



Estimation of time-specific intervention effects on continuously distributed time-to-event outcomes by targeted maximum likelihood estimation

Rytgaard, Helene Charlotte Wiese; Eriksson, Frank; Laan, Mark van der

Published in:
arXiv

Publication date:
2021

Citation for published version (APA):
Rytgaard, H. C. W., Eriksson, F., & Laan, M. V. D. (2021). Estimation of time-specific intervention effects on continuously distributed time-to-event outcomes by targeted maximum likelihood estimation. *arXiv*.

Estimation of time-specific intervention effects on continuously distributed time-to-event outcomes by targeted maximum likelihood estimation

Helene C. W. Rytgaard^{1,*}, F. Eriksson¹, and Mark J. van der Laan²

¹Section of Biostatistics, University of Copenhagen, Denmark

²Devison of Biostatistics, University of California, Berkeley

June 22, 2021

Abstract

Targeted maximum likelihood estimation is a general methodology combining flexible ensemble learning and semiparametric efficiency theory in a two-step procedure for estimation of causal parameters. Proposed targeted maximum likelihood procedures for survival and competing risks analysis have so far focused on events taken values in discrete time. We here present a targeted maximum likelihood estimation procedure for event times that take values in \mathbb{R}_+ . We focus on the estimation of intervention-specific mean outcomes with stochastic interventions on a time-fixed treatment. For data-adaptive estimation of nuisance parameters, we propose a new flexible highly adaptive lasso estimation method for continuous-time intensities that can be implemented with L_1 -penalized Poisson regression. In a simulation study the targeted maximum likelihood estimator based on the highly adaptive lasso estimator proves to be unbiased and achieve proper coverage in agreement with the asymptotic theory and further displays efficiency improvements relative to a Kaplan-Meier approach.

Keywords: Targeted maximum likelihood estimation, survival analysis, treatment effects, semi-parametric model, efficient estimation, causal inference

1 Introduction

In recent years, semiparametric efficient and doubly robust estimators (van der Laan and Robins, 2003; Tsiatis, 2007) have gained great popularity in many fields and are increasingly used to draw inference about the effect of a treatment (Robins, 1986; Hernan and Robins, 2020). In survival analysis, such methods provide an alternative to the widely used Cox regression model (Cox, 1972) that relies on the assumption of proportional hazards and requires correct specification of main effects and interaction effects of treatment and confounders. We consider a standard survival analysis setting where a population of subjects are followed over a period of time until an event of interest occurs (Andersen et al., 1993). We suppose that covariates are measured and a treatment decision, either randomized or conditional on covariates, is made at the beginning of the follow-up period, and that we are interested in the effect of the treatment decision on the time until the event of interest happens. The data are further characterized by right-censoring, saying that a subject is right-censored if the event of interest did not occur within the subject-specific follow-up period. In this work we target parameters of the g-computation formula that have a clear causal interpretation under a

set of structural assumptions (Robins, 1986; Gill and Robins, 2001; van der Laan and Robins, 2003; Hernan and Robins, 2020) and consider the general case with stochastic interventions (Robins et al., 2004; Dawid and Didelez, 2010; Gill and Robins, 2001, Sections 6 and 7) on the treatment decision on time-to-event outcomes; an important special case is the average treatment effect on the τ -year risk of death.

Our overall goal for estimation is to impose as few assumptions as possible on the data-generating mechanism while providing efficient and double robust estimation by optimizing the estimation procedure for the target parameter specifically. Targeted maximum likelihood estimation (van der Laan and Rubin, 2006; van der Laan and Rose, 2011, 2018) provides a general methodology that combines cross-validated machine learning (van der Vaart et al., 2006; van der Laan et al., 2006; van der Laan et al., 2007) and semiparametric efficiency theory (Bickel et al., 1993) for constructing asymptotically efficient estimators for low-dimensional parameters in infinite-dimensional models. Targeted maximum likelihood methods have been extensively developed for treatment effect estimation in survival analysis settings where events are observed on a discrete time-scale (Moore et al., 2009a,b,c; Stitelman et al., 2011a,b; van der Laan and Gruber, 2012; Benkeser et al., 2018; Cai et al., 2019), but requires artificial discretization to be applied to events observed in continuous time and may lead to instability and high memory usage (Sofrygin et al., 2019). We propose a generalization of existing targeted maximum likelihood methods to continuously measured time-to-event outcomes for survival and competing risks analysis based on a specialization of the work by Rytgaard et al. (2020), and, further, derive exact expressions for the second-order remainders to establish double robustness properties of the considered survival and competing risks estimation problems.

Generally, targeted maximum likelihood estimation proceeds in two steps. Initial estimators are constructed for the nuisance parameters in the first step which is then followed by a second step (the targeting step) that reduces bias for the initial estimators, improves precision and ensures reliable statistical inference in terms of confidence intervals and p-values. We suggest a targeting step based on a proportional hazards type submodel for the intensity of the event process, iteratively updating the hazard estimators carried out in a smoothed fashion across time until the optimal score equation is solved. We further construct a flexible highly adaptive lasso estimator (van der Laan, 2017) for continuous-time hazards that can be implemented with L_1 -penalized regression (Tibshirani, 1996) by utilizing Poisson regression modeling techniques (Andersen et al., 1993; Lindsey, 1995). The highly adaptive lasso is a nonparametric estimation method proven to converge faster than $n^{-1/4}$ with respect to the square-root of the loss-based dissimilarity under the only assumption that the true function is càdlàg (right-continuous with left limits) with a finite sectional variation norm (van der Laan, 2017).

Our estimation methods have applications in observational as well as experimental data settings such as randomized clinical trials. In trial settings, it is standard to analyze time-to-event data in each intervention arm using a Kaplan-Meier estimator. It is well-known that the Kaplan-Meier estimator, an unadjusted estimator, may yield biased estimation under covariate dependent censoring, and may, furthermore, be inefficient in presence of predictive covariates (Lu et al., 2008; Rubin et al., 2008; Moore et al., 2009c; Rotnitzky et al., 1997; Diaz et al., 2019). We revisit these issues in our simulation study, where we demonstrate that our targeted maximum likelihood estimator based on the highly adaptive lasso estimator improves the precision and robustness of findings while providing accurate confidence intervals without any parametric model specification.

2 Setting and notation

Consider unit-specific observed data on the form $O = (L, A, \tilde{T}, \Delta)$. These data contain a vector of information on pre-treatment characteristics ($L \in \mathcal{L} = \mathbb{R}^d$), tell us what treatment the unit was exposed to at study entry ($A \in \mathcal{A} = \{0, 1\}$), how long the unit was under observation ($\tilde{T} \in \mathbb{R}_+$) and if the event of interest was observed for the unit or not ($\Delta \in \{0, 1\}$). We introduce the variables $T \in \mathbb{R}_+$ and $C \in \mathbb{R}_+$ as the times to event and censoring, respectively, such that the observed time is $\tilde{T} = \min(T, C)$ and the event indicator can be written $\Delta = \mathbb{1}\{T \leq C\}$. The corresponding competing risks setting is treated in Appendix A. Suppose that we observe a sample $\{O_i\}_{i=1}^n$ of independent and identically distributed observations of $O \sim P_0$. Let $N(t) = \mathbb{1}\{\tilde{T} \leq t, \Delta = 1\}$ and $N^c(t) = \mathbb{1}\{\tilde{T} \leq t, \Delta = 0\}$ denote the observed counting processes (Andersen et al., 1993). The hazards for the distributions of T and C are defined by

$$\begin{aligned}\lambda_0(t | a, \ell) &= \lim_{h \rightarrow 0} h^{-1} P(T \leq t + h | \tilde{T} \geq t, A = a, L = \ell), \\ \lambda_0^c(t | a, \ell) &= \lim_{h \rightarrow 0} h^{-1} P(C \leq t + h | \tilde{T} \geq t, A = a, L = \ell);\end{aligned}$$

let Λ_0, Λ_0^c denote the corresponding cumulative hazards; thus, the compensators of N, N^c are characterized by

$$\mathbb{E}_{P_0}[N(dt) | \mathcal{F}_{t-}] = \mathbb{1}\{\tilde{T} \geq t\} \Lambda_0(dt | A, L), \text{ and, } \mathbb{E}_{P_0}[N^c(dt) | \mathcal{F}_{t-}] = \mathbb{1}\{\tilde{T} \geq t\} \Lambda_0^c(dt | A, L),$$

where $\mathcal{F}_t = \sigma((N(s), N^c(s)) : s \leq t, A, L)$ is the filtration generated by the observed data in $[0, t]$. Let further μ_0 be the density of L with respect to an appropriate dominating measure ν and $\pi_0(\cdot | L)$ be the conditional distribution of A given L . Under coarsening at random (van der Laan and Robins, 2003), λ, λ^c , the distribution for the observed data can now be represented as $dP_0(o) = p_0(o) d\nu(\ell) dt$ where the density p_0 is given by:

$$p_0(o) = \mu_0(\ell) \pi_0(a | \ell) (\lambda_0(t | a, \ell))^\delta S_0(t- | a, \ell) (\lambda_0^c(t | a, \ell))^{1-\delta} S_0^c(t- | a, \ell); \quad (1)$$

here $o = (\ell, a, t, \delta)$ and S_0, S_0^c denote

$$S_0(t | a, \ell) = \prod_{s \leq t} (1 - \Lambda_0(ds | a, \ell)) = \exp\left(-\int_0^t \lambda_0(s | a, \ell) ds\right), \quad (2)$$

$$S_0^c(t | a, \ell) = \prod_{s \leq t} (1 - \Lambda_0^c(ds | a, \ell)) = \exp\left(-\int_0^t \lambda_0^c(s | a, \ell) ds\right), \quad (3)$$

the survival functions for T and C , respectively. In the above display, \prod denotes the product integral (Gill and Johansen, 1990; Andersen et al., 1993).

