



Published in final edited form as:

Stat Med. 2016 August 30; 35(19): 3319–3332. doi:10.1002/sim.6929.

Estimating restricted mean treatment effects with stacked survival models

Andrew Wey^{a,b}, David M. Vock^c, John Connett^c, and Kyle Rudser^c

^aMinneapolis Medical Research Foundation, Minneapolis, Minnesota

^bBiostatistics and Data Management Core, John A. Burns School of Medicine, University of Hawaii, Honolulu, Hawaii

^cDivision of Biostatistics, School of Public Health, University of Minnesota, Minneapolis, Minnesota

Abstract

The difference in restricted mean survival times between two groups is a clinically relevant summary measure. With observational data, there may be imbalances in confounding variables between the two groups. One approach to account for such imbalances is estimating a covariate-adjusted restricted mean difference by modeling the covariate-adjusted survival distribution, and then marginalizing over the covariate distribution. Since the estimator for the restricted mean difference is defined by the estimator for the covariate-adjusted survival distribution, it is natural to expect that a better estimator of the covariate-adjusted survival distribution is associated with a better estimator of the restricted mean difference. We therefore propose estimating restricted mean differences with stacked survival models. Stacked survival models estimate a weighted average of several survival models by minimizing predicted error. By including a range of parametric, semi-parametric, and non-parametric models, stacked survival models can robustly estimate a covariate-adjusted survival distribution and, therefore, the restricted mean treatment effect in a wide range of scenarios. We demonstrate through a simulation study that better performance of the covariate-adjusted survival distribution often leads to better mean-squared error of the restricted mean difference although there are notable exceptions. In addition, we demonstrate that the proposed estimator can perform nearly as well as Cox regression when the proportional hazards assumption is satisfied and significantly better when proportional hazards is violated. Finally, the proposed estimator is illustrated with data from the United Network for Organ Sharing to evaluate post-lung transplant survival between large and small-volume centers.

Keywords

Bias-variance tradeoff; proportional hazards assumption; restricted mean difference; stacked survival models; survival analysis

*Correspondence to: Andrew Wey, Minneapolis Medical Research Foundation, Minneapolis, Minnesota. awey@cdrg.org.

1. Introduction

Patients with end-stage lung disease may be eligible for lung transplantation after other treatment options fail. Unfortunately, post-lung transplant survival is poor, especially in comparison to other solid organ transplants, with one and three-year graft survival of 79% and 64%, respectively. Given the poor prognosis, understanding the factors related to post-transplant survival remains an important but controversial task. For example, previous research has suggested that transplant center volume, which we define as the number of lung transplants performed at a center over the preceding two years, is associated with post-transplant survival [1, 2]. However, previous research has relied on estimators based on the proportional hazards assumption, which is potentially violated in lung transplantation [2]. We therefore develop a causal estimator of the restricted mean treatment effect between low and high volume centers that extends beyond proportional hazard scenarios.

In the absence of censoring, the difference in post-transplant survival time between high and low volume centers could be summarized by the difference in mean survival time. The mean for a non-negative random variable is defined as $E\{T\} = \int_0^\infty S(t)dt$, where $S(t) = P(T > t)$ is the survival function of the random variable T . However, the estimate of the mean is not defined when $\hat{S}(t) > 0$ for all observed t , a situation regularly experienced with even light censoring. Furthermore, no consistent estimator exists for $E\{T\}$ when $S(t_{max}) > 0$, where t_{max} is the maximum follow-up time. Thus, a different summary measure is required as substantial censoring is experienced in lung transplantation.

The τ -restricted mean is an alternative summary measure which, by truncating observations at some time point τ , is always estimable in the presence of right-censored survival times [i.e., $E\{\min(T, \tau)\} = \int_0^\tau S(t)dt$] [3]. By choosing a value for τ within the observed follow-up time, the restricted mean is an estimable summary measure with a direct interpretation and close relationship to the mean. For example, the average difference in post-transplant survival over one year between high-volume and low-volume centers is the difference in one-year restricted means. However, estimating the difference of the restricted mean survival time with data from observational studies, such as the data available in the lung transplantation example, is difficult due to potential confounding. In particular, the difference in the area under the Kaplan-Meier survival curve up to time τ is not necessarily a consistent estimator of the causal restricted mean difference between the two treatment groups.

To account for imbalances in potential confounders between the two treatments, several researchers have proposed estimating the covariate-adjusted restricted mean difference by modeling the covariate-adjusted survival distribution (i.e., estimating the conditional survival function), and then marginalizing over the covariate distribution to estimate the restricted mean difference. This is referred to as the “regression” approach throughout the rest of the paper. For such an approach, the model for the covariate-adjusted survival distribution defines the estimator of the restricted mean difference.

The covariate-adjusted survival distribution is often assumed to follow a proportional hazards model. For example, Karrison [4] proposed modeling the survival time distribution

as a proportional hazards model with a piecewise constant baseline hazard function. As a natural alternative, Zucker [5] proposed a proportional hazards model with an unspecified baseline hazard function, i.e., a Cox proportional hazards model [6]. Both Karrison and Zucker assume that the proportional effect of the covariates on the hazard is the same for each treatment. Chen and Tsiatis [7] relax this assumption by estimating separate baseline hazard functions and covariate effects for each treatment.

Unfortunately, the proportional hazards assumption may not hold in many applications. For example, centers with greater lung transplant volume are more likely to perform bilateral lung transplants as opposed to single lung transplants, and the type of lung transplant is well-known to violate the proportional hazards assumption [8]. As such, current approaches, which rely on the proportional hazards assumption, may produce biased and inefficient estimates of the restricted mean difference between high volume and low volume centers. As an alternative, we could estimate the survival distribution with an accelerated failure time model, but the estimator would then be biased if the accelerated failure time assumption is violated. Rather than rely on a single parametric or semi-parametric model, we pursue a flexible estimator of the conditional survival function and, therefore, of the restricted mean difference that performs well across a wide range of situations for a given sample size.

In particular, we investigate ‘stacked survival models’ for estimating the conditional survival function. Stacking finds a weighted average of several conditional survival function estimators by minimizing predicted error [9]. Since the minimization is based on predicted error, stacking can include parametric models, semi-parametric models and non-parametric models. This allows more weight to be given to the model that most accurately estimates the underlying survival function for a given situation and sample size. Wey et al. [9] demonstrated that stacked survival models is competitive with parametric models and the Cox proportional hazards model for estimating the conditional survival function when assumptions are satisfied, but performs better when assumptions are violated. Since better estimators of the conditional survival function should lead to a better estimator of the restricted mean difference, we show that the mean-squared error of the restricted mean difference estimator is bounded, in part, by the mean-squared error of the conditional survival function estimator. Therefore, our goal is improving the estimation of the restricted mean difference in a wide range of scenarios by estimating the conditional survival function with stacked survival models.

Section 2 introduces the estimator of the restricted mean treatment effect and bounds the mean-squared error (MSE) of the restricted mean treatment effect by the mean-squared error of the conditional survival function. Section 3 outlines the application of stacked survival models to restricted mean treatment effect estimation. A simulation study evaluates the finite sample performance of the proposed estimator in Section 4, which is then applied in Section 5 to an observational registry of post-lung transplantation survival from the United Network for Organ Sharing. Concluding remarks are presented in Section 6.

