



A Comprehensive Framework for Evaluating Time to Event Predictions using the Restricted Mean Survival Time

Ariane Cwiling, Vittorio Perduca, Olivier Bouaziz

► To cite this version:

Ariane Cwiling, Vittorio Perduca, Olivier Bouaziz. A Comprehensive Framework for Evaluating Time to Event Predictions using the Restricted Mean Survival Time. 2023. hal-04143419

HAL Id: hal-04143419

<https://hal.science/hal-04143419>

Preprint submitted on 27 Jun 2023

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.

A Comprehensive Framework for Evaluating Time to Event Predictions using the Restricted Mean Survival Time

Ariane Cwiling*, Vittorio Perduca*, and Olivier Bouaziz*

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

Abstract. The restricted mean survival time (RMST) is a widely used quantity in survival analysis due to its straightforward interpretation. For instance, predicting the time to event based on patient attributes is of great interest when analyzing medical data. In this paper, we propose a novel framework for evaluating RMST estimations. Our criterion estimates the mean squared error of an RMST estimator using Inverse Probability Censoring Weighting (IPCW). A model-agnostic conformal algorithm adapted to right-censored data is also introduced to compute prediction intervals and to evaluate variable importance. Our framework is valid for any RMST estimator that is asymptotically convergent and works under model misspecification.

Keywords. RMST, prediction, IPCW, Brier score, conformal, prediction intervals, variable importance.

1 Introduction

In survival analysis, in the context of right-censored data, it is common to model the effect of covariates on the hazard rate, using the ubiquitous Cox model (see Cox, 1975). This model is interpreted in terms of hazard ratios and is widely used to analyze incomplete data in medical applications. However, it relies on the strong assumption of proportional hazard (PH). As a result, approaches that focus on other estimands, such as the restricted mean survival time (RMST), have been proposed. This quantity represents the expected duration of the minimum between the occurrence of an event of interest and a predefined time horizon. It is clinically meaningful (as an expected time) and has gained considerable attention in recent years due to its simple interpretability. While initial works on this topic still relied on the PH assumption (see Karrison, 1987; Zucker, 1998), new approaches have been developed to directly model the RMST without making any assumptions (see Andersen et al., 2004; Tian et al., 2014; Wang and Schaubel, 2018).

Also, in time to event analysis, it might be of interest to predict the future occurrence of the event of interest using an estimation of the RMST. This is the case for instance, when clinicians aim at predicting the time to relapse, cancer occurrence or death of a patient. In recent years, new methods have been developed in this context (see Zhao, 2021) and there is thus a need for evaluating the performance of those learning models. This is usually a challenge in the presence of right-censored data because the censored times are not observed and it is therefore difficult to assess the performance of the learning model for those times. To address this challenge, the C-index (see Heagerty and Zheng, 2005) has emerged as a widely used metric, particularly as it

has been adapted to censored data (see Gerds et al., 2013). However, it has been shown not to be a proper scoring function (see Blanche et al., 2019). In contrast, the time-dependent area under the ROC curve is a proper score, but it is also based on the evaluation of the rank preservation of the predictions. When it comes to quantitative measures, the mean squared error (MSE) is a proper one, but it is not readily available due to censoring. Gerds and Schumacher (2006) proposed a measure designed to approximate the MSE of a survival function estimator, but, prior to our work, this had not been extended to the context of RMST.

Another important topic in this context concerns the construction of prediction intervals, which evaluate the degree of confidence in the prediction by taking into account the individual variability. The conformal approach originally proposed by Vovk et al. (2005) and later expanded and popularized by Lei et al. (2018), offers a way to build prediction intervals with guaranteed coverage. Although this approach has been adapted to right-censored data, it is still subject to significant constraints. Boström et al. (2019), Chen (2020), Teng et al. (2021) proposed model-specific conformal inference algorithms for random survival forests, DeepHit, and Cox-MLP, respectively. Candès et al. (2023) proposed a model-agnostic algorithm to build a prediction lower bound for right-censored data, but with no upper bound and only for censoring of type I.

Finally, being able to interpret the output from a learning model is crucial, especially when using black-box models. To that end, one possibility is to determine the variables importance by using permutations as developed in Breiman (2001) in the context of random forests. This technique is widely used and has been extended to right-censored data, for instance in Ishwaran et al., 2008 and Zhao, 2021. Recently, model-agnostic importance measures such as LIME and SHAP have been adapted to the estimation of the survival function (see Kovalev et al., 2020; Krzyżiński et al., 2023) but no extensions for the RMST have been developed yet. Although the leave-one-covariate-out (LOCO) conformal approach introduced by Lei et al. (2018) provides an alternative method for exploring variable importance, it has not yet been extended to right-censored data.

In this work, our goal is to propose a new measure designed to approximate the MSE of an RMST estimator, relying on Inverse Probability Censoring Weighting (IPCW) to take into account right-censoring. This measure is similar to the Brier score introduced in Gerds and Schumacher (2006). In addition, a new conformal algorithm based on IPCW for the construction of prediction intervals for restricted times to event is developed, which is inspired from the split algorithm in Lei et al. (2018). It is further extended to study local and global variable importance within a learning model. In particular, a statistical test for global variable importance is proposed. Those methods are based on the LOCO procedure from Lei et al. (2018) and take into account right-censored data using IPCW. All our methods are proved to be asymptotically valid. They are illustrated on simulations and on a real data analysis on breast cancer.

In Section 2 we give the main notations used in the following sections. In Sections 3, 4, 5 we present the new methods described previously, respectively the mean squared error measure, the prediction intervals algorithms and the variable importance measures. The results on simulations are presented in Section 6 and on real data in Section 7.

2 Notations and assumptions

In the context of right-censored data, we denote by T^* the variable of interest, C the censoring time, $T = \min(T^*, C)$ the observed variable and $\delta = \mathbf{1}\{T^* \leq C\}$ 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 introduce the following notations for the cumulative distribution and survival functions: F , G , H and L are the cumulative distribution functions of T^* , C , T and Z , respectively. $S = 1 - F$

is the survival function of T^* . For all these functions, the same notations are used for the joint and conditional cumulative distribution functions with respect to Z , for instance $F(t | z) = \mathbb{P}(T^* \leq t | Z = z) = 1 - S(t | z)$. Finally, we note $P(t, \cdot, z) = \sum_{\delta=0,1} \mathbb{P}(T \leq t, \delta, Z \leq z)$.

Let $\tau_H = \inf\{t > 0 : 1 - H(t | Z) = 0 \text{ a.s.}\}$. The RMST is defined for a fixed time horizon $\tau \leq \tau_H$ as

$$\mu_\tau^*(Z) = \mathbb{E}[T^* \wedge \tau | Z] = \int_0^\tau S(t | Z) dt. \quad (1)$$

We suppose that the i.i.d. observations are divided into a training set $D_n = \{O_i : i \in \mathcal{I}\}$ of size n and a test set $D_m = \{O_j : j \in \mathcal{J}\}$ of size m . We will call a learning model/algorithm \mathcal{A} , a function $\mathcal{A} : D_n \mapsto \hat{\mu}_{\tau,n} \in \mathcal{M}_{\mathcal{A}}$ that maps a training set D_n to an RMST estimator $\hat{\mu}_{\tau,n}$, also referred to as a predictor. $\mathcal{M}_{\mathcal{A}}$ is the space of possible outcomes $\mu_\tau(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}_+$ of \mathcal{A} . We impose that these functions verify $\sup_{z \in \mathbb{R}^d} |\mu_\tau(z)| \leq K$ for a positive constant K . In the following, when the split conformal intervals will be introduced (see Section 4), the training set will be further split in two parts, that is, \mathcal{I} will be divided into two subsets \mathcal{I}_1 and \mathcal{I}_2 of sizes n_1 and n_2 respectively, such that $n_1 + n_2 = n$. In that case, $D_{n_k} = \{O_i : i \in \mathcal{I}_k\}$, $k = 1, 2$, will be the corresponding subsets of D_n . Estimators of the functions F , S , G and μ_τ^* computed on the training set D_n are written \hat{F}_n , \hat{S}_n , \hat{G}_n , and $\hat{\mu}_{\tau,n}$ respectively. If they are computed on one of the subsets D_{n_k} , $k = 1, 2$, they are denoted \hat{F}_{n_k} , \hat{S}_{n_k} , \hat{G}_{n_k} and $\hat{\mu}_{\tau,n_k}$, respectively.

Unless mentioned otherwise, we will assume conditional independence in the following sense:

$$T^* \perp\!\!\!\perp C | Z. \quad (2)$$

We will also assume the RMST estimator to be convergent in the following sense: for all $\tau \leq \tau_H$, there exists $\tilde{\mu}_\tau \in \mathcal{M}_{\mathcal{A}}$ such that

$$\int_{\mathbb{R}^d} |\hat{\mu}_{\tau,n}(z) - \tilde{\mu}_\tau(z)| L(dz) \xrightarrow{n \rightarrow \infty} 0 \text{ in probability.} \quad (3)$$

Moreover we say that the model \mathcal{A} is correctly specified if $\mu_\tau^* \in \mathcal{M}_{\mathcal{A}}$ and $\tilde{\mu}_\tau = \mu_\tau^*$. It should be noted that this assumption allows for model misspecification since we do not impose $\tilde{\mu}_\tau$ to be equal to the true RMST μ_τ^* (and in practice this will usually not be the case). Finally, we will assume the censoring estimator \hat{G}_n to be consistent using two different definitions. A censoring estimator \hat{G}_n is said to be *strongly consistent* if for all $\tau < \tau_H$,

$$\sup_{s \leq \tau, z \in \mathbb{R}^d} |\hat{G}_n(s | z) - G(s | z)| \xrightarrow{n \rightarrow \infty} 0 \text{ a.s.} \quad (4)$$

and *weakly consistent* if

$$\int_{\mathbb{R}^d} \int_0^{\tau_H} |\hat{G}_n(s | z) - G(s | z)| P(ds, \cdot, dz) \xrightarrow{n \rightarrow \infty} 0 \text{ a.s.} \quad (5)$$

We emphasize that, in those two definitions, we impose the censoring estimator to converge towards the true G . This is a strong assumption since it imposes that the censoring model is correctly specified.

3 Performance criterion for the estimation of the RMST

When the data are fully observed, a classical quantity to measure the prediction performance of an estimator $\hat{\mu}_{\tau,n}$ is the Residual Sum of Squares (RSS)

$$\text{RSS}(\hat{\mu}_{\tau,n}, D_m) = \frac{1}{m} \sum_{j \in \mathcal{J}} \left(T_j^* \wedge \tau - \hat{\mu}_{\tau,n}(Z_j) \right)^2, \quad (6)$$

where $\hat{\mu}_{\tau,n}$ is computed on the training set and the RSS is evaluated on the test set. When the RMST estimator is convergent (see Equation (3)) then, as n and m go to infinity, the RSS will converge to the Mean Squared Error (MSE)

$$\text{MSE}(\tilde{\mu}_\tau) = \mathbb{E}\left[\left(T^* \wedge \tau - \tilde{\mu}_\tau(Z)\right)^2\right].$$

However, in the context of right-censored data, the event times are not all observed and the score (6) cannot be computed. This issue has been addressed in Gerds and Schumacher, 2006 when the goal is to estimate the survival function. In our work, we use a similar approach to estimate the MSE of an RMST estimator, based on Inverse Probability Censoring Weighting (IPCW). We define:

$$\text{WRSS}(\hat{\mu}_{\tau,n}, \hat{G}_n, D_m) = \frac{1}{m} \sum_{j \in \mathcal{J}} \left(T_j \wedge \tau - \hat{\mu}_{\tau,n}(Z_j)\right)^2 \hat{\omega}_j, \quad (7)$$

where

$$\hat{\omega}_j = \frac{\mathbb{1}\{T_j \leq \tau\} \delta_j}{1 - \hat{G}_n(T_j - | Z_j)} + \frac{\mathbb{1}\{T_j > \tau\}}{1 - \hat{G}_n(\tau | Z_j)}, \quad (8)$$

and \hat{G}_n is a consistent estimator of the censoring cumulative distribution function. We have the following result.

Theorem 1. *Let \hat{G}_n be a consistent estimator in the weak sense, as defined by Equation (5). Then, under conditional independence (see Equation (2)) we have:*

$$\sup_{\tau \leq \tau_H} |\text{WRSS}(\hat{\mu}_{\tau,n}, \hat{G}_n, D_m) - \text{MSE}(\tilde{\mu}_\tau)| \rightarrow 0 \text{ as } n, m \rightarrow \infty \text{ in probability.}$$

Note that the validity of the result relies on Equation (5). If the latter is verified with the estimator of the censoring distribution function computed on the training set only, then it also holds on the pooled data including both the training and the test set. In applications we recommend to use all the data to estimate G . Also, a straightforward result from Theorem 1 is that our WRSS estimator is asymptotically equivalent to the RSS estimator defined in Equation (6) since both estimators converge towards $\text{MSE}(\tilde{\mu}_\tau)$.

