



HAL
open science

Pseudo-Observations and Super Learner for the Estimation of the Restricted Mean Survival Time

Ariane Cwiling, Vittorio Perduca, Olivier Bouaziz

► **To cite this version:**

Ariane Cwiling, Vittorio Perduca, Olivier Bouaziz. Pseudo-Observations and Super Learner for the Estimation of the Restricted Mean Survival Time. 2024. hal-04559520

HAL Id: hal-04559520

<https://hal.science/hal-04559520>

Preprint submitted on 25 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Pseudo-Observations and Super Learner for the Estimation of the Restricted Mean Survival Time

Ariane Cwiling, Vittorio Perduca, and Olivier Bouaziz

Université Paris Cité, CNRS, MAP5, F-75006 Paris, France

Abstract

In the context of right-censored data, we study the problem of predicting the restricted time to event based on a set of covariates. Under a quadratic loss, this problem is equivalent to estimating the conditional Restricted Mean Survival Time (RMST). To that aim, we propose a flexible and easy-to-use ensemble algorithm that combines pseudo-observations and super learner. The classical theoretical results of the super learner are extended to right-censored data, using a new definition of pseudo-observations, the so-called split pseudo-observations. Simulation studies indicate that the split pseudo-observations and the standard pseudo-observations are similar even for small sample sizes. The method is applied to maintenance and colon cancer datasets, showing the interest of the method in practice, as compared to other prediction methods. We complement the predictions obtained from our method with our RMST-adapted risk measure, prediction intervals and variable importance measures developed in a previous work.

Keywords: Right-censoring, RMST, prediction, stacking, pseudo-observations, super learner.

1 Introduction

Predicting the time to an event of interest based on a set of covariates is a relevant goal in applications. For example, in medical applications, it could be interesting to predict the time to onset of cancer, relapse or death of a patient, while in industrial applications, we might be interested in predicting the time to failure of a mechanical part. Due to tail issues caused by right-censoring, it is common to focus on the restricted time to event instead of the time to event itself (see Eaton et al., 2020). In practice, this prediction problem is equivalent to estimating the Restricted Mean Survival Time (RMST) when a quadratic loss is used. The RMST is a clinically meaningful quantity that has gained attention over the years for its simplicity and interpretability. The RMST can be easily retrieved by integrating an estimator of the survival function, yet new approaches have been developed to directly model it (see Andersen et al., 2004; Tian et al., 2014; Wang and Schaubel, 2018). Remarkably, these methods avoid strong modeling assumptions such as the proportional hazard assumption from the Cox model.

In particular, pseudo-observations, introduced by Andersen et al. (2004), have enabled the application of a large range of RMST estimation models. Pseudo-observations are a transformation of the incomplete observed times that have the strong following property (see Jacobsen and Martinussen, 2016): Their conditional expectation is equal, up to a remainder term, to the conditional RMST. As a result, they can be fed into any prediction model adapted to uncensored data, from generalized linear models (Andersen et al., 2004) to neural networks (Zhao, 2021).

In practical applications, the choice of the best learner among all of available machine learning algorithms for uncensored data is a difficult question, as the prediction performance of the different algorithms may depend on the application. As a result, an attractive alternative is to use ensemble methods, such as the super learner algorithm (see Van Der Laan et al., 2007). When dealing with uncensored data, it has been proven that, under mild assumptions, the learner with the lowest cross-validated risk in a user-defined library of algorithms, the so-called discrete super learner, performs as well as the best algorithm in the library up to an error term (Van Der Laan and Dudoit, 2003; Dudoit and Van Der Laan, 2005; Van Der Vaart et al., 2006). This result can be exploited to construct the best linear combination of all the learners in the library, the so-called continuous super learner. This stacking method has the theoretical guarantee to perform asymptotically at least as well as any of the candidate learners. The continuous super learner is simply termed “super learner” in what follows.

The goal of this work is to adapt the super learner algorithm for RMST estimation with right-censored data. The super learner for analyzing censored data has already been the subject of previous research. For instance Golmakani and Polley (2020) proposed a survival super learner that takes censored data as input and applies a library of survival algorithms. However, the models used in this super learner must verify the proportional hazard assumption, which is a strong limitation of the approach. Another example is the super learner with Inverse Probability Censoring Weight (IPCW) loss (see Van Der Laan and Dudoit, 2003; Keles et al., 2004; Gonzalez Ginestet et al., 2021; Devaux et al., 2022), which allows to use any survival model as long as the conditional independence assumption holds and the censoring distribution is consistently estimated. Finally, Sachs et al. (2019) and Gonzalez Ginestet et al. (2022) proposed two super learner methods with a pseudo-observations-based-AUC loss. However, to our knowledge, none of the pseudo-observations-based super learner methods are provided with theoretical convergence guarantees.

In this work, we present a novel approach combining the super learner with pseudo-observations for the estimation of the RMST. In order to derive theoretical results similar to those of Van Der Laan et al. (2007) in the uncensored case, we first present a new type of pseudo-observations, namely the *split pseudo-observations*. Those are introduced, along with the standard pseudo-observations, in Section 2. In Section 3, we present in details the super learner based on both standard and split pseudo-observations. Theoretical results for the combination of split pseudo-observations with the super learner are also derived. The performance of our method is studied extensively through simulations in Section 4 and on two real datasets in Section 5.

2 Pseudo-observations

In the context of right-censored data, we denote by T^* the time to the event, C the censoring time, $T = T^* \wedge C$ the observed time and $\delta = \mathbb{1}\{T^* \leq C\}$ the censoring indicator. An observation is then represented by the vector $O = (T, \delta, Z)$ where $Z \in \mathbb{R}^d$ is a covariate vector. We note $S(t | Z) = \mathbb{P}(T^* > t | Z)$ the survival function of T^* conditionally on the covariates Z . Let $\tau_H = \inf\{t > 0 : \mathbb{P}(T > t | Z) = 0 \text{ a.s.}\}$. The RMST is defined for a fixed time horizon $\tau < \tau_H$, conditionally on the covariates, as

$$\mathbb{E}[T^* \wedge \tau | Z] = \int_0^\tau S(t | Z) dt.$$

Given this definition, the RMST can be estimated for instance by integrating an estimator of the conditional survival function between 0 and τ , or by regressing the restricted event times on covariates. In the second case, censoring must be taken into account since the times T^* are

not observed for all individuals. This can be achieved by using pseudo-observations. Consider censored observations $D_n = \{O_i = (T_i, \delta_i, Z_i), i = 1, \dots, n\}$, where n is the sample size. Classical pseudo-observations are computed in the following way: For a given $\tau < \tau_H$ and all $i = 1, \dots, n$,

$$\Gamma_i := n \int_0^\tau \hat{S}(t) dt - (n-1) \int_0^\tau \hat{S}^{-i}(t) dt, \quad (1)$$

where \hat{S} is the Kaplan-Meier estimator of the survival function computed on all data and \hat{S}^{-i} is the same estimator computed on all data but the i -th observation. The interest of pseudo-observations for regression purposes lies in the following result by Jacobsen and Martinussen (2016):

$$\Gamma_i = \mathbb{E}[T^* \wedge \tau] + \int_0^\tau \dot{\phi}(O_i)(t) dt + \xi_n,$$

where $\xi_n = o_{\mathbb{P}}(1)$, ϕ is the first order influence function in the Von Mises expansion of the Kaplan-Meier estimator, and

$$\mathbb{E} \left[\int_0^\tau \dot{\phi}(O_i)(t) dt \mid Z_i \right] = \mathbb{E}[T^* \wedge \tau \mid Z_i] - \mathbb{E}[T^* \wedge \tau],$$

so that

$$\mathbb{E}[\Gamma_i \mid Z_i] = \mathbb{E}[T^* \wedge \tau \mid Z_i] + \mathbb{E}[\xi_n \mid Z_i]. \quad (2)$$

This result is valid under the following independent censoring assumption.

Assumption 1 (Independent censoring). *The censoring time C and the pair of variables (T^*, Z) are independent.*

Pseudo-observations are, by construction, correlated with each other, which makes it difficult to study their theoretical properties. To deal with this issue, we propose a new type of pseudo-observations, called *split pseudo-observations*. The idea is to split the data in two subsets D_{n_1} and D_{n_2} of size n_1 and $n_2 = n - n_1$, respectively. The former is used to compute the Kaplan-Meier estimator and the latter for the pseudo-observations. We then define a new type of pseudo-observations as follows:

$$\Gamma_i(D_{n_1}) = \Gamma_{O_i}(D_{n_1}) := (n_1 + 1) \int_0^\tau \hat{S}_{D_{n_1}}^{+i}(t) dt - n_1 \int_0^\tau \hat{S}_{D_{n_1}}(t) dt, \quad (3)$$

where O_i is an observation in D_{n_2} , $\hat{S}_{D_{n_1}}$ is the Kaplan-Meier estimator of the survival function computed on the n_1 data points in D_{n_1} and $\hat{S}_{D_{n_1}}^{+i}$ is the same estimator computed on the $n_1 + 1$ data points obtained by adding O_i to the sample D_{n_1} . The main advantage of this construction is that the new pseudo-observations constructed for all the observations in D_{n_2} are independent conditionally on D_{n_1} . A result similar to Equation (2) can then be easily derived for those split pseudo-observations. Under Assumption 1:

$$\mathbb{E}[\Gamma_i(D_{n_1}) \mid Z_i, D_{n_1}] = \mathbb{E}[T_i^* \wedge \tau \mid Z_i] + \mathbb{E}[\xi_{n_1} \mid Z_i, D_{n_1}], \quad (4)$$

where $\xi_{n_1} = o_{\mathbb{P}}(1)$. From a theoretical standpoint, we establish in Section 3 finite sample and asymptotic results for the super learner coupled with split pseudo-observations. In Section 4, we observe on simulated data that split and traditional pseudo-observations are very similar, and the choice between them has minimal impact on the prediction quality. Therefore, while split pseudo-observations are easier to study for theoretical results, split and traditional pseudo-observations can be used interchangeably in applications.

3 Super Learner

In this section we introduce a super learner algorithm for right-censored data based on pseudo-observations. Our aim is to predict the restricted time to event with a quadratic loss - the Mean Squared Error (MSE) - using the super learner. It is well known that, under this loss, the best prediction model is the conditional expectation of the restricted time. Therefore, our super learner algorithm can also be seen as an estimator of the RMST. We start by reintroducing the classical super learner for uncensored observations along with a fundamental theoretical result from Dudoit and Van Der Laan (2005) and we then derive a new result for our proposed approach.

We can formalize the problem in the following way. Let $\psi : \mathcal{M} \rightarrow \mathcal{D}(\mathcal{Z})$ be a parameter mapping a distribution from a statistical model \mathcal{M} to an element in the space $\mathcal{D}(\mathcal{Z})$ of real-valued functions defined on a d -dimensional Euclidean set $\mathcal{Z} \subset \mathbb{R}^d$. A realization $\psi = \psi(P)$ of ψ for a given $P \in \mathcal{M}$ belongs to the parameter space $\Psi := \{\psi(P) : P \in \mathcal{M}\} \subset \mathcal{D}(\mathcal{Z})$. Consider a loss function

$$L : (\psi, O) \mapsto L(\psi, O) \in \mathbb{R}, \text{ for } \psi \in \Psi, O \sim P.$$

The quantity of interest is the *risk* of the parameter ψ for a distribution P ,

$$\Theta(\psi, P) = \int L(\psi, o) dP(o).$$

Given this definition, the risk minimizer is defined as

$$\psi^* = \psi^*(P) = \arg \min_{\psi \in \Psi} \Theta(\psi, P) = \arg \min_{\psi \in \Psi} \int L(\psi, o) dP(o),$$

and characterizes the optimal risk

$$\theta^* = \Theta(\psi^*, P) = \min_{\psi \in \Psi} \Theta(\psi, P) = \min_{\psi \in \Psi} \int L(\psi, o) dP(o). \quad (5)$$

For instance, if $L(\psi, O) = (T^* \wedge \tau - \psi(Z))^2$ is the quadratic loss function and if $P(t, z) = \mathbb{P}(T^* \leq t, Z \leq z)$ is the joint law of (T^*, Z) , then the RMST minimizes this risk, i.e. $\psi^*(Z) = \mathbb{E}[T^* \wedge \tau \mid Z]$. In practice, the risk minimizer is typically unknown and we aim to approximate it as faithfully as possible. Suppose that a dataset of n i.i.d. observations $\{O_i, i = 1, \dots, n\}$, $O_i \sim P \in \mathcal{M}$, of empirical law P_n , is available to estimate a parameter $\psi(P)$. An estimator mapping $\hat{\psi}$ can be viewed simply as an algorithm one applies to data, i.e. to empirical distributions P_n . A realization of this mapping applied to a particular empirical distribution P_n is denoted $\hat{\psi}(P_n)$ and belongs to $\mathcal{D}(\mathcal{Z})$.