2. Proposed Estimator

Throughout the paper, random variables and observed variables are distinguished by capital and lower case letters, respectively. The treatment or condition, which the restricted mean survival is compared across, is denoted by a_i , where i denotes the subject, and follows the Bernoulli random variable A (i.e., $A = \{0, 1\}$). Additional covariates, denoted by vector \mathbf{x}_i , are measured at the beginning of the study and follow the distribution of the random variable \mathbf{X} . For this paper, we define the non-negative survival time random variable as $T = T^0 \cdot I(A = 0) + T^1 \cdot I(A = 1)$, where T^0 and T^1 are the possibly unobservable survival time random variables had a patient received treatment 0 and 1, respectively. We assume that there are no unmeasured confounders; that is, the set of potential outcomes, (T^0, T^1) , is conditionally independent of A given \mathbf{X} , i.e., $(T^0, T^1) \perp A \mid \mathbf{X}$, where \perp denotes statistical independence. The censoring time is c_i , which follows the distribution of the continuous non-negative random variable C and is assumed to be conditionally independent of (T^0, T^1) (i.e., $(T^0, T^1) \perp C \mid \{\mathbf{X}, A\}$). Hence a sample of right censored survival data for n patients is $\{y_i, \delta_i, a_i, \mathbf{x}_i\}$, $i = 1, \dots, n$, where $y_i = \min(t_i, c_i)$ and $\delta_i = I(t_i < c_i)$.

Now let $S^{(A=a)}(t \mid \mathbf{X} = \mathbf{x}) = P(T > t \mid \mathbf{X} = \mathbf{x}, A = a)$ and $G^{(A=a)}(t \mid \mathbf{X} = \mathbf{x}) = P(C > t \mid \mathbf{X} = \mathbf{x}, A = a)$ be, respectively, the treatment-specific conditional survival functions for the covariate-adjusted survival and censoring distributions for treatment a . For brevity, we write throughout that $S^{(A=a)}(t \mid \mathbf{X} = \mathbf{x}) = S^{(a)}(t \mid \mathbf{x})$ and $G^{(A=a)}(t \mid \mathbf{X} = \mathbf{x}) = G^{(a)}(t \mid \mathbf{x})$.

2.1. Restricted Mean Treatment Effects

Following the outline of Chen and Tsiatis [7], we estimate the restricted mean for each treatment group with the “regression” approach, which involves modeling the covariate-adjusted survival time distribution. In particular, the restricted mean for treatment a is defined as

$$\mu(\tau, a) \equiv E\{\min(T^a, \tau)\} = E_{\mathbf{X}}[E\{\min(T^a, \tau) \mid \mathbf{X} = \mathbf{x}\}] = E_{\mathbf{X}}[E\{\min(T, \tau) \mid \mathbf{X} = \mathbf{x}, A = a\}] = E_{\mathbf{X}}\left\{\int_0^{\tau} S^{(a)}(t \mid \mathbf{x}) dt\right\}, \tag{1}$$

where (1) holds due to the assumption that $(T^0, T^1) \perp A \mid \mathbf{X}$. It is important to note that the outer expectation in (1) is taken with respect to the marginal, rather than conditional, covariate distribution.

After estimating the treatment-specific conditional survival functions, $S^{(a)}(t \mid \mathbf{x})$, we estimate the expectation over the covariate space with the empirical covariate distribution. Thus, the estimator for the τ -restricted mean of the potential outcome T^a is

$$\hat{\mu}(\tau, a) = \frac{1}{n} \sum_{i=1}^n \int_0^\tau \hat{S}^{(a)}(t|\mathbf{x}_i) dt, \tag{2}$$

where $\hat{S}^{(a)}(t|\mathbf{x}_i)$ is the estimate of $S^{(a)}(t|\mathbf{x}_i)$. In practice, a closed form solution of equation (2) may not exist, and we approximate equation (2) by

$$\hat{\mu}(\tau, a) \approx \frac{1}{n} \sum_{l=1}^n \sum_{j=1}^{N_\tau} (t_{(j)} - t_{(j-1)}) \times \hat{S}^{(a)}(t_{(j-1)}|\mathbf{x}_i), \tag{3}$$

where $t_{(j)}$ are the ordered event times, i.e., $t_{(j)} - t_{(j-1)} \geq 0$ for all $j = 1, \dots, n$, and N_τ is one more than the number of event times less than τ . If no event time equals τ (i.e., $t_{(j)} < \tau$ for all $j = 1, \dots, n$), then $t_{(N_\tau)} = \tau$ and $t_{(N_\tau-1)}$ is the largest event time less than τ . For two treatments, the estimated difference in restricted mean survival time is

$$\hat{\gamma}(\tau) = \hat{\mu}(\tau, a=1) - \hat{\mu}(\tau, a=0), \tag{4}$$

which also corresponds to the difference in the area under the estimated survival curves for the two potential outcomes up to time τ .

2.2. Influence of Treatment-Specific Conditional Survival Functions

There is a clear connection between the restricted mean treatment effect and the treatment-specific conditional survival functions. We formalize this connection by placing an upper bound on the MSE of the restricted mean treatment effect estimator (abbreviated as MSE_{RM}) that depends, in part, on the MSE of the estimators for the treatment-specific conditional survival functions (abbreviated as MSE_{CSF}).

Theorem 1 Define the MSE_{RM} for γ as $MSE[\hat{\gamma}(\tau)] = E[\hat{\gamma}(\tau) - \gamma(\tau)]^2$, then

$$MSE[\hat{\gamma}(\tau)] \leq \tau \times \left[MSE_\tau\{\hat{S}^{(0)}\} + MSE_\tau\{\hat{S}^{(1)}\} - 2 \times [Cov_\tau\{\hat{S}^{(0)}, \hat{S}^{(1)}\} + Bias_\tau\{\hat{S}^{(0)}, \hat{S}^{(1)}\}] \right],$$

where $MSE_\tau\{\hat{S}^{(a)}\} = E \int_0^\tau [\hat{S}^{(a)}(t|\mathbf{x}) - S^{(a)}(t|\mathbf{x})]^2 dt$ is the MSE_{CSF} for treatment a , $Cov_\tau\{\hat{S}^{(0)}, \hat{S}^{(1)}\} = E_{\mathbf{X}} \int_0^\tau [\hat{S}^{(1)}(t|\mathbf{x}) - E_{LS|\mathbf{X}}\{\hat{S}^{(1)}(t|\mathbf{x})\}][\hat{S}^{(0)}(t|\mathbf{x}) - E_{LS|\mathbf{X}}\{\hat{S}^{(0)}(t|\mathbf{x})\}] dt$ is the covariance between the treatment-specific conditional survival functions, and $Bias_\tau\{\hat{S}^{(0)}, \hat{S}^{(1)}\} = E_{\mathbf{X}} \int_0^\tau [E_{LS|\mathbf{X}}\{\hat{S}^{(1)}(t|\mathbf{x})\} - S^{(1)}(t|\mathbf{x})][E_{LS|\mathbf{X}}\{\hat{S}^{(0)}(t|\mathbf{x})\} - S^{(0)}(t|\mathbf{x})] dt$ is the interaction of bias between the treatment-specific conditional survival functions. The unconditional expectations are over the random variable for the covariate space (\mathbf{X}) and the sampling distribution of the conditional survival function estimator (LS).