3 Target of estimation

We formulate our target parameter under hypothetical interventions on the baseline treatment decision governed by π_0 and the censoring mechanism governed by λ_0^c . For this purpose, we define:

$$g_0(o) = \pi_0(a | \ell) (\lambda_0^c(t | a, \ell))^{1-\delta} S_0^c(t- | a, \ell), \quad (4)$$

$$q_0(o) = \mu_0(\ell) (\lambda_0(t | a, \ell))^\delta S_0(t- | a, \ell). \quad (5)$$

We refer to (4) as the *interventional* and to (5) as the *non-interventional* part of the likelihood. The distribution P_0 that factorizes as in (1) can now be parametrized by (4)–(5): We write P_{q_0, g_0} . We consider the statistical model \mathcal{M} for P_0 consisting of distributions that admit an equivalent parametrization:

$$\mathcal{M} = \left\{ P : P = P_{q, g}, q \in \mathcal{Q}, g \in \mathcal{G} \right\},$$

with parameter space \mathcal{Q} for the non-interventional part and \mathcal{G} for the interventional part. An intervention $g \mapsto g^*$ is imposed on $P \in \mathcal{M}$ by substituting g^* for g , so that P_{q, g^*} defines the post-interventional distribution, also known as the g -computation formula (Robins, 1986). Such an intervention g^* can be specified as follows:

$$g^*(o) = \delta \pi^*(a \mid \ell), \quad (6)$$

imposing the treatment decision π^* at baseline and no censoring throughout the follow-up period. To define the average treatment effect, for example, we specify the interventions g^{a^*} , $a^* = 0, 1$, where

$$g^{a^*}(o) = \delta \mathbb{1}\{a = a^*\}, \quad (7)$$

i.e., π^* is a degenerate distribution that puts all mass in $A = a^*$, $a^* = 0, 1$.

With an intervention g^* and corresponding post-interventional distribution P_{q, g^*} as defined in Section 3, we now define our target parameter $\Psi_\tau : \mathcal{M} \rightarrow \mathbb{R}$ as the intervention-specific mean outcome:

$$\Psi_\tau(P) = \mathbb{E}_{P_{q, g^*}}[\mathbb{1}\{\tilde{T} \leq \tau\}] = 1 - \int_{\mathcal{L}} \sum_{a=0,1} S(\tau \mid a, \ell) \pi^*(a \mid \ell) \mu(\ell) d\nu(\ell), \quad (8)$$

under imposing the hypothetical intervention g^* on the distribution $P \in \mathcal{M}$. We denote by $\psi_0 = \Psi_\tau(P_0)$ the true value. Under causal assumptions reviewed in Remark 1 below, the target parameter can be interpreted as the risk we would observe had subjects been treated according to the treatment strategy π^* . Importantly, the average treatment effect can be defined in terms of a contrast between two intervention-specific mean outcomes, namely:

$$\Psi_\tau^{\text{ATE}}(P) = \Psi_\tau^1(P) - \Psi_\tau^0(P), \quad \text{where,} \quad \Psi_\tau^a(P) = \mathbb{E}_{P_{q, g^a}}[\mathbb{1}\{\tilde{T} \leq \tau\}], \quad \text{for } a = 0, 1, \quad (9)$$

which, under the causal assumptions listed below, is interpreted as the average treatment effect on the τ -year risk had we randomized subjects to treatment or no treatment.

Remark 1 (Causal interpretation) Define the hypothetical event time T^{g^*} had the subject been treated according to the treatment strategy π^* , irrespective of what it actually was. Particularly, let T^a be the hypothetical event time that would result if the treatment decision had been $A = a$. We formulate the causal assumptions as follows.

Assumption 1 (Consistency) $T = T^a$ on the event that $A = a$, for $a = 0, 1$.

Assumption 2 (Conditionally independent censoring) $T \perp\!\!\!\perp C \mid A, L$.

Assumption 3 (No unmeasured confounding) $T^a \perp\!\!\!\perp A \mid L$, for $a = 0, 1$.

Assumption 4 (Positivity) $P(C \geq \tau \mid a, L) \pi(a \mid L) > \eta > 0$, a.e., for $a = 0, 1$.

Under Assumptions 1–4, the target parameter can be interpreted as the τ -year risk we would observe had we in fact imposed the intervention π^* in the real world, i.e.,

$$\Psi_\tau(P) = P(T^g \leq \tau). \quad (10)$$

Particularly, the average treatment effect, defined by Equation (9), is equal to $\Psi_\tau^{\text{ATE}}(P) = P(T^1 \leq \tau) - P(T^0 \leq \tau)$.

We show in the supplementary material that (10) follows under Assumptions 1–3.

4 Efficient estimation

An estimator $\hat{\psi}_n^* = \Psi_\tau(\hat{P}_n^*)$ for the target parameter ψ_0 is obtained by providing an estimator \hat{P}_n^* for (the relevant components of) P_0 and plugging this into the parameter mapping (8). The target parameter only depends on P through λ and the marginal distribution μ of L . We will estimate the latter by the empirical mean so that an estimator for ψ_0 can be based solely on an estimator $\hat{\lambda}$ for λ . Particularly, an estimator $\hat{\lambda}_n$ is mapped to an estimator \hat{S}_n for the conditional survival function by evaluating Equation (2), and next to an estimator for the target parameter by:

$$1 - \frac{1}{n} \sum_{i=1}^n \left(\sum_{a=0,1} \hat{S}_n(\tau | a, L_i) \pi^*(a | L_i) \right). \quad (11)$$

Efficient estimation, on the other hand, also requires estimation of the interventional part of the likelihood. Here, and throughout, \mathbb{P}_n denotes the empirical distribution of the data $\{O_i\}_{i=1}^n$ and $o_P(1)$ is a term that converges to zero in probability. An estimator $\hat{\psi}_n$ for the target parameter is asymptotically linear with influence curve equal to the efficient influence curve $D_\tau^*(P_0)$ (Bickel et al., 1993; van der Vaart, 2000) if and only if

$$\sqrt{n}(\hat{\psi}_n - \psi_0) = \sqrt{n} \mathbb{P}_n D_\tau^*(P_0) + o_P(1). \quad (12)$$

The efficient influence function $D_\tau^*(P)$ for our statistical model \mathcal{M} and target parameter $\Psi_\tau : \mathcal{M} \rightarrow \mathbb{R}$ is well-known from the literature (Robins et al., 1992; van der Laan and Robins, 2003; Moore et al., 2009a,b; van der Laan and Rose, 2011). A sketch of its derivation can be found in the supplementary material; we present an expression for $D_\tau^*(P)$ in Section 5. Specifically, the efficient influence function is a mapping of $P \in \mathcal{M}$ through π , λ^c and λ , thus comprising the nuisance parameters for our estimation problem.

4.1 Conditions for asymptotically linear and efficient estimation

Below we give sufficient conditions for establishing (12); the proof follows similar work (see, e.g., van der Laan and Rubin, 2006; van der Laan, 2017; van der Laan and Rose, 2011, Theorem A.5) but is included for completeness in the supplementary material. Define the second-order remainder:

$$R(P, P_0) := \Psi_\tau(P) - \Psi_\tau(P_0) + P_0 D_\tau^*(P), \quad P \in \mathcal{M}. \quad (13)$$

Consider an estimator \hat{P}_n^* for P_0 which solves the efficient influence curve equation:

$$\mathbb{P}_n D_\tau^*(\hat{P}_n^*) = o_P(n^{-1/2}). \quad (14)$$

Now, if the following conditions (i) and (ii) hold true,

(i) $R(\hat{P}_n^*, P_0) = o_P(n^{-1/2})$,

(ii) $D_\tau^*(\hat{P}_n^*)$ belongs to a Donsker class, and $P_0(D_\tau^*(\hat{P}_n^*) - D_\tau^*(P_0))^2$ converges to zero in probability,

then (12) holds for $\hat{\psi}_n^* = \Psi_\tau(\hat{P}_n^*)$, that is, $\Psi_\tau(\hat{P}_n^*)$ is asymptotically linear at P_0 with influence curve $D_\tau^*(P_0)$.

To construct an efficient estimator for the target parameter, the nuisance parameters π , λ^c and λ must be estimated such that the conditions above are met. To shed light on conditions (i) and (ii), we derive the exact second-order remainder $R(P, P_0)$ as displayed in Section 4.2 below; in the subsequent Section 4.3 we present the overall smoothness restrictions on the statistical model (Assumption 5) that combined with highly adaptive lasso estimation provides the basis for establishing conditions (i) and (ii).

4.2 Double robustness of the second-order remainder

As we show in the supplementary material, the second-order remainder defined in Display (13) can be written out explicitly as follows:

$$R(P, P_0) = \int_{\mathcal{O}} \mathbb{1}\{t \leq \tau\} \left(\frac{S_0^c(t- | a, \ell)\pi_0(a | \ell) - S^c(t- | a, \ell)\pi(a | \ell)}{S^c(t- | a, \ell)\pi(a | \ell)} \right) S_0(t- | a, \ell)(\lambda_0(t | a, \ell) - \lambda(t | a, \ell)) dt \frac{S(\tau | a, \ell)}{S(t | a, \ell)} \pi^*(a | \ell) \mu_0(\ell) d\nu(\ell).$$

Importantly, this second-order remainder displays a double robustness structure: We see that $R(\hat{P}_n, P_0) = 0$ if either

1. $\pi_0(a | \ell) = \hat{\pi}_n(a | \ell)$ and $S_0^c(t | a, \ell) = \hat{S}_n^c(t | a, \ell)$ for all $t \in (0, \tau)$,

or,

2. $\lambda_0(t | a, \ell) = \hat{\lambda}_n(t | a, \ell)$ for all $t \in (0, \tau)$.