Finally, we recall that the mean squared error can be decomposed in the following way:

$$\text{MSE}(\tilde{\mu}_\tau) = \mathbb{E}\left[\left(\mu_\tau^*(Z) - \tilde{\mu}_\tau(Z)\right)^2\right] + \mathbb{E}\left[\left(T^* \wedge \tau - \mu_\tau^*(Z)\right)^2\right], \quad (9)$$

where on the right-hand side of the equation, the first quantity represents an imprecision term and the second one, an inseparability term (see Gerds and Schumacher, 2006). If the model is correctly specified, i.e. $\tilde{\mu}_\tau = \mu_\tau^*$, then the imprecision term will vanish. In that case, the WRSS estimator will converge to the inseparability term, as n and m go to infinity.

4 Prediction intervals

In this section, we explain how a prediction interval can be built from the RMST estimator. Our method is based on the conformal intervals method, developed initially by Vovk et al. (2005) and extended, among others, by Lei et al. (2018). In the latter article, the authors provide algorithms to construct prediction intervals that have finite-sample validity without making any assumptions on the distribution of the observations. More specifically, for a given confidence

level $1 - \alpha$ and a new individual with covariate Z , the aim is to construct a prediction interval $\mathcal{C}(Z) \subseteq \mathbb{R}$ such that

$$\mathbb{P}(T^* \wedge \tau \in \mathcal{C}(Z)) \geq 1 - \alpha.$$

Our method adapts the conformal prediction approach from Lei et al. (2018) using the IPCW technique to deal with right-censored data.

4.1 IPCW Split Conformal algorithm

Since the original conformal prediction algorithm is computationally intensive, we rely on the so-called *split conformal prediction* developed by Lei et al. (2018) as an alternative approach. Its computational cost is a small fraction of the full conformal method and finite-sample guarantees are very similar. It operates as follows. First divide \mathcal{I} into two subsets \mathcal{I}_1 and \mathcal{I}_2 of size n_1 and n_2 respectively, such that $n_1 + n_2 = n$. Let $D_{n_k} = \{O_i : i \in \mathcal{I}_k\}$, $k = 1, 2$, be the corresponding subsets of D_n . Train the learning algorithm on D_{n_1} . The resulting estimator $\hat{\mu}_{\tau, n_1}$ provides predictions for the data in D_{n_2} , namely $\{\hat{\mu}_{\tau, n_1}(Z_i), i \in \mathcal{I}_2\}$.

For the ease of presentation, consider first a situation where the data are uncensored. In this setting, the residuals

$$R_i^* = |T_i^* \wedge \tau - \hat{\mu}_{\tau, n_1}(Z_i)|, i \in \mathcal{I}_2,$$

can be directly computed from the data. As a result, the cumulative distribution function of the residuals, defined for all $t \geq 0$ by $\mathcal{R}^*(t) := \mathbb{P}(R^* \leq t)$, can be approximated by the empirical estimator

$$\hat{\mathcal{R}}_{n_2}^*(t) = \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \mathbb{1}\{R_i^* \leq t\}.$$

Finally, the prediction interval for a new individual with covariate Z is defined as

$$\mathcal{C}_n^{\text{split}}(Z) = [\hat{\mu}_{\tau, n_1}(Z) - \hat{q}_{n_2}^*, \hat{\mu}_{\tau, n_1}(Z) + \hat{q}_{n_2}^*],$$

with $\hat{q}_{n_2}^* = \inf\{t : \hat{\mathcal{R}}_{n_2}^*(t) \geq 1 - \alpha\} = R_{([n_2(1-\alpha)])}^*$ the $1 - \alpha$ quantile of the empirical distribution $\hat{\mathcal{R}}_{n_2}^*$, where $R_{(1)}^* \leq \dots \leq R_{(n_2)}^*$ denotes the order statistics of $R_1^*, \dots, R_{n_2}^*$. By exchangeability between the residual at (T^*, Z) and those at $(T_i^*, Z_i), i \in \mathcal{I}_2$, we have

$$\mathbb{P}(T^* \wedge \tau \in \mathcal{C}_n^{\text{split}}(Z)) = \mathbb{P}(R^* \leq \hat{q}_{n_2}^*) \geq 1 - \alpha.$$

However, when the data are censored, the residuals can no longer be computed. Clearly, computing the residuals from the observed times

$$R_i = |T_i \wedge \tau - \hat{\mu}_{\tau, n_1}(Z_i)|, i \in \mathcal{I}_2,$$

will induce a bias. To correct for this bias, we therefore propose to adjust the estimator of the cumulative distribution function of the residuals by IPCW, using the same weights as in Equation (8). The distribution estimator now becomes

$$\hat{\mathcal{R}}_{n_2}^{\hat{G}}(t) = \frac{1}{\sum_{i \in \mathcal{I}_2} \hat{\omega}_i} \sum_{i \in \mathcal{I}_2} \mathbb{1}\{R_i \leq t\} \hat{\omega}_i.$$

We present below the algorithm that summarizes the procedure. Then, Theorem 2 ensures that the prediction intervals produced by Algorithm 1 have asymptotically valid coverage.

Algorithm 1 IPCW Split Conformal prediction

Input: Data $D_n = \{O_i : i \in \mathcal{I}\}$, miscoverage level $\alpha \in (0, 1)$, regression algorithm \mathcal{A} for the RMST, regression algorithm \mathcal{B} for the censoring function G , split coefficient $\rho \in (0, 1)$

Output: Prediction interval, over $z \in \mathbb{R}^d$

$$\hat{G}_n = \mathcal{B}(D_n)$$

Randomly split \mathcal{I} into subsets $\mathcal{I}_1, \mathcal{I}_2$ of sizes $n_1 = \lfloor \rho n \rfloor$ and $n_2 = n - n_1$ such that $D_{n_1} = \{O_i : i \in \mathcal{I}_1\}$ and $D_{n_2} = \{O_i : i \in \mathcal{I}_2\}$

$$\hat{\mu}_{\tau, n_1} = \mathcal{A}(D_{n_1})$$

$$R_i = |T_i \wedge \tau - \hat{\mu}_{\tau, n_1}(Z_i)| \text{ and } \hat{\omega}_i = \frac{\mathbb{1}\{T_i \leq \tau\} \delta_i}{1 - \hat{G}_n(T_i - |Z_i|)} + \frac{\mathbb{1}\{T_i > \tau\}}{1 - \hat{G}_n(\tau |Z_i|)}, i \in \mathcal{I}_2$$

$\hat{q}_{n_2} = \inf\{t, \hat{\mathcal{R}}_{n_2}^{\hat{G}}(t) \geq 1 - \alpha\}$ the $(1 - \alpha)$ -quantile of the empirical cumulative distribution function of the residuals defined for all $t \in \mathbb{R}$ by

$$\hat{\mathcal{R}}_{n_2}^{\hat{G}}(t) = \frac{1}{\sum_{i \in \mathcal{I}_2} \hat{\omega}_i} \sum_{i \in \mathcal{I}_2} \mathbb{1}\{R_i \leq t\} \hat{\omega}_i$$

return $\mathcal{C}_n^{\text{split}}(z) = [\hat{\mu}_{\tau, n_1}(z) - \hat{q}_{n_2}, \hat{\mu}_{\tau, n_1}(z) + \hat{q}_{n_2}]$ for all $z \in \mathbb{R}^d$

Theorem 2. Suppose that \hat{G}_n is strongly consistent in the sense defined by Equation (4). If the observations O_i , $i \in \mathcal{I}$, are i.i.d., then for a new i.i.d. individual with event time T^* and covariate Z , under conditional independence (see Equation (2)),

$$\lim_{n \rightarrow \infty} \mathbb{P}(T^* \wedge \tau \in \mathcal{C}_n^{\text{split}}(Z)) \geq 1 - \alpha$$

for the split conformal prediction interval $\mathcal{C}_n^{\text{split}}$ constructed by Algorithm 1. In addition, if the residuals have a continuous distribution, then

$$\lim_{n \rightarrow \infty} \mathbb{P}(T^* \wedge \tau \in \mathcal{C}_n^{\text{split}}(Z)) = 1 - \alpha.$$

Remark 1. The validity of the results only relies on the accuracy of the estimate of the cumulative distribution function of the residuals, which depends on n_2 . However, a higher value of n_1 yields a better predictor $\hat{\mu}_{\tau, n_1}$, resulting in smaller residuals and shorter prediction intervals, which is desired in applications.

Remark 2. In most cases the residuals will have a continuous distribution. However, it may happen that the residuals distribution put a positive mass on discrete values when some predictions are identical and the corresponding observations $T_i \wedge \tau$ are also equal. This can occur if some individuals have the exact same covariate values or if they only differ for some covariates that are not used by the learning model. It is also common to use a reference model that does not take into account the covariates at all, in which case all predictions will be identical. When the residuals have a positive mass on some discrete values it may not be possible to exactly reach the desired confidence level and, in that case, the quantile of the residuals \hat{q}_{n_2} will be chosen such that the coverage exceeds $1 - \alpha$. Another approach consists in randomly breaking the ties in the residuals, as in Kuchibhotla (2021).

4.2 IPCW In-Sample Split Conformal algorithm

The split algorithm introduced in the previous section aims at producing a prediction interval with the correct coverage for a new individual independent from the training set. However it

might also be of interest to construct prediction intervals for the individuals in the training set itself. A simple, yet computationally inefficient, way to compute the prediction interval $\mathcal{C}(Z_i)$ for each $i \in \mathcal{I}$, is to apply Algorithm 2 using all observations but O_i as training set. A more interesting approach is the Rank-One-Out (ROO) Split Conformal algorithm introduced by Lei et al. (2018) to construct valid prediction intervals for all training data without requiring much more calculation. Similarly to Section 4.1, we adapt the ROO algorithm to the right-censoring framework using IPCW. Our weighted leave-one-out method is presented in Algorithm 2. Asymptotic guarantees as in Theorem 2 can be similarly derived for Algorithm 2 but have been omitted for the sake of conciseness.

Algorithm 2 IPCW Rank-One-Out Split Conformal algorithm

Input: Data $D_n = \{O_i : i \in \mathcal{I}\}$, miscoverage level $\alpha \in (0, 1)$, regression algorithm \mathcal{A} for the RMST, regression algorithm \mathcal{B} for the censoring function G , split coefficient $\rho \in (0, 1)$

Output: Prediction intervals at each $Z_i, i \in \mathcal{I}$

$\hat{G}_n = \mathcal{B}(D_n)$

Randomly split \mathcal{I} into two equal-sized subsets $\mathcal{I}_1, \mathcal{I}_2$ of sizes $n_1 = \lfloor n/2 \rfloor$ and $n_2 = n - n_1$ such that $D_{n_1} = \{O_i : i \in \mathcal{I}_1\}$ and $D_{n_2} = \{O_i : i \in \mathcal{I}_2\}$

for $k \in \{1, 2\}$ **do**

$\hat{\mu}_{\tau, n_k} = \mathcal{A}(D_{n_k})$

$R_i = |T_i \wedge \tau - \hat{\mu}_{\tau, n_k}(Z_i)|$ and $\hat{\omega}_i = \frac{\mathbb{1}\{T_i \leq \tau\} \delta_i}{1 - \hat{G}_n(T_i - |Z_i|)} + \frac{\mathbb{1}\{T_i > \tau\}}{1 - \hat{G}_n(\tau | Z_i)}, i \notin \mathcal{I}_k$

for $i \notin \mathcal{I}_k$ **do**

$\hat{q}_{n_k, i} = \inf\{t, \hat{\mathcal{R}}_{n_k, i}^{\hat{G}}(t) \geq 1 - \alpha\}$ the $(1 - \alpha)$ -quantile of the empirical cumulative distribution function of the residuals defined for all $t \in \mathbb{R}$ by

$$\hat{\mathcal{R}}_{n_k, i}^{\hat{G}}(t) = \frac{1}{\sum_{j \neq i, j \notin \mathcal{I}_k} \hat{\omega}_j} \sum_{j \neq i, j \notin \mathcal{I}_k} \mathbb{1}\{R_j \leq t\} \hat{\omega}_j$$

$\hat{\mathcal{C}}_{n_k, i}^{\text{roo}}(Z_i) = [\hat{\mu}_{\tau, n_k}(Z_i) - \hat{q}_{n_k, i}, \hat{\mu}_{\tau, n_k}(Z_i) + \hat{q}_{n_k, i}]$

end for

end for

return $\hat{\mathcal{C}}_n^{\text{roo}}(Z_i)$ for all $i \in \mathcal{I}$

5 Variable importance

Conformal inference can also be used to assess the importance of each variable in the learning model. In particular, the Leave-One-Covariate-Out (LOCO) inference method described by Lei et al. (2018) can be adapted to the estimation of the RMST with censored data based on the new algorithms described in Section 4. To evaluate the importance of the k th variable, $k \in \{1, \dots, d\}$, the approach involves comparing the accuracy of the predictor trained with or without the k th variable. The magnitude of the performance difference indicates the variable significance in the model. Specifically, if we denote $\hat{\mu}_{\tau, n_1}^{(-k)}$ the predictor trained on data D_{n_1} without the k th variable, we are interested in the random variable

$$\Delta_k(Z, T^*) = |T^* \wedge \tau - \hat{\mu}_{\tau, n_1}^{(-k)}(Z)| - |T^* \wedge \tau - \hat{\mu}_{\tau, n_1}(Z)| \quad (10)$$

that measures the increase in prediction accuracy resulting from the inclusion of the k th variable in the model. The higher its value above zero, the greater the variable importance. This analysis

can be conducted globally to assess the variable overall importance in the model or locally to identify specific variable values that have a more significant impact on the outcome.

