

Inference for treatment-specific survival curves using machine learning

Ted Westling

Department of Mathematics and Statistics
University of Massachusetts Amherst
twestling@umass.edu

Peter B. Gilbert

Vaccine and Infectious Disease Division
Fred Hutchinson Cancer Center
pgilbert@fredhutch.org

Alex Luedtke

Department of Statistics
University of Washington
aluedtke@uw.edu

Marco Carone

Department of Biostatistics
University of Washington
mcarone@uw.edu

Abstract

In the absence of data from a randomized trial, researchers may aim to use observational data to draw causal inference about the effect of a treatment on a time-to-event outcome. In this context, interest often focuses on the treatment-specific survival curves, that is, the survival curves were the population under study to be assigned to receive the treatment or not. Under certain conditions, including that all confounders of the treatment-outcome relationship are observed, the treatment-specific survival curve can be identified with a covariate-adjusted survival curve. In this article, we propose a novel cross-fitted doubly-robust estimator that incorporates data-adaptive (e.g. machine learning) estimators of the conditional survival functions. We establish conditions on the nuisance estimators under which our estimator is consistent and asymptotically linear, both pointwise and uniformly in time. We also propose a novel ensemble learner for combining multiple candidate estimators of the conditional survival estimators. Notably, our methods and results accommodate events occurring in discrete or continuous time, or an arbitrary mix of the two. We investigate the practical performance of our methods using numerical studies and an application to the effect of a surgical treatment to prevent metastases of parotid carcinoma on mortality.

1 Introduction

The gold standard for assessing the causal effect of a binary treatment on a time-to-event outcome is a randomized controlled trial in which participants are randomly assigned to treatment or control and followed over time. The effect of treatment may then be assessed by comparing the fraction of participants who experience the event by the end of the study in the treatment and control arms. However, some participants' outcomes may be unknown for various reasons, such as dropping out of the study or moving away from the study site, which is known as right-censoring of the event time. If the time of right-censoring is independent of the event time conditional on treatment arm, then contrasts of the stratified Kaplan-Meier estimators can be used to assess the treatment effect (Kaplan and Meier, 1958).

Randomizing treatment status is often infeasible or unethical, or preliminary evidence may be needed to justify the cost of conducting a randomized trial. In such cases, researchers may turn to observational data — obtained, for example, from cohort studies, registries, or electronic medical records. In such contexts, the treatment or exposure is not randomized, but instead assigned or selected according to an unknown mechanism. Assessing the causal effect of a treatment on a time-to-event outcome with observational data is challenging due to confounding of the treatment-outcome relationship. When there are confounding variables that affect the treatment selection or assignment process and also impact the outcome, simple approaches such as contrasts of stratified Kaplan-Meier estimators typically cannot be interpreted causally. Any observed differences (or lack thereof) in the outcome between those who received treatment and those who did not may be due to the confounding variables rather than the treatment. Even if the treatment is randomized, dependence of the event and censoring times can also render the Kaplan-Meier estimator inconsistent and resulting inference invalid.

If the recorded covariates are rich enough to de-confound the treatment-outcome, treatment-censoring, and outcome-censoring relationships, then a causal effect may still be recovered, and there are a variety of existing methods for doing so. The most common approach consists

of fitting a time-to-event regression model, such as a Cox proportional hazards model (Cox, 1972). If the Cox model holds, then the exponentiated regression coefficient corresponding to treatment can be interpreted as a conditional hazard ratio comparing treated and control patients. However, if the model does not hold, this interpretation fails. Even if the proportional hazards assumption holds for the treatment, but fails for the other covariates, the estimator of the hazard ratio of treatment based on Cox regression can be severely biased (Strandberg et al., 2014). Furthermore, the causal interpretation of the Cox model is complicated (Hernán, 2010; Martinussen, 2022), which has motivated various modifications of the model. Vansteelandt et al. (2022) considered an estimand that reduces to the hazard ratio under the Cox model, but also has a useful interpretation when the Cox model does not hold. Dukes et al. (2019) and Seaman et al. (2021) considered a class of semiparametric additive hazards models in which the causal survival ratio is a simple transformation of the Euclidean target parameter. Hou et al. (2023) considered doubly-robust estimation of the hazard ratio under an additive hazards model. Other authors have considered the mean residual life as causal parameter of interest (Mansourvar et al., 2016; Mansourvar and Martinussen, 2017).

Alternatively, any time-to-event regression model can be marginalized using the G-formula to obtain estimated treatment-specific survival curves corresponding to the hypothetical scenarios in which all patients are assigned to treatment or control (Makuch, 1982). However, if the model is misspecified, the resulting marginalized survival curves will typically be inconsistent. As an alternative to outcome regression models, inverse probability weighting may be used (Cole and Hernán, 2004). However, their consistency hinges on consistent estimation of both the treatment assignment mechanism and the censoring distribution. Finally, if the event time of interest is known to take values on a finite grid of time-points, then methods for longitudinal data can be used — see, e.g., Rotnitzky et al. (2012) and references therein. However, using a discrete-time approximation for an event truly occurring in continuous time generally yields inconsistent estimators (Ferreira Guerra et al., 2020).

Doubly-robust estimators combine regression and weighting estimators in such a way

that the bias of the resulting estimator is a product of the biases of the outcome regression and weighting function estimators. As a result, doubly-robust estimators are consistent if either the outcome regression or weighting function estimators are consistent. Furthermore, doubly-robust estimators can converge in distribution to a normal limit at the parametric rate even when flexible (e.g., machine learning) procedures are used to construct the outcome regression and weighting function estimators.

Several doubly-robust estimators of treatment-specific survival curves in continuous time have been proposed. Zeng (2004) proposed an estimator that is consistent as long as either the conditional time-to-event or censoring distributions follows a Cox model. Zhang and Schaubel (2012) proposed an estimator that is consistent as long as either the conditional time-to-event distribution follows a Cox model or the treatment assignment mechanism follows a logistic regression model. Finally, Hubbard et al. (2000) and Bai et al. (2013) proposed doubly-robust estimators based on semiparametric efficiency theory. However, comprehensive theoretical results regarding these estimators have not been developed, and both articles suggested using common semiparametric regression models to estimate the conditional time-to-event and censoring distributions. To the best of our knowledge, the use of machine learning techniques for doubly-robust estimation of treatment-specific survival curves (and contrasts thereof) permitting events occurring in continuous time has not yet been studied.

In this article, we fill this gap in the literature. Specifically, we make the following contributions: (1) we propose a novel cross-fitted one-step estimator of the treatment-specific survival curve (Section 3); (2) we provide detailed conditions, which permit the use of machine learning for nuisance function estimation, under which our estimator is (uniformly) consistent and (uniformly) asymptotically linear (Section 4); (3) we propose methods for pointwise and uniform inference (Section 5); (4) we propose a novel ensemble learner for combining multiple candidate estimators of the conditional survival functions (Section 6). In addition, we conduct a numerical study and apply our results to assess the effect of elective neck dissection on all-cause mortality using an observational cohort of patients with

parotid carcinomas (Sections 7 and 8). Importantly, the methods we propose and conditions we derive all permit the event and censoring times to be discrete, continuous, or an arbitrary mixture of the two. Furthermore, the nature of the time scale does not need to be known or specified. This is important in many applications. For instance, censoring often involves a mix of loss to follow-up occurring in continuous time and administrative censoring occurring in discrete time. Similarly, composite event times defined as the earlier of two or more event times may have both discrete and continuous components.

We have made the estimator and associated inferential procedures proposed here available through the R package `CFsurvival` (<https://github.com/tedwestling/CFsurvival>), and we have implemented the method proposed in Section 6 for estimating conditional survival functions in the package `survSuperLearner` (<https://github.com/tedwestling/survSuperLearner>).

2 Statistical setting and parameters of interest

2.1 Ideal and observed data structures

We now define the ideal data structure we consider in temporal order. As we discuss below, we only observe a coarsening of this ideal data structure. First, we record a vector W of baseline covariates taking values in $\mathcal{W} \subseteq \mathbb{R}^d$. Throughout, the dimension d of the covariates is fixed. After recording W , but prior to time $t = 0$, we observe a binary exposure $A \in \{0, 1\}$. Adopting the Neyman-Rubin potential outcomes framework (Neyman, 1923; Rubin, 1974), we let $T(a)$ be the event time of interest under assignment to exposure $A = a$. We assume that for $a \in \{0, 1\}$, $T(a)$ takes values in $(0, \infty]$. Since we assume that $T(a) > 0$, all patients start the study without having experienced the event of interest, and since we allow $T(a) = \infty$, some patients may never experience the event. We then let $C(a)$ be a right-censoring time under assignment to exposure $A = a$, and we assume that $C(a) \in [0, \infty]$. Since we allow $C(a) = 0$, patients may be censored immediately if, for instance, a patient is lost to follow-up just after the exposure A is recorded. We define $O_F := (W, A, T(0), T(1), C(0), C(1))$ to be

the ideal data unit, and denote by $P_{0,F}$ the distribution of O_F . We assume that each patient's potential event and censoring times are independent of all other patients' exposures.

We now describe the coarsened version of O_F actually observed. We denote by $T := T(A)$ and $C := C(A)$ the event and censoring times corresponding to the exposure received. We assume that the right-censored time $Y := \min\{T, C\}$ and the event indicator $\Delta := I(T \leq C)$ are observed for each patient. Thus, the available data consist of n independent and identically distributed observations O_1, O_2, \dots, O_n of the observed data unit $O := (W, A, Y, \Delta)$. We denote by P_0 the distribution of the observed data unit, as induced by the distribution $P_{0,F}$ of the ideal data unit.

We denote summaries of P_0 with the subscript 0, e.g., $E_0[f(O)] := E_{P_0}[f(O)]$, and summaries of $P_{0,F}$ with subscript 0, F . In cases where f is a random function, the expectation $E_0[f(O)]$ should be understood as being taken with respect to the distribution of the random unit O , but not the function f . In addition, we let $a \wedge b$ denote $\min\{a, b\}$, \mathbb{P}_n be the empirical distribution corresponding to O_1, O_2, \dots, O_n , and $Pf := \int f(o) dP(o)$ for any probability measure P and P -measurable function f . For a function $f : \mathbb{R} \rightarrow \mathbb{R}$ that is left-continuous at $x \in \mathbb{R}$, we let $f(x-) := \lim_{u \uparrow x} f(u)$. Finally, we use the convention $0/0 := 1$.