The inequality (see the Supporting Information for the proof) illustrates that MSE_{RM} should be associated with MSE_{CSF} . However, the last term indicates potential exceptions to the relationship that depend on the association between the estimators for the treatment-specific conditional survival functions. For example, a positive correlation between the treatment-specific conditional survival functions will tighten the bound on MSE_{RM} , which could result in a surprisingly good MSE_{RM} despite a relatively poor MSE_{CSF} . In contrast, a negative correlation between the treatment-specific conditional survival functions may result in a surprisingly poor MSE_{RM} despite relatively good MSE_{CSF} . Thus, the behavior of the separate treatment-specific conditional survival functions may loosen the association between MSE_{RM} and MSE_{CSF} . Finally, mean-squared error decomposes into a squared bias term and a variance term, which implies the following corollary

Corollary 2 Define $Bias[\hat{\gamma}(t)] = E\{\hat{\gamma}(t) - \gamma(t)\}$ and $Var[\hat{\gamma}(t)] = E\{\hat{\gamma}(t) - E\hat{\gamma}(t)\}^2$ as, respectively, the bias and variance of the restricted mean treatment effect, then

$$Bias[\hat{\gamma}(\tau)]^2 \leq \tau \times \left[MSE_{\tau}\{\hat{S}^{(0)}\} + MSE_{\tau}\{\hat{S}^{(1)}\} - 2 \times [Cov_{\tau}\{\hat{S}^{(0)}, \hat{S}^{(1)}\} + Bias_{\tau}\{\hat{S}^{(0)}, \hat{S}^{(1)}\}] \right] - Var[\hat{\gamma}(\tau)] \\ < \tau \times \left[MSE_{\tau}\{\hat{S}^{(0)}\} + MSE_{\tau}\{\hat{S}^{(1)}\} - 2 \times [Cov_{\tau}\{\hat{S}^{(0)}, \hat{S}^{(1)}\} + Bias_{\tau}\{\hat{S}^{(0)}, \hat{S}^{(1)}\}] \right],$$

since $Var[\hat{\gamma}(t)] > 0$.

Thus, the performance of the treatment-specific conditional survival functions places an upper bound on both traditional measures of performance (i.e., bias and MSE) for the restricted mean treatment effect. The bound on the squared bias, or absolute bias, is less tight than the bound on MSE_{RM} due to a positive, and potentially large, variance term. Therefore, we would expect a stronger association between MSE_{RM} and MSE_{CSF} than the association between the bias of the restricted mean treatment effect and MSE_{CSF} . The simulation study in Section 4 presents an empirical demonstration of this relationship.

3. Stacked Survival Models

Since the performance of the restricted mean treatment effect estimator is related to the performance of the treatment-specific conditional survival function estimator, we propose estimating restricted mean treatment effects with stacked survival models. By minimizing predicted squared survival error, stacked survival models estimate a weighted combination of survival models that can span parametric (e.g., Weibull model), semi-parametric (e.g., Cox proportional hazards model), and non-parametric models (e.g., random survival forests). For estimating conditional survival functions, non-parametric estimators can be preferred to parametric and semi-parametric estimators due to relaxed assumptions. However, even when misspecified, parametric and semi-parametric estimators can possess better operating characteristics in small sample sizes due to smaller variance than non-parametric estimators. Fundamentally, this is a bias-variance tradeoff situation and, by minimizing predicted error, stacking estimates an optimal combination of survival models that balances the bias-variance trade-off of each estimator for a given sample size. For example, Wey et al. [9] illustrate that

stacked survival models effectively estimate a conditional survival function across a wide range of situations.

In uncensored settings, stacking estimates the optimally weighted average of several candidate models by minimizing predicted squared error [10]. However, predicted squared error is often poorly defined in censored settings due to potentially unobserved event times. Therefore, a different loss function is required that is tailored to censored data. Wey et al. [9] show that the Brier Score has a direct relationship with the definition of mean-squared error for a conditional survival function estimate presented in Section 2.2 and evaluate stacked survival models that minimize the Brier Score [11] over a grid of time points. In addition, they show that, under certain conditions, stacking survival models with the Brier Score is uniformly consistent for the true conditional survival function. In contrast to Wey et al. [9], this paper appropriately modifies stacked survival models to estimate restricted mean treatment effects within an inference setting. In addition, we assess the improvement in the restricted mean treatment effect through stacked survival models.

The Brier Score [11] measures the predicted squared error of a conditional survival function at a particular time point. Following Lostritto et al. [12], the estimated Brier Score for a given estimator of the conditional survival function for treatment a at a single time point t can be written as

$$\widehat{\text{BS}}^{(a)}(t) = \frac{1}{|\Gamma_a|} \sum_{i \in \Gamma_a} \frac{\Delta_i(t)}{\hat{G}^{(a)}(\min\{t_i, t\} | \mathbf{x}_i)} \times \{Z_i(t) - \hat{S}^{(a)}(t | \mathbf{x}_i)\}^2, \quad (5)$$

where $Z_i(t) = \mathbb{I}(t_i > t)$, $\Delta_i(t) = \mathbb{I}(\min\{t_i, t\} < c_j)$, $\hat{G}^{(a)}(\cdot | \mathbf{x}_i)$ is the estimated conditional survival function of the censoring distribution for subjects that received the a^{th} treatment, Γ_a is the set of patients that received treatment a , and $|\Gamma_a|$ is the number of patients that received treatment a . For a fixed time t , censored observations with $c_j > t$ will contribute to the Brier Score, but the censored observations with $c_j < t$ will only contribute to the Brier Score indirectly through the estimation of $G^{(a)}(T_i(t) | \mathbf{x}_i)$.

The conditional survival function for each treatment group is estimated by a separate stacked survival model with the same set of m candidate models, i.e., the models included in the stack. Each model has a corresponding treatment-specific conditional survival function

estimate, say $\hat{S}_k^{(a)}(t | \mathbf{x})$, for $k = 1, \dots, m$. The set of candidate survival models influences the performance of the final conditional survival function estimator. Thus, as recommended by Breiman [10] and Wey et al. [9], we include a ‘diverse set’ of candidate survival models. We note, however, that the stacked survival model could include other combinations of survival models. Since the goal is estimating the entire conditional survival function up to time τ ,

stacked survival models minimize the sum of $\widehat{\text{BS}}^{(a)}(t)$ over a set of time points, say $\{t_1, \dots, t_s\}$. The estimated stacking weights for the m models are the solution to the weighted least

squares problem with the additional constraints that $\hat{\alpha}_k^{(a)} \geq 0$ for $k = 1, \dots, m$ and

$$\sum_{k=1}^m \hat{\alpha}_k^{(a)} = 1;$$

$$\hat{\alpha}^{(a)} = \arg \min_{\alpha^{(a)}, \alpha_k^{(a)} \geq 0} \sum_{r=1}^s \sum_{i \in \Gamma_a} \frac{\Delta_i(t_r)}{\hat{G}^{(a)}(\min\{t_i, t_r\} | \mathbf{x}_i)} \times \left\{ Z_i(t_r) - \sum_{k=1}^m \alpha_k \hat{S}_k^{(a,-i)}(t_r | \mathbf{x}_i) \right\}^2, \tag{6}$$

where $\hat{S}_k^{(a,-i)}(t | \mathbf{x}_i)$ is the survival estimate from the k^{th} model while leaving the i^{th} observation out during the fitting process, which is implemented by leaving the i^{th} observation out in the fitting process, i.e., leave-one-out cross-validation. This ensures that stacking does not reward models that over-fit the data.

After minimizing equation (6), the stacked estimate of the conditional survival function for treatment a is

$$\hat{S}^{(a)}(t | \mathbf{x}) = \sum_{k=1}^m \hat{\alpha}_k^{(a)} \hat{S}_k^{(a)}(t | \mathbf{x}), \tag{7}$$

where $\hat{S}_k^{(a)}(t | \mathbf{x})$ is the k^{th} survival model estimated with all observations on treatment a . The treatment-specific restricted means and the restricted mean treatment effects are then estimated with equations (3) and (4), respectively.