5.1 Local measure of variable importance

The aim is to construct a prediction interval for $\Delta_k(Z, T^*)$. Let \mathcal{C}_n be a prediction interval constructed with the split procedure for T^* given Z with coverage $1 - \alpha$. For all $k = 1, \dots, d$, define

$$W_{n,k}(z) = \{|t \wedge \tau - \hat{\mu}_{\tau, n_1}^{(-k)}(z)| - |t \wedge \tau - \hat{\mu}_{\tau, n_1}(z)| : t \in \mathcal{C}_n(z)\}.$$

Then from the asymptotic coverage of the prediction interval \mathcal{C}_n , we have

$$\mathbb{P}(\Delta_k(Z, T^*) \in W_{n,k}(Z), \text{ for all } k = 1, \dots, d) \xrightarrow[n \rightarrow \infty]{} 1 - \alpha. \quad (11)$$

It should be stressed that the result is not given conditionally on $Z = z$, however varying the value of the covariate will nevertheless have an impact on the intervals. Therefore, it will be of interest to evaluate the effect of the k th variable by plotting the intervals $W_{n,k}(Z_j)$ for $j = 1, \dots, m$.

5.2 Global measure of variable importance

In this section, we consider that the data set D_{n_1} is fixed. As opposed to the previous sections, we make the additional assumption that the censoring C is independent from the time to event T^* and from the covariates Z . We use the IPCW technique to construct the statistical test below, specifying that the weights use the Kaplan-Meier estimator for the censoring distribution based on the sample D_{n_2} . Formally, with $1 - \hat{G}_{n_2}$ denoting the Kaplan-Meier estimator of the censoring survival function, the weights $\hat{\omega}_i$, $i \in \mathcal{I}_2$, become

$$\hat{\omega}_i = \frac{\mathbf{1}\{T_i \leq \tau\} \delta_i}{1 - \hat{G}_{n_2}(T_i -)} + \frac{\mathbf{1}\{T_i > \tau\}}{1 - \hat{G}_{n_2}(\tau)}. \quad (12)$$

The global measure of the importance of the k th variable is constructed from the distribution of $\Delta_k(Z, T^*)$ marginally over (Z, T^*) . We consider the cumulative distribution function of $\Delta_k(Z, T^*)$ conditional on T^* being lower than the threshold τ , with $\tau < \tau_H$. For simplicity, we suppose this distribution to be continuous and we are interested in inferring its median:

$$m_k = \text{median} [\Delta_k(Z, T^*) \mid T^* \leq \tau].$$

Specifically, we want to perform the following test

$$H_0 : m_k \leq 0 \text{ versus } H_1 : m_k > 0,$$

which is equivalent to the test

$$H_0 : p_k \leq 1/2 \text{ versus } H_1 : p_k > 1/2.$$

with $p_k = \mathbb{P}(\Delta_k(Z, T^*) \geq 0 \mid T^* \leq \tau)$. Let $\Phi_k(t, z) = \mathbf{1}\{|t - \hat{\mu}_{\tau, n_1}^{(-k)}(z)| - |t - \hat{\mu}_{\tau, n_1}(z)| \geq 0, 0 \leq t \leq \tau\}$. In particular, note that

$$p_k = \frac{1}{1 - S(\tau)} \int \Phi_k(u, z) dF(u, z).$$

We introduce the new test statistic

$$T_{n_2} = \sqrt{\frac{n_2}{\hat{\sigma}^2(\Phi_k)}} \left(\frac{1}{1 - \hat{S}_{n_2}(\tau)} \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \Phi_k(T_i, Z_i) \hat{\omega}_i - \frac{1}{2} \right), \quad (13)$$

where \hat{S}_{n_2} denotes the Kaplan-Meier estimator of the survival function, the censoring weights $\hat{\omega}_i$, $i \in \mathcal{I}_2$ are defined as in Equation (12), and

$$\begin{aligned} \hat{\sigma}^2(\Phi_k) = & \frac{(1 - \hat{S}_{n_2}(\tau))^{-2}}{n_2} \sum_{i \in \mathcal{I}_2} \mathbb{1}\{T_i \leq \tau\} \left(\Phi_k(T_i, Z_i) \hat{\omega}_i \right. \\ & \left. - \frac{\delta_i}{\hat{Y}_{n_2}(T_i)} \frac{1}{n_2} \sum_{j \in \mathcal{I}_2} \left(\mathbb{1}\{T_i \leq T_j\} - \frac{\hat{S}_{n_2}(\tau)}{1 - \hat{S}_{n_2}(\tau)} \right) \Phi_k(T_j, Z_j) \hat{\omega}_j \right)^2, \end{aligned}$$

with $\hat{Y}_{n_2}(t) = \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \mathbb{1}\{T_i \geq t\}$. This is a modification of the standard sign test statistic suited for right-censored data. In Theorem 3, we show that this test statistic follows asymptotically a Gaussian distribution with variance equal to 1 under H_0 . This allows us to construct the statistical test with predefined level. The construction of an asymptotic confidence interval for p_k is also proposed in Theorem 3.

Theorem 3. *Let D_{n_1} be a fixed data set. Let \hat{S}_{n_2} , $1 - \hat{G}_{n_2}$ denote the Kaplan-Meier estimators of the functions S , $1 - G$ respectively, computed on a data set D_{n_2} independent from D_{n_1} . Let the censoring be independent from the time to event and from the covariates. We consider the weights $\hat{\omega}_i$ as defined in Equation (12). Then, under H_0 , the test statistic T_{n_2} defined in Equation (13) follows asymptotically a Gaussian distribution with variance equal to 1. Let $q_{1-\alpha}^{\mathcal{N}(0,1)}$ denote the $1 - \alpha$ quantile of this distribution, then*

$$\lim_{n_2 \rightarrow \infty} \mathbb{P}_{H_0} \left(T_{n_2} > q_{1-\alpha}^{\mathcal{N}(0,1)} \right) \geq \alpha.$$

We also have

$$\lim_{n_2 \rightarrow \infty} \mathbb{P} \left(p_k \in \left[\frac{1}{1 - \hat{S}_{n_2}(\tau)} \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \Phi_k(T_i, Z_i) \hat{\omega}_i \pm \sqrt{\frac{\hat{\sigma}^2(\Phi_k)}{n_2}} q_{1-\alpha/2}^{\mathcal{N}(0,1)} \right] \right) = 1 - \alpha. \quad (14)$$

Remark 3. *To simplify the proof, Theorem 3 is presented with the censoring estimator being calculated on D_{n_2} . We were not able to prove that the results would still hold if the censoring estimator was calculated from all the data but our simulations suggest that this would be the case.*

Remark 4. *In some cases, the inclusion or exclusion of the k th variable in the learning model may not affect the predictions. This may occur for models that are independent of covariates such as the Kaplan-Meier model, or in learning models that include variable selection steps like the LASSO. When this happens, $\Delta_k(Z, T^*)$ equals 0 and p_k is equal to 1, although $p_k = 0.5$ would better indicate that the covariate has no predictive power. One way to address this issue is to introduce a different definition of p_k based on a jittered version of Δ_k :*

$$p_k = \mathbb{P}(\Delta_k(Z, T^*) + \xi U \geq 0 \mid T^* \leq \tau) \quad (15)$$

where $U \sim U[-1, 1]$ and $0 < \xi \ll 1$. This jittering approach randomly allocates a sign to $\Delta_k(Z, T^*)$ when it is equal to 0, which avoids putting mass on 0 and guarantees that $p_k = 0.5$ for the k th variable if it is not used by the predictor.

6 Simulations

This section illustrates all the methods and results described in Sections 3, 4 and 5 through simulated experiments. In Section 7 an application to the German Breast Cancer Study Group (GBCSG) is also presented. In both sections, the following learning algorithms are implemented and then evaluated using our methods.

Integrated Kaplan-Meier: A Kaplan-Meier estimator is fitted to D_n to estimate the survival function. By integrating this curve on the interval $[0, \tau]$, we obtain an estimation of the RMST that is identical for all individuals in the test sample D_m . This represents a naive algorithm since it does not take into account the covariates.

Integrated Cox: A Cox model (see Cox, 1975) is fitted to D_n . Unless mentioned otherwise, no interaction between covariates is included in the model. The fitted model provides an estimation of the survival curve for each observation O_j in D_m , adjusted with respect to the covariate Z_j . The estimation of the RMST is obtained by integrating this curve with respect to time on the interval $[0, \tau]$.

Integrated Random Survival Forest (RSF): The procedure is identical to the one described above with the difference that an RSF (see Ishwaran et al., 2008) is fitted to D_n instead of a Cox model.

Pseudo-observations and linear model: We transform the censored data in D_n into pseudo-observations (see Andersen et al., 2004). We use the linear model as the link function to obtain a linear model for the RMST. Unless mentioned otherwise, no interaction term is included in the linear model.

We stress that, in the simulations, the RMST estimations obtained with those algorithms will never be truncated even if they exceed τ . This avoids creating ties in the residuals as mentioned in Remark 2. We consider three different simulations schemes.

Scheme A: Following the setting from Wang and Schaubel (2018), the event times are simulated according to the following linear model:

$$T_i^* = \tilde{\beta}_0^\top Z_i + \varepsilon_i,$$

where $\tilde{\beta}_0 = (5.5, 2.5, 2.5)^\top$, the covariates are denoted $Z_i = (1, Z_i^1, Z_i^2)^\top$ with $Z_i^1, Z_i^2 \sim \mathcal{B}(0.5)$ and $\varepsilon_i \sim U[-3, 3]$ is a random noise. From this model we obtain the following closed form for the RMST:

$$\mu_\tau^*(Z) = \mathbb{E}[T^* \wedge \tau \mid Z] = \beta_{00} + \beta_{01}Z^1(1 - Z^2) + \beta_{10}Z^2(1 - Z^1) + \beta_{11}Z^1Z^2, \quad (16)$$

where $\beta_0 = (\beta_{00}, \beta_{01}, \beta_{10}, \beta_{11})^\top$ is computed from 10 million Monte-Carlo samples. The value of τ is fixed to 8.8 which yields $\beta_0 = (5.5, 2.097, 2.097, 3.16)^\top$. Next, two different types of censoring are considered. In scheme **A1**, the censoring is simulated independently from the covariates according to an exponential law with parameter $\lambda = 0.07$, leading to 42% of censored data. In scheme **A2**, the censoring is simulated from a Cox model $\lambda(t \mid Z) = \lambda_0(t) \exp(\beta_1 Z^1 + \beta_2 Z^2)$ with Weibull baseline hazard $\mathcal{W}(\nu, \kappa)$ defined as

$$\lambda_0(t) = \frac{\nu}{\kappa} \left(\frac{t}{\kappa} \right)^{\nu-1}.$$

We set $\kappa = 12$, $\nu = 6$, and $\beta_1 = 2$, $\beta_2 = 1$, leading to 44% of censored data.

Scheme B: The 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 $k = 1, 2, 3$, with Cox regression parameters $\beta = (\beta_1, \beta_2, \beta_3)^\top$. Note that the survival function can be expressed as

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

and the RMST can be obtained from Equation (1). Parameters are set to $\kappa = 2$, $\nu = 6$, $a = 5$, $(\beta_1, \beta_2, \beta_3) = (2, 1, 0)$. The 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 gives $\tau = 3.6$.