3.1 Uncensored data

For the sake of simplicity, suppose first that we observe the complete data $\{O_i^* = (T_i^*, Z_i), i = 1, \dots, n\}$, and we have access to K_n candidate estimators $\hat{\psi}_k, k = 1, \dots, K_n$, of ψ^* . We also call these estimators candidate learners. We then wish to select the best estimator among all the candidate learners, that is the one that minimizes the quadratic loss. Van Der Laan and Dudoit (2003), Dudoit and Van Der Laan (2005) and Van Der Vaart et al. (2006) theoretically validated the use of cross-validation to select an optimal learner among many candidates and protect against overfitting. Cross-validation starts with splitting the data into a training set and a validation set. Candidate estimators are then constructed on the training set and evaluated on the validation set. Formally, we divide the data according to an independent random vector

$B_n = (B_n(i) : i = 1, \dots, n) \in \{0, 1\}^n$ into a training set $\{O_i^* : i, B_n(i) = 0\}$ of size $n_0 = n - \lfloor np_n \rfloor$ and a validation set $\{O_i^* : i, B_n(i) = 1\}$ of size $n_1 = \lfloor np_n \rfloor$ for $p_n \in (0, 1)$. Several cross-validation schemes, i.e. distributions for B_n , can be used. In the following, we will focus on V -fold cross-validation, where data are divided into V subsets, or folds, of approximately same size. One by one, each fold serves as validation set while the remaining folds constitute the training set. The associated distribution of B_n assigns a mass of $1/V$ to each of the resulting V binary vectors.

We suppose that all observations O_i^* are i.i.d. of law $P^* \in \mathcal{M}$, and we denote $P_{B_n}^0$ and $P_{B_n}^1$ the empirical distributions of the training and validation sets respectively. The cross-validated risk estimator for the k -th candidate learner is then defined as

$$\begin{aligned}\widehat{\theta}_n(k) &= \mathbb{E}_{B_n} \Theta(\widehat{\psi}_k(P_{B_n}^0), P_{B_n}^1) = \mathbb{E}_{B_n} \int L(\widehat{\psi}_k(P_{B_n}^0), o) dP_{B_n}^1(o) \\ &= \mathbb{E}_{B_n} \frac{1}{n_1} \sum_{i: B_n(i)=1} L(\widehat{\psi}_k(P_{B_n}^0), O_i^*),\end{aligned}$$

and the cross-validated selector is denoted

$$\widehat{k} = \arg \min_{k \in \{1, \dots, K_n\}} \widehat{\theta}_n(k).$$

Optimality results are based on the comparison between the cross-validated selector and the selector which, for each given dataset, makes the best choice, knowing the true full data distribution. This cross-validated oracle selector minimizes the cross-validated conditional risk

$$\widetilde{\theta}_n(k) = \mathbb{E}_{B_n} \Theta(\widehat{\psi}_k(P_{B_n}^0), P^*) = \mathbb{E}_{B_n} \int L(\widehat{\psi}_k(P_{B_n}^0), o) dP^*(o), \quad (6)$$

and is denoted

$$\widetilde{k} = \arg \min_{k \in \{1, \dots, K_n\}} \widetilde{\theta}_n(k). \quad (7)$$

In Dudoit and Van Der Laan (2005), it is proven that the cross-validation selector performs asymptotically as well as the oracle cross-validation selector in terms of performance measure such as the MSE for regression or the AUC for classification tasks. The result for the MSE in the setting with uncensored data is stated in Theorem 1 below.

Theorem 1. *Let O_1^*, \dots, O_n^* be a random sample from a data generating distribution P^* , where each $O_i^* = (T_i^*, Z_i)$ consists of two components, a univariate outcome $T_i^* \in \mathbb{R}^+$ and a d -dimensional covariate vector $Z_i \in \mathbb{R}^d$. Let $\{\widehat{\psi}_k : k = 1, \dots, K_n\}$ denote a sequence of K_n candidate estimators for the RMST, $\psi^*(Z) = \mathbb{E}[T^* \wedge \tau \mid Z]$, which is the risk minimizer for the quadratic loss function $L(\psi, O^*) = (T^* \wedge \tau - \psi(Z))^2$. Suppose there exists M such that $\tau \leq M < \infty$ and*

$$\sup_{Z \in \mathcal{Z}, \psi \in \Psi} |\psi(Z)| \leq M \text{ almost surely.}$$

Finite sample result. *Let $M_1 = 8M^2$, $M_2 = 16M^2$ and $c(M, \gamma) = 2(1 + \gamma)^2(M_1/3 + M_2/\gamma)$.*

For all $\gamma > 0$,

$$0 \leq \mathbb{E}[\widetilde{\theta}_n(\widehat{k}) - \theta^*] \leq (1 + 2\gamma)\mathbb{E}[\widetilde{\theta}_n(\widetilde{k}) - \theta^*] + 2c(M, \gamma) \frac{1 + \log(K_n)}{np_n}.$$

Asymptotic results. *If*

$$\frac{\log(K_n)}{np_n \mathbb{E}[\widetilde{\theta}_n(\widetilde{k}) - \theta^*]} \rightarrow 0 \text{ as } n \rightarrow \infty,$$

then

$$\frac{\mathbb{E}[\tilde{\theta}_n(\hat{k}) - \theta^*]}{\mathbb{E}[\tilde{\theta}_n(\tilde{k}) - \theta^*]} \rightarrow 1 \text{ as } n \rightarrow \infty.$$

Similarly, if

$$\frac{\log(K_n)}{np_n(\tilde{\theta}_n(\tilde{k}) - \theta^*)} \rightarrow 0 \text{ in probability as } n \rightarrow \infty,$$

then

$$\frac{\tilde{\theta}_n(\hat{k}) - \theta^*}{\tilde{\theta}_n(\tilde{k}) - \theta^*} \rightarrow 1 \text{ in probability as } n \rightarrow \infty.$$

Theoretical guarantees for the discrete super learner algorithm, which selects the learner minimizing the cross-validated risk among all candidates, are outlined in Theorem 1. However, instead of simply selecting one candidate estimator in the library, it is also possible to consider the continuous super learner algorithm which fits a weighted combination of the candidate learners. During the V -fold cross-validation process, a new data matrix is created, where each row i consists in the set of predictions obtained from Z_i by every candidate learners, together with the true outcome $T_i^* \wedge \tau$. Outcomes are then regressed onto the predictions, using another algorithm chosen by the user. This puts weights on the candidate learners. Such a continuous super learner will perform asymptotically at least as well as the best candidate learner. This is a direct consequence of Theorem 1 by considering any combination of the candidate learners as a candidate learner itself, see Van Der Laan et al. (2007).

3.2 Right-censored data

3.2.1 Super learner on standard pseudo-observations




Figure 1: Diagram of the super learner based on standard pseudo-observations for right-censored data, see Equation (1). Pseudo-observations are computed once and for all at the beginning.

In order to take into account right-censored data, and motivated by the asymptotic result in Equation (2), we propose a new method that consists in feeding pseudo-observations directly in the super learner, using the quadratic loss applied to those pseudo-observations. Our algorithm is therefore identical to the super learner described in Van Der Laan et al. (2007), with an additional first step for computing pseudo-observations. A diagram illustrating the method is provided in Figure 1. The corresponding algorithm is detailed in Algorithm 1. It describes both the discrete and the continuous pseudo-observations-based super learners, which differ from step 5 onwards.

Algorithm 1 Pseudo-observations-based Super Learner

Input: Data $D_n = \{O_i = (T_i, \delta_i, Z_i) : i = 1, \dots, n\}$, number of folds $V \geq 2$, library of candidate algorithms for estimating the RMST based on pseudo-observations $\{\widehat{\psi}_k : k = 1, \dots, K_n\}$.

Output: A trained algorithm for the prediction of the restricted time to event.

1. Compute standard pseudo-observations for the whole dataset: $\{\Gamma_i : i = 1, \dots, n\}$ (see Equation (1)).
2. Split the dataset into V mutually exclusive blocks.
- for** $v \in \{1, \dots, V\}$ **do**
 3. Divide observations according to the vector B_n verifying, for all $i = 1, \dots, n$,

$$B_n(i) = \begin{cases} 1 & \text{if the } i\text{-th observation belongs to the } v\text{-th fold,} \\ 0 & \text{else,} \end{cases}$$

such that $\{(\Gamma_i, Z_i) : B_n(i) = 0\}$ is the training set, of empirical law $P_{B_n}^0$, and $\{(\Gamma_i, Z_i) : B_n(i) = 1\}$ is the validation set, of empirical law $P_{B_n}^1$. Train all candidate learners $\{\widehat{\psi}_k : k = 1, \dots, K_n\}$ on the training set, resulting in the trained predictors $\{\widehat{\psi}_k(P_{B_n}^0) : k = 1, \dots, K_n\}$.

4. Predict the restricted time to event for the data in the validation set with each trained candidate learner: $\{\widehat{\psi}_{k,i} = \widehat{\psi}_k(P_{B_n}^0)(Z_i) : k = 1, \dots, K_n, B_n(i) = 1\}$.

end for

Discrete super learner

5. Identify the discrete super learner that minimizes the cross-validated risk: $\widehat{k} = \arg \min_{k \in \{1, \dots, K_n\}} \frac{1}{n} \sum_{i=1}^n (\Gamma_i - \widehat{\psi}_{k,i})^2$.

6. Train the discrete super learner $\widehat{\psi}_{\widehat{k}}$ on the entire dataset $\{(\Gamma_i, Z_i) : i = 1, \dots, n\}$.

Continuous super learner

5. Regress the pseudo-observations $\{\Gamma_i : i = 1, \dots, n\}$ onto the predictions $\{\widehat{\psi}_{k,i} : k = 1, \dots, K_n, i = 1, \dots, n\}$ with a parametric regression model in order to assign weights to the candidate learners.

6. Train all candidate learners $\{\widehat{\psi}_k : k = 1, \dots, K_n\}$ on the entire dataset $\{(\Gamma_i, Z_i) : i = 1, \dots, n\}$.

7. Combine the candidate learners trained on the whole dataset at step 6 with the weights obtained at step 5 to form the continuous super learner.
-

3.2.2 Super learner on split pseudo-observations

We also propose a second algorithm based on split pseudo-observations. It requires to use a subset of the data, that we call the Kaplan-Meier (or KM) set, to compute the pseudo-observations in the validation set. A diagram illustrating this second method is provided in Figure 2. The algorithm is detailed in Algorithm 2, for both discrete and continuous super learners. The interest in this second algorithm lies in the conditional independence structure

of the split pseudo-observations which allows to derive theoretical results similar to Theorem 1. Thus, the theoretical results below are provided with regard to the algorithm based on split pseudo-observations.

Algorithm 2 Split Pseudo-observations-based Super Learner

Input: Data $D_n = \{O_i = (T_i, \delta_i, Z_i) : i = 1, \dots, n\}$, number of folds $V \geq 3$, library of candidate algorithms for estimating the RMST based on right-censored data $\{\widehat{\psi}_k : k = 1, \dots, K_n\}$.

Output: A trained algorithm for the prediction of the restricted time to event.

1. Split the dataset into V mutually exclusive blocks.

for $v \in \{1, \dots, V\}$ **do**

2. Divide observations according to the vector B_n verifying, for all $i = 1, \dots, n$,

$$B_n(i) = \begin{cases} 2 & \text{if the } i\text{-th observation belongs to the } v\text{-th fold,} \\ 1 & \text{if the } i\text{-th observation belongs to the } (v+1)\text{-th fold (or first fold if } v = V), \\ 0 & \text{else,} \end{cases}$$

such that $\{O_i : B_n(i) = 0\}$ is the training set, of empirical law $P_{B_n}^0$, $\{O_i : B_n(i) = 1\}$ is the KM set, of empirical law $P_{B_n}^1$, and $\{O_i : B_n(i) = 2\}$ is the validation set, of empirical law $P_{B_n}^2$. Compute the split pseudo-observations for the data in the validation set, based on the KM set: $\{\Gamma_i^S = \Gamma_i(P_{B_n}^1) : B_n(i) = 2\}$ (see Equation (3)).

3. Train all candidate learners $\{\widehat{\psi}_k : k = 1, \dots, K_n\}$ on the training set, resulting in the trained predictors $\{\widehat{\psi}_k(P_{B_n}^0) : k = 1, \dots, K_n\}$.

4. Predict the time to event for the data in the validation set with each trained candidate learner: $\{\widehat{\psi}_{k,i} = \widehat{\psi}_k(P_{B_n}^0)(Z_i) : k = 1, \dots, K_n, B_n(i) = 1\}$.

end for

Discrete super learner

5. Identify the discrete super learner that minimizes the cross-validated risk: $\widehat{k} = \arg \min_{k \in \{1, \dots, K_n\}} \frac{1}{n} \sum_{i=1}^n (\Gamma_i^S - \widehat{\psi}_{k,i})^2$.

