

WHAT MAKES FOREST-BASED HETEROGENEOUS TREATMENT EFFECT ESTIMATORS WORK?

BY SUSANNE DANDL^{1,a} , CHRISTIAN HASLINGER^{2,b} , TORSTEN HOTHORN^{3,c} ,
HEIDI SEIBOLD^{4,d} , ERIK SVERDRUP^{5,e} , STEFAN WAGER^{5,f}  AND
ACHIM ZEILEIS^{6,g} 

¹Institut für Statistik, Ludwig-Maximilians-Universität München, ^aSusanne.Dandl@stat.uni-muenchen.de

²Klinik für Geburtshilfe, Universitätsspital und Universität Zürich, ^bChristian.Haslinger@usz.ch

³Institut für Epidemiologie, Biostatistik und Prävention, Universität Zürich, ^cTorsten.Hothorn@R-project.org

⁴Institute for Globally Distributed Open Research and Education (IGDORE), München, ^dheidi@seibold.co

⁵Stanford Graduate School of Business, Stanford University, ^eerikcs@stanford.edu, ^fswager@stanford.edu

⁶Faculty of Economics and Statistics, Universität Innsbruck, ^gAchim.Zeileis@R-project.org

Estimation of heterogeneous treatment effects (HTE) is of prime importance in many disciplines, from personalized medicine to economics among many others. Random forests have been shown to be a flexible and powerful approach to HTE estimation in both randomized trials and observational studies. In particular “causal forests” introduced by Athey, Tibshirani and Wager (*Ann. Statist.* **47** (2019) 1148–1178), along with the R implementation in package *grf* were rapidly adopted. A related approach, called “model-based forests” that is geared toward randomized trials and simultaneously captures effects of both prognostic and predictive variables, was introduced by Seibold, Zeileis and Hothorn (*Stat. Methods Med. Res.* **27** (2018) 3104–3125) along with a modular implementation in the R package *model4you*.

Neither procedure is directly applicable to the estimation of individualized predictions of excess postpartum blood loss caused by a cesarean section in comparison to vaginal delivery. Clearly, randomization is hardly possible in this setup, and thus model-based forests lack clinical trial data to address this question. On the other hand, the skewed and interval-censored postpartum blood loss observations violate assumptions made by causal forests. Here we present a tailored model-based forest for skewed and interval-censored data to infer possible predictive prepartum characteristics and their impact on excess postpartum blood loss caused by a cesarean section.

As a methodological basis, we propose a unifying view on causal and model-based forests that goes beyond the *theoretical* motivations and investigates which *computational* elements make causal forests so successful and how these can be blended with the strengths of model-based forests. To do so, we show that both methods can be understood in terms of the same parameters and model assumptions for an additive model under L_2 loss. This theoretical insight allows us to implement several flavors of “model-based causal forests” and dissect their different elements *in silico*.

The original causal forests and model-based forests are compared with the new blended versions in a benchmark study exploring both randomized trials and observational settings. In the randomized setting, both approaches performed akin. If confounding was present in the data-generating process, we found local centering of the treatment indicator with the corresponding propensities to be the main driver for good performance. Local centering of the outcome was less important and might be replaced or enhanced by simultaneous split selection with respect to both prognostic and predictive effects. This lays the foundation for future research combining random forests for HTE estimation with other types of models.

Received September 2022; revised June 2023.

Key words and phrases. Causal forests, heterogeneous treatment effects, observational data, personalized medicine, postpartum hemorrhage, random forest.

1. Introduction.

1.1. *Challenges in treatment effect estimation for cesarean sections.* Cesarean section is the most frequent surgical procedure performed in young and healthy women, currently with one out of three babies in the U.S.A. being born that way (Antoine and Young (2021)). Short-term postpartum benefits and the perceived safety of the procedure explain the increase in popularity over the last 50 years, including the rise of electively performed cesarean sections. At the same time, maternal mortality and morbidity increased globally (WHO (2012), Say et al. (2014)). More recently, adverse long-term effects, including gynecological and obstetrical complications in mothers as well as potential and controversially discussed immune disorders in their children, have gained attention (Antoine and Young (2021)). Lack of clinical trial data directly comparing outcomes of natural births with those following cesarean sections render characterization and quantification of such effects challenging. Postpartum hemorrhage (PPH), defined as blood loss ≥ 500 mL within 24 hours after delivery by the WHO (2012), is a short-term complication associated with maternal morbidity and mortality worldwide. The prevalence of PPH is increasing in industrialized countries (for the U.S.A., see MacDorman et al. (2016)).

Management of PPH requires identification of at risk parturients, and calls went out to the statistics, machine learning, and artificial intelligence communities to develop and evaluate prognostic models (Ende (2022)). Typically, models for dichotomized PPH prognosis were created aiming at either women giving birth by vaginal delivery (Erickson and Carlson (2020), Akazawa et al. (2021)) or at women scheduled for a cesarean section (Kawakita et al. (2019)). Models trained on data from both modes of delivery are rare, for example, in Venkatesh et al. (2020) the mode of delivery was not taken into account as risk factor. Because of the often elective nature of the decision to undergo cesarean section, a quantification of the *additional* amount of hemorrhaging caused by surgery is relevant for the decision process; however, such information is hard to extract from stratified prognostic models. This is true even more considering the possibility of unplanned cesarean deliveries following attempted vaginal deliveries. From a statistical perspective, estimation of a heterogeneous cesarean section effect is nontrivial for a number of reasons. First, potential risk factors for PPH, such as age of the mother, estimated birth weight, gestational age, previous PPH, suspected placental disorders, or multifetal pregnancy might have an impact on both the decision to undergo a cesarean section (treatment) and postpartum blood loss (outcome). Randomization of mode of delivery is impossible, and thus effects have to be estimated from observational data. Second, it is hard to obtain exact measurements of postpartum blood loss in the often hectic environment of a delivery ward, and thus imprecise assessments via interval-censored observations are only available. Third, one has to expect a high level of skewness and extreme values in blood loss measurements, rendering strong distributional assumptions questionable. Lastly, the association of prognostic factors and blood loss is expected to be complex, including nonlinear and interaction terms.

1.2. *Heterogeneous treatment effect estimation and random forests.* In the statistical literature, methods for the estimation of such heterogeneous treatment effects (HTEs) from randomized trials or observational studies has been receiving a lot of attention during the past decade, triggered by an increasing demand from personalized medicine and the need for refined methods in causal inference. In particular, different variations of random forests (Breiman (2001)) have been suggested for HTE estimation and seem promising candidates for addressing the statistical challenges we are facing here. Random forest variants for HTE estimation can be roughly grouped in two classes.

The *first class* of methods employs random forests to estimate the expected outcomes, given covariates separately in the treatment groups. The conditional average treatment effect (CATE) then corresponds to the difference in estimated mean factual and counterfactual outcomes. Notably, the virtual twins method (Foster, Taylor and Ruberg (2011)) has adopted this approach using random forests. Improvements can be obtained by additionally considering treatment-covariate-interactions or fitting separate (synthetic) forests for each treatment group (Foster, Taylor and Ruberg (2011), Dasgupta et al. (2014), Ishwaran and Malley (2014)). Moreover, Lu et al. (2018) proposed a bivariate imputation approach, which uses a bivariate splitting rule (Ishwaran et al. (2008), Tang and Ishwaran (2017)) that simultaneously considers the expected outcome under both treatments. In a more general setup, Kuenzel et al. (2019) introduced X-learners, a class of meta-algorithms which build upon any supervised/regression algorithm, including random forests, Bayesian regression trees (BART, Chipman, George and McCulloch (2010), Hill (2011), Starling et al. (2021)), or neural networks. Most forest methods were initially developed for randomized controlled trials and have later been adapted to be more robust to confounding. For example, the pollinated transformed outcome forests of Powers et al. (2018) build a single forest on propensity score weighted outcomes instead of the original outcomes to account for confounding.

The subject of this paper is the *second class* of random forest-type algorithms aiming at the *direct* estimation of HTEs in a model-driven way. Two such approaches, “causal forests” (Athey, Tibshirani and Wager (2019)) and “model-based forests” (Seibold, Zeileis and Hothorn (2018)), have recently been proposed. “Causal forests” by Athey, Tibshirani and Wager (2019) implement a divide-and-conquer strategy, also referred to as “local centering” or “orthogonalization” for the direct estimation of HTEs from observational data. They first account for the dependence of both the marginal mean of the outcome and the treatment propensity on the available covariates. Subsequently, they exclusively focus on the estimation of the HTEs. In terms of distributional assumptions, causal forests have been developed for continuous outcomes and corresponding conditional means, and the squared error loss plays an important role in the motivation of this algorithm. Cui et al. (2023) also applied causal forests to survival data, and Mayer et al. (2020) discussed strategies to handle missing values. We note that earlier causal tree and forest algorithms, described in Athey and Imbens (2016) and Wager and Athey (2018), do not involve such a local centering step. In this paper we use the term causal forests to describe the algorithm from Athey, Tibshirani and Wager (2019); see also Athey and Wager (2019). Causal forests are implemented in the R package *grf* (Tibshirani et al. (2021)).

“Model-based forests” by Seibold, Zeileis and Hothorn (2018) simultaneously estimate prognostic effects and HTEs. They do so by leveraging model-based recursive partitioning (“MOB”, Zeileis, Hothorn and Hornik (2008)), a technique for learning model trees in which all relevant parameters are reestimated in each subset of a tree. MOB is not a specific model but rather a general framework for model construction where the adaptation to different types of models often still necessitates working out the details of parameter interpretation or model assessment, etc. Seibold, Zeileis and Hothorn (2016) have adapted MOB to model-based trees for HTE, working out the details for Gaussian regression models as well as censored survival models (parametric Weibull model and semiparametric Cox model). Subsequently, Seibold, Zeileis and Hothorn (2018) have extended this work to model-based forests for HTEs, again working out the details of Gaussian regression and censored Weibull survival modeling. Other authors have adapted the general MOB idea to outcome variables on other scales and/or subject to censoring and truncation, for example, as in survival data (Korepanova et al. (2020)), ordinal data (Buri and Hothorn (2020)), generalized mixed models (Fokkema et al. (2018)), or transformation models (Hothorn and Zeileis (2021a)). So far, model-based forests have only been developed for HTE estimation based on randomized trial data.

1.3. *Model-based causal forests for postpartum blood loss.* Neither of the random forest approaches from Section 1.2 is directly applicable to the estimation of heterogeneous cesarean section effects, described in Section 1.1. Our main contribution is, therefore, a novel random forest model that combines the strengths of the existing methods to tackle the challenges in the cesarean section data. We approach this problem by first studying the similarities and differences between causal forests and model-based forests theoretically and empirically. In a second step, we identify the key drivers for good HTE estimation performance in observational data on the one hand and for asymmetric and potentially interval-censored outcomes on the other hand. Lastly, we derive and apply the novel “blended” HTE random forest for PPH by combining the elements identified as being instrumental.

Given that both causal forests and model-based forests encompass additive models under L_2 loss, we adopt this modeling framework to investigate the specific elements that explain both the success of causal forests for observational studies and the flexibility of model-based forests for randomized trials. Specifically, the question of how the disparate strategies for handling the prognostic and confounding effects differ—or how they can be combined—is of both theoretical and practical interest. For obtaining some answers to this question, we employ the modular computational toolbox for tree induction and forest inference in the R package *model4you* (Seibold, Zeileis and Hothorn (2019)), which allows to “mix & match” the elements of both model-based and causal forests.