2.2 Causal parameter of interest and identification

In this article, we are interested in the causal survival curves $t \mapsto \theta_{0,F}(t, a) := P_{0,F}(T(a) > t)$ for $a \in \{0, 1\}$ and $t \in [0, \tau]$ for some positive $\tau < \infty$. Thus defined, $\theta_{0,F}(t, 0)$ represents the population probability that a patient would experience the event later than time t if, possibly contrary to fact, the patient received the control exposure ($A = 0$), and $\theta_{0,F}(t, 1)$ represents the same if the patient received the treatment under study ($A = 1$). In addition to the exposure-specific survival curves $t \mapsto \theta_{0,F}(t, 0)$ and $t \mapsto \theta_{0,F}(t, 1)$, we are interested in survival contrasts including the survival difference $t \mapsto \theta_{0,F}(t, 1) - \theta_{0,F}(t, 0)$, survival ratio $t \mapsto \theta_{0,F}(t, 1)/\theta_{0,F}(t, 0)$, and risk ratio $t \mapsto [1 - \theta_{0,F}(t, 1)]/[1 - \theta_{0,F}(t, 0)]$ functions. Under certain conditions, we can identify the causal parameter $\theta_{0,F}(t, a)$ in terms of the distribution

P_0 of the observed data unit. These causal identifications have been established before, so we only briefly review them. A full set of identification conditions and additional discussion may be found in Supplementary Material.

We define $S_0(t | a, w) := \mathcal{P}_{(0,t]} \{1 - \Lambda_0(du | a, w)\}$, where \mathcal{P} denotes the Riemann-Stieltjes product integral (Gill and Johansen, 1990) and $\Lambda_0(t | a, w) := \int_0^t \frac{F_{0,1}(du | a, w)}{R_0(u | a, w)}$ for $F_{0,\delta}(t | a, w) := P_0(Y \leq t, \Delta = \delta | A = a, W = w)$ and $R_0(t | a, w) := P_0(Y \geq t | A = a, W = w)$. If the covariates W contain all common causes of the exposure-outcome, exposure-censoring, and outcome-censoring relationships, and certain positivity assumptions hold, then

$$\theta_{0,F}(t, a) = \theta_0(t, a) := E_0[S_0(t | a, W)] \quad (1)$$

for every $t \in [0, \tau]$ and $a \in \{0, 1\}$. This result is a combination of the G-formula (also known as the backdoor or the regression standardization formula) from causal inference (Robins, 1986; Gill and Robins, 2001) and the identification of a survival function in the context of dependent censoring (Beran, 1981; Dabrowska, 1989). The product integral reduces to a product in the case of events occurring in discrete time, and to $\exp\{-\Lambda_0(t | a, w)\}$ in the case of events occurring in continuous time. We use the product integral in order to permit either of these cases or a mix of the two. We note that this identification does not permit time-varying common causes of the event and censoring times because time-varying covariates would occur post-exposure, and hence possibly be on the causal pathway between A and T or A and C . To permit time-varying common causes of the event and censoring times, an alternative approach is to use an iterated G-formula based on sequential exchangeability as in longitudinal studies (Rotnitzky et al., 2012). If the event and censoring times do not naturally occur on a discrete grid, this approach would typically require choosing a discretization of time.

The parameter $\theta_0(t, a)$ is referred to as the *G-computed probability* that the event T occurs after time t given that exposure A is set to a . This parameter measures the survival

probability under exposure $A = a$ while adjusting for potential confounding between the exposure and the event of interest and for dependence between the event and censoring times. The curves $\{\theta_0(t, 0) : t \in [0, \tau]\}$ and $\{\theta_0(t, 1) : t \in [0, \tau]\}$, and contrasts thereof, are the observed-data statistical parameters we focus on.

3 Estimation

3.1 Efficiency calculations

In this section, we describe the proposed methodology for nonparametric efficient estimation of the treatment-specific G-computed survival function $\{\theta_0(t, a) : t \in [0, \tau]\}$. The definition of our estimator involves three steps. First, we characterize the behavior of nonparametric efficient estimators of the estimand of interest; this is done by deriving the nonparametric efficient influence function (EIF) of $\theta_0(t, a)$ for each t and a . Second, we use the EIF to construct an efficient estimator of $\theta_0(t, a)$ for each t and a . Finally, we describe a simple procedure to ensure that the resulting survival curves are monotone.

We define $\pi_0(a | w) := P_0(A = a | W = w)$. Below, we make use of the fact that, for any (a, w) such that $S_0(\tau - | a, w) > 0$, $P_{0,F}(C \geq t | A = a, W = w)$ can be identified under the causal identifiability conditions as $P_{0,F}(C \geq t | A = a, W = w) = G_0(t | a, w) := \prod_{[0,t)} \{1 - H_0(du | a, w)\}$ for any $t \in [0, \tau]$, where $H_0(u | a, w) := \int_{[0,u]} \left\{ \frac{S_0(s | a, w)}{S_0(s- | a, w)} \right\} \frac{F_{0,0}(ds | a, w)}{R_0(s | a, w)}$. We emphasize that G_0 is defined as the left-continuous conditional survival function of C , whereas S_0 is defined as the right-continuous conditional survival function of T . Even when $t = \tau$, it is sufficient that $S_0(t - | a, w) > 0$ in the above identification because the product integral is taken up to, but not including, t .

In Theorem 1, we present the nonparametric efficient influence function of $\theta_0(t, a_0)$, where we use a_0 rather than a to denote the exposure value of interest in order not to confuse values of the random variable A with the specific a_0 at which we want to evaluate θ_0 .

Theorem 1. *If there exists $\eta > 0$ such that $\min\{\pi_0(a_0 | w), G_0(t | a_0, w)\} \geq \eta$ for P_0 -almost every w such that $S_0(t | a_0, w) > 0$, then $\theta_0(t, a_0)$ is a pathwise differentiable parameter in a nonparametric model with efficient influence function $\phi_{0,t,a_0}^* := \phi_{0,t,a_0} - \theta_0(t, a_0)$, where $\phi_{0,t,a_0}(y, \delta, a, w)$ equals*

$$S_0(t | a_0, w) \left[1 - \frac{I(a = a_0)}{\pi_0(a | w)} \left\{ \frac{I(y \leq t, \delta = 1)}{S_0(y | a, w)G_0(y | a, w)} - \int_0^{t \wedge y} \frac{\Lambda_0(du | a, w)}{S_0(u | a, w)G_0(u | a, w)} \right\} \right].$$

Hubbard et al. (2000) and Bai et al. (2013) also derived the efficient influence function of $\theta_0(t, a_0)$ in a nonparametric model. Their influence functions have a different form than ours; in particular, their influence functions involve a conditional martingale residual process. However, since there is only one influence function in a nonparametric model, the influence function provided in Theorem 1 is algebraically equivalent to those of Hubbard et al. (2000) and Bai et al. (2013).

3.2 Cross-fitted one-step estimator

The efficient influence function ϕ_{0,t,a_0}^* involves three variation-independent nuisance functions: S_0 , G_0 and π_0 . We discuss estimation of these functions in Section 6. We note that Λ_0 and S_0 are in one-to-one correspondence with one another, so estimating S_0 gives an estimator of Λ_0 and vice-versa. Given estimators S_n , G_n and π_n of S_0 , G_0 and π_0 , respectively, there are multiple possible asymptotically linear and efficient estimators of $\theta_0(t, a)$. Denoting by $\phi_{n,t,a}$ the function $\phi_{0,t,a}$ with S_0 , Λ_0 , G_0 and π_0 replaced by their respective estimators, the standard one-step estimator would be $\mathbb{P}_n \phi_{n,t,a}$, which is also an estimating equations-based estimator in this case because the influence function is linear in $\theta_0(t, a)$. This was the approach taken by Hubbard et al. (2000) and Bai et al. (2013).

Asymptotic linearity of estimators of this type depend on nuisance estimators in two important ways. First, negligibility of a so-called *second-order* remainder term requires that the nuisance parameters converge at fast enough rates to their true counterparts. Second, negli-

gibility of an *empirical process* remainder term can be guaranteed if the nuisance estimators fall in sufficiently small function classes with probability tending to one. In observational studies, researchers can rarely specify correct parametric models for nuisance parameters a priori, which motivates the use of data-adaptive estimators. However, data-adaptive estimators typically fail to fall in small function classes. This poses a challenge in simultaneously achieving negligibility of these two remainder terms. Cross-fitting has been found to resolve this challenge by removing this constraint on the complexity of nuisance estimators (see, e.g. Bickel, 1982; Robins et al., 2008; Zheng and van der Laan, 2011; Díaz, 2019a). Therefore, we will use a cross-fitted version of the one-step estimator stated above, which we now define.

For a deterministic integer $K \in \{2, 3, \dots, \lfloor n/2 \rfloor\}$, we randomly partition the indices $\{1, 2, \dots, n\}$ into K disjoint sets $\mathcal{V}_{n,1}, \mathcal{V}_{n,2}, \dots, \mathcal{V}_{n,K}$ with cardinalities n_1, n_2, \dots, n_K . We require that these sets be of as close to equal sizes as possible, so that $|n_k - n/K| \leq 1$ for each k , and that the number of folds K be bounded as n grows. For each $k \in \{1, 2, \dots, K\}$, we define $\mathcal{T}_{n,k} := \{O_i : i \notin \mathcal{V}_{n,k}\}$ as the *training set* for fold k . We then define $S_{n,k}, G_{n,k}, \pi_{n,k}$ and $\Lambda_{n,k}$ as nuisance estimators estimated using only the observations from the training set $\mathcal{T}_{n,k}$, and $\phi_{n,k,t,a}$ as the function $\phi_{0,t,a}$ in which these nuisance estimators have substituted their true counterparts. We then define the cross-fitted one-step estimator θ_n pointwise as

$$\theta_n(t, a) := \frac{1}{n} \sum_{k=1}^K \sum_{i \in \mathcal{V}_{n,k}} \phi_{n,k,t,a}(O_i) . \quad (2)$$

Once the nuisance functions are estimated, $\theta_n(t, a)$ can be efficiently computed for many time-points t because the same nuisance function estimators can be re-used for each t .

How the integral in $\phi_{n,k,t,a}$ is computed depends on the form of $S_{n,k}$. If $S_{n,k}$ is defined as a step function, then the integral reduces to a sum. Otherwise, the integral can be approximated as a sum. As illustrated in numerical studies in Supplementary Material, a very fine grid can be used for this purpose with little impact on computational cost.

3.3 Enforcing monotonicity of the proposed estimator

The function $t \mapsto \theta_0(t, a)$ is necessarily monotone non-increasing for each $a \in \{0, 1\}$ and takes values in $[0, 1]$. However, the proposed estimator $\theta_n(t, a)$ is generally neither guaranteed to lie in $[0, 1]$ nor to be monotone in t in any given sample.