6. Train the discrete super learner $\widehat{\psi}_{\widehat{k}}$ on the entire dataset $\{O_i : i = 1, \dots, n\}$.

Continuous super learner

5. Regress the split pseudo-observations $\{\Gamma_i^S : i = 1, \dots, n\}$ onto the predictions $\{\widehat{\psi}_{k,i} : k = 1, \dots, K_n, i = 1, \dots, n\}$ with a parametric regression model in order to assign weights to the candidate learners.

6. Train all candidate learners $\{\widehat{\psi}_k : k = 1, \dots, K_n\}$ on the entire dataset $\{O_i : i = 1, \dots, n\}$.

7. Combine the candidate learners trained on the whole dataset at step 6 with the weights obtained at step 5 to form the continuous super learner.

Formally, consider the censored observations $\{O_i = (T_i, \delta_i, Z_i), i = 1, \dots, n\}$, i.i.d. of joint law $P \in \mathcal{M}$, with $T_i \in \mathbb{R}^+$, $\delta_i \in \{0, 1\}$, $Z_i \in \mathbb{R}^d$. As previously, K_n candidate estimator mappings $\widehat{\psi}_k$, $k = 1, \dots, K_n$ are considered, among which we wish to select the best in terms of quadratic error, using cross-validation. Combining split pseudo-observations with cross-validation imposes to divide the data in three subsets instead of two. Observations are divided according to an independent random vector $B_n = (B_n(i) : i = 1, \dots, n) \in \{0, 1, 2\}^n$ into a first training set $\{O_i : i, B_n(i) = 0\}$ of size $n_0 = n - \lfloor np_{1,n} \rfloor - \lfloor np_{2,n} \rfloor$ for $p_{1,n}, p_{2,n}, p_{1,n} + p_{2,n} \in (0, 1)$, a second training set $\{O_i : i, B_n(i) = 1\}$ of size $n_1 = \lfloor np_{1,n} \rfloor$ and a validation set $\{O_i : i, B_n(i) = 2\}$ of size $n_2 = \lfloor np_{2,n} \rfloor$. The first training set is used to compute the candidate estimators of the RMST. The second training set is used to compute the Kaplan-Meier estimator which in turn is used for the computation of the pseudo-observations. We refer to this set as the Kaplan-Meier

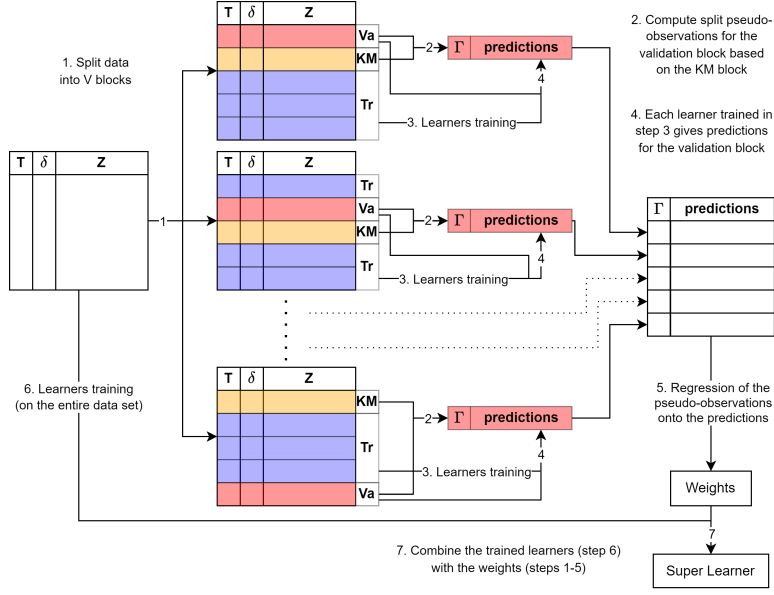


Figure 2: Diagram of the super learner based on split pseudo-observations for right-censored data, see Equation (3). Pseudo-observations are computed during the cross-validation step, for each validation block, based on an additional subset of the data (KM).

(KM) set. Pseudo-observations are computed for the data in the validation set. We denote $P_{B_n}^0$, $P_{B_n}^1$ and $P_{B_n}^2$ the empirical distributions of the three subsets. In this context, Equation (4) can be rewritten as

$$\mathbb{E}[\Gamma_O(P_{B_n}^1) | Z, P_{B_n}^1, B_n] = \mathbb{E}[T^* \wedge \tau | Z] + \mathbb{E}[\xi_{n_1} | Z, P_{B_n}^1, B_n], \quad (8)$$

where $\Gamma_O(P_{B_n}^1)$ is the split pseudo-observation constructed for the observation O based on the distribution $P_{B_n}^1$ of the KM set. It turns out that $\mathbb{E}[\xi_{n_1} | Z, P_{B_n}^1, B_n] = o_{\mathbb{P}}(1)$, so that the pseudo-observation plays a role of substitute to the true event time in the loss function. Consequently, the loss function also depends on the KM set used to compute the pseudo-observation,

$$L^{\text{po}} : (\psi, P_{B_n}^1, O) \mapsto L(\psi, (\Gamma_O(P_{B_n}^1), Z)) \in \mathbb{R}, \text{ for } \psi \in \Psi, P_{B_n}^1 \sim P^{\otimes n_1}, O \sim P.$$

As above, we can define the risk of the parameter ψ for distributions P and $P_{B_n}^1$,

$$\Theta^{\text{po}}(\psi, P_{B_n}^1, P) = \mathbb{E}_{B_n} \int L^{\text{po}}(\psi, P_{B_n}^1, o) dP(o).$$

The risk minimizer now also depends on the empirical distribution $P_{B_n}^1$,

$$\psi_1^* = \psi_1^*(P_{B_n}^1, P) = \arg \min_{\psi \in \Psi} \Theta^{\text{po}}(\psi, P_{B_n}^1, P) = \arg \min_{\psi \in \Psi} \mathbb{E}_{B_n} \int L^{\text{po}}(\psi, P_{B_n}^1, o) dP(o),$$

and the corresponding optimal risk can be written as

$$\theta_1^* = \Theta^{\text{po}}(\psi_1^*, P_{B_n}^1, P) = \min_{\psi \in \Psi} \Theta^{\text{po}}(\psi, P_{B_n}^1, P) = \min_{\psi \in \Psi} \mathbb{E}_{B_n} \int L^{\text{po}}(\psi, P_{B_n}^1, o) dP(o). \quad (9)$$

We can rewrite the cross-validated risk estimator,

$$\begin{aligned}\widehat{\theta}_n^{\text{p}o}(k) &= \Theta^{\text{p}o}(\widehat{\psi}_k(P_{B_n}^0), P_{B_n}^1, P_{B_n}^2) \\ &= \mathbb{E}_{B_n} \int L^{\text{p}o}(\widehat{\psi}_k(P_{B_n}^0), P_{B_n}^1, o) dP_{B_n}^2(o) \\ &= \mathbb{E}_{B_n} \frac{1}{n_2} \sum_{i: B_n(i)=2} L^{\text{p}o}(\widehat{\psi}_k(P_{B_n}^0), P_{B_n}^1, O_i),\end{aligned}$$

the cross-validated selector,

$$\widehat{k}^{\text{p}o} = \arg \min_{k \in \{1, \dots, K_n\}} \widehat{\theta}_n^{\text{p}o}(k), \quad (10)$$

the cross-validated conditional risk,

$$\widetilde{\theta}_n^{\text{p}o}(k) = \Theta^{\text{p}o}(\widehat{\psi}_k(P_{B_n}^0), P_{B_n}^1, P) = \mathbb{E}_{B_n} \int L^{\text{p}o}(\widehat{\psi}_k(P_{B_n}^0), P_{B_n}^1, o) dP(o),$$

and the cross-validated oracle selector

$$\widetilde{k}^{\text{p}o} = \arg \min_{k \in \{1, \dots, K_n\}} \widetilde{\theta}_n^{\text{p}o}(k).$$

We compare the risk differences $\widetilde{\theta}_n^{\text{p}o}(\widehat{k}^{\text{p}o}) - \theta_1^*$ and $\widetilde{\theta}_n^{\text{p}o}(\widetilde{k}^{\text{p}o}) - \theta_1^*$ to demonstrate optimality in a similar manner as in Theorem 1. This result is stated in Theorem 2 below.

Theorem 2. *Let O_1, \dots, O_n be a random sample from a data generating distribution P , where each $O_i = (T_i, \delta_i, Z_i)$ consists of three components, a univariate outcome $T_i \in \mathbb{R}^+$, a binary censoring indicator $\delta_i \in \{0, 1\}$ and a d -dimensional covariate vector $Z_i \in \mathbb{R}^d$. Let $\tau < \tau_H$. Let $\{\widehat{\psi}_k : k = 1, \dots, K_n\}$ denote a sequence of K_n candidate estimators for the RMST, $\mathbb{E}[T^* \wedge \tau | Z]$, which is the risk minimizer for the quadratic loss function $L^{\text{p}o}(\psi, P_{B_n}^1, O) = (\Gamma_O(P_{B_n}^1) - \psi(Z))^2$, up to an asymptotically negligible term (see Equation (8)). Suppose there exists M such that $\tau \leq M < \infty$, and*

$$|\Gamma_O(P_{B_n}^1)| \leq M \text{ and } \sup_{Z \in \mathcal{Z}, \psi \in \Psi} |\psi(Z)| \leq M \text{ almost surely.} \quad (11)$$

Suppose that Assumption 1 holds.

Finite sample result. *Let $M_1 = 8M^2$, $M_2 = 16M^2$ and $c(M, \gamma) = 2(1 + \gamma)^2(M_1/3 + M_2/\gamma)$. For all $\gamma > 0$,*

$$0 \leq \mathbb{E}[\widetilde{\theta}_n^{\text{p}o}(\widehat{k}^{\text{p}o}) - \theta_1^*] \leq (1 + 2\gamma)\mathbb{E}[\widetilde{\theta}_n^{\text{p}o}(\widetilde{k}^{\text{p}o}) - \theta_1^*] + 2c(M, \gamma) \frac{1 + \log(K_n)}{np_{2,n}}. \quad (12)$$

Asymptotic results. *If*

$$\frac{\log(K_n)}{np_{2,n}\mathbb{E}[\widetilde{\theta}_n^{\text{p}o}(\widetilde{k}^{\text{p}o}) - \theta_1^*]} \rightarrow 0 \text{ as } n \rightarrow \infty,$$

then

$$\frac{\mathbb{E}[\widetilde{\theta}_n^{\text{p}o}(\widehat{k}^{\text{p}o}) - \theta_1^*]}{\mathbb{E}[\widetilde{\theta}_n^{\text{p}o}(\widetilde{k}^{\text{p}o}) - \theta_1^*]} \rightarrow 1 \text{ as } n \rightarrow \infty. \quad (13)$$

Similarly, if

$$\frac{\log(K_n)}{np_{2,n}(\tilde{\theta}_n^{\text{po}}(\tilde{k}^{\text{po}}) - \theta_1^*)} \rightarrow 0 \text{ in probability as } n \rightarrow \infty,$$

then

$$\frac{\tilde{\theta}_n^{\text{po}}(\hat{k}^{\text{po}}) - \theta_1^*}{\tilde{\theta}_n^{\text{po}}(\tilde{k}^{\text{po}}) - \theta_1^*} \rightarrow 1 \text{ in probability as } n \rightarrow \infty. \quad (14)$$

With the additional result from Jacobsen and Martinussen (2016) in Equation (2), adapted to split pseudo-observations in Equation (8), we can also compare the risk differences $\tilde{\theta}_n(\hat{k}^{\text{po}}) - \theta^*$ and $\tilde{\theta}_n(\tilde{k}) - \theta^*$. In other words, we can compare the cross-validated selector on pseudo-observations (10) to the oracle selector on true event times (7), in terms of risks on true event times, see (5) and (6). The asymptotic result is given as a difference in Corollary 2.1 below. The result in the form of a ratio requires stronger assumptions and is provided in the Supplementary Material.