The results lay the foundation for future research that further expands potential synergies in HTE estimation, using *model-based causal forests*, by blending model-based and causal forests to leverage the strengths of both approaches. To demonstrate this in practice, we investigate the effect of cesarean section on postpartum blood loss in comparison to vaginal deliveries based on a prospective observational study from Switzerland. In this application there is a need for a model-based approach that can deal with the skewed outcome distribution, which is also interval-censored due to the lack of precise measurement techniques. Thus, we showcase a model-based transformation forest applicable to this observational setting. Our contributions here are three-fold: First, we provide a unified understanding of causal forests and model-based forests for HTE estimation in Section 2. Second, based on these theoretical insights, we introduce novel “blended” random forest models in Section 3. Third, we evaluate why causal and model-based forests work in different scenarios and identify the key drivers for good HTE estimation performance in the observational setting (Section 4). Based on the insights gained theoretically and empirically, we pool the key components from causal and model-based forests to derive a model specifically designed for blood loss prediction (Section 5).

2. Models and forest algorithms. In this section we first outline similarities and differences between causal forests and model-based forests theoretically, using the basic setup of regression for real-valued outcomes. Subsequently, three novel blended approaches are introduced that adapt HTE estimation with model-based forests to observational data.

2.1. *The interaction model.* We are interested in the conditional mean of a real-valued outcome $Y \in \mathbb{R}$, given covariates $\mathbf{X} \in \mathcal{X}$ under a specific binary treatment or intervention $W \in \{0, 1\}$, corresponding to control vs. treatment. Under the assumptions that a binomial model $W | \mathbf{X} = \mathbf{x} \sim \text{B}(1, \pi(\mathbf{x}))$ with propensities $\pi(\mathbf{x}) = \mathbb{P}(W = 1 | \mathbf{X} = \mathbf{x}) = \mathbb{E}(W | \mathbf{X} = \mathbf{x})$ describes treatment assignment, and residuals are given by an error term σZ with $\mathbb{E}(Z | \mathbf{X}, W) = 0$ and standard deviation $\sigma > 0$; the model reads

$$(1) \quad Y = \mu(\mathbf{X}) + \tau(\mathbf{X})W + \sigma Z$$

with conditional mean function

$$\mathbb{E}(Y | \mathbf{X} = \mathbf{x}) = \mu(\mathbf{x}) + \tau(\mathbf{x})\pi(\mathbf{x}) =: m(\mathbf{x}).$$

Covariates \mathbf{x} with impact on the prognostic effect $\mu(\mathbf{x})$ are called *prognostic*, while covariates affecting the treatment effect $\tau(\mathbf{x})$ are called *predictive*. Treatment assignment is assumed to be nondeterministic; that is, propensity scores have to be bounded away from zero and one

$$0 < \pi(\mathbf{x}) = \mathbb{P}(W = 1 \mid \mathbf{X} = \mathbf{x}) = \mathbb{E}(W \mid \mathbf{X} = \mathbf{x}) < 1.$$

Personalized medicine and causal inference, in general, focus on the estimation of the heterogeneous treatment effect $\tau(\mathbf{x})$ and thus on the impact of predictive variables on treatment success; accurate estimation of $\tau(\mathbf{x})$ is the main goal of all methods discussed in this paper.

As discussed in Nie and Wager (2021), the interaction model (1) is closely connected to a treatment model with potential outcomes (Imbens and Rubin (2015)), where we posit potential outcomes $Y(0)$ and $Y(1)$, corresponding to the outcome a unit would have experienced without or with treatment, respectively, and assume that we observe $Y = Y(W)$. Then under unconfoundedness (Rosenbaum and Rubin (1983))

$$(Y(0), Y(1)) \perp\!\!\!\perp W \mid \mathbf{X} = \mathbf{x},$$

we can define residuals σZ in (1) such that the interaction model is observationally equivalent to the specification using potential outcomes, and

$$\tau(\mathbf{x}) = \text{CATE}(\mathbf{x}) = \mathbb{E}(Y(1) - Y(0) \mid \mathbf{X} = \mathbf{x})$$

can be interpreted as the conditional average treatment effect. We note that in a uniformly randomized trial, we have $W \perp\!\!\!\perp \{\mathbf{X}, Y(0), Y(1)\}$, so unconfoundedness is always satisfied, and the propensity scores $\pi(\mathbf{x}) \equiv \pi$ are constant by design.

2.2. *Causal forests.* For developing causal forests, Athey, Tibshirani and Wager (2019) rewrite equation (1) as

$$\begin{aligned} (2) \quad (Y \mid \mathbf{X} = \mathbf{x}) &= m(\mathbf{x}) - m(\mathbf{x}) + \mu(\mathbf{x}) + \tau(\mathbf{x})W + \sigma Z \\ &= m(\mathbf{x}) + \tau(\mathbf{x})(W - \pi(\mathbf{x})) + \sigma Z, \end{aligned}$$

which motivates their algorithmic approach of eliminating the marginal mean $m(\mathbf{x}) = \mathbb{E}(Y \mid \mathbf{X} = \mathbf{x})$ and propensities $\pi(\mathbf{x}) = \mathbb{E}(W \mid \mathbf{X} = \mathbf{x})$ first before estimating the heterogeneous treatment effect $\tau(\mathbf{x})$. This orthogonalization (introduced by Robinson (1988)) is also called “local centering” because both outcome $Y - \hat{m}(\mathbf{x})$ and treatment indicator $W - \hat{\pi}(\mathbf{x})$ are centered before $\tau(\mathbf{x})$ is estimated. This approach leads to more robustness to confounding effects in case of observational data because it regresses out the effect of covariates \mathbf{X} on Y and W (Nie and Wager (2021)). While, in principle, any nonparametric regression technique could be applied to estimate $m(\mathbf{x})$ and $\pi(\mathbf{x})$, Athey, Tibshirani and Wager (2019) chose regression forests.

In the second step of causal forests, treatment effects $\tau(\mathbf{x})$ in the model

$$(Y \mid \mathbf{X} = \mathbf{x}, W = w) = \hat{m}(\mathbf{x}) + \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x})) + \sigma Z$$

are then estimated by minimizing the L_2 loss

$$\ell_{\text{cf}}(\tau(\mathbf{x})) := 1/2(Y - \hat{m}(\mathbf{x}) - \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x})))^2$$

w.r.t. τ , the only unknown quantity in this loss function.

Specifically, when splitting a (parent) node, cut-point estimation for causal trees relies first on estimating a constant treatment effect $\hat{\tau}$ in the parent node minimizing $\ell_{\text{cf}}(\tau)$ by solving the score equation

$$(3) \quad s_{\text{cf}}(\tau) = -\frac{\partial \ell_{\text{cf}}(\tau)}{\partial \tau} = (Y - \hat{m}(\mathbf{x}) - \tau(w - \hat{\pi}(\mathbf{x}))) (w - \hat{\pi}(\mathbf{x})) = 0$$

and second on regressing the resulting score

$$s_{cf}(\hat{\tau}) = (Y - \hat{m}(\mathbf{x}) - \hat{\tau}(w - \hat{\pi}(\mathbf{x})))(w - \hat{\pi}(\mathbf{x}))$$

on \mathbf{x} by means of a simple cut-point model. The classical simultaneous analysis-of-variance (ANOVA) selection of split variable and cut-point is implemented. Causal forests are robust to confounding because the score equation (3) is Neyman-orthogonal, in the sense of Chernozhukov et al. (2018), thus enabling it to accurately target $\tau(\mathbf{x})$, even when estimators for the nuisance components $\pi(\mathbf{x})$ or $\mu(\mathbf{x})$ may be somewhat imprecise (Nie and Wager (2021)). Of course, causal forests can be also applied to randomized data in which case treatment should be centered by the true randomization probability π .

2.3. *Model-based forests.* In contrast to the marginal model (1) motivating local centering in causal forests, model-based forests (Seibold, Zeileis and Hothorn (2018)) for real-valued outcomes are based on a model which, in addition to \mathbf{x} , also conditions on treatment assignment $W = w$,

$$(4) \quad (Y | \mathbf{X} = \mathbf{x}, W = w) = \mu(\mathbf{x}) + \tau(\mathbf{x})w + \sigma Z.$$

The main difference between causal forests and model-based forests is that the latter aims to estimate both $\mu(\mathbf{x})$ and $\tau(\mathbf{x})$ simultaneously, whereas the former applies local centering in a two-step approach, that is, treating the prognostic effect $\mu(\mathbf{x})$ as a nuisance parameter. More specifically, by using model (4) instead of model (2), $(\mu(\mathbf{x}), \tau(\mathbf{x}))^\top$ is *simultaneously* estimated by minimizing the L_2 loss

$$(5) \quad \ell_{\text{mob}}(\mu(\mathbf{x}), \tau(\mathbf{x})) = 1/2(Y - \mu(\mathbf{x}) - \tau(\mathbf{x})w)^2$$

w.r.t. μ and τ , the two unknown quantities in this loss function.

Model-based forests separate split-variable and cut-point selection in a way inspired by unbiased recursive partitioning procedures. Specifically, in each node constants $(\hat{\mu}, \hat{\tau})^\top$ are estimated by minimizing

$$\ell_{\text{mob}}(\mu, \tau) := 1/2(Y - \mu - \tau w)^2$$

w.r.t. both μ and τ . A split variable is selected by a bivariate permutation test relying on a quadratic test statistic for the null hypothesis that μ and τ are constant and independent of any split variable \mathbf{X} . For splitting, the variable is selected that has the lowest p -value. Afterward, a cut-point is found by regressing the bivariate score

$$(6) \quad s_{\text{mob}}(\hat{\mu}, \hat{\tau}) := (Y - \hat{\mu} - \hat{\tau}w)(1, w)^\top$$

on covariates \mathbf{x} by a simple bivariate cut-point model. A cut-point is selected as the point that results in the largest discrepancy between the score functions in the two resulting subgroups (details are given in Appendix 2, Seibold, Zeileis and Hothorn (2018)). The core idea of this tree-induction method originates from unbiased recursive partitioning (Hothorn, Hornik and Zeileis (2006)) and the introduction of multiple model-based scores (Zeileis, Hothorn and Hornik (2008)) in this framework. Section 1 in the Supplementary Material A (Dandl et al. (2024a)) provides a more detailed comparison of the cut-point selection of model-based forests with causal forests.

As a side-effect, heterogeneous treatment contrasts $\tau_{2-1}(\mathbf{x}), \tau_{3-1}(\mathbf{x}), \dots, \tau_{K-1}(\mathbf{x})$ of $K > 2$ treatment groups $W | \mathbf{X} = \mathbf{x} \sim M(K, \pi(\mathbf{x}))$ from a multinomial distribution can be estimated by model-based forests. In each node the criterion

$$\frac{1}{2} \left(Y - \mu(\mathbf{x}) - \sum_{k=2}^K \tau_{k-1}(\mathbf{x})w_{k-1} \right)^2$$

is then minimized w.r.t. μ and all treatment contrasts τ_{k-1} for $k = 2, \dots, K$ simultaneously. This allows the comparison of the effects of different treatments or one treatment with various doses to a placebo (application examples could be found in Feng et al. (2012), Schnell et al. (2017), Zanutto, Lu and Hornik (2005)).