We estimate confidence intervals with the jackknife [13], which can be estimated during the leave-one-out cross-validation step for minimizing equation (6). Thus, we only need to minimize equation (6) an additional n times to estimate the standard error of the restricted mean treatment effect. In contrast, a bootstrap estimator for $\chi(\hat{\tau})$ would require running the entire process to minimize equation (6), including the estimation of $\hat{S}_k^{(a,-i)}(\cdot | \mathbf{x}_i)$, on each bootstrapped data set, which is computationally expensive. Finally, let $\chi(\hat{\tau})^{(-i)}$ be the restricted mean treatment effect from leaving the i^{th} observation out during the fitting process, then the jackknife variance estimate of the restricted mean treatment effect is

$$\widehat{\text{Var}}\{\hat{\gamma}(\tau)\} = \frac{n-1}{n} \sum_{i=1}^n \left(\hat{\gamma}(\tau)^{(-i)} - \frac{1}{n} \sum_{j=1}^n \hat{\gamma}(\tau)^{(-j)} \right)^2.$$

A 95% confidence interval is then estimated through a Normal approximation with $\hat{\gamma}(\tau) \pm 1.96 \cdot \widehat{\text{Var}}\{\hat{\gamma}(\tau)\}^{\frac{1}{2}}$.

Remark 1. The Brier Score measures agreement at only one particular time point. As such, the set of values, $\{t_1, \dots, t_s\}$, over which the Brier Score is minimized [see equation (6)] has implications for performance. In particular, care should be taken to avoid picking only very small, or very large values. Wey et al. [9] recommend at least nine evenly spaced quantiles of the observed event distribution to ensure good estimation of the conditional survival function. The effect of the set of time points over which the Brier Score is minimized is investigated in the Supporting Information.

Remark 2. Wey et al. [9] show that, given the stack contains a uniformly consistent estimator of the conditional survival function, the stacked estimator is uniformly consistent for the underlying conditional survival function. Therefore, when at least one model within the set of candidate survival models is correctly specified, $\hat{\gamma}(\tau)$ is consistent for the true restricted mean treatment effect by the dominated convergence theorem [14]. The proposed estimator is, therefore, consistent in a wider range of scenarios than previous methods that assume a proportional hazards model to estimate the conditional survival distribution.

4. Simulation Study

4.1. Set-up

The simulation study evaluates the finite sample performance for estimating the causal restricted mean treatment effect with stacked survival models. We consider six different data-generating scenarios, indexed by q , for the covariate-adjusted survival distribution of the potential outcomes. When $q = 1, 2$ then $T_q^A \sim \text{Exp}[\exp\{\lambda_q^A\}]$, where

$E\{T_q^A\} = 1/\exp\{\lambda_q^A\}$; when $q = 3, 4$ then $T_q^A \sim \text{Gamma}[\text{scale}=\exp\{\lambda_q^A\}, \text{shape}=2.5]$, where

$E\{T_q^A\} = 2.5 \cdot \exp\{\lambda_q^A\}$; when $q = 5, 6$ then $T_q^A \sim \exp\{N(\lambda_q^A, 0.5)\}$ which corresponds to a log-Normal distribution with a variance of 0.5. The covariate effects are

$$\begin{aligned}\lambda_1^A &= -4.50 \times I\{A=0\} - 3.50 \times I\{A=1\} - 0.125 \cdot \{x_1+x_2+x_3+x_4\} \\ \lambda_2^A &= -0.70 \times I\{A=0\} - 1.70 \times I\{A=1\} - 1.000 \cdot \{\Phi(4 \times x_1)+\Phi(4 \times x_2)+\Phi(4 \times x_3)+\Phi(4 \times x_4)\} \\ \lambda_3^A &= 3.50 \times I\{A=0\}+3.00 \times I\{A=1\} - 0.125 \cdot \{x_1+x_2+x_3+x_4\} \\ \lambda_4^A &= 4.50 \times I\{A=0\}+4.00 \times I\{A=1\} - 0.500 \cdot \{\Phi(4 \times x_1)+\Phi(4 \times x_2)+\Phi(4 \times x_3)+\Phi(4 \times x_4)\} \\ \lambda_5^A &= 3.75 \times I\{A=0\}+3.25 \times I\{A=1\} - 0.125 \cdot \{x_1+x_2+x_3+x_4\} \\ \lambda_6^A &= 4.00 \times I\{A=0\}+3.50 \times I\{A=1\} - 0.500 \cdot \{\Phi(4 \times x_1)+\Phi(4 \times x_2)+\Phi(4 \times x_3)+\Phi(4 \times x_4)\},\end{aligned}$$

where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal distribution (i.e., the non-linear effect is a ‘smooth step function’). The censoring distributions are defined

similarly with λ_q^A replaced by β_q^A , and are designed to achieve a marginal censoring rate of approximately 20% to 30%. The censoring distributions are

$$\begin{aligned}\beta_1^A &= -4.895 \times I\{A=0\} - 5.395 \times I\{A=1\} - 0.0625 \cdot \{x_1+x_2+x_3+x_4\} \\ \beta_2^A &= -3.680 \times I\{A=0\} - 4.680 \times I\{A=1\} - 0.5000 \cdot \{\Phi(4 \times x_1)+\Phi(4 \times x_2)+\Phi(4 \times x_3)+\Phi(4 \times x_4)\} \\ \beta_3^A &= 3.780 \times I\{A=0\}+3.280 \times I\{A=1\} - 0.0625 \cdot \{x_1+x_2+x_3+x_4\} \\ \beta_4^A &= 4.765 \times I\{A=0\}+4.265 \times I\{A=1\} - 0.5000 \cdot \{\Phi(4 \times x_1)+\Phi(4 \times x_2)+\Phi(4 \times x_3)+\Phi(4 \times x_4)\} \\ \beta_5^A &= 3.900 \times I\{A=0\}+3.400 \times I\{A=1\} - 0.0625 \cdot \{x_1+x_2+x_3+x_4\} \\ \beta_6^A &= 4.150 \times I\{A=0\}+3.650 \times I\{A=1\} - 0.5000 \cdot \{\Phi(4 \times x_1)+\Phi(4 \times x_2)+\Phi(4 \times x_3)+\Phi(4 \times x_4)\}.\end{aligned}$$

For brevity, we refer to scenarios 1 and 2 as, respectively, the linear and non-linear exponential scenarios; scenarios 3 and 4 as, respectively, the linear and non-linear gamma scenarios; and scenarios 5 and 6 as, respectively, the linear and non-linear log-normal scenarios.

The covariate distribution is a four-dimensional multivariate normal with mean zero, unit variances, and a positive $AR(1)$ correlation structure with $\rho = 0.4$. To mimic observational studies with confounding, the treatment assignment depends on the covariate distribution. In particular, the probability of receiving treatment, i.e., $P(a_i = 1 | \mathbf{x}) = p_i$ is defined as $\text{logit}(p_i) = 0.5 \times (x_1 + x_2 + x_3 + x_4)$, where $\text{logit}(x) = 1/[1 + \exp(-x)]$.

Each simulation scenario evaluates performance for $\tau = 50$ with 1,000 replications and a sample size of 300. All simulations were run in R version 3.0.0 [15]. The stacking weights are estimated by minimizing the Brier Score over nine equally spaced quantiles of the observed events with the constrained minimization problem solved by the `alabama` package, which uses an augmented Lagrangian and adaptive barrier algorithm for minimizing equation (6) [16]. For the simulations in this paper, $\hat{G}^{(a)}(\cdot | \mathbf{x}_j)$ is estimated with an unconditional treatment-specific Kaplan-Meier survival estimator. In these scenarios, stacked survival models remain consistent; see the Supporting Information for a sketch of the proof.