Corollary 2.1. *Same setup and assumptions as in Theorem 2.*

Finite sample result. *Let $M_1 = 8M^2$, $M_2 = 16M^2$ and $c(M, \gamma) = 2(1 + \gamma)^2(M_1/3 + M_2/\gamma)$. For all $k = 1, \dots, K_n$, let*

$$\begin{aligned} \phi_{n_1}(k) &= \mathbb{E}_{B_n} \left[\mathbb{E} \left[\mathbb{E}[\xi_{n_1} \mid Z, P_{B_n}^1, B_n]^2 \right. \right. \\ &\quad \left. \left. + 2(\psi^*(Z) - \hat{\psi}_k(P_{B_n}^0)(Z)) \mathbb{E}[\xi_{n_1} \mid Z, P_{B_n}^1, B_n] \mid P_{B_n}, B_n \right] \right]. \end{aligned}$$

Then, for all $\gamma > 0$,

$$\begin{aligned} 0 &\leq \mathbb{E}[\tilde{\theta}_n(\hat{k}^{\text{po}}) - \theta^*] \\ &\leq (1 + 2\gamma) \mathbb{E}[\tilde{\theta}_n(\tilde{k}^{\text{po}}) - \theta^*] + \mathbb{E}[(1 + 2\gamma)\phi_{n_1}(\tilde{k}^{\text{po}}) - \phi_{n_1}(\hat{k}^{\text{po}})] + 2c(M, \gamma) \frac{1 + \log(K_n)}{np_{2,n}}. \end{aligned}$$

Asymptotic results. *If*

$$\frac{\log(K_n)}{np_{2,n}} \rightarrow 0 \text{ as } n \rightarrow \infty,$$

then

$$\mathbb{E}[\tilde{\theta}_n(\hat{k}^{\text{po}}) - \tilde{\theta}_n(\tilde{k})] \rightarrow 0 \text{ as } n \rightarrow \infty, \quad (15)$$

and

$$\tilde{\theta}_n(\hat{k}^{\text{po}}) - \tilde{\theta}_n(\tilde{k}) \rightarrow 0 \text{ in probability as } n \rightarrow \infty. \quad (16)$$

4 Simulations

In this section, we conduct simulation studies to assess the validity of our approach on finite sample sizes. Our aim is threefold: first, we want to compare the values of the pseudo-observations constructed from the standard procedure and from our new procedure, the split pseudo-observations. Second, we want to compare the performance of our two proposed super learner algorithms, as described in Section 3.2, the one constructed from the standard pseudo-observations and the one constructed from the split pseudo-observations. Finally, we want to assess and compare the performances of our super learner algorithm based on standard pseudo-observations and two competitors: the Cox model and the Random Survival Forests (RSF)

model (see Ishwaran et al., 2008). The simulations are implemented from two different Cox based models with increasing complexity. The two models were originally introduced in Cwiling et al. (2023).

Scheme 1: Event times are simulated according to a Cox model with Weibull baseline hazard $\mathcal{W}(\nu, \kappa)$ and three covariates $Z = (Z^1, Z^2, Z^3)^\top$, where $Z^k \sim U[-a, a]$ for $k = 1, 2, 3$. Note that the survival function can be expressed as

$$S(t | Z) = \exp \left[- \left(\frac{t}{\kappa} \right)^\nu \exp(\beta^\top Z) \right],$$

with Cox regression parameters $\beta = (\beta_1, \beta_2, \beta_3)^\top$. Parameters are set to $\kappa = 2$, $\nu = 6$, $a = 5$, $(\beta_1, \beta_2, \beta_3) = (2, 1, 0)$. Censoring is simulated independently according to an exponential law with parameter $\lambda = 0.3$, leading to 47% censored data. The time horizon τ is chosen as the 90th percentile of the observed times T which corresponds to $\tau = 3.6$ in this setting.

Scheme 2: Event times are simulated according to a Cox model with Weibull baseline hazard $\mathcal{W}(\nu, \kappa)$, $\kappa = 2$, $\nu = 6$, this time with $\lambda(t | Z) = \lambda_0(t) \exp(g(Z))$ and

$$g(Z) = Z^3 - 3Z^5 + 2Z^1 Z^{10} + 4Z^2 Z^7 + 3Z^4 Z^5 - 5Z^6 Z^{10} \\ + 3Z^8 Z^9 + Z^1 Z^4 - 2Z^6 Z^9 - 4Z^3 Z^4 - Z^7 Z^8,$$

so that the survival function is expressed as

$$S(t | Z) = \exp \left[- \left(\frac{t}{\kappa} \right)^\nu \exp(g(Z)) \right].$$

Let $Z = (Z^1, \dots, Z^{15})^\top$. We simulate the covariates $Z = (Z^1, \dots, Z^{15})^\top$ with $Z^j \sim \mathcal{B}(0.4)$ for $j \in \{2, 4, 6, 9, 11, 12\}$ and $Z^j \sim U[0, 1]$, for $j \in \{1, 3, 5, 7, 8, 10, 13, 14, 15\}$. As a result, only the first 10 covariates are associated with the event times. However, the other 5 covariates, that are independent from the event times, will still be included in our regression models when analyzing those data. The censoring distribution is the same as in scheme 1, leading to 47% censored data. The time horizon τ is chosen as the 90th percentile of the observed times which corresponds to $\tau = 2.8$ in this setting.

4.1 Standard versus split pseudo-observations

In this section, we compare the values of the standard pseudo-observations and the split pseudo-observations constructed as described in Section 2. We construct samples of different sizes $n \in \{20, 50, 100\}$ according to the simulation scheme 1 and we vary the proportion of the two samples D_{n_1} and D_{n_2} used for the construction of the split pseudo-observations by setting $n_2 = \rho n$ with $\rho \in \{0.2, 0.5, 0.8\}$. The process is repeated 50 times. For each split pseudo-observation, a corresponding standard pseudo-observation can also be computed and compared. For clarity, 200 pseudo-observations are randomly drawn for each scenario and their scatter plots are displayed in Figure 3. Overall, we observe a very close agreement between the values of the standard and split pseudo-observations. As the sample size increases, the similarity between the two methods increases even more. A larger discrepancy occurs when n_2 is larger. This is a consequence of the sample size n_1 used for the Kaplan-Meier estimator in the construction of the split pseudo-observations which is smaller when n_2 is larger. However, even in the worst scenario ($\rho = 0.8$ and $n = 20$), the discrepancy between the two methods remains moderate.

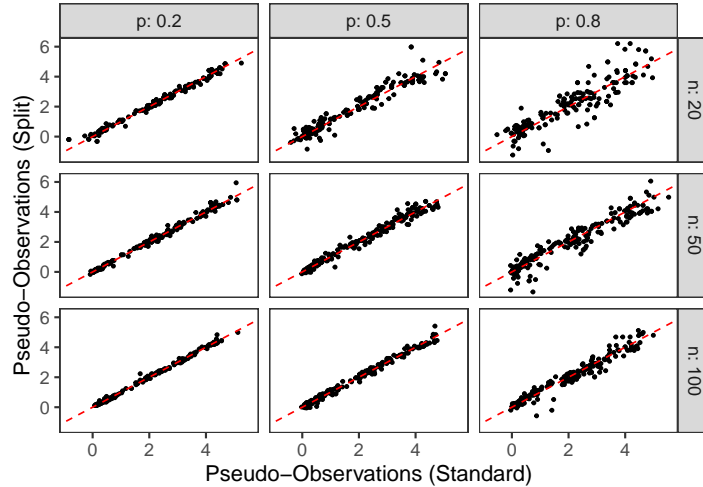




Figure 3: Comparison of standard and split pseudo-observations on data simulated with the first simulation scheme, with varying sample sizes n and $n_2 = \rho n$ for the split pseudo-observations.

4.2 Super learner implemented with standard versus split pseudo-observations

In this section, we want to compare the implementation of the two super learner algorithms, as introduced in Section 3.2: The first one is based on standard pseudo-observations and the second one is based on the split pseudo-observations. While the two types of pseudo-observations are very similar (as illustrated in the previous section), the two algorithms are still different in their implementations. The first one applies candidate learners directly on the pseudo-observations, treating those observations as uncensored data. On the contrary, in the second approach, the candidate learners are applied to the censored observations and require algorithms that can handle right-censored observations. In order to be able to compare the two approaches, we only consider learning algorithms based on pseudo-observations. More precisely, for the second approach, pseudo-observations are computed from the training set before applying the learning algorithms. The algorithms considered for the two methods are: the Linear Model (LM), the Lasso, the Generalized Additive Model (GAM), the Random Forests (RF) and the Neural Network (NN). For both methods, once the weights attributed to the candidate learners are computed, the algorithms are trained one last time on the whole dataset and combined using those weights, see Figure 1 and Figure 2. In our simulation setting, this last step is identical for both methods, i.e. standard pseudo-observations are computed on the whole dataset and used to train LM, Lasso, GAM, RF and NN. The only difference between the two approaches results in the computation of the weights in the construction of the super learner. The performances of the two methods are further compared with the oracle method that applies the same algorithms directly to the unobserved true event times. All super learners are trained using 6-fold cross-validation. The results are displayed in Figure 4, for the three different methods in the two simulation schemes with increasing sample sizes $n \in \{100, 200, 300, 400, 500\}$. The MSEs of the super learner algorithms and of every candidate learners are computed on an independent test set of size 1000 and the whole procedure is repeated 80 times for each sample size. The median and the first and third quantiles of the MSEs are reported in the figure. Exact values can be found in Table 1 in the Supplementary Material.



(a) Simulation scheme 1



(b) Simulation scheme 2

Figure 4: Mean Squared Errors (MSE) for the restricted time prediction with algorithms implemented on the true uncensored observations (first column) and the pseudo-observations (second column). The considered algorithms are: the linear model (LM), the Lasso, the generalized additive model (GAM), the random forests (RF), the neural network (NN) and the super learner (SL) combining all these algorithms. In the second column, standard and split pseudo-observations-based super learners are implemented (SL and SLsplit). Their curves overlap almost entirely. The sample size of the training set ranges from $n = 100$ to $n = 500$ and two simulation schemes are considered in the top and bottom rows. The MSEs of all the algorithms are computed on an independent test set of size 1000 and the whole procedure is repeated 80 times. Values of the median (cross), first and third quantiles (dashed lines) of the MSEs are reported.

We first observe that the two pseudo-observations-based super learner algorithms share similar performances for all sample sizes, regardless of the simulation scheme. This indicates that the two algorithms are almost identical in practice. Also, their performance is only slightly worse than the oracle algorithm. For the three algorithms, the super learner outperforms all individual learners. This is an illustration of Theorem 1 in Van Der Laan et al. (2007) for the oracle algorithm and of Corollary 2.1 for the split pseudo-observations algorithm. While we were not able to derive a similar theoretical result for the super learner based on standard pseudo-observations, the simulations clearly support this conjecture. Finally, the distribution of the weights assigned to each learner can be found in the Supplementary Material in Figure 1. We observe in the simulation scheme 1 that the largest weight is assigned to the NN algorithm with an increasing value of the weight when the sample size increases. For instance, for $n = 500$, the median weight allocated to the NN algorithm is equal to 0.906, 0.762 and 0.727 for the oracle algorithm, the super learner combined with standard pseudo-observations and the super learner combined with split pseudo-observations, respectively. This is in accordance with the results in Figure 4 where the NN algorithm outperforms all other individual algorithms and is only slightly less performing than the super learner. In the simulation scheme 2, the weights are more equally shared among all learners. The RF tend to have the largest weights with an increasing value when the sample size increases. For example, when $n = 500$, the median weights for the RF, LM and NN are equal to 0.525, 0.232, 0.172 respectively for the oracle algorithm, equal to 0.467, 0.345, 0.071 respectively for the super learner combined with standard pseudo-observations and equal to 0.462, 0.370, 0.014 respectively for the super learner combined with split pseudo-observations.

Since the two implementations of the super learner are almost identical, we only study the super learner based on standard pseudo-observations in the next section. This algorithm has the advantage to have a lower computational complexity and to allocate more data to the training set.

4.3 Prediction performance of the super learner and comparison with other methods

In this section, we compare our super learner algorithm combined with the standard pseudo-observations to the standard Cox model (without including interactions) and the Random Survival Forests (RSF) algorithm (see Ishwaran et al., 2008). For the super learner, we employ the same candidate learners as in the previous section and still use 6-fold cross-validation. For the Cox and RSF algorithms, the RMST is obtained by integrating the estimated survival curve between 0 and τ . We simulated data of increasing size $n \in \{200, 500, 1000, 1500, 2000\}$ in the two simulation schemes. As previously, the MSEs of all prediction models are computed on an independent test set of size 1000 and the procedure is repeated 80 times for each sample size. The median, first and third quantiles of the MSEs are reported in Figure 5.

The first simulation scheme consists in simulating time to events according to a standard Cox model without interactions between covariates. Hence, it is not surprising that Figure 5a shows a better performance for the Cox model compared to the other methods. However, the RSF and our method perform only slightly worse than the Cox model, with an increasing performance as the sample size increases. For small sample sizes, the super learner is clearly more performing than the RSF. On the other hand, the second simulation scheme is a Cox model with complex interactions between the covariates. In the results presented in Figure 5b, our method outperforms the Cox model for $n \geq 500$ and also exhibits better performances than the RSF for all sample sizes.