2.4. Aggregation and honesty. Once multiple trees have been fitted to subsamples of the data, causal forests and model-based forests apply the same local maximum likelihood aggregation scheme based on nearest neighbor weights for the estimation of heterogeneous treatment effects $\tau(\mathbf{x})$ (Hothorn et al. (2004), Meinshausen (2006), Lin and Jeon (2006), Athey, Tibshirani and Wager (2019), Hothorn and Zeileis (2021a)). First, nearest neighbor weights $\alpha_i(\mathbf{x})$ are derived from the B trees in a forest fitted to observations $(Y_i, \mathbf{x}_i, w_i), i = 1, \dots, N$. These weights measure the relevance of a training observation i for estimating $\tau(\mathbf{x})$. For a forest with B trees, $\alpha_i(\mathbf{x})$ for an observation \mathbf{x} is equal to the frequency with which the i th training sample falls in the same leaf as \mathbf{x} over all B trees. In a second step, $\tau(\mathbf{x})$ is estimated, using the reweighted training data, by minimizing

$$\hat{\tau}(\mathbf{x}) = \arg \min_{\tau} \sum_{i=1}^n \alpha_i^{\text{cf}}(\mathbf{x}) \ell_{\text{cf},i}(\tau)$$

in causal forests and

$$(\hat{\mu}(\mathbf{x}), \hat{\tau}(\mathbf{x}))^{\top} = \arg \min_{\mu, \tau} \sum_{i=1}^n \alpha_i^{\text{mob}}(\mathbf{x}) \ell_{\text{mob},i}(\mu, \tau)$$

in model-based forests, where $\ell_{\text{cf},i}$ and $\ell_{\text{mob},i}$ denote the loss for the i th observation and α_i^{cf} and α_i^{mob} are the weights obtained from a causal forest and a model-based forest, respectively.

Wager and Athey (2018) additionally recommend a subsample splitting technique called honesty, “A tree is honest if, for each training example i , it only uses the response Y_i to estimate the within-leaf treatment effect τ [...] or to decide where to place the splits, but not both.” They empirically and theoretically proved that honesty is necessary to accomplish valid statistical inference. This technique is independent of both tree-induction and forest aggregation and can be applied in both causal forests and model-based forests. In the following we refer to the *adaptive version* of a tree fitting process, when no sample splitting is conducted, and we refer to the *honest version*, when honesty is performed.

2.5. Model generalizations. When heterogeneous treatment effects shall be estimated for an outcome variable Y that is not well described by model (1), adaptations to both causal forests and model-based forests are necessary. Causal forests rely on reformulations of the corresponding estimation problems such that the squared error loss can also be applied in other contexts, for example, in survival analysis (Cui et al. (2023)). For model-based forests, the loss function ℓ_{mob} (5) changes from squared error to the negative log-likelihood of some appropriate model (see Seibold, Zeileis and Hothorn (2016), Seibold, Zeileis and Hothorn (2018), Korepanova et al. (2020), Buri and Hothorn (2020), Fokkema et al. (2018), Hothorn and Zeileis (2021a)).

As a simple example, consider count observations $(Y | \mathbf{X} = \mathbf{x}, W = w) \sim \text{Po}(\exp(\mu(\mathbf{x}) + \tau(\mathbf{x})w))$ from a conditional Poisson distribution. A “Poisson forest” for HTE estimation can be implemented by replacing the squared error loss (5) with the corresponding Poisson negative log-likelihood

$$\ell_{\text{mob}}(\mu(\mathbf{x}), \tau(\mathbf{x})) = \exp(\mu(\mathbf{x}) + \tau(\mathbf{x})w) - (\mu(\mathbf{x}) + \tau(\mathbf{x})w)Y.$$

When it is appropriate to assume $Z \sim N(0, 1)$ with cumulative distribution function Φ , the conditional distribution $(Y \mid \mathbf{X} = \mathbf{x}, W = w) \sim N(\mu(\mathbf{x}) + \tau(\mathbf{x})w, \sigma^2)$ is also normal with cumulative distribution function

$$\mathbb{P}(Y \leq y \mid \mathbf{X} = \mathbf{x}, W = w) = \Phi\left(\frac{y - \mu(\mathbf{x}) - \tau(\mathbf{x})w}{\sigma}\right).$$

For an observed interval $\underline{y} < Y \leq \bar{y}$, model-based forests, equipped with the negative log-likelihood

$$\ell_{\text{mob}}(\mu(\mathbf{x}), \tau(\mathbf{x}), \sigma) = -\log\left(\Phi\left(\frac{\bar{y} - \mu(\mathbf{x}) - \tau(\mathbf{x})w}{\sigma}\right) - \Phi\left(\frac{\underline{y} - \mu(\mathbf{x}) - \tau(\mathbf{x})w}{\sigma}\right)\right),$$

allow us to implement a variant of model-based forests applicable to imprecise interval-censored observations. In a Tobit model, this is the negative log-likelihood contributed by an observation $(-\infty, 0]$ left-censored at zero (Schlosser et al. (2019), equation (2.1)). A similar likelihood, however, without the strict normal assumption will be introduced for interval-censored blood loss in Section 5.1. In this sense model-based forests can be understood as a conceptual and computational framework for method construction rather than a model with a special domain of application.

3. Strategies and research questions for blended approaches. When applied to data well described by the additive model (1) in the randomized setting, the principles underlying causal forests and model-based forests are conceptually the same; the only difference is that causal forests follow a sequential two-step approach and model-based forests implement a simultaneous approach to parameter estimation. We are now interested in assessing the impact of implementation details in causal forests and model-based forests on HTE estimation performance by the two algorithms. The theoretical understanding from Section 2 motivates straightforward adaptations to model-based forests such that the procedure can also be applied to observational studies. The flexibility of its implementation in *model4you* allows us to define and evaluate blended estimation approaches transferring the concept of local centering from causal forests to model-based forests. Along with these new algorithms, we propose a set of five research questions, which we investigate empirically in Section 4. An overview of the questions is given in Table 1. We begin with the standard implementations of causal forests (cf) and model-based (mob) forests without centering:

RQ 1. How do cf and mob, as implemented in the two R add-on packages *grf* (for cf) and *model4you* (for mob), compare to each other in randomized and observational settings?

After addressing RQ 1, the question remains if and to what extent local centering inherent in cf leads to more robustness against confounding effects. To answer that, we will incorporate orthogonalization in mob, as explained in the following. Causal forests apply local centering to both the outcome Y and treatment indicator w , and mob do not center locally at all. To bring cf and mob closer, we define a method which applies mob to the model

$$\mathbb{E}(Y \mid \mathbf{X} = \mathbf{x}, W = w) = \hat{m}(\mathbf{x}) + \tilde{\mu}(\mathbf{x}) + \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x})),$$

that is, after centering the treatment indicator w and the outcome Y . By using $\tilde{\mu}(\mathbf{x})$ instead of $\mu(\mathbf{x})$, we emphasize that $\tilde{\mu}(\mathbf{x})$ is now the prognostic effect for the *centered* Y .

The rationale is to estimate the marginal mean and propensities $\pi(\mathbf{x})$, as in cf first, and then apply mob to the centered treatment $w - \hat{\pi}(\mathbf{x})$ and centered outcome $Y - \hat{m}(\mathbf{x})$ to obtain the prognostic and predictive effect. We call this approach $\text{mob}(\hat{W}, \hat{Y})$. The bivariate score function for mob is changed from (6) to

$$s_{\text{mob}(\hat{W}, \hat{Y})}(\hat{\mu}, \hat{\tau}) := (Y - \hat{m}(\mathbf{x}) - \hat{\mu} - \hat{\tau}(w - \hat{\pi}(\mathbf{x}))(1, w - \hat{\pi}(\mathbf{x}))^\top.$$

TABLE 1
Overview of research questions

RQ	Question	Methods	Linear predictors
1	Comparison of causal forests and model-based forests	cf mob	$\hat{m}(\mathbf{x}) + \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x}))$ $\mu(\mathbf{x}) + \tau(\mathbf{x})w$
2	Effect of splitting only in $\tau(\mathbf{x})$ vs. in $\tau(\mathbf{x})$ and $\tilde{\mu}(\mathbf{x})$	mobcf mob(\hat{W} , \hat{Y})	$\hat{m}(\mathbf{x}) + \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x}))$ $\hat{m}(\mathbf{x}) + \tilde{\mu}(\mathbf{x}) + \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x}))$
3	Comparison of causal forests implemented in grf vs. model4you	cf mobcf	$\hat{m}(\mathbf{x}) + \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x}))$ $\hat{m}(\mathbf{x}) + \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x}))$
4	Effect of locally centering W in model-based forests	mob(\hat{W}) mob	$\mu(\mathbf{x}) + \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x}))$ $\mu(\mathbf{x}) + \tau(\mathbf{x})w$
5	Effect of additionally centering Y in model-based forests centering W	mobcf mob(\hat{W} , \hat{Y}) mob(\hat{W})	$\hat{m}(\mathbf{x}) + \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x}))$ $\hat{m}(\mathbf{x}) + \tilde{\mu}(\mathbf{x}) + \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x}))$ $\mu(\mathbf{x}) + \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x}))$

In cases where local centering of Y effectively regresses out the effect of X on Y , $\tilde{\mu}(\mathbf{x})$ will be close to 0. Since removing $\tilde{\mu}$ leads to the conditional mean function underlying cf

$$\mathbb{E}(Y \mid X = \mathbf{x}, W = w) = \hat{m}(\mathbf{x}) + \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x})),$$

we call this version “mobcf.” Both the outcome and the treatment indicator are centered and only splitting with respect to scores corresponding to the treatment effect τ is performed, while intercept scores are ignored in this process. The only difference between mobcf and mob(\hat{W} , \hat{Y}) is that simultaneous splitting in both the intercept and treatment effect parameters is performed by the latter, whereas the intercept is ignored in the former.

RQ 2. How does mob(\hat{W} , \hat{Y}) perform compared to mobcf?

The mobcf approach helps us to directly compare the different more technical aspects, such as variable and split point selection or stopping criteria, of tree induction implemented in *grf* and *model4you*, because it can be seen as a reimplementaion of cf using the computational infrastructure of the *model4you* package.

RQ 3. How does mobcf perform compared to cf implemented in *grf*?

Centering the response is straightforward under L_2 loss but more difficult under other forms of the likelihood, as discussed in Section 2.5. The questions arise if and to what extent solely centering of the treatment indicator w already improves the estimation accuracy in observational settings. To answer that, we define a “hybrid approach” mob(\hat{W}) that applies mob to models parameterized by $\mu(\mathbf{x}) + \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x}))$, that is, after solely centering the w but not the outcome Y . The score function for mob is changed from (6) to

$$s_{\text{mob}(\hat{W})}(\hat{\mu}, \hat{\tau}) := (Y - \hat{\mu} - \hat{\tau}(w - \hat{\pi}(\mathbf{x}))(1, w - \hat{\pi}(\mathbf{x}))^\top.$$

RQ 4. How does solely centering of the treatment indicator (mob(\hat{W})) influence the performance of mob without centering in settings with confounding?

The final research question is whether additional outcome centering improves upon a forest with treatment centering and simultaneous splits in prognostic and predictive effects as implemented by mob(\hat{W}).

RQ 5. How does mob(\hat{W}) perform compared to mob that center both treatment and outcome (mobcf, and mob(\hat{W} , \hat{Y}))?

4. Empirical evaluation. In this section we provide answers to the research questions defined in Section 3 by evaluating the performance of *cf* and *mob* as well as the different blended versions in a simulation study for normal outcomes, different predictive and prognostic effects, and a varying number of observations and covariates. The reference implementations in the *grf* and *model4you* R add-on packages were used for the original *cf* and *mob* algorithms. Moreover, the blended approaches from Section 3 are implemented using *model4you*, that is, by fitting model-based forests after centering of treatment indicators ($\text{mob}(\hat{W})$) and, additionally, of outcomes ($\text{mob}(\hat{W}, \hat{Y})$) and *mobcf*, with and without explicitly accounting for μ , respectively).

