



Published in final edited form as:

Stat Med. 2021 September 20; 40(21): 4691–4713. doi:10.1002/sim.9090.

Estimating heterogeneous survival treatment effect in observational data using machine learning

Liangyuan Hu^{1,2}, Jiayi Ji¹, Fan Li^{3,4}

¹Department of Population Health Science and Policy, Icahn School of Medicine at Mount Sinai, New York, NY

²Department of Biostatistics and Epidemiology, Rutgers School of Public Health, Piscataway, NJ

³Department of Biostatistics, Yale University School of Public Health, New Haven, Connecticut

⁴Center for Methods in Implementation and Prevention Science, Yale University School of Public Health, New Haven, Connecticut

Abstract

Methods for estimating heterogeneous treatment effect in observational data have largely focused on continuous or binary outcomes, and have been relatively less vetted with survival outcomes. Using flexible machine learning methods in the counterfactual framework is a promising approach to address challenges due to complex individual characteristics, to which treatments need to be tailored. To evaluate the operating characteristics of recent survival machine learning methods for the estimation of treatment effect heterogeneity and inform better practice, we carry out a comprehensive simulation study presenting a wide range of settings describing confounded heterogeneous survival treatment effects and varying degrees of covariate overlap. Our results suggest that the nonparametric Bayesian Additive Regression Trees within the framework of accelerated failure time model (AFT-BART-NP) consistently yields the best performance, in terms of bias, precision, and expected regret. Moreover, the credible interval estimators from AFT-BART-NP provide close to nominal frequentist coverage for the individual survival treatment effect when the covariate overlap is at least moderate. Including a nonparametrically estimated propensity score as an additional fixed covariate in the AFT-BART-NP model formulation can further improve its efficiency and frequentist coverage. Finally, we demonstrate the application of flexible causal machine learning estimators through a comprehensive case study examining the heterogeneous survival effects of two radiotherapy approaches for localized high-risk prostate cancer.

Keywords

Bayesian additive regression trees; causal inference; machine learning; observational studies; survival treatment effect heterogeneity

Correspondence Liangyuan Hu, 683 Hoes Lane West, Piscataway, NJ 08854. liangyuan.hu@rutgers.edu.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

1 | INTRODUCTION

Observational data sources are expanding to include large registries and electronic health records, providing new opportunities to obtain comparative effectiveness evidence among broader target patient populations. Although traditional methods for observational studies focus on estimating the average treatment effect in a fixed target population, the rise of personalized healthcare and the increasing complexity of observational data have generated a more patient-centric view for drawing causal inferences about treatment effect. As a result, treatments need to be tailored to the covariate distribution of the target population, which necessitates the evaluation of *treatment effect heterogeneity* (TEH). If TEH presents, there exists a partition of the population into subgroups across which the treatment effect differs. Because the assessment of TEH requires addressing challenges due to complex and oftentimes high-dimensional individual characteristics, using flexible modeling techniques such as machine learning methods within the counterfactual framework is a promising approach and has received increasing attention in the statistical literature. For example, Foster et al¹ studied the virtual twins approach that first uses random forests to estimate the individual treatment effect (ITE). They obtained the difference between the predicted counterfactual outcomes under different treatment condition for each unit, and then regressed the predicted ITEs on covariates to identify subgroups with a treatment effect that is considerably different than the average treatment effect. Wager and Athey² developed the causal forests for ITE estimation. Lu et al³ compared several random forests methods, including Breiman forests,^{4,5} causal forests and synthetic forests,⁶ for predicting ITE and found that the synthetic RF had the best performance. Hill⁷ adapted the Bayesian Additive Regression Trees (BART) for causal inference and described methods for visualizing the modes of the predicted ITEs to gain insights into the underlying number of subgroups. Anoke et al⁸ a priori set the number of subgroups and group observations based on percentiles of the empirical distribution of either the ITEs or propensity scores and estimated a treatment effect within each subgroup. They compared several machine learning methods including BART, gradient boosted models and the facilitating score method, for the estimation of TEH and found BART performed the best in their simulation studies.

While these previous studies focus either on continuous or binary outcomes, time-to-event outcomes are of primary interest in many biomedical studies.^{9,10} Besides the average survival treatment effect, it is also important to understand whether there is survival TEH in order to make informed personalized treatment decisions.¹¹ The detection of survival TEH, however, may be complicated by the right censoring (eg, due to loss-to-follow-up) as well as challenges in precisely relating survival curves to individual covariates. Extending BART to survival outcomes, Henderson et al¹² developed a nonparametric accelerated failure time (AFT) model to flexibly capture the relationship between covariates and the failure times for the estimation of ITE. Tian et al¹³ approached the TEH problem by modeling interactions between a treatment and a large number of covariates. Both methods were designed for analyzing data from randomized clinical trials, with little empirical evidence within the context of observational studies. Shen et al¹⁴ used random forests with weighted bootstrap to estimate the optimal personalized treatment strategy; however, identification of TEH was not the focus of their paper. Cui et al¹⁵ extended the causal forests of Wager and Athey²

for ITE estimation to accommodate right censoring, and used bootstrap for inference. Tabib and Larocque¹⁶ also proposed a new random forests splitting rule to partition the survival data with the goal of maximizing the heterogeneity of the ITE estimates at the left and right nodes from a split. Neither of the last two methods has made software publicly available.

Despite burgeoning literature in causal machine learning methods for survival data, their comparative performance for estimating the survival TEH remains underexplored. The practical implementation of machine learning methods for detecting possible survival TEH also merits additional study. To address these gaps and facilitate better use of machine learning for estimating survival TEH, we carry out a series of simulations to investigate the operating characteristics of nine state-of-the-art machine learning techniques, including

1. nonparametric BART model within the AFT framework (AFT-BART-NP);¹²
2. a variant of AFT-BART-NP where the estimated propensity score is included as a covariate (AFT-BART-NP-PS);
3. semiparametric BART model within the AFT framework (AFT-BART-SP);¹²
4. random survival forests (RSF);¹⁷
5. Cox regression based on deep neural network (DeepSurv);¹⁸
6. deep learning with doubly robust censoring unbiased loss function (DR-DL);¹⁹
7. deep learning with Buckley-James censoring unbiased loss function (DL-BJ);^{19,20}
8. targeted survival heterogeneous effect estimation (TSHEE);²¹
9. generalized additive proportional hazards model (GAPH).²²

These techniques are representative of the recent machine learning advancement specifically designed for estimating individual-specific survival functions with right-censored failure time data; details of each method are reviewed in Section 2. In particular, recent work of Dorie et al²³ and Hahn et al²⁴ suggests that inclusion of an estimated propensity score in the BART model can improve the estimation of causal effects with noncensored outcomes. To investigate whether the utility of this strategy is transferable to censored survival outcomes, we considered a version of the nonparametric AFT BART model, AFT-BART-NP-PS, where the estimated propensity score was included as a covariate in model formulation. We contributed sets of simulations representing a variety of causal inference settings, and evaluated the empirical performance of these machine learning methods, with emphasis on new insights that can be gained, relative to the average survival treatment effects. We also compared these machine learning methods with two commonly used parametric and semiparametric regression methods, the AFT and Cox proportional hazards model, and elucidated the ramifications of potentially restrictive parametric model assumptions. Finally, to improve our understanding of the survival effect under two radiotherapy approaches for treating high-risk localized prostate cancer, we applied these machine learning methods to an observational dataset drawn from the national cancer registry database (NCDB). We focused on 7330 prostate cancer patients¹⁰ who received either external beam radiotherapy (EBRT) combined with androgen deprivation (AD), EBRT + AD, or EBRT plus brachytherapy with

or without AD, EBRT + brachy ± AD, with at least 7920 cGy EBRT dose, and explored the possible survival TEH in relation to individual covariate information.

The remainder of the article is organized as follows. Section 2 introduces causal assumptions and estimand of interest; it further provides a concise overview of each machine learning method considered. In Section 3, we describe the data generating processes and parameter constellations in our simulation study. Section 4 summarizes the comparative findings from simulation studies. In Section 5, we provide a case study where we estimated the heterogeneous survival treatment effects of two radiotherapy approaches for high-risk localized prostate cancer patients. Section 6 concludes with a discussion.

2 | STATISTICAL METHODS

2.1 | Notation and causal estimand

Consider an observational study with two treatment conditions and n observations, $i = 1, \dots, n$. Let Z_i be a binary treatment variable with $Z_i = 1$ for treatment and $Z_i = 0$ for control, and X_i be a vector of confounders measured before treatment assignment. We proceeded in the counterfactual framework, and defined the causal effect for individual i as a comparison of $T_i(1)$ and $T_i(0)$, representing the counterfactual survival times that would be observed under treatment and control, respectively. We similarly define $C_i(1)$ and $C_i(0)$ as the counterfactual censoring times under each treatment condition. We further write T_i as the observed survival time and C_i the observed censoring time. The observed outcome consists of $Y_i = \min(T_i, C_i)$ and censoring indicator $\delta_i = \mathbb{I}(T_i < C_i)$. Throughout, we maintained the standard assumptions for drawing causal inference with observational survival data:²⁵

(A1) Consistency: The observed failure time and censoring time $T_i = Z_i T_i(1) + (1 - Z_i) T_i(0)$ and $C_i = Z_i C_i(1) + (1 - Z_i) C_i(0)$;

(A2) Weak unconfoundedness: $T_i(z) \perp\!\!\!\perp Z_i | X_i$ for $Z_i = z, z \in \{0, 1\}$;

(A3) Positivity: the propensity score for treatment assignment $e(X_i) = P(Z_i = 1 | X_i)$ is bounded away from 0 and 1;

(A4) Covariate-dependent censoring: $T_i(z) \perp\!\!\!\perp C_i(z) | \{X_i, Z_i\}$, for $Z_i = z, z \in \{0, 1\}$.

Assumption (A1) maps the counterfactual outcomes and treatment assignment to the observed outcomes. It allows us to write $Y_i = Z_i Y_i(1) + (1 - Z_i) Y_i(0)$, where $Y_i(z) = \min\{T_i(z), C_i(z)\}$. Likewise, $\delta_i = Z_i \delta_i(1) + (1 - Z_i) \delta_i(0)$, where $\delta_i(z) = \mathbb{I}(T_i(z) < C_i(z))$. Assumption (A2) is referred to as the “no unmeasured confounders” assumption. Because this assumption is generally not testable without additional data, sensitivity analysis may be required to help interpret the causal effect estimates.²⁶ Assumption (A3) is also referred to as the overlap assumption,²⁷ which requires that the treatment assignment is not deterministic within each strata formed by the covariates. This assumption can be directly assessed by visualizing the distribution of estimated propensity scores. Finally, Assumption (A4) states that the counterfactual survival time is independent of the counterfactual censoring time given the pretreatment covariates and treatment. This condition directly

implies $T_i \perp\!\!\!\perp C_i \{X_i, Z_i\}$, and is akin to the “independent censoring” assumption in the traditional survival analysis literature.

For survival outcomes, the population-level survival treatment effect can be characterized by the survival quantile effect,²⁸

$$\omega^q = \theta^q(1) - \theta^q(0) = \sup \{t: P(T(1) \leq t) \leq q\} - \sup \{t: P(T(0) \leq t) \leq q\}, \quad (1)$$

where q is a prespecified number between 0 and 1. The estimand (1) compares the q th quantile of the marginal counterfactual survival distribution under the treated and the control condition. Implicit in our notation, we assume that the treatment strategy includes a hypothetical intervention that removes censoring until the end of follow-up. It has been suggested that the survival quantile effect should be used when the proportional hazards assumption does not hold.^{29,30} In this article, we focus on $q = 1/2$ and $\omega^q = 1/2$ corresponds to the causal contrast in median survival times, which is a natural summary statistics of the entire causal survival curve and of clinical relevance.^{28,31} Extending the population-level treatment effect (1), we define the *individual survival treatment effect* (ISTE) as

$$\omega^q = 1/2(x) = \theta^q = 1/2(1, x) - \theta^q = 1/2(0, x), \quad (2)$$

where $\theta^q = 1/2(z, x) = \sup \{t: P(T(z) \leq t \mid X = x) \leq 1/2\}$. In other words, $\theta^q = 1/2(z, x)$ represents the median survival time of the conditional counterfactual survival function $P(T(z) > t \mid X = x)$ and $\omega^q = 1/2(x)$ is alternatively called the *conditional survival treatment effect*. Because only the outcome, (Y_i, δ_i) , corresponding to the actual treatment assignment is observed, the ISTE as a function of the unknown counterfactuals is not directly identifiable without additional assumptions. However, under Assumptions (A1) to (A3), we have

$$P(T(z) > t \mid X = x) = P(T(z) > t \mid Z = z, X = x) = P(T > t \mid Z = z, X = x), \quad (3)$$

where $P(T > t \mid Z = 1, X = x)$ becomes the group-specific survival functions conditional on $X = x$ and is estimable from observed data. Specifically, the identification condition (3) allows us to estimate $\omega^q = 1/2(x)$ by fitting two survival outcome regression models, \mathcal{M}_0 and \mathcal{M}_1 , to two sets of the observed data $\{X_i, Y_i, \delta_i \mid Z_i = 0\}$ and $\{X_i, Y_i, \delta_i \mid Z_i = 1\}$. We then predict, for each data point, two individual-specific counterfactual survival curves conditional on individual covariate profile X_i , $\hat{P}(T_i(z) > t \mid X_i) = \hat{S}_{\mathcal{M}_z}(t, X_i)$ from model $\mathcal{M}_z, z = 0, 1$. The general form of the ISTE estimator is then given by

$$\hat{\omega}^q = 1/2(X_i) = \hat{\theta}_{\mathcal{M}_1}^q = 1/2(X_i) - \hat{\theta}_{\mathcal{M}_0}^q = 1/2(X_i), \quad (4)$$

where $\hat{S}_{\mathcal{M}_1}(\hat{\theta}_{\mathcal{M}_1}^q = 1/2(X_i), X_i) = \hat{S}_{\mathcal{M}_0}(\hat{\theta}_{\mathcal{M}_0}^q = 1/2(X_i), X_i) = 1/2$.