Scheme C: Similarly to scheme B, the event times are simulated according to a Cox model with Weibull baseline hazard $\mathcal{W}(\nu, \kappa)$, $\lambda(t | Z) = \lambda_0(t) \exp(g(Z))$, with

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

and $\kappa = 2$, $\nu = 6$. Let $Z = (Z^1, \dots, Z^{15})$, we simulate the covariates such that $Z^j \sim \mathcal{B}(0.4)$ for $j \in \{2, 4, 6, 9, 11, 12\}$ and $Z^j \sim U[0, 1]$, $j \in \{1, 3, 5, 7, 8, 10, 13, 14, 15\}$. As a result, only the first 10 covariates are associated with the event times, but the other 5 covariates (that are non-informative) will still be included in our regression models. The survival function is expressed as:

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

and the RMST can be obtained from Equation (1). The censoring distribution is the same as in scheme B, leading to 47% censored data. The time horizon τ is chosen as the 90th percentile of the observed times T which gives $\tau = 2.8$.

6.1 Illustration of the WRSS estimator

In this section, we want to illustrate Theorem 1. It states that, if the censoring estimator is consistent, then our prediction performance criterion, called WRSS (see Equation (7)), converges towards the MSE as the sample size goes to infinity. We recall that $\text{MSE}(\tilde{\mu}_\tau) = \mathbb{E}[(T^* \wedge \tau - \tilde{\mu}_\tau(Z))^2]$ and that $\tilde{\mu}_\tau$ is the limit of the predictor, as defined in Equation (3). We also recall that $\text{MSE}(\tilde{\mu}_\tau)$ can be decomposed into an inseparability and imprecision terms, as shown in Equation (9). We will consider the simulation schemes **A1**, **A2** and two learning models. The first one is a linear model that is directly fitted on the minimum between the true event times and the time horizon τ , using the correct link function (see Equation (16)). It is considered as the oracle model. The second one is based on pseudo-observations with linear link function. The latter is implemented without interaction terms, i.e. only the covariates Z^1, Z^2 are included.

In the scenario **A1**, the censoring is independent of the covariates and we use the Kaplan-Meier estimator to estimate its cumulative distribution function which is a consistent estimator (in the weak sense). Thus, the WRSS should converge to the MSE (see Theorem 1). In addition, for the oracle model, the imprecision term should vanish as the model is correctly specified and the WRSS should therefore converge to the inseparability term. On the other hand, for the model based on pseudo-observations, the imprecision term will not vanish and the WRSS should then be larger than with the oracle model. Since the RMST has an explicit form, the inseparability term

can be easily computed using Monte Carlo simulations. For the imprecision term of the model based on pseudo-observations, $\tilde{\mu}_\tau$ was approximated by a predictor $\hat{\mu}_{\tau,n}$ trained on a sample of size 20,000 and the expectation was calculated using a million Monte-Carlo simulations. In Figure 6.1, we represent the WRSS based on train and test samples of equal size 100, 500 and 1,000 for those two learning algorithms. The boxplots are obtained from 1,000 repetitions. We clearly see that the oracle estimator converges towards the inseparability term, displayed in red in the figure. On the contrary, we see that, with the estimator based on pseudo-observations, the WRSS converges towards a value greater than the inseparability term, as expected. With the pseudo-observations model, we observe that the imprecision term is relatively small compared to the inseparability term which suggests that, while the regression model is incorrect, it still provides predictions that are close to the ones obtained using the oracle model.




Figure 6.1: Distribution of 1,000 replications of the WRSS estimator in the scenario **A1** and illustration of its convergence towards $\text{MSE}(\tilde{\mu}_\tau)$ (see Equation (9)), where $\tilde{\mu}_\tau$ represents the limit defined in Equation (3). Two learning models are compared. On the left panel, the oracle model (16) is a linear model fitted on the minimum between the true event times and τ , using the correct link function. On the right panel, a linear model is implemented based on pseudo-observations, including all covariates without interaction terms. The red dotted line illustrates the inseparability term. It also represents the $\text{MSE}(\tilde{\mu}_\tau)$ for the oracle model, whose imprecision term is null. The blue dotted line represents the $\text{MSE}(\tilde{\mu}_\tau)$ for the model based on pseudo-observations, whose imprecision term is non-zero.

In the scenario **A2**, the censoring depends on the covariates through a Cox model. In this setting we also compare the performance of the two learning algorithms using different censoring models: a Kaplan-Meier method, a Cox model and an RSF model are implemented. We stress that the Kaplan-Meier method is no longer consistent while the Cox model is a consistent estimator for the censoring distribution. The results are illustrated in Figure 6.2 where the boxplots are also obtained from 1,000 repetitions. We clearly observe that the Kaplan-Meier method for the censoring distribution provides biased estimates of the MSE. On the other hand, the Cox and RSF models provide accurate estimations of the MSE. As in the scenario **A1**, the imprecision term is seen to vanish with the oracle model and has a relatively small value compared to the inseparability term, when using the pseudo-observations model. Those results also suggest that the RSF model for the censoring distribution is a consistent estimator (in the weak sense) since the results are almost identical with the results from the Cox model.

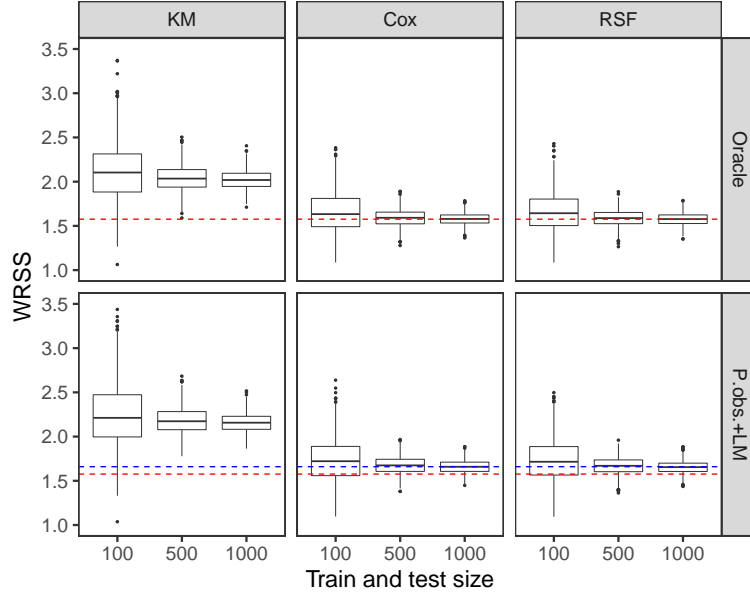


Figure 6.2: Distribution of 1,000 replications of the WRSS estimator in the scenario **A2** and illustration of its convergence towards $\text{MSE}(\tilde{\mu}_\tau)$ (see Equation (9)), where $\tilde{\mu}_\tau$ represents the limit defined in Equation (3). Two learning models are compared. On the top panel, the oracle model (16) is a linear model fitted on the minimum between the true event times and τ , using the correct link function. On the bottom panel, a linear model is implemented based on pseudo-observations, including all covariates without interaction terms. In addition, three censoring estimators are compared. From left to right, a Kaplan-Meier method, a Cox model and an RSF model. The red dotted line illustrates the inseparability term. It also represents the $\text{MSE}(\tilde{\mu}_\tau)$ for the oracle model, whose imprecision term is null. The blue dotted line represents the $\text{MSE}(\tilde{\mu}_\tau)$ for the model based on pseudo-observations, whose imprecision term is non-zero.

6.2 Illustration of the IPCW Split Conformal algorithm

In this section, we illustrate the construction of prediction intervals using the IPCW split conformal algorithm introduced in Section 4.1 (see also Algorithm 1). We simulate data using the simulation scheme **B** and we train all four learning models introduced at the beginning of Section 6. We first start by displaying in Figure 6.3 the prediction intervals at level $1 - \alpha = 0.9$ on a sample of 10 individuals while the algorithms were trained on a single independent sample of size 4,000.

Then, the coverage of the intervals, as claimed by Theorem 2, is assessed in Figure 6.4, with $1 - \alpha$ equal to 0.8, 0.9 or 0.95. This time, the learning algorithms were trained on samples of size $n = 100, 500$ and $1,000$ and the testing set is of size $m = 500$. The simulations were repeated 200 times. We observe that for all learning models and confidence levels, the empirical coverage converges to $1 - \alpha$, except for the integrated Kaplan-Meier algorithm which converges to a level greater than $1 - \alpha$. Since this algorithm gives an RMST estimation that is identical for all individuals (as it does not take into account the covariates), the observations T_i that are greater than τ will all have the same residual value. These ties make the distribution of the residuals discrete and therefore the empirical quantile of order $1 - \alpha$ is such that the coverage exceeds $1 - \alpha$ (see Remark 2).




Figure 6.3: Prediction intervals at the 90% level constructed with Algorithm 1 for four learning models: the Kaplan-Meier estimator, the Cox model, the RSF model and the linear model based on pseudo-observations. The training size is $n = 4,000$ and the prediction intervals are constructed for 10 individuals independent from the test set. All data are simulated according to the scenario **B**. The grey dotted line represents the time horizon $\tau = 3.6$. The red segments are placed at the minimum between the true event times of the test set and τ .




Figure 6.4: Empirical coverage for the prediction intervals constructed with Algorithm 1 for four learning models: the Kaplan-Meier estimator, the Cox model, the RSF model and the linear model based on pseudo-observations. All data were simulated according to the scenario **B**.

6.3 Illustration of the LOCO variable importance measures

In this section, we provide illustrations of the use and performance of the LOCO variable importance measures introduced in Section 5. We consider the simulation scheme **B** and, for all four learning algorithms, we want to test $H_0 : p_k \leq 1/2$ versus $H_1 : p_k > 1/2$, for $k = 1, 2, 3$. We use the definition (15) for p_k so as to get $p_k = 0.5$, $k = 1, 2, 3$, for the Kaplan-Meier model as explained in Remark 4. In this scenario, Z^3 has no effect on the outcome since $\beta_3 = 0$ in the Cox model used to generate the data. When using the Kaplan-Meier estimator, p_k is equal to 0.5 for all k since adding or removing a variable will have no impact on the prediction result. However, the value of p_k is unknown when using other algorithms and in particular, we do not know in advance if $p_k \leq 0.5$ (H_0 is true) or if $p_k > 0.5$ (H_1 is true). Their values are thus approximated via Monte-Carlo simulations. We first simulate a training set D_{n_1} of size $n_1 = 500$ which remains unchanged throughout the whole simulations (note that Theorem 3 holds for a fixed D_{n_1}). Next, we train the learning algorithms on this data set, simulate 10^5 pairs (T^*, Z) and compute p_k from the distribution of the corresponding $\Delta_k(Z, T^*)$. Table 1 shows the resulting values.

Learning model	p_1	p_2	p_3
Kaplan-Meier	0.50	0.50	0.50
Cox	0.87	0.79	0.49
Random Survival Forests	0.82	0.71	0.44
Pseudo-observations and linear model	0.84	0.70	0.46

Table 1: Values of p_k , $k \in \{1, 2, 3\}$ for a fixed sample D_{n_1} , generated with $n_1 = 500$ according to the scenario **B**, for four learning models: the Kaplan-Meier estimator, the Cox model, the RSF model and the linear model based on pseudo-observations.

In practice, global variable importance can be determined using Theorem (3). In Figure (6.5), using Equation (14), the confidence intervals at the 90% level for p_k , $k = 1, 2, 3$ for the four learning algorithms are displayed based on D_{n_1} and a data set D_{n_2} of size $n_2 = 500$ simulated independently according to scenario **B**. We observe that the Cox model, the RSF model and the linear model based on pseudo-observations seem to all agree that only the first two covariates are important. For the Kaplan-Meier estimator, the three confidence intervals cross 0.5.




Figure 6.5: Confidence intervals at the 90% level for p_k , $k = 1, 2, 3$ (see Equation (14)), whose values are reported in Table (1). The intervals were computed with the global variable importance measure applied to the fixed data set D_{n_1} and a data set D_{n_2} of size $n_2 = 500$ simulated independently according to scenario **B**. Four learning models are considered: the Kaplan-Meier estimator, the Cox model, the RSF model and the linear model based on pseudo-observations.

Using the same data sets D_{n_1} and D_{n_2} , the local variable importance is illustrated for all four models and all three covariates in Figure 6.6. The Kaplan-Meier model is set as a reference as none of the three covariates play a role in it. The other learning algorithms all consider the first covariate to have importance in its high and low values and the third covariate to have no importance in the predictions, but they do not reach the same conclusions about the second covariate. The Cox model detects its importance for high and low values of the variable while the other two do not detect any importance of the variable. When the size of the data set increases, our simulations show that the RSF also comes to detect importance in high and low values (data not shown). The Cox and RSF models provide then very similar conclusions for all three covariates in terms of local variable importance.