4.1. *Data-generating process.* The comparison is based on the study settings of Nie and Wager (2021). The authors proposed four study settings—referred to as Setups A, B, C, and D. For Setup A explanatory variables were sampled by $X \sim U([0, 1]^P)$, and for the other three setups, they used $X \sim N(0, \mathbb{1}_{P \times P})$ – with $P = \{10, 20\}$ (5 informative and $P - 5$ noise variables). Treatment was sampled by $W \mid X = \mathbf{x} \sim B(1, \pi(\mathbf{x}))$ with propensity function $\pi(\mathbf{x})$ that varied among the four considered setups,

$$\pi(\mathbf{x}) = \begin{cases} \pi_A(x_1, x_2) = \max\{0.1, \min\{\sin(\pi x_1 x_2), 1 - 0.1\}\}, \\ \pi_B \equiv 0.5, \\ \pi_C(x_2, x_3) = 1/(1 + \exp(x_2 + x_3)), \\ \pi_D(x_1, x_2) = 1/(1 + \exp(-x_1) + \exp(-x_2)). \end{cases}$$

For Setup B probability $\pi \equiv 0.5$ referred to a randomized study. The conditional average treatment effect function for each setup was given as

$$\tau(\mathbf{x}) = \begin{cases} \tau_A(x_1, x_2) = (x_1 + x_2)/2, \\ \tau_B(x_1, x_2) = x_1 + \log(1 + \exp(x_2)), \\ \tau_C \equiv 1, \\ \tau_D(x_1, x_2, x_3, x_4, x_5) = \max\{x_1 + x_2 + x_3, 0\} - \max\{x_4 + x_5, 0\}. \end{cases}$$

For Setup C the treatment effect was constant. The prognostic effects were defined as

$$\mu(\mathbf{x}) = \begin{cases} \mu_A(x_1, x_2, x_3, x_4, x_5) = \sin(\pi x_1 x_2) + 2(x_3 - 0.5)^2 + x_4 + 0.5x_5, \\ \mu_B(x_1, x_2, x_3, x_4, x_5) = \max\{x_1 + x_2, x_3, 0\} + \max\{x_4 + x_5, 0\}, \\ \mu_C(x_1, x_2, x_3) = 2 \log(1 + \exp(x_1 + x_2 + x_3)), \\ \mu_D(x_1, x_2, x_3, x_4, x_5) = (\max\{x_1 + x_2 + x_3, 0\} + \max\{x_4 + x_5, 0\})/2. \end{cases}$$

Overall, Setup A has complicated confounding that needs to be overcome before a relatively simple treatment effect function $\tau(\mathbf{x})$ can be estimated. In Setup B it is possible to accurately estimate τ without explicitly controlling for confounding. Setup C has strong confounding, but the propensity score function is easier to estimate than the prognostic effect while the treatment effect is constant. In Setup D the treatment and control arms are unrelated, in the sense that $\mathbb{E}[Y \mid X, W = 1]$ and $\mathbb{E}[Y \mid X, W = 0]$ are uncorrelated and there is no benefit to jointly learn them.

As in Nie and Wager (2021), we studied a normal linear regression model

$$(Y \mid X = \mathbf{x}, W = w) \sim N(\mu(\mathbf{x}) + \tau(\mathbf{x})(w - 0.5), 1),$$

where half of the predictive effect was added to the prognostic effect.

All procedures were applied to 100 learning samples of size $N \in \{800, 1600\}$ and number of explanatory variables $P \in \{10, 20\}$. In order to minimize the impact of different implementation details, *cf*, *mob* and the blended versions were grown with the same hyperparameter

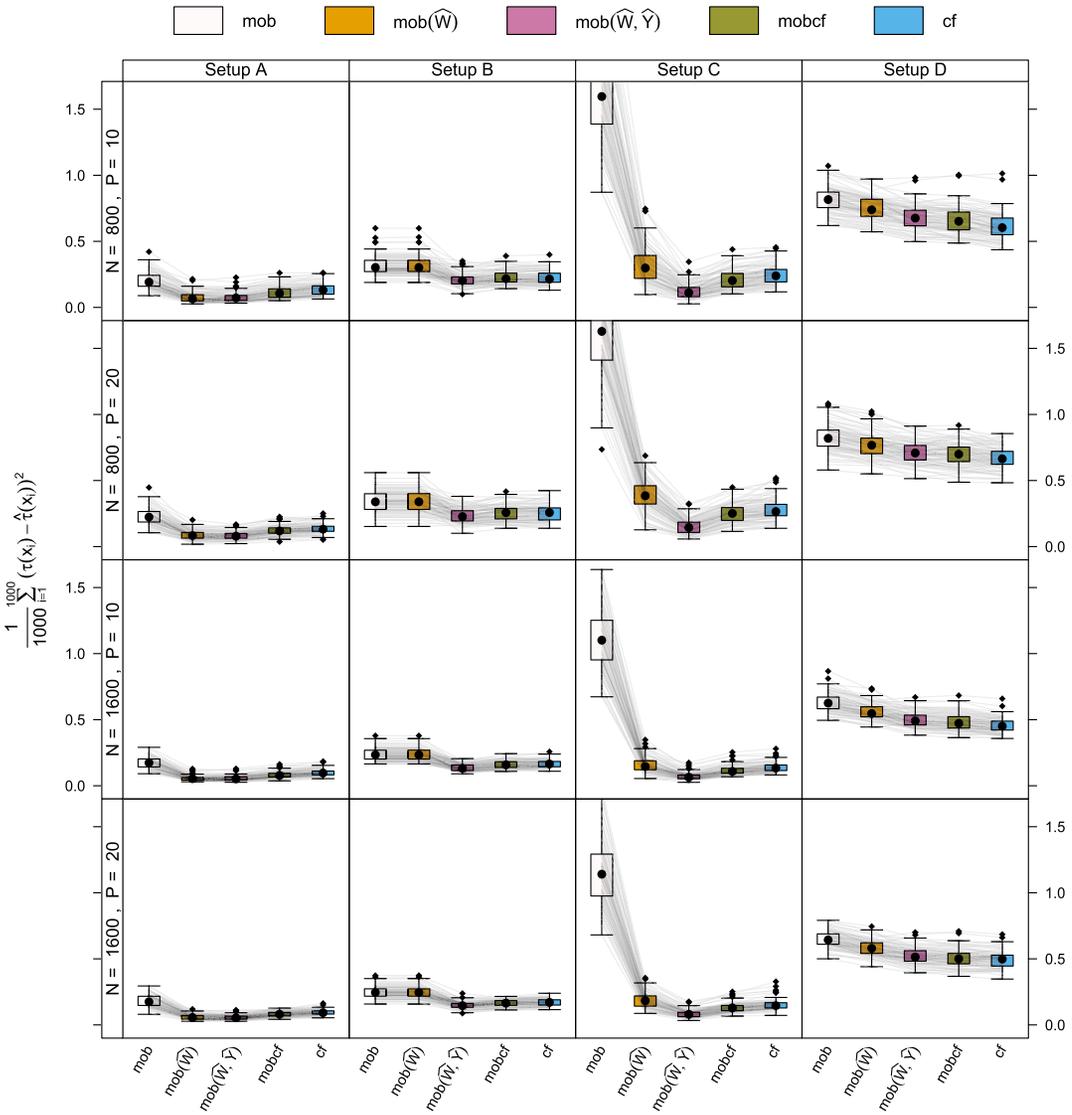


FIG. 1. Results for the experimental setups 4.1. Direct comparison of the adaptive versions of causal forests (*cf*), model-based forests without centering (*mob*), *mob* imitating causal forests (*mobcf*), *mob* with centered W ($mob(\hat{W})$), and additional of Y ($mob(\hat{W}, \hat{Y})$).

options; see Section 7. Propensities $\pi(\mathbf{x})$ and means $m(\mathbf{x})$ were estimated by *grf* regression forests for local centering in all forest variants. For the causal forest, the outcome was always centered by $\hat{m}(\mathbf{x})$. In case of randomized data (Setup B), the treatment indicator was centered by $\pi \equiv 0.5$; in all other settings, estimated propensities $\hat{\pi}(\mathbf{x})$ were used.

Performance was assessed by the ability of the methods to estimate the predictive effect $\tau(\mathbf{x})$. The mean squared error $\mathbb{E}_X\{(\hat{\tau}(X) - \tau(X))^2\}$, evaluated on a test sample of size 1000, was used to compare the predictive performance of all candidate models in the 16 different scenarios. The results are shown in Figure 1.

The results were also analyzed statistically by means of a normal linear mixed model with log-link, explaining the estimated mean squared error for $\hat{\tau}(\mathbf{x})$ by a four-way interaction of data-generating process, sample size N , dimension P , and random forest variant. We estimated the mean squared error ratios between *cf* and *mob* (RQ 1), between *mobcf* and $mob(\hat{W}, \hat{Y})$ (RQ 2), between *cf* and the *mobcf* approach (RQ 3), between *mob* with centered

W ($\text{mob}(\hat{W})$) and without (mob) (RQ 4), and between $\text{mob}(\hat{W})$ and mobcf or $\text{mob}(\hat{W}, \hat{Y})$ (RQ 5). For each simulation run, the model featured a corresponding random intercept reflecting the paired simulation design. Simultaneous 95% confidence intervals for the mean squared error ratios are presented along with the estimates. For example, the ratio of the mean squared errors of cf and mob in the first line of Table 2 was 0.663 with confidence interval (0.596, 0.738). This is in line with the performance error of cf being at least 59.6% and at most 73.8% of the performance error of mob , with 66.3% denoting the estimate. Bold, italic, and normal fonts are used to indicate superior, inferior, and equivalent prediction performance.

4.2. Results. The results for adaptive forests are presented in Figure 1. In Section 2 of the Supplementary Material A (Dandl et al. (2024a)), we report on the effect of honesty on predictive error as well as the mean squared differences in performance to cf for the adaptive and honest versions (Figures S.1 and S.2). The statistical analysis of the results is given in Table 2 for the adaptive version of forests and in Table S.1 of the Supplementary Material A for the honest version:

RQ 1. mob vs. cf. In all setups cf outperformed mob . Especially in Setup C, mob was unable to overcome the strong confounding effect and, therefore, did not provide accurate estimates for the (constant) treatment effect.

RQ 2. mob(\hat{W}, \hat{Y}) vs. mobcf. The $\text{mob}(\hat{W}, \hat{Y})$ approach performed better than the mobcf approach in almost all scenarios, except for Setup D. (However, uncorrelated treatment and control arms rarely occur in reality. All methods had a higher MSE than in the other setups.) These performance differences suggest that splitting by treatment *and* prognostic effect is beneficial.

RQ 3. mobcf vs. cf. Despite the fundamentally different internal splitting and stopping criteria, the original implementation of cf from package *grf* had very similar performance to our reimplementations mobcf from package *model4you* in Setup A and B. In Setup C with strong confounding, the mobcf approach performed slightly better than cf , while in Setup D cf performed slightly better.

RQ 4. mob(\hat{W}) vs. mob. In case of confounding (Setup A, C), local centering of W ($\text{mob}(\hat{W})$) significantly improved the performance of mob . In Setup B without confounding, both approaches performed equally since $\text{mob}(\hat{W})$ is equal to mob applied to $w - 0.5$.

RQ 5. Methods centering the outcome (mobcf, mobmob(\hat{W}, \hat{Y})) vs. mob(\hat{W}). By centering the outcome Y in addition to the treatment W , $\text{mob}(\hat{W}, \hat{Y})$ and mobcf performed better than $\text{mob}(\hat{W})$, except for Setup A—centering the outcome did not further improve the results. The improvements by additionally centering Y were relatively small for mob , compared to the improvements due to centering the treatment W (see RQ 4).

Overall, our results reveal treatment effect centering ($\text{mob}(\hat{W})$) as the most relevant ingredient to random forests for HTE estimation in observational studies. If possible, additional centering Y in combination with simultaneous estimation of predictive and prognostic effects ($\text{mob}(\hat{W}, \hat{Y})$) is recommended.