2.2 | Estimating individual-specific counterfactual survival curves

Because the estimation of the causal estimand (2) critically depends on the estimation of individual-specific counterfactual survival functions (or conditional survival functions),

$\hat{S}_{\mathcal{M}_z}(t, X)$ for $z = 0, 1$, below we briefly review the machine learning methods for estimating $\hat{S}_{\mathcal{M}_z}(t, X)$. For simplicity, we drop the subscript \mathcal{M}_z in the survival function, and only note here that the flexible outcome model will be fitted for each treatment group in a symmetric fashion, and the causal interpretation will be based on (3).

2.2.1 | Accelerated failure time BART model—The AFT-BART model relates the failure time T to the covariates through $\log T = f(X) + W$, where $f(X)$ is a sum-of-trees BART model, $f(X) = \sum_{j=1}^J g(X; \mathcal{T}_j, \mathcal{M}_j)$, and W is the residual term. The sum-of-trees model is fundamentally an additive model with multivariate components, which can more naturally incorporate interaction effects than generalized additive models and more easily incorporate additive effects than a single tree model.³² If W is assumed to follow a Gaussian distribution, the resultant model is semiparametric, termed as AFT-BART-SP. Henderson et al.¹² model the distribution of W as a location-mixture of Gaussian densities by using the centered Dirichlet process (CDP), leading to a Bayesian nonparametric specification, which we refer to as AFT-BART-NP. The individual trees \mathcal{T}_j and bottom nodes \mathcal{M}_j are tree model parameters. A regularization prior is placed on $\{\mathcal{T}_j, \mathcal{M}_j\}$ to allow each individual tree model to contribute only a small part to the overall fit and therefore avoids overfitting.³² Furthermore, prior distributions are assumed for the distributional parameters of the residual term W , including the normal variance components, or the location and scale parameters for the CDP assumption. The posterior distributions of AFT model parameters are computed using Markov chain Monte Carlo (MCMC)—Metropolis-within-Gibbs sampler for updating tree parameters and then Gibbs sampler for parameters related to the residual distribution. The unobserved survival times are imputed from a truncated normal distribution $N(\log Y_i, \infty)(f(X_i), \sigma^2)$ in each Gibbs iteration, where σ^2 is the variance or common scale of the residual term. The individual-specific survival curve can then be calculated by summarizing the posterior predictive distribution of

$$S^{\text{AFT-BART-SP}}(t, X_i) = 1 - \Phi\left(\frac{\log t - f(X_i)}{\sigma}\right)$$

for AFT-BART-SP, where $\Phi(\cdot)$ is the cumulative standard normal distribution and σ is the variance of the residual term. For AFT-BART-NP, the individual-specific survival curve is calculated by summarizing the posterior predictive distribution of

$$S^{\text{AFT-BART-NP}}(t, X_i) = 1 - \sum_{l=1}^L \pi_l \Phi\left(\frac{\log t - f(X_i) - \tau_l}{\sigma}\right),$$

where π_l is the mixture proportion of cluster l ($l = 1, \dots, L$), τ_l is the cluster-specific location parameter, σ is the common scale parameter. Additional technical details are available in Henderson et al.¹² A recent review of BART methods can also be found in Tan and Roy.³³

A direct implication from the weak unconfoundedness assumption (A2) is that $T_i(z) \perp\!\!\!\perp Z_i | \epsilon(X_i)$ for $Z_i = 0, 1$, suggesting that controlling for the one-dimensional propensity score

is sufficient to remove confounding in estimating the average causal effect.³⁴ For the purpose of estimating ITE with noncensored outcomes, Dorie et al²³ and Hahn et al²⁴ considered including an estimated propensity score as an additional covariate in BART for improved robustness and estimation efficiency. Assuming \hat{e}_i is an estimate of $e(X_i)$, we further consider a version of AFT-BART-NP that includes both X_i and \hat{e}_i in the model formulation. To mitigate concerns on model misspecification, we obtain the propensity score estimate, \hat{e}_i , nonparametrically from a Super Learner³⁵ and treat it as a fixed covariate in the subsequent AFT-BART-NP model formulation. The individual-specific survival curve from AFT-BART-NP-PS is then obtained by summarizing the posterior predictive distribution of

$$S^{\text{AFT-BART-NP-PS}}(t, X_i, \hat{e}_i) = 1 - \sum_{l=1}^{L^*} \pi_l \Phi\left(\frac{\log t - f(X_i, \hat{e}_i) - \tau_l}{\sigma}\right),$$

where π_l is the mixture proportion of cluster l ($l = 1, \dots, L^*$), τ_l is the cluster-specific location parameter, and σ is the common scale parameter.

2.2.2 | Random survival forests—The RSF approach extends Breiman random forests to right-censored survival data.¹⁷ The key idea of the RSF algorithm is to grow a binary survival tree \mathcal{H} and construct the ensemble cumulative hazard function (CHF). Each terminal node of the survival tree contains the observed survival times and the binary censoring information for individuals who fall in that terminal node $h \in \mathcal{H}$. Let $t_{1,h} < t_{2,h} < \dots, t_{N(h),h}$ be the $N(h)$ distinct event times. The CHF estimate for node h is the Nelson-Aalen estimator $\hat{\Lambda}_h = \sum_{l: t_{l,h} \leq t} d_{l,h} / R_{l,h}$, where $d_{l,h}$ and $R_{l,h}$ are the number of events and size of the risk set at time $t_{l,h}$. All individuals within node h then have the same CHF estimate. To compute the CHF for an individual with covariates X_i , one can drop X_i down the tree, and obtain the Nelson-Aalen estimator for X_i 's terminal node, $\Lambda(t | X_i) = \hat{\Lambda}_h(t)$ if $X_i \in h$. The out-of-bag (OOB) ensemble CHF for each individual, $\lambda^{\text{OOB}}(t | X_i)$, can be calculated by dropping OOB data down a survival tree built from in-bag data, recording i 's terminal node and its CHF and taking the average of all of the recorded CHFs. By a conversation-of-events principle,¹⁷ the individual-specific survival curve can be computed as one minus the estimated ensemble OOB CHF summed over survival times (event time or censoring time) until t , conditional on X_i , namely

$$S^{\text{RSF}}(t, X_i) = 1 - \sum_{\{j: T_j \leq t\}} \hat{\Lambda}^{\text{OOB}}(T_j | X_i).$$

2.2.3 | Targeted survival heterogeneous effect estimation—TSHEE was recently proposed based on the doubly robust estimation method, the targeted maximum likelihood estimation (TMLE), for survival causal inference.²¹ At the first step, a stylized counting process representation of the survival data is created, then the Super Learner³⁶ is applied to estimate, for each individual at each survival time point, the conditional hazard rate $h(t | X_i)$. The initial individual survival function is then derived via the probability chain rule,

$$S(t, X_i) = \prod_{j: t_j \leq t} q(t_j | X_i).$$

In the second step, the initial estimator is adjusted with conditional censoring probability $P(C > t | Z, X)$ and the propensity score $e(X)$ using the one-step TMLE.³⁷ To carry out the one-step TMLE, one replaces the univariate conditional hazard function $q(t|X)$ —a covariate in a logistic regression working (fluctuation) model for the failure event conditional hazard—with a high dimensional vector $\vec{q}(\bullet | X) = (q(t_1 | X), \dots, q(t_{\max} | X))$, where $\{t_1, \dots, t_{\max}\}$ corresponds to the set of ordered unique survival times. Fitting the high-dimensional logistic regression, the one-step TMLE updates in small steps locally along the working model to iteratively update the conditional survival function of failure event to form a targeted estimator $S^{\text{TSHEE}}(t, X_i)$. We refer to Zhu and Gallego²¹ for full technical details on TSHEE.

2.2.4 | Cox regression based on deep neural network—DeepSurv is a Cox regression tool based on deep neural network.¹⁸ It uses a deep feed-forward neural network to predict the effects of individual covariates on their hazard rate parameterized by the weights of the network. The input to the network is each individual's baseline covariates X_i . The network propagates the inputs through a number of hidden layers with weights ξ . The weights parameter ξ is optimized to minimize the average negative log partial likelihood subject to ℓ_2 regularization,

$$l^{\text{PL}}(\xi) = -\frac{1}{N_E} \sum_{i: \Delta_i = 1} \left[h_{\xi}(X_i) - \log \sum_{j \in R(Y_i)} \exp\{h_{\xi}(X_j)\} \right] + \kappa \|\xi\|_2^2,$$

where N_E is the number of observed events, Y_i 's are the unique observed survival times, $R(Y_i)$ is the risk set at time Y_i , $h_{\xi}(\cdot)$ is the log hazard function, and κ is L_2 regularization parameter. The hidden layers comprise fully connected layer of nodes, followed by a dropout layer. As a fully connected layer occupies most of the parameters, neurons develop codependency among each other during the training phase, which leads to overfitting of training data. By dropping out individual nodes, that is, neurons with certain probability during the training phase, a reduced network is retained to prevent overfitting. The output of the network estimates the log hazard function $\hat{h}_{\xi}(X)$ within the Cox model framework. To compute the individual-specific survival curve, we first calculate the Nelson-Aalen estimate of the baseline CHF using $\hat{h}_{\xi}(X_i)$ outputted by DeepSurv,

$$\hat{\Lambda}_0(t) = \sum_{i: Y_i \leq t} \frac{\Delta_i}{\sum_{j \in R(Y_i)} \exp\{\hat{h}_{\xi}(X_j)\}}.$$

We then relate the survival function to the CHF and estimate the individual-specific survival curve by

$$S^{\text{DeepSurv}}(t, X_i) = \exp\left[-\hat{\Lambda}_0(t) \exp\{\hat{h}_{\xi}(X_i)\}\right].$$

2.2.5 | Censoring unbiased deep learning survival model—A recent deep learning (DL) survival modeling technique builds upon the DL algorithms for uncensored data, but develops censoring unbiased loss functions to replace the unobserved full data loss, resulting in two DL algorithms for censored survival outcomes—one with doubly robust censoring unbiased loss (DR-CL) and one with Buckley-James censoring unbiased loss (BJ-DL).¹⁹ The doubly robust loss function consists of two terms, one based on uncensored outcomes and one constructed to inversely weight observed outcome information with functionals of the event or censoring probability. The loss function is doubly robust in the sense that it is a consistent estimator for the expected loss if one of the models for failure time T or censoring time C is correctly specified but not necessary both. A related estimator BJ-DL requires only modeling the conditional survival function for T for the term of censored outcomes in the loss function, which does not share the property of double robustness with the DR-CL estimator. A survival prediction model, $S(t, X)$, can be obtained by replacing the full data loss function with either the doubly robust or Buckley-James loss function and minimizing the Brier risk.¹⁹ Specifically, $S^{\text{DR-DL}}(t, X_i)$ minimizes the doubly robust censoring unbiased loss

$$\mathcal{L}^{\text{DR}}(t) = \frac{1}{n} \sum_{i=1}^n \frac{\delta_i(t)(I(Y_i \geq t) - S(t, X_i))^2}{\hat{G}(Y_i | X_i)} + \frac{1}{n} \sum_{i=1}^n \left\{ \frac{(1 - \delta_i(t)) \mathcal{K}_{L_2, t}(Y_i, X_i, S(t, X_i))}{\hat{G}(Y_i | X_i)} - \int_0^{Y_i} \frac{\mathcal{K}_{L_2, t}(Y_i, X_i, S(t, X_i)) d\hat{\Lambda}_G(u | X_i)}{\hat{G}(u | X_i)} \right\},$$

and $S^{\text{BJ-DL}}(t, X_i)$ minimizes the Buckley-James censoring unbiased loss

$$\mathcal{L}^{\text{BJ}}(t) = \frac{1}{n} \sum_{i=1}^n \delta_i(t)(I(Y_i \geq t) - S(t, X_i))^2 + (1 - \delta_i(t)) \mathcal{K}_{L_2, t}(C_i, X_i, S(t, X_i)),$$

where $\delta_i(t)$ is the event indicator at time t for individual i , $G(u|X)$ is the conditional survival function for C , $\mathcal{K}_{L_2, t}(u, X_i, S(t, X_i)) = E\left[(I(Y_i \geq t) - S(t, X_i))^2 | T \geq u, X_i\right]$ is the Brier risk, and $\hat{\Lambda}_G(u | X_i)$ is the conditional cumulative hazard function for C . We refer the reader to Steingrimsdottir and Morrison¹⁹ for additional technical details.