We ensure that our final estimator satisfies the above parameter constraints as follows. First, we construct $\theta_n(t, a)$ as defined above for each $t \in \mathcal{T}_n$, where \mathcal{T}_n is the set of unique values of Y_1, Y_2, \dots, Y_n . Second, for each $t \in \mathcal{T}_n$ and $a \in \{0, 1\}$, we define $\theta_n^+(t, a) = \theta_n(t, a)$ if $\theta_n(t, a) \in [0, 1]$, $\theta_n^+(t, a) = 1$ if $\theta_n(t, a) > 1$, and $\theta_n^+(t, a) = 0$ if $\theta_n(t, a) < 0$. Next, for each $a \in \{0, 1\}$, we define $\{\theta_n^\circ(t, a) : t \in \mathcal{T}_n\}$ as the projection of $\{\theta_n^+(t, a) : t \in \mathcal{T}_n\}$ onto the space of non-increasing functions using isotonic regression. For any $t \in (0, \tau]$, we then define $\theta_n^\circ(t, a)$ as the evaluation of the right-continuous stepwise interpolation of $\{\theta_n^\circ(t, a) : t \in \mathcal{T}_n\}$. The projected estimator θ_n° is guaranteed to be no farther from θ_0 than θ_n in every finite sample, and if the true function is strictly decreasing, then the initial and projected estimators are asymptotically equivalent (Westling et al., 2020). Therefore, in what follows, we focus on providing large-sample results for θ_n , since results for the isotonized estimator θ_n° are identical in view of the general results of Westling et al. (2020).

4 Large-sample properties

4.1 Consistency

In this section, we study the large-sample properties of our estimator. First, we provide conditions under which $\theta_n(t, a)$ is consistent for $\theta_0(t, a)$ for fixed t and uniformly over t .

(B1) There exist π_∞ , G_∞ and S_∞ such that:

$$(a) \max_k E_0 \left[\frac{1}{\pi_{n,k}(a|W)} - \frac{1}{\pi_\infty(a|W)} \right]^2 \xrightarrow{\text{P}} 0;$$

$$(b) \max_k E_0 \left[\sup_{u \in [0, t]} \left| \frac{1}{G_{n,k}(u|a,W)} - \frac{1}{G_\infty(u|a,W)} \right| \right]^2 \xrightarrow{\text{P}} 0;$$

$$(c) \max_k E_0 \left[\sup_{u \in [0,t]} \left| \frac{S_{n,k}(t \mid a, W)}{S_{n,k}(u \mid a, W)} - \frac{S_\infty(t \mid a, W)}{S_\infty(u \mid a, W)} \right| \right]^2 \xrightarrow{\text{P}} 0.$$

(B2) There exists $\eta > 0$ such that, with probability tending to one, for P_0 -almost all w , $\pi_{n,k}(a \mid w) \geq 1/\eta$, $\pi_\infty(a \mid w) \geq 1/\eta$, $G_{n,k}(t \mid a, w) \geq 1/\eta$, and $G_\infty(t \mid a, w) \geq 1/\eta$.

(B3) For P_0 -almost all w , there exist measurable sets $\mathcal{S}_w, \mathcal{G}_w \subseteq [0, t]$ such that $\mathcal{S}_w \cup \mathcal{G}_w = [0, t]$ and $\Lambda_0(u \mid a, w) = \Lambda_\infty(u \mid a, w)$ for all $u \in \mathcal{S}_w$ and $G_0(u \mid a, w) = G_\infty(u \mid a, w)$ for all $u \in \mathcal{G}_w$. In addition, if \mathcal{S}_w is a strict subset of $[0, t]$, then $\pi_0(a \mid w) = \pi_\infty(a \mid w)$ as well.

(B4) It holds that $\max_k E_0 \left[\sup_{u \in [0,t]} \sup_{v \in [0,u]} \left| \frac{S_{n,k}(u \mid a, W)}{S_{n,k}(v \mid a, W)} - \frac{S_\infty(u \mid a, W)}{S_\infty(v \mid a, W)} \right| \right]^2 \xrightarrow{\text{P}} 0$.

Theorem 2 (Consistency). *If conditions (B1)–(B3) hold, then $\theta_n(t, a) \xrightarrow{\text{P}} \theta_0(t, a)$. If condition (B4) also holds, then $\sup_{u \in [0,t]} |\theta_n(u, a) - \theta_0(u, a)| \xrightarrow{\text{P}} 0$.*

Condition (B1) require that the estimated functions converge in an appropriate sense to fixed limit functions, which is used to control certain empirical process terms. Condition (B4) requires a slightly stronger condition on the convergence of $S_{n,k}$ to its limit S_∞ for uniform consistency. Condition (B1)(c) depends on t through the numerators $S_{n,k}(t \mid a, W)$ and $S_\infty(t \mid a, W)$. For uniform convergence, this term needs to be controlled uniformly over t , so the extra supremum is needed in condition (B4). We note again that the expectations in conditions (B1) and (B4) are with respect to W , and not with respect to the randomness of the nuisance estimators. Condition (B2) ensures that the estimated propensity and censoring functions are bounded uniformly away from zero in all subpopulation of patients defined by W . In practice, this can be guaranteed by truncating the estimated propensities and censoring probabilities. We note that there is no restriction on complexity of these nuisance function estimators; the lack of such a condition is due to the use of cross-fitting.

Condition (B3) requires that, for almost all (t, w) , either $S_\infty(t \mid a, w) = S_0(t \mid a, w)$ or *both* $G_\infty(t \mid a, w) = G_0(t \mid a, w)$ and $\pi_\infty(a \mid w) = \pi_0(a \mid w)$. In combination with condition (B1), this implies that for almost all (t, w) , either $S_n(t \mid a, w)$ or *both* $G_n(t \mid a, w)$ and $\pi_n(a \mid w)$

are consistent. This is a form of *double-robustness* of the estimator θ_n to estimation of the nuisances S_0 and (G_0, π_0) , because, in particular, θ_n is consistent when either S_n is consistent everywhere or both G_n and π_n are consistent everywhere. However, condition (B3) is a relaxed form of doubly-robustness since none of the limit functions need to be identically equal to their true counterparts. Instead, consistency of S_n or (G_n, π_n) may vary with t and w . This is akin to *sequential doubly-robustness* or 2^K -*robustness* in longitudinal studies, where there are 2^K possible ways to achieve consistency for a G -computation parameter in a longitudinal study with K time-points (Tchetgen Tchetgen, 2009; Molina et al., 2017; Luedtke et al., 2018; Rotnitzky et al., 2017; Díaz, 2019b). In our setting, there are infinitely many ways to achieve consistency due to the use of continuous time. In longitudinal studies with a time-varying treatment, specialized methods are needed to achieve 2^K -robustness. In contrast, here, we achieve our version of 2^K -robustness with a standard one-step estimator due to our consideration of a baseline, rather than time-varying, treatment. With a baseline treatment, the counterfactual mean is identified with a single G -computation rather than K iterated G -computations as in the case of a time-varying treatment over K time points. Our simpler observed-data parameter yields a simpler efficient influence function that permits 2^K -robustness without specialized methodology.

4.2 Asymptotic linearity

We now present additional conditions under which $\theta_n(t, a)$ is asymptotically linear for fixed t and uniformly over t . We define

$$\begin{aligned} r_{n,t,a,1} &:= \max_k E_0 |\{\pi_{n,k}(a | W) - \pi_0(a | W)\}\{S_{n,k}(t | a, W) - S_0(t | a, W)\}|; \\ r_{n,t,a,2} &:= \max_k E_0 \left| S_{n,k}(t | a, W) \int_0^t \left\{ \frac{G_0(u | a, W)}{G_{n,k}(u | a, W)} - 1 \right\} \left(\frac{S_0}{S_{n,k}} - 1 \right) (du | a, W) \right|. \end{aligned}$$

Based on these quantities, we introduce additional conditions for asymptotic linearity:

(B5) It holds that $r_{n,t,a,1} = o_P(n^{-1/2})$ and $r_{n,t,a,2} = o_P(n^{-1/2})$.

(B6) It holds that $\sup_{u \in [0, t]} r_{n,u,a,1} = o_P(n^{-1/2})$ and $\sup_{u \in [0, t]} r_{n,u,a,2} = o_P(n^{-1/2})$.

We have the following result concerning the asymptotic linearity of $\theta_n(t, a)$.

Theorem 3 (Asymptotic linearity). *If conditions (B1)–(B2) hold with $S_\infty = S_0$, $G_\infty = G_0$ and $\pi_\infty = \pi_0$ and condition (B5) also holds, then $\theta_n(t, a) = \theta_0(t, a) + \mathbb{P}_n \phi_{0,t,a}^* + o_P(n^{-1/2})$. In particular, $n^{1/2}[\theta_n(t, a) - \theta_0(t, a)]$ then converges in distribution to a normal random variable with mean zero and variance $\sigma_0^2(t, a) := P_0 \phi_{0,t,a}^{*2}$. If in addition conditions (B4) and (B6) also hold, then*

$$\sup_{u \in [0, t]} |\theta_n(u, a) - \theta_0(u, a) - \mathbb{P}_n \phi_{0,u,a}^*| = o_P(n^{-1/2}).$$

In particular, $\{n^{1/2}[\theta_n(u, a) - \theta_0(u, a)] : u \in [0, t]\}$ then converges weakly as a process in the space $\ell^\infty([0, t])$ of uniformly bounded functions on $[0, t]$ to a tight mean zero Gaussian process with covariance function $(u, v) \mapsto P_0(\phi_{0,u,a}^ \phi_{0,v,a}^*)$.*

Condition (B5) requires roughly that the rates of convergence of $(S_n - S_0)(\pi_n - \pi_0)$ and $(S_n - S_0)(G_n - G_0)$ to zero be faster than $n^{-1/2}$. One approach to satisfying this condition is to assume that these nuisance functions fall in known parametric or semiparametric families such that existing estimators achieve the stipulated rates. For instance, if S_0 and G_0 follow the Cox proportional hazard model (Cox, 1972) and π_0 the logistic regression model, and model-based maximum likelihood estimators are used to obtain S_n , G_n and π_n , the required rates will be achieved for W of any fixed dimension. In some cases, scientific knowledge of the problem at hand can be used to guide selection of nuisance estimators guaranteed to achieve given rates of convergence. For example, in the context of a randomized experiment, the true propensity score π_0 is known, so it can be used directly. Alternatively, in some settings, information about the censoring mechanism may be available, such as knowledge that patients are only censored for administrative reasons. In such cases, it may be possible to design a parametric or semiparametric model guaranteed to contain the censoring distribution. If the propensity and censoring estimators both achieve $n^{-1/2}$ rates of convergence, then S_n only needs to be consistent for the conditions of Theorem 3 to hold.

In many cases, it is not possible to construct parametric or semiparametric models that are known to be correctly specified. Importantly, condition (B5) can be satisfied when using data-adaptive estimators since $(S_n - S_0)(\pi_n - \pi_0)$ and $(S_n - S_0)(G_n - G_0)$ can converge faster than $n^{-1/2}$ even if S_n , π_n , and/or G_n converge slower than $n^{-1/2}$. This is a primary benefit of constructing an estimator based on the nonparametric efficient influence function. However, achieving these rates of convergence is not guaranteed when using data-adaptive estimators. Whether the rate is achieved depends on the dimension of the covariates, structure such as smoothness, additivity, or sparsity of S_0 , G_0 , and π_0 , and the extent to which the nuisance estimators adapt to this structure. Since the true structure of these functions is often unknown, we recommend combining multiple candidate parametric, semiparametric, and nonparametric estimators using cross-validation, which has the potential to achieve the same rate as the best candidate estimator. We discuss this more in Section 6.

