al., 2020), and streaming data analysis for the natural sciences (Abdulsalam et al., 2010).

The inference procedures developed in this paper extend to the online setting, allowing for uncertainty quantification in streaming data applications. However, care must be taken in situations where the underlying distribution may change over time, or where validity of the inferential procedures is required to hold uniformly over the data arrival times. Developing rigorous statistical

Table 1. Simulation results with $d = 1$, $n = 1{,}000$, and $B = 800$, over 3,000 repeats

	J	LS	LM	λ	RMSE	Bias	SD	Bias/SD	$\widehat{\mathrm{SD}}$	$\hat{\sigma}^2$	ARMSE	ABias	ASD	CR	CIW
No debiasing	0	$\hat{\lambda}_0$	1.0	14.73	0.0351	−0.0250	0.0247	1.0123	0.0236	0.0931	0.0369	−0.0270	0.0232	82.5%	0.093
		λ_0	1.2	23.10	0.0307	−0.0092	0.0293	0.3130	0.0292	0.0894	0.0306	−0.0092	0.0292	93.6%	0.114
			1.1	21.18	0.0300	−0.0109	0.0280	0.3888	0.0280	0.0897	0.0300	−0.0110	0.0279	93.4%	0.110
			1.0	19.25	0.0297	−0.0131	0.0267	0.4909	0.0267	0.0901	0.0298	−0.0133	0.0266	92.9%	0.105
			0.9	17.33	0.0300	−0.0160	0.0253	0.6326	0.0254	0.0907	0.0301	−0.0164	0.0253	90.7%	0.100
			0.8	15.40	0.0312	−0.0201	0.0238	0.8438	0.0241	0.0916	0.0316	−0.0208	0.0238	87.5%	0.095
Debiasing	1	$\hat{\lambda}_1$	1.0	11.14	0.0301	−0.0031	0.0300	0.1036	0.0302	0.1002	0.0296	−0.0026	0.0287	95.0%	0.119
		λ_1	1.2	7.86	0.0255	−0.0070	0.0246	0.2835	0.0269	0.1103	0.0245	−0.0038	0.0242	95.9%	0.106
			1.1	7.21	0.0255	−0.0095	0.0236	0.4031	0.0263	0.1147	0.0238	−0.0053	0.0232	95.2%	0.103
			1.0	6.55	0.0264	−0.0135	0.0227	0.5950	0.0256	0.1198	0.0235	−0.0078	0.0221	94.3%	0.100
			0.9	5.90	0.0288	−0.0191	0.0216	0.8817	0.0249	0.1259	0.0241	−0.0119	0.0210	90.6%	0.097
			0.8	5.24	0.0343	−0.0274	0.0206	1.3346	0.0240	0.1329	0.0275	−0.0191	0.0198	82.0%	0.094
Robust BC	1	$\hat{\lambda}_0$	1.0	14.73	0.0334	−0.0014	0.0333	0.0405	0.0339	0.0940	0.0336	−0.0011	0.0330	95.3%	0.133
		λ_0	1.2	23.10	0.0420	−0.0004	0.0420	0.0105	0.0419	0.0898	0.0415	−0.0001	0.0415	94.8%	0.164
			1.1	21.18	0.0401	−0.0003	0.0401	0.0078	0.0402	0.0901	0.0398	−0.0001	0.0398	95.0%	0.158
			1.0	19.25	0.0381	−0.0004	0.0381	0.0115	0.0383	0.0905	0.0379	−0.0001	0.0379	94.7%	0.150
			0.9	17.33	0.0362	−0.0003	0.0362	0.0084	0.0365	0.0912	0.0360	−0.0002	0.0360	95.0%	0.143
			0.8	15.40	0.0341	−0.0005	0.0341	0.0139	0.0346	0.0922	0.0339	−0.0003	0.0339	95.3%	0.136

Table 2. Simulation results with $d = 2$, $n = 1{,}000$, and $B = 800$, over 3,000 repeats