2.2.6 | Generalized additive proportional hazards model—In the absence of censoring, a flexible approach for outcome modeling and prediction is the generalized additive model (GAM).²² This approach has also been extended to accommodate survival data. Specifically, the generalized additive proportional hazards model (GAPH) model²² represents the hazard ratio function of the Cox proportional hazards model via a sum of smooth functions of covariates, and can be used to estimate the group-specific survival functions conditional on $X = x$. In our notation, the GAPH model is written as $\lambda(t; X) = \lambda_0(t) \exp\{\eta(X)\}$, where $\eta(X) = \sum_{j=1}^p f_j(X_j)$ is the sum of smooth functions of each confounder X_j . Estimation of model functions η proceeds by maximizing the log partial likelihood coupled with the second-order smoothness penalties for each function. In other words, the objective function is

$$l^{\text{PL}}(f_1, \dots, f_p) + \sum_{j=1}^p \lambda_j \int \{f_j''(x)\}^2 dx,$$

where l^{PH} generically denote the log partial likelihood, and λ_j is the regularization parameter (as $\lambda_j \rightarrow \infty$, $f_j(x)$ degenerates to the identity function). Similar to the DeepSurv, we can obtain the Nelson-Aalen estimate of the baseline CHF $\widehat{\Lambda}_0(t)$ and compute the individual-specific survival curve by

$$S^{\text{GAPH}}(t, X_i) = \exp\{-\widehat{\Lambda}_0(t)\exp[\widehat{\eta}(X_i)]\}.$$

2.2.7 | Parametric accelerated failure time model—A standard parametric model for analyzing survival data is the accelerated failure time (AFT) model. The parametric AFT model³⁸ postulates an explicit functional relationship between the covariates and the failure time T , $\log T = \mu + X^T \beta + \sigma W$, where β is a vector of coefficients describing the effects of covariates X on the log transformed survival time, μ and σ are intercept and scale parameters, respectively, W is the error term, for which a variety of distributions can be assumed. In our simulations, we assume a standard extreme value distribution and a normal distribution, for which the survival time has a Weibull distribution and a log-normal distribution, respectively. The parameters can be estimated by maximizing the likelihood function for censored data, and are denoted by $\widehat{\mu}$, $\widehat{\beta}$, $\widehat{\sigma}$. The individual-specific survival curve can then be calculated as

$$S^{\text{AFT-Lognormal}}(t, X_i) = 1 - \Phi\left(\frac{\log t - \widehat{\mu} - X_i^T \widehat{\beta}}{\widehat{\sigma}}\right),$$

and

$$S^{\text{AFT-Weibull}}(t, X_i) = \exp\left\{-\exp\left(\frac{\log t - \widehat{\mu} - X_i^T \widehat{\beta}}{\widehat{\sigma}}\right)\right\}.$$

A good review of parametric AFT model and the associated computational algorithms can be found in Wei.³⁹

2.2.8 | Cox proportional hazards model—A Cox proportional hazards model⁴⁰ relates the hazard function and the covariates X through $\lambda(t|X) = \lambda_0(t) \exp(X^T \gamma)$, where γ is a vector of parameters describing the multiplicative effects of pretreatment covariates on the hazard function. The model assumes that the hazard ratio is constant across the time span (the so-called proportional hazards assumption). The parameter γ can be estimated by maximizing the partial likelihood. Similar to the procedure described for DeepSurv, we can obtain the Nelson-Aalen estimate of the baseline CHF $\widehat{\Lambda}_0(t)$ and compute the individual-specific survival curve by

$$S^{\text{CoxPH}}(t, X_i) = \exp\left\{-\widehat{\Lambda}_0(t)\exp\left(X_i^T \widehat{\gamma}\right)\right\}.$$

Even though the Cox model includes a nonparametric specification of the baseline hazard, the performance of this approach critically depends on the parametric specification of the linear component $X_i^T \widehat{\gamma}$. When the true survival outcome surface involves complicated functional forms of the covariates, both the parametric AFT model and the Cox model may be subject to bias due to inaccurate model assumptions. We use these two conventional modeling techniques as a benchmark to quantify the bias that may occur in estimating the ISTE due to model misspecification.

3 | SIMULATION STUDIES

3.1 | Simulation design

We carried out a series of simulation studies to compare the above methods for estimating the individual counterfactual survival curves and hence the ISTE in the context of observational studies. Throughout we considered two levels of total sample size: $n = 500$ and $n = 5000$, representing small and large studies. We simulated $p = 10$ independent covariates, where X_1, \dots, X_5 were drawn from the standard $\mathcal{N}(0, 1)$, and X_6, \dots, X_{10} from Bernoulli(0.5). The treatment assignment was generated from $Z|X \sim \text{Bernoulli}(e(X))$, where the true propensity score follows

$$\text{logit}\{e(X)\} = \alpha_0 + \alpha_1 X_1 + \alpha_2 X_2 + \alpha_3 X_3 + \alpha_4 X_5 + \alpha_5 X_6 + \alpha_6 X_7 + \alpha_7 X_9 + \alpha_8 X_{10},$$

and $(\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_6, \alpha_7, \alpha_8) = (-0.1\psi, -0.9\psi, -0.3\psi, -0.1\psi, -0.2\psi, -0.4\psi, 0.5\psi)$. We chose $\psi \in \{1, 2.5, 5\}$ to represent strong, moderate, and weak levels of confounding, leading to strong, moderate, and weak covariate overlap.^{41,42} The distribution of the true propensity scores is shown in Figure 1. In particular, the degree of overlap determines the effective sample size for estimating population ATE,⁴³ but may have important implications for estimating ISTE. The intercept in the true propensity score model α_0 was chosen such that the proportion of treated remained approximately 0.5.

We simulated the true counterfactual survival times T from the Weibull distribution,

$$P(T(z) > t | X = x) = S(t | X = x, Z = z) = \exp\left[-\{d_z \exp(m_z(x)t)\}^\eta\right] \quad \forall z \in \{0, 1\}, \quad (5)$$

where d_z is the treatment-group-specific parameter, η is the shape parameter and $m_z(x)$ represents a generic functional form of covariates on survival time; specifications of these parameters will be discussed below. We considered the Weibull distribution because it has both the AFT and Cox proportional hazards representations, and therefore allows us to simulate scenarios where either one of these assumptions is valid. Using the inverse transform sampling, we generated counterfactual survival times by,

$$T(z) = \left\{ \frac{-\log U}{d_z \exp[m_z(X)]} \right\}^{1/\eta}, \quad \forall z \in \{0, 1\},$$

where $U \sim \text{Unif}(0, 1)$ was a random variable following a uniform distribution on the unit interval $[0, 1]$, $d_z = 1200$ for $z = 0$, and $d_z = 2000$ for $z = 1$, and η was the shape parameter used to induce nonproportional hazards. We considered three structures for $m_z(X)$, leading to four treatment effect heterogeneity settings (HS) with increasing complexities:

- i. $m_1(X) = -0.2 + 0.1/(1 + e^{-X_1}) - 0.8\sin(X_3) - 0.1X_5^2 - 0.3X_6 - 0.2X_7$, and $m_0(X) = 0.2 - 0.5X_1 - 0.8X_3 - 1.8X_5 - 0.9X_6 - 0.1X_7$,
- ii. $m_1(X) = -0.2 + 0.1/(1 + e^{-X_1}) - 0.8\sin(X_3) - 0.1X_5^2 - 0.3X_6 - 0.2X_7$, and $m_0(X) = -0.1 + 0.1X_1^2 - 0.2\sin(X_3) + 0.2/(1 + e^{-X_5}) + 0.2X_6 - 0.3X_7$,
- iii. $m_1(X) = 0.5 - 0.1/(1 + e^{-X_2}) + 0.1\sin(X_3) - 0.1X_4^2 + 0.2X_4 - 0.1X_5^2 - 0.3X_6$, and $m_0(X) = -0.1 + 0.1X_1^2 - 0.2\sin(X_3) + 0.2/(1 + e^{-X_5}) + 0.2X_6 - 0.3X_7$,
- iv. $m_1(X) = 0.5 - 0.1/(1 + e^{-X_2}) + 0.1\sin(X_3) - 0.1X_4^2 + 0.2X_4 - 0.1X_5^2 - 0.3X_6$, and $m_0(X) = -0.2 + 0.5\sin(\pi X_1 X_3) + 0.2/(1 + e^{-X_5}) + 0.2X_6 - 0.3X_7$.

Across all four outcome generating processes, X_1, X_3, X_5, X_6, X_7 were confounders that were related to both the treatment assignment and the potential survival outcome. In scenario (i), there is a nonlinear covariate effect only in the treated group. In scenario (ii), the nonlinear covariate effects are present in both treatment groups. Scenario (iii) is similar to scenario (ii) except that a nonoverlapping set of covariates are included in the treated and control groups; this additional complexity represents an additional source of TEH. To avoid concerns that these settings may favor tree-based models, in scenario (iv), we further increased the data complexity by introducing a correlation of 0.6 between X_2 and X_4 in $m_1(X)$ and by including a nonadditive term $\sin(\pi X_1 X_3)$ in $m_0(X)$. Finally, when simulating the counterfactual survival times, the parameter η was set to be 2 to satisfy the proportional hazards assumption. To induce nonproportional hazards, η was set to be $\exp\{0.7 - 1.8X_3 + 0.8X_7\}$ for $Z = 0$, and $\exp\{0.9 - 0.5X_1 + 0.5X_2\}$ for $Z = 1$. For simplicity, we generated the censoring time C independently from an exponential distribution with rate parameters 0.007 and 0.02, corresponding to two levels of censoring rate (CR), 20% and 60%. In Section 4.5, we additionally investigated the robustness of findings when censoring depends on observed covariates. An exploratory sensitivity analysis of the estimated ISTE to departure from the weak unconfoundedness assumption (A2) for all methods considered was conducted in Section 4.6. The simulated counterfactual survival curves are presented in Web Figure 1.

3.2 | Implementation

We used the `abart` function from the R package `bart` with the default settings to implement AFT-BART-SP, and used the R package `AFTrees` to implement AFT-BART-NP. We used 1100 draws with the first 100 discarded as burn-in. `DeepSurv` was

implemented in Theano with the Python package `Lasagne`. To fit RSF, we used the `randomForestSRC` R package with 1000 trees, `mtry = p/3` and a nodesize of 3. For the TSHEE method, we used the R code available at <https://github.com/EliotZhu/Targeted-Survival>. To implement DR-DL and BJ-DL, we used R codes provided at <https://github.com/jonsteingrimsson/CensoringDL>. The GAPH model was implemented using the `gam` function with `family=cox.ph` from the R package `mgcv`. Super Learner was implemented using the R package `SuperLearner` to estimate the propensity score \hat{e}_i for the method AFT-BART-NP-PS. Models in the Super Learner ensemble library included `BART`, `GAM` and `XGBoost`. To facilitate the application of these state-of-the-art machine learning methods to estimate ISTE, we provide example code to replicate our simulation studies in the GitHub page of the first author <https://github.com/liangyuanhu/ML-DL-TEH>.

3.3 | Performance metrics

We first assessed the overall precision in the estimation of heterogeneous effects (PEHE) for each method considered using

$$\text{PEHE} = \sqrt{\frac{1}{n} \sum_{i=1}^n \{\hat{\omega}^{q=1/2}(X_i) - \omega^{q=1/2}(X_i)\}^2},$$

where $\omega^{q=1/2}(X_i) = \theta^{q=1/2}(1, X_i) - \theta^{q=1/2}(0, X_i)$ is the difference between two counterfactual median survival times calculated from each data point, and $\hat{\omega}^{q=1/2}(X_i)$ is the estimated ISTE for individual i . This metric was also considered in Hill⁷ and Dorie et al,²³ and a smaller value of PEHE indicates better accuracy and is considered favorable. For computational considerations, we replicated each of our simulations independently $B = 250$ times for $n = 5000$ and $B = 1000$ times for $n = 500$. We summarized the mean and standard deviation (SD) of the PEHE across B replications.