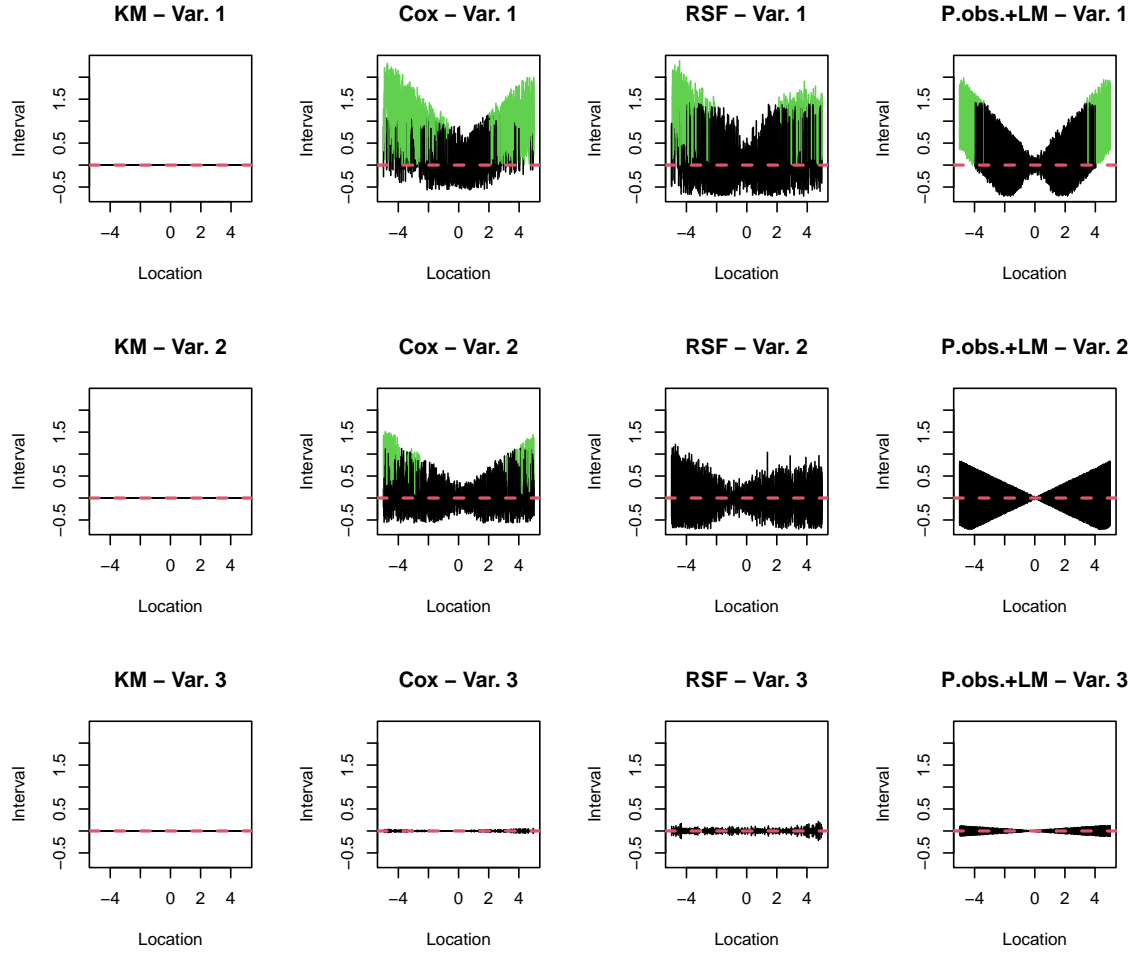


Figure 6.6: Local variable importance computed on the fixed data set D_{n_1} and a data set D_{n_2} of size $n_2 = 500$ simulated independently according to scenario the **B**, for four learning models: the Kaplan-Meier estimator, the Cox model, the RSF model and the linear model based on pseudo-observations. For the k th variable the vertical segments show the intervals $W_{n,k}(Z_j)$, $j = 1, \dots, m$, containing, with probability $1 - \alpha = 0.9$, the measure of the impact of Z_j^k on predictions (see (11)). Vertical segments lying above zero are colored in green.

Still using the same fixed data set D_{n_1} , we empirically assess the calibration and power of our test for global importance by simulating 1,000 samples D_{n_2} of size $n_2 = 500$ and by computing for each one the p-value for the statistical test. The histograms of those p-values for each value of k and all four algorithms are displayed in Figure 6.7. We first observe that the p-values for the Kaplan-Meier estimator are uniformly distributed, since in this case, $p_k = 0.5$ for all k (H_0 is true). For the other algorithms, when $k = 3$, we observe a skewed distribution of the p-values towards 1. This was also expected since the H_0 hypothesis is composite and, according to Table 1, H_0 is true for $k = 3$ for all models. Finally, when $k = 1, 2$ (H_1 is true) we observe that all three algorithms have a very strong power (all p-values are less than 0.01).

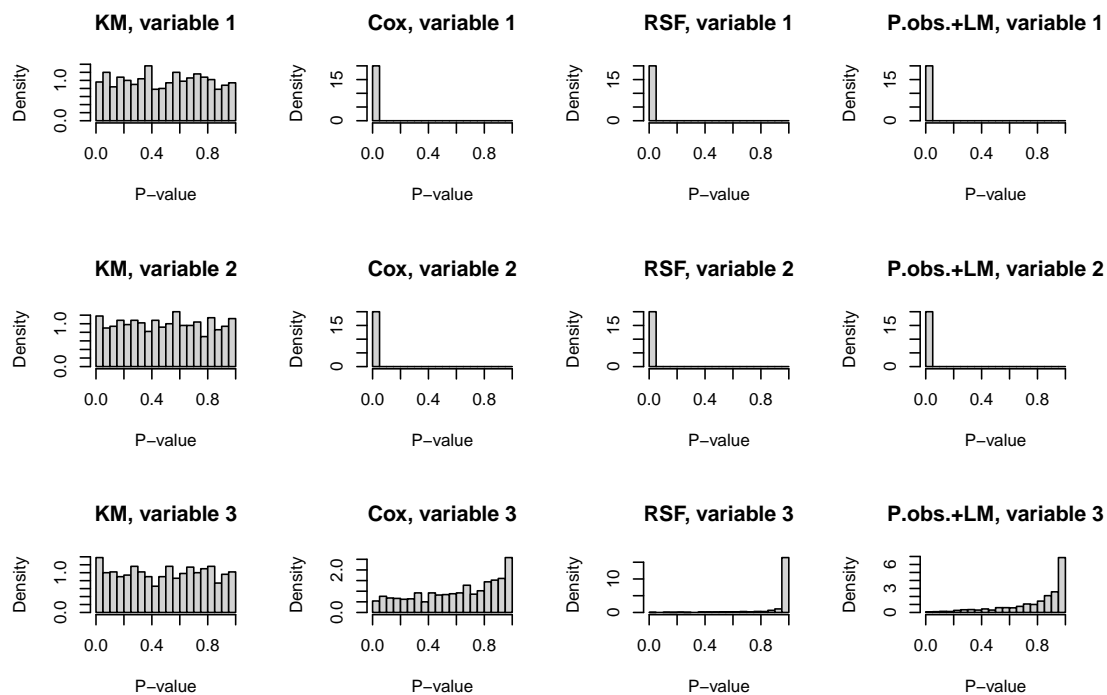


Figure 6.7: Distribution of the p-values from 1,000 repetitions of the LOCO global variable importance test, for four learning models: the Kaplan-Meier estimator, the Cox model, the RSF model and the linear model based on pseudo-observations. The sample D_{n_1} was generated with $n_1 = 500$ and remained fixed while D_{n_2} was simulated 1,000 times with $n_2 = 500$ in order to obtain the distribution of the p-values. All data were simulated according to the scenario **B**.

6.4 Multi-splitting

We want to emphasise that variable importance results depend not only on the learning algorithm but also on the split of the data. In particular when it comes to assessing global variable importance, results may vary drastically depending on the split, for instance when the link between covariates and outcome is complex such as in the simulation scheme **C**. As an illustration, the LOCO global variable importance test is performed for all four learning models, on a data set of size 1,000 simulated with the simulation scheme **C**. The oracle Cox model

taking interactions into account is added for comparison. Figure 6.8 presents the outputs of the procedure repeated twice, the only difference being the data split. Some variables seem to improve the accuracy of the model on one split, and on the contrary to degrade it on the other.

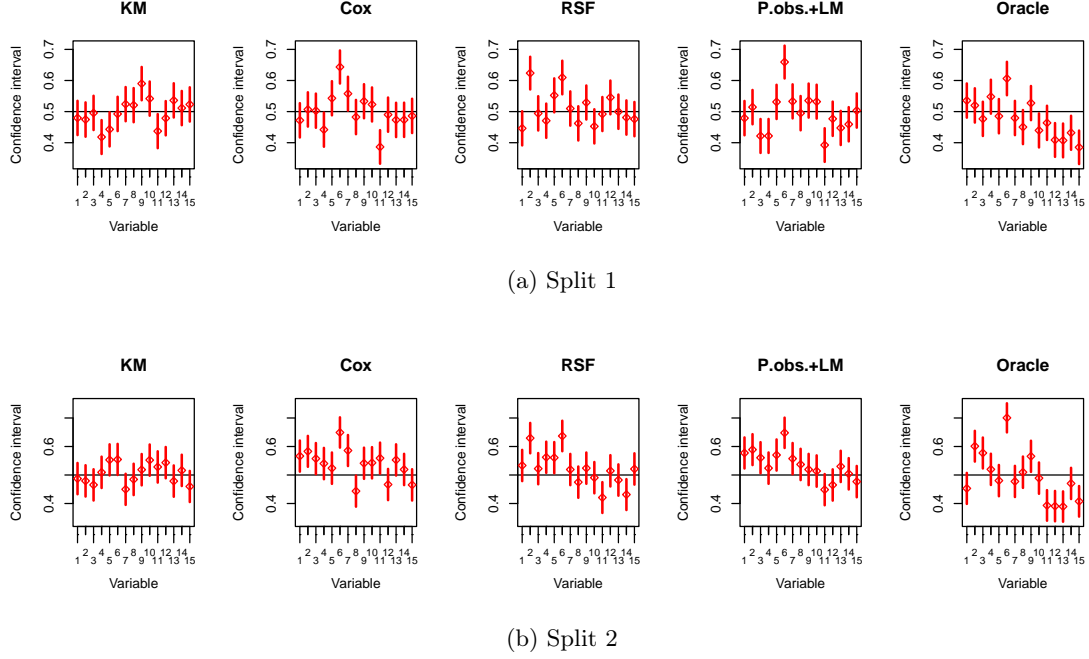


Figure 6.8: Confidence intervals at the 90% level for p_k , $k = 1, \dots, 15$ (see Equation (14)), computed with the global variable importance measure on a data set of size $n = 1,000$ simulated using the scenario **C**, for five learning models: the Kaplan-Meier estimator, the Cox model, the RSF model, the linear model based on pseudo-observations and the oracle Cox model taking interactions into account. The procedure is conducted on two different splits of the data.

To reduce the effect of the split, methods like multi splitting described in DiCiccio et al. (2020) can be considered. Specifically the authors of this paper propose to choose M splits of the data, and then compute a p-value for each split. Twice the resulting median or average p-value provides a valid p-value for the overall test. The overall Type 1 error can be controlled and bounded under α , however this comes with a loss of power compared to a single split. We selected $M = 50$ splits of our data set simulated according to simulation scheme **C**. For each split, each learning algorithm and each variable, we computed a p-value. Twice the median p-value is computed for each setting and reported in Table 2. No variables have a notable impact in the Kaplan-Meier model. For the other models, binary variables such as Z^2 , Z^6 and Z^9 seem to play a significant role on the predictions. There is a consensus on variables Z^{11} to Z^{15} , which are uninformative and indeed seem to have no impact. The effect of other variables is elusive, and can vary depending on the model.