5 Pointwise and uniform inference

5.1 Pointwise inference

Theorem 3 can be used to conduct asymptotically valid pointwise and uniform inference for $\theta_0(t, 0)$, $\theta_0(t, 1)$ and contrasts thereof. Specifically, $\theta_n^\circ(t, a) \pm z_{1-\alpha/2} n^{-1/2} \sigma_n(t, a)$ is a Wald-type asymptotic $(1 - \alpha)$ -level confidence interval for $\theta_0(t, a)$, where z_p denotes the p -quantile of the standard normal distribution and $\sigma_n^2(t, a) := \frac{1}{n} \sum_{k=1}^K \sum_{i \in \mathcal{V}_{n,k}} [\phi_{n,k,t,a}(O_i) - \theta_n^\circ(t, a)]^2$ is a cross-fitted influence function-based estimator of the asymptotic variance $\sigma_0^2(t, a)$. However, since constructing Wald-type intervals on the logistic probability scale has been found to improve finite-sample coverage in classical settings (Anderson et al., 1982), we suggest this approach as well. Defining $\text{expit}(x) := \exp(x) / \{1 + \exp(x)\}$ for any $x \in \mathbb{R}$ and $\text{logit}(u) := \log(u) - \log(1 - u)$ for any $u \in (0, 1)$, and defining $\tilde{\sigma}_n(t, a) := \sigma_n(t, a) / \{\theta_n^\circ(t, a)[1 - \theta_n^\circ(t, a)]\}$, we propose the transformed Wald-type interval $[\ell_n(t, a), u_n(t, a)] := \text{expit}\{\text{logit}[\theta_n^\circ(t, a)] \pm z_{1-\alpha/2} n^{-1/2} \tilde{\sigma}_n(t, a)\}$ for any (t, a) for which $\theta_n^\circ(t, a) \in (0, 1)$. If $\theta_n^\circ(t, a) = 0$, we set $\ell_n(t, a) := 0$

and $u_n(t, a) := \min_s \{u_n(s, a) : u_n(s, a) > 0\}$, whereas if $\theta_n^\circ(t, a) = 1$, we set $\ell_n(t, a) := \max_s \{\ell_n(s, a) : \ell_n(s, a) < 1\}$ and $u_n(t, a) := 1$. The endpoints of this interval will be strictly contained between 0 and 1 for any (t, a) such that $\theta_n^\circ(t, a) \in (0, 1)$.

5.2 Uniform inference

If the uniform statement of Theorem 3 holds with $t = \tau$, then it can be used to construct asymptotically valid uniform confidence bands for $t \mapsto \theta_0(t, a)$ over $t \in [0, \tau]$, that is, to construct functions $t \mapsto \ell_n(t, a)$ and $t \mapsto u_n(t, a)$ such that $P_0\{\ell_n(t, a) \leq \theta_0(t, a) \leq u_n(t, a) \text{ for all } t \in [0, \tau]\}$ converges to $1 - \alpha$. The simplest such band is a fixed-width band with endpoints $\theta_n^\circ(t, a) \pm n^{-1/2}c_{n,a,\alpha}$. Here, $c_{n,a,\alpha}$ is any consistent estimator of the $(1 - \alpha)$ -quantile $c_{0,a,\alpha}$ of the supremum of the absolute value of the Gaussian process to which $\{n^{1/2}[\theta_n(t, a) - \theta_0(t, a)] : t \in [0, \tau]\}$ converges weakly, that is, a mean zero Gaussian process with covariance function $(u, v) \mapsto \Sigma_0(u, v, a) := P_0(\phi_{0,u,a}^* \phi_{0,v,a}^*)$. To obtain $c_{n,a,\alpha}$, we simulate sample paths of a Gaussian process on $[0, \tau]$ with covariance function given by the cross-fitted covariance estimator $(u, v) \mapsto \Sigma_n(u, v, a) := \frac{1}{n} \sum_{k=1}^K \sum_{i \in \mathcal{V}_{n,k}} [\phi_{n,k,u,a}(O_i) - \theta_n^\circ(u, a)][\phi_{n,k,v,a}(O_i) - \theta_n^\circ(v, a)]$. We then set $c_{n,a,\alpha}$ as the sample $(1 - \alpha)$ -quantile of the uniform norm over $[0, \tau]$ of these sample paths. Finally, we ensure monotonicity of these bands using isotonic regression, which can only increase their coverage, as established in Westling et al. (2020). While this fixed-width band is appealing in its simplicity, it does not reflect the variability in the uncertainty around $\theta_n^\circ(t, a)$ for different t . For instance, there is typically less uncertainty near $t = 0$, when few patients have been censored and the survival probability remains close to one, than elsewhere. An equal-width band will not reflect this.

An alternative confidence band that adapts to the variability in uncertainty over $[0, \tau]$ and is guaranteed to lie strictly within $(0, 1)$ can be formed by use of standard error scaling. The proposed variable-width confidence band is given by $\text{expit}\{\text{logit}[\theta_n^\circ(t, a)] \pm \tilde{c}_{n,\alpha} n^{-1/2} \tilde{\sigma}_n(t, a)\}$. Here, $\tilde{c}_{n,\alpha}$ is the $(1 - \alpha)$ -quantile of the uniform norm over $[0, \tau]$ of the sample paths of a mean zero Gaussian process with covariance function $(u, v) \mapsto \tilde{\Sigma}_n(u, v, a) :=$

$\Sigma_n(u, v, a)/\{\tilde{\sigma}_n(u, a)\tilde{\sigma}_n(v, a)\}$. However, since $\lim_{t \rightarrow 0} \sigma_n(t, a) = \lim_{t \rightarrow t^+} \sigma_n(t, a) = 0$ for $t^+ = \inf\{t : \theta_n(t, a) = 0\}$, these sample paths are unbounded near $t = 0$ and $t = t^+$. Therefore, this method of constructing confidence bands can only produce asymptotically valid bands on intervals of the form $[t_0, t_1]$ for $t_0 > 0$ and $t_1 < t^+$. Given $[t_0, t_1]$, we then proceed in constructing the band using the approximate critical value $\tilde{c}_{n,\alpha}$ obtained as the sample $(1 - \alpha)$ -quantile of the uniform norms over $[t_0, t_1]$ of the above sample paths. As before, we ensure monotonicity of these bands using isotonic regression. In practice, we suggest choosing t_0 and t_1 based on the quantiles of the observed event times.

5.3 Inference on causal effects

If the pointwise statement of Theorem 3 holds for both $a = 0$ and $a = 1$, then $n^{1/2}[\theta_n(t, 0) - \theta_0(t, 0)]$ and $n^{1/2}[\theta_n(t, 1) - \theta_0(t, 1)]$ converge jointly to a mean zero bivariate normal distribution. This fact can be used in conjunction with the delta method to perform inference on causal effects of the form $h(\theta_0(t, 0), \theta_0(t, 1))$ for any differentiable h . Similarly, if the uniform statement of Theorem 3 holds for both $a = 0$ and $a = 1$, then the processes $\{n^{1/2}[\theta_n(t, 0) - \theta_0(t, 0)] : t \in [0, \tau]\}$ and $\{n^{1/2}[\theta_n(t, 1) - \theta_0(t, 1)] : t \in [0, \tau]\}$ converge jointly as processes to correlated Gaussian process limits, so that uniform confidence bands can be constructed for causal effects in much the same way as described above. For risk and survival ratios, these confidence bands are only valid on intervals over which the denominator is bounded away from zero.

To test the null hypothesis $H_0 : \theta_0(t, 0) = \theta_0(t, 1)$ for all $t \in [0, \tau]$ against the complementary alternative, we propose using a test statistic of the form $n^{1/2} \int_0^\tau |\theta_n^\circ(t, 1) - \theta_n^\circ(t, 0)| \Omega_n(dt)$, where Ω_n is a user-specified, possibly data-dependent weight function. Under the null hypothesis, this test statistic converges in distribution to $\int_0^\tau |G_0(t)| \Omega_0(dt)$ by the continuous mapping theorem for G_0 denoting the limiting Gaussian process of $\{n^{1/2}[\{\theta_n(t, 1) - \theta_0(t, 1)\} - \{\theta_n(t, 0) - \theta_0(t, 0)\}] : t \in [0, \tau]\}$ and Ω_0 the deterministic in-probability limit of Ω_n . This limit distribution can be estimated by simulating Gaussian processes using the estimated

covariance matrices in a similar manner as discussed above, which can then be used to find a p -value for the test using the observed test statistic. The user-specified weight function Ω_n can be chosen to improve power against particular alternatives that may be expected based on the scientific context, such as early or late differences in survival, as has been done in the context of log-rank tests for uninformative censoring (see, e.g. Harrington and Fleming, 1982; Wu and Gilbert, 2002).

Our results can also be used to make inference on functionals of the treatment-specific survival functions. For example, a natural estimator of the treatment-specific restricted mean survival time $r_{0,a} := \int_0^\tau \theta_0(t, a) dt$ is given by $r_{n,a} := \int_0^\tau \theta_n^\circ(t, a) dt$. Uniform consistency of $\theta_n^\circ(\cdot, a)$ on $[0, \tau]$, as implied by Theorem 2, implies consistency of $r_{n,a}$. In view of an application of the functional delta method, the weak convergence of $\{n^{1/2}[\theta_n^\circ(t, a) - \theta_0(t, a)] : t \in [0, \tau]\}$, as implied by Theorem 3, implies that $n^{1/2}(r_{n,a} - r_{0,a})$ is asymptotically linear with influence function $o \mapsto \int_0^\tau \phi_{0,t,a}^*(o) dt$. Inference for contrasts of treatment-specific restricted mean survival times can be obtained analogously.

6 Data-adaptive estimation of nuisance functions

As discussed above, our proposed estimator requires estimation of three nuisance parameters: the conditional survival functions S_0 and G_0 of the event time and censoring distributions, respectively, given exposure and covariates, and the propensity π_0 of exposure given covariates. We note that π_0 can be estimated using any regression estimator for a binary outcome. We recommend leveraging multiple parametric, semiparametric and nonparametric regression strategies using the SuperLearner algorithm (Breiman, 1996; van der Laan et al., 2007).