To further evaluate the operating characteristics of each method in recovering treatment effect within subpopulations, we subclassified the individuals into $\mathcal{G} = \{G_1, \dots, G_K\}$ based on K quantiles of the true propensity score distribution, and calculated relative bias and root mean squared error (RMSE) for each method. This strategy was firstly considered in Lu et al³ in their simulations with noncensored outcomes, and adapted in our evaluation. We choose $K = 50$ to represent a large enough number of subclasses with adequate resolution. By investigating the performance of each method in distinct covariate regions with different proportions of treatment assignment and overlap, we can assess the robustness of each method to data sparsity. More concretely, given an ISTE estimator $\hat{\omega}^{q=1/2}(X)$ we define the relative bias (or percent bias) in subclass G_k by

$$\text{RelBias}(k) = \frac{E\{\hat{\omega}^{q=1/2}(X_i) - \omega^{q=1/2}(X_i) \mid X_i \in G_k\}}{E\{\omega^{q=1/2}(X_i) \mid X_i \in G_k\}}, \quad k = 1, \dots, K.$$

While Lu et al³ report the bias of each estimator on its original scale, we assess the bias on the relative scale with respect to the true treatment effect (which is constant across different

methods) so that the simulation results of different methods are more comparable even across different data generating processes. Under the same data generating process, however, the relative rankings of methods should be the same regardless of using the bias or relative bias, and the choice between these two metrics are immaterial.

We also define the root mean squared error (RMSE) of an ISTE estimator for subclass G_k as

$$\text{RMSE}(k) = \sqrt{E\left\{\left(\hat{\omega}^{q=1/2}(X_i) - \omega^{q=1/2}(X_i)\right)^2 \mid X_i \in G_k\right\}}, \quad k = 1, \dots, K.$$

Finally, given the objective of ISTE estimation is often associated with informing individualized treatment rule (ITR), we also consider regret as a performance metric to better connect with the literature on ITR.⁴⁴ Specifically, the expected regret (or sometimes referred to as risk) in subclass G_k is defined as

$$\text{Regret}(k) = E\left\{\left[I(\hat{z}_{\text{opt}}(X_i) \neq z^*(X_i)) \times \omega^{q=1/2}(X_i)\right] \mid X_i \in G_k\right\}, \quad k = 1, \dots, K,$$

where $\hat{z}_{\text{opt}}(X_i) = I\{\hat{\omega}^{q=1/2}(X_i) > 0\}$ is the recommended optimal ITR based on the estimated ISTE, $z^*(X_i) = I\{\omega^{q=1/2}(X_i) > 0\}$ is the “true” (infeasible) optimal ITR given X_i . Similar concepts were also discussed in Manski,⁴⁵ Hirano and Porter,⁴⁶ and references therein. Intuitively, the expected regret is interpreted as the average treatment benefit loss due to incorrect treatment recommendation based on the estimated ISTE, and a smaller value indicates more precise treatment recommendations.

To estimate these performance metrics, let $G_{k,b}$ contain individuals within the q_k quantile of the true propensity score from replication b , and $N_{G_{k,b}} = \sum_{i=1}^n I(X_i \in G_{k,b})$ denote its size.

We estimate the relative bias for subclass G_k by

$$\widehat{\text{RelBias}}(k) = \frac{\left\{ \frac{1}{B} \sum_{b=1}^B \frac{1}{N_{G_{k,b}}} \sum_{i=1}^n I(X_i \in G_{k,b}) \left(\hat{\omega}_b^{q=1/2}(X_i) - \omega_b^{q=1/2}(X_i) \right) \right\}}{\left\{ \frac{1}{B} \sum_{b=1}^B \frac{1}{N_{G_{k,b}}} \sum_{i=1}^n I(X_i \in G_{k,b}) \omega_b^{q=1/2}(X_i) \right\}},$$

where the true ISTE $\omega_b^{q=1/2}(X_i)$ is computed as the difference in the median survival times of two true counterfactual survival curves for each individual with covariate profile X_i in the b th simulation iteration, and $\hat{\omega}_b^{q=1/2}(X_i)$ is the corresponding estimate based the b th simulated data set. The RMSE for subclass G_k was estimated using

$$\widehat{\text{RMSE}}(k) = \sqrt{\frac{1}{B} \sum_{b=1}^B \frac{1}{N_{G_{k,b}}} \sum_{i=1}^n I(X_i \in G_{k,b}) \left(\hat{\omega}_b^{q=1/2}(X_i) - \omega_b^{q=1/2}(X_i) \right)^2}.$$

Finally, the expected regret for subclass G_k was estimated by

$$\widehat{\text{Regret}}(k) = \frac{1}{B} \sum_{b=1}^B \frac{1}{N_{G_{k,b}}} \sum_{i=1}^n \left\{ I(X_i \in G_{k,b}) I(\hat{z}_{\text{opt}}(X_i) \neq z^*(X_i)) \right\} \omega^q = 1/2(X_i).$$

Of note, while it is generally challenging to obtain valid confidence intervals for the ISTE using the frequentist machine learning methods directly from the model output, the Bayesian methods naturally produce coherent posterior credible intervals for the ISTE. Thereby, we summarized in Section 4.4 the empirical coverage probability for two “winners” of the AFT-BART methods.

4 | SIMULATION RESULTS

4.1 | Performance under strong overlap

A total of 12 methods were considered. We first evaluated the performance of each method without the challenge due to lack of overlap. Table 1 summarizes, for sample size $n = 5000$, the mean and standard deviation of PEHE under each configuration with strong covariate overlap (see panel A in Figure 1). Several patterns are immediate. First, all machine learning methods yielded substantially smaller PEHE than the parametric AFT and Cox regression models across all simulation scenarios. Second, with uncorrelated confounders and no interaction between confounders (HS (i) to (iii)), as the heterogeneity setting became more complex, machine learning methods tended to have better performance evidenced by decreasing mean and variation of the PEHE. When correlation and interaction between confounders were introduced in HS (iv), all machine learning methods but GAPH yielded similar results as for HS (ii), whereas GAPH had increased PEHE. In sharp contrast, parametric regression methods tended to perform worse, with increasing average PEHE as the heterogeneity setting became more complex. All methods show decreasing precision in estimating ISTE when the proportional hazards assumption fails to hold and under a higher censoring rate. The Cox proportional hazards model results in the largest increase in PEHE when the proportional hazards assumption is violated. In comparison, all machine learning estimators have a much smaller increase in PEHE under nonproportional hazards. This is possibly due to more complex data generating processes (the parameter η depends on a nonlinear form of the covariates) and slightly reduced effective sample size. Third, among the nine machine learning approaches, AFT-BART-NP had the best overall performance across simulation configurations; AFT-BART-NP-PS had essentially the same PEHE results as AFT-BART-NP. Web Table 1 presents the PEHE for $n = 500$. The findings are qualitatively similar except that the mean and variance of PEHE increase with a smaller sample size.

The relative bias, RMSE and expected regret within the $K = 50$ propensity score subclasses are summarized in Web Figures 2 to 7. The parametric AFT and Cox models had substantially larger bias, RMSE and regret compared to the class of machine learning methods across all scenarios; and even more so when the underlying model assumptions were violated. Specifically, the Cox model had the largest relative bias, RMSE and regret under the nonproportional hazards setting and AFT-Lognormal had the largest relative bias, RMSE and regret when the true outcomes were simulated from the Weibull curves.

Among the machine learning methods, AFT-BART-NP and AFT-BART-NP-PS generally maintained the lowest bias, RMSE and regret with a larger sample size ($n = 5000$) and under nonproportional hazards. AFT-BART-NP-PS, while producing the similar median values of the bias, RMSE and regret as AFT-BART-NP, had less variable performance measures demonstrated by the shorter length of boxplots of each metric. This efficiency improvement due to the inclusion of a propensity score estimate echoes the findings in Dorie et al²³ with noncensored outcomes. Overall, the tightness of the range of relative bias values for machine learning methods under the most complex HS (iii) and HS (iv) when $n = 5000$ demonstrate their superior performance over the traditional parametric implementations.

4.2 | Impact of covariate overlap

When there is strong covariate overlap, AFT-BART-NP demonstrates moderately better performance than AFT-BART-SP, RSF, DeepSurv, DR-DL, BJ-DL, TSHEE, and GAPH. As the degree of covariate overlap decreases, we find that the advantage of AFT-BART-NP becomes more apparent. In addition, inclusion of the estimated propensity score in the AFT-BART-NP model considerably reduced the variability of all performance measures, regardless of degree of overlap. Figure 2 shows that for the most challenging scenario (HS (iv) + nonproportional hazards + 60% censoring), the three performance measures (relative bias, RMSE and expected regret) increase as the lack of overlap increases for all machine learning methods, with both $n = 500$ and $n = 5000$. However, AFT-BART-NP and its variant AFT-BART-NP-PS appeared to be the most robust approach with the median relative bias around only 7% in the most extreme scenario ($\psi = 5$); DeepSurv appeared to have the largest increase in all performance metrics as the degree of overlap decreased. The two DL methods based on the censoring unbiased loss function, DR-DL and BJ-DL, showed similar performance across different simulation configurations, and performed better than the other DL method, DeepSurv, TSHEE as well as GAPH. AFT-BART-SP compared unfavorably with AFT-BART-NP, underscoring the benefit of more flexible nonparametric modeling of the residual term.

4.3 | Overall performance assessment

In Figure 3, we provided an overall assessment of the machine learning methods across the simulation scenarios following the approach taken in Tabib et al.¹⁶ Specifically, in each of the B simulation iterations, we computed the percent increase in PEHE with respect to the best performer for this iteration. The percent increase in PEHE for method s is given by

$$\left(\frac{\text{PEHE}_s}{\text{PEHE}_{\min}} - 1 \right) \times 100,$$

where PEHE_{\min} corresponds to the PEHE from the best performer (hence with the smallest PEHE). For a given iteration, only the best method will correspond to a percent increase of exact zero while the others lead to positive values. Figure 3 presents the boxplots of percent increase in PEHE over all 3500 simulations for $n = 5000$ and 14000 simulations for $n = 500$ (Details of the scenarios covered are explained in the figure title). Evidently, the AFT-BART-NP method outperformed the other approaches (with AFT-BART-NP-PS slightly trailing behind) for having the smallest median and interquartile range of the percent

increase in PEHE; the advantage of AFT-BART-NP and AFT-BART-NP-PS is much more pronounced for a larger sample size $n = 5000$.

4.4 | Frequentist coverage for AFT-BART-NP

Because we identified the AFT-BART-NP and its variant AFT-BART-NP-PS to be the “best performer” in our simulations, we further assessed the accuracy of associated frequentist inference based on this flexible outcome model. An additional advantage of AFT-BART-NP is that it naturally produces coherent posterior intervals in contrast to other machine learning methods, for which it may be difficult to derive simple and valid interval estimates for ISTE. Table 2 provides the mean frequentist coverage of the AFT-BART-NP credible interval estimators of the ISTE within the propensity score subclasses, under HS (iv), 60% censoring rate and three different degrees of covariate overlap. In this most challenging HS scenario, when there is strong or moderate level of overlap, AFT-BART-NP leads to valid inference for ISTE because it generally provides close to nominal frequentist coverage across all subclasses for both sample sizes. Furthermore, under weak overlap for which it is challenging to correctly estimate or make valid inference even for the population average causal effects,⁴³ the AFT-BART-NP approach could still provide close to nominal frequentist coverage among the data region where treatment assignment is approximately balanced (eg, the 25th subclass where there is sufficient *local overlap*). However, the coverage probability decreases as we move toward the tail regions of the propensity score distribution. This interesting finding on frequentist coverage echoes the previous findings for estimating the population level causal effect, but on a finer scale as we focused on individual-specific heterogeneous treatment effect. In addition, we find inclusion of the estimated propensity score in the AFT-BART model formulation further improved the mean frequentist coverage across all scenarios, and more considerably so near the tail regions of the propensity score distribution, that is, in subclasses G_k 's away from G_{25} . A full summary of frequentist coverage across different scenarios is visualized in Web Figure 8 and conveys essentially the same message.

4.5 | Sensitivity to covariate-dependent censoring

We assessed the sensitivity of our findings to more complex censoring mechanisms. Specifically, we replicated the simulation for the most challenging scenario (HS (iv) + nonproportional hazards) using the data generating process outlined in Section 3.1, except that we generate the counterfactual censoring times $C(1) = C(0)$ from an exponential distribution with rate parameter $\rho(X)$, where

$$\rho(X) = 0.02 + 0.1X_1^2 + 0.1\sin(X_3) + 0.2e^{X_5}/(1 + e^{X_5}) + 0.02X_6 - 0.03X_7.$$

The marginal censoring proportion implied by this model remains 60%. Web Figure 9 presents the relative bias, RMSE and expected regret across all methods under strong, moderate, and weak overlap, and with sample sizes $n = 500$ and $n = 5000$. The results appeared almost identical to those in Figure 2, and AFT-BART-NP and AFT-BART-NP-PS were still identified to be the best approaches in terms of bias and precision. Web Table 2 presents the mean frequentist coverage of the AFT-BART-NP and AFT-BART-NP-PS

credible interval estimators of the ISTE within five propensity score subclasses, in parallel to Table 2. The results confirm that the frequentist coverage probability of these approaches is unaffected by covariate-dependent censoring.