7 Application on real data

We illustrate our comprehensive framework for evaluating RMST estimators on the classic German Breast Cancer Study Group (GBCSG) data set, which was first introduced in Schu-

Variable	P-value KM	P-value Cox	P-value RSF	P-value P.obs.+LM	P-value Oracle
1	1	0.62	1	0.588	1
2	0.982	0.007 **	0 ***	0.012 *	0.001 **
3	0.963	0.545	1	0.489	0.467
4	0.891	1	0.863	1	0.145
5	1	0.105	0.466	0.054 .	0.102
6	1	0 ***	0.001 **	0 ***	0 ***
7	1	0.453	0.54	0.804	0.225
8	1	0.958	1	0.677	0.576
9	1	0.087 .	0.05 .	0.067 .	0.046 *
10	0.893	0.992	1	1	0.46
11	1	1	1	1	1
12	1	1	1	1	1
13	1	1	1	1	1
14	1	1	1	1	1
15	0.98	1	1	0.893	1

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

Table 2: Multi splitting on a data set of size $n = 1,000$ simulated according to scenario the **C** for five learning models: the Kaplan-Meier estimator, the Cox model, the RSF model, the linear model based on pseudo-observations and the oracle Cox model taking interactions into account. $M = 50$ splits of the data are randomly chosen. The global variable importance test is performed on each one. The resulting M p-values are aggregated by twice the median value, serving as p-value of level α for the overall test.

macher et al. (1994) and is now available in the **survival** R package. The GBCSG gathered patients with node-positive breast cancer. The study was conducted from 1984 to 1989 and aimed to investigate the impact of hormone treatment on recurrence-free survival time. The event of interest is then the recurrence of a cancer, which was observed for 299 of the 686 patients. 246 patients were treated with additional hormonal therapy. Finally, prognostic factors were collected on all patients: age, menopausal status, tumour size, tumour grade, number of positive nodes, progesterone receptor and estrogen receptor. We estimated the censoring survival function using the Kaplan-Meier model. We set the time horizon τ at the 90th percentile of the observed times distribution (2014 days) and were interested in the recurrence time prediction up to that time limit. We tried all four of the learning algorithms described in Section 6. For each one, the WRSS is computed based on 20-fold cross-validation in Figure 7.1. The Kaplan-Meier model is less performing than the covariate-dependent models, indicating that the chosen prognostic factors indeed play a role on the recurrence-free survival time. The RSF seems to have a slightly better performance than the Cox model and the linear model based on pseudo-observations, though the variability of the results makes it difficult to identify clearly which algorithm is best suited to the data. Results of the global variable importance test using multi-splitting on $M = 40$ splits are displayed in Table 3. Interestingly, no variables turn out to be important in the RSF predicting the recurrence time, despite its observed overall good predictive performance. Hormonal therapy, progesterone receptors, and tumor grade are important in the prediction with both the Cox model and the linear model based on pseudo-observations.




Figure 7.1: Estimation of the mean squared error on the GBCSG data set using the WRSS based on 20-fold cross-validation, for four learning models: the Kaplan-Meier estimator, the Cox model, the RSF model and the linear model based on pseudo-observations.

Variable	P-value KM	P-value Cox	P-value RSF	P-value P.obs.+LM
hormon	1	0.001 **	0.166	0.001 **
age (years)	1	1	1	1
menopausal status	1	0.203	1	0.093 .
tumour size (mm)	1	0.989	0.498	1
number of positive nodes	1	1	0.329	1
progesterone receptor (fmol)	1	0 ***	0.359	0 ***
estrogen receptor (fmol)	1	1	1	1
tumour grade ≥ 2	1	0 ***	0.516	0 ***

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

Table 3: Multi splitting on the GBCSG data set for four learning models: the Kaplan-Meier estimator, the Cox model, the RSF model and the linear model based on pseudo-observations. $M = 40$ splits of the data are randomly chosen. The global variable importance test is performed on each split. The resulting M p-values are aggregated by twice the median value, serving as p-value of level α for the overall test.

8 Conclusion

One way of assessing the fit of a model and its prediction ability is to evaluate its MSE based on a test sample. This is a general measure that allows to compare a wide range of models without making any specific model assumption. In the presence of right-censoring, due to tail estimation issues, it is natural to focus on the prediction of the restricted time $T^* \wedge \tau$ based on a clinically relevant fixed time horizon τ (see Eaton et al., 2020). Over the recent years, there has been a wide variety of new models that can handle censored data from the random survival forest method to neural networks based methods (see Ishwaran et al., 2008; Zhao, 2021). Those methods make few assumptions as compared to the ubiquitous Cox model

that relies on the proportional hazard assumption and it is interesting to evaluate and compare their prediction abilities. Based on the MSE it is well known that the best prediction model is the conditional expectation of the restricted time and therefore predicting the restricted time amounts to estimating the RMST.

Methods for evaluating RMST estimation models remain limited. Prior to our work, concordance measures were used to compare predictive performances of RMST estimators but no quantitative metric was suggested. In addition, no model-agnostic methods were available to construct prediction intervals. In this context, we propose a complete framework for evaluating a time to event predictor using the RMST. Specifically, we propose methods to assess its prediction ability and variability, and to estimate variable importance, globally and locally. The steps of the implementation can be summarized as follows.

1. Evaluate the predictive performance of the predictor quantitatively with the WRSS (7), an estimator of the MSE based on IPCW. To reduce the variability induced by the splitting procedure, V -fold cross-validation is recommended.
2. Construct prediction intervals with the IPCW Conformal Algorithm 1 or 2, as illustrated in Figure 6.3.
3. For each variable of interest, evaluate its global importance in the regression model with the statistical test (13) and construct the corresponding confidence interval given by (14). Note that this measure is inherent in the model and in the data split and can differ from the true significance of the variable.
4. Compute the local variable importance when it is relevant, that is for variables with a sufficient amount of possible values. This measure is also inherent in the model and in the data split.

In a context of high-dimensional data, one can form groups of variables that are relevant according to domain-specific knowledge. To this aim, in the definition (10), the k which represents the k th variable can equivalently represent the k th group of variables. Otherwise, we recommend to focus only on a subset of relevant covariates by first applying a pre-selection step on the data, typically with selection variables methods such as the Lasso or based on an AIC criterion (see for instance Kojima, 2022). These methods allow to reduce the computational cost of our procedure.

Under consistency assumptions defined in Section 2, we proved that all the methods described above are asymptotically valid. We want to stress that the whole analysis procedure is model-agnostic, thus robust against model misspecification, but not against censoring model misspecification. This is a classical limitation when using IPCW methods: in particular, the same assumption on the censoring model is required for the standard Brier score (see Gerds and Schumacher, 2006). Another limitation concerns the independence assumption between censoring and covariate for the global variable importance test, in Theorem 3. Unfortunately, we were not able to extend this result to a more general dependence relation. The distribution of the test statistic relies on properties of Kaplan-Meier integrals and new theoretical results would need to be derived for a test statistic that would be based on a general estimator of the censoring cumulative distribution. This is left to future research.

9 Technical proofs

9.1 Proofs for Section 3

Proof of Theorem 1. The proof follows the outline of the proof of Theorem 5.2 in Gerds and Schumacher (2006). It essentially relies on the assumptions of weak consistency of the censoring estimator defined by Equation (5), and of $\hat{\mu}_{\tau,n} \in \mathcal{M}_{\mathcal{A}}$ and its limit $\tilde{\mu}_{\tau} \in \mathcal{M}_{\mathcal{A}}$ being bounded. \square

9.2 Proofs for Section 4

Let $\{\omega_i, i \in \mathcal{I}_2\}$ be the inverse probability censoring weights computed with the unknown function G , i.e. for all i

$$\omega_i = \frac{\mathbb{1}\{T_i \leq \tau\}\delta_i}{1 - G(T_i - | Z_i)} + \frac{\mathbb{1}\{T_i > \tau\}}{1 - G(\tau | Z_i)}.$$

Lemma 1. *Let (T, δ, Z) be an observation independent of D_{n_1} . Let*

$$\begin{aligned} f : \mathbb{R}^+ \times \mathbb{R}^d &\rightarrow \mathbb{R} \\ (t, z) &\mapsto f(t, z, D_{n_1}), \end{aligned}$$

such that $\mathbb{E}[f(T^ \wedge \tau, Z, D_{n_1}) | D_{n_1}]$ is almost surely finite. Then*

$$\mathbb{E}[f(T \wedge \tau, Z, D_{n_1})\omega | D_{n_1}] = \mathbb{E}[f(T^* \wedge \tau, Z, D_{n_1}) | D_{n_1}].$$

Proof. We have:

$$\begin{aligned} &\mathbb{E}[f(T \wedge \tau, Z, D_{n_1})\omega | D_{n_1}] \\ &= \mathbb{E}\left[\mathbb{1}\{T \leq \tau\} \frac{f(T \wedge \tau, Z, D_{n_1})\delta}{1 - G(T - | Z)} + \mathbb{1}\{T > \tau\} \frac{f(T \wedge \tau, Z, D_{n_1})}{1 - G(\tau | Z)} \mid D_{n_1}\right] \\ &= \mathbb{E}\left[\mathbb{1}\{T^* \leq \tau\} \frac{f(T^*, Z, D_{n_1})}{1 - G(T^* - | Z)} \mathbb{E}[\mathbb{1}\{T^* < C\} | T^*, Z, D_{n_1}] \mid D_{n_1}\right] \\ &\quad + \mathbb{E}\left[\mathbb{1}\{T^* > \tau\} \frac{f(\tau, Z, D_{n_1})}{1 - G(\tau | Z)} \mathbb{E}[\mathbb{1}\{C > \tau\} | T^*, Z, D_{n_1}] \mid D_{n_1}\right] \\ &= \mathbb{E}[\mathbb{1}\{T^* \leq \tau\} f(T^*, Z, D_{n_1}) + \mathbb{1}\{T^* > \tau\} f(\tau, Z, D_{n_1}) | D_{n_1}] \\ &= \mathbb{E}[f(T^* \wedge \tau, Z, D_{n_1}) | D_{n_1}], \end{aligned}$$

where in the third equality we used the independence assumption between C and T^* conditional on (Z, D_{n_1}) . \square

Lemma 2. *For (T^*, Z) independent of D_{n_1} , let f be a function verifying $\mathbb{E}[|f(T^* \wedge \tau, Z, D_{n_1})| | D_{n_1}] < \infty$ a.s. for all $\tau < \tau_H$. If \hat{G}_n is strongly consistent in the sense defined by (4) then*

$$\frac{1}{n_2} \sum_{i \in \mathcal{I}_2} f(T_i \wedge \tau, Z_i, D_{n_1}) \hat{\omega}_i \xrightarrow[n_2 \rightarrow \infty]{a.s.} \mathbb{E}[f(T^* \wedge \tau, Z, D_{n_1}) | D_{n_1}].$$

Proof.

$$\begin{aligned} \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} f(T_i \wedge \tau, Z_i, D_{n_1}) \hat{\omega}_i &= \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} f(T_i \wedge \tau, Z_i, D_{n_1}) \omega_i \\ &\quad + \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} f(T_i \wedge \tau, Z_i, D_{n_1}) (\hat{\omega}_i - \omega_i) \end{aligned}$$

As the $f(T_i \wedge \tau, Z_i, D_{n_1})$, $i \in \mathcal{I}_2$, are conditionally independent given D_{n_1} with the same conditional distribution, then by the strong conditional law of large numbers (see Theorem 4.2 in Majerek et al., 2005) and by Lemma 1,

$$\frac{1}{n_2} \sum_{i \in \mathcal{I}_2} f(T_i \wedge \tau, Z_i, D_{n_1}) \omega_i \xrightarrow[n_2 \rightarrow \infty]{a.s.} \mathbb{E}[f(T^* \wedge \tau, Z, D_{n_1}) \mid D_{n_1}].$$

On the other hand, as that there exists $\epsilon > 0$ such that $1 - \hat{G}_n(\tau \mid Z_i) > \epsilon$ for n high enough and $\tau < \tau_H$,

$$\begin{aligned} & \left| \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} f(T_i \wedge \tau, Z_i, D_{n_1}) (\hat{\omega}_i - \omega_i) \right| \\ &= \left| \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} f(T_i \wedge \tau, Z_i, D_{n_1}) \mathbb{1}\{T_i \leq \tau\} \delta_i \frac{G(T_i - \mid Z_i) - \hat{G}_n(T_i - \mid Z_i)}{(1 - \hat{G}_n(T_i - \mid Z_i))(1 - G(T_i - \mid Z_i))} \right. \\ & \quad \left. + \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} f(T_i \wedge \tau, Z_i, D_{n_1}) \mathbb{1}\{T_i > \tau\} \frac{G(\tau \mid Z_i) - \hat{G}_n(\tau \mid Z_i)}{(1 - \hat{G}_n(\tau \mid Z_i))(1 - G(\tau \mid Z_i))} \right| \\ &\leq \epsilon^{-1} \sup_{s \leq \tau, z \in \mathbb{R}^d} |\hat{G}_n(s \mid z) - G(s \mid z)| \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} |f(T_i \wedge \tau, Z_i, D_{n_1})| \omega_i \\ &\xrightarrow[n_2 \rightarrow \infty]{a.s.} 0, \end{aligned}$$

where the convergence to zero follows from the strong consistency of the censoring estimator defined by (4) and the law of large numbers for $|f(T_i \wedge \tau, Z_i, D_{n_1})| \omega_i$. \square