There are several existing strategies for estimating S_0 and G_0 . The most widely-used regression model for survival outcomes is the Cox proportional hazard model (Cox, 1972), which can be used in conjunction with the Breslow estimator (Breslow, 1972) or parametric estimators of the baseline cumulative hazard function to obtain estimates of S_0 and G_0 . The

accelerated failure time model can also be used as a semiparametric estimator of S_0 and G_0 (Wei, 1992). Alternatively, various other semiparametric and nonparametric regression techniques for survival data have been proposed, including, to name a few, additive Cox models (Hastie and Tibshirani, 1990), piecewise constant hazard models (Friedman, 1982), survival random forests (Ishwaran et al., 2008), gradient boosting (Hothorn et al., 2005), and deep neural networks (Rava and Bradic, 2020). In practice, it may not be *a priori* clear to the researcher which of these or other algorithms are most appropriate in a given setting. An ensemble algorithm would allow the researcher to select from or combine multiple candidate estimators in a data-adaptive manner. However, ensemble algorithms for regression do not immediately extend to our setting due to the right-censored data structure.

Here, we propose an iterative SuperLearner ensemble algorithm for combining multiple candidate nuisance estimators of S_0 and G_0 . We are aware of several existing approaches to ensemble learning with right-censored data. van der Laan and Dudoit (2003), Keles et al. (2004), and Polley and van der Laan (2011) proposed ensemble learners for a conditional survival function at a fixed point t , for regression and conditional quantile functions, and for conditional density and hazard functions, assuming they exist. While it builds on these previous works, our procedure accomplishes several goals that, to the best of our knowledge, these previous works did not. First, we target the entire survival function on an interval rather than a summary such as the survival at a single point t , the mean, or a quantile. Second, unlike methods that target the conditional density or hazard functions, we do not require that the event occurs on either a fully discrete or fully continuous scale, but rather allow both of these possibilities as well as mixed distributions. Third, we target both S_0 and G_0 together rather than one or the other by iterating between optimization of S_n^* and G_n^* , which has the potential to improve estimation of both.

We recall that if identification conditions presented in Supplementary Material hold for some $a \in \{0, 1\}$ and $\tau \in (0, \infty)$, then $S_0(t | a, w) = P_{0,F}(T(a) > t | W = w)$ and $G_0(t | a, w) = P_{0,F}(C(a) \geq t | W = w)$ for any $t \in [0, \tau]$. Central to our ensemble method

are representations of S_0 and G_0 as minimizers of oracle risk functions, as stated in the next result. For this result, we define \mathcal{C}_τ as the set of functions from $[0, \tau] \times \{0, 1\} \times \mathcal{W}$ to $[0, 1]$.

Theorem 4. *Let S^* be a minimizer of $S \mapsto P_0 L_{S, G_0}$ over $S \in \mathcal{C}_\tau$ and G^* be a minimizer of $G \mapsto P_0 M_{G, S_0}$ over $G \in \mathcal{C}_\tau$, where we define the loss functions*

$$L_{S, G} : (w, a, y, \delta) \mapsto \int_0^\tau S(t | a, w) \left[S(t | a, w) - 2 \left\{ 1 - \frac{\delta I(y \leq t)}{G(y | a, w)} \right\} \right] dt ;$$

$$M_{G, S} : (w, a, y, \delta) \mapsto \int_0^\tau G(t | a, w) \left[G(t | a, w) - 2 \left\{ 1 - \frac{(1 - \delta)I(y < t)}{S(y | a, w)} \right\} \right] dt .$$

If conditions (A1)–(A5) in Supplementary Material hold for each $a \in \{0, 1\}$, then $S^(t | a, w) = S_0(t | a, w)$ for P_0 -almost every (a, w) and all $t \leq \tau$, and $G^*(t | a, w) = G_0(t | a, w)$ for P_0 -almost every (a, w) and all $t \leq \tau$ such that $S_0(t | a, w) > 0$.*

We now provide some intuition for the loss function $M_{G, S}$; analogous intuition applies to $L_{S, G}$. Up to a constant not depending on G , the loss function $M_{G, S}$ is equal to $\int_0^\tau [G(t | a, w) - f_S(t, o)]^2 dt$, where $f_S(t, o) := 1 - (1 - \delta)I(y < t)/S(y | a, w)$. The term involving f_S^2 is constant with respect to G and is therefore omitted from $M_{S, G}$ because f_S is only square-integrable under slightly stronger conditions. Under the conditions of Theorem 4, $E_0[f_{S_0}(t, O) | A = a, W = w] = G_0(t | a, w)$, so that for each fixed (t, a, w) , ignoring integrability issues, $G_0(t | a, w)$ minimizes $\gamma \mapsto E_0[\{\gamma - f_{S_0}(t, O)\}^2 | A = a, W = w]$. The result essentially follows by taking the expectation over A and W and integrating over t .

Were G_0 known, an optimal weighted combination of p candidate estimators $S_n^{(1)}, S_n^{(2)}, \dots, S_n^{(p)}$ of S_0 could be found by minimizing the cross-validated empirical risk $\mathbb{P}_n L_{S, G_0}$ over S in the set Π_S of convex combinations $\sum_{j=1}^p \alpha_j S_n^{(j)}$ for α in the p -dimensional simplex. Here, by *cross-validated* we mean that the sample is split into K folds, candidate estimators are each trained holding out each fold, evaluated on the held-out fold, and these held-out evaluations are used to compute the empirical mean $\mathbb{P}_n L_{S, G}$ (see, e.g., van der Laan et al., 2007 or van der Laan and Rose, 2011 for additional details). Were S_0 known, an analogous procedure could be used to find an optimal weighted combination of q candidate estimators

$G_n^{(1)}, G_n^{(2)}, \dots, G_n^{(q)}$ of G_0 in the set Π_G of convex combinations $\sum_{j=1}^q \alpha_j G_n^{(j)}$ for α in the q -dimensional simplex. Since S_0 and G_0 are not known in practice, we propose the following iterative strategy:

Step 0: Obtain an initial estimator $G_{n,0}^*$ of G_0 using a nonparametric procedure.

Step k: Compute $S_{n,k}^* := \operatorname{argmin}_{S \in \Pi_S} \mathbb{P}_n L_{S,G_{n,k-1}^*}$ and $G_{n,k}^* := \operatorname{argmin}_{G \in \Pi_G} \mathbb{P}_n M_{G,S_{n,k}^*}$.

The procedure can be terminated, for example, when $\|S_{n,k}^* - S_{n,k-1}^*\|_\infty$ and $\|G_{n,k}^* - G_{n,k-1}^*\|_\infty$ both fall below some pre-specified threshold. In practice, we can evaluate the integrals in $L_{S,G}$ and $M_{G,S}$ using a Riemann sum over a fine grid.

Obtaining the cross-validated estimates $S_n^{(1)}, S_n^{(2)}, \dots, S_n^{(p)}$ and $G_n^{(1)}, G_n^{(2)}, \dots, G_n^{(q)}$, a requirement for any ensemble learner, is the most computationally expensive step of the above procedure. In our proposed procedure, these estimates only need to be obtained once. The only computational burden of the algorithm beyond that of an ordinary SuperLearner is the possibly multiple optimization steps to find the optimal convex combinations of the candidate learners, which is typically much less computationally expensive than obtaining the cross-validated estimates of the candidate learners. Therefore, the algorithm outlined above is not substantially more computationally expensive than an ordinary SuperLearner.

7 Numerical studies

We conducted a numerical study to evaluate the finite-sample performance of our methods. For brevity, we summarize the design of this study; full details can be found in Supplementary Material. We simulated a vector $W := (W_1, W_1, W_3)$ of three continuous covariates. We then set $\operatorname{logit} P_0(A = 1 | W = w) = -1 + \log \{1 + \exp(-20 + \frac{w_1}{10}) + \exp(-3 + \frac{w_3}{2})\}$. We considered two processes for generating the event and censoring times. First, we simulated both T and C from proportional hazards models with main terms as well as interactions between A each component of W and between W_1 and W_3 . Second, we simulated T and

C from non-proportional hazards models. For both simulation settings, we truncated C at $\tau = 24$, and parameters were chosen to yield an average censoring rate of $E_0[P_0(C \leq 12 | A = 0, W)] = 0.2$, an average observed event rate of $E_0[P_0(T \leq C | A = 0, W)] = 0.15$, and a counterfactual risk ratio of 0.7 at $t = 12$.

We simulated 1000 datasets using the above processes for $n = 250, 500, \dots, 1500$. For each dataset, we estimated the propensity score using SuperLearner with a library detailed in Supplementary Material. We estimated the conditional survival curves in four ways: a correctly-specified Cox proportional hazards, an incorrectly-specified proportional hazards model with main terms only, survival random forest (Ishwaran et al., 2008), and the iterative SuperLearner described in Section 6. For the iterative SuperLearner, we used a combination of parametric survival models, semiparametric proportional hazard models, generalized additive Cox models, and survival random forest. In addition, we used five-fold cross-validation, survival random forest as the initial estimator, and limited the recursive procedure to fifteen iterations. For each candidate estimator S_n of S_0 , we computed the t -specific average root mean squared error (RMSE) of the estimator as $\{\frac{1}{n} \sum_{i=1}^n [S_n(t | A_i, W_i) - S_0(t | A_i, W_i)]^2\}^{1/2}$ for each $t \in \{0.5, 1, \dots, 12\}$. We did the same for estimators of G_0 . We estimated the counterfactual survival curves by (1) G-computation of each candidate estimator of the conditional survival function of the event using the empirical distribution of the covariates, and (2) using our cross-fitted one-step estimator with various combinations of the conditional survival estimators. For each method, we recorded the estimated control and treatment survival probabilities and the risk ratio at time $t = 12$. For our method, we also computed pointwise confidence intervals at $t = 12$ and uniform confidence bands over $t \in [0, 12]$.