4.6 | Sensitivity to unmeasured confounding

While the flexible machine learning approaches relaxed the parametric assumptions for outcome modeling, their validity for estimating ISTE using observational data still critically depends on the structural causal assumptions (A1) to (A4) in Section 2.1. Among them, a key structural assumption is the availability of pretreatment confounders so that the weak unconfoundedness holds. To assess the sensitivity of the estimated ISTE to various degrees of departure from the weak unconfoundedness assumption, we conducted an *exploratory* sensitivity analysis using our simulation structure. We generated potential survival outcomes under the heterogeneous setting (iii), in which all of the machine learning methods had relatively better performance compared to other heterogeneous settings. We sequentially increased the degree of unmeasured confounding by removing access to confounders 1) X_3 , 2) X_3 and X_5 and 3) X_3 , X_5 , and X_6 in the analysis stage for each of the 12 methods. The PEHE results for $n = 5000$ are provided in Table 3. Results for $n = 500$ convey the same message and were therefore omitted. In the presence of unmeasured confounders, the performance of all methods deteriorated, demonstrated by the larger mean and standard deviation of the PEHE under increasing degree of unmeasured confounding. However, the relative performance rankings remain the same as in the situation when the weak unconfoundedness assumption holds, with AFT-BART-NP and AFT-BART-NP-PS being the top performing methods. Results on the relative bias, RMSE and expected regret are summarized in Web Figures 10 to 12, and further support the findings in Table 3.

5 | CASE STUDY: ESTIMATING HETEROGENEOUS SURVIVAL CAUSAL EFFECT OF TWO RADIOTHERAPY APPROACHES FOR PATIENTS WITH HIGH-RISK PROSTATE CANCER

Two popular radiotherapy based approaches for treating high-risk localized prostate cancer are EBRT + AD and EBRT + brachy ± AD.¹⁰ It is unclear whether EBRT + AD and EBRT + brachy ± AD have significantly different survival effect among those who received at least 7920 cGY EBRT dose, and whether there exists TEH. To better understand the comparative overall and heterogeneous survival causal effects of these two radiotherapy approaches, we pursued a two-stage causal analysis with the NCDB data including 7730 high-risk localized prostate cancer patients who were diagnosed between 2004 and 2013 and were treated with either EBRT + AD or EBRT + brachy ± AD with no less than 7920 cGY EBRT dose. The pretreatment confounders included in this analysis were age, prostate-specific antigen (PSA), clinical T stage, Charlson-Deyo score, biopsy Gleason score, year of diagnosis, insurance status, median income level (quartiles on the basis of zip code of residence), education (based on the proportion of residents in the patient's zip code who did not graduate high school), race, and ethnicity. Web Table 3 summarizes the baseline characteristics of patients in the two treatment groups. Clearly, there are a number of covariates that are unbalanced between the two groups, suggesting the need to adjust for

confounding. Web Figure 13 shows the distribution of the estimated propensity scores by fitting a Super Learner³⁵ on the treatment and covariates data. Relatively strong covariate overlap is suggested from the propensity score distribution.

Our analyses proceed in a two-stage fashion. In the first stage, we applied the machine learning methods to estimate the ISTE within the counterfactual framework. The estimated ISTE were then used, in the second stage, as dependent variables in an exploratory linear regression analysis to generate some further insights. As argued in Lu et al,³ the estimated regression coefficients can be interpreted as how treatment effects are modulated by corresponding covariates. To estimate standard errors and uncertainty intervals for the regression coefficients, we drew for each individual 1000 posterior MCMC samples of ISTE from the AFT-BART models. For the frequentist machine learning approaches (RSF, DeepSurv, TSHEE, DR-DL and BJ-DL), we subsampled the entire procedure to obtain uncertainty intervals for the linear regression coefficients. Specifically, we drew a sample of size $n/5$ without replacement, refit each model based on the subsampled data and used the resulting ISTE estimates as dependent variable in the exploratory linear regression. The subsampling process was repeated 1000 times. Similar to Politis et al⁴⁷ and Lu et al,³ we used the subsampling approach instead of bootstrap due to considerations on computational speed and robustness. Web Table 4 presents the regression coefficients and 95% credible intervals for all covariates and the intercept from the AFT-BART-NP-PS model; the intercept could be considered as an overall measure of the treatment effect for the “baseline” population (with all covariates set to be the reference category or mean value). Results from using ISTE estimated by frequentist machine learning approaches were similar and therefore omitted. The intercept estimates suggest that the population average survival treatment effect of EBRT + brachy ± AD vs EBRT + AD is not statistically significant at the 0.05 level. However, PSA and age significantly modulated the comparative effect of the two radiotherapy approaches on median survival. For example, the estimated coefficient was 3.99 (95% CI: 3.60 to 4.38) for PSA; this means that difference in the counterfactual median survival under EBRT + brachy ± AD vs EBRT + AD, which was positive, became greater for patients with higher PSA. Similarly, the negative coefficient estimate for age suggests that the median survival benefit from EBRT + brachy ± AD were more pronounced among younger patients. In Figure 4, we created 300 subgroups based on percentiles of the distribution of posterior mean ISTE, and provide the scatter plots of subgroup-specific average covariate (PSA, age, clinical T stage, and insurance status) vs subgroup-specific average survival treatment effect. While the subgroup survival treatment effect estimates are randomly scattered across values of clinical T stage and insurance status, there is a clear nonrandom pattern between subgroup-specific survival treatment effect estimates and PSA or age, with directions in close agreement to findings from the exploratory linear regression analysis (Web Table 4).

To further explore the subgroups that may have enhanced survival treatment effect from either of the two radiotherapy approaches, we leveraged the posterior MCMC draws of the ISTE from AFT-BART-NP, and employed the “fit-the-fit” strategy to explore the potential TEH in relation to the covariate subspaces. The “fit-the-fit” strategy has been used to identify possible subgroups defined by combination rules of covariates that have differential treatment effects; see for example, Foster et al¹ and Logan et al⁴⁴ Instead of a linear

regression analysis, the classification and regression tree (CART) model was used to regress the posterior mean of the ISTE on the covariates. A sequence of CART models were fit, where covariates were sequentially added to the CART model in a stepwise manner to improve the model fit, as measured by R^2 . At each step, the variable leading to the largest R^2 improvement was selected into the model and the procedure ended until the percent improvement in R^2 was less than 1%. This complexity parameter controls the size of the final CART tree. Subgroup treatment effects were estimated by averaging the ISTE among individuals falling into each node of the final CART model, and the branch decision rules (ie, binary splits of the covariate space) suggested combination rules of covariates. Figure 5 shows the results of the final tree estimates. The final R^2 between the tree fit and the posterior mean treatment effect of EBRT + brachy ± AD vs EBRT + AD was 80%. The first splitting variable was PSA. Patients with PSA ≥ 10 ng/mL had approximately 7.6 (95% CI: 0.6 to 14.5) months longer median survival with EBRT + brachy ± AD, while patients with PSA < 10 ng/mL did not experience significantly higher median survival from either treatment. A second level of variable splitting provided further resolution on the magnitude of the treatment benefit for patients with higher PSA ≥ 10 ng/mL; the most beneficial subgroup was those with PSA ≥ 25 ng/mL and had approximately 10.9 months longer median survival. Among patients with lower PSA < 10 ng/mL, older patients aged above 70 years experienced treatment benefit from EBRT + AD with median survival approximately 2.2 months longer; on the contrary, younger patients aged below 70 years could benefit from EBRT + brachy ± AD with approximately 4.6 months longer median survival. This alternative analysis strategy generated evidence that converged to the exploratory linear regression analysis on the ISTE estimates, and suggested interpretable patient subpopulations with beneficial survival treatment effects due to EBRT + brachy ± AD. We further explored the sensitivity of the summarizing CART model fit to the tuning parameters of both CART and AFT-BART-NP. When the complexity parameter value of CART ranges from 0.6% to 1.1%, the final tree remains the same. With a value > 1.1%, the final summarizing tree includes only PSA and the final R^2 drops to 0.67. When the complexity parameter value is < 0.6%, the final summarizing tree gets more complex by including an additional variable biopsy Gleason score with only a slight increase (0.005) in the final R^2 . Varying values for the hyperparameters of the AFT-BART-NP model did not change the results of the CART model fit. Hence, the summarizing CART model fit in this case study is robust to the tuning parameters.

6 | DISCUSSION

Identification of TEH has been conventionally carried out by first enumerating potential effect modifiers with subject-matter experts and then estimating the average treatment effect within each subgroup.^{48,49} In randomized clinical trials, this approach could also facilitate prespecification and therefore is particularly suitable to *confirmatory TEH analysis*.⁴⁸ In observational data with complex heterogeneity of treatment effect, such a priori specification that separates the issues of confounding and TEH is often practically infeasible, especially with a large number of pretreatment covariates. Such complications motivate *exploratory THE analysis* for generating scientifically meaningful hypotheses and treatment effect discovery, which is the focus of this current article.^{48–50} Exploiting

on the counterfactual framework, flexibly modeling of the outcome has been shown to be well-suited for exploratory TEH analysis. However, ongoing discussions about causal methods for estimating TEH in observation data have largely focused on nonsurvival data. For example, tree-based methods such as BART and random forests have recently emerged as popular approaches to causal effect estimation for continuous or binary outcomes.^{3,7} On the other hand, several modern deep learning methods (DeepSurv, DR-DL, BJ-DL) have been developed for censored survival data,^{18,19} and a TMLE based method has recently been proposed to estimate the survival heterogeneous effect.²¹ The GAPH model is adept at flexibly fitting complex main effects, which are often of great interest in health studies.²² We adapt these machine learning modeling tools for estimating ISTE and provide new empirical evidence using simulations representative of complex confounding and heterogeneity settings with right-censored survival data.

Our results indicate that all machine learning methods substantially outperformed the traditional AFT and Cox proportional hazards models in estimating TEH. This is not surprising given that the traditional models often encode strong parametric and linearity assumptions, which are prone to bias when model assumptions fail to hold. Among the class of machine learning methods, the nonparametric version of the AFT-BART model¹² generally outperformed the other methods, both from the traditional prediction perspective (bias and RMSE) and the decision perspective (expected regret). The advantage of AFT-BART-NP becomes more pronounced with a larger sample size and increased data complexity. On the other hand, AFT-BART-NP also conveniently enables posterior inference by providing credible intervals for ISTE and the associated subgroup-specific treatment effects. In particular, we demonstrated that when there is strong or moderate overlap, AFT-BART-NP could provide close to nominal frequentist coverage for almost all average causal effect among subpopulations defined by propensity score strata. Under weak overlap, AFT-BART-NP still provides satisfactory frequentist coverage for the subpopulation near the centroid of the propensity score distribution. While the overall population may exhibit weak overlap, the central region of the propensity score distribution may still present satisfying *local overlap* and in fact admits valid inference for the causal effects. However, AFT-BART-NP could understate the uncertainty of ISTE estimates in regions with lack of overlap. This observation was also discussed in an invited Commentary to Hahn et al²⁴ by Papadogeorgou and Li,⁵¹ who illustrated that BART for continuous outcomes produced overly narrow confidence band in the region of poor overlap. As a potential improvement, we found that including a nonparametrically estimated propensity score in the AFT-BART-NP model formulation leads to more accurate estimation and inference for ISTE, evidenced by less variable subgroup-specific performance measures, as well as higher frequentist coverage under weak overlap. Because previous simulations for TEH rarely considered the effect of overlap, our results offer new insights. Importantly, our findings can provide new perspectives into the estimation of ISTE relative to average treatment effects. At the minimum, we noticed that even though it is challenging to estimate ISTE for units at the tails of the propensity score distribution, it remains practically feasible to make accurate inference for ISTE within the centroid of the propensity score distribution. In fact, the centroid region of the propensity score distribution includes individuals at clinical equipoise (and for whom the treatment decisions are mostly unclear) and resemble those recruited

in a randomized controlled trial. In the propensity score literature, this subpopulation is sometimes referred to as the *overlap population*, for which there exists an efficient weighting estimator for the average causal effect.^{27,43,52,53}

In general, the flexible outcome modeling represents a simple and powerful approach for causal effect estimation, with the caveat that accurate estimation critically depends on correct model specification of the outcome surface.⁵⁴ With the rapid advancement of machine learning techniques, more precise modeling of the outcome surface can be achieved, which can serve as the building block for improved individualized causal effect estimation. The ISTE estimates can also be utilized to explore the degree of TEH. For example, as in the second stage of the virtual twins approach, the ISTE estimates were used as the outcome and a tree diagram was fitted to identify predictors that explain away differences in the ISTE estimates. In Lu et al.,³ the ISTE estimates were used as the dependent variable in a standard regression model to explain between-subgroup differences in treatment effect. As suggested in Anoke et al.,⁸ averaging ISTE estimates within a prespecified subgroup (eg, female vs male) can also facilitate exploration of the TEH. Given these existing tools, the ISTE from a flexible outcome model can yield insights not only for the population average treatment effect but also for the detection of TEH. Another interesting practical point regards how to use the estimated ISTE to inform treatment decisions for future patients when only part of the covariates X^{obs} will be available at the decision time, and X^{mis} will not be collected. In such situations, a plausible strategy is to marginalize the product of the ISTE functional $f(\omega|x)$ and the conditional distribution $f(X^{\text{mis}}|X^{\text{obs}})$, over the probability distribution of X^{mis} . However, the estimation of the conditional $f(X^{\text{mis}}|X^{\text{obs}})$ can be computationally challenging, especially when there are a large number of covariates.