Let $G(t | a, \ell) := S^c(t- | a, \ell)\pi(a | \ell)$. When G is bounded away from zero by some $\eta > 0$ (Assumption 4), the denominator of the first factor of $R(P, P_0)$ is bounded from above by η^{-1} . Furthermore, $S_0(t- | a, \ell)$ is bounded by 1, and, since $t \leq \tau$, we have that $S(\tau | a, \ell)/S(t | a, \ell) \leq 1$. Now, the product structure of the remainder $R(P, P_0)$ yields, by the Cauchy-Schwartz inequality an upper bound as follows:

$$\begin{aligned} |R(P, P_0)| &\leq \eta^{-1} \int_{\mathcal{L}} \sum_{a=0,1} \int_0^\tau (G_0(t | a, \ell) - G(t | a, \ell)) (\lambda_0(t | a, \ell) - \lambda(t | a, \ell)) dt \pi^*(a | \ell) d\mu_0(\ell) \\ &\leq \eta^{-1} \|G_0 - G\|_{\pi^* \otimes \mu_0 \otimes \rho} \|\lambda_0 - \lambda\|_{\pi^* \otimes \mu_0 \otimes \rho}, \end{aligned}$$

where $\|f\|_{\pi^* \otimes \mu_0 \otimes \rho} = \sqrt{\int f^2 d(\pi^* \otimes \mu_0 \otimes \rho)}$ and ρ denotes the Lebesgue measure on $[0, \tau]$. We see that the required convergence rate, $R(\hat{P}_n^*, P_0) = o_P(n^{-1/2})$, for example is achieved if we estimate both G_0 and λ_0 at a rate faster than $o_P(n^{-1/4})$ with respect to the $L_2(\pi_0 \otimes \mu_0 \otimes \rho)$ -norm.

4.3 Efficient estimation under weak conditions on \mathcal{M}

As we sketch below, the following weak conditions assumed for the statistical model \mathcal{M} combined with the use of highly adaptive lasso estimation (Section 6) for the nuisance parameters provides the basis for conditions (i) and (ii) of Section 4.1.

Assumption 5 (Conditions on \mathcal{M}) *Assume that the nuisance parameters λ, λ^c, π can be parametrized by functions that are càdlàg and have finite sectional variation norm (Gill et al., 1995; van der Laan, 2017), that positivity holds (Assumption 4) and further that $S(\tau) > \eta'$ for some $\eta' > 0$.*

Assumption 5 has the following two important implications. First, the class of càdlàg functions with finite variation is a Donsker class (van der Vaart and Wellner, 1996). Since the efficient influence function is a well-behaved mapping of the nuisance parameters, it inherits the Donsker properties. Second, it constitutes the basis for construction of the highly adaptive lasso estimator shown to converge at a rate faster than $n^{-1/4}$ to its true counterpart belonging to the class of càdlàg functions with finite sectional variation (van der Laan, 2017). The convergence is with respect to the square-root of the loss-based dissimilarity which behaves as the $L_2(\pi_0 \otimes \mu_0 \otimes \rho)$ -norm (Appendix C). In light of the double robustness properties of the second-order remainder (Section 4.2), we see that Assumption 5 combined with highly adaptive lasso estimation provides the basis for establishing conditions (i) and (ii) of Section 4.1.

5 Targeting algorithm

Suppose we have at hand estimators $\hat{\pi}_n, \hat{\lambda}_n^c$ and $\hat{\lambda}_n$. The overall idea of targeted maximum likelihood estimation is to perform an update of $\hat{\lambda}_n \mapsto \hat{\lambda}_{n,*}$ for the given estimators $\hat{\pi}_n, \hat{\lambda}_n^c$, such as to solve the efficient influence curve equation. The estimator $\hat{\lambda}_{n,*}$ is then mapped to an estimator $\hat{S}_{n,*}$ according to (11) to construct an estimator for the target parameter.

In the following Definitions 1 and 2, we define so-called ‘clever weights’ and ‘clever covariates’ that will allow for a concise representation of the efficient influence function that we utilize in the construction of our targeting algorithm.

Definition 1 (Clever weights) *Define clever weights by:*

$$w_t(O) = \mathbb{1}\{\tilde{T} \geq t\} \frac{\pi^*(A | L)}{\pi(A | L)} \frac{\mathbb{1}\{t \leq \tau\}}{S^c(t - | A, L)}, \quad t > 0.$$

Notably, the clever weights depend on the interventional part of the data-generating distribution and on the choice of intervention π^* . In our simulation study (Section 7), we focus on the average treatment effect displayed in (8) which we target directly. In this case the clever weights are defined in terms of both interventions $\pi^1(A | L) = \mathbb{1}\{A = 1\}$ and $\pi^0(A | L) = \mathbb{1}\{A = 0\}$ as follows:

$$w_t^{\text{ATE}}(O) = \mathbb{1}\{\tilde{T} \geq t\} \left(\frac{\mathbb{1}\{A = 1\}}{\pi(1 | L)} - \frac{\mathbb{1}\{A = 0\}}{\pi(0 | L)} \right) \frac{\mathbb{1}\{t \leq \tau\}}{S^c(t - | A, L)}. \quad (15)$$

The clever covariates defined below, to the contrary, only depend on the non-interventional part and are thus fixed across choices of interventions. Estimators for the clever weights remain constant, whereas estimators for the clever weights are updated as part of the targeting algorithm.

Definition 2 (Clever covariates) Define clever covariates by:

$$h_t(O) = \frac{S(\tau | A, L)}{S(t | A, L)}, \quad t > 0.$$

With w_t and h_t defined by Definitions 1 and 2, the efficient influence function for our statistical model \mathcal{M} and target parameter $\Psi_\tau : \mathcal{M} \rightarrow \mathbb{R}$ can now be represented on the concise form:

$$\begin{aligned} D_\tau^*(P)(O) = & \int_0^\tau w_t(O)h_t(O) (N(dt) - \Lambda(dt | A, L)) \\ & + 1 - \sum_{a=0,1} S(\tau | a, L)\pi^*(a | L) - \Psi_\tau(P). \end{aligned} \quad (16)$$

Any estimator on the form (11) solves all but the first term of the efficient influence curve equation, so what remains is the first term, corresponding to the first term of (16), giving rise to the equation:

$$\mathbb{P}_n D_{1,\tau}^*(\hat{P}_n) = o_P(1), \quad \text{where,} \quad D_{1,\tau}^*(P)(O) := \int_0^\tau w_t(O)h_t(O) (N(dt) - \Lambda(dt | A, L)). \quad (17)$$

Our algorithm is formed by iterative update steps for $\hat{\lambda}_n$ along a path defined by a one-dimensional fluctuation model, defining a sequence of estimators,

$$\hat{\lambda}_n = \hat{\lambda}_{n,k=0}, \quad \hat{\lambda}_{n,k=1}, \quad \hat{\lambda}_{n,k=2}, \quad \dots \quad (18)$$

such that \hat{P}_n^* , characterized by $\hat{\lambda}_n^c$, $\hat{\pi}_n$ and a final estimator $\hat{\lambda}_{n,*} = \hat{\lambda}_{n,k=k^*}$ from the sequence (18), solves (17). We refer to $\hat{\lambda}_n = \hat{\lambda}_{n,k=0}$ as the *initial estimator* and to $\hat{\lambda}_{n,*} = \hat{\lambda}_{n,k=k^*}$ as the *targeted estimator*.

5.1 Fluctuation model and update algorithm for λ

The update steps for each element $\hat{\lambda}_{n,k}$ of the sequence of estimators (18) are performed along a one-dimensional fluctuation model through $\hat{\lambda}_{n,k}$. To this end, we define the multiplicative hazards type fluctuation model:

$$\lambda(t; \varepsilon) = \lambda(t) \exp(\varepsilon w_t h_t), \quad \varepsilon \in \mathbb{R}. \quad (19)$$

The estimators $\hat{\pi}_n$ and $\hat{\lambda}_n^c$ together define an estimator \hat{w}_t for the clever weights, whereas $\hat{\lambda}_n$ defines an estimator \hat{h}_t for the clever covariate. Now, plugging the evaluation of (19) in the current estimator $\hat{\lambda}_{n,k}$, the estimator \hat{w}_t for the clever weights and the estimator $\hat{h}_{t,k}$, obtained from $\hat{\lambda}_{n,k}$, for the clever covariate into (17) defines an equation in ε :

$$\begin{aligned} 0 = & \frac{1}{n} \sum_{i=1}^n \left(\int_0^\tau \hat{w}_t(O_i) \hat{h}_{t,k}(O_i) N_i(dt) \right. \\ & \left. - \int_0^\tau \hat{w}_t(O_i) \hat{h}_{t,k}(O_i) \exp(\varepsilon \hat{w}_t(O_i) \hat{h}_{t,k}(O_i)) \hat{\lambda}_{n,k}(t | A_i, L_i) dt \right). \end{aligned} \quad (20)$$

The solution $\hat{\varepsilon}_k$ to (A.4) defines the update of $\hat{\lambda}_{n,k}$ along the fluctuation model:

$$\hat{\lambda}_{n,k+1} := \hat{\lambda}_{n,k}(t) \exp(\hat{\varepsilon}_k \hat{w}_t \hat{h}_{t,k}), \quad k \geq 0.$$

The steps from k to $k+1$ are repeated until $\hat{\varepsilon}_{k^*} \approx 0$, or, more precisely, the corresponding estimator \hat{P}_n^* , consisting of $\hat{\lambda}_n^c$, $\hat{\pi}_n$ and $\hat{\lambda}_{n,*} = \hat{\lambda}_{n,k^*}$, solves

$$|\mathbb{P}_n D_\tau^*(\hat{P}_n^*)| \leq s_n,$$

for the stopping criterion $s_n = \hat{\sigma}_n / (n^{-1/2} \log n)$, where $\hat{\sigma}_n^2$ is the estimated variance of the efficient influence function.