5. Effect of cesarean section on postpartum blood loss. In this section we discuss random forest-based HTEs expressing the additional amount of blood loss explained by prepartum variables, comparing cesarean sections with vaginal deliveries. We analyze data from 1309 women who participated in a prospective study conducted from October 2015 to November 2016 at the University Hospital Zurich (details and data are available from Haslinger et al. (2020)). The outcome is defined as measured blood loss (MBL) in mL, and

TABLE 2

Results for the experimental setups 4.1 for the adaptive versions of the methods. Comparison of mean squared errors for $\hat{\tau}(\mathbf{x})$ in the different scenarios. Estimates and simultaneous 95% confidence intervals were obtained from a normal linear mixed model with log-link. Cells printed in bold font correspond to a superior reference (mob in the first and fourth columns, $\text{mob}(\hat{W}, \hat{Y})$ in the second column, mobcf in the third and fifth column, and $\text{mob}(\hat{W})$ in the last column), cells printed in italics indicate an inferior reference

			Mean squared error ratio					
DGP	N	P	(RQ 1)	(RQ 2)	(RQ 3)	(RQ 4)	(RQ 5)	(RQ 5)
			cf vs. mob	mobcf vs. $\text{mob}(\hat{W}, \hat{Y})$	cf vs. mobcf	$\text{mob}(\hat{W})$ vs. mob	$\text{mob}(\hat{W})$ vs. mobcf	$\text{mob}(\hat{W})$ vs. $\text{mob}(\hat{W}, \hat{Y})$
Setup A	800	10	<i>0.663 (0.596, 0.738)</i>	1.446 (1.206, 1.734)	1.165 (1.017, 1.335)	<i>0.392 (0.334, 0.461)</i>	<i>0.689 (0.574, 0.826)</i>	0.996 (0.806, 1.230)
		20	<i>0.596 (0.537, 0.662)</i>	1.465 (1.228, 1.747)	1.111 (0.972, 1.270)	<i>0.385 (0.331, 0.446)</i>	<i>0.717 (0.604, 0.850)</i>	1.050 (0.859, 1.284)
	1600	10	<i>0.575 (0.499, 0.663)</i>	1.458 (1.123, 1.893)	1.201 (0.991, 1.455)	<i>0.324 (0.258, 0.408)</i>	<i>0.677 (0.521, 0.881)</i>	0.988 (0.727, 1.342)
		20	<i>0.517 (0.447, 0.598)</i>	1.453 (1.117, 1.889)	1.150 (0.944, 1.401)	<i>0.328 (0.265, 0.407)</i>	<i>0.730 (0.567, 0.940)</i>	1.061 (0.788, 1.428)
Setup B	800	10	<i>0.707 (0.662, 0.756)</i>	1.099 (1.015, 1.190)	0.987 (0.914, 1.065)	1.000 (0.947, 1.056)	1.395 (1.306, 1.491)	1.533 (1.429, 1.646)
		20	<i>0.745 (0.701, 0.791)</i>	1.093 (1.018, 1.174)	1.001 (0.935, 1.071)	1.000 (0.951, 1.052)	1.345 (1.266, 1.428)	1.470 (1.380, 1.567)
	1600	10	<i>0.695 (0.635, 0.762)</i>	1.166 (1.036, 1.313)	1.034 (0.929, 1.152)	1.000 (0.929, 1.076)	1.487 (1.355, 1.633)	1.734 (1.563, 1.924)
		20	<i>0.683 (0.625, 0.746)</i>	1.110 (0.992, 1.243)	1.037 (0.934, 1.152)	1.000 (0.932, 1.073)	1.518 (1.387, 1.662)	1.686 (1.529, 1.859)
Setup C	800	10	<i>0.148 (0.141, 0.156)</i>	1.693 (1.514, 1.893)	1.150 (1.067, 1.240)	<i>0.197 (0.190, 0.205)</i>	1.529 (1.429, 1.636)	2.589 (2.335, 2.870)
		20	<i>0.170 (0.162, 0.177)</i>	1.673 (1.520, 1.841)	1.123 (1.051, 1.199)	<i>0.236 (0.229, 0.244)</i>	1.563 (1.474, 1.657)	2.615 (2.395, 2.856)
	1600	10	<i>0.124 (0.113, 0.136)</i>	1.651 (1.348, 2.023)	1.184 (1.032, 1.359)	<i>0.143 (0.132, 0.155)</i>	1.368 (1.201, 1.558)	2.258 (1.868, 2.731)
		20	<i>0.131 (0.121, 0.142)</i>	1.573 (1.320, 1.875)	1.166 (1.030, 1.320)	<i>0.163 (0.153, 0.174)</i>	1.452 (1.295, 1.628)	2.284 (1.943, 2.684)
Setup D	800	10	<i>0.756 (0.737, 0.775)</i>	<i>0.970 (0.945, 0.996)</i>	<i>0.934 (0.909, 0.960)</i>	<i>0.917 (0.897, 0.938)</i>	1.133 (1.105, 1.162)	1.099 (1.072, 1.127)
		20	<i>0.807 (0.788, 0.826)</i>	0.983 (0.959, 1.008)	<i>0.958 (0.933, 0.982)</i>	<i>0.926 (0.906, 0.947)</i>	1.100 (1.074, 1.126)	1.081 (1.056, 1.107)
	1600	10	<i>0.720 (0.696, 0.744)</i>	0.970 (0.936, 1.005)	<i>0.939 (0.904, 0.974)</i>	<i>0.885 (0.859, 0.912)</i>	1.155 (1.116, 1.194)	1.120 (1.083, 1.157)
		20	<i>0.763 (0.739, 0.787)</i>	0.967 (0.935, 1.001)	0.982 (0.949, 1.018)	<i>0.894 (0.869, 0.920)</i>	1.151 (1.114, 1.189)	1.113 (1.078, 1.149)

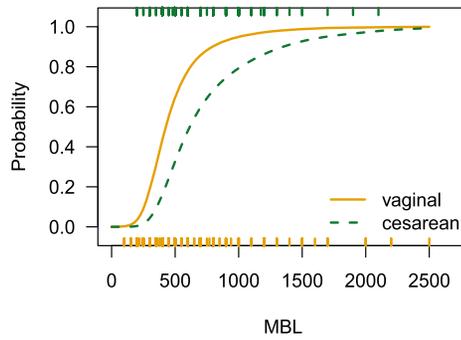


FIG. 2. Marginal distribution of measured blood loss (mL) for cesarean section and vaginal delivery. Rugs indicate measured blood loss observations.

the authors ensured application of a standardized measurement procedure for all study participants (Kahr et al. (2018)). For our study we removed one outlier observation with a blood loss of 5700 mL and eight observations with missing values for BMI so that a sample of size $N = 1300$ remains. MBL was recorded as an interval-censored variable, because it is impossible to exactly determine the amount of blood loss in the sometimes hectic environment of a delivery ward (Kahr et al. (2018)). Potential inaccuracies in the measuring process are represented by an interval width of 50 mL for blood losses ≤ 1 L and an interval width of 100 mL when the mother lost more than one liter of blood. Measured blood loss can a priori be considered a positive real and right-skewed variable (Figure 2). Table 3 gives a summary of the eight considered prepartum characteristics ($P = 8$).

As the outcome variable MBL is skewed and interval-censored not all assumptions for causal forests are fulfilled as they estimate a conditional mean of some continuous outcome optimizing L_2 risk. The extensibility of model-based forests discussed in Section 2.5 allows us to take into account the structural assumptions of MBL by substituting ℓ_{mob} in (5) with the negative log-likelihood of a more appropriate model. We set up a model-based transformation forest with treatment centering by combining the $\text{mob}(\hat{W})$ approach using local centering of the treatment indicator within a transformation model.

5.1. *Transformation base model.* The reasoning in Section 2 is based on the normal linear model (4) and its corresponding likelihood (5) for absolutely continuous observations. While the latter can easily be adapted to interval-censored observations, more effort is needed for allowing skewness in the response distribution. Adopting a standard normal distribution for the error term Z , like in Section 2.5, model (4) can be written as a conditional distribution

TABLE 3
Prepartum characteristics

Variable	Description	Range
GA	Gestational age	177–297 (days)
AGE	Maternal age	18–48 (years)
MULTIPAR	Multiparity	no/yes
BMI	Body mass index	15.4–66
MULTIFET	Multifetal pregnancy	no/yes
NW	Neonatal weight	360–4630 (g)
IOL	Induction of labor	no/yes
AIS	Chorioamnionitis	no/yes

function

$$\mathbb{P}(\text{MBL} \leq y \mid \mathbf{X} = \mathbf{x}, W = w) = \Phi\left(\frac{y - \mu(\mathbf{x}) - \tau(\mathbf{x})w}{\sigma}\right).$$

In this model symmetry is achieved by a linear transformation of the y argument on the probit scale. Replacement of this linear transformation by a potentially nonlinear one gives rise to transformation models. In combination with the probit link, this model is a Box–Cox-type linear regression model that transforms the skewed outcome variable to normality. Instead of using the traditional Box–Cox power transformation, we estimate a suitable transformation of MBL by means of a flexible polynomial in Bernstein form (Hothorn, Möst and Bühlmann (2018)). Ignoring covariates and the local centering of W for a moment, our transformation model describes the conditional distribution of the positive skewed real variable MBL using mode of delivery W as treatment indicator for vaginal delivery ($W = 0$) vs. cesarean section ($W = 1$),

$$\mathbb{P}(\text{MBL} \leq y \mid W = w) = \Phi(h(y) - \mu - \tau w).$$

Deviations from normality are captured by the nonlinear transformation function h in this model. Because the transformation function h contains an intercept term, the parameter μ is not identified. We thus estimate the transformation base model under the constraint $\mu \equiv 0$. The intercept function h varies with the chosen MBL cut-off y and is smooth and monotonically increasing; a polynomial in Bernstein form of order six was used to parameterize this function. The parameter $\tau = \mathbb{E}(h(Y(1)) - h(Y(0)))$ is not identical to an average treatment effect on the untransformed scale, which could be interpreted directly in terms of the original units of the outcome (here blood loss in mL). Nevertheless, τ in our transformation model has an intuitive interpretation corresponding to Cohen's d : the units of the treatment effect correspond to standard deviations under the normal model.

The parameters of the transformation base model were estimated by minimization of the negative log-likelihood for an interval-censored observation $(y, \bar{y}]$

$$\begin{aligned} \ell_{\text{Trafo}}(\mu, \tau, \boldsymbol{\vartheta}) &= -\log(\mathbb{P}(y < Y \leq \bar{y} \mid W = w)) \\ &= -\log(\Phi(h(\bar{y} \mid \boldsymbol{\vartheta}) - \mu - \tau w) - \Phi(h(y \mid \boldsymbol{\vartheta}) - \mu - \tau w)), \end{aligned}$$

where all parameters, including $\boldsymbol{\vartheta}$ for the transformation function, are estimated in each node. A parameterisation of h in terms of a polynomial in Bernstein form $h(\cdot \mid \boldsymbol{\vartheta})$ ensures uniform convergence to any continuous unknown transformation function h on some interval by Weierstrass' approximation theorem (Farouki (2012)).