In the NCDB case study, we showed how the counterfactual ISTE estimates from machine learning models could be utilized to understand whether there might be TEH among localized high-risk prostate cancer patients who were treated with two different radiotherapy approaches. These two radiotherapy techniques did not show significantly different average treatment effects in the overall study population, but we found age and PSA are two key patient factors that modulated the comparative treatment effect. We further exploited the posterior MCMC samples of the ISTE from AFT-BART-NP and the binary tree model CART to explore interesting connections between radiotherapy techniques, patient characteristics and median counterfactual survival. This causal analysis revealed that patients with very high PSA or older patients with low PSA may have enhanced treatment benefit from EBRT + brachy ± AD or EBRT + AD. The results could facilitate treatment effect discovery in subpopulations and may inform personalized treatment strategy and the planning of future confirmatory randomized trials.

While we have defined our individual-level causal estimand based on the median counterfactual survival times in Section 2, it is straightforward to extend this definition to alternative individual-level causal estimands. For example, we can take different values for $q \in (0, 1)$ in Equation (2) than 1/2 and obtain the survival causal effect conditional on the individual profile $X = x$ for other quantiles of the distribution of failure times. While the median survival times are frequently used for its simple interpretation in clinical practice,

studying individual quantile survival causal effects could provide additional insights on how treatment effect varies across the distribution of failure times. However, when q departs from $1/2$, the effective sample size for estimating the individual causal effect may decrease due to data sparsity, and the resulting interval estimates could become wider. Beyond the quantile survival causal effect, we can also define individual causal estimands by contrasting the conditional survival probability up to a fixed time t , conditional restricted mean survival times (RMST), and conditional (cumulative) hazard rate, extending the population-level estimands introduced in, for example, Chen and Tsiatis,²⁵ Zhang et al,^{55,56} Bai et al,⁵⁷ and Mao et al.²⁸ Because we embedded the machine learning approaches considered in the outcome modeling framework and were able to estimate the individual counterfactual survival curves, it is convenient to leverage the estimated individual-specific survival curves and obtain TEH represented by these alternative definitions. We have not examined the corresponding empirical results, and therefore recommend future investigations to identify the best approaches for estimating these alternative individual survival causal effect estimands.

We acknowledge that there are several limitations of our study, which could stimulate future research in the field of (frequentist or Bayesian) causal machine learning. First, while our simulations aimed at providing new empirical evidence about the operating characteristics of state-of-the-art machine learning techniques for estimating TEH with survival data, we only provided frequentist coverage probability for AFT-BART-NP because the credible intervals are readily available from the MCMC output. On the other hand, it may be challenging to precisely estimate the variance of the ISTE using DL, RSF, and TSHEE, without resorting to more computationally intensive sample-splitting methods.^{2,15,58,59} Developing and investigating new methods for the variance and interval estimation for frequentist machine learning represent an important avenue for future research. Second, an anonymous reviewer pointed out that one may further improve the estimation of ISTE by parameterizing the counterfactual AFT mean model as $E[\log(T(z))|X = x] = \alpha(x) + (z - \epsilon(x))\tau(x)$, where the ISTE can be explicitly represented as a function of $\tau(x)$. This apparent parameterization orthogonalizes the model into two components: a prognostic score, $\alpha(x) - \epsilon(x)\tau(x)$, as well as a treatment effect model, $z\tau(x)$, and anchors the propensity score in model formulation to build in robustness. While this idea represents a promising approach, developing a formal estimation procedure by placing separate regularization priors on these two model components along the lines of Hahn et al²⁴ is beyond the scope of this work, and will be pursued in future work. Third, we have assumed conditional noninformative censoring so that the counterfactual censoring time is independent of the counterfactual survival time given baseline covariates. It would be interesting to further develop these machine learning to address informative censoring such that the censoring process is a function of some time-varying covariates.⁶⁰ Combining the inverse probability of censoring weights^{61,62} and the flexible outcome models also represents a promising approach to obtain valid causal estimates in this scenario but requires further study in the context of TEH. Finally, as in all causal inference work with observational data, we require an untestable assumption that the observed pretreatment covariates are sufficient to deconfound the relationship between treatment and outcome. We conducted an *exploratory* sensitivity analysis to demonstrate that, the violation of this assumption will lead to more biased ISTE estimates even for

machine learning methods. To wit, using flexible modeling techniques alone will not address violations of structural assumptions (such as weak unconfoundedness) required for causal inference. Developing a formal sensitivity analysis approach to capture the sensitivity of an ISTE estimation method to the potential magnitude of departure from the unconfoundedness assumption would be a worthwhile and important contribution.^{26,63}

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGEMENTS

This work was supported in part by the National Cancer Institute under grant NIH NCI R21CA245855 and grant NIH NCI P30CA196521-01, and by award ME_2017C3_9041 from the Patient-Centered Outcomes Research Institute (PCORI). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health or PCORI. The authors thank the associate editor and two anonymous referees, whose suggestions have substantially improved the exposition of this work.

Funding information

National Cancer Institute, Grant/Award Number: R21CA245855; Patient-Centered Outcomes Research Institute, Grant/Award Number: ME_2017C3_9041

DATA AVAILABILITY STATEMENT

The simulation codes that generate the data supporting the findings of the simulation study are openly available on GitHub at <https://github.com/liangyuanhu/ML-DL-TEH>. The NCDB data used in the case study is publicly available upon approval of the NCDB Participant User File application.

REFERENCES

1. Foster JC, Taylor JM, Ruberg SJ. Subgroup identification from randomized clinical trial data. *Stat Med.* 2011;30(24):2867–2880. [PubMed: 21815180]
2. Wager S, Athey S. Estimation and inference of heterogeneous treatment effects using random forests. *J Am Stat Assoc.* 2018;113(523):1228–1242.
3. Lu M, Sadiq S, Feaster DJ, Ishwaran H. Estimating individual treatment effect in observational data using random forest methods. *J Comput Graph Stat.* 2018;27(1):209–219. [PubMed: 29706752]
4. Leo Breiman. *Machine Learning.* 2001;45(1):5–32.
5. Mazumdar M, Lin J, Zhang W, Li L, Liu M, Dharmarajan K, Sanderson M, Isola L, Hu L. Comparison of statistical and machine learning models for healthcare cost data: a simulation study motivated by Oncology Care Model (OCM) data. *BMC Health Services Research.* 2020;20(1):350. [PubMed: 32334595]
6. Ishwaran H, Malley JD. Synthetic learning machines. *BioData Mining.* 2014;7(1):28. [PubMed: 25614764]
7. Hill JL. Bayesian nonparametric modeling for causal inference. *J Comput Graph Stat.* 2011;20(1):217–240.
8. Anoke SC, Normand SL, Zigler CM. Approaches to treatment effect heterogeneity in the presence of confounding. *Stat Med.* 2019;38(15):2797–2815. [PubMed: 30931547]
9. Hu L, Hogan JW, Mwangi AW, Siika A. Modeling the causal effect of treatment initiation time on survival: application to HIV/TB co-infection. *Biometrics.* 2018;74(2):703–713. [PubMed: 28960243]

10. Ennis RD, Hu L, Ryemon SN, Lin J, Mazumdar M. Brachytherapy-based radiotherapy and radical prostatectomy are associated with similar survival in high-risk localized prostate cancer. *J Clin Oncol*. 2018;36(12):1192–1198. [PubMed: 29489433]
11. Chen RC. Challenges of interpreting registry data in prostate cancer: interpreting retrospective results along with or in absence of clinical trial data. *J Clin Oncol*. 2018;36(12):1181–1183. [PubMed: 29489432]
12. Henderson NC, Louis TA, Rosner GL, Varadhan R. Individualized treatment effects with censored data via fully nonparametric Bayesian accelerated failure time models. *Biostatistics*. 2020;21(1):50–68. [PubMed: 30052809]
13. Tian L, Alizadeh AA, Gentles AJ, Tibshirani R. A simple method for estimating interactions between a treatment and a large number of covariates. *J Am Stat Assoc*. 2014;109(508):1517–1532. [PubMed: 25729117]
14. Shen J, Wang L, Daignault S, Spratt DE, Morgan TM, Taylor JM. Estimating the optimal personalized treatment strategy based on selected variables to prolong survival via random survival forest with weighted bootstrap. *J Biopharm Stat*. 2018;28(2):362–381. [PubMed: 28934002]
15. Cui Y, Kosorok MR, Wager S, Zhu R. Estimating heterogeneous treatment effects with right-censored data via causal survival forests; 2020. arXiv preprint arXiv:2001.09887.
16. Tabib S, Larocque D. Non-parametric individual treatment effect estimation for survival data with random forests. *Bioinformatics*. 2020;36(2):629–636. [PubMed: 31373350]
17. Ishwaran H, Kogalur UB, Blackstone EH, Lauer MS. Random survival forests. *Ann Appl Stat*. 2008;2(3):841–860.
18. Katzman JL, Shaham U, Cloninger A, Bates J, Jiang T, Kluger Y. DeepSurv: personalized treatment recommender system using a Cox proportional hazards deep neural network. *BMC Med Res Methodol*. 2018;18(1):24. [PubMed: 29482517]
19. Steingrimsson JA, Morrison S. Deep learning for survival outcomes. *Stat Med*. 2020;39(17):2339–2349. [PubMed: 32281672]
20. Buckley J, James I. Linear regression with censored data. *Biometrika*. 1979;66(3):429–436.
21. Zhu J, Gallego B. Targeted estimation of heterogeneous treatment effect in observational survival analysis. *J Biomed Inform*. 2020;107:103474. [PubMed: 32562899]
22. Hastie TJ, Tibshirani RJ. *Generalized Additive Models*. Boca Raton, FL: Chapman & Hall; 1990.
23. Dorie V, Hill J, Shalit U, Scott M, Cervone D. Automated versus do-it-yourself methods for causal inference: lessons learned from a data analysis competition. *Stat Sci*. 2019;34(1):43–68.
24. Hahn PR, Murray JS, Carvalho CM. Bayesian regression tree models for causal inference: regularization, confounding, and heterogeneous effects (with discussion). *Bayesian Anal*. 2020;15(3):965–1056.
25. Chen PY, Tsiatis AA. Causal inference on the difference of the restricted mean lifetime between two groups. *Biometrics*. 2001;57(4):1030–1038. [PubMed: 11764241]
26. Hogan J, Daniels M, Hu L. A Bayesian perspective on assessing sensitivity to assumptions about unobserved data. *Handbook of Missing Data Methodology*. Boca Raton, FL: CRC Press; 2014.
27. Li F, Morgan KL, Zaslavsky AM. Balancing covariates via propensity score weighting. *J Am Stat Assoc*. 2018;113(521):390–400.
28. Mao H, Li L, Yang W, Shen Y. On the propensity score weighting analysis with survival outcome: estimands, estimation, and inference. *Stat Med*. 2018;37(26):3745–3763. [PubMed: 29855060]
29. Uno H, Claggett B, Tian L, et al. Moving beyond the hazard ratio in quantifying the between-group difference in survival analysis. *J Clin Oncol*. 2014;32(22):2380. [PubMed: 24982461]
30. Uno H, Wittes J, Fu H, et al. Alternatives to hazard ratios for comparing the efficacy or safety of therapies in noninferiority studies. *Ann Intern Med*. 2015;163(2):127–134. [PubMed: 26054047]
31. Hu L, Hogan JW. Causal comparative effectiveness analysis of dynamic continuous-time treatment initiation rules with sparsely measured outcomes and death. *Biometrics*. 2019;75(2):695–707. [PubMed: 30638268]
32. Chipman HA, George EI, McCulloch RE. BART: Bayesian additive regression trees. *Ann Appl Stat*. 2010;4(1):266–298.