In the first simulation setting, correctly-specified proportional hazards models yield conditional survival estimators converging at the parametric rate $n^{-1/2}$, so condition (B5) is satisfied in this case. The library of our proposed SuperLearner includes a correctly-specified proportional hazards estimator, so if our method is able to correctly select this estimator from the candidate library, it should also achieve the necessary rate of convergence. In the

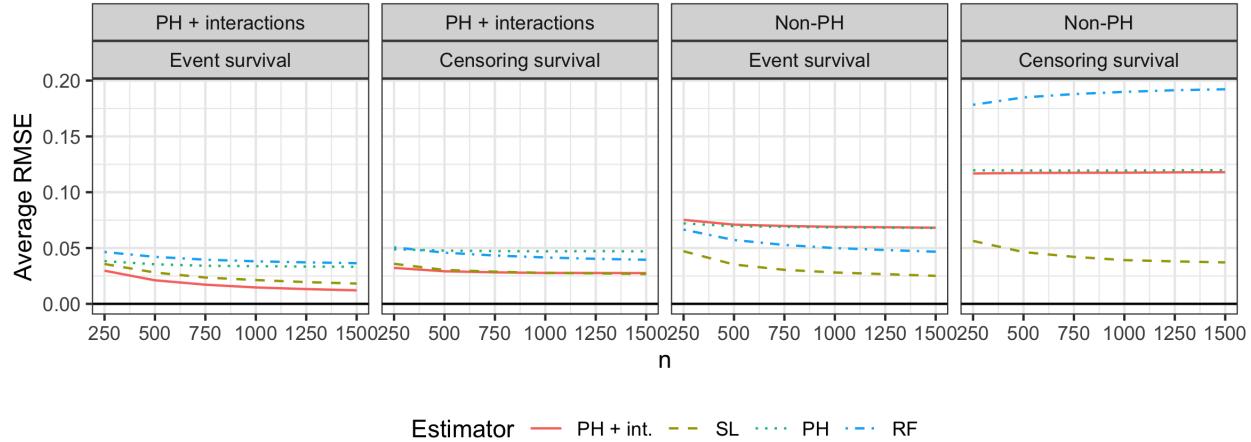


Figure 1: Average root mean squared error (RMSE), as defined in the text, of the four conditional survival estimators as a function of sample size n . The four panels correspond to the two simulation settings and to which conditional survival is being considered. In the legend, “PH” refers to Cox proportional hazards model, “int.” refers to interactions, “RF” refers to survival random forest, and “SL” refers to our iterative SuperLearner.

second setting, the proportional hazards estimators are both inconsistent. The censoring time is generated from a generalized additive proportional hazards model, and the library of our proposed SuperLearner includes a generalized additive model estimator, so if our estimator is able to adapt, then it should obtain the necessary rate of convergence. However, the event time is generated from a complicated non-proportional hazards mechanism that is not included in common semiparametric survival models. Hence, we did not include a correctly-specified semiparametric estimator in our SuperLearner library, so it is unclear what rate of convergence our estimator will attain. While Cui et al. (2022) recently developed rates of convergence of survival random forest, it appears that these results do not apply directly to our setting because they concern a bias-corrected estimator rather than that of Ishwaran et al. (2008). Hence, it is unclear what rate of convergence survival random forest will attain.

We now turn to the results of the numerical study. Figure 1 displays the average RMSE over the 1000 simulations and over $t \in \{0.5, 1, \dots, 12\}$ for the four estimators of the conditional survival of event and censoring and for the two simulation settings. A figure comparing the estimators for each t is provided in Supplementary Material, and does not change the

following observations. For the proportional hazards setting (left two panels), the correctly-specified Cox proportional hazards estimator with interactions had the best RMSE of the four estimators. The SuperLearner proposed here had slightly higher RMSE for the event survival and very similar RMSE for the censoring survival, indicating that it did a good job selecting the correctly-specified estimator from the candidate library. The RMSE of the incorrectly-specified Cox estimator and the survival random forest (RF) were larger, but that of RF decreased with n . For the non-proportional hazards setting (right two panels), the RMSE of our SuperLearner estimator was the best of the four estimators for both the event and censoring survivals. The RMSE of both Cox model estimators did not decrease substantially with n because neither estimator was correctly specified. The RMSE of RF was slightly higher than that of the SuperLearner for the survival of the event, but was much worse for the censoring survival. We conclude that our iterative SuperLearner was able to adapt to proportional hazards in the setting where that was the correct model, and was also able to leverage other candidate estimators to outperform the proportional hazards estimators in the setting where the data were not generated from a proportional hazards model.

Figure 2 displays the properties of the estimators and confidence intervals for the control survival curve. Analogous plots for the treatment survival curve and risk ratio are presented in Supplementary Material. For ease of viewing, only a subset of the estimators considered are included in Figure 2; the remaining estimators are also shown in Supplementary Material. We first discuss the results for the case where the data were generated from a proportional hazards model with interactions (left column). The biases of the proposed method (first row) using the Cox model with interactions and the G-computed Cox estimator with interactions were within Monte Carlo error of zero for all sample sizes. The bias of the proposed method using the iterative SuperLearner for nuisance estimators was less than 0.5% for all sample sizes. The biases of the G-computed Cox estimator without interactions and the G-computed survival random forest were over 1%, and the former relatively constant as a function of n ,

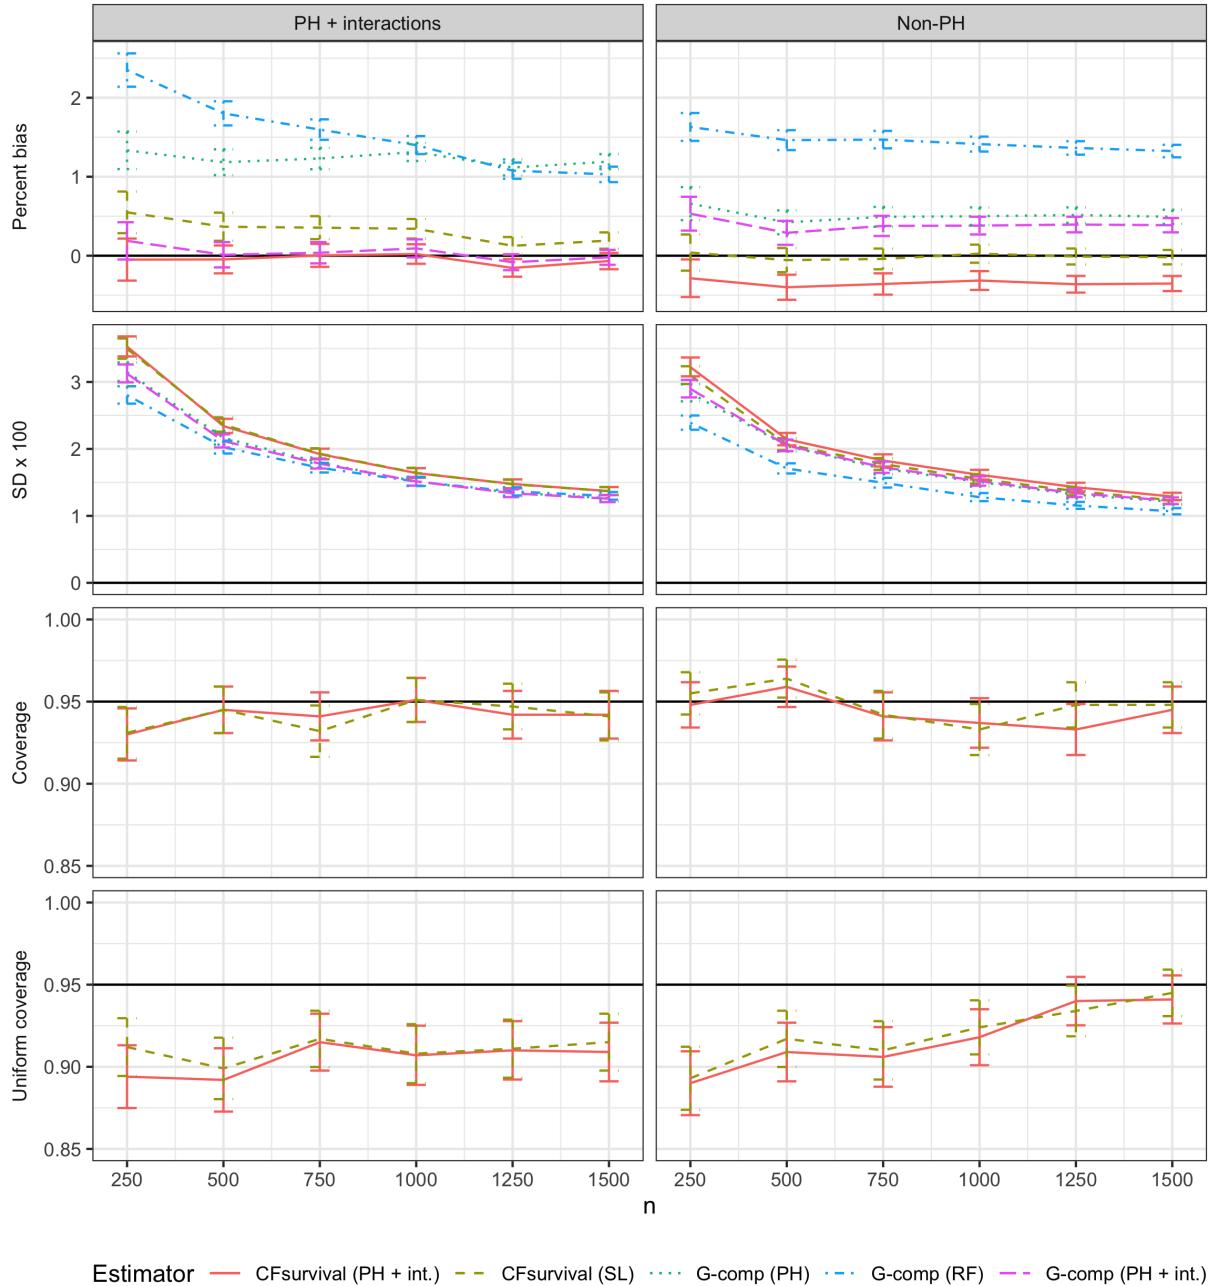


Figure 2: Properties of five of the estimators of the counterfactual control survival as a function of sample size. Columns correspond to the two simulation settings. From top to bottom, the rows contain: percent bias, standard deviation, pointwise coverage, and uniform coverage. The first three rows correspond to inference at time $t = 12$. “CFsurvival” is the method developed here, and “G-comp” is G-computation. Paratheticals indicate the estimator used for the conditional survival(s), with shorthand defined in the Figure 1 caption. Vertical bars represent 95% confidence intervals taking into account uncertainty due to conducting a finite number of simulations.

suggesting that the method is inconsistent. This was expected because the true conditional survival curves include interactions that the estimator omitted. Estimators based on the one-step method had larger standard deviation (second row) than those based on G-computation. As a result, G-computation of a correctly specified Cox model estimator yielded the smallest mean squared error. However, this relies heavily on correct specification of the Cox model. The pointwise coverage (third row) of confidence intervals constructed using our method was within Monte Carlo error of the nominal 95% for all sample sizes 500 and larger. The uniform coverage of our method (bottom row) was between 90 and 93%, indicating slight undercoverage even at large sample sizes. We believe this is due to poor coverage of small values of t , and reflects the challenge of constructing equi-precision confidence bands when the standard deviation is small at the boundary. The uniform coverage is at or slightly above the nominal level for the risk ratio, as shown in Supplementary Material.

We now discuss the results for the case where the data were generated from a non-proportional hazards model (right column of Figure 2). In this case, the bias (first row) of the proposed method using the iterative SuperLearner for nuisance estimators was within Monte Carlo error of zero for all sample sizes. This is somewhat surprising because the SuperLearner library does not contain an estimator of the conditional survival of the event that is correctly specified. The G-computed Cox estimators and the proposed method using the Cox estimator with interactions for nuisance estimators were biased because these nuisances were inconsistent in this setting (see Figure 1). The G-computed random forest had the largest bias, though it did decrease with sample size. The G-computed random forest had the smallest standard deviation of the estimators considered (second row), which resulted in comparable mean squared errors of the estimators. The pointwise coverage of our estimator (third row) was very good for all sample sizes considered. It is surprising that our estimator had good coverage when using the Cox estimator for conditional survivals, since these estimators were inconsistent, and we do not generally expect the coverage to be good when the nuisances are inconsistent. We expect the coverage would worsen at larger sample

sizes. The uniform coverage of our estimator (bottom row) was slightly below nominal for small sample sizes, but increased to the nominal level as the sample size increased.