For the stacked survival models, we consider a mixture of parametric, semi-parametric, and non-parametric candidate survival models. The parametric models are the Weibull model and log-Normal model with only linear main effects. Both models are special cases of an accelerated failure time model, while the Weibull is also a special case of a proportional hazards model. The semi-parametric models are two versions of the Cox model. The first Cox model has only linear main effects, while the second Cox model uses penalized splines for main effects with the roughness penalty set to 0.5. The `survival` package estimates both the parametric and semi-parametric models [17]. The non-parametric model is random survival forests (RSF), which is estimated with the `randomForestSRC` package with 1,000 trees grown [18]. For RSF, the confidence intervals are estimated with the jackknife after bootstrap [19], which saves computational resources by not regrowing the forest for each observation and instead uses the trees that do not include the i^{th} observation. We consider two different versions of the stacked estimator: the set of candidate survival models with and without RSF. This is important as RSF substantially increases the computational burden of stacked survival models. Thus, we want to characterize the level of improvement in estimating the restricted mean treatment effect by including RSF in the set of candidate survival models.

The different methods for estimating the restricted mean differ in their approach to estimating the treatment-specific conditional survival functions in equation (3). We are most interested in comparing the Stacked estimator to a Cox proportional hazards model with linear main effects (referred to as the ‘Cox estimator’) and a Cox proportional hazards model with penalized splines (referred to as the ‘Splines estimator’) due to their role in previous methods for estimating restricted mean treatment effects. The Cox estimator was proposed by Chen and Tsiatis [7], while the Splines estimator is a straightforward extension of the Cox estimator that should be more robust in a variety of situations. For the sake of completeness, we also present the performance of the restricted mean treatment effect for the other models in the set of candidate survival models.

The methods are compared on the basis of percent relative bias, i.e., $100 \times [E\hat{\alpha}(t) - \alpha(t)]/\alpha(t)$, and the ratio of mean squared error (MSE) to the Cox estimator, where $MSE = E\{\hat{\alpha}(t) - \alpha(t)\}^2$. Confidence interval performance is assessed with two measures: the ratio of average jackknife standard error to the Cox estimator and coverage probability. We also present the ratio of ‘integrated squared survival error’ (ISSE) to the Cox estimator for each method: $ISSE_{\tau}\{\hat{S}\} = \frac{1}{2} \times E \int_0^{\tau} \{[\hat{S}^{(0)}(t|\mathbf{x}) - S^{(0)}(t|\mathbf{x})]^2 + [\hat{S}^{(1)}(t|\mathbf{x}) - S^{(1)}(t|\mathbf{x})]^2\} dt$, which corresponds to the average of the mean-squared error of the treatment-specific conditional survival functions presented in Section 2.2. All expectations are approximated by averaging across 1,000 replications.

4.2. Results

Tables 1 and 2 present the results of the exponential and gamma settings, respectively. The Stacked estimator without random survival forests (RSF) possesses similar bias and MSE than the Stacked estimator with RSF across the different scenarios with the non-linear log-Normal scenario possessing the largest difference. Since RSF has substantial computational costs, we focus on comparing the Stacked estimator without RSF to the competing estimators.

In the exponential setting (Table 1), the Stacked estimator possesses similar, or slightly more, bias than the Cox estimator for the linear scenario and similar, or more, bias than the Splines estimator for both linear and non-linear scenarios. The Stacked estimator has approximately $\sim 4\%$ less bias and 20% lower MSE than the misspecified Cox estimator in the non-linear scenario. In addition, the Stacked estimator possesses approximately 15 – 20% lower MSE than the Splines estimator. Finally, the Stacked estimator possesses an approximately 3 – 10% lower standard deviation than the Splines estimator for both linear and non-linear scenarios. The Stacked estimator is therefore competitive in the exponential setting with linear effects and more efficient than both the Cox and Splines estimators in the non-linear scenario.

In the gamma setting (Table 2), the Stacked estimator possesses 5%–9% less relative bias than the Cox estimator and 3% – 5% less relative bias than the Splines estimator for both linear and non-linear scenarios. In addition, the Stacked estimator possesses 10% – 20% lower MSE than the Cox estimator and 25% lower MSE than the Splines estimator. The Stacked estimator has approximately 3% lower standard deviation than the Cox estimator and 10 – 15% lower standard deviation than the Splines estimator. The performance of the Stacked estimator is similar in the log-Normal setting (Table 3): the Stacked estimator has similar, or lower, bias than both the Cox and Splines estimators with 10 – 25% lower MSE than the Cox or Splines estimators. The results of the gamma and log-Normal settings demonstrate the ability of the Stacked estimator to also perform well in non-proportional hazard settings.

Wey et al. [9] motivated stacking with the goal of balancing, for a given sample size, the low variance of potentially misspecified parametric models with robust but potentially inefficient semi-parametric and non-parametric models. This is illustrated in the simulation study here. In the linear exponential scenario, the Stacked estimator possesses similar MSE as the

correctly specified Cox and Splines estimators, which is likely due to the inclusion of a correctly specified parametric Weibull model. Yet, even when the Weibull model is misspecified in the non-linear gamma scenario, the Stacked estimator still performs better than both the Cox and Splines estimators despite the lack of a correctly specified model. This ability to adaptively find a good balance between the low variance of potentially misspecified parametric models with the more robust, but still potentially misspecified, semi-parametric models (e.g., a Cox model with penalized splines) is the most appealing aspect of stacked survival models.

Section 2.2 argues that performance of the treatment-specific conditional survival functions will be more tightly associated with MSE of the restricted mean treatment effect than bias. In general, this relationship was observed in the simulation study for the parametric and semi-parametric models although RSF possessed significant departures. In particular, RSF occasionally had the best MSE for the restricted mean treatment effect despite possessing worse ISSE than the Cox model in every scenario. These exceptions are likely due to, for example, bias terms in the same direction for the restricted means of individual treatments, which leads to the bias terms canceling out for the restricted mean treatment effect despite relatively large MSE_{CSF} terms for the treatment-specific conditional survival functions. It is well recognized that non-parametric approaches usually possess lower bias while having higher variance. However, when taking a difference such as with a restricted mean treatment effect, the bias can cancel between the treatment-specific conditional survival function estimates which may place higher priority on conditional survival function estimators with a lower variance.

Remark 3. The Supporting Information contains several additional investigations into the simulation study: a larger sample size, the choice of time points over which the stacking weights are determined, and the impact of a misspecified censored distribution.

5. Effect of Center Volume in Lung Transplantation

It is well-known that higher center volume is associated with better graft survival in lung transplantation (i.e., time to death or retransplantation) [1, 2, 8, 20]. However, previous studies have only estimated the hazard ratio outside of a causal framework. Therefore, we estimate the restricted mean treatment effect of post-transplant survival between high volume centers (more than 100 lung transplants over the past two years [20]) and low volume centers (less than, or equal to, 100 lung transplants over the past two years) with data from an observational registry of post-lung transplant survival.

The United Network for Organ Sharing (UNOS) collects patient information, donor information and survival status of every solid organ transplant performed in the United States. This analysis only includes lung transplants performed between January 1, 2008 and December 31, 2011 in adult recipients receiving their first lung-only transplantation. Many of the covariates collected by UNOS possess imbalances between low volume and high volume centers (see Table 1 in the Supporting Information), which indicates potential confounding of the restricted mean treatment effect. We therefore adjust for potential confounding with several patient related covariates: gender, age, lung allocation score, native

disease grouping (obstructive, vascular, cystic and restrictive), distance walked in six minutes, ventilator use, level of oxygen use, and type of transplant (single versus bilateral lung transplant). In addition, we adjust for several donor related covariates: age over 55 years, African American race, smoking history greater than 20 pack years, and height difference between donor and recipient. The event of interest is time to death or retransplantation. A total of 5,499 transplanted patients were included in this analysis with approximately 76% censoring.