	J	LS	LM	λ	RMSE	Bias	SD	Bias/SD	$\widehat{\text{SD}}$	$\hat{\sigma}^2$	ARMSE	ABias	ASD	CR	CIW
No debiasing	0	$\hat{\lambda}_0$	1.0	12.35	0.0805	−0.0646	0.0481	1.3432	0.0481	0.0989	0.0828	−0.0666	0.0479	71.1%	0.189
		λ_0	1.2	18.39	0.0758	−0.0310	0.0692	0.4481	0.0627	0.0882	0.0771	−0.0292	0.0714	88.2%	0.246
			1.1	16.85	0.0735	−0.0361	0.0640	0.5650	0.0593	0.0898	0.0741	−0.0347	0.0654	87.0%	0.233
			1.0	15.32	0.0726	−0.0427	0.0587	0.7280	0.0558	0.0919	0.0728	−0.0420	0.0595	84.9%	0.219
			0.9	13.79	0.0743	−0.0518	0.0532	0.9740	0.0520	0.0947	0.0746	−0.0519	0.0535	80.8%	0.204
			0.8	12.26	0.0796	−0.0637	0.0477	1.3346	0.0478	0.0985	0.0811	−0.0657	0.0476	71.6%	0.188
Debiasing	1	$\hat{\lambda}_1$	1.0	9.20	0.0726	−0.0144	0.0712	0.2020	0.0746	0.1277	0.0723	−0.0086	0.0691	95.1%	0.292
		λ_1	1.2	7.18	0.0584	−0.0217	0.0542	0.3999	0.0672	0.1490	0.0550	−0.0108	0.0540	96.3%	0.263
			1.1	6.58	0.0577	−0.0283	0.0503	0.5620	0.0644	0.1602	0.0518	−0.0154	0.0495	95.8%	0.252
			1.0	5.99	0.0596	−0.0381	0.0459	0.8299	0.0613	0.1733	0.0503	−0.0225	0.0450	94.0%	0.240
			0.9	5.39	0.0664	−0.0516	0.0418	1.2332	0.0578	0.1879	0.0530	−0.0343	0.0405	90.4%	0.227
			0.8	4.79	0.0797	−0.0704	0.0373	1.8873	0.0538	0.2044	0.0656	−0.0549	0.0360	79.6%	0.211
Robust BC	1	$\hat{\lambda}_0$	1.0	12.35	0.0889	−0.0053	0.0888	0.0598	0.0854	0.1047	0.0928	−0.0014	0.0927	94.9%	0.335
		λ_0	1.2	18.39	0.1208	−0.0032	0.1208	0.0265	0.0971	0.0925	0.1381	−0.0003	0.1381	89.4%	0.380
			1.1	16.85	0.1135	−0.0040	0.1134	0.0351	0.0953	0.0941	0.1266	−0.0004	0.1266	90.8%	0.373
			1.0	15.32	0.1056	−0.0039	0.1055	0.0367	0.0927	0.0964	0.1151	−0.0005	0.1151	92.6%	0.363
			0.9	13.79	0.0974	−0.0042	0.0973	0.0427	0.0893	0.0994	0.1036	−0.0008	0.1036	93.8%	0.350
			0.8	12.26	0.0883	−0.0047	0.0882	0.0532	0.0853	0.1041	0.0921	−0.0013	0.0921	94.9%	0.334

inference tools for online Mondrian forests in those more complicated time-dependent regimes is an interesting direction for future work.

The core of our computational approaches for batch and online learning comprise several main ideas; these enable substantial improvements over naive algorithms based on the equations presented in previous sections. The first of these is to keep track of which data points are 'local' to the evaluation point x, according to the forest $(T_{br}(x) : 1 \le b \le B, 0 \le r \le J)$. Define the *union cell* $U(x) \subseteq [0, 1]^d$ and *active indices* $I(x) \subseteq \{1, \ldots, n\}$ by

$$U(x) = \prod_{j=1}^{d} \bigcup_{b=1}^{B} \bigcup_{r=0}^{J} T_{br}(x)_j \quad \text{and} \quad I(x) = \{1 \le i \le n : X_i \in U(x)\} \tag{12}$$

respectively, noting that any data point contributing to $\hat{\mu}_{\mathrm{d}}(x)$ or $\hat{\Sigma}_{\mathrm{d}}(x)$ satisfies $X_i \in U(x)$ and $i \in I(x)$. As the lifetime parameter λ grows, the volume of $U(x)$ and the proportion of contributing samples $|I(x)|/n$ both converge to zero in expectation, lowering the effective sample size and significantly decreasing the amount of computation necessary. Further, $U(x)$ can be efficiently computed with a divide-and-conquer approach whenever multiple parallel processors are available.