Proof of Theorem 2. Applying Lemma 2 to the functions $(u, z, D_{n_1}) \mapsto 1$ and $(u, z, D_{n_1}) \mapsto \mathbb{1}\{|u - \hat{\mu}_{\tau, n_1}(z)| \leq t\}$ for $t \in \mathbb{R}$ gives

$$\frac{\sum_{i \in \mathcal{I}_2} \hat{\omega}_i}{n_2} \xrightarrow[n_2 \rightarrow \infty]{a.s.} 1$$

and for all $t \in \mathbb{R}$,

$$\frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \mathbb{1}\{R_i \leq t\} \hat{\omega}_i \xrightarrow[n_2 \rightarrow \infty]{a.s.} \mathcal{R}_1^*(t) := \mathbb{P}(R^* \leq t \mid D_{n_1}).$$

Then for all $t \in \mathbb{R}$,

$$\hat{\mathcal{R}}_{n_2}^{\hat{G}}(t) = \frac{1}{\sum_{i \in \mathcal{I}_2} \hat{\omega}_i} \sum_{i \in \mathcal{I}_2} \mathbb{1}\{R_i \leq t\} \hat{\omega}_i \xrightarrow[n_2 \rightarrow \infty]{a.s.} \mathcal{R}_1^*(t).$$

For all $\alpha \in (0, 1)$, let $\hat{q}_{n_2}^\alpha = \inf\{t, \hat{\mathcal{R}}_{n_2}^{\hat{G}}(t) \geq 1 - \alpha\}$ and $q^\alpha = \inf\{t, \mathcal{R}_1^*(t) \geq 1 - \alpha\}$ be the $1 - \alpha$ quantiles of the cumulative distribution functions $\hat{\mathcal{R}}_{n_2}^{\hat{G}}$ and \mathcal{R}_1^* respectively. Then from Lemma 21.2 in Van Der Vaart (1998) we have that for all $\alpha \in (0, 1)$

$$\hat{q}_{n_2}^\alpha \xrightarrow[n_2 \rightarrow \infty]{a.s.} q^\alpha.$$

Finally, applying the continuous mapping theorem to the function \mathcal{R}_1^* gives

$$\mathcal{R}_1^*(\hat{q}_{n_2}^\alpha) \xrightarrow[n_2 \rightarrow \infty]{a.s.} \mathcal{R}_1^*(q^\alpha)$$

i.e.

$$\mathbb{P}(R^* \leq \hat{q}_{n_2}^\alpha \mid D_{n_1}) = \mathcal{R}_1^*(\hat{q}_{n_2}^\alpha) \xrightarrow[n_2 \rightarrow \infty]{a.s.} \mathcal{R}_1^*(q^\alpha) = \mathbb{P}(R^* \leq q^\alpha \mid D_{n_1}) \geq 1 - \alpha.$$

In particular, if the residuals have a continuous distribution given D_{n_1} , then $\mathbb{P}(R^* \leq q^\alpha \mid D_{n_1}) = 1 - \alpha$. \square

9.3 Proofs for Section 5

We provide here a slightly more general result than the one given by Theorem 3, where the function Φ_k can be any bounded function with support on $[0, \tau]$. For clarity, we state the theorem in its general form below.

Proposition 1. *Let $\tau < \tau_H$ and $\Phi_k(t, z) : (t, z) \in \mathbb{R}^+ \times \mathbb{R}^d \rightarrow \mathbb{R}^+$ be a uniformly bounded function with support $[0, \tau]$. Let*

$$\theta_k = \frac{1}{1 - S(\tau)} \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z).$$

Suppose that $C \perp\!\!\!\perp (T^, Z)$. Let $\hat{S}_{n_2} = 1 - \hat{F}_{n_2}$, and $1 - \hat{G}_{n_2}$ denote the Kaplan-Meier estimators of the functions $S = 1 - F$ and $1 - G$ computed on D_{n_2} . The censoring weights $\hat{\omega}_i$, $i \in \mathcal{I}_2$ are defined as in Equation (12). Let*

$$\hat{\theta}_k = \frac{1}{1 - \hat{S}_{n_2}(\tau)} \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) d\hat{F}_{n_2}(u, z) = \frac{1}{1 - \hat{S}_{n_2}(\tau)} \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \Phi_k(T_i, Z_i) \hat{\omega}_i$$

be an estimator of θ_k . Then

$$\sqrt{n_2} (\hat{\theta}_k - \theta_k) \xrightarrow[n_2 \rightarrow \infty]{} \mathcal{N}(0, \sigma^2(\Phi_k)) \text{ in distribution,}$$

where

$$\begin{aligned} \sigma^2(\Phi_k) = & \frac{1}{(1 - S(\tau))^2} \int_0^\tau \int_{\mathbb{R}^d} \left(\frac{\Phi_k(u, z)}{1 - G(u-)} - \frac{\bar{\Phi}_k(u)}{1 - H(u)} \right. \\ & \left. - \frac{S(\tau)}{1 - S(\tau)} \frac{\int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z)}{1 - H(u)} \right)^2 (1 - G(u-)) dF(u, z), \end{aligned}$$

with $\bar{\Phi}_k(t) = \int_t^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z)$. Moreover, the estimator

$$\begin{aligned} \hat{\sigma}^2(\Phi_k) = & \frac{(1 - \hat{S}_{n_2}(\tau))^{-2}}{n_2} \sum_{i \in \mathcal{I}_2} \mathbb{1}\{T_i \leq \tau\} \left(\Phi_k(T_i, Z_i) \hat{\omega}_i \right. \\ & \left. - \frac{\delta_i}{\hat{Y}_{n_2}(T_i)} \frac{1}{n_2} \sum_{j \in \mathcal{I}_2} \left(\mathbb{1}\{T_i \leq T_j\} - \frac{\hat{S}_{n_2}(\tau)}{1 - \hat{S}_{n_2}(\tau)} \right) \Phi_k(T_j, Z_j) \hat{\omega}_j \right)^2, \end{aligned}$$

where $\hat{Y}_{n_2}(t) = \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \mathbb{1}\{T_i \geq t\}$, is consistent for the variance $\sigma^2(\Phi_k)$.

Proof. For all $i \in \mathcal{I}_2$, we introduce the martingale residuals

$$\begin{aligned} M_i^F(t) &= N_i(t) - \int_0^t \mathbf{1}\{T_i \geq u\} \frac{dF(u)}{1 - F(u-)}, \text{ with } N_i(t) = \mathbf{1}\{T_i \leq t, \delta_i = 1\}, \\ M_i^G(t) &= N_i^G(t) - \int_0^t \mathbf{1}\{T_i \geq u\} \frac{dG(u)}{1 - G(u-)}, \text{ with } N_i^G(t) = \mathbf{1}\{T_i \leq t, \delta_i = 0\}, \\ M_+^X(t) &= \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} M_i^X(t) \text{ for } X = F, G. \end{aligned}$$

First note that the weights $\hat{\omega}_i, i \in \mathcal{I}_2$ defined in Equation (12) simplify into $\hat{\omega}_i = \delta_i / (1 - \hat{G}_{n_2}(T_i -))$ because Φ_k has support $[0, \tau]$. Then

$$\begin{aligned} & \sqrt{n_2}(\hat{\theta}_k - \theta_k) \\ &= \sqrt{n_2} \left(\frac{\int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) d\hat{F}_{n_2}(u, z)}{1 - \hat{S}_{n_2}(\tau)} - \frac{\int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z)}{1 - S(\tau)} \right) \\ &= \frac{\sqrt{n_2}}{1 - S(\tau)} \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) d(\hat{F}_{n_2} - F)(u, z) \\ &\quad + \sqrt{n_2} \frac{\hat{S}_{n_2}(\tau) - S(\tau)}{(1 - S(\tau))(1 - \hat{S}_{n_2}(\tau))} \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z) \\ &\quad + \sqrt{n_2} \frac{\hat{S}_{n_2}(\tau) - S(\tau)}{(1 - S(\tau))(1 - \hat{S}_{n_2}(\tau))} \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) d(\hat{F}_{n_2} - F)(u, z) \\ &=: A + B + C. \end{aligned}$$

From Lopez (2007), A is a Kaplan-Meier integral and we have

$$\begin{aligned} A &= \frac{\sqrt{n_2}}{1 - S(\tau)} \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \left(\Phi_k(T_i, Z_i) \frac{\delta_i \mathbf{1}\{T_i \leq \tau\}}{1 - G(T_i -)} - \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z) + \gamma_1(\Phi_k; T_i, \delta_i) \right) \\ &\quad + O_{\mathbb{P}}(n_2^{-1/2}), \end{aligned}$$

where for all $i \in \mathcal{I}_2$,

$$\gamma_1(\Phi_k; T_i, \delta_i) = \int_0^\tau \frac{\bar{\Phi}_k(y) dM_i^G(y)}{1 - H(y)}.$$

Moreover, using the martingale decomposition of the Kaplan-Meier estimator $\hat{S}_{n_2}(\tau) - S(\tau) = -S(\tau) \int_0^\tau (dM_+^F(u)) / (1 - H(u)) + o_{\mathbb{P}}(n_2^{-1/2})$ (see for instance Andersen et al., 1993), we can rewrite B as

$$B = -\sqrt{n_2} \frac{S(\tau)}{(1 - S(\tau))^2} \int_0^\tau \frac{dM_+^F(u)}{1 - H(u)} \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z) + o_{\mathbb{P}}(1).$$

Finally, the C term tends to 0 in probability as n_2 tends to infinity since $\sqrt{n_2} \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) d(\hat{F}_{n_2} - F)(u, z) = O_{\mathbb{P}}(1)$ (see Proposition 2.3.1 in Lopez, 2007) and $\hat{S}_{n_2}(\tau) - S(\tau) = o_{\mathbb{P}}(1)$ from the consistency of the Kaplan-Meier estimator. Finally, we obtain the following centered and i.i.d

representation of $\sqrt{n_2}(\hat{\theta}_k - \theta_k)$:

$$\begin{aligned}\sqrt{n_2}(\hat{\theta}_k - \theta_k) = & \sqrt{n_2} \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \frac{1}{1 - S(\tau)} \left(\Phi_k(T_i, Z_i) \frac{\delta_i \mathbb{1}\{T_i \leq \tau\}}{1 - G(T_i -)} \right. \\ & - \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z) + \gamma_1(\Phi_k; T_i, \delta_i) \\ & \left. - \frac{S(\tau)}{1 - S(\tau)} \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z) \int_0^\tau \frac{dM_i^F(u)}{1 - H(u)} \right) + o_{\mathbb{P}}(1).\end{aligned}$$

The convergence in law follows from the central limit theorem.

Next, we compute the asymptotic variance, denoted $\sigma^2(\Phi_k)$, in the following way:

$$\begin{aligned}(1 - S(\tau))^2 \sigma^2(\Phi_k) &= \text{Var} \left(\Phi_k(T, Z) \frac{\delta \mathbb{1}\{T \leq \tau\}}{1 - G(T -)} + \gamma_1(\Phi_k; T, \delta) \right) \\ &+ \text{Var} \left(\frac{S(\tau)}{1 - S(\tau)} \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z) \int_0^\tau \frac{dM^F(u)}{1 - H(u)} \right) \\ &+ 2 \text{Cov} \left(\Phi_k(T, Z) \frac{\delta \mathbb{1}\{T \leq \tau\}}{1 - G(T -)} + \gamma_1(\Phi_k; T, \delta), \frac{S(\tau)}{1 - S(\tau)} \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z) \int_0^\tau \frac{dM^F(u)}{1 - H(u)} \right) \\ &=: D + E + 2F.\end{aligned}$$

First,

$$D = \int_0^\tau \int_{\mathbb{R}^d} \frac{\left(\Phi_k(u, z) - \frac{\bar{\Phi}_k(u)}{1 - F(u)} \right)^2}{1 - G(u -)} dF(u, z)$$

by Proposition 2.3.1 in Lopez (2007). Second, using Theorem 2.4.4 in Fleming and Harrington (2005),

$$E = \frac{S(\tau)^2}{(1 - S(\tau))^2} \left(\int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z) \right)^2 \int_0^\tau \frac{dF(u)}{(1 - F(u))(1 - H(u))}.$$

Third,

$$\begin{aligned}F &= \frac{S(\tau)}{1 - S(\tau)} \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z) \\ &\times \mathbb{E} \left[\left(\Phi_k(T, Z) \frac{\delta \mathbb{1}\{T \leq \tau\}}{1 - G(T -)} + \int_0^\tau \frac{\bar{\Phi}_k(y) dM^G(y)}{1 - H(y)} \right) \int_0^\tau \frac{dM^F(u)}{1 - H(u)} \right].\end{aligned}$$