5.2 Final estimation

The final estimator $\hat{\lambda}_{n,*}$ defines a corresponding estimator $\hat{S}_{n,*}$ for the survival function based on which we can construct an estimator of the target parameter by:

$$\hat{\psi}_n^* = 1 - \frac{1}{n} \sum_{i=1}^n \left(\sum_{a=0,1} \hat{S}_{n,*}(\tau | a, L_i) \pi^*(a | L_i) \right).$$

Under conditions (i) and (ii) of Section 4.1, we can use the asymptotic normal distribution

$$\sqrt{n} (\hat{\psi}_n^* - \psi_0) \xrightarrow{\mathcal{D}} \mathcal{N}(0, P_0 D_\tau^*(P_0)^2),$$

to provide an approximate two-sided confidence interval. The asymptotic variance of the estimator is equal to the variance of the efficient influence function and can be estimated by $\hat{\sigma}_n^2 = \mathbb{P}_n (D_\tau^*(\hat{P}_n^*))^2$.

6 Initial estimation

To carry out our targeting algorithm, we need initial estimators for the conditional hazards λ^c and λ of the censoring process and the event process, respectively, and further for the conditional distribution of treatment given covariates π . Estimation of π can be done by any binary regression method, including logistic regression and a large variety of machine learning algorithms. Thus, our focus is here on the estimation of a general continuous-time conditional hazard, denoted $\lambda(t | Z)$, where, for our purposes, Z consists of the baseline treatment and covariates, $Z = (A, L) \in \mathbb{R}^{d+1}$. Particularly, following previous work (Benkeser et al., 2016; van der Laan, 2017; van der Laan and Rose, 2018, Chapter 6,7), Section 6.1 below presents our highly adaptive lasso estimator for such continuous-time conditional hazards.

6.1 Highly adaptive lasso estimation of hazards

To construct our highly adaptive lasso estimator for the conditional hazard $\lambda(t | Z)$, we propose the reparametrization as follows:

$$\lambda(t | Z) = \exp(f(t, Z)), \quad f : [0, \tau] \times \mathbb{R}^{d+1} \rightarrow \mathbb{R}, \quad (21)$$

and denote by $[0, \kappa] \subset \mathbb{R}^{d+1}$ the support of $z \mapsto f(t, z)$ with $\kappa = \infty$ allowed. The steps to define the highly adaptive lasso estimator for λ involves some technical parts following earlier work (van der Laan, 2017), and, specifically, we need notation for the sectional variation of f that we now

present. For a subset of indices $\mathcal{S} \subset \{1, \dots, d+1\}$, we denote by $z_{\mathcal{S}}$ the \mathcal{S} -specific coordinates of $z \in \mathbb{R}^{d+1}$ and by $z \mapsto f_{\mathcal{S}}(t, z) = f(t, z_{\mathcal{S}}, 0_{\mathcal{S}^c})$ the \mathcal{S} -specific section of f that sets the coordinates in the complement of \mathcal{S} equal to zero. As we state below, we will make the assumption that the function $f(t, z)$ is càdlàg (i.e., right-continuous with left limits), then, following Gill et al. (1995); van der Laan (2017), the sectional variation norm of f is

$$\|f\|_v = |f(0, 0)| + \int_{(0, \tau]} |f(dt, 0)| + \sum_{\mathcal{S} \subset \{1, \dots, d+1\}} \left(\int_{(0_{\mathcal{S}}, \kappa_{\mathcal{S}}]} |f_{\mathcal{S}}(0, dz)| + \int_{(0, \tau]} \int_{(0_{\mathcal{S}}, \kappa_{\mathcal{S}}]} |f_{\mathcal{S}}(dt, dz)| \right).$$

The key to defining the highly adaptive lasso estimator is the following assumption on the function class containing f . Let $\mathcal{F}_{\mathcal{M}}$ denote the class of càdlàg functions with sectional variation norm bounded by a constant $\mathcal{M} < \infty$. We assume that $f \in \mathcal{F}_{\mathcal{M}}$.

Highly adaptive lasso estimation of f is defined by the infinite-dimensional minimization problem over all $f \in \mathcal{F}_{\mathcal{M}}$:

$$\min_{f \in \mathcal{F}_{\mathcal{M}}} \mathbb{P}_n \mathcal{L}(f), \quad (22)$$

for a loss function $(O, f) \mapsto \mathcal{L}(f)(O)$. The practical construction consists of approximating the minimizer over all $\mathcal{F}_{\mathcal{M}}$ in (22) by the minimizer over discrete measures in $\mathcal{F}_{\mathcal{M}}$. Particularly, the assumption that $f \in \mathcal{F}_{\mathcal{M}}$ yields a representation for $f \in \mathcal{F}_{\mathcal{M}}$ in terms of its measures over sections (Gill et al., 1995)

$$f(t, z) = f(0, 0) + \int_{(0, \tau]} \mathbb{1}\{s \leq t\} f(ds, 0) + \sum_{\mathcal{S} \subset \{1, \dots, d+1\}} \left(\int_{(0_{\mathcal{S}}, \kappa_{\mathcal{S}}]} \mathbb{1}\{u \leq z_{\mathcal{S}}\} f_{\mathcal{S}}(0, du) + \int_{(0, \tau]} \int_{(0_{\mathcal{S}}, \kappa_{\mathcal{S}}]} \mathbb{1}\{s \leq t\} \mathbb{1}\{u \leq z_{\mathcal{S}}\} f_{\mathcal{S}}(ds, du) \right),$$

which for an approximation over a finite support becomes a finite linear combinations of indicators functions and corresponding coefficients being the pointmass assigned to support points.

Let us consider a grid of time-points partitioning $[0, \tau]$, $0 = t_0 < t_1 < \dots < t_R < t_{R+1} = \tau$, and further a partitioning $\mathcal{Z}_1 \cup \dots \cup \mathcal{Z}_M$ of the sample space of Z into $(d+1)$ -dimensional cubes,

$$\cup_{m=1}^M \mathcal{Z}_m = [0, \kappa];$$

let $z_m \in \mathcal{Z}_m$ be the midpoint of the cube \mathcal{Z}_m . We introduce the indicator basis functions $\phi_r(t) = \mathbb{1}\{t_r \leq t\}$ and $\phi_{\mathcal{S}, m}(z) = \mathbb{1}\{z_{m, \mathcal{S}} \leq z_{\mathcal{S}}\}$ which are central components for the following. Indeed, the discrete approximation f_{β} of f with support over these points admits a representation as follows

$$f_{\beta}(t, z) = \sum_{r=0}^R \phi_r(t) \beta_r + \sum_{r=0}^R \sum_{\mathcal{S} \subset \{1, \dots, d+1\}} \sum_{m=1}^M \phi_r(t) \phi_{\mathcal{S}, m}(z) \beta_{r, \mathcal{S}, m}. \quad (23)$$

Here, we have that $\beta_{r, \mathcal{S}, m} = f_{\mathcal{S}}(dt_r, dz_{m, \mathcal{S}})$, for each $r \in \{1, \dots, R\}$, $\mathcal{S} \subset \{1, \dots, d+1\}$, $m = 1, \dots, M$, is the point-mass that the \mathcal{S} -specific section of f_{β} assigns to the point defined by z_m and t_{r-1} . Furthermore, $\beta_r = f_{\beta}(dt_r, 0)$, for $r = 1, \dots, R$, are the increments along the time axis alone, and $\beta_{0, \mathcal{S}, m} = f_{\beta}(0, dz_{m, \mathcal{S}})$, for $\mathcal{S} \subset \{1, \dots, d+1\}$, $m = 1, \dots, M$, the increments along the

z -axis alone. Lastly, $\beta_0 = f_\beta(0, 0)$ is the point-mass assigned by f_β to zero. We refer to the stacked vector of $\beta_r, \beta_{r,\mathcal{S},m}$ as the vector of parameter coefficients and note that this vector completely characterizes the behavior of f_β . Particularly, the sectional variation norm of f_β becomes a sum over the absolute values of its coefficients

$$\|f_\beta\|_v = \sum_{r=0}^R |\beta_r| + \sum_{r=0}^R \sum_{\mathcal{S} \subset \{1, \dots, d+1\}} \sum_{m=1}^M |\beta_{r,\mathcal{S},m}| = \|\beta\|_1,$$

i.e., the sectional variation norm of f_β equals to the L_1 -norm of the coefficient vector. We may now define the highly adaptive lasso estimator as follows.

Definition 3 (Highly adaptive lasso estimator) *The highly adaptive lasso estimator for f is obtained as $\hat{f}_n = f_{\hat{\beta}_n}$ where:*

$$\hat{\beta}_n = \underset{\beta}{\operatorname{argmin}} \mathbb{P}_n \mathcal{L}(f_\beta), \quad \text{s.t.}, \quad \|\beta\|_1 \leq \mathcal{M}. \quad (24)$$

Notably, (24) corresponds to an L_1 -penalized regression with indicator functions $\phi_r(t)$ and $\phi_r(t)\phi_{\mathcal{S},m}(z)$ as covariates and $\beta_r, \beta_{r,\mathcal{S},m}$ as corresponding coefficients.

Choosing the support of f_β fine enough, the L_1 -norm of the coefficient vector approximates the sectional variation norm of f and the solution $f_{\hat{\beta}_n}$ to the minimization problem defined by (24) in Definition 3 approximates the infinite-dimensional minimization problem in (22) over all $f \in \mathcal{F}_{\mathcal{M}}$ (van der Laan, 2017, Appendix D).

We consider particularly the log-likelihood loss function for our purposes, i.e., we define $\mathcal{L}(f)(O) = -\ell_{\log\text{lik}}(f)(O)$ where $(O, f) \mapsto \ell_{\log\text{lik}}(f)(O)$ denotes the log-likelihood

$$\ell_{\log\text{lik}}(f)(O) = \int_0^\tau f(t, Z) N(dt) - \int_0^\tau \mathbb{1}\{\tilde{T} \geq t\} \exp(f(t, Z)) dt.$$

As we review in Appendix B, one way to solve the minimization problem in (24) in practice with this loss functions is by standard L_1 -penalized Poisson regression software, exploiting the correspondence between $\ell_{\log\text{lik}}$ and the Poisson log-likelihood (Andersen et al., 1993; Lindsey, 1995). The highly adaptive lasso estimator $\hat{f}_n = f_{\hat{\beta}_n}$ can next be plugged into (21), providing an estimator for the hazard itself. Note that we use cross-validation to data-adaptively select the bound on the variation norm; by the oracle properties of cross-validation (van der Laan and Dudoit, 2003; van der Vaart et al., 2006), we only need at least one of the bounds considered as a candidate in the library to be larger than the true variation norm.