Figure 5: Mean Squared Errors (MSE) for the restricted time prediction with three different methods: the Cox model (with no interactions between covariates), the RSF and the super learner combined with standard pseudo-observations. The algorithms are trained on samples of size $n \in \{200, 500, 1000, 1500, 2000\}$ and two simulation schemes are considered in the left and right columns. The MSEs are computed on an independent test set of size 1000 and the whole procedure is repeated 80 times. Values of the median (cross), first and third quartiles (dashed lines) of the MSEs are reported.

5 Applications on real data

5.1 The maintenance dataset

We first study a maintenance dataset publicly available in the `PySurvival` Python package. It gathers data on 1000 machines, for which the number of weeks in activity has been recorded. The aim is to study the failure time of the machines. Overall 40% of the collected times are right-censored, and no failures occurred before 60 weeks of operation. Observed times range from 1 to 93 weeks, and we set the time horizon to the 90th quantile of the observed times which corresponds to $\tau = 88$. The dataset contains 5 predictors: three continuous covariates, pressure, moisture, temperature and two categorical variables, the team using the machine (three levels) and the manufacturer (four levels).

We first investigate the possible dependence between censoring and covariates. To that aim we fit a Cox model and a RSF to the data using the censoring time as the outcome. Both methods indicate no evidence of dependence: the global p-value for the Cox model is equal to 0.604 while the variable importance measure per permutation in the RSF does not exceed 0.005 for any of the variables (see Table 2 in Supplementary Material).

Next, we aim at predicting the time to failure using several models. We were not able to implement the Cox model from the `coxph` function because the Newton-Raphson algorithm ran

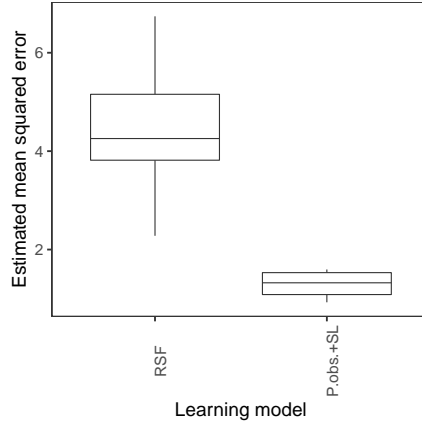


Figure 6: Estimation of the MSE on the `maintenance` dataset using the WRSS based on 10-fold cross-validation, for the RSF and the super learner based on pseudo-observations.

out of iterations without converging. We use instead the RSF and the super learner based on pseudo-observations. For the latter, cross-validation is fixed to 6 folds and the library was composed of LM, Lasso, GAM and RF. Both learning models are compared and analyzed using the methods developed in Cwiling et al. (2023).

First, the MSE is evaluated with the weighted residual sum of squares (WRSS), an inverse probability censoring weights (IPCW) estimator derived from the Brier score. Since our previous analysis supports the independent censoring assumption, we simply use the Kaplan-Meier estimator to compute the censoring weights. Results obtained on 10-fold cross-validation are displayed in Figure 6, indicating reasonable prediction errors in both cases considering the time range, with better performances for our method.

Next, a prediction interval for the true restricted time to failure of level 90% is computed for each item. These intervals are computed with the IPCW Rank-One-Out Split Conformal algorithm. This method relies on the distribution of the weighted residuals where, again, the Kaplan-Meier estimator is used for the censoring weights. The median length of the prediction intervals equals 7.217 and 6.789 for the RSF and the super learner combined with pseudo-observations, respectively. An illustration of these intervals, for some machines chosen randomly in the dataset, is provided in the Supplementary Material in Figure 2.

Finally, a global variable importance test is implemented based on the Leave-One-Covariate-Out conformal methodology. This test indicates whether removing a variable from the training set significantly deteriorates the prediction quality. This test is conducted by splitting the data into a training set (used for the implementation of the RMST estimator) and a validation set (used for the implementation of the test itself). Since this procedure is sensitive to the split, we conduct a multi-splitting procedure to stabilize the results along with a method for controlling the familywise rejection rate (see DiCiccio et al., 2020, for more details on this method). Consequently, the test is implemented at the 10% level on $M = 40$ different splits of the data and the M p-values are aggregated by computing twice the median value. This ensures to obtain a 10% level of the overall test. The results are displayed in Table 1. They show clear evidence of the influence of the team and provider variables on the failure time, while pressure does not alter the prediction quality and moisture seems to only have very limited impact in the RSF. However, both models do not reach the same conclusion for temperature, which seems to have a strong predictive power in the RSF but not in the super learner.

Variable	P-value RSF		P-value P.obs.+SL	
Pressure	1		1	
Moisture	0.257		1	
Temperature	0.003	**	1	
Team	0	***	0	***
Provider	0	***	0	***

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Table 1: Variable importance of the `maintenance` dataset for the RSF and the super learner based on pseudo-observations using a multi-splitting procedure with $M = 40$ splits. The global variable importance test is performed on each split and the M p-values are aggregated by computing twice the median value in order to obtain a 10% overall test.

5.2 The colon cancer dataset




Figure 7: Estimation of the MSE on the `colon` dataset using the WRSS based on 10-fold cross-validation, for the Cox model, the RSF and the super learner based on pseudo-observations.

We study now a dataset from a clinical trial on the effect of chemotherapy for colon cancer. The `colon` data are available in the R `survival` package. The complete case analysis contains 888 patients whose follow-up times range from 8 days to 3329 days (9.12 years). The goal of the study is to analyze the progression free survival (PFS) defined as the time elapsed from randomization to the minimum between recurrence and death. For this endpoint, 46% of the observations are censored, that is neither recurrence or death is observed. We set the time horizon to the 90th quantile of the observed times which corresponds to $\tau = 2672$ days. We use as predictors the type of treatment administrated (three levels), sex (binary), age (in years), obstruction indicator of the colon by the tumor (binary), whether the colon was perforated or not (binary), whether or not it adhered to nearby organs (binary), number of lymph nodes with detectable cancer (integer value that ranges from 0 to 33), level of differentiation of the tumor (three levels), extent of local spread (four levels), and whether the time from surgery to registration was short or long (binary).

We first investigate the possible dependence between censoring and covariates. As for the previous dataset, we fit a Cox model and a RSF to the data using the censoring time as the outcome. Both methods indicate no evidence of dependence: the global p-value for the Cox model is equal to 0.383 while the variable importance measure per permutation in the RSF does

not exceed 0.015 for any of the variables (see Table 3 in Supplementary Material).

Next, we predict the time to progression with the Cox model, the RSF and the pseudo-observations-based super learner. For the latter, 6-fold cross-validation is implemented with the library composed of LM, Lasso, GAM and RF. As previously, the performances of the learning models are analyzed using the methods developed in Cwiling et al. (2023) and the censoring distribution is estimated using the Kaplan-Meier estimator for all the methods involving IPCW.

In terms of MSE, all three models seem to have similar performances (see Figure 7). The prediction error is particularly large, emphasizing the difficulty for all the learners to provide accurate individual predictions. This is in accordance with the length of the 90% prediction intervals, displayed in Figure 3 in Supplementary Material, whose median equals 2961 for the Cox model, 3098 for the RSF and 3050 for the super learner combined with pseudo-observations. Finally, the variable importance test is implemented based on the multi-split approach with 40 different splits and overall rejection rate equal to 10%. The results are presented in Table 2. We observe that the importance of the type of treatment and extent of local spread variables are highly significant for the Cox and super learner algorithms while moderately significant for the RSF (p-value equals 0.084 and 0.063, respectively). The importance of all the other variables is far from significance for all algorithms.

Variable	P-value Cox	P-value RSF	P-value P.obs.+SL
Treatment	0 ***	0.084 .	0.002 **
Sex	0.981	0.724	0.622
Age	1	1	1
Obstruction	1	1	1
Perforation	1	1	1
Adherence	1	1	1
Nodes	1	0.637	0.662
Differentiation	1	1	1
Spread	0 ***	0.063 .	0.002 **
Surgery	1	1	1

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Table 2: Variable importance of the colon dataset for the Cox model, the RSF and the super learner based on pseudo-observations using a multi-splitting procedure with $M = 40$ splits. The global variable importance test is performed on each split and the M p-values are aggregated by computing twice the median value in order to obtain a 10% overall test.

6 Conclusion

Predicting the restricted time to event is of great interest in many applications. When using a quadratic loss function, this problem amounts to estimating the RMST. In this work we focused on the super learner combined with pseudo-observations to achieve this goal. The algorithm is simple to implement as it only requires an extra initial step for constructing the pseudo-observations. From there, the standard super learner is simply applied to the pseudo-observations and allows to get access to all the prediction algorithms that are classically used when dealing with uncensored observations. It also enjoys the nice property of the super learner: It will automatically select the best learner among all the candidates (for the discrete super learner) or it will provide the best combination of the candidates (for the continuous super learner). This is an important feature in applications since the best candidate can vary from one

study to another or when the sample size increases. Those results were proved in our context of right-censored data, when the split pseudo-observations are used. We also observed this property on simulated data and we empirically showed that combining the super learner with split or standard pseudo-observations provides very similar results. This is why we recommend to combine the super learner with standard pseudo-observations in practice, as it provides a simpler method and it allows to allocate more data to the training of candidate learners.

Interestingly, the use of split pseudo-observations might go beyond the scope of the super learner when studying theoretical properties of an estimator based on pseudo-observations. Indeed, they are almost identical to standard pseudo-observations, yet they do not suffer from dependency problems which makes them much easier to study in theoretical problems.

Our method relies on the independent censoring assumption inherent to pseudo-observations. When dealing with categorical covariates it is still possible to use the approach of Andersen and Pohar Perme (2010) (see Section 2.2) based on a mixture of Kaplan-Meier estimators to relax this assumption. Finally, our method could easily be extended to other situations such as recurrent events, competing risks, or left truncation by using a different definition of pseudo-observations adapted to these types of data (see Binder et al., 2014; Grand et al., 2019; Furberg et al., 2023). This is left to future research work.

In practice, several methods exist to implement pseudo-observations-based super learners. It would be of interest to provide some guidelines to practitioners, for instance on the choice between our approach and the ones based on IPCW or AUC losses. The main novelty of our work, as compared to previous methods, relies on theoretical results about the pseudo-observations-based super learner. Empirical studies comparing the different methods are warranted.

Acknowledgments

We thank Professor Antoine Chambaz for his helpful suggestions regarding the application of the super learner to pseudo-observations. His insights greatly improved the quality of our work.

7 Technical Results

Lemma 1. *If we consider the quadratic loss function $L^{\text{po}}(\psi, P_{B_n}^1, O) = (\Gamma_O(P_{B_n}^1) - \psi(Z))^2$, with O independent from $P_{B_n}^1$ and B_n , then the risk minimizer is $\psi_1^*(Z) = \mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n]$.*

Proof. We have

$$\begin{aligned}
& \mathbb{E} \left[L^{\text{po}}(\psi, P_{B_n}^1, O) \mid P_{B_n}^1, B_n \right] \\
&= \mathbb{E} \left[\left(\Gamma_O(P_{B_n}^1) - \psi(Z) \right)^2 \mid P_{B_n}^1, B_n \right] \\
&= \mathbb{E} \left[\left(\Gamma_O(P_{B_n}^1) - \mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] + \mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] - \psi(Z) \right)^2 \mid P_{B_n}^1, B_n \right] \\
&= \mathbb{E} \left[\left(\Gamma_O(P_{B_n}^1) - \mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] \right)^2 \mid P_{B_n}^1, B_n \right] \\
&\quad + \mathbb{E} \left[\left(\mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] - \psi(Z) \right)^2 \mid P_{B_n}^1, B_n \right] \\
&\quad + 2\mathbb{E} \left[\left(\Gamma_O(P_{B_n}^1) - \mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] \right) \left(\mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] - \psi(Z) \right) \mid P_{B_n}^1, B_n \right]
\end{aligned}$$

and

$$\begin{aligned}
& \mathbb{E} \left[\left(\Gamma_O(P_{B_n}^1) - \mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] \right) \left(\mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] - \psi(Z) \right) \mid P_{B_n}^1, B_n \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[\left(\Gamma_O(P_{B_n}^1) - \mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] \right) \right. \right. \\
&\quad \left. \left. \left(\mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] - \psi(Z) \right) \mid Z, P_{B_n}^1, B_n \right] \mid P_{B_n}^1, B_n \right] \\
&= \mathbb{E} \left[\left(\mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] - \mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] \right) \right. \\
&\quad \left. \left(\mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] - \psi(Z) \right) \mid P_{B_n}^1, B_n \right] \\
&= 0.
\end{aligned}$$

Hence

$$\begin{aligned}
& \arg \min_{\psi \in \Psi} \mathbb{E}_{B_n} \mathbb{E} [L^{\text{Po}}(\psi, P_{B_n}^1, O) \mid P_{B_n}^1, B_n] \\
&= \arg \min_{\psi \in \Psi} \mathbb{E}_{B_n} \mathbb{E} \left[\left(\mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] - \psi(Z) \right)^2 \mid P_{B_n}^1, B_n \right].
\end{aligned}$$

□

We prove that Lemma 3 in Dudoit and Van Der Laan (2005) is still valid with the pseudo-observation quadratic loss function.