The second main idea is to observe that the estimators $\hat{\mu}_{\mathrm{d}}(x)$ and $\hat{\sigma}^2(x)$ can be expressed as ratios of sums. More precisely, firstly define

$$\begin{aligned} N_{br}(x) &= \sum_{i \in I(x)} \mathbb{I}\{X_i \in T_{br}(x)\}, \qquad S_{br}(x) = \sum_{i \in I(x)} Y_i \mathbb{I}\{X_i \in T_{br}(x)\}, \\ V_{br}(x) &= \sum_{i \in I(x)} Y_i^2 \mathbb{I}\{X_i \in T_{br}(x)\}, \end{aligned} \tag{13}$$

which are efficient to update as new samples arrive; furthermore, they can be computed separately for each b and r in parallel. Then,

$$\hat{\mu}_{\mathrm{d}}(x) = \sum_{r=0}^{J} \omega_r \frac{1}{B} \sum_{b=1}^{B} \frac{S_{br}(x)}{N_{br}(x)} \quad \text{and} \quad \hat{\sigma}^2(x) = \frac{1}{B} \sum_{b=1}^{B} \frac{V_{b0}(x)}{N_{b0}(x)}. \tag{14}$$

The third observation is that the estimators depend on the trees only through the cell $T_{br}(x)$. Since Mourtada et al. (2020) characterize the exact distribution of this quantity, it can be sampled without needing to grow an entire Mondrian tree. Further, the memoryless property of the exponential distribution (and thus also of the Mondrian process) means that in the online setting, only a small fraction of the cells typically need to be updated.

The fourth and final concept is to avoid fitting the relatively computationally expensive $\hat{\Sigma}_{\mathrm{d}}(x)$ too often. This estimator does not readily admit a 'ratio of sums' formulation, and hence is not efficient to update incrementally. Our recommendation is to instead only update this term after $K \ge 1$ new data points have arrived on average. Note however that using the active indices $I(x)$ still permits an improvement over the naive approach, since

$$\hat{\Sigma}_{\mathrm{d}}(x) = \hat{\sigma}^2(x) \frac{n}{\lambda^d} \sum_{i \in I(x)} \left(\sum_{r=0}^{J} \omega_r \frac{1}{B} \sum_{b=1}^{B} \frac{\mathbb{I}\{X_i \in T_{br}(x)\}}{N_{br}(x)} \right)^2. \tag{15}$$

Before discussing the online learning setting in more detail, we present our efficient procedure for batch estimation and inference in Algorithm 1.

The following lemma bounds the average case time complexity of our batch learning procedure (Algorithm 1), under the same assumptions made throughout the paper.

Lemma 7 (Computational complexity of batch learning). Suppose Assumptions 1 and 2 hold. Then the average case time complexity of Algorithm 1 is

$$\mathbb{E}[\mathcal{T}_{\mathrm{b}}] \lesssim d(J+1)\big(nd(J+1) + B\big) + \frac{nBd(J+1)\log(2B(J+1))^d}{\lambda^d}.$$

We now turn to the online learning setting, making the following assumptions. Firstly, suppose that a (debiased) Mondrian random forest with B trees has already been fitted to a data set with n samples, using a lifetime of λ, and that this original data set is still available. Assume that the union cell $U(x)$, the index set $I(x)$, and the point estimates $\hat{\mu}_{\mathrm{d}}(x)$ and $\hat{\Sigma}_{\mathrm{d}}(x)$ have been computed, as well as the trees $T_{br}(x)$ and the quantities $S_{br}(x)$, $N_{br}(x)$ and $V_{br}(x)$ for $1 \le b \le B$ and $0 \le r \le J$. A new data set with k samples then arrives, where $1 \le k \le n$, and we must produce updated estimates $\hat{\mu}^*_{\mathrm{d}}(x)$, $\hat{\Sigma}^*_{\mathrm{d}}(x)$ and $\mathrm{CI}^*_{\mathrm{d}}(x)$ based on all $n+k$ samples. Our randomized procedure for doing this is described below, using a star to indicate updated quantities, and summarized in Algorithm 2.