33. Tan YV, Roy J. Bayesian additive regression trees and the general BART model. *Stat Med*. 2019;38(25):5048–5069. [PubMed: 31460678]
34. Rosenbaum PR, Rubin DB. The central role of the propensity score in observational studies for causal effects. *Biometrika*. 1983;70(1):41–55.
35. Pirracchio R, Petersen ML, Van der Laan M. Improving propensity score estimators' robustness to model misspecification using super learner. *Am J Epidemiol*. 2015;181(2):108–119. [PubMed: 25515168]
36. Van der Laan MJ, Polley EC, Hubbard AE. Super learner. *Stat Appl Genet Mol Biol*. 2007;6(1):25.
37. Cai W, Van der Laan MJ. One-step targeted maximum likelihood estimation for time-to-event outcomes. *Biometrics*. 2020;76(3):722–733. [PubMed: 31729005]
38. Kalbfleisch JD, Prentice RL. *The Statistical Analysis of Failure Time Data*. Vol 360. Hoboken, NJ: John Wiley & Sons; 2011.
39. Wei LJ. The accelerated failure time model: a useful alternative to the Cox regression model in survival analysis. *Stat Med*. 1992;11(14–15):1871–1879. [PubMed: 1480879]
40. Cox DR. Regression models and life-tables. *J Royal Stat Soc Ser B (Methodol)*. 1972;34(2):187–202.
41. Hu L, Gu C, Lopez M, Ji J, Wisnivesky J. Estimation of causal effects of multiple treatments in observational studies with a binary outcome. *Stat Methods Med Res*. 2020;29(11):3218–3234. [PubMed: 32450775]
42. Hu L, Gu C. Estimation of causal effects of multiple treatments in healthcare database studies with rare outcomes. *Health Serv Outcomes Res Methodol*. 2021. 10.1007/s10742-020-00234-4.
43. Li F, Thomas LE, Li F. Addressing extreme propensity scores via the overlap weights. *Am J Epidemiol*. 2019;188(1):250–257. [PubMed: 30189042]
44. Logan BR, Sparapani R, McCulloch RE, Laud PW. Decision making and uncertainty quantification for individualized treatments using Bayesian additive regression trees. *Stat Methods Med Res*. 2019;28(4):1079–1093. [PubMed: 29254443]
45. Manski CF. Statistical treatment rules for heterogeneous populations. *Econometrica*. 2004;72(4):1221–1246.
46. Hirano K, Porter JR. Asymptotics for statistical treatment rules. *Econometrica*. 2009;77(5):1683–1701.
47. Politis DN, Romano JP, Wolf M. *Subsampling*. Berlin, Germany: Springer Science & Business Media; 1999.
48. Kent DM, Paulus JK, Van Klaveren D, et al. The predictive approaches to treatment effect heterogeneity (PATH) statement. *Ann Intern Med*. 2020;172(1):35–45. [PubMed: 31711134]
49. Wang R, Lagakos SW, Ware JH, Hunter DJ, Drazen JM. Statistics in medicine-reporting of subgroup analyses in clinical trials. *N Engl J Med*. 2007;357(21):2189–2194. [PubMed: 18032770]
50. Assmann SF, Pocock SJ, Enos LE, Kasten LE. Subgroup analysis and other (mis) uses of baseline data in clinical trials. *Lancet*. 2000;355(9209):1064–1069. [PubMed: 10744093]
51. Papadogeorgou G, Li F. Invited discussion for “Bayesian regression tree models for causal inference: regularization, confounding, and heterogeneous effects”. *Bayesian Anal*. 2020;15(3):1007–1013.
52. Li F, Li F. Propensity score weighting for causal inference with multiple treatments. *Ann Appl Stat*. 2019;13(4):2389–2415.
53. Li F Comment: stabilizing the doubly-robust estimators of the average treatment effect under positivity violations. *Stat Sci*. 2020;35(3):503–510.
54. Antonelli J, Daniels MJ. Discussion of PENCOMP. *J Am Stat Assoc*. 2019;114(525):24–27. [PubMed: 34305210]
55. Zhang M, Schaubel DE. Contrasting treatment-specific survival using double-robust estimators. *Stat Med*. 2012;31(30):4255–4268. [PubMed: 22807175]
56. Zhang M, Schaubel DE. Double-robust semiparametric estimator for differences in restricted mean lifetimes in observational studies. *Biometrics*. 2012;68(4):999–1009. [PubMed: 22471876]

57. Bai X, Tsiatis AA, O'Brien SM. Doubly-robust estimators of treatment-specific survival distributions in observational studies with stratified sampling. *Biometrics*. 2013;69(4):830–839. [PubMed: 24117096]
58. Díaz I Machine learning in the estimation of causal effects: targeted minimum loss-based estimation and double/debiased machine learning. *Biostatistics*. 2020;21(2):353–358. [PubMed: 31742333]
59. Chernozhukov V, Chetverikov D, Demirer M, et al. Double/debiased machine learning for treatment and structural parameters. *Econ J*. 2018;21(1):C1–C68.
60. Siannis F, Copas J, Lu G. Sensitivity analysis for informative censoring in parametric survival models. *Biostatistics*. 2005;6(1):77–91. [PubMed: 15618529]
61. Robins JM, Finkelstein DM. Correcting for noncompliance and dependent censoring in an AIDS clinical trial with inverse probability of censoring weighted (IPCW) log-rank tests. *Biometrics*. 2000;56(3):779–788. [PubMed: 10985216]
62. Zhang M, Schaibel DE. Estimating differences in restricted mean lifetime using observational data subject to dependent censoring. *Biometrics*. 2011;67(3):740–749. [PubMed: 21039400]
63. Hu L, Zou J, Gu C, Ji J, Lopez M, Kale M. A flexible sensitivity analysis approach for unmeasured confounding with multiple treatments and a binary outcome. 2021;arXiv preprint arXiv:2012.06093 [stat.ME].




FIGURE 1.

Distributions of true propensity scores generated from three treatment assignment mechanisms with an overall treatment prevalence equal to 0.5. The unshaded bars indicate the treated group; the gray shaded bars indicate the control group. Panels A to C correspond to setting $\psi = 1$, $\psi = 2.5$, and $\psi = 5$, respectively




FIGURE 2. Relative bias, panel A, RMSE, panel B, and expected regret, panel C results for each of nine machine learning methods under strong, moderate, and weak covariate overlap, and for the complex scenario of heterogeneous setting (iv), nonproportional hazards and 60% censoring. Each boxplot visualizes the distribution of a performance measure across the propensity score subgroups G_k , $k = 1, \dots, 50$




FIGURE 3.

Overall comparison of nine machine learning methods across all simulation configurations. The boxplots represent the distribution of the percent increase in PEHE with respect to the best performer for each simulation run for all runs. Smaller values are better. The total number of simulation runs for each boxplot equals to the number of scenarios multiplied by B ($B = 250$ for $n = 5000$ and $B = 1000$ for $n = 500$). There are a total of 18 scenarios considered, including (a) 2 proportional hazards assumptions \times 2 censoring proportions \times 4 heterogeneity settings under strong overlap; plus (b) 1 proportional hazards assumptions (nonproportional hazards) \times 1 censoring proportions (CR= 60%) \times 1 heterogeneity settings (HS (iv)) under moderate and weak overlap

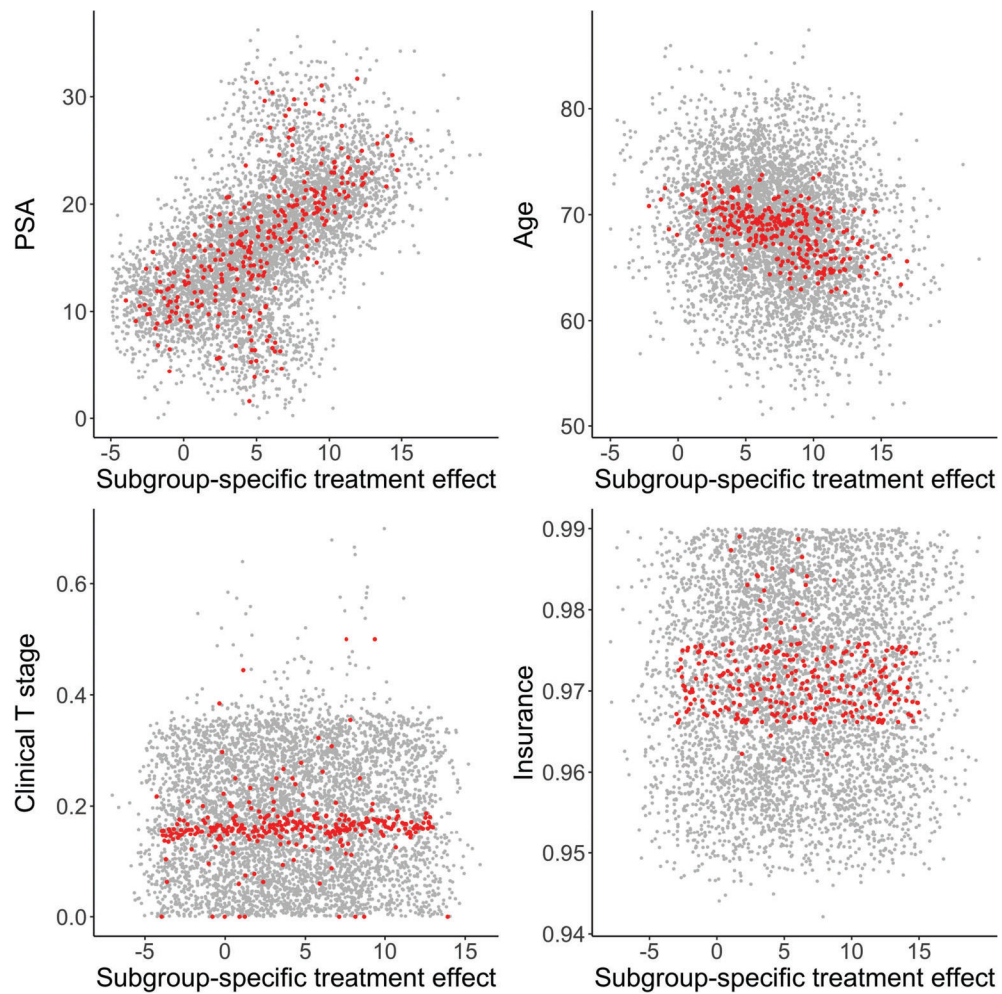
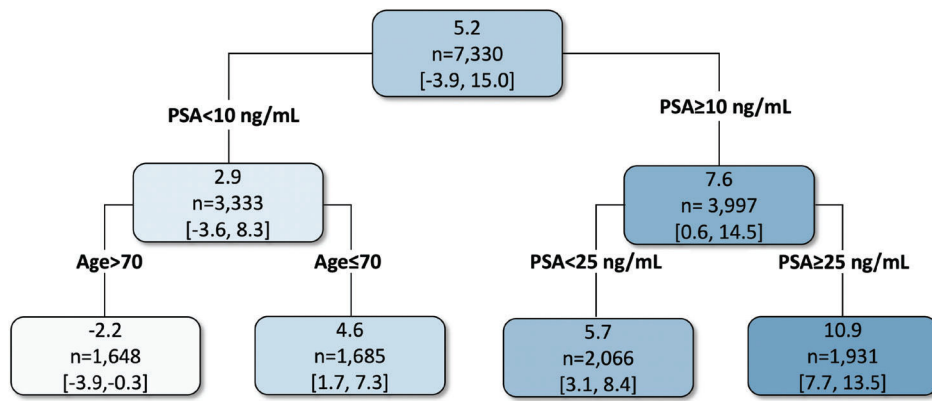


FIGURE 4.

Visualization of results from the NCDB data analysis. Covariates considered are PSA, age, clinical T stage, and insurance status. The x -axis represents the subgroup-specific posterior average survival treatment effect (median counterfactual survival in months), and the y -axis represents the subgroup-specific average value of the covariate. A red dot represents a subgroup, generated by dividing percentiles of the distribution of posterior mean individual survival treatment effects (ISTE). There are 300 subgroups in total. Gray dots represent a random subset of 100 (out of 1000) posterior draws of the ISTE with the subgrouping process applied

**FIGURE 5.**

Final CART model fit to the posterior mean of the individual survival treatment effect (measured by median survival in months) comparing EBRT plus brachytherapy with or without AD (EBRT + brachy ± AD) vs EBRT plus AD (EBRT + AD). Values in each node correspond to the posterior mean and 95% credible intervals for the average survival treatment effect for the subgroup of individuals represented in that node

TABLE 1

Mean (SD) of PEHE for each of 12 methods and each of 8 simulation configurations for $n = 5000$ with strong overlap