Lemma 2. *Same setup and assumptions as in Theorem 2. Conditional on the training sets empirical distributions $P_{B_n}^0$, $P_{B_n}^1$ and split vector B_n , define, for any $k = 1, \dots, K_n$, the random variable $X_k = L^{\text{Po}}(\hat{\psi}_k(P_{B_n}^0), P_{B_n}^1, O) - L^{\text{Po}}(\psi_1^*, P_{B_n}^1, O)$, where O is independent from $P_{B_n}^0$, $P_{B_n}^1$, B_n , and L^{Po} denotes the quadratic loss function $L^{\text{Po}}(\psi, P_{B_n}^1, O) = (\Gamma_O(P_{B_n}^1) - \psi(Z))^2$. Then*

$$\text{Var}(X_k \mid P_{B_n}^0, P_{B_n}^1, B_n) \leq M_2 \mathbb{E}[X_k \mid P_{B_n}^0, P_{B_n}^1, B_n].$$

Proof. For the quadratic loss function, $\psi_1^*(Z) = \mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n]$ from Lemma 1. As

$$X_k = (\psi_1^*(Z) - \hat{\psi}_k(P_{B_n}^0)(Z)) (2\Gamma_O(P_{B_n}^1) - \hat{\psi}_k(P_{B_n}^0)(Z) - \psi_1^*(Z)),$$

we have

$$\begin{aligned}
& \mathbb{E}[X_k \mid P_{B_n}^0, P_{B_n}^1, B_n] \\
&= \mathbb{E} \left[\mathbb{E}[X_k \mid Z, P_{B_n}^0, P_{B_n}^1, B_n] \mid P_{B_n}^0, P_{B_n}^1, B_n \right] \\
&= \mathbb{E} \left[\left(\psi_1^*(Z) - \hat{\psi}_k(P_{B_n}^0)(Z) \right) \right. \\
&\quad \left. \left(2\mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] - \hat{\psi}_k(P_{B_n}^0)(Z) - \psi_1^*(Z) \right) \mid P_{B_n}^0, P_{B_n}^1, B_n \right] \\
&= \mathbb{E} \left[(\psi_1^*(Z) - \hat{\psi}_k(P_{B_n}^0)(Z))^2 \mid P_{B_n}^0, P_{B_n}^1, B_n \right].
\end{aligned}$$

As $|2\Gamma_O(P_{B_n}^1) - \widehat{\psi}_k(P_{B_n}^0)(Z) - \psi_1^*(Z)| \leq 4M$ by assumption (11), and with $M_2 = (4M)^2$, we find

$$\begin{aligned} \text{Var}(X_k | P_{B_n}^0, P_{B_n}^1, B_n) &\leq \mathbb{E}[X_k^2 | P_{B_n}^0, P_{B_n}^1, B_n] \\ &\leq (4M)^2 \mathbb{E}\left[\left(\psi_1^*(Z) - \widehat{\psi}_k(P_{B_n}^0)(Z)\right)^2 | P_{B_n}^0, P_{B_n}^1, B_n\right] \\ &= M_2 \mathbb{E}[X_k | P_{B_n}^0, P_{B_n}^1, B_n]. \end{aligned}$$

□

7.1 Proof of Theorem 2

We write

$$\begin{aligned} 0 &\leq \widetilde{\theta}_n^{\text{po}}(\widehat{k}^{\text{po}}) - \theta_1^* \\ &= \mathbb{E}_{B_n} \int \left(L^{\text{po}}(\widehat{\psi}_{\widehat{k}^{\text{po}}}(P_{B_n}^0), P_{B_n}^1, o) - L^{\text{po}}(\psi_1^*, P_{B_n}^1, o) \right) dP(o) \\ &\leq \mathbb{E}_{B_n} \int \left(L^{\text{po}}(\widehat{\psi}_{\widetilde{k}^{\text{po}}}(P_{B_n}^0), P_{B_n}^1, o) - L^{\text{po}}(\psi_1^*, P_{B_n}^1, o) \right) dP(o) \\ &\quad - (1 + \gamma) \mathbb{E}_{B_n} \int \left(L^{\text{po}}(\widehat{\psi}_{\widetilde{k}^{\text{po}}}(P_{B_n}^0), P_{B_n}^1, o) - L^{\text{po}}(\psi_1^*, P_{B_n}^1, o) \right) dP_{B_n}^2(o) \\ &\quad + (1 + \gamma) \mathbb{E}_{B_n} \int \left(L^{\text{po}}(\widehat{\psi}_{\widetilde{k}^{\text{po}}}(P_{B_n}^0), P_{B_n}^1, o) - L^{\text{po}}(\psi_1^*, P_{B_n}^1, o) \right) dP_{B_n}^2(o) \\ &\quad - (1 + 2\gamma) \mathbb{E}_{B_n} \int \left(L^{\text{po}}(\widehat{\psi}_{\widetilde{k}^{\text{po}}}(P_{B_n}^0), P_{B_n}^1, o) - L^{\text{po}}(\psi_1^*, P_{B_n}^1, o) \right) dP(o) \\ &\quad + (1 + 2\gamma) \mathbb{E}_{B_n} \int \left(L^{\text{po}}(\widehat{\psi}_{\widetilde{k}^{\text{po}}}(P_{B_n}^0), P_{B_n}^1, o) - L^{\text{po}}(\psi_1^*, P_{B_n}^1, o) \right) dP(o) \end{aligned}$$

where the first inequality follows by definition of the optimal risk θ_1^* (9) and the second by definition of \widehat{k}^{po} (10), which implies $\widetilde{\theta}_n^{\text{po}}(\widehat{k}^{\text{po}}) \leq \widetilde{\theta}_n^{\text{po}}(\widetilde{k}^{\text{po}})$. Denote the first two terms in the last expression by $R_{\widehat{k}^{\text{po}}, n}$ and the third and fourth terms by $T_{\widetilde{k}^{\text{po}}, n}$; the last term is the benchmark risk difference $(1 + 2\gamma)(\widetilde{\theta}_n^{\text{po}}(\widetilde{k}^{\text{po}}) - \theta_1^*)$. Hence,

$$0 \leq \widetilde{\theta}_n^{\text{po}}(\widehat{k}^{\text{po}}) - \theta_1^* \leq (1 + 2\gamma)(\widetilde{\theta}_n^{\text{po}}(\widetilde{k}^{\text{po}}) - \theta_1^*) + R_{\widehat{k}^{\text{po}}, n} + T_{\widetilde{k}^{\text{po}}, n}. \quad (17)$$

Let

$$\begin{aligned} \widetilde{H}_k &= \int \left(L^{\text{po}}(\widehat{\psi}_k(P_{B_n}^0), P_{B_n}^1, o) - L^{\text{po}}(\psi_1^*, P_{B_n}^1, o) \right) dP(o) \\ \widehat{H}_k &= \int \left(L^{\text{po}}(\widehat{\psi}_k(P_{B_n}^0), P_{B_n}^1, o) - L^{\text{po}}(\psi_1^*, P_{B_n}^1, o) \right) dP_{B_n}^2(o) \end{aligned}$$

and

$$\begin{aligned} R_{k, n}(B_n) &= (1 + \gamma)(\widetilde{H}_k - \widehat{H}_k) - \gamma \widetilde{H}_k \\ T_{k, n}(B_n) &= (1 + \gamma)(\widehat{H}_k - \widetilde{H}_k) - \gamma \widetilde{H}_k. \end{aligned}$$

Note that $\widetilde{H}_k \geq 0$ by definition of ψ_1^* as the risk minimizer. With these notations, we set $R_{k, n} = \mathbb{E}_{B_n}[R_{k, n}(B_n)]$ and $T_{k, n} = \mathbb{E}_{B_n}[T_{k, n}(B_n)]$.

We have

$$\begin{aligned}
& \mathbb{P}(R_{\hat{k}^{\text{p}o},n}(B_n) > s \mid P_{B_n}^0, P_{B_n}^1, B_n) \\
&= \mathbb{P}\left(\tilde{H}_{\hat{k}^{\text{p}o}} - \hat{H}_{\hat{k}^{\text{p}o}} > \frac{1}{1+\gamma}(s + \gamma\tilde{H}_{\hat{k}^{\text{p}o}}) \mid P_{B_n}^0, P_{B_n}^1, B_n\right) \\
&\leq K_n \max_{k \in \{1, \dots, K_n\}} \mathbb{P}\left(\tilde{H}_k - \hat{H}_k > \frac{1}{1+\gamma}(s + \gamma\tilde{H}_k) \mid P_{B_n}^0, P_{B_n}^1, B_n\right).
\end{aligned}$$

Consider the random variables $X_k = L^{\text{p}o}(\hat{\psi}_k(P_{B_n}^0), P_{B_n}^1, O) - L^{\text{p}o}(\psi_1^*, P_{B_n}^1, O)$, where O is independent from $P_{B_n}^0, P_{B_n}^1, B_n$. Let $X_{k,i} = L^{\text{p}o}(\hat{\psi}_k(P_{B_n}^0), P_{B_n}^1, O_i) - L^{\text{p}o}(\psi_1^*, P_{B_n}^1, O_i)$, $i = 1, \dots, n_2$. Conditionally on $P_{B_n}^0, P_{B_n}^1, B_n$, the $X_{k,i}$'s are n_2 i.i.d. copies of X_k in the validation set. Note that $\hat{H}_k = \sum_{i=1}^{n_2} X_{k,i}/n_2$ and $\tilde{H}_k = \mathbb{E}[X_k \mid P_{B_n}^0, P_{B_n}^1, B_n]$, so that $\tilde{H}_k - \hat{H}_k = \mathbb{E}[X_k \mid P_{B_n}^0, P_{B_n}^1, B_n] - \sum_{i=1}^{n_2} X_{k,i}/n_2$ represents an empirical mean of n_2 centered conditionally i.i.d. random variables. For the quadratic loss function, $|X_k| \leq 8M^2$ almost surely, from Assumption (11). From Lemma 2,

$$\sigma_k^2 := \text{Var}(X_k \mid P_{B_n}^0, P_{B_n}^1, B_n) \leq M_2 \mathbb{E}[X_k \mid P_{B_n}^0, P_{B_n}^1, B_n] = M_2 \tilde{H}_k.$$

For $s > 0$, $M_1 = 8M^2$, by Bernstein's inequality (see Lemma 1 in Dudoit and Van Der Laan, 2005), as the $X_{k,i}$'s are independent conditionally on $P_{B_n}^0, P_{B_n}^1, B_n$,

$$\begin{aligned}
\mathbb{P}(R_{k,n}(B_n) > s \mid P_{B_n}^0, P_{B_n}^1, B_n) &= \mathbb{P}\left(\tilde{H}_k - \hat{H}_k > \frac{1}{1+\gamma}(s + \gamma\tilde{H}_k) \mid P_{B_n}^0, P_{B_n}^1, B_n\right) \\
&\leq \mathbb{P}\left(\tilde{H}_k - \hat{H}_k > \frac{1}{1+\gamma}(s + \gamma\sigma_k^2/M_2) \mid P_{B_n}^0, P_{B_n}^1, B_n\right) \\
&\leq \exp\left(-\frac{n_2}{2(1+\gamma)^2} \frac{(s + \gamma\sigma_k^2/M_2)^2}{\sigma_k^2 + \frac{M_1}{3(1+\gamma)}(s + \gamma\sigma_k^2/M_2)}\right).
\end{aligned}$$

Note that

$$\frac{(s + \gamma\sigma_k^2/M_2)^2}{\sigma_k^2 + \frac{M_1}{3(1+\gamma)}(s + \gamma\sigma_k^2/M_2)} = \frac{s + \gamma\sigma_k^2/M_2}{\frac{\sigma_k^2}{s + \gamma\sigma_k^2/M_2} + \frac{M_1}{3(1+\gamma)}} \geq \frac{s + \gamma\sigma_k^2/M_2}{\frac{M_2}{\gamma} + \frac{M_1}{3(1+\gamma)}} \geq \frac{s}{\frac{M_2}{\gamma} + \frac{M_1}{3(1+\gamma)}}.$$

Then for $s > 0$ and $c(M, \gamma) = 2(1+\gamma)^2 \left(\frac{M_1}{3} + \frac{M_2}{\gamma}\right)$,

$$\mathbb{P}(R_{\hat{k}^{\text{p}o},n}(B_n) > s \mid P_{B_n}^0, P_{B_n}^1, B_n) \leq K_n \exp\left(-\frac{n_2}{c(M, \gamma)}s\right).$$

The same bound applies to the marginal probabilities $\mathbb{P}(R_{\hat{k}^{\text{p}o},n}(B_n) > s)$. For all $u \geq 0$,

$$\mathbb{E}[R_{\hat{k}^{\text{p}o},n}] \leq u + \int_u^\infty K_n \exp\left(-\frac{n_2}{c(M, \gamma)}s\right) ds.$$

This function of u achieves a minimum of $c(M, \gamma)(1 + \log(K_n))/n_2$ at $u = c(M, \gamma) \log(K_n)/n_2$. Thus,

$$\mathbb{E}[R_{\hat{k}^{\text{p}o},n}] \leq c(M, \gamma) \frac{1 + \log(K_n)}{n_2}.$$

The same bound applies to $\mathbb{E}[T_{\hat{k}^{\text{p}o},n}]$. Taking the expectation in Equation (17) gives the first asymptotic result stated in Equation (13). This result, combined with Lemma 2 in Dudoit and Van Der Laan (2005), allows to derive the second asymptotic result given in Equation (14).