We then study the F term. From the definition of M^F and the Fubini theorem, we have

$$\begin{aligned}\mathbb{E} \left[\Phi_k(T, Z) \frac{\delta \mathbb{1}\{T \leq \tau\}}{1 - G(T -)} \int_0^\tau \frac{dM^F(u)}{1 - H(u)} \right] &= \int_0^\tau \int_{\mathbb{R}^d} \frac{\Phi_k(u, z)}{1 - H(u)} dF(u, z) - \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) \int_0^\tau \frac{\mathbb{1}\{u \geq y\}}{1 - H(y)} \frac{dF(y)}{1 - F(y)} dF(u, z) \\ &= \int_0^\tau \int_{\mathbb{R}^d} \frac{\Phi_k(u, z)}{1 - H(u)} dF(u, z) - \int_0^\tau \frac{\bar{\Phi}_k(y)}{1 - H(y)} \frac{dF(y)}{1 - F(y)}.\end{aligned}$$

In addition, using the fact that N^G and N^F cannot jump at the same time, we have

$$\begin{aligned}
& \mathbb{E} \left[\int_0^\tau \frac{\bar{\Phi}_k(y) dM^G(y)}{1-H(y)} \int_0^\tau \frac{dM^F(u)}{1-H(u)} \right] \\
&= \mathbb{E} \left[\int_0^\tau \int_0^\tau \frac{\bar{\Phi}_k(y)}{1-H(y)} \frac{1}{1-H(u)} \left(-dN^G(y) \mathbb{1}\{T \geq u\} \frac{dF(u)}{1-F(u-)} \right. \right. \\
&\quad \left. \left. - dN^F(u) \mathbb{1}\{T \geq y\} \frac{dG(y)}{1-G(y-)} + \mathbb{1}\{T \geq u \vee y\} \frac{dF(u)}{1-F(u-)} \frac{dG(y)}{1-G(y-)} \right) \right] \\
&= -\mathbb{E} \left[\int_0^\tau \frac{\bar{\Phi}_k(C)}{1-H(C)} \frac{(1-\delta) \mathbb{1}\{C \leq \tau\}}{1-H(u)} \mathbb{1}\{C \geq u\} \frac{dF(u)}{1-F(u-)} \right] \\
&\quad - \mathbb{E} \left[\int_0^\tau \frac{\bar{\Phi}_k(y)}{1-H(y)} \frac{\delta \mathbb{1}\{T^* \leq \tau\}}{1-H(T^*)} \mathbb{1}\{T^* \geq y\} \frac{dG(y)}{1-G(y-)} \right] \\
&\quad + \int_0^\tau \int_0^\tau \frac{\bar{\Phi}_k(y)}{1-H(y)} \frac{1-H(u \vee y)}{1-H(u)} \frac{dF(u)}{1-F(u-)} \frac{dG(y)}{1-G(y-)} .
\end{aligned}$$

We now take the conditional expectation with respect to C for the first term, with respect to T^* for the second term and we use the relations $\mathbb{E}[1-\delta \mid C] = 1-F(C)$, $\mathbb{E}[\delta \mid T^*] = 1-G(T^*-)$ in order to obtain the following result

$$\begin{aligned}
\mathbb{E} \left[\int_0^\tau \frac{\bar{\Phi}_k(y) dM^G(y)}{1-H(y)} \int_0^\tau \frac{dM^F(u)}{1-H(u)} \right] &= - \int_0^\tau \int_0^\tau \frac{\bar{\Phi}_k(y)}{1-H(u)} \mathbb{1}\{y \geq u\} \frac{dF(u)}{1-F(u-)} \frac{dG(y)}{1-G(y-)} \\
&\quad - \int_0^\tau \int_0^\tau \frac{\bar{\Phi}_k(y)}{1-H(y)} \mathbb{1}\{u \geq y\} \frac{dF(u)}{1-F(u-)} \frac{dG(y)}{1-G(y-)} \\
&\quad + \int_0^\tau \int_0^\tau \frac{\bar{\Phi}_k(y)}{1-H(y)} \frac{1-H(u \vee y)}{1-H(u)} \frac{dF(u)}{1-F(u-)} \frac{dG(y)}{1-G(y-)} \\
&= 0.
\end{aligned}$$

This proves that $\mathbb{E}[\gamma_1(\Phi_k; T, \delta) \int_0^\tau (dM^F(u))/(1-H(u))] = 0$. To conclude, we have shown that

$$F = \frac{S(\tau)}{1-S(\tau)} \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z) \left(\int_0^\tau \int_{\mathbb{R}^d} \frac{\Phi_k(u, z)}{1-H(u)} dF(u, z) - \int_0^\tau \frac{\bar{\Phi}_k(y)}{1-H(y)} \frac{dF(y)}{1-F(y)} \right).$$

Finally,

$$\begin{aligned}
& (1-S(\tau))^2 \sigma^2(\Phi_k) \\
&= \int_0^\tau \int_{\mathbb{R}^d} \frac{\left(\Phi_k(u, z) - \frac{\bar{\Phi}_k(u)}{1-F(u)} \right)^2}{1-G(u-)} dF(u, z) \\
&\quad + \frac{S(\tau)^2}{(1-S(\tau))^2} \left(\int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z) \right)^2 \int_0^\tau \frac{dF(u)}{(1-F(u))(1-H(u))} \\
&\quad + \frac{2 S(\tau)}{1-S(\tau)} \int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z) \left(\int_0^\tau \int_{\mathbb{R}^d} \frac{\Phi_k(u, z)}{1-H(u)} dF(u, z) - \int_0^\tau \frac{\bar{\Phi}_k(y)}{1-H(y)} \frac{dF(y)}{1-F(y)} \right) \\
&= \int_0^\tau \int_{\mathbb{R}^d} \left(\frac{\Phi_k(u, z)}{1-G(u-)} - \frac{\bar{\Phi}_k(u)}{1-H(u)} - \frac{S(\tau)}{1-S(\tau)} \frac{\int_0^\tau \int_{\mathbb{R}^d} \Phi_k(u, z) dF(u, z)}{1-H(u)} \right)^2 (1-G(u-)) dF(u, z).
\end{aligned}$$

The consistency of the variance estimator follows from the consistency of the Kaplan-Meier integrals (see Lopez, 2007), the consistency of \hat{S} and the consistency of \hat{Y}_{n_2} . \square

References

- Andersen, Per Kragh, Ørnulf Borgan, Richard D. Gill, and Niels Keiding (1993). *Statistical Models Based on Counting Processes*. Springer Series in Statistics. New York, NY: Springer US.
- Andersen, Per Kragh, Mette Gerster Hansen, and John P. Klein (2004). “Regression Analysis of Restricted Mean Survival Time Based on Pseudo-Observations”. In: *Lifetime Data Analysis* 10.4, pp. 335–350.
- Blanche, Paul, Michael W. Kattan, and Thomas A. Gerds (2019). “The c-index is not proper for the evaluation of t -year predicted risks”. In: *Biostatistics* 20.2, pp. 347–357.
- Boström, Henrik, Ulf Johansson, and Anders Vesterberg (2019). “Predicting with Confidence from Survival Data”. In: *Proceedings of the Eighth Symposium on Conformal and Probabilistic Prediction and Applications*. PMLR, pp. 123–141.
- Breiman, Leo (2001). “Random Forests”. In: *Machine Learning* 45.1, pp. 5–32.
- Candès, Emmanuel, Lihua Lei, and Zhimei Ren (2023). “Conformalized survival analysis”. In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 85.1, pp. 24–45.
- Chen, George H. (2020). “Deep Kernel Survival Analysis and Subject-Specific Survival Time Prediction Intervals”. In: *Proceedings of the 5th Machine Learning for Healthcare Conference*. PMLR, pp. 537–565.
- Cox, David R. (1975). “Partial Likelihood”. In: *Biometrika* 62.2, pp. 269–276.
- DiCiccio, Cyrus J., Thomas J. DiCiccio, and Joseph P. Romano (2020). “Exact tests via multiple data splitting”. In: *Statistics & Probability Letters* 166, p. 108865.
- Eaton, Anne, Terry Therneau, and Jennifer Le-Rademacher (2020). “Designing clinical trials with (restricted) mean survival time endpoint: Practical considerations”. In: *Clinical Trials* 17.3, pp. 285–294.
- Fleming, Thomas R. and David P. Harrington (2005). *Counting Processes and Survival Analysis*. Wiley Series in Probability and Statistics. Hoboken, NJ, USA: John Wiley & Sons, Inc.
- Gerds, Thomas A., Michael W. Kattan, Martin Schumacher, and Changhong Yu (2013). “Estimating a time-dependent concordance index for survival prediction models with covariate dependent censoring”. In: *Statistics in Medicine* 32.13, pp. 2173–2184.
- Gerds, Thomas A. and Martin Schumacher (2006). “Consistent Estimation of the Expected Brier Score in General Survival Models with Right-Censored Event Times”. In: *Biometrical Journal* 48.6, pp. 1029–1040.
- Heagerty, Patrick J. and Yingye Zheng (2005). “Survival Model Predictive Accuracy and ROC Curves”. In: *Biometrics* 61.1, pp. 92–105.
- Ishwaran, Hemant, Udaya B. Kogalur, Eugene H. Blackstone, and Michael S. Lauer (2008). “Random survival forests”. In: *The Annals of Applied Statistics* 2.3.
- Karrison, Theodore (1987). “Restricted Mean Life With Adjustment for Covariates”. In: *Journal of the American Statistical Association* 82.400, pp. 1169–1176.
- Kojima, Masahiro (2022). *Variable Selection using Inverse Survival Probability Weighting*.
- Kovalev, Maxim S., Lev V. Utkin, and Ernest M. Kasimov (2020). “SurvLIME: A method for explaining machine learning survival models”. In: *Knowledge-Based Systems* 203, p. 106164.
- Krzyżiński, Mateusz, Mikołaj Spytek, Hubert Baniecki, and Przemysław Biecek (2023). “SurvSHAP(t): Time-dependent explanations of machine learning survival models”. In: *Knowledge-Based Systems* 262, p. 110234.
- Kuchibhotla, Arun Kumar (2021). *Exchangeability, Conformal Prediction, and Rank Tests*.
- Lei, Jing, Max G’Sell, Alessandro Rinaldo, Ryan J. Tibshirani, and Larry Wasserman (2018). “Distribution-Free Predictive Inference for Regression”. In: *Journal of the American Statistical Association* 113.523, pp. 1094–1111.

- Lopez, Olivier (2007). “Réduction de dimension en présence de données censurées”. PhD thesis. ENSAE ParisTech.
- Majerek, Dariusz, Wioletta Nowak, and Wies Ziba (2005). “Conditional strong law of large number”. In: *International Journal of Pure and Applied Mathematics* 20.
- Schumacher, M., G. Bastert, H. Bojar, K. Hübner, M. Olschewski, W. Sauerbrei, C. Schmoor, C. Beyerle, R. L. Neumann, and H. F. Rauschecker (1994). “Randomized 2 x 2 trial evaluating hormonal treatment and the duration of chemotherapy in node-positive breast cancer patients. German Breast Cancer Study Group.” In: *Journal of Clinical Oncology* 12.10, pp. 2086–2093.
- Teng, Jiaye, Zeren Tan, and Yang Yuan (2021). “T-SCI: A Two-Stage Conformal Inference Algorithm with Guaranteed Coverage for Cox-MLP”. In: *Proceedings of the 38th International Conference on Machine Learning*. PMLR, pp. 10203–10213.
- Tian, Lu, Lihui Zhao, and L. J. Wei (2014). “Predicting the restricted mean event time with the subject’s baseline covariates in survival analysis”. In: *Biostatistics* 15.2, pp. 222–233.
- Van Der Vaart, Aad W. (1998). *Asymptotic statistics*. Cambridge series in statistical and probabilistic mathematics. Cambridge, UK: Cambridge University Press.
- Vovk, Vladimir, A. Gammerman, and Glenn Shafer (2005). *Algorithmic learning in a random world*. New York: Springer.
- Wang, Xin and Douglas E. Schaubel (2018). “Modeling restricted mean survival time under general censoring mechanisms”. In: *Lifetime Data Analysis* 24.1, pp. 176–199.
- Zhao, Lili (2021). “Deep neural networks for predicting restricted mean survival times”. In: *Bioinformatics* 36.24, pp. 5672–5677.
- Zucker, David M. (1998). “Restricted Mean Life with Covariates: Modification and Extension of a Useful Survival Analysis Method”. In: *Journal of the American Statistical Association* 93.442, pp. 702–709.