The Supplementary Material contains additional results from the simulation study, including an analogue of Figure 2 for the treatment survival curve and risk ratio, a figure illustrating the double-robust properties of our estimator, and a figure illustrating the effect of cross-fitting. Cross-fitting reduces bias and improves coverage of confidence intervals. The Supplementary Material also contains a second simulation study comparing the effect of grid choice in the computation of the integral in our estimator and in a method designed for discrete-time survival data. This study demonstrates that while both methods suffer from bias when the grid is too coarse, increasing the size of the grid has little impact on the run time of our methods, but it has an enormous impact on the run time of methods for discrete-time data.

8 Effect of elective neck dissection on mortality

In this section, we use the methods developed in this article to assess the effect of elective neck dissection (END) on survival among patients with clinically node-negative, high-grade parotid carcinoma. END consists of surgical removal of lymph nodes to prevent metastatic spread via the lymphatic system, and has been the subject of controversy among surgeons and oncologists. On one hand, lymph node metastases are common among patients with high-grade oral carcinomas, and END is an effective treatment for preventing these metastases. On the other hand, END is more invasive and leads to higher morbidity than radiation therapy, which can also be used to treat and prevent metastases. We refer the reader to Jalisi (2005) and Kowalski and Sanabria (2007) for a more detailed discussion of END.

We analyzed a retrospective cohort consisting of $n = 1547$ patients in the National Cancer Database who were diagnosed with clinically node-negative, high-grade parotid cancer between January 1, 2004 and December 31, 2013, and followed until the latter date. The

exposure level $A = 1$ here corresponded to receipt of END at diagnosis, and the outcome of interest was all-cause mortality up to five years post-diagnosis. Mortality was subject to right-censoring because patients could be lost to follow-up or still alive on December 31, 2013. The baseline covariate vector W consisted of patient age, sex, race, tumor stage, histology, comorbidity, and payor, as well as the average income, education, county of residence, and treatment facility type. Additional details of the cohort construction and demographics may be found in Harbison et al. (2020).

An unadjusted analysis yielded stratified Kaplan-Meier survival estimates of 56.4% (95% CI: 52.8–60.3) for patients receiving END and 48.6% (43.4–54.5) for those not receiving end at $t = 5$ years post-diagnosis. The survival curves were deemed to be significantly different using a log-rank test ($p < 0.0001$). These results suggest that END has a significant positive association with survival. However, since the data are observational, these results cannot be interpreted causally. By using the methods proposed here, we can adjust for baseline confounding flexibly while still reporting survival curves and contrasts thereof, which provide a simple interpretation that is familiar for many clinicians and scientists.

We used the methods presented here to estimate the treatment-specific G-computed survival functions $\theta_0(t, 0)$ and $\theta_0(t, 1)$. If the untestable causal conditions (A1)–(A5) presented in Supplementary Material hold, then these curves correspond to the counterfactual survival functions under assignment of all patients in the target population to no END and END, respectively. In particular, (A1)–(A5) require that the covariate vector W be sufficient to control for confounding between receipt of END and mortality, and that A and W together be sufficient to control for the dependence between mortality and censoring. We also estimated the survival difference, survival ratio, and risk ratio functions.

We estimated the treatment propensity using SuperLearner (van der Laan et al., 2007) with a library consisting of generalized linear models, generalized additive models, multivariate adaptive regression splines, random forests, and extreme gradient boosting. We estimated the conditional survival and censoring functions using the novel SuperLearner de-

fined in Section 6 with a library consisting of the treatment group-specific Kaplan-Meier estimators, parametric survival models, Cox proportional hazard models, generalized additive models, and piecewise constant hazard models. Additional details on the libraries used for nuisance estimation and the estimated SuperLearner coefficients may be found in the Supplementary Material.

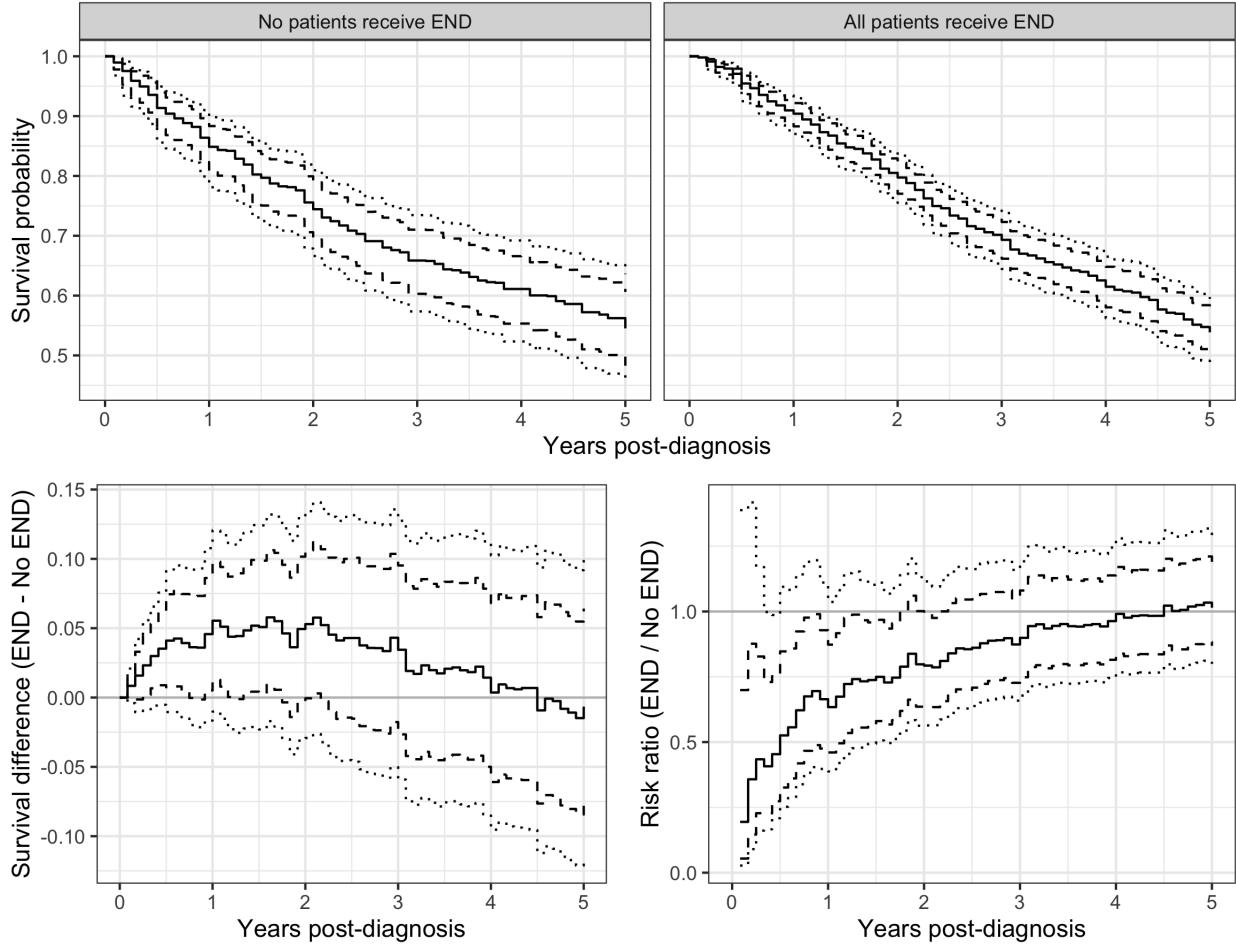


Figure 3: Results of the analysis of the effect of elective neck dissection (END) on all-cause mortality. The top row shows the estimated treatment-specific survival curves were all (right panel) and no (left panel) patients to receive END. The bottom row shows the estimated survival difference (left panel) and risk ratio (right panel) functions. In all figures, pointwise 95% confidence intervals are shown as dashed lines, and uniform 95% confidence bands are shown as dotted lines.

The same scientific question addressed here was studied in Harbison et al. (2020) using a preliminary version of the methods developed here. However, the estimator used for the

analysis presented in Harbison et al. (2020) did not use cross-fitting, and only used random forests to estimate the conditional survival and censoring functions. In addition, in Harbison et al. (2020), uniform confidence bands or contrasts of the survival functions, which are both important for comparing the survival functions uniformly in time, were not provided.

Figure 3 displays the results of the analysis. The top row displays the estimated counterfactual survival functions corresponding to receiving END (left) versus not receiving END (right) along with pointwise and uniform confidence regions. We estimate that 53.9% (95% CI: 50.1–57.5) of patients would be alive five years post-diagnosis if undergoing END, while 54.5% (48.3–60.6) would be alive if not undergoing END. The bottom row displays the estimated survival difference and risk ratio functions. The estimated survival difference was positive, and the estimated risk ratio was less than 1 between 0 and 4 years post-diagnosis, suggesting that END possibly improves short-term survival. However, both confidence bands included the null effect throughout this time period, and the p-value of the test of the null hypothesis that $\theta_0(t, 0) = \theta_0(t, 1)$ for all $t \in [0, 5]$ is 0.12. Thus, we cannot reject the null hypothesis that END does not impact overall survival through five years. The estimated survival ratio function was very similar in form to the estimated survival difference function. We estimate the restricted mean survival time through five years to be 3.62 years (95% CI: 3.42–3.81) under no END and 3.76 years (95% CI: 3.65–3.87) under END, with an estimated difference of 0.14 years (95%CI: -0.07–0.36). Therefore, after adjusting for baseline confounding, the data no longer provide evidence of an effect of END on survival.

9 Concluding remarks

In this article, we proposed a doubly-robust estimator of the treatment-specific survival curve in the presence of baseline confounders that permits the use of data-adaptive estimators of nuisance functions. In addition, we proposed an ensemble learner for combining multiple candidate estimators of the conditional event and censoring survival functions. We provided

general sufficient conditions for consistency and asymptotic linearity, both pointwise and uniformly, of the proposed estimator, and used these results to construct confidence intervals, confidence bands and tests. The proposed methods permit event and censoring distributions that may be continuous, discrete, or mixed continuous-discrete. This is important because in many applications the event and censoring distributions may have either a continuous or mixed continuous-discrete support, whereas most existing methods for counterfactual survival estimation are tailored either to the fully continuous or fully discrete setting.

The methods discussed here can also be used for analyzing data from randomized trials with time-to-event outcomes. In such settings, in view of randomization, the treatment-outcome and treatment-censoring relationships are unconfounded, but the outcome-censoring relationship may still be confounded. Adjusting for baseline covariates can reduce bias due to such dependent censoring, and our methods provide a way to do so without assuming any particular form for the conditional survival and censoring functions.