7 Simulation study

In this section we consider a simulation study of estimation of the average treatment effect parameter $\Psi_\tau^{\text{ATE}} : \mathcal{M} \rightarrow \mathbb{R}$ from (9) specifically. Accordingly, we work with the clever weights w_t^{ATE} from (15), and the efficient influence function is given by:

$$D_\tau^{\text{ATE}}(P)(O) = \int_0^\tau w_t^{\text{ATE}}(O) h_t(O) (N(dt) - \Lambda(dt | A, L)) + S(\tau | 1, L) - S(\tau | 0, L) - \Psi_\tau^{\text{ATE}}(P).$$

Our simulation setting imitates the setting of a randomized trial in which trial participants are randomized to a treatment and followed over time until either the event or right-censoring happens. Particularly, $L = (L_1, L_2, L_3)$ are baseline covariates and $A \in \{0, 1\}$ is the randomized treatment. Altogether, we consider two different versions of the censoring mechanism:

$$\begin{aligned} \text{Covariate independent censoring: } & \lambda^c(t | A, L) = \lambda_0^c(t), \\ \text{Covariate dependent censoring: } & \lambda^c(t | A, L) = \lambda_0^c(t) \exp(-0.8L_3 + 1.2L_1A). \end{aligned}$$

Two of the covariates, L_1, L_2 , are uniformly distributed on $(-1, 1)$, whereas L_3 is uniform on $(0, 1)$. Baseline hazards, λ_0^c, λ_0 , correspond to Weibull distributions with shape parameter 0.7 and scale parameter 1.7 and are the same across all simulations. The hazard of the event distribution is given by:

$$\lambda(t | A, L) = \lambda_0(t) \exp(0.7\mathbb{1}\{t < t'\}A - 0.225\mathbb{1}\{t \geq t'\}A + 1.2L_1^2),$$

with changepoint $t' = 0.7$. We focus on the survival difference beyond $\tau = 1.2$ years of follow-up and consider the following different estimation procedures for comparison: 1) A substitution estimator based on a Kaplan-Meier estimator for the survival curve in each treatment arm, 2) a targeted maximum likelihood estimator based on a misspecified Cox model (including main effects for A and L_1) for initial estimation, and 3) a targeted maximum likelihood estimator based on a Poisson-based highly adaptive lasso estimator for initial estimation. A grid of ten time-points was used for the time axis and eight knot-points for each covariate. The upper bound for the sectional variation norm was selected with cross-validation. In the targeting step, we use a correctly specified Cox model for the hazard of the censoring distribution. The results are presented in Table 1, showing that the targeted maximum likelihood estimator based on the flexible Poisson-based highly adaptive lasso initial estimation improves precision compared to both the misspecified Cox model and the Kaplan-Meier estimator.

Table 1: Results from the simulation study. ‘HAL-TMLE’ uses the Poisson-based highly adaptive lasso estimator for initial estimation in the targeted maximum likelihood estimation algorithm. ‘Cox-TMLE’ uses a misspecified Cox model for initial estimation in the targeted maximum likelihood estimation algorithm. ‘KM’ uses a Kaplan-Meier approach.

	Covariate dependent censoring			Covariate independent censoring		
	HAL-TMLE	Cox-TMLE	KM	HAL-TMLE	Cox-TMLE	KM
Bias	-0.0001	-0.0005	-0.0059	0.0000	0.0001	0.0002
Cov (95%)	0.9540	0.9500	0.9320	0.9360	0.9320	0.9300
$\sqrt{\text{MSE}}$	0.0315	0.0321	0.0326	0.0310	0.0316	0.0316
rel. MSE	0.9344	0.9703	1.0000	0.9600	0.9988	1.0000

8 Concluding remarks

Our simulations show that the targeted maximum likelihood estimator based on the implemented Poisson-based highly adaptive lasso estimator performs in agreement with the asymptotic theory, improving precision relative to the Kaplan-Meier approach and achieving proper coverage based on the efficient influence function.

To fully optimize the estimation of all nuisance parameters, which are comprised by the hazard for the censoring and event distributions and the conditional distribution of treatment, we recommend to apply loss-function based cross-validation to combine the highly adaptive lasso estimator with other estimators. This procedure of selecting the best estimator from a prespecified library of candidate algorithms by minimizing the cross-validated empirical risk is often referred to as super learning (van der Laan et al., 2007) and the general oracle inequality for loss-based cross-validation (van der Laan and Dudoit, 2003; van der Vaart et al., 2006) yields that the super learner will achieve the minimal rate of convergence of the estimators in the library. In future work, we plan to establish the oracle inequality for the considered loss function, involving understanding that a discrete baseline hazard is really the appropriate approach (see also Appendix B).

Supplementary material

Supplementary material includes I) a sketch of the proof of causal interpretability; II) the proof for the sufficiency of the conditions stated in Section 4.1; III) derivations of efficient influence functions and of second-order remainders for the survival analysis setting; IV) derivations of efficient influence functions and of second-order remainders for the competing risks setting; V) descriptions of implementations; and VI) additional simulations for the competing risks setting.

Appendix A

Targeted maximum likelihood estimation for the competing risks setting

We here consider the competing risks analogue of the observed data setting considered in the main text. Let $A \in \{0, 1\}$, $L \in \mathbb{R}^d$ and the time under observation $\tilde{T} \in \mathbb{R}_+$ be as in Section 2. The variables $T \in \mathbb{R}_+$ and $C \in \mathbb{R}_+$ represent the times to event, now one of $J \geq 1$ types, and censoring, respectively, such that the observed time is $\tilde{T} = \min(T, C)$. Together, $O = (L, A, \tilde{T}, \tilde{\Delta})$ constitutes the observed data. We further define an event indicator $\Delta \in \{1, 2, \dots, J\}$ telling us the type of event happening. Note that Δ is subject to right-censoring such that we only observe $\tilde{\Delta} = \mathbb{1}\{\tilde{T} \geq C\}\Delta$. We define the counting processes $N(t) = (N_1(t), \dots, N_J(t)) = (\mathbb{1}\{\tilde{T} \leq t, \tilde{\Delta} = 1\}, \dots, \mathbb{1}\{\tilde{T} \leq t, \tilde{\Delta} = J\})$ and $N^c(t) = \mathbb{1}\{\tilde{T} \leq t, \Delta = 0\}$. Let $\lambda_{0,j}$ denote the cause j specific hazard, defined as

$$\lambda_{0,j}(t | a, \ell) = \lim_{h \rightarrow 0} h^{-1} P_0(T \leq t + h, \Delta = j | \tilde{T} \geq t, A = a, L = \ell),$$

and $\Lambda_{0,j}$ the corresponding cumulative hazard. The interventional part of the likelihood is unchanged, whereas the non-interventional part is given by:

$$g_0(o) = \mu_0(\ell) \prod_{j=1}^J (\lambda_{0,j}(t | a, \ell))^{\mathbb{1}\{\delta=j\}} S_0(t- | a, \ell),$$

where $S_0(t | a, \ell) = \exp(-\int_0^t \sum_{j=1}^J \lambda_{0,j}(s | a, \ell) ds)$. With no loss of generality, we assume that $J = 2$. We define the target parameter $\Psi_\tau : \mathcal{M} \rightarrow \mathbb{R}$ as the intervention-specific absolute risk beyond time τ :

$$\Psi_\tau(P) = \int_{\mathcal{L}} \sum_{a=0,1} F_1(\tau | a, \ell) \pi^*(a | L) \mu(\ell) d\nu(\ell), \quad (\text{A.1})$$

where

$$F_1(t | a, \ell) = \int_0^t S(s- | a, \ell) \Lambda_1(ds | a, \ell),$$

is the risk function for the event of interest ($j = 1$) (Gray, 1988). We define clever weights and covariates according to Definitions 4 and 5 as follows.

Definition 4 (Clever weights) *Define clever weights by:*

$$w_t(O) = \mathbb{1}\{\tilde{T} \geq t\} \frac{\pi^*(A | L)}{\pi(A | L)} \frac{\mathbb{1}\{t \leq \tau\}}{S^c(t- | A, L)}.$$

The clever weights are the same as for the survival analysis setting (Definition 1).

Definition 5 (Clever covariates) *Define clever covariates by:*

$$h_{1,t}(O) = 1 - \frac{F_1(\tau | A, L) - F_1(t | A, L)}{S(t | A, L)}$$

$$h_{2,t}(O) = -\frac{F_1(\tau | A, L) - F_1(t | A, L)}{S(t | A, L)},$$

for $t > 0$.