The new sample size is $n^* = n + k$, so the first step is to update B. As recommended in Section 5, we take $B \asymp n^{\xi}$ for some $\xi \in (0, 1)$; therefore set $B^* = \lfloor (n+k)^{\xi} B / n^{\xi} \rfloor$. Next, we update the lifetime parameter λ. To avoid excessive computation, we suggest the following: with probability $1 \wedge (k/K)$, use the methods from Section 5.2 to compute a new lifetime parameter $\lambda^* \ge \lambda$ using all of the data. Otherwise, note that $\lambda \asymp n^{\zeta}$ for some $\zeta \in (0, 1/d)$ (for example $\zeta = 1/(d+2\beta)$ under the conditions of Theorem 4) and set $\lambda^* = (n+k)^{\zeta}\lambda / n^{\zeta}$. Next, to update the trees $T_{br}(x)$, sample E_{brj1} and E_{brj2} i.i.d. Exp(1), and set

$$T^*_{br}(x)_j^- = T_{br}(x)_j^- \vee \left(x_j - \frac{E_{brj1}}{\lambda^* - \lambda}\right), \quad T^*_{br}(x)_j^+ = T_{br}(x)_j^+ \wedge \left(x_j + \frac{E_{brj2}}{\lambda^* - \lambda}\right). \tag{16}$$

Since $B^* \ge B$, we also generate new trees $T^*_{br}(x)$ for $B+1 \le b \le B^*$ and $0 \le r \le J$ using

$$T^*_{br}(x)_j^- = 0 \vee \left(x_j - \frac{E_{brj1}}{a_r \lambda}\right), \quad T^*_{br}(x)_j^+ = 1 \wedge \left(x_j + \frac{E_{brj2}}{a_r \lambda}\right). \tag{17}$$

Computing $U^*(x)$ is simple, applying (12) to $T^*_{br}(x)$. To update $I(x)$, set

$$I^*(x) = \begin{cases} \{i \in I(x) \cup \{n+1, \ldots, n+k\} : X_i \in U^*(x)\} & \text{if } U^*(x) \subseteq U(x), \\ \{1 \le i \le n+k : X_i \in U^*(x)\} & \text{otherwise.} \end{cases} \tag{18}$$

For $N_{br}(x)$, and analogously for $S_{br}(x)$ and $V_{br}(x)$, apply the following method:

$$N^*_{br}(x) = \begin{cases} N_{br}(x) + \displaystyle\sum_{i \in V^*(x), i > n} \mathbb{I}\{X_i \in T_b(x)\} & \text{if } b \le B \text{ and } T_{br}(x) = T^*_{br}(x) \\ \displaystyle\sum_{i \in I^*(x)} \mathbb{I}\{X_i \in T^*_{br}(x)\} & \text{otherwise.} \end{cases} \tag{19}$$

Finally, $\hat{\mu}^*_{\mathrm{d}}(x)$ and $\hat{\sigma}^{2*}(x)$ are computed using (14). With probability $1 \wedge (k/K)$, recalculate $\hat{\Sigma}^*_{\mathrm{d}}(x)$ with (15); otherwise set $\hat{\Sigma}^*_{\mathrm{d}}(x) = \hat{\Sigma}_{\mathrm{d}}(x)$. The confidence interval $\mathrm{CI}^*_{\mathrm{d}}(x)$ can then be constructed with (11). The following algorithm summarizes our online methodology.

Lemma 8 bounds the average case time complexity of our online computational procedure presented in Algorithm 2.

Lemma 8 (Computational complexity of online learning). Suppose Assumptions 1 and 2 hold. Then the average case time complexity of Algorithm 2 is

$$\mathbb{E}[\mathcal{T}_{\mathrm{o}}] \lesssim d(J+1)\left(\frac{knd(J+1)}{K} + kd + B\right) + \frac{d(J+1)\log(2B(J+1))^d}{\lambda^d}\left(n + Bk + \frac{nB}{K}\right).$$

Lemma 7 already demonstrated that Algorithm 1 is more efficient than the naive approach of computing $\hat{\mu}_{\mathrm{d}}(x)$ and $\hat{\Sigma}_{\mathrm{d}}(x)$ directly with (8) and (10), respectively, which each have a time

complexity of $n(J+1)B$. The reason for this is that by first constructing the active indices $I(x)$, we avoid iterating over the entire sample for each tree in Algorithm 1. Lemma 8 formalizes the improvement achieved by Algorithm 2 in online settings, relative to the batch estimation approach of Algorithm 1. Most importantly, the terms involving the product nB are reduced to $n + Bk + nB/K$, offering a substantial speed-up in large forests when the new sample size k is much smaller than that of the existing data n, and when K is large to avoid regularly estimating the lifetime λ and variance $\hat{\Sigma}_{\mathrm{d}}(x)$.