7.2 Proof of Corollary 2.1

In the calculations below, the expectations are written with respect to the variables $O^* = (T^*, Z)$ and $O = (T^* \wedge C, \mathbb{1}\{T^* \leq C\}, Z)$, for a triplet (T^*, C, Z) independent from the data distribution B_n and from the dataset $P_{B_n} = \{P_{B_n}^0, P_{B_n}^1, P_{B_n}^2\}$. The results essentially rely on Equation (8). On the one hand,

$$\begin{aligned}
& \tilde{\theta}_n^{\text{po}}(\widehat{k}^{\text{po}}) - \tilde{\theta}_n(\widehat{k}^{\text{po}}) \\
&= \mathbb{E}_{B_n} \left[\mathbb{E} \left[L^{\text{po}}(\widehat{\psi}_{\widehat{k}^{\text{po}}}(P_{B_n}^0), P_{B_n}^1, O) \mid P_{B_n}, B_n \right] \right] - \mathbb{E}_{B_n} \left[\mathbb{E} \left[L(\widehat{\psi}_{\widehat{k}^{\text{po}}}(P_{B_n}^0), O^*) \mid P_{B_n}, B_n \right] \right] \\
&= \mathbb{E}_{B_n} \left[\mathbb{E} \left[(\Gamma_O(P_{B_n}^1) - \widehat{\psi}_{\widehat{k}^{\text{po}}}(P_{B_n}^0)(Z))^2 - (T^* \wedge \tau - \widehat{\psi}_{\widehat{k}^{\text{po}}}(P_{B_n}^0)(Z))^2 \mid P_{B_n}, B_n \right] \right] \\
&= \mathbb{E}_{B_n} \left[\mathbb{E} \left[\Gamma_O(P_{B_n}^1)^2 - (T^* \wedge \tau)^2 - 2\widehat{\psi}_{\widehat{k}^{\text{po}}}(P_{B_n}^0)(Z)(\Gamma_O(P_{B_n}^1) - T^* \wedge \tau) \mid P_{B_n}, B_n \right] \right] \\
&= \mathbb{E}_{B_n} \left[\mathbb{E} \left[\mathbb{E} \left[\Gamma_O(P_{B_n}^1)^2 \mid Z, P_{B_n}^1, B_n \right] - \mathbb{E} \left[(T^* \wedge \tau)^2 \mid Z \right] \right. \right. \\
&\quad \left. \left. - 2\widehat{\psi}_{\widehat{k}^{\text{po}}}(P_{B_n}^0)(Z) (\mathbb{E} \left[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n \right] - \mathbb{E} \left[T^* \wedge \tau \mid Z \right]) \mid P_{B_n}, B_n \right] \right] \\
&= \mathbb{E}_{B_n} \left[\mathbb{E} \left[\mathbb{E} \left[\Gamma_O(P_{B_n}^1)^2 \mid Z, P_{B_n}^1, B_n \right] - \mathbb{E} \left[(T^* \wedge \tau)^2 \mid Z \right] \right. \right. \\
&\quad \left. \left. - 2\widehat{\psi}_{\widehat{k}^{\text{po}}}(P_{B_n}^0)(Z) \mathbb{E} \left[\xi_{n_1} \mid Z, P_{B_n}^1, B_n \right] \mid P_{B_n}, B_n \right] \right].
\end{aligned}$$

Similarly,

$$\begin{aligned}
\tilde{\theta}_n^{\text{po}}(\tilde{k}^{\text{po}}) - \tilde{\theta}_n(\tilde{k}^{\text{po}}) &= \mathbb{E}_{B_n} \left[\mathbb{E} \left[\mathbb{E} \left[\Gamma_O(P_{B_n}^1)^2 \mid Z, P_{B_n}^1, B_n \right] - \mathbb{E} \left[(T^* \wedge \tau)^2 \mid Z \right] \right. \right. \\
&\quad \left. \left. - 2\widehat{\psi}_{\tilde{k}^{\text{po}}}(P_{B_n}^0)(Z) \mathbb{E} \left[\xi_{n_1} \mid Z, P_{B_n}^1, B_n \right] \mid P_{B_n}, B_n \right] \right].
\end{aligned}$$

On the other hand,

$$\begin{aligned}
& \theta_1^* - \theta^* \\
&= \mathbb{E}_{B_n} \left[\mathbb{E} \left[L^{\text{po}}(\psi_1^*, P_{B_n}^1, O) \mid P_{B_n}, B_n \right] \right] - \mathbb{E} \left[L(\psi^*, O^*) \right] \\
&= \mathbb{E}_{B_n} \left[\mathbb{E} \left[(\Gamma_O(P_{B_n}^1) - \psi_1^*(Z))^2 - (T^* \wedge \tau - \psi^*(Z))^2 \mid P_{B_n}, B_n \right] \right] \\
&= \mathbb{E}_{B_n} \left[\mathbb{E} \left[\Gamma_O(P_{B_n}^1)^2 - (T^* \wedge \tau)^2 + \psi_1^*(Z)^2 - \psi^*(Z)^2 \right. \right. \\
&\quad \left. \left. - 2\Gamma_O(P_{B_n}^1)\psi_1^*(Z) + 2(T^* \wedge \tau)\psi^*(Z) \mid P_{B_n}, B_n \right] \right] \\
&= \mathbb{E}_{B_n} \left[\mathbb{E} \left[\mathbb{E} \left[\Gamma_O(P_{B_n}^1)^2 \mid Z, P_{B_n}^1, B_n \right] - \mathbb{E} \left[(T^* \wedge \tau)^2 \mid Z \right] - (\psi_1^*(Z)^2 - \psi^*(Z)^2) \mid P_{B_n}, B_n \right] \right] \\
&= \mathbb{E}_{B_n} \left[\mathbb{E} \left[\mathbb{E} \left[\Gamma_O(P_{B_n}^1)^2 \mid Z, P_{B_n}^1, B_n \right] - \mathbb{E} \left[(T^* \wedge \tau)^2 \mid Z \right] \right. \right. \\
&\quad \left. \left. - \mathbb{E} \left[\xi_{n_1} \mid Z, P_{B_n}^1, B_n \right]^2 - 2\psi^*(Z) \mathbb{E} \left[\xi_{n_1} \mid Z, P_{B_n}^1, B_n \right] \mid P_{B_n}, B_n \right] \right].
\end{aligned}$$

Thus, if we let

$$\begin{aligned}
\phi_{n_1}(k) &= \mathbb{E}_{B_n} \left[\mathbb{E} \left[\mathbb{E} \left[\xi_{n_1} \mid Z, P_{B_n}^1, B_n \right]^2 \right. \right. \\
&\quad \left. \left. + 2(\psi^*(Z) - \widehat{\psi}_k(P_{B_n}^0)(Z)) \mathbb{E} \left[\xi_{n_1} \mid Z, P_{B_n}^1, B_n \right] \mid P_{B_n}, B_n \right] \right],
\end{aligned}$$

we have, from Equation (12),

$$\begin{aligned}
0 &\leq \mathbb{E}[\tilde{\theta}_n^{\text{po}}(\widehat{k}^{\text{po}}) - \theta_1^*] \leq (1 + 2\gamma)\mathbb{E}[\tilde{\theta}_n^{\text{po}}(\tilde{k}^{\text{po}}) - \theta_1^*] + 2c(M, \gamma)\frac{1 + \log(K_n)}{n_2} \\
\iff 0 &\leq \mathbb{E}[\tilde{\theta}_n(\widehat{k}^{\text{po}}) - \theta^* + \phi_{n_1}(\widehat{k}^{\text{po}})] \\
&\leq (1 + 2\gamma)\mathbb{E}[\tilde{\theta}_n(\tilde{k}^{\text{po}}) - \theta^* + \phi_{n_1}(\tilde{k}^{\text{po}})] + 2c(M, \gamma)\frac{1 + \log(K_n)}{n_2} \\
\iff 0 &\leq \mathbb{E}[\tilde{\theta}_n(\widehat{k}^{\text{po}}) - \theta^*] \\
&\leq (1 + 2\gamma)\mathbb{E}[\tilde{\theta}_n(\tilde{k}^{\text{po}}) - \theta^*] + \mathbb{E}[(1 + 2\gamma)\phi_{n_1}(\tilde{k}^{\text{po}}) - \phi_{n_1}(\widehat{k}^{\text{po}})] + 2c(M, \gamma)\frac{1 + \log(K_n)}{n_2}.
\end{aligned}$$

To obtain the asymptotic results, we rely on Lemma 3 below. It implies that $\tilde{\theta}_n(\tilde{k}^{\text{po}}) \rightarrow \tilde{\theta}_n(\tilde{k})$ in distribution as $n \rightarrow \infty$. In addition, we assumed the pseudo-observations and candidate learners to be bounded (Assumption (11)). Thus $\mathbb{E}[\tilde{\theta}_n(\tilde{k}^{\text{po}}) - \tilde{\theta}_n(\tilde{k})] \rightarrow 0$ as $n \rightarrow \infty$. Besides, $\phi_{n_1}(k) = o_{\mathbb{P}}(1)$ and is bounded almost surely for all k , so that we also have $\mathbb{E}[(1 + 2\gamma)\phi_{n_1}(\tilde{k}^{\text{po}}) - \phi_{n_1}(\widehat{k}^{\text{po}})] \rightarrow 0$ as $n \rightarrow \infty$. Hence the first asymptotic result (Equation (15)) which, combined with Lemma 2 in Dudoit and Van Der Laan (2005), allows to derive the second asymptotic result (Equation (16)).

Lemma 3. $\tilde{\theta}_n(\tilde{k}^{\text{po}}) \rightarrow \tilde{\theta}_n(\tilde{k})$ in probability as $n \rightarrow \infty$.