5.2. Personalized transformation model. The results of Section 2–4 motivate the application of model-based forests to a Box–Cox-type transformation model for the estimation of HTEs of cesarean sections on PPH. The transformation base model provides skewness and interval-censoring, whereas the locally centered treatment indicator controls for potential confounding. In more detail we used a $\text{mob}(\hat{W})$ forest in combination with the transformation base model, that is, with local centered treatment indicator \hat{w} , to compute personalized treatment effects $\tau(\mathbf{x})$ and prognostic effects $\mu(\mathbf{x})$ of the model

$$(7) \quad \mathbb{P}(\text{MBL} \leq y \mid \mathbf{X} = \mathbf{x}, W = w) = \Phi(h(y) - \mu(\mathbf{x}) - \tau(\mathbf{x})(w - \hat{\pi}(\mathbf{x}))).$$

As in the simulation study, a regression forest was applied to estimate propensities $\pi(\mathbf{x})$. We only used locally centered propensities because the empirical results of Section 4 showed that centering W was the main driver for good performance in observational settings. Furthermore, while centering W is straightforward for the transformation model at hand, implementing centering on the outcome Y is less clear.

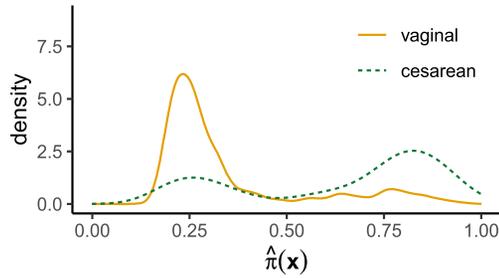


FIG. 3. Estimates of propensity scores $\pi(\mathbf{x})$ returned by the regression forest for orthogonalization of the treatment indicator.

Figure 3 shows that the distribution functions of $\hat{\pi}(\mathbf{x})$ for each treatment group greatly differ. This indicates that prepartum characteristics indeed influence the mode of delivery and that the treated and control group are dissimilar with respect to these characteristics.

We first fitted the transformation base model without covariates but with propensity-centered mode of delivery to estimate a constant effect adjusted for potential confounding. The corresponding effect $\hat{\tau}$, that is, the marginal Cohen’s d , was 0.823 ($CI_{0.95} = (0.686, 0.959)$), indicating that women giving birth by cesarean section have a higher postpartum blood loss compared to women giving birth by vaginal delivery.

The model-based transformation forest was fitted with the same hyperparameter settings, as in the simulation study (Section 7). We did not adjust the hyperparameters because random forests have been shown to be insensitive to hyperparameter changes (Probst, Boulesteix and Bischl (2019)). Figure S. 3 in the Supplementary Material A (Dandl et al. (2024a)) demonstrates this for the $mtry$ parameter—the number of chosen variables per split. We only analysed the $mtry$ parameter since Probst, Wright and Boulesteix (2019) found that the “ $mtry$ parameter is most influential [...]” while “[s]ample size and node size have a minor influence on the performance [...]”.

Figure 4 depicts the distribution of the estimated out-of-bag (OOB) heterogeneous treatment effects $\hat{\tau}(\mathbf{x})$ of cesarean section compared to vaginal delivery. The distribution is unimodal and slightly left-skewed. For almost all births, a cesarean section increases the risk for higher blood losses, compared to vaginal delivery. For comparison the average treatment effect of $\hat{\tau} = 0.823$ of the transformation base model is included.

The interval-censored negative log-likelihood of the transformation base model was 3613.972. The model-based transformation forest improved upon this, yielding a likelihood of 3413.989 (estimated in-bag to make it comparable to the transformation base model).

5.3. *Dependence plots.* The dependency of the treatment effect τ on the prepartum variables is visualized by dependence plots (Figure 5). Scatter plots are used for continuous co-

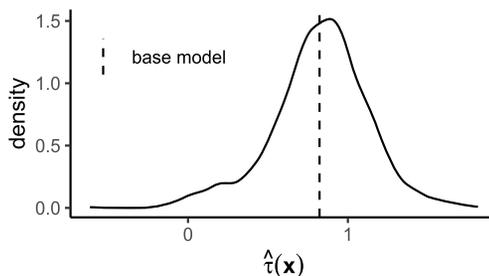


FIG. 4. Kernel density estimates of the personalized treatment estimates of the model-based transformation forest. The dashed line presents the estimated effect of the transformation base model.

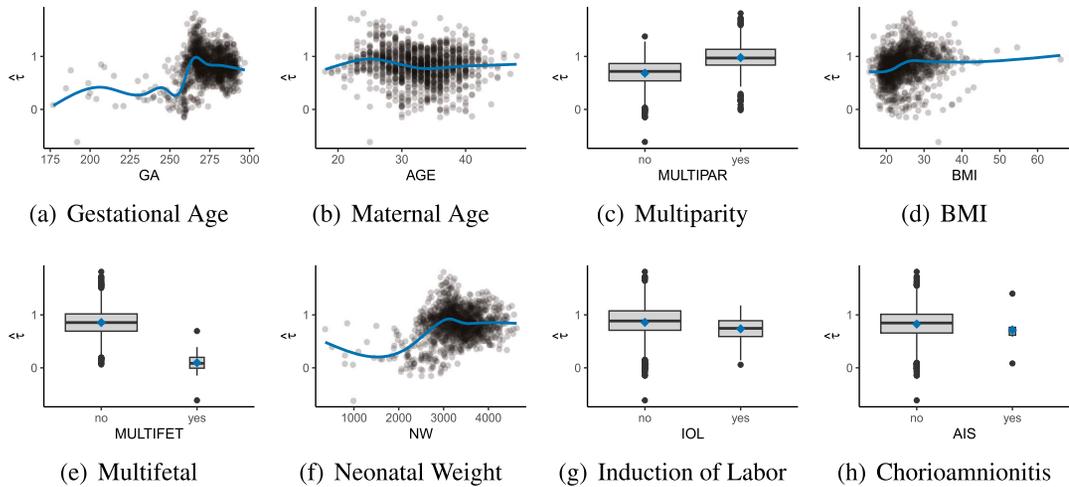


FIG. 5. Dependency plots of the individual treatment effects calculated by the model-based transformation forest. Values $\hat{\tau} > 0$ mean that cesarean section increases the blood loss, compared to vaginal delivery. Lines and diamond points depict (smooth conditional) mean effects.

variates and boxplots for categorical covariates. We also provide mean effects per group for categorical covariates and the smooth conditional mean effect function for continuous covariates. The latter was estimated by a generalized additive model (GAM) with a single smooth term depending on the considered variable. Births with higher gestational age, higher neonatal weight, and singleton pregnancy have a higher risk for elevated blood loss due to cesarean section compared to vaginal delivery. The effect differences were most pronounced between multifetal and singleton births. For multifetal pregnancies treatment effects are closer to zero than for singleton pregnancies. For a very premature multifetal birth (gestational age of 192 days) of a 25-year-old mother with an elevated BMI of 33.7, a cesarean section was determined to be most effective ($\hat{\tau} = -0.614$). Because the distribution of the gestational age (GA) is left-skewed, the curve of the smoothed conditional mean effects is somewhat erratic. It might also indicate that GA was often used as a splitting variable. While interpreting these results, it should be noted that violations of the unconfoundedness assumption do not seem implausible.

5.4. Model interpretation and communication. Interpretation and risk communication in terms of predicted $\hat{\tau}(\mathbf{x})$ is difficult because the effect is defined by Cohen's d on a transformed latent normal scale in model (7). However, the model allows conditional quantiles to be computed, and thus information about the conditional MBL distribution for given prepartum covariates and propensities $\hat{\pi}(\mathbf{x})$ can be expressed on the quantile scale for both modes of delivery.

To assess the prognostic effects on MBL, we computed median measured blood losses for $W = 0$ (vaginal delivery), given the covariates and propensities. Figure 6 indicates that a gestational age of about 270 days, a birth weight around 3050 g, and singleton births are associated with small median postpartum blood losses for vaginal deliveries.

The predictive effect of a cesarean section on MBL in such a low-risk group can be communicated by comparing the MBL distributions under vaginal delivery and cesarean section. The median blood loss for a hypothetical woman in this low-risk group (aged 32.7 years with a BMI of 24.7, the mean values in the study population) is predicted to increase from 329 mL (vaginal delivery, 80% prediction interval 209–507 mL) to 470 mL (cesarean section, 80% prediction interval 305–817 mL) by our model. The asymmetric prediction intervals reflect

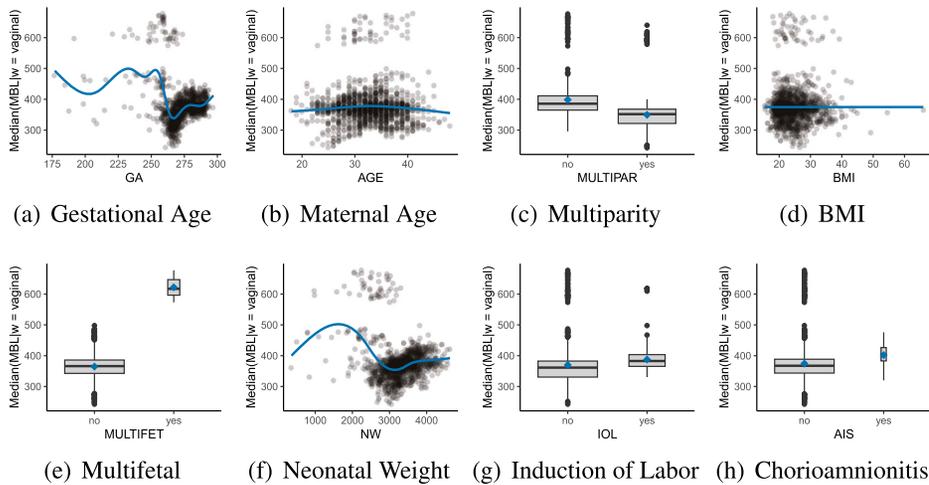


FIG. 6. *Dependency plots of median measured blood losses calculated by the model-based transformation for-est. Higher values mean higher blood loss. Lines and diamond points depict (smooth conditional) means.*

skewness in the MBL distribution, and the wider interval for a cesarean section suggests variance heterogeneity is captured by the model. The risk of PPH (defined by the 500 mL cut-off) is small for vaginal deliveries but substantial under a cesarean section.

6. Discussion and outlook.

6.1. *Effects of cesarean sections of postpartum blood loss.* The lives of many of us have been, or will be, impacted by a cesarean section directly or indirectly. Empowering women for making an informed decision, especially in an elective setting, crucially relies on evidence about the short- and long-term consequences for them and their children (Antoine and Young (2021)). Providing an estimate of the individual predicted excess blood loss caused by a cesarean section, in comparison to a vaginal delivery, to pregnant women and their obstetricians not only offers the possibility to decide based on a personalized risk assessment but also has the potential to help the overarching goal of reducing the prevalence of cesarean sections. The question to perform a cesarean section or not is less imminent in women with obvious risk factors, which make a cesarean section inevitable (e.g., prematurity and multiple fetus pregnancy), but is of utmost clinical interest in women with a prepartum low-risk profile (singleton pregnancy at term with normal fetal weight estimation). To the best of our knowledge, this is the first study to predict excess postpartum blood loss in low-risk women. Our approach of modeling the continuous blood loss distribution for arbitrary cut-off values is also unique in the sense that published prognostic models provide risk estimates for events $MBL > 500$ mL, or other prespecified cut-off values, only.

Our results were estimated based on data originating from a prospective study employing a standardized and validated assessment of blood loss under both modes of delivery. Such efforts can only be successfully implemented in a controlled setting and hardly apply to retrospective collections of routine clinical data from multiple study centers. However, the detection of smaller but still relevant patterns in HTEs might require more information than available from the $N = 1300$ study participants. The random forest methodology would allow differentiation between planned and unplanned cesarean sections (Section 2.3) in a single model; however, the sample sizes in the present study seem too limited for such an analysis. It remains to be seen if refined analyses of large-scale routine clinical data will provide results similar to those reported here.