7. Conclusion

We presented a Berry–Esseen theorem under mild conditions for the Mondrian random forest estimator, and showed how it can be used to perform statistical inference on an unknown non-parametric regression function. We introduced debiased versions of Mondrian random forests, exploiting higher-order smoothness, and demonstrated their advantages for statistical inference and their minimax optimality properties. We discussed tuning parameter selection, enabling fully feasible and practical estimation and inference procedures, and demonstrated the empirical performance of our proposed methodology. Finally, we developed efficient algorithms for batch and online settings.

There are several potential avenues for future work on inference with Mondrian random forests. The development of data-adaptive partitioning schemes is one such important direction, and could be implemented perhaps by allowing the lifetime parameter λ to vary across different covariates, yielding the d-dimensional parameter $(\lambda_1, \ldots, \lambda_d)$. One approach to designing such methodology might involve adapting sparse, greedy algorithms for non-parametric regression, similar to those described by Lafferty and Wasserman (2008), to the context of axis-aligned partitioning estimators. Specifically, by examining how changes in each λ_j affect the Mondrian forest estimator, e.g. via an estimate of $\frac{\partial}{\partial \lambda_j}\mathbb{E}[\hat{\mu}(x)]$, these parameters can be dynamically adjusted to more effectively learn low-dimensional structure in the regression function. Alternatively, one might formulate a Goldenshluger–Lepski-type procedure (Goldenshluger & Lepski, 2008) for multiple tuning parameter selection. Another potential line of research would consist of proposing further strategies for debiasing Mondrian random forests (and related estimators); an approach based on within-cell local polynomial smoothing, for example, may serve to eliminate both design bias and boundary bias, as well as allowing for less restrictive conditions on the covariate density function and the regression function.

Acknowledgments

The authors would like to thank the Editor, two anonymous reviewers, Rajita Chandak, Jianqing Fan, Henry Reeve, Richard Samworth, and Sanjeev Kulkarni for insightful comments and suggestions. Klusowski would like to thank the participants at the 11th Bernoulli-IMS World Congress in Probability and Statistics in Bochum, Germany for valuable feedback. Underwood would like to thank the attendees at the University of Pittsburgh, University of Illinois Urbana-Champaign, and University of Michigan Statistics Seminars, as well as the members of the University of Cambridge Statistical Laboratory.

Conflicts of interest: None declared.

Funding

M.D.C. was supported in part by the National Science Foundation through DMS-2210561 and SES-2241575. J.M.K. was supported in part by the National Science Foundation through CAREER DMS-2239448, DMS-2054808, and HDR TRIPODS CCF-1934924.

Supplementary material

Supplementary material is available online at *Journal of the Royal Statistical Society: Series B*.

References

Abdulsalam H., Skillicorn D. B., & Martin P. (2010). Classification using streaming random forests. *IEEE Transactions on Knowledge and Data Engineering*, *23*(1), 22–36. https://doi.org/10.1109/TKDE.2010.36

Arnould L., Boyer C., & Scornet E. (2023). Is interpolation benign for random forest regression? In *International Conference on Artificial Intelligence and Statistics* (pp. 5493–5548). PMLR.

Baptista R., O'Reilly E., & Xie Y. (2024). 'TrIM: Transformed iterative Mondrian forests for gradient-based dimension reduction and high-dimensional regression', arXiv, arXiv:2407.09964v1, preprint: not peer reviewed.

Biau G. (2012). Analysis of a random forests model. *Journal of Machine Learning Research*, *13*, 1063–1095.

Birgé L. (2001). An alternative point of view on Lepski's method. *Lecture Notes–Monograph Series*, *36*, 113–133. https://doi.org/10.1214/lnms/1215090065

Boucheron S., Lugosi G., & Massart P. (2013). *Concentration inequalities: A nonasymptotic theory of independence*. Oxford University Press.

Breiman L. (2001). Random forests. *Machine Learning*, *45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

Breiman L., Friedman J., Olshen R., & Stone C. J. (1984). *Classification and regression trees*. Chapman and Hall/CRC.