This is a particularly demonstrative example due to the questionable applicability of the proportional hazards assumption in lung transplantation. In particular, the type of lung transplant is well-known to violate the proportional hazards assumption, e.g., see Thabut et al. [2]. In addition, proportional hazard tests for our particular data set suggests non-proportional hazards among the native disease grouping, height difference between donor and recipient, and ventilator use. Thus, we may improve estimation of the restricted mean treatment effect by allowing for models that extend beyond proportional hazards through stacked survival models.

The stacked survival model includes two versions of the Cox model, a Weibull model, a log-Normal model, and random survival forests. Each model includes the entire set of covariates due to expected associations with survival. The height difference between donor and recipient is expected *a priori* to possess a non-linear relationship with survival. In particular, donors and recipients with similar heights are expected to possess the best survival while larger height differences are expected to possess worse survival. Thus, the Weibull, log-Normal, and the first Cox model fit linear main effects to all continuous covariates except the height difference between donor and recipient, which is fit with a quadratic main effect. A second Cox model fits penalized splines to each continuous covariate, which provides additional flexibility in the case that covariates possess unexpected non-linearity. For RSF, the minimum number of observations in each node is selected by minimizing predicted error as implemented in the `randomForestSRC` package whereas the package default is used for the number of randomly selected splits. The 95% confidence intervals are estimated with the leave-one-out jackknife.

Table 4 presents the one-year and three-year estimated causal restricted mean treatment effects between high volume and low volume centers in addition to the estimated treatment-specific model weights for each survival model. There is notable variability in the estimated model weights between the two treatment groups. Specifically, the Cox model without splines receives no weight for low volume centers but a weight of 0.339 for high volume centers. In contrast, the Cox model with splines consistently receives weight regardless of center volume: 0.225 and 0.275 for low and high volume centers, respectively. Similarly, RSF receives a large proportion of weight regardless of center volume: 0.532 and 0.386 for low and high volume centers, respectively. Additionally, the log-Normal model receives no weight for high volume centers but receives a weight of 0.243 for low volume centers. Regardless, as noted by Wey et al. [9], the estimated model weights are not an indication of a correct model due to potential instability in the minimization procedure from similar estimated conditional survival functions. Instead, the model weights help describe the process of combining individual survival models within the stack.

The stack estimates a seven day and twenty-five day difference between large and small centers for, respectively, the one-year and three-year restricted mean treatment effects. This treatment effect is neither the smallest nor the largest restricted mean treatment effect compared to the individual survival models, which reflects in part the weight given to the smallest difference and largest differences among the candidate survival models. The stacked estimator also balances the wide confidence intervals of the non-parametric RSF estimator with the more narrow confidence intervals of the parametric and semi-parametric estimators.

We are particularly interested in comparing the stacked estimator to the Cox model without splines due to its role in previously proposed methods [7]. The stack estimates an approximately 30% larger difference between large and small centers although the confidence interval is wider than the Cox model for both one-year and three-year restricted mean treatment effects. The larger difference between large and small centers for the one-year restricted mean treatment effect is clinically meaningful as earlier survival is indicative of peri-operative and early post-operative mortality, which is relatively high for lung transplantation. However, the wide confidence interval that includes zero contrasts with previous studies that have found differences between centers based on center volume [1, 2]. Broadly, this difference could be due to estimating the restricted mean treatment effect instead of the hazard ratio, or the estimation of the causal rather than associational effects. In addition, the stacked estimator gives substantial weight to the non-parametric RSF estimator which possesses wider confidence intervals than the parametric and semi-parametric models. Thus, the stacked estimator could also be capturing additional variability. Regardless, the stacked estimator suggests that large centers possess better survival over one and three years although there is more uncertainty than suggested by previous studies.

6. Concluding Remarks

We explore improving the estimation of causal restricted mean treatment effects through better estimation of the conditional survival function. In most application areas, there is little *a priori* information to suggest an appropriate distributional assumption for the survival time or functional form of the covariates. This motivates flexibly estimating restricted mean treatment effects for observational studies with stacked survival models, which can effectively estimate the conditional survival function in a wide range of scenarios [9]. The simulation study illustrates that stacked survival models achieve similar mean-squared error (MSE) of the restricted mean treatment effect in proportional hazards scenarios as the estimator based on a Cox proportional hazards model [7]. In addition, when the proportional hazards assumption is violated, stacked survival models can substantially reduce the bias and MSE of the causal restricted mean treatment effect compared to the estimator based on a Cox proportional hazards model. Thus, stacked survival models can robustly estimate restricted mean treatment effects in a wide range of scenarios.

As noted in Section 4.2, the Stacked estimator can gain additional flexibility by including random survival forests (RSF) although inclusion of RSF did not noticeably improve the performance of the restricted mean treatment effect in the simulation study. The Supporting Information demonstrates that similar differences in performance exist at larger sample sizes when RSF is included in the stack for the scenarios considered in this paper. Regardless, the

relatively good performance of the Stacked estimator without RSF and the high computational costs associated with RSF precludes inclusion in every situation. Further investigation is warranted on the selection of survival models to include in the stack.

We also note that RSF requires decisions that have a potentially large impact on the performance of the conditional survival function estimate. For example, RSF requires selecting tuning parameters (e.g., the minimum size of each node) that can have large impact on the performance of the conditional survival function. In addition, the `randomForestSRC` and `randomSurvivalForest` packages have different definitions for the tuning parameter that controls the size of individual nodes. In particular, the `randomForestSRC` package requires that each node have a minimum number of unique time points (ignoring censoring status) in each node, while the `randomSurvivalForest` package requires that each node have a minimum number of unique *event* times in each node. This difference can lead to noticeably different performance in the same scenarios with the same tuning parameter values. The impact of these tuning parameters and associated definitions deserve further investigation.

There are limitations for estimating restricted mean treatment effects with stacked survival models. First, it is difficult to evaluate the final stacked survival model as the minimization procedure acts as a ‘black box’. Thus, the fit of the final stacked survival model can be difficult to assess in practice. Second, stacked survival models are potentially computationally expensive depending on the models included in the stack. For example, random survival forests significantly inflate the computational costs of stacked survival models compared to including only parametric and semi-parametric survival models. Wey et al. [9] provides a more detailed description of the potential computational costs associated with stacked survival models. Third, stacked survival models based on the Brier Score require specifying a model for the censoring distribution to estimate the stacking weights. This is avoided by current estimators of the restricted mean treatment effect such as the method proposed by Chen and Tsiatis [7]. However, the simulation study demonstrates that the stacked survival models based on the Brier Score perform relatively well even when the censoring distribution is incorrectly modeled. Finally, an issue for any estimator of restricted mean treatment effects, the value of τ for restricted mean treatment effects can be difficult to select in practice although emphasis should be given to values that are clinically meaningful.

The proposed approach for stacked survival models minimizes the Brier Score for combining different estimators of the conditional survival function. However, different loss functions could be minimized for combining estimators of the conditional survival function. For example, Polley and van der Laan [21] note estimators could be combined based on discrete conditional hazard functions, which indirectly estimates the conditional survival function. This approach would also avoid specification of the censoring distribution although extensions to continuous survival times present additional challenges due, in part, to the unbounded nature of the hazard function.