CR	Method	Proportional hazards				Nonproportional hazards			
		HS (i)	HS (ii)	HS (iii)	HS (iv)	HS (i)	HS (ii)	HS (iii)	HS (iv)
20%	AFT-L	3.63 (0.14)	4.17 (0.15)	4.13 (0.14)	4.81 (0.15)	3.79 (0.13)	4.35 (0.16)	4.41 (0.17)	4.89 (0.18)
	AFT-W	2.11 (0.12)	2.58 (0.12)	2.65 (0.12)	2.94 (0.13)	2.31 (0.13)	2.84 (0.13)	2.81 (0.13)	3.08 (0.15)
	CoxPH	2.16 (0.11)	2.61 (0.12)	2.61 (0.13)	2.98 (0.13)	3.81 (0.13)	4.38 (0.14)	4.39 (0.13)	4.86 (0.17)
	ABN	0.51 (0.03)	0.21 (0.02)	0.12 (0.02)	0.23 (0.03)	0.72 (0.07)	0.41 (0.04)	0.33 (0.04)	0.42 (0.05)
	ABNPS	0.52 (0.03)	0.22 (0.02)	0.13 (0.02)	0.24 (0.03)	0.73 (0.07)	0.42 (0.04)	0.34 (0.04)	0.43 (0.05)
	ABS	0.56 (0.04)	0.24 (0.03)	0.15 (0.02)	0.26 (0.03)	0.78 (0.08)	0.44 (0.05)	0.36 (0.04)	0.45 (0.05)
	RSFs	0.53 (0.03)	0.22 (0.02)	0.13 (0.02)	0.24 (0.03)	0.73 (0.07)	0.42 (0.04)	0.33 (0.04)	0.43 (0.05)
	DeepSurv	0.62 (0.04)	0.29 (0.03)	0.22 (0.03)	0.29 (0.04)	0.82 (0.08)	0.47 (0.06)	0.41 (0.05)	0.49 (0.06)
	DR-DL	0.58 (0.04)	0.26 (0.02)	0.16 (0.02)	0.27 (0.04)	0.78 (0.07)	0.44 (0.05)	0.36 (0.04)	0.45 (0.05)
	BJ-DL	0.57 (0.05)	0.25 (0.02)	0.15 (0.02)	0.26 (0.04)	0.77 (0.07)	0.43 (0.05)	0.35 (0.04)	0.44 (0.05)
	TSHEE	0.60 (0.04)	0.27 (0.03)	0.20 (0.03)	0.26 (0.04)	0.80 (0.08)	0.46 (0.06)	0.40 (0.04)	0.48 (0.06)
	GAPH	0.60 (0.04)	0.28 (0.03)	0.20 (0.03)	0.36 (0.05)	0.80 (0.07)	0.48 (0.05)	0.40 (0.04)	0.58 (0.07)
CR	Method	Proportional hazards				Nonproportional hazards			
		HS (i)	HS (ii)	HS (iii)	HS (iv)	HS (i)	HS (ii)	HS (iii)	HS (iv)
60%	AFT-L	4.14 (0.14)	4.57 (0.15)	4.45 (0.15)	5.19 (0.20)	4.39 (0.18)	4.89 (0.19)	4.91 (0.18)	5.88 (0.22)
	AFT-W	2.62 (0.12)	3.22 (0.13)	3.14 (0.13)	4.01 (0.16)	2.89 (0.16)	3.71 (0.14)	3.72 (0.13)	4.66 (0.18)
	CoxPH	2.63 (0.12)	3.21 (0.13)	3.22 (0.13)	3.89 (0.18)	4.41 (0.18)	4.83 (0.18)	4.83 (0.18)	5.93 (0.20)
	ABN	0.81 (0.04)	0.31 (0.04)	0.21 (0.02)	0.31 (0.03)	0.99 (0.09)	0.51 (0.06)	0.41 (0.04)	0.50 (0.05)
	ABNPS	0.82 (0.04)	0.32 (0.04)	0.22 (0.02)	0.32 (0.03)	1.00 (0.09)	0.52 (0.06)	0.42 (0.04)	0.51 (0.05)
	ABS	0.84 (0.04)	0.34 (0.04)	0.24 (0.02)	0.35 (0.03)	1.04 (0.09)	0.56 (0.06)	0.45 (0.04)	0.54 (0.05)
	RSFs	0.82 (0.05)	0.32 (0.05)	0.22 (0.02)	0.33 (0.03)	1.01 (0.09)	0.52 (0.06)	0.41 (0.04)	0.52 (0.05)
	DeepSurv	0.85 (0.04)	0.38 (0.04)	0.26 (0.03)	0.37 (0.04)	1.09 (0.11)	0.59 (0.07)	0.51 (0.05)	0.62 (0.06)
	DR-DL	0.84 (0.04)	0.35 (0.04)	0.24 (0.02)	0.40 (0.03)	1.04 (0.09)	0.55 (0.06)	0.45 (0.04)	0.56 (0.05)
	BJ-DL	0.83 (0.04)	0.35 (0.04)	0.23 (0.02)	0.38 (0.03)	1.03 (0.09)	0.53 (0.06)	0.43 (0.04)	0.54 (0.05)
	TSHEE	0.85 (0.05)	0.37 (0.05)	0.25 (0.03)	0.36 (0.04)	1.08 (0.10)	0.57 (0.07)	0.48 (0.05)	0.58 (0.06)
	GAPH	0.86 (0.04)	0.38 (0.04)	0.25 (0.02)	0.47 (0.05)	1.08 (0.09)	0.58 (0.06)	0.48 (0.05)	0.69 (0.07)

Abbreviations: AB, AFT-BART-NP; ABNPS, AFT-BART-NP-PS; ABS, AFT-BART-SP; AFT-L, AFT-Lognormal; AFT-W, AFT-Weibull; CR, censoring rate; HS, heterogeneity setting.

TABLE 2

Summary of average frequentist coverage probability of the Bayesian credible intervals from the AFT-BART-NP (denoted by X) and AFT-BART-NP-PS (denoted by $X + PS$) for five propensity score subclasses, for each of two sample sizes, proportional and nonproportional hazards assumptions, and under varying degrees of overlap, in the scenario of HS (iv) and 60% censoring

Sample size	G_k	Nonproportional hazards						Proportional hazards					
		Strong overlap		Moderate overlap		Weak overlap		Strong overlap		Moderate overlap		Weak overlap	
		X	$X + PS$	X	$X + PS$	X	$X + PS$	X	$X + PS$	X	$X + PS$	X	$X + PS$
$n = 5000$	1	0.93	0.94	0.90	0.90	0.50	0.55	0.92	0.94	0.91	0.92	0.49	0.56
	12	0.92	0.94	0.90	0.91	0.70	0.75	0.92	0.93	0.92	0.94	0.68	0.77
	25	0.94	0.95	0.92	0.93	0.92	0.93	0.93	0.93	0.92	0.94	0.92	0.94
	37	0.94	0.94	0.90	0.92	0.71	0.75	0.92	0.95	0.91	0.93	0.67	0.74
	50	0.94	0.95	0.90	0.91	0.50	0.56	0.91	0.94	0.90	0.92	0.50	0.57
$n = 500$	1	0.93	0.94	0.86	0.90	0.52	0.62	0.92	0.93	0.89	0.91	0.52	0.58
	12	0.90	0.94	0.90	0.90	0.70	0.73	0.93	0.94	0.90	0.92	0.67	0.74
	25	0.91	0.95	0.91	0.92	0.90	0.93	0.94	0.93	0.92	0.93	0.90	0.93
	37	0.90	0.94	0.89	0.90	0.70	0.77	0.92	0.94	0.92	0.93	0.66	0.72
	50	0.91	0.93	0.86	0.87	0.50	0.58	0.94	0.95	0.87	0.88	0.50	0.57

Note: The subclasses $G_k = 1$ and $G_k = 50$ include units with the most extreme propensity scores, with the propensity scores closest to zero in $G_k = 1$ and the propensity scores closest to one in $G_k = 50$. The numbers in each cell represent the mean coverage of the ISTE for the corresponding subclass.

TABLE 3

Mean (SD) of PEHE for each of 12 methods

CR		Proportional hazards				Nonproportional hazards			
		No U	U_1	U_2	U_3	No U	U_1	U_2	U_3
20%	AFT-L	4.13 (0.14)	5.03 (0.16)	5.82 (0.19)	6.52 (0.19)	4.41 (0.17)	5.16 (0.20)	5.77 (0.20)	6.47 (0.20)
	AFT-W	2.65(0.12)	3.13 (0.14)	3.73 (0.14)	4.43 (0.14)	2.81 (0.13)	3.35 (0.16)	4.03 (0.16)	4.71 (0.16)
	CoxPH	2.61 (0.13)	3.03 (0.14)	3.63 (0.14)	4.33 (0.14)	4.39 (0.13)	5.09 (0.20)	5.78 (0.20)	6.48 (0.20)
	ABN	0.12 (0.02)	0.31 (0.04)	0.46 (0.04)	0.66 (0.04)	0.33 (0.04)	0.50 (0.06)	0.65 (0.06)	0.85 (0.06)
	ABNPS	0.13 (0.02)	0.32 (0.04)	0.47 (0.04)	0.67 (0.04)	0.34 (0.04)	0.51 (0.06)	0.66 (0.06)	0.86 (0.06)
	ABS	0.15 (0.02)	0.35 (0.04)	0.50 (0.04)	0.70 (0.04)	0.36 (0.04)	0.53 (0.06)	0.68 (0.06)	0.88 (0.06)
	RSFs	0.13 (0.02)	0.32 (0.04)	0.47 (0.04)	0.67 (0.04)	0.33 (0.04)	0.51 (0.06)	0.66 (0.06)	0.86 (0.06)
	DeepSurv	0.22 (0.03)	0.38 (0.05)	0.53 (0.05)	0.73 (0.05)	0.41 (0.05)	0.58 (0.07)	0.73 (0.07)	0.93 (0.07)
	DR-DL	0.16 (0.02)	0.35 (0.04)	0.50 (0.04)	0.70 (0.04)	0.36 (0.04)	0.53 (0.06)	0.68 (0.06)	0.88 (0.06)
	BJ-DL	0.15 (0.02)	0.34 (0.04)	0.49 (0.04)	0.69 (0.04)	0.35 (0.04)	0.53 (0.06)	0.68 (0.06)	0.88 (0.06)
	TSHEE	0.20 (0.03)	0.37 (0.05)	0.52 (0.05)	0.72 (0.05)	0.40 (0.04)	0.57 (0.07)	0.72 (0.07)	0.92 (0.07)
	GAPH	0.20 (0.03)	0.37 (0.05)	0.52 (0.05)	0.72 (0.05)	0.40 (0.04)	0.57 (0.07)	0.72 (0.07)	0.92 (0.07)
CR		Proportional hazards				Nonproportional hazards			
		No U	U_1	U_2	U_3	No U	U_1	U_2	U_3
60%	AFT-L	4.45 (0.15)	5.33 (0.20)	6.04 (0.22)	6.78 (0.25)	4.91 (0.18)	5.99 (0.22)	6.56 (0.22)	7.27 (0.22)
	AFT-W	3.14 (0.13)	4.22 (0.16)	4.93 (0.19)	5.66 (0.22)	3.72 (0.13)	4.75 (0.18)	5.45 (0.18)	6.05 (0.18)
	CoxPH	3.22 (0.13)	3.94 (0.18)	4.52 (0.20)	5.13 (0.22)	4.83 (0.18)	6.02 (0.20)	6.64 (0.20)	7.24 (0.20)
	ABN	0.21 (0.02)	0.34 (0.03)	0.49 (0.04)	0.69 (0.05)	0.41 (0.04)	0.56 (0.05)	0.71 (0.05)	0.92 (0.05)
	ABNPS	0.22 (0.02)	0.35 (0.03)	0.50 (0.04)	0.70 (0.05)	0.42 (0.04)	0.57 (0.05)	0.73 (0.05)	0.93 (0.05)
	ABS	0.24 (0.02)	0.38 (0.03)	0.52 (0.04)	0.72 (0.05)	0.45 (0.04)	0.59 (0.05)	0.73 (0.05)	0.94 (0.05)
	RSFs	0.22 (0.02)	0.36 (0.03)	0.51 (0.04)	0.71 (0.05)	0.41 (0.04)	0.60 (0.05)	0.75 (0.05)	0.94 (0.05)
	DeepSurv	0.26 (0.03)	0.40 (0.04)	0.55 (0.05)	0.75 (0.06)	0.51 (0.05)	0.68 (0.06)	0.83 (0.06)	1.02 (0.06)
	DR-DL	0.24 (0.02)	0.44 (0.03)	0.59 (0.04)	0.79 (0.05)	0.45 (0.04)	0.65 (0.05)	0.80 (0.05)	1.00 (0.05)
	BJ-DL	0.23 (0.02)	0.42 (0.03)	0.56 (0.04)	0.76 (0.05)	0.43 (0.04)	0.64 (0.05)	0.79 (0.05)	0.99 (0.05)
	TSHEE	0.25 (0.02)	0.41 (0.04)	0.56 (0.05)	0.77 (0.06)	0.48 (0.05)	0.67 (0.06)	0.82 (0.06)	1.02 (0.06)
	GAPH	0.25 (0.02)	0.41 (0.04)	0.56 (0.05)	0.77 (0.06)	0.48 (0.05)	0.67 (0.06)	0.82 (0.06)	1.02 (0.06)

Note: Data are generated under heterogeneous setting (iii) with $n = 5000$ and strong overlap. U : unmeasured confounders; U_1 : X_3 is the unmeasured confounder; U_2 : X_3 and X_5 are the unmeasured confounders; U_3 : X_3 , X_5 , and X_6 are the unmeasured confounders.

Abbreviations: ABN, AFT-BART-NP; ABNPS, AFT-BART-NP-PS; ABS, AFT-BART-SP; AFT-L, AFT-Lognormal; AFT-W, AFT-Weibull; CR: censoring rate; HS: heterogeneity setting.