When the exposure varies over time rather than being fixed, it is typically necessary to adjust for time-varying confounders in order to recover causal parameters, since the change in exposure status may be related to changes in patient characteristics that are also related to the outcome. Even when the exposure does not vary over time, there may be time-varying common causes of the event and censoring times. We are unaware of an extension of the identification result we used to the setting with time-varying confounders. In the context of discrete-time longitudinal data, the nested G-formula provides an identification of the counterfactual survival probabilities (Robins, 1986). It is unclear how or whether the methods proposed here would extend to estimation of treatment-specific survival curves in continuous time with time-varying confounders. This is a topic of ongoing research.

The iterative SuperLearner proposed in Section 6 performed well in numerical studies, but its utility could be improved with future research. First, demonstrating convergence of the iterative algorithm would provide an important computational guarantee for the procedure. Second, proving an oracle inequality for the estimator using, e.g., the results of Dudoit and

van der Laan (2005) would provide an important theoretical guarantee for the estimator. These are important areas of future research.

Acknowledgements

The authors gratefully acknowledge support from the University of Massachusetts Department of Mathematics and Statistics startup fund (TW), NSF Award 2113171 (TW), NIH grant DP2-LM013340 (AL), and NHLBI grant HL137808 (MC).

Conflicts of Interest

The authors report there are no competing interests to declare.

References

Anderson, J. R., Bernstein, L., and Pike, M. C. (1982). Approximate Confidence Intervals for Probabilities of Survival and Quantiles in Life-Table Analysis. *Biometrics*, 38(2):407–416.

Bai, X., Tsiatis, A. A., and O’Brien, S. M. (2013). Doubly-robust estimators of treatment-specific survival distributions in observational studies with stratified sampling. *Biometrics*, 69(4):830–839.

Beran, R. (1981). Nonparametric regression with randomly censored survival data. Technical Report.

Bickel, P. J. (1982). On adaptive estimation. *The Annals of Statistics*, 10(3):647–671.

Breiman, L. (1996). Stacked regressions. *Machine Learning*, 24(1):49–64.

Breslow, N. E. (1972). Discussion of the paper by D. R. Cox. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2):216–217.

Cole, S. R. and Hernán, M. A. (2004). Adjusted survival curves with inverse probability weights. *Computer Methods and Programs in Biomedicine*, 75(1):45–49.

Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2):187–202.

Cui, Y., Zhu, R., Zhou, M., and Kosorok, M. (2022). Consistency of survival tree and forest models: splitting bias and correction. *Statistica Sinica*, 32:1245–1267.

Dabrowska, D. M. (1989). Uniform consistency of the kernel conditional Kaplan-Meier estimate. *The Annals of Statistics*, 17(3):1157–1167.

Díaz, I. (2019a). Machine learning in the estimation of causal effects: targeted minimum loss-based estimation and double/debiased machine learning. *Biostatistics*, 21(2):353–358.

Díaz, I. (2019b). Statistical inference for data-adaptive doubly robust estimators with survival outcomes. *Statistics in Medicine*, 38(15):2735–2748.

Dudoit, S. and van der Laan, M. J. (2005). Asymptotics of cross-validated risk estimation in estimator selection and performance assessment. *Statistical Methodology*, 2(2):131–154.

Dukes, O., Martinussen, T., Tchetgen Tchetgen, E. J., and Vansteelandt, S. (2019). On doubly robust estimation of the hazard difference. *Biometrics*, 75(1):100–109.

Ferreira Guerra, S., Schnitzer, M. E., Forget, A., and Blais, L. (2020). Impact of discretization of the timeline for longitudinal causal inference methods. *Statistics in Medicine*, 39(27):4069–4085.

Friedman, M. (1982). Piecewise Exponential Models for Survival Data with Covariates. *The Annals of Statistics*, 10(1):101–113.

Gill, R. D. and Johansen, S. (1990). A survey of product-integration with a view toward application in survival analysis. *The Annals of Statistics*, 18(4):1501–1555.

Gill, R. D. and Robins, J. M. (2001). Causal inference for complex longitudinal data: The continuous case. *The Annals of Statistics*, 29(6):1785–1811.

Harbison, A. R., Gray, A. J., et al. (2020). The Role of Elective Neck Dissection in High-Grade Parotid Malignancy: A Hospital-Based Cohort Study. *The Laryngoscope*, 130(6):1487–1495.

Harrington, D. P. and Fleming, T. R. (1982). A Class of Rank Test Procedures for Censored Survival Data. *Biometrika*, 69(3):553–566.

Hastie, T. J. and Tibshirani, R. J. (1990). *Generalized Additive Models*, volume 43. CRC press.

Hernán, M. A. (2010). The Hazards of Hazard Ratios. *Epidemiology*, 21(1):13–15.

Hothorn, T., Bühlmann, P., Dudoit, S., Molinaro, A., and van der Laan, M. J. (2005). Survival ensembles. *Biostatistics*, 7(3):355–373.

Hou, J., Bradic, J., and Xu, R. (2023). Treatment Effect Estimation Under Additive Hazards Models With High-Dimensional Confounding. *Journal of the American Statistical Association*, 118(541):327–342.

Hubbard, A. E., van der Laan, M. J., and Robins, J. M. (2000). Nonparametric locally efficient estimation of the treatment specific survival distribution with right censored data and covariates in observational studies. *IMA Volumes in Mathematics and Its Applications*, 116:135–178.

Ishwaran, H., Kogalur, U. B., Blackstone, E. H., and Lauer, M. S. (2008). Random survival forests. *Annals of Applied Statistics*, 2(3):841–860.

Jalisi, S. (2005). Management of the Clinically Negative Neck in Early Squamous Cell Carcinoma of the Oral Cavity. *Otolaryngologic Clinics of North America*, 38(1):37–46.

Kaplan, E. L. and Meier, P. (1958). Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53(282):457–481.

Keles, S., van der Laan, M., and Dudoit, S. (2004). Asymptotically optimal model selection method with right censored outcomes. *Bernoulli*, 10(6):1011–1037.

Kowalski, L. P. and Sanabria, A. (2007). Elective neck dissection in oral carcinoma: a critical review of the evidence. *Acta Otorhinolaryngologica Italica*, 27(3):113.

Luedtke, A., Sofrygin, O., van der Laan, M., and Carone, M. (2018). Sequential Double Robustness in Right-Censored Longitudinal Models. *ArXiv e-prints*, 1705.02459.

Makuch, R. W. (1982). Adjusted survival curve estimation using covariates. *Journal of Chronic Diseases*, 35(6):437–443.

Mansourvar, Z. and Martinussen, T. (2017). Estimation of average causal effect using the restricted mean residual lifetime as effect measure. *Lifetime Data Analysis*, 23:426–438.

Mansourvar, Z., Martinussen, T., and Scheike, T. H. (2016). An Additive–Multiplicative Restricted Mean Residual Life Model. *Scandinavian Journal of Statistics*, 43(2):487–504.

Martinussen, T. (2022). Causality and the Cox Regression Model. *Annual Review of Statistics and Its Application*, 9(1):249–259.

Molina, J., Rotnitzky, A., Sued, M., and Robins, J. (2017). Multiple robustness in factorized likelihood models. *Biometrika*, 104(3):561–581.

Neyman, J. (1923). On the application of probability theory to agricultural experiments. Essay on principles (with discussion). *Statistical Science*, 5:465–480.

Polley, E. C. and van der Laan, M. J. (2011). Super Learning for Right-Censored Data. In van der Laan, M. and Rose, S., editors, *Targeted Learning: Causal Inference for Observational and Experimental Data*, chapter 16, pages 249–258. Springer-Verlag New York, New York.

Rava, D. and Bradic, J. (2020). DeepHazard: neural network for time-varying risks. *arXiv e-prints*, 2007.13218.

Robins, J. (1986). A new approach to causal inference in mortality studies with a sustained exposure period – application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7(9):1393–1512.

Robins, J., Li, L., Tchetgen, E., van der Vaart, A., et al. (2008). Higher order influence functions and minimax estimation of nonlinear functionals. In Nolan, D. and Speed, T., editors, *Probability and Statistics: Essays in Honor of David A. Freedman*, pages 335–421. Institute of Mathematical Statistics.

Rotnitzky, A., Lei, Q., Sued, M., and Robins, J. M. (2012). Improved double-robust estimation in missing data and causal inference models. *Biometrika*, 99(2):439–456.

Rotnitzky, A., Robins, J., and Babino, L. (2017). On the multiply robust estimation of the mean of the g-functional. *ArXiv e-prints*, 1705.08582.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701.

Seaman, S. R., Keogh, R. H., Dukes, O., and Vansteelandt, S. (2021). Using generalized linear models to implement g-estimation for survival data with time-varying confounding. *Statistics in Medicine*, 40(16):3779–3790.

Strandberg, E., Lin, X., and Xu, R. (2014). Estimation of main effect when covariates have non-proportional hazards. *Communications in Statistics - Simulation and Computation*, 43(7):1760–1770.

Tchetgen Tchetgen, E. J. (2009). A commentary on G. Molenberghs's review of missing data methods. *Drug Information Journal*, 43(4):433–435.

van der Laan, M. and Rose, S. (2011). *Targeted learning: causal inference for observational and experimental data*. Springer-Verlag New York.

van der Laan, M. J. and Dudoit, S. (2003). Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples. Working Paper Series Working Paper 130, U.C. Berkeley Division of Biostatistics.

van der Laan, M. J., Polley, E. C., and Hubbard, A. E. (2007). Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1).

Vansteelandt, S., Dukes, O., Lancker, K. V., and Martinussen, T. (2022). Assumption-Less Cox Regression. *Journal of the American Statistical Association*. Forthcoming.

Wei, L. J. (1992). The accelerated failure time model: A useful alternative to the cox regression model in survival analysis. *Statistics in Medicine*, 11(14-15):1871–1879.

Westling, T., van der Laan, M. J., and Carone, M. (2020). Correcting an estimator of a multivariate monotone function with isotonic regression. *Electronic Journal of Statistics*, 14(2):3032–3069.

Wu, L. and Gilbert, P. B. (2002). Flexible weighted log-rank tests optimal for detecting early and/or late survival differences. *Biometrics*, 58(4):997–1004.

Zeng, D. (2004). Estimating marginal survival function by adjusting for dependent censoring using many covariates. *Annals of Statistics*, 32(4):1533–1555.

Zhang, M. and Schaubel, D. E. (2012). Contrasting treatment-specific survival using double-robust estimators. *Statistics in Medicine*, 31(30):4255–4268.

Zheng, W. and van der Laan, M. J. (2011). Cross-Validated Targeted Minimum-Loss Based Estimation. In van der Laan, M. and Rose, S., editors, *Targeted Learning: Causal Inference for Observational and Experimental Data*, chapter 27, pages 459–473. Springer-Verlag New York, New York.