There are two main approaches to estimating the causal restricted mean treatment effect: the “regression approach” pursued in this paper and an approach based on inverse-probability weighting (IPW) for treatment assignment and censoring [22, 23]. The IPW approach

requires forming models for the censoring and treatment distributions. The “regression” approach is a more efficient estimator of the restricted mean difference when the conditional survival model has been correctly specified. However, the IPW approach is sometimes preferred because standard methods that estimate the conditional survival distribution may be overly restrictive (e.g., the Cox proportional hazards model). The flexibility of stacked survival models may mitigate some of the concerns of the “regression” approach by allowing additional flexibility across a variety of data generating mechanisms.

There has been recent research on using model averaging to account for uncertainty in the confounders of the treatment-outcome relationship [24, 25]. However, previous work has assumed that the structure of the relationship between the covariates and outcome was known (e.g., linear relationship between covariates and log-hazard) although in practice there is usually little evidence to support *a priori* assumptions on the survival distribution and the functional form of the covariate. We demonstrate that principally averaging different model structures can lead to substantially better performance in the estimation of the causal restricted mean treatment effect. Thus, an interesting avenue for future research would be to consider the selection of covariates based on both the outcome and treatment models under varying distributional assumptions and functional forms for the outcome model.

A conditional survival function is required by many methods besides restricted mean treatment effects: for example, censored quantile regression [26], time-dependent ROC curves [27], inverse probability-of-censoring weighted estimators [28], model-free contrast approaches [29], and dynamic treatment regime methods [30]. Similar to restricted mean treatment effect estimation, all of these methods have traditionally used a single Cox model or a non-parametric method to estimate the conditional survival function. The improvement in estimation of restricted mean treatment effects shown here demonstrates that stacked survival models deserve consideration for a wide spectrum of methods.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Contract/grant sponsor: Research reported in this publication was supported in part by NIH grants UL1TR000114, U54MD007584, G12MD007601, and P20GM103466.

References

1. Weiss ES, Allen JG, Meguid RA, Patel ND, Merlo CA, Orens JB, Baumgartner WA, Conte JV, Shah AS. The impact of center volume on survival in lung transplantation: An analysis of more than 10,000 cases. *The Annals of Thoracic Surgery*. 2009; 88:1062–1070. [PubMed: 19766782]
2. Thabut G, Christie JD, Kremers WK, Fournier M, Halpern SD. Survival differences following lung transplantation among us transplant centers. *Journal of the American Medical Association*. 2010; 304:53–60. [PubMed: 20606149]
3. Royston P, Parmar MK. Restricted mean survival time: an alternative to the hazard ratio for the design and analysis of randomized trials with a time-to-event outcome. *BMC Medical Research Methodology*. 2013; 13:1–15. [PubMed: 23297754]

4. Karrison T. Restricted mean life with adjustment for covariates. *Journal of the American Statistical Association*. 1987; 82:1169–1176.
5. Zucker DM. Restricted mean life with covariates: Modification and extension of a useful survival analysis method. *Journal of the American Statistical Association*. 1998; 83:702–709.
6. Cox DR. Regression models and life-tables. *Journal of the Royal Statistical Society, Series B*. 1972; 34:187–220.
7. Chen PY, Tsiatis AA. Causal inference on the difference of the restricted mean lifetime between two groups. *Biometrics*. 2001; 57:1030–1038. [PubMed: 11764241]
8. Thabut G, Christie JD, Ravaud P, Castier Y, Dauriat G, Jebrak G, Fournier M, Leseche G, Porcher Raphael, et al. Survival after bilateral versus single-lung transplantation for idiopathic pulmonary fibrosis. *Annals of Internal Medicine*. 2009; 151:767–774. [PubMed: 19949142]
9. Wey A, Connett J, Rudser K. Combining parametric, semi-parametric, and non-parametric survival models with stacked survival models. *Biostatistics*. 2015; 16:537–549. [PubMed: 25662068]
10. Breiman L. Stacked regressions. *Machine Learning*. 1996; 24:49–64.
11. Graf E, Schmoor C, Sauerbrei W, Schumacher M. Assessment and comparison of prognostic classification schemes for survival data. *Statistics in Medicine*. 1999; 18:2529–2545. [PubMed: 10474158]
12. Lostritto K, Strawderman RL, Molinaro AM. A partitioning deletion/substitution/addition algorithm for creating survival risk groups. *Biometrics*. 2012; 68:1146–1156. [PubMed: 22519965]
13. Efron, B.; Tibshirani, R. *An introduction to the bootstrap*. Chapman and Hall; 1993.
14. Ferguson, TS. *A course in large sample theory*. Chapman and Hall/CRC; 1996.
15. R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing; Vienna, Austria: 2014. URL <http://www.R-project.org/>
16. Varadhan, R. *alabama: Constrained nonlinear optimization*. 2012. URL <http://CRAN.R-project.org/package=alabama>, rpackage version 2011.9-1
17. Therneau, T. *Survival analysis, including penalized likelihood*. 2013. URL <http://CRAN.R-project.org/package=survival>, r package version 2.37-4
18. Ishwaran, H.; Kogalur, U. *Random Forests for Survival, Regression and Classification (RF-SRC)*. 2015. URL <http://cran.r-project.org/web/packages/randomForestSRC/>, r package version 1.6.0
19. Wager S, Hastie T, Efron B. Confidence intervals for random forests: The jackknife and the infinitesimal jackknife. *Journal of Machine Learning Research*. 2014; 15:1625–1651. [PubMed: 25580094]
20. Tsuang WM, Vock DM, Copeland CAF, Lederer DJ. An acute change in lung allocation score and survival after lung transplantation: a cohort study. *Annals of Internal Medicine*. 2013; 158:650–657. [PubMed: 23648947]
21. Polley, EC.; Van der Laan, M. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer; 2011. Super learning for right-censored data.
22. Schaubel DE, Wei G. Double inverse-weighted estimation of cumulative treatment effects under nonproportional hazards and dependent censoring. *Biometrics*. 2011; 67:29–38. [PubMed: 20560935]
23. Zhang M, Schaubel DE. Double-robust semiparametric estimator for differences in restricted mean lifetimes in observations studies. *Biometrics*. 2012; 68:999–1009. [PubMed: 22471876]
24. Wang C, Parmigiani G, Dominici F. Bayesian effect estimation accounting for adjustment uncertainty. *Biometrics*. 2012; 68:661–686. [PubMed: 22364439]
25. Cefalu, M.; Dominici, F.; Parmigiani, G. *Model averaged double robust estimation Technical Report 149*. Division of Biostatistics, Harvard University, Harvard University; 2013.
26. Wey A, Wang L, Rudser K. Censored quantile regression with recursive partitioning based weights. *Biostatistics*. 2014; 15:170–181. [PubMed: 23975800]
27. Zheng Y, Heagerty P. Semiparametric estimation of time-dependent roc curves for longitudinal marker data. *Biostatistics*. 2004; 5:615–632. [PubMed: 15475423]
28. Fine JP, Gray RJ. A proportional hazards model for the subdistribution of a competing risk. *Journal of the American Statistical Association*. 1999; 94:496–509.