Now we can define a targeted maximum likelihood algorithm for estimation of the target parameter (A.1). The algorithm involves targeting steps for both the intensity of the event process of interest λ_1 and then intensity of the competing event process λ_2 . We define the fluctuation models as follows:

$$\lambda_1(t; \varepsilon_1) = \lambda_1(t) \exp(\varepsilon_1 w_t h_{1,t}), \quad \varepsilon_1 \in \mathbb{R}, \quad (\text{A.2})$$

$$\lambda_2(t; \varepsilon_2) = \lambda_2(t) \exp(\varepsilon_2 w_t h_{2,t}), \quad \varepsilon_2 \in \mathbb{R}. \quad (\text{A.3})$$

As the setting of Section 5, $\hat{\pi}_n$ and $\hat{\lambda}_n^c$ together define an estimator \hat{w}_t for the clever weights, but now estimators of the clever covariates $\hat{h}_{1,t}$, $\hat{h}_{2,t}$ need both $\hat{\lambda}_{1,n}$ and $\hat{\lambda}_{2,n}$. In each round of iterations, the evaluation of (A.2) in the current estimator $\hat{\lambda}_{1,n,k}$ and the evaluation of (A.3) in the current estimator $\hat{\lambda}_{2,n,k}$ together with the estimator \hat{w}_t for the clever weights defines one equation in ε_1 and one in ε_2 :

$$0 = \frac{1}{n} \sum_{i=1}^n \left(\int_0^\tau \hat{w}_t(O_i) \hat{h}_{j,t,k}(O_i) N_{j,i}(dt) \right. \\ \left. - \int_0^\tau \hat{w}_t(O_i) \hat{h}_{j,t,k}(O_i) \exp(\varepsilon_j \hat{w}_t(O_i) \hat{h}_{j,t,k}(O_i)) \hat{\lambda}_{j,n,k}(t | A_i, L_i) dt \right), \quad j = 1, 2. \quad (\text{A.4})$$

The solution $\hat{\varepsilon}_{j,k}$, $j = 1, 2$, defines the update of $\hat{\lambda}_{j,n,k}$, $j = 1, 2$, along the corresponding fluctuation model:

$$\hat{\lambda}_{j,n,k+1}(t) := \hat{\lambda}_{j,n,k}(t) \exp(\hat{\varepsilon}_{j,k} \hat{w}_t \hat{h}_{j,t,k}), \quad j = 1, 2, \quad k \geq 0.$$

The steps from k to $k + 1$ are repeated until convergence.

Appendix B

In this appendix, we sketch a result for the highly adaptive lasso estimator presented in Section 6.1 of the main text: That we can implement the highly adaptive lasso estimator (Definition 3) in practice with L_1 -penalized Poisson regression software.

The fact that the likelihood for a proportional hazards model corresponds to that of a certain Poisson regression is a well-known result from the survival analysis literature (Andersen et al., 1993; Lindsey, 1995). The technique is commonly applied in large-scale observational studies (Lin et al., 1998; Grøn, 2016) to approximate Cox regression models in a way that can potentially save considerable computation time and memory usage. Put shortly, the Poisson formulation is just a different formulation of a proportional hazards model with a baseline rate modeled by a parameter over a grid of time-points, assuming a constant rate in each interval between time-points times which allows for the use of standard Poisson regression software. We emphasize that the Poisson regression is only used as a tool; never is the Poisson model in fact assumed for the data. Furthermore, although the formulation is over a discrete time-points, the exact continuous event times are still used in the estimation, specifically in the aggregated risk time used as an offset in the regression.

Highly adaptive lasso estimation implemented as a Poisson regression

We here clarify how one may in practice solve the minimization problem of Definition 3 for the highly adaptive lasso estimator by standard L_1 -penalized Poisson regression software. The loss function used in Definition 3 to define the highly adaptive lasso for a hazard on the form $\lambda(t | Z) = \exp(f(t, Z))$ is the log-likelihood loss function $\mathcal{L}(f)(O) = -\ell_{\text{loglik}}(f)(O)$ where $(O, f) \mapsto \ell_{\text{loglik}}(f)(O)$ is given by

$$\ell_{\text{loglik}}(f)(O) = \int_0^\tau f(t, Z) N(dt) - \int_0^\tau \mathbb{1}\{\tilde{T} \geq t\} \exp(f(t, Z)) dt.$$

Particularly, for f_β admitting the representation in (23),

$$f_\beta(t, z) = \sum_{r=0}^R \phi_r(t) \beta_r + \sum_{r=0}^R \sum_{\mathcal{S} \subset \{1, \dots, d+1\}} \sum_{j \in \mathcal{I}_\mathcal{S}} \phi_r(t) \phi_{\mathcal{S}, j}(z) \beta_{r, \mathcal{S}, j},$$

we have that

$$\begin{aligned} \ell_{\text{loglik}}(f_\beta)(O) &= \sum_{r=0}^R (N(t_{r+1}) - N(t_r)) f_\beta(t_r, Z) \\ &\quad - \sum_{r=0}^R \mathbb{1}\{\tilde{T} \geq t_r\} \exp(f_\beta(t_r, Z)) (\min(\tilde{T}, t_{r+1}) - t_r). \end{aligned} \tag{B.1}$$

Applying our notation from Section 6.1 of the main text with a partitioning $\mathcal{Z}_1 \cup \dots \cup \mathcal{Z}_M$ of the sample space of Z into $(d+1)$ -dimensional cubes, m is an index that runs through distinct values of the vector $(\phi_{\mathcal{S}, m}(Z) : \mathcal{S}, m)$, i.e., for each $m = 1, \dots, M$,

$$\forall z^1, z^2 \in \mathcal{Z}_m : \phi_{\mathcal{S}, m}(z^1) = \phi_{\mathcal{S}, m}(z^2), \quad \text{for all } \mathcal{S}, m.$$

Let again $z_m \in \mathcal{Z}_m$ be the midpoint of the cube \mathcal{Z}_m . Now, observe on the one hand that

$$\begin{aligned} & \sum_{i=1}^n \sum_{r=0}^R (N_i(t_{r+1}) - N_i(t_r)) f_\beta(t_r, Z_i) \\ &= \underbrace{\sum_{m=1}^M \sum_{r=0}^R \sum_{i=1}^n (N_i(t_{r+1}) - N_i(t_r)) \mathbb{1}\{Z_i \in \mathcal{Z}_m\} f_\beta(t_r, z(m))}_{=:\mathcal{D}_{m,r}}, \end{aligned}$$

and, on the other hand, that

$$\begin{aligned} & \sum_{i=1}^n \sum_{r=0}^R \mathbb{1}\{\tilde{T}_i \geq t_r\} (\min(\tilde{T}_i, t_{r+1}) - t_r) \exp(f_\beta(t_r, Z_i)) \\ &= \underbrace{\sum_{m=1}^M \sum_{r=0}^R \sum_{i=1}^n \mathbb{1}\{\tilde{T}_i \geq t_r\} (\min(\tilde{T}_i, t_{r+1}) - t_r) \mathbb{1}\{Z_i \in \mathcal{Z}_m\} \exp(f_\beta(t_r, z(m)))}_{=:\mathcal{R}_{m,r}}, \end{aligned}$$

from which we see that

$$\begin{aligned} & \operatorname{argmin}_{\beta} -\mathbb{P}_n \ell_{\log\text{lik}}(f_\beta) \\ &= \operatorname{argmin}_{\beta} - \left(\sum_{i=1}^n \sum_{r=0}^R (N(t_{r+1}) - N(t_r)) f_\beta(t_r, Z_i) \right. \\ & \quad \left. - \sum_{r=0}^R \mathbb{1}\{\tilde{T}_i \geq t_r\} (\min(\tilde{T}_i, t_{r+1}) - t_r) \exp(f_\beta(t_r, Z_i)) \right) \\ &= \operatorname{argmin}_{\beta} - \sum_{m=1}^M \sum_{r=0}^R (\mathcal{D}_{m,r} f_\beta(t_r, z_m) - \mathcal{R}_{m,r} \exp(f_\beta(t_r, z_m))). \end{aligned} \quad (\text{B.2})$$

Importantly, (B.2) corresponds to minimizing a Poisson likelihood loss function with event counts $\mathcal{D}_{m,r}$, mean $\exp(f_\beta(t_r, z_m))$ and offset $\log \mathcal{R}_{m,r}$. So, in practice we can fit this Poisson regression model, working with the aggregated data (where, for each combination of covariates values, z_1, \dots, z_M , and intervals $[t_r, t_{r+1})$, we count events and we sum up the total risk time).

Convergence of the highly adaptive lasso estimator implemented as a Poisson regression

For the log-likelihood loss function $\mathcal{L}(f)(O) = -\ell_{\log\text{lik}}(f)(O)$, define as follows

$$\hat{f}_n^* = \operatorname{argmin}_{f \in \mathcal{F}_{\mathcal{M}}} \mathbb{P}_n \mathcal{L}(f). \quad (\text{B.3})$$

For a constant $\mathcal{M}' < \infty$ we make the assumptions that

$$\sup_{f \in \mathcal{F}_{\mathcal{M}}} \frac{P_0(\mathcal{L}(f) - \mathcal{L}(f_0))^2}{P_0(\mathcal{L}(f) - \mathcal{L}(f_0))} \leq \mathcal{M}', \quad (\text{B.4})$$

and,

$$\sup_{f \in \mathcal{F}_{\mathcal{M}}} \frac{\|\mathcal{L}(f)\|_v}{\|f\|_v} < \infty. \quad (\text{B.5})$$

The following arguments follow previous work (van der Laan, 2017; van der Laan and Rose, 2018, Chapters 6,7), but are repeated here for completeness. Consider the following bound for the log-likelihood based dissimilarity:

$$\begin{aligned} P_0(\mathcal{L}(\hat{f}_n^*) - \mathcal{L}(f_0)) &= -(\mathbb{P}_n - P_0)(\mathcal{L}(\hat{f}_n^*) - \mathcal{L}(f_0)) + \mathbb{P}_n(\mathcal{L}(\hat{f}_n^*) - \mathcal{L}(f_0)) \\ &\leq -(\mathbb{P}_n - P_0)(\mathcal{L}(\hat{f}_n^*) - \mathcal{L}(f_0)), \end{aligned} \quad (\text{B.6})$$

where, at the inequality, we used (B.3). It is a straightforward consequence of Assumption 5 and the bound from (B.5) that $\mathcal{L}(\hat{f}_n^*) - \mathcal{L}(f_0)$ belongs to a Donsker class. Thus, we have that

$$-(\mathbb{P}_n - P_0)(\mathcal{L}(\hat{f}_n^*) - \mathcal{L}(f_0)) = o_P(n^{-1/2}),$$

which, combined with (B.6) and the bound from (B.4) implies that

$$(\mathbb{P}_n - P_0)(\mathcal{L}(\hat{f}_n^*) - \mathcal{L}(f_0)) = o_P(n^{-1/2}),$$

by van der Vaart (2000, Lemma 19.24). Thus, again by (B.6), it follows that

$$P_0(\mathcal{L}(\hat{f}_n^*) - \mathcal{L}(f_0)) = o_P(n^{-1/2}). \quad (\text{B.7})$$