6.2. *Forest-based HTE estimation.* From a statistical perspective, estimating heterogeneous treatment effects (HTEs) is a difficult task, both when data from randomized trials and observational studies are analyzed. Based on a common theoretical understanding of two strands of random forest algorithms for HTE estimation, we hypothesized that centering the treatment with corresponding propensities helps to address confounding. The empirical results suggest that this simple modification of the data is instrumental for the analysis of observational and thus potentially confounded data.

Centering the outcome is equally simple in models for conditional means but may be much harder in other models. Empirically, we found that the combination of centered treatment and simultaneous split selection (with respect to both prognostic and predictive effects) performed at least as well as explicit outcome centering. This may seem surprising from a theoretical point of view, because a nuisance parameter is dealt with in two completely different ways. Even more interesting is the overall strong performance of a variant employing both principles at the same time: The $\text{mob}(\hat{W}, \hat{Y})$ forest is grown on centered outcomes and treatments and, additionally, also splits nodes with respect to both prognostic and predictive effects, leading to a performance at least as well as the best-performing competitor. Other aspects of tree and forest induction, such as exhaustive search vs. association tests for variable selection, internal stopping criteria based on sample-size constraints etc., did not explain much variability in performance.

Based on our current theoretical and empirical understanding of the elements of both model-based and causal random forests for HTE estimation, we can make the following recommendations for their application in practice, especially when the conditional mean of a numeric outcome captures all relevant aspects: Data from randomized trials can be analyzed by causal forests (with outcome centering and known treatment probability π for treatment indicator centering) or model-based forests (with or without outcome centering) under the intention-to-treat principle. Under potential confounding, it is important to accurately model treatment propensities, as in causal forests (with outcome and treatment centering). When combined with treatment centering, model-based forests will lead to approximately the same results. Additionally, centering the outcome may even offer a small performance gain compared to standard causal forests.

The empirical performances, reported in Section 4, coupled with established asymptotic results for causal random forests with treatment centering (Athey, Tibshirani and Wager (2019)) and the benign asymptotic behavior of other ingredients, such as transformation models (Hothorn, Möst and Bühlmann (2018)) or uniform convergence of polynomials in Bernstein form, suggest favorable asymptotic properties for special flavors of model-based forests. We leave the presentation of formal results to future work.

6.3. *Outlook.* The blending of model-based and causal forests discussed here seems to be a promising approach for HTE estimation beyond mean regression. Under potential confounding with binary, ordinal, count, or survival outcomes, it is easy to combine model-based forests with treatment centering ($\text{mob}(\hat{W})$), following the path outlined in Section 2.5. For example, for a binary outcome $Y \in \{0, 1\}$, a logistic regression-based causal forest can estimate models of the form

$$\text{logit}(\mathbb{P}(Y = 1 \mid X = \mathbf{x}, W = w)) = \mu(\mathbf{x}) + \tau(\mathbf{x})w.$$

The HTE $\tau(\mathbf{x})$ can then be interpreted as a covariate-dependent log-odds ratio. In practice, this model can be estimated by package *model4you*, with appropriate treatment centering being the only modification necessary (under the usual assumptions, of course). We leave an in-depth analysis and evaluation of this principle to future research, which should also address the question of how to achieve outcome centering in such models similar to $\text{mob}(\hat{Y}, \hat{W})$.

Finally, going beyond these recommendations and insights, our results are interesting from two further perspectives. First, the empirical application to postpartum blood loss in Section 5 has shown that blended model-based causal forests can be tailored to specific setups by adapting the underlying loss function. Second, we empirically demonstrated that two independent implementations of random forests for HTE estimation performed akin in comparable settings. This form of external software validation is important in its own right because the underlying algorithms and implementations are rather complex, and external validity can only be assessed with the help of an independent implementation. In case of *grf* and *model4you*, past, current, and future users of these software packages can have higher confidence in HTEs estimated using either package.

7. Computational details. All computations were performed using R version 4.1.1 (R Core Team (2021)) with the following add-on packages: *grf* (Tibshirani et al. (2021)), *model4you* (Seibold, Zeileis and Hothorn (2021)), *trtf* (Hothorn (2021)), and *partykit* (Hothorn and Zeileis (2015, 2021b)).

In all empirical experiments, both causal forests and all variants of model-based forests were grown with $M = 500$ trees (`model4you::pmforest` default) with minimum node size of `node = 14`, number of chosen variables per split `mtry = P`, and subsampling (the latter two being `causal_forest` defaults for $P = 10, 20$). We chose a minimum node size of 14 because the default of `partykit::ctree_control` (which *model4you* is based on) is 7, but we require this minimum node size for each of the two treatment groups. For adaptive forests 50 % of data were used to build each tree, and for honest forests subsamples were further cut in half (25 % to determine splits, 25 % for estimation, all *grf* defaults). To implement local centering of W in case of randomized data for causal forests, we set `W.hat` to 0.5 within `grf::causal_forest`.

We used the transformation forest implementation of the *trtf* package (Hothorn (2021), Hothorn and Zeileis (2021a)) for fitting the transformation-based forest in Section 5.

Ratios and confidence intervals, presented in Table 2 and Table S.1 (Supplementary Material A, Dandl et al. (2024a)), were computed by generalized linear mixed models fitted by the *glmmTMB* package (Brooks et al. (2021)) and post hoc inference was performed by the *multcomp* package (Hothorn, Bretz and Westfall (2021)).

We implemented all study settings in a dedicated R package called *htesim*. We also included the code and performance results of the empirical study as well as the code and dataset on postpartum blood loss. This should facilitate full reproducibility of all findings in this paper. The package is available as a supplement to this paper (Supplementary Material B, Dandl et al. (2024b)) and is also published on Github: <https://github.com/dandls/htesim>.

Acknowledgments. We thank the Handling Editor, Dr. Griffin, and two anonymous referees for helpful criticism of the initial manuscript.

Funding. The work of Susanne Dandl was funded by the Munich Center for Machine Learning (MCML). Torsten Hothorn received funding from the Swiss National Science Foundation, Grant No. 200021_184603, Horizon 2020 Research and Innovation Programme of the European Union under grant agreement number 681094 and is supported by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number 15.0137.

SUPPLEMENTARY MATERIAL

A. Details and additional results (DOI: [10.1214/23-AOAS1799SUPPA](https://doi.org/10.1214/23-AOAS1799SUPPA); .pdf). The document provides details about the cut-point selection of model-based forests and causal forests, the comparative results of adaptive and honest forests for the simulations study of Section 4,

and a sensitivity analysis on the number of chosen variables per split for the model-based transformation forest of Section 5.

B. Code and data (DOI: [10.1214/23-AOAS1799SUPPB](https://doi.org/10.1214/23-AOAS1799SUPPB); .zip). The tar.gz file contains the R package *htesim* to replicate the study settings. It also includes the code and performance results of Section 4 as well as the code and dataset of Section 5.

REFERENCES

- AKAZAWA, M., HASHIMOTO, K., KATSUHIKO, N. and KANAME, Y. (2021). Machine learning approach for the prediction of postpartum hemorrhage in vaginal birth. *Sci. Rep.* **11** 22620. <https://doi.org/10.1038/s41598-021-02198-y>
- ANTOINE, C. and YOUNG, B. K. (2021). Cesarean section one hundred years 1920–2020: The good, the bad and the ugly. *J. Perinat. Med.* **49** 5–16. <https://doi.org/10.1515/jpm-2020-0305>
- ATHEY, S. and IMBENS, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proc. Natl. Acad. Sci. USA* **113** 7353–7360. [MR3531135 https://doi.org/10.1073/pnas.1510489113](https://doi.org/10.1073/pnas.1510489113)
- ATHEY, S., TIBSHIRANI, J. and WAGER, S. (2019). Generalized random forests. *Ann. Statist.* **47** 1148–1178. [MR3909963 https://doi.org/10.1214/18-AOS1709](https://doi.org/10.1214/18-AOS1709)
- ATHEY, S. and WAGER, S. (2019). Estimating treatment effects with causal forests: An application. *Obs. Stud.* **5** 37–51. <https://doi.org/10.1353/obs.2019.0001>
- BREIMAN, L. (2001). Random forests. *Mach. Learn.* **45** 5–32. <https://doi.org/10.1023/a:1010933404324>
- BROOKS, M., BOLKER, B., KRISTENSEN, K., MAECHLER, M., MAGNUSSON, A., SKAUG, H., NIELSEN, A., BERG, C. and VAN BENTHAM, K. (2021). *glmmTMB*: Generalized linear mixed models using template model builder. R package version 1.1.2. Available at <https://CRAN.R-project.org/package=glmmTMB>.
- BURI, M. and HOTHORN, T. (2020). Model-based random forests for ordinal regression. *Int. J. Biostat.* **16** 20190063. <https://doi.org/10.1515/ijb-2019-0063>
- CHERNOZHUKOV, V., CHETVERIKOV, D., DEMIRER, M., DUFLO, E., HANSEN, C., NEWEY, W. and ROBINS, J. (2018). Double/debiased machine learning for treatment and structural parameters. *Econom. J.* **21** C1–C68. [MR3769544 https://doi.org/10.1111/ectj.12097](https://doi.org/10.1111/ectj.12097)
- CHIPMAN, H. A., GEORGE, E. I. and MCCULLOCH, R. E. (2010). BART: Bayesian additive regression trees. *Ann. Appl. Stat.* **4** 266–298. [MR2758172 https://doi.org/10.1214/09-AOAS285](https://doi.org/10.1214/09-AOAS285)
- CUI, Y., KOSOROK, M. R., SVERDRUP, E., WAGER, S. and ZHU, R. (2023). Estimating heterogeneous treatment effects with right-censored data via causal survival forests. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **85** 179–211. <https://doi.org/10.1093/jrsssb/qkac001>
- DANDL, S., HASLINGER, C., HOTHORN, T., SEIBOLD, H., SVERDRUP, E., WAGER, S. and ZEILEIS, A. (2024a). Details and additional results. Supplement A to “What makes forest-based heterogeneous treatment effect estimators work?” <https://doi.org/10.1214/23-AOAS1799SUPPA>
- DANDL, S., HASLINGER, C., HOTHORN, T., SEIBOLD, H., SVERDRUP, E., WAGER, S. and ZEILEIS, A. (2024b). Code and data. Supplement B to “What makes forest-based heterogeneous treatment effect estimators work?” <https://doi.org/10.1214/23-AOAS1799SUPPB>
- DASGUPTA, A., SZYMCAK, S., MOORE, J. H., BAILEY-WILSON, J. E. and MALLEY, J. D. (2014). Risk estimation using probability machines. *BioData Min.* **7** 2. <https://doi.org/10.1186/1756-0381-7-2>
- ENDE, H. B. (2022). Risk assessment tools to predict postpartum hemorrhage. *Bailliere's Best Pract. Res., Clin. Anaesthesiol.* **36** 341–348. <https://doi.org/10.1016/j.bpa.2022.08.003>
- ERICKSON, E. N. and CARLSON, N. S. (2020). Predicting postpartum hemorrhage after low-risk vaginal birth by labor characteristics and oxytocin administration. *J. Obstet. Gynecol. Neonatal Nurs.* **49** 549–563. <https://doi.org/10.1016/j.jogn.2020.08.005>
- FAROUKI, R. T. (2012). The Bernstein polynomial basis: A centennial retrospective. *Comput. Aided Geom. Design* **29** 379–419. [MR2921860 https://doi.org/10.1016/j.cagd.2012.03.001](https://doi.org/10.1016/j.cagd.2012.03.001)
- FENG, P., ZHOU, X.-H., ZOU, Q.-M., FAN, M.-Y. and LI, X.-S. (2012). Generalized propensity score for estimating the average treatment effect of multiple treatments. *Stat. Med.* **31** 681–697. [MR2900870 https://doi.org/10.1002/sim.4168](https://doi.org/10.1002/sim.4168)
- FOKKEMA, M., SMITS, N., ZEILEIS, A., HOTHORN, T. and KELDERMAN, H. (2018). Detecting treatment-subgroup interactions in clustered data with generalized linear mixed-effects model trees. *Behav. Res. Methods* **50** 2016–2034. <https://doi.org/10.3758/s13428-017-0971-x>
- FOSTER, J. C., TAYLOR, J. M. G. and RUBERG, S. J. (2011). Subgroup identification from randomized clinical trial data. *Stat. Med.* **30** 2867–2880. [MR2844689 https://doi.org/10.1002/sim.4322](https://doi.org/10.1002/sim.4322)
- HASLINGER, C., KORTE, W., HOTHORN, T., BRUN, R., GREENBERG, C. and ZIMMERMANN, R. (2020). The impact of prepartum factor XIII activity on postpartum blood loss. *J. Thromb. Haemost.* **18** 1310–1319. <https://doi.org/10.1111/jth.14795>