Calonico S., Cattaneo M. D., & Farrell M. H. (2018). On the effect of bias estimation on coverage accuracy in nonparametric inference. *Journal of the American Statistical Association*, *113*(522), 767–779. https://doi.org/10.1080/01621459.2017.1285776

Calonico S., Cattaneo M. D., & Farrell M. H. (2022). Coverage error optimal confidence intervals for local polynomial regression. *Bernoulli*, *28*(4), 2998–3022. https://doi.org/10.3150/21-BEJ1445

Caruana R., Niculescu-Mizil A., Crew G., & Ksikes A. (2004). Ensemble selection from libraries of models. In *Proceedings of the Twenty-First International Conference on Machine Learning* (p. 18). ACM.

Cattaneo M. D., Chandak R., & Klusowski J. M. (2024). Convergence rates of oblique regression trees for flexible function libraries. *Annals of Statistics*, *52*(2), 466–490. https://doi.org/10.1214/24-AOS2354

Cattaneo M. D., Farrell M. H., & Feng Y. (2020). Large sample properties of partitioning-based series estimators. *Annals of Statistics*, *48*(3), 1718–1741. https://doi.org/10.1214/19-AOS1865

Cattaneo M. D., Klusowski J. M., & Tian P. M. (2024). 'On the pointwise behavior of recursive partitioning and its implications for heterogeneous causal effect estimation', arXiv, arXiv:2211.10805v3, preprint: not peer reviewed.

Cattaneo M. D., Klusowski J. M., & Underwood W. G. (2025). Supplement to "Inference with Mondrian random forests". *Submitted*.

Cattaneo M. D., Klusowski J. M., & Yu R. R. (2025). 'The honest truth about causal trees: Accuracy limits for heterogeneous treatment effect estimation', arXiv, arXiv:2509.11381v1, preprint: not peer reviewed.

Chi C.-M., Vossler P., Fan Y., & Lv J. (2022). Asymptotic properties of high-dimensional random forests. *Annals of Statistics*, *50*(6), 3415–3438. https://doi.org/10.1214/22-AOS2234

Duroux R., & Scornet E. (2018). Impact of subsampling and tree depth on random forests. *ESAIM: Probability and Statistics*, *22*, 96–128. https://doi.org/10.1051/ps/2018008

Gao W., Xu F., & Zhou Z.-H. (2022). Towards convergence rate analysis of random forests for classification. *Artificial Intelligence*, *313*(18), Article 103788. https://doi.org/10.1016/j.artint.2022.103788

Goldenshluger A., & Lepski O. (2008). Universal pointwise selection rule in multivariate function estimation. *Bernoulli*, *14*(4), 1150–1190. https://doi.org/10.3150/08-BEJ144

Gomes H. M., Bifet A., Read J., Barddal J. P., Enembreck F., Pfharinger B., Holmes G., & Abdessalem T. (2017). Adaptive random forests for evolving data stream classification. *Machine Learning*, *106*(9–10), 1469–1495. https://doi.org/10.1007/s10994-017-5642-8

Huang J. Z. (2003). Local asymptotics for polynomial spline regression. *Annals of Statistics*, *31*(5), 1600–1635. https://doi.org/10.1214/aos/1065705120

Ibragimov I. A., & Linnik Y. V. (1975). *Independent and stationary sequences of random variables*. Wolters–Noordhoff.

Klusowski J. M. (2021). Sharp analysis of a simple model for random forests. In *International Conference on Artificial Intelligence and Statistics* (pp. 757–765). PMLR.

Klusowski J. M., & Tian P. (2024). Large scale prediction with decision trees. *Journal of the American Statistical Association*, *119*(545), 525–537. https://doi.org/10.1080/01621459.2022.2126782

Krauss C., Do X. A., & Huck N. (2017). Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500. *European Journal of Operational Research*, *259*(2), 689–702. https://doi.org/10.1016/j.ejor.2016.10.031

Lafferty J., & Wasserman L. (2008). Rodeo: Sparse, greedy nonparametric regression. *Annals of Statistics*, *36*(1), 28–63. https://doi.org/10.1214/009053607000000811

Lakshminarayanan B., Roy D. M., & Teh Y. W. (2014). Mondrian Forests: Efficient Online Random Forests. In *Advances in Neural Information Processing Systems* (Vol. *27*). Curran Associates.