29. Rudser KD, LeBlanc ML, Emerson SS. Distribution-free inference on contrasts of arbitrary summary measures of survival. *Statistics in Medicine*. 2012; 31:1722–1737. [PubMed: 22362470]
30. Zhao Y, Zeng D, Socinski MA, Kosorok MR. Reinforcement learning strategies for clinical trials in nonsmall cell lung cancer. *Biometrics*. 2011; 67:1422–1433. [PubMed: 21385164]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Simulation results for the exponential distributed scenarios: $N = 300$, $N_{SM} = 1000$, and a marginal censoring rate of 30% for the linear scenario and 20% for the non-linear scenario. The confidence intervals are estimated with the leave-one-out jackknife. ‘Percent Relative Bias’ is the ratio of the bias and true restricted mean difference. ‘MSE Ratio’ is the ratio of MSE relative to the Cox estimator. ‘SD Ratio’ is the ratio of average jackknife standard deviation relative to the Cox estimator. ‘Cov’ is the coverage probability for the confidence interval. ‘ISSE(0)’ and ‘ISSE(1)’ are the treatment-specific ratios of integrated squared survival error, which corresponds to the mean-squared of the treatment-specific conditional survival function, relative to the Cox estimator. $\gamma(50)$ is the true restricted mean treatment effects for $\tau = 50$.

Table 1

Estimator	Percent Relative Bias	MSE Ratio	SD Ratio	Cov.	ISSE(0)	ISSE(1)
Cox	1%	1.00	1.00	0.96	1.00	1.00
Splines	1%	1.19	1.09	0.96	4.41	3.92
log-Normal	4%	0.94	0.95	0.96	1.53	1.16
Weibull	1%	0.91	0.95	0.95	0.88	0.89
RSF	21%	2.22	0.89	0.87	5.31	4.21
Stacked with RSF	3%	0.98	0.97	0.98	1.29	1.12
Stacked without RSF	1%	0.95	0.97	0.99	1.15	1.05
Cox	10%	1.00	1.00	0.93	1.00	1.00
Splines	4%	0.92	1.05	0.98	1.19	1.19
log-Normal	6%	0.79	0.95	0.95	1.11	1.07
Weibull	9%	0.92	0.97	0.94	0.95	0.95
RSF	-4%	0.59	0.83	0.98	2.24	2.08
Stacked with RSF	5%	0.81	0.98	0.98	0.77	0.80
Stacked without RSF	6%	0.79	0.94	0.97	0.75	0.77

Linear $\gamma(50) = -12.318$

Non-Linear $\gamma(50) = -10.334$

Simulation results for the gamma distributed scenarios: $N = 300$, $N_{SM} = 1000$, and a marginal censoring rate of 30% for the linear scenario and 30% for the non-linear scenario. The confidence intervals are estimated with the leave-one-out jackknife. ‘Percent Relative Bias’ is the ratio of the bias and true restricted mean difference. ‘MSE Ratio’ is the ratio of MSE relative to the Cox estimator. ‘SD Ratio’ is the ratio of average jackknife standard deviation relative to the Cox estimator. ‘Cov’ is the coverage probability for the confidence interval. ‘ISSE(0)’ and ‘ISSE(1)’ are the treatment-specific ratios of integrated squared survival error, which corresponds to the mean-squared of the treatment-specific conditional survival function, relative to the Cox estimator. $\gamma(50)$ is the true restricted mean treatment effects for $\tau = 50$.

Table 2

Estimator	Percent Relative Bias	MSE Ratio	SD Ratio	Cov.	ISSE(0)	ISSE(1)
Cox	-6%	1.00	1.00	0.94	1.00	1.00
Splines	-5%	1.28	1.15	0.95	4.21	4.17
log-Normal	-3%	1.01	1.02	0.95	1.16	1.22
Weibull	3%	0.80	0.91	0.94	0.85	0.89
RSF	-5%	0.97	0.99	0.98	3.75	5.41
<hr/>						
Stacked with RSF	0%	0.88	0.97	0.97	1.11	1.20
Stacked without RSF	0%	0.90	0.97	0.97	1.07	1.09
<hr/>						
Cox	-13%	1.00	1.00	0.92	1.00	1.00
Splines	-7%	1.00	1.08	0.96	1.66	1.34
log-Normal	-6%	0.78	0.96	0.94	0.88	0.98
Weibull	1%	0.64	0.89	0.95	0.92	0.94
RSF	-5%	0.72	0.93	0.98	1.51	1.98
<hr/>						
Stacked with RSF	-4%	0.73	0.94	0.97	0.86	0.83
Stacked without RSF	-4%	0.75	0.96	0.97	0.88	0.81

Non-Linear $\gamma(50) = -6.407$

Simulation results for the log-Normal distributed scenarios: $N = 300$, $N_{SYM} = 1000$, and a marginal censoring rate of 30% for the linear scenario and 30% for the non-linear scenario. The confidence intervals are estimated with the leave-one-out jackknife. 'Percent Relative Bias' is the ratio of the bias and true restricted mean difference. 'MSE Ratio' is the ratio of MSE relative to the Cox estimator. 'SD Ratio' is the ratio of average jackknife standard deviation relative to the Cox estimator. 'Cov' is the coverage probability for the confidence interval. 'ISSE(0)' and 'ISSE(1)' are the treatment-specific ratios of integrated squared survival error, which corresponds to the mean-squared of the treatment-specific conditional survival function, relative to the Cox estimator. $\gamma(50)$ is the true restricted mean treatment effects for $\tau = 50$.

Table 3

Estimator	Percent Relative Bias	MSE Ratio	SD Ratio	Cov.	ISSE(0)	ISSE(1)
Cox	-4%	1.00	1.00	0.94	1.00	1.00
Splines	-5%	1.18	1.07	0.96	3.05	3.50
log-Normal	1%	0.84	0.95	0.94	0.66	0.66
Weibull	0%	0.92	0.99	0.95	0.92	1.03
RSF	4%	0.68	0.81	0.97	3.10	5.06
<hr/>						
Stacked with RSF	1%	0.86	0.96	0.96	0.84	0.88
Stacked without RSF	1%	0.88	0.97	0.95	0.81	0.81
<hr/>						
Cox	-2%	1.00	1.00	0.96	1.00	1.00
Splines	3%	1.00	1.00	0.96	1.00	1.13
log-Normal	-9%	1.14	0.96	0.93	0.87	0.80
Weibull	-8%	1.21	1.02	0.93	0.97	0.89
RSF	-5%	0.82	0.88	0.98	3.19	3.27
<hr/>						
Stacked with RSF	-1%	1.20	1.10	0.96	0.71	0.70
Stacked without RSF	-2%	0.90	0.95	0.95	0.68	0.67

Non-Linear $\gamma(50) = -8.153$

Table 4

The estimated one-year and three-year restricted mean treatment effect (in days) with 95% confidence intervals (in parenthesis) and the stacking weights by center volume [ie., $\hat{\alpha}_k^{(a)}$] for each survival model in the lung transplant example. The stacked model includes random survival forests (RSF).

	Model	Restricted Mean Treatment Effect with 95% CI	Estimated Model Weights	
			Low Volume Centers	High Volume Centers
One Year	Cox	5.0 (-0.4, 10.4)	0.000	0.339
	Cox with Splines	6.0 (0.6, 11.3)	0.225	0.275
	log-Normal	8.0 (2.4, 13.6)	0.243	0.000
	Weibull	6.5 (1.2, 11.8)	0.000	0.000
	RSF	5.5 (-14.1, 25.0)	0.532	0.386
	Stacked	7.3 (-4.0, 18.7)	n/a	n/a
Three Year	Cox	21.3 (-3.4, 46.0)	0.000	0.339
	Cox with Splines	27.1 (2.2, 52.1)	0.225	0.275
	log-Normal	24.8 (1.2, 48.4)	0.243	0.000
	Weibull	23.9 (-0.3, 48.0)	0.000	0.000
	RSF	26.1 (-51.2, 102.1)	0.532	0.386
	Stacked	25.8 (-17.4, 68.3)	n/a	n/a