We next show that the highly adaptive lasso estimator from Definition 3 fulfills the same convergence condition for a partitioning that is chosen fine enough. For the following we let $h > 0$ denote the maximal side length of each cube of the partitioning for a given partitioning of $[0, \tau]$ and $[0, \kappa]$. Then, for given h and a constant $\mathcal{M} < \infty$, we define as follows:

$$\begin{aligned} \mathcal{F}_{\mathcal{M}}^h = \left\{ \sum_{r=0}^R \phi_r(t) \beta_r + \sum_{r=0}^R \sum_{\mathcal{S} \subset \{1, \dots, d+1\}} \sum_{m=1}^M \phi_r(t) \phi_{\mathcal{S}, m}(z) \beta_{r, \mathcal{S}, m} : \right. \\ \left. \|\beta\|_1 = \sum_{r=0}^R |\beta_r| + \sum_{r=0}^R \sum_{\mathcal{S} \subset \{1, \dots, d+1\}} \sum_{m=1}^M |\beta_{r, \mathcal{S}, m}| \leq \mathcal{M} \right\}. \end{aligned} \quad (\text{B.8})$$

Now the highly adaptive lasso estimator from Definition 3 can be written as follows

$$\hat{f}_n^h = \underset{f \in \mathcal{F}_{\mathcal{M}}^h}{\operatorname{argmin}} \mathbb{P}_n \mathcal{L}(f). \quad (\text{B.9})$$

Consider

$$\begin{aligned} P_0(\mathcal{L}(\hat{f}_n^h) - \mathcal{L}(f_0)) &= -(\mathbb{P}_n - P_0)(\mathcal{L}(\hat{f}_n^h) - \mathcal{L}(f_0)) + \mathbb{P}_n(\mathcal{L}(\hat{f}_n^h) - \mathcal{L}(f_0)) \\ &= -(\mathbb{P}_n - P_0)(\mathcal{L}(\hat{f}_n^h) - \mathcal{L}(f_0)) + \mathbb{P}_n(\mathcal{L}(\hat{f}_n^h) - \mathcal{L}(\hat{f}_n^*)) + \mathbb{P}_n(\mathcal{L}(\hat{f}_n^*) - \mathcal{L}(f_0)) \\ &\leq -(\mathbb{P}_n - P_0)(\mathcal{L}(\hat{f}_n^h) - \mathcal{L}(f_0)) + \mathbb{P}_n(\mathcal{L}(\hat{f}_n^h) - \mathcal{L}(\hat{f}_n^*)). \end{aligned}$$

Following the arguments from (B.6) to (B.7) above with \hat{f}_n^h substituted for \hat{f}_n^* yields that

$$(\mathbb{P}_n - P_0)(\mathcal{L}(\hat{f}_n^h) - \mathcal{L}(f_0)) = o_P(n^{-1/2}).$$

Furthermore, van der Laan (2017, Lemma 11) gives that

$$\mathbb{P}_n(\mathcal{L}(\hat{f}_n^h) - \mathcal{L}(\hat{f}_n^*)) \rightarrow 0, \quad \text{as } h \rightarrow 0;$$

particularly, if $m = o(n^{-1/2})$ then $\mathbb{P}_n(\mathcal{L}(\hat{f}_n^h) - \mathcal{L}(\hat{f}_n^*)) = o_P(n^{-1/2})$ and it follows that

$$P_0(\mathcal{L}(\hat{f}_n^h) - \mathcal{L}(f_0)) = o_P(n^{-1/2}),$$

as desired. In Appendix C we outline the arguments to show that convergence in terms of the loss-based dissimilarity implies the same convergence rate for the squared $L_2(\pi^* \otimes \mu_0 \otimes \rho)$; i.e., we show that

$$P_0(\mathcal{L}(\hat{f}_n^h) - \mathcal{L}(f_0)) = o_P(n^{-1/2}),$$

implies for $\lambda_{f^h}(t | Z) = \exp(f^h(t, Z))$ that

$$\|\lambda_{f^h} - \lambda_0\|_{\pi^* \otimes \mu_0 \otimes \rho}^2 = o_P(n^{-1/2}),$$

which was exactly what we needed to control the second-order remainder, see Section 4.2 of the main text.

Appendix C

Rate of convergence

It is a general result that the likelihood loss-based dissimilarity, which is really the Kullback-Leibler dissimilarity, behaves as a square of an $L_2(P_0)$ -norm (see, e.g. van der Laan, 2017, Lemma 4) when the density p_0 of the distribution P_0 is bounded. We here repeat the arguments to demonstrate that the highly adaptive lasso convergence in terms of the loss-based dissimilarity also implies the needed convergence to control the second-order remainder (Section 4.2).

The general arguments follow van der Vaart (2000, p. 62) for the density p_0 of the observed data distribution P_0 . Indeed, the Kullback-Leibler dissimilarity for a $P \in \mathcal{M}$ with density p with respect to the same dominating measure ν can be bounded as follows:

$$\int \log \left(\frac{p}{p_0} \right) (o) dP_0(o)$$

here, we use that $\log x \leq 2(\sqrt{x} - 1)$ for $x \geq 0$;

$$\begin{aligned} &\leq 2 \int \left(\frac{\sqrt{p}}{\sqrt{p_0}} - 1 \right) (o) dP_0(o) \\ &= 2 \int \left(\frac{\sqrt{p}}{\sqrt{p_0}} \right) (o) p_0(o) d\nu(o) - 2 \\ &= 2 \int (\sqrt{p}\sqrt{p_0}) (o) d\nu(o) - 2 \\ &\leq - \int (\sqrt{p} - \sqrt{p_0})^2 (o) d\nu(o), \end{aligned} \tag{B.10}$$

here, for the last inequality we used that

$$\begin{aligned}
& \int (\sqrt{p} - \sqrt{p_0})^2(o) d\nu(o) \\
&= \int p(o) d\nu(o) + \int p_0(o) d\nu(o) - 2 \int (\sqrt{pp_0})(o) d\nu(o) \\
&\leq 2 - 2 \int (\sqrt{pp_0})(o) d\nu(o).
\end{aligned}$$

On the other hand we have that

$$\begin{aligned}
\int (p - p_0)^2(o) d\nu(o) &= \int (\sqrt{p} - \sqrt{p_0})^2 \underbrace{(\sqrt{p} + \sqrt{p_0})^2(o)}_{\leq \mathcal{M}'} d\nu(o) \\
&\leq \mathcal{M}' \int (\sqrt{p} - \sqrt{p_0})^2(o) d\nu(o) \leq -\mathcal{M}' \int \log\left(\frac{p}{p_0}\right)(o) dP_0(o),
\end{aligned} \tag{B.11}$$

using (B.10) and that the density is bounded. We see that (B.11) confirms the general claim.

Next, consider our observed data distribution P_0 of $O = (L, A, \tilde{T}, \Delta)$. For the following, we consider the conditional distribution of (\tilde{T}, Δ) given A, L which we shall denote by \tilde{P}_0 with density \tilde{p}_0 , i.e.,

$$\tilde{p}_0(o) = \underbrace{(\lambda_0(t | a, \ell))^\delta S_0(t | a, \ell)}_{=: \tilde{q}_0(t, \delta | a, \ell)} \underbrace{(\lambda_0^c(t | a, \ell))^{1-\delta} S_0^c(t | a, \ell)}_{=: \tilde{g}_0(t, \delta | a, \ell)} = \tilde{q}_0(t, \delta | a, \ell) \tilde{g}_0(t, \delta | a, \ell).$$

For the following, we further denote by

$$p_\lambda(t | a, \ell) = \lambda(t | a, \ell) S(t | a, \ell),$$

the conditional density of the distribution of T . Note that $p_\lambda(t | a, \ell) = \tilde{q}(t, 1 | a, \ell)$.

Repeating the arguments of (B.10) above conditional on fixed a, ℓ yields the following bound in terms of the Kullback-Leibler dissimilarity

$$\sum_{\delta=0,1} \int_0^\tau (\sqrt{\tilde{p}} - \sqrt{\tilde{p}_0})^2(t, \delta | a, \ell) dt \leq - \sum_{\delta=0,1} \int_0^\tau \log\left(\frac{\tilde{p}}{\tilde{p}_0}\right)(t, \delta | a, \ell) \tilde{p}_0(t, \delta | a, \ell) dt.$$

Particularly, note that we only care about the \tilde{q}_0 -factor; due to the factorization of \tilde{p}_0 displayed above, we can act as if \tilde{g}_0 is known, i.e., $\tilde{p} = \tilde{q}\tilde{g}_0$ as well as $\tilde{p}_0 = \tilde{q}_0\tilde{g}_0$. For the left hand side of the above, we thus have that

$$\begin{aligned}
- \sum_{\delta=0,1} \int_0^\tau (\sqrt{\tilde{p}} - \sqrt{\tilde{p}_0})^2(t, \delta | a, \ell) dt &= - \sum_{\delta=0,1} \int_0^\tau (\sqrt{\tilde{q}} - \sqrt{\tilde{q}_0})^2(t, \delta | a, \ell) \tilde{g}_0(t, \delta | a, \ell) dt \\
&\stackrel{*}{\leq} -\tilde{\eta} \sum_{\delta=0,1} \int_0^\tau (\sqrt{\tilde{q}} - \sqrt{\tilde{q}_0})^2(t, \delta | a, \ell) dt,
\end{aligned}$$

where $*$ follows under the assumption that \tilde{g}_0 is bounded away from zero for all a, ℓ by $\tilde{\eta} > 0$. Now