Lepski O. V. (1992). Asymptotically minimax adaptive estimation I: Upper bounds. Optimally adaptive estimates. *Theory of Probability & its Applications*, *36*(4), 682–697. https://doi.org/10.1137/1136085
Li X., Wang Z., Wang L., Hu R., & Zhu Q. (2018). A multi-dimensional context-aware recommendation approach based on improved random forest algorithm. *IEEE Access: Practical Innovations, Open Solutions*, *6*, 45071–45085. https://doi.org/10.1109/Access.6287639
Ma H., Ghojogh B., Samad M. N., Zheng D., & Crowley M. (2020). Isolation Mondrian forest for batch and online anomaly detection. In *2020 IEEE International Conference on Systems, Man, and Cybernetics* (pp. 3051–3058). IEEE.
Martindale N., Ismail M., & Talbert D. A. (2020). Ensemble-based online machine learning algorithms for network intrusion detection systems using streaming data. *Information*, *11*(6), 315. https://doi.org/10.3390/info11060315
Miftachov R., & Reiß M. (2025). 'Early stopping for regression trees', arXiv, arXiv:2502.04709, preprint: not peer reviewed.
Mourtada J., Gaïffas S., & Scornet E. (2020). Minimax optimal rates for Mondrian trees and forests. *Annals of Statistics*, *48*(4), 2253–2276. https://doi.org/10.1214/19-AOS1886
Mourtada J., Gaïffas S., & Scornet E. (2021). AMF: Aggregated Mondrian forests for online learning. *Journal of the Royal Statistical Society: Series B, Statistical Methodology*, *83*(3), 505–533. https://doi.org/10.1111/rssb.12425
Näf J., Emmenegger C., Bühlmann P., & Meinshausen N. (2023). Confidence and uncertainty assessment for distributional random forests. *Journal of Machine Learning Research*, *24*(366), 1–77.
O'Reilly E. (2024). 'Statistical advantages of oblique randomized decision trees and forests', arXiv, arXiv:2407.02458v2, preprint: not peer reviewed.
O'Reilly E., & Tran N. M. (2022). Stochastic geometry to generalize the Mondrian process. *SIAM Journal on Mathematics of Data Science*, *4*(2), 531–552. https://doi.org/10.1137/20M1354490
O'Reilly E., & Tran N. M. (2024). Minimax rates for high-dimensional random tessellation forests. *Journal of Machine Learning Research*, *25*(89), 1–32.
Osborne C., & O'Reilly E. (2025). 'The uniformly rotated Mondrian kernel', arXiv, arXiv:2502.04323v2, preprint: not peer reviewed.
Petrov V. V. (1995). *Limit theorems of probability theory: Sequences of independent random variables*, Volume 4 of Oxford Studies in Probability. Oxford University Press.
Roy D. M., & Teh Y. W. (2008). The Mondrian Process. In *Advances in Neural Information Processing Systems* (Vol. *21)*. Curran Associates.
Schucany W. R., & Sommers J. P. (1977). Improvement of kernel type density estimators. *Journal of the American Statistical Association*, *72*(358), 420–423. https://doi.org/10.1080/01621459.1977.10481012
Scillitoe A., Seshadri P., & Girolami M. (2021). Uncertainty quantification for data-driven turbulence modelling with Mondrian forests. *Journal of Computational Physics*, *430*(4), Article 110116. https://doi.org/10.1016/j.jcp.2021.110116
Scornet E., Biau G., & Vert J.-P. (2015). Consistency of random forests. *Annals of Statistics*, *43*(4), 1716–1741. https://doi.org/10.1214/15-AOS1321
Stone C. J. (1982). Optimal global rates of convergence for nonparametric regression. *Annals of Statistics*, *10*(4), 1040–1053. https://doi.org/10.1214/aos/1176345969
Vicuna M., Khannouz M., Kiar G., Chatelain Y., & Glatard T. (2021). Reducing numerical precision preserves classification accuracy in Mondrian forests. In *2021 IEEE International Conference on Big Data* (pp. 2785–2790). IEEE.
Wager S., & Athey S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, *113*(523), 1228–1242. https://doi.org/10.1080/01621459.2017.1319839
Zhan H., Wang J., & Xia Y. (2024). 'Non-asymptotic properties of generalized Mondrian forests in statistical learning', arXiv, arXiv:2406.00660v3, preprint: not peer reviewed.
Zhou Z.-H., & Feng J. (2019). Deep forest. *National Science Review*, *6*(1), 74–86. https://doi.org/10.1093/nsr/nwy108