Proof. Let $f_k(n) = \mathbb{E}_{B_n}[\mathbb{E}[(T^* \wedge \tau - \widehat{\psi}_k(P_{B_n}^0)(Z))^2 \mid P_{B_n}^0, B_n]]$. We have $f_k(n) = \tilde{\theta}_n(k)$, so that

$$\tilde{k} = \arg \min_{k \in \{1, \dots, K_n\}} f_k(n).$$

On the other hand, using arguments similar to the proof of Lemma 1, we find

$$\begin{aligned}
\tilde{k}^{\text{po}} &= \arg \min_{k \in \{1, \dots, K_n\}} \tilde{\theta}_n^{\text{po}}(k) \\
&= \arg \min_{k \in \{1, \dots, K_n\}} \mathbb{E}_{B_n} \left[\mathbb{E} \left[(\Gamma_O(P_{B_n}^1) - \widehat{\psi}_k(P_{B_n}^0)(Z))^2 \mid P_{B_n}^0, P_{B_n}^1, B_n \right] \right] \\
&= \arg \min_{k \in \{1, \dots, K_n\}} \mathbb{E}_{B_n} \left[\mathbb{E} \left[\left(\mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] - \widehat{\psi}_k(P_{B_n}^0)(Z) \right)^2 \mid P_{B_n}^0, P_{B_n}^1, B_n \right] \right].
\end{aligned}$$

Using Equation (8), we write

$$\begin{aligned}
&\mathbb{E} \left[\left(\mathbb{E}[\Gamma_O(P_{B_n}^1) \mid Z, P_{B_n}^1, B_n] - \widehat{\psi}_k(P_{B_n}^0)(Z) \right)^2 \mid P_{B_n}^0, P_{B_n}^1, B_n \right] \\
&= \mathbb{E} \left[\left(\mathbb{E}[T^* \wedge \tau \mid Z] + \mathbb{E}[\xi_{n_1} \mid Z, P_{B_n}^1, B_n] - \widehat{\psi}_k(P_{B_n}^0)(Z) \right)^2 \mid P_{B_n}^0, P_{B_n}^1, B_n \right] \\
&= \mathbb{E} \left[\left(\mathbb{E}[T^* \wedge \tau \mid Z] - T^* \wedge \tau \right)^2 \mid P_{B_n}^0, P_{B_n}^1, B_n \right] \\
&\quad + \mathbb{E} \left[\left(T^* \wedge \tau + \mathbb{E}[\xi_{n_1} \mid Z, P_{B_n}^1, B_n] - \widehat{\psi}_k(P_{B_n}^0)(Z) \right)^2 \mid P_{B_n}^0, P_{B_n}^1, B_n \right] \\
&\quad + 2\mathbb{E} \left[\left(\mathbb{E}[T^* \wedge \tau \mid Z] - T^* \wedge \tau \right) \left(T^* \wedge \tau + \mathbb{E}[\xi_{n_1} \mid Z, P_{B_n}^1, B_n] - \widehat{\psi}_k(P_{B_n}^0)(Z) \right) \mid P_{B_n}^0, P_{B_n}^1, B_n \right].
\end{aligned}$$

The second term can be developed into

$$\begin{aligned} & \mathbb{E} \left[\left(T^* \wedge \tau + \mathbb{E}[\xi_{n_1} \mid Z, P_{B_n}^1, B_n] - \widehat{\psi}_k(P_{B_n}^0)(Z) \right)^2 \mid P_{B_n}^0, P_{B_n}^1, B_n \right] \\ &= \mathbb{E} \left[\left(T^* \wedge \tau - \widehat{\psi}_k(P_{B_n}^0)(Z) \right)^2 \mid P_{B_n}^0, B_n \right] + \mathbb{E} \left[\mathbb{E}[\xi_{n_1} \mid Z, P_{B_n}^1, B_n]^2 \mid P_{B_n}^1, B_n \right] \\ & \quad + 2\mathbb{E} \left[(T^* \wedge \tau) \mathbb{E}[\xi_{n_1} \mid Z, P_{B_n}^1, B_n] \mid P_{B_n}^1, B_n \right] - 2\mathbb{E} \left[\widehat{\psi}_k(P_{B_n}^0)(Z) \xi_{n_1} \mid P_{B_n}^0, P_{B_n}^1, B_n \right], \end{aligned}$$

while the third term simplifies into

$$\begin{aligned} & \mathbb{E} \left[\left(\mathbb{E}[T^* \wedge \tau \mid Z] - T^* \wedge \tau \right) \left(T^* \wedge \tau + \mathbb{E}[\xi_{n_1} \mid Z, P_{B_n}^1, B_n] - \widehat{\psi}_k(P_{B_n}^0)(Z) \right) \mid P_{B_n}^0, P_{B_n}^1, B_n \right] \\ &= \mathbb{E} \left[\left(\mathbb{E}[T^* \wedge \tau \mid Z] - T^* \wedge \tau \right) (T^* \wedge \tau) \right] \\ & \quad + \mathbb{E} \left[\left(\mathbb{E}[T^* \wedge \tau \mid Z] - T^* \wedge \tau \right) \left(\mathbb{E}[\xi_{n_1} \mid Z, P_{B_n}^1, B_n] - \widehat{\psi}_k(P_{B_n}^0)(Z) \right) \mid P_{B_n}^0, P_{B_n}^1, B_n \right] \\ &= \mathbb{E} \left[\left(\mathbb{E}[T^* \wedge \tau \mid Z] - T^* \wedge \tau \right) (T^* \wedge \tau) \right]. \end{aligned}$$

Removing all terms independent from k , we find

$$\widetilde{k}^{\text{Po}} = \arg \min_{k \in \{1, \dots, K_n\}} \left\{ f_k(n) + \varepsilon_k(n) \right\}$$

where $\varepsilon_k(n) = -2\mathbb{E}_{B_n} \left[\mathbb{E} \left[\widehat{\psi}_k(P_{B_n}^0)(Z) \xi_{n_1} \mid P_{B_n}^0, P_{B_n}^1, B_n \right] \right] = o_{\mathbb{P}}(1)$ as the $\widehat{\psi}_k(P_{B_n}^0)$'s are bounded by assumption and $\xi_{n_1} = o_{\mathbb{P}}(1)$.

Let $f_{(1)_n}(n) < \dots < f_{(K_n)_n}(n)$ denote the order statistics of $f_1(n), \dots, f_{K_n}(n)$. The order statistics depends on n both by the dependence of f_1, \dots, f_{K_n} on n and by the number K_n increasing with n . Asymptotically, either the $\{f_k(n) + \varepsilon_k(n)\}$'s are ordered in the same way as the $f_k(n)$'s, i.e. $\widetilde{k}^{\text{Po}} = \widetilde{k}$, or the inversions are negligible in terms of risk. Let $C > 0$. First suppose that there exist k, k' and n such that $|f_k(n) - f_{k'}(n)| < C$. This is equivalent to

$$|\widetilde{\theta}_n(k) - \widetilde{\theta}_n(k')| < C.$$

Now suppose that for all k, k' and for all n , $|f_k(n) - f_{k'}(n)| > C$. For all $k = 1, \dots, K_n$, $\varepsilon_k(n) = o_{\mathbb{P}}(1)$, thus there exists $N > 0$ such that for all $n \geq N$, for all $k = 1, \dots, K_n$, $|\varepsilon_k(n)| < C/2$ with probability one. Thus, for all $n \geq N$, if $f_{(k)_n}(n) < f_{(k')_n}(n)$, then

$$\begin{aligned} f_{(k)_n}(n) + \varepsilon_{(k)_n}(n) &< f_{(k)_n}(n) + \frac{C}{2} < \frac{f_{(k')_n}(n) + f_{(k)_n}(n)}{2} \\ f_{(k')_n}(n) + \varepsilon_{(k')_n}(n) &> f_{(k)_n}(n) - \frac{C}{2} > \frac{f_{(k')_n}(n) + f_{(k)_n}(n)}{2}, \end{aligned}$$

i.e. $f_{(k)_n}(n) + \varepsilon_{(k)_n}(n) < f_{(k')_n}(n) + \varepsilon_{(k')_n}(n)$. The order is preserved asymptotically, i.e. $\widetilde{k}^{\text{Po}} = \widetilde{k}$, with probability one. We conclude by making C tend towards 0. \square

References

- P. K. Andersen and M. Pohar Perme. Pseudo-observations in survival analysis. *Statistical Methods in Medical Research*, 19(1):71–99, Feb. 2010. ISSN 0962-2802, 1477-0334. doi: 10.1177/0962280209105020.
- P. K. Andersen, M. G. Hansen, and J. P. Klein. Regression Analysis of Restricted Mean Survival Time Based on Pseudo-Observations. *Lifetime Data Analysis*, 10(4):335–350, Dec. 2004. ISSN 1380-7870, 1572-9249. doi: 10.1007/s10985-004-4771-0.
- N. Binder, T. A. Gerds, and P. K. Andersen. Pseudo-observations for competing risks with covariate dependent censoring. *Lifetime Data Analysis*, 20(2):303–315, Apr. 2014. ISSN 1380-7870, 1572-9249. doi: 10.1007/s10985-013-9247-7.
- A. Cwiling, V. Perduca, and O. Bouaziz. A Comprehensive Framework for Evaluating Time to Event Predictions using the Restricted Mean Survival Time, 2023.
- A. Devaux, R. Genuer, K. Peres, and C. Proust-Lima. Individual dynamic prediction of clinical endpoint from large dimensional longitudinal biomarker history: a landmark approach. *BMC Medical Research Methodology*, 22(1):188, 2022. ISSN 1471-2288. doi: 10.1186/s12874-022-01660-3.
- C. J. DiCiccio, T. J. DiCiccio, and J. P. Romano. Exact tests via multiple data splitting. *Statistics & Probability Letters*, 166:108865, Nov. 2020. ISSN 01677152. doi: 10.1016/j.spl.2020.108865.
- S. Dudoit and M. J. Van Der Laan. Asymptotics of cross-validated risk estimation in estimator selection and performance assessment. *Statistical Methodology*, 2(2):131–154, 2005. ISSN 1572-3127. doi: 10.1016/j.stamet.2005.02.003.
- A. Eaton, T. Therneau, and J. Le-Rademacher. Designing clinical trials with (restricted) mean survival time endpoint: Practical considerations. *Clinical Trials*, 17(3):285–294, 2020. ISSN 1740-7745, 1740-7753. doi: 10.1177/1740774520905563.
- J. K. Furberg, P. K. Andersen, S. Korn, M. Overgaard, and H. Ravn. Bivariate pseudo-observations for recurrent event analysis with terminal events. *Lifetime Data Analysis*, 29(2):256–287, Apr. 2023. ISSN 1380-7870, 1572-9249. doi: 10.1007/s10985-021-09533-5.
- M. K. Golmakani and E. C. Polley. Super Learner for Survival Data Prediction. *The International Journal of Biostatistics*, 16(2):20190065, Nov. 2020. ISSN 2194-573X, 1557-4679. doi: 10.1515/ijb-2019-0065.
- P. Gonzalez Ginestet, A. Kotalik, D. M. Vock, J. Wolfson, and E. E. Gabriel. Stacked Inverse Probability of Censoring Weighted Bagging: A Case Study In the InfCareHIV Register. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 70(1):51–65, 2021. ISSN 0035-9254, 1467-9876. doi: 10.1111/rssc.12448.
- P. Gonzalez Ginestet, E. E. Gabriel, and M. C. Sachs. Survival stacking with multiple data types using pseudo-observation-based-AUC loss. *Journal of Biopharmaceutical Statistics*, 32(6):858–870, Nov. 2022. ISSN 1054-3406, 1520-5711. doi: 10.1080/10543406.2022.2041655.
- M. K. Grand, H. Putter, A. Allignol, and P. K. Andersen. A note on pseudo-observations and left-truncation. *Biometrical Journal. Biometrische Zeitschrift*, 61(2):290–298, Mar. 2019. ISSN 1521-4036. doi: 10.1002/bimj.201700274.

- H. Ishwaran, U. B. Kogalur, E. H. Blackstone, and M. S. Lauer. Random survival forests. *The Annals of Applied Statistics*, 2(3), Sept. 2008. ISSN 1932-6157. doi: 10.1214/08-AOAS169.
- M. Jacobsen and T. Martinussen. A Note on the Large Sample Properties of Estimators Based on Generalized Linear Models for Correlated Pseudo-observations. *Scandinavian Journal of Statistics*, 43(3):845–862, 2016. ISSN 0303-6898, 1467-9469. doi: 10.1111/sjos.12212.
- S. Keles, M. J. Van Der Laan, and S. Dudoit. Asymptotically optimal model selection method with right censored outcomes. *Bernoulli*, 10(6), 2004. ISSN 1350-7265. doi: 10.3150/bj/1106314848.
- M. Sachs, A. Discacciati, A. Everhov, O. Olén, and E. Gabriel. Ensemble Prediction of Time-to-Event Outcomes with Competing Risks: A Case-Study of Surgical Complications in Crohn’s Disease. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 68, July 2019. doi: 10.1111/rssc.12367.
- L. Tian, L. Zhao, and L. J. Wei. Predicting the restricted mean event time with the subject’s baseline covariates in survival analysis. *Biostatistics*, 15(2):222–233, Apr. 2014. ISSN 1465-4644, 1468-4357. doi: 10.1093/biostatistics/kxt050.
- M. J. Van Der Laan and S. Dudoit. Unified Cross-Validation Methodology For Selection Among Estimators and a General Cross-Validated Adaptive Epsilon-Net Estimator: Finite Sample Oracle Inequalities and Examples. *U.C. Berkeley Division of Biostatistics Working Paper Series*, Nov. 2003.
- M. J. Van Der Laan, E. C. Polley, and A. E. Hubbard. Super Learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1), Jan. 2007. ISSN 1544-6115, 2194-6302. doi: 10.2202/1544-6115.1309.
- A. W. Van Der Vaart, S. Dudoit, and M. J. Van Der Laan. Oracle inequalities for multi-fold cross validation. *Statistics & Decisions*, 24(3):351–371, Dec. 2006. ISSN 0721-2631. doi: 10.1524/stdn.2006.24.3.351.
- X. Wang and D. E. Schaubel. Modeling restricted mean survival time under general censoring mechanisms. *Lifetime Data Analysis*, 24(1):176–199, Jan. 2018. ISSN 1380-7870, 1572-9249.
- L. Zhao. Deep neural networks for predicting restricted mean survival times. *Bioinformatics*, 36(24):5672–5677, Apr. 2021. ISSN 1367-4803, 1460-2059. doi: 10.1093/bioinformatics/btaa1082.