- HILL, J. L. (2011). Bayesian nonparametric modeling for causal inference. *J. Comput. Graph. Statist.* **20** 217–240. MR2816546 <https://doi.org/10.1198/jcgs.2010.08162>
- HOTHORN, T. (2021). *trtf*: Transformation trees and forests. R package version 0.3-8. Available at <https://CRAN.R-project.org/package=trtf>.
- HOTHORN, T., BRETZ, F. and WESTFALL, P. (2021). *multcomp*: Simultaneous inference in general parametric models. R package version 1.4-17. Available at <https://CRAN.R-project.org/package=multcomp>.
- HOTHORN, T., HORNIK, K. and ZEILEIS, A. (2006). Unbiased recursive partitioning: A conditional inference framework. *J. Comput. Graph. Statist.* **15** 651–674. MR2291267 <https://doi.org/10.1198/106186006X133933>
- HOTHORN, T., LAUSEN, B., BENNER, A. and RADESPIEL-TRÖGER, M. (2004). Bagging survival trees. *Stat. Med.* **23** 77–91. <https://doi.org/10.1002/sim.1593>
- HOTHORN, T., MÖST, L. and BÜHLMANN, P. (2018). Most likely transformations. *Scand. J. Stat.* **45** 110–134. MR3764288 <https://doi.org/10.1111/sjos.12291>
- HOTHORN, T. and ZEILEIS, A. (2015). *partykit*: A modular toolkit for recursive partytioning in R. *J. Mach. Learn. Res.* **16** 3905–3909. MR3450556
- HOTHORN, T. and ZEILEIS, A. (2021a). Predictive distribution modeling using transformation forests. *J. Comput. Graph. Statist.* **30** 1181–1196. MR4356613 <https://doi.org/10.1080/10618600.2021.1872581>
- HOTHORN, T. and ZEILEIS, A. (2021b). *partykit*: A toolkit for recursive partytioning. R package version 1.2-15. Available at <https://CRAN.R-project.org/package=partykit>.
- IMBENS, G. W. and RUBIN, D. B. (2015). *Causal Inference—For Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge Univ. Press, New York. MR3309951 <https://doi.org/10.1017/CBO9781139025751>
- ISHWARAN, H., KOGALUR, U. B., BLACKSTONE, E. H. and LAUER, M. S. (2008). Random survival forests. *Ann. Appl. Stat.* **2** 841–860. MR2516796 <https://doi.org/10.1214/08-AOAS169>
- ISHWARAN, H. and MALLEY, J. D. (2014). Synthetic learning machines. *BioData Min.* **7** 28. <https://doi.org/10.1186/s13040-014-0028-y>
- KAHR, M. K., BRUN, R., ZIMMERMANN, R., FRANKE, D. and HASLINGER, C. (2018). Validation of a quantitative system for real-time measurement of postpartum blood loss. *Arch. Gynecol. Obstet.* **298** 1071–1077. <https://doi.org/10.1007/s00404-018-4896-0>
- KAWAKITA, T., MOKHTARI, N., HUANG, J. C. and LANDY, H. J. (2019). Evaluation of risk-assessment tools for severe postpartum hemorrhage in women undergoing Cesarean delivery. *Obstet. Gynecol.* **134** 1308–1316. <https://doi.org/10.1097/AOG.0000000000003574>
- KOREPANOVA, N., SEIBOLD, H., STEFFEN, V. and HOTHORN, T. (2020). Survival forests under test: Impact of the proportional hazards assumption on prognostic and predictive forests for amyotrophic lateral sclerosis survival. *Stat. Methods Med. Res.* **29** 1403–1419. MR4097153 <https://doi.org/10.1177/0962280219862586>
- KÜNZEL, S. R., SEKHON, J. S., BICKEL, P. J. and YU, B. (2019). Metalearners for estimating heterogeneous treatment effects using machine learning. *Proc. Natl. Acad. Sci. USA* **116** 4156–4165. <https://doi.org/10.1073/pnas.1804597116>
- LIN, Y. and JEON, Y. (2006). Random forests and adaptive nearest neighbors. *J. Amer. Statist. Assoc.* **101** 578–590. MR2256176 <https://doi.org/10.1198/016214505000001230>
- LU, M., SADIQ, S., FEASTER, D. J. and ISHWARAN, H. (2018). Estimating individual treatment effect in observational data using random forest methods. *J. Comput. Graph. Statist.* **27** 209–219. MR3788313 <https://doi.org/10.1080/10618600.2017.1356325>
- MACDORMAN, M. F., DECLERCQ, E., CABRAL, H. and MORTON, C. (2016). Recent increases in the U.S. maternal mortality rate: Disentangling trends from measurement issues. *Obstet. Gynecol.* **128** 447–455. <https://doi.org/10.1097/AOG.0000000000001556>
- MAYER, I., SVERDRUP, E., GAUSS, T., MOYER, J.-D., WAGER, S. and JOSSE, J. (2020). Doubly robust treatment effect estimation with missing attributes. *Ann. Appl. Stat.* **14** 1409–1431. MR4152139 <https://doi.org/10.1214/20-AOAS1356>
- MEINSHAUSEN, N. (2006). Quantile regression forests. *J. Mach. Learn. Res.* **7** 983–999. MR2274394
- NIE, X. and WAGER, S. (2021). Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika* **108** 299–319. MR4259133 <https://doi.org/10.1093/biomet/asaa076>
- POWERS, S., QIAN, J., JUNG, K., SCHULER, A., SHAH, N. H., HASTIE, T. and TIBSHIRANI, R. (2018). Some methods for heterogeneous treatment effect estimation in high dimensions. *Stat. Med.* **37** 1767–1787. MR3799840 <https://doi.org/10.1002/sim.7623>
- PROBST, P., BOULESTEIX, A.-L. and BISCHL, B. (2019). Tunability: Importance of hyperparameters of machine learning algorithms. *J. Mach. Learn. Res.* **20** Paper No. 53, 32 pp. MR3948093
- PROBST, P., WRIGHT, M. N. and BOULESTEIX, A.-L. (2019). Hyperparameters and tuning strategies for random forest. *WIREs Data Min. Knowl. Discov.* **9** e1301. <https://doi.org/10.1002/widm.1301>
- R CORE TEAM (2021). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. Available at <https://www.R-project.org>.

- ROBINSON, P. M. (1988). Root- N -consistent semiparametric regression. *Econometrica* **56** 931–954. MR0951762 <https://doi.org/10.2307/1912705>
- ROSENBAUM, P. R. and RUBIN, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika* **70** 41–55. MR0742974 <https://doi.org/10.1093/biomet/70.1.41>
- SAY, L., CHOU, D., GEMMILL, A., TUNÇALP, O., MOLLER, A.-B., DANIELS, J., GÜLMEZOĞLU, A. M., TEMMERMAN, M. and ALKEMA, L. (2014). Global causes of maternal death: A WHO systematic analysis. *Lancet Glob. Health* **2** e323–e333. [https://doi.org/10.1016/S2214-109X\(14\)70227-X](https://doi.org/10.1016/S2214-109X(14)70227-X)
- SCHLOSSER, L., HOTHORN, T., STAUFFER, R. and ZEILEIS, A. (2019). Distributional regression forests for probabilistic precipitation forecasting in complex terrain. *Ann. Appl. Stat.* **13** 1564–1589. MR4019150 <https://doi.org/10.1214/19-AOAS1247>
- SCHNELL, P., TANG, Q., MÜLLER, P. and CARLIN, B. P. (2017). Subgroup inference for multiple treatments and multiple endpoints in an Alzheimer’s disease treatment trial. *Ann. Appl. Stat.* **11** 949–966. MR3693553 <https://doi.org/10.1214/17-AOAS1024>
- SEIBOLD, H., ZEILEIS, A. and HOTHORN, T. (2016). Model-based recursive partitioning for subgroup analyses. *Int. J. Biostat.* **12** 45–63. MR3505686 <https://doi.org/10.1515/ijb-2015-0032>
- SEIBOLD, H., ZEILEIS, A. and HOTHORN, T. (2018). Individual treatment effect prediction for amyotrophic lateral sclerosis patients. *Stat. Methods Med. Res.* **27** 3104–3125. MR3855657 <https://doi.org/10.1177/0962280217693034>
- SEIBOLD, H., ZEILEIS, A. and HOTHORN, T. (2019). *model4you*: An R package for personalised treatment effect estimation. *J. Open Res. Softw.* **7** 1–6. <https://doi.org/10.5334/jors.219>
- SEIBOLD, H., ZEILEIS, A. and HOTHORN, T. (2021). *model4you*: Stratified and personalised models based on model-based trees and forests. R package version 0.9-7. Available at <https://CRAN.R-project.org/package=model4you>.
- STARLING, J. E., MURRAY, J. S., LOHR, P. A., AIKEN, A. R. A., CARVALHO, C. M. and SCOTT, J. G. (2021). Targeted smooth Bayesian causal forests: An analysis of heterogeneous treatment effects for simultaneous vs. interval medical abortion regimens over gestation. *Ann. Appl. Stat.* **15** 1194–1219. MR4316646 <https://doi.org/10.1214/20-aoas1438>
- TANG, F. and ISHWARAN, H. (2017). Random forest missing data algorithms. *Stat. Anal. Data Min.* **10** 363–377. MR3733611 <https://doi.org/10.1002/sam.11348>
- TIBSHIRANI, J., ATHEY, S., SVERDRUP, E. and WAGER, S. (2021). *grf*: Generalized random forests. R package version 2.0.2. Available at <https://CRAN.R-project.org/package=grf>.
- VENKATESH, K. K., STRAUSS, R. A., GROTEGUT, C. A., HEINE, R. P., CHESCHEIR, N. C., STRINGER, J. S. A., STAMILIO, D. M., MENARD, K. M. and JELOVSEK, J. E. (2020). Machine learning and statistical models to predict postpartum hemorrhage. *Obstet. Gynecol.* **135** 935–944. <https://doi.org/10.1097/AOG.0000000000003759>
- WAGER, S. and ATHEY, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *J. Amer. Statist. Assoc.* **113** 1228–1242. MR3862353 <https://doi.org/10.1080/01621459.2017.1319839>
- WHO (2012). WHO recommendations for the prevention and treatment of postpartum haemorrhage. World Health Organization, Geneva, Switzerland.
- ZANUTTO, E., LU, B. and HORNIK, R. (2005). Using propensity score subclassification for multiple treatment doses to evaluate a national antidrug media campaign. *J. Educ. Behav. Stat.* **30** 59–73. <https://doi.org/10.3102/10769986030001059>
- ZEILEIS, A., HOTHORN, T. and HORNIK, K. (2008). Model-based recursive partitioning. *J. Comput. Graph. Statist.* **17** 492–514. MR2439970 <https://doi.org/10.1198/106186008X319331>