we see that

$$\begin{aligned}
\sum_{\delta=0,1} \int_0^\tau (\tilde{q} - \tilde{q}_0)^2(t, \delta | a, \ell) dt &= \sum_{\delta=0,1} \int_0^\tau (\sqrt{\tilde{q}} - \sqrt{\tilde{q}_0})^2 \underbrace{(\sqrt{\tilde{q}} + \sqrt{\tilde{q}_0})^2}_{\leq \tilde{\mathcal{M}}'}(t, \delta | a, \ell) dt \\
&\leq \tilde{\mathcal{M}}' \sum_{\delta=0,1} \int_0^\tau (\sqrt{\tilde{q}} - \sqrt{\tilde{q}_0})^2(t, \delta | a, \ell) dt \\
&\leq -\tilde{\mathcal{M}}' \tilde{\eta}^{-1} \sum_{\delta=0,1} \int_0^\tau \log\left(\frac{\tilde{p}}{\tilde{p}_0}\right)(t, \delta | a, \ell) \tilde{p}_0(t, \delta | a, \ell) dt,
\end{aligned}$$

i.e., we have that

$$\begin{aligned}
&\int_{\mathcal{L}} \sum_{a=0,1} \left(\sum_{\delta=0,1} \int_0^\tau (\tilde{q} - \tilde{q}_0)^2(t, \delta | a, \ell) dt \right) \pi_0(a | \ell) d\mu_0(\ell) \\
&\leq \tilde{\mathcal{M}}' \tilde{\eta}^{-1} \int_{\mathcal{L}} \sum_{a=0,1} \left(\sum_{\delta=0,1} \int_0^\tau \log\left(\frac{\tilde{p}}{\tilde{p}_0}\right)(t, \delta | a, \ell) \tilde{p}_0(t, \delta | a, \ell) dt \right) \pi_0(a | \ell) d\mu_0(\ell).
\end{aligned} \tag{B.12}$$

By the general result for highly adaptive lasso estimation (van der Laan, 2017) we have the following convergence with respect to the Kullback-Leibler dissimilarity

$$\int_{\mathcal{L}} \sum_{a=0,1} \left(\sum_{\delta=0,1} \int_0^\tau \log\left(\frac{\tilde{p}}{\tilde{p}_0}\right)(t, \delta | a, \ell) \tilde{p}_0(t, \delta | a, \ell) dt \right) \pi_0(a | \ell) d\mu_0(\ell) = o_P(n^{-1/2}), \tag{B.13}$$

particularly, observe that

$$\sum_{\delta=0,1} \int_0^\tau (\tilde{q} - \tilde{q}_0)^2(t, \delta | a, \ell) dt = \int_0^\tau (p_\lambda - p_{\lambda_0})^2(t | a, \ell) dt + \int_0^\tau (S - S_0)^2(t | a, \ell) dt,$$

so that combining (B.12) and (B.13) yields that $\|p_\lambda - p_{\lambda_0}\|_{\pi^* \otimes \mu_0 \otimes \rho} = o_P(n^{-1/4})$. To realize that this also implies that $\|\lambda - \lambda_0\|_{\pi^* \otimes \mu_0 \otimes \rho} = o_P(n^{-1/4})$, observe that

$$\|\lambda - \lambda_0\|_{\pi^* \otimes \mu_0 \otimes \rho} = \int_{\mathcal{L}} \sum_{a=0,1} \left(\int_0^\tau (\lambda - \lambda_0)^2(t | a, \ell) dt \right) \pi^*(a | \ell) \mu_0(\ell) d\nu(\ell)$$

since $p_\lambda = \lambda S$, this is the same as

$$\begin{aligned}
&= \int_{\mathcal{L}} \sum_{a=0,1} \left(\int_0^\tau \left(\frac{p_\lambda}{S} - \frac{p_{\lambda_0}}{S_0} \right)^2(t | a, \ell) dt \right) \pi^*(a | \ell) \mu_0(\ell) d\nu(\ell) \\
&= \int_{\mathcal{L}} \sum_{a=0,1} \left(\int_0^\tau \left(\frac{p_\lambda - p_{\lambda_0}}{S_0} + \frac{p_\lambda}{S} - \frac{p_\lambda}{S_0} \right)^2(t | a, \ell) dt \right) \pi^*(a | \ell) \mu_0(\ell) d\nu(\ell) \\
&= \int_{\mathcal{L}} \sum_{a=0,1} \left(\int_0^\tau \left(\frac{p_\lambda - p_{\lambda_0}}{S_0} + p_\lambda \frac{S_0 - S}{SS_0} \right)^2(t | a, \ell) dt \right) \pi^*(a | \ell) \mu_0(\ell) d\nu(\ell)
\end{aligned}$$

here we can use that S, S_0 are bounded away from zero by some $\eta', \eta'_0 > 0$ on the bounded interval $[0, \tau]$ and that the density is bounded, so that we get the upper bound

$$\leq \int_{\mathcal{L}} \sum_{a=0,1} \left(\int_0^\tau ((p_\lambda - p_{\lambda_0})^2 + \eta_0'^{-1} p_\lambda (S_0 - S))^2(t | a, \ell) dt \right) \pi^*(a | \ell) \mu_0(\ell) d\nu(\ell),$$

so that $\|\lambda - \lambda_0\|_{\pi^* \otimes \mu_0 \otimes \rho} = o_P(n^{-1/4})$ follows from $\|p - p_0\|_{\pi^* \otimes \mu_0 \otimes \rho} = o_P(n^{-1/4})$.

Appendix D

One-step targeted maximum likelihood for the target parameter

The one-step TMLE (Laan et al., 2016; van der Laan and Rose, 2018) uses a universal least favorable submodel to solve the efficient influence curve equation after only one update of the initial estimator. Instead of running the iterative TMLE as in Section 5, we can use the fluctuation model (19) to generate a corresponding universal loss function model. The general recipe for the construction of a universal least favorable submodel that solves the efficient influence curve equation can be summarized largely as follows. Say we have already defined a local least favorable submodel; in our case, this is the fluctuation model defined in (19). Now, instead of solving the score equation of interest (A.4) (Section 5), we track the local least favorable submodel recursively with small step $d\varepsilon$ until the desired influence curve equation is solved. First it is checked which direction ($\pm d\varepsilon$) that decreases the value of the score equation. The first step then moves $d\varepsilon$ in the relevant direction along the fluctuation model (19) evaluated in the initial estimator $\hat{\lambda}_n$. This gives $\hat{\lambda}_n(\cdot; d\varepsilon)$. We then consider the fluctuation model through $\hat{\lambda}_n(\cdot; d\varepsilon)$ and again move $d\varepsilon$ in the relevant direction to obtain $\hat{\lambda}_n(\cdot; 2d\varepsilon)$. This process is continued iteratively, at each step tracking the score equation at zero fluctuation. We stop when we reach k^* such that $\hat{\lambda}_n(\cdot; k^*d\varepsilon)$ solves the efficient influence curve equation. The corresponding TMLE is obtained by plugging in $\hat{\lambda}_n(\cdot; k^*d\varepsilon)$.

References

- Andersen, J. T., M. Petersen, E. Jimenez-Solem, K. Broedbaek, E. W. Andersen, N. L. Andersen, S. Afzal, C. Torp-Pedersen, N. Keiding, and H. E. Poulsen (2013). Trimethoprim use in early pregnancy and the risk of miscarriage: a register-based nationwide cohort study. *Epidemiology & Infection* 141(8), 1749–1755.
- Andersen, P. K., O. Borgan, R. D. Gill, and N. Keiding (1993). *Statistical models based on counting processes*. Springer, New York.
- Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)* 34(2), 187–202.
- Stitelman, O. M. and Wester, C. W. and De Gruttola, V. and van der Laan, M. J. (2011). Targeted maximum likelihood estimation of effect modification parameters in survival analysis. *The international journal of biostatistics* 7(1)1–34.
- Stitelman, O. M. and van der Laan, M. J. (2011). Targeted maximum likelihood estimation of time-to-event parameters with time-dependent covariates. *Technical Report, Division of Biostatistics, University of California, Berkeley*.
- Benkeser, D. and Carone, M. and Gilbert, P. B. (2018). Improved estimation of the cumulative incidence of rare outcomes. *Statistics in medicine* 37(2)280–293.
- Cai, W. and van der Laan, M. J. (2019). One-step targeted maximum likelihood estimation for time-to-event outcomes. *Biometrics*.
- Sofrygin, O. and Zhu, Z. and Schmittdiel, J. A. and Adams, A. S. and Grant, R. W. and van der Laan, M. J. and Neugebauer, R. (2019). Targeted learning with daily EHR data. *Statistics in medicine* 38(16)3073–3090.
- Lindsey, J. K. (1995). Fitting parametric counting processes by using log-linear models. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 44(2)201–212.
- Lu, X. and Tsiatis, A. A. (2008). Improving the efficiency of the log-rank test using auxiliary covariates. *Biometrika* 95(3)679–694.
- Rotnitzky, A. and Lei, Q. and Sued, M. and Robins, J. M. (1997). Improved double-robust estimation in missing data and causal inference models. *Biometrika* 99(2)439–456.
- Rubin, D. B. and van der Laan, M. J. (2008). Empirical efficiency maximization: Improved locally efficient covariate adjustment in randomized experiments and survival analysis. *The International Journal of Biostatistics* 4(1).
- Díaz, I. and Colantuoni, E. and Hanley, D. F. and Rosenblum, M. (2019). Improved precision in the analysis of randomized trials with survival outcomes, without assuming proportional hazards. *Lifetime data analysis* 25(3)439–468.
- Robins, J. M. and Rotnitzky, A. (1992). Recovery of information and adjustment for dependent censoring using surrogate markers. *AIDS epidemiology*, 297–331.
- Benkeser, D. and van der Laan, M. J. (2016). The highly adaptive lasso estimator. *Proceedings of the... International Conference on Data Science and Advanced Analytics. IEEE International Conference on Data Science and Advanced Analytics 2016*.

- Gray, R. J. (1988). A class of K -sample tests for comparing the cumulative incidence of a competing risk. *The Annals of statistics* 16(3)1141–1154.
- Lin, D. Y. and Psaty, B. M. and Kronmal, R. A (1998). Assessing the sensitivity of regression results to unmeasured confounders in observational studies. *Biometrics*, 948–963.
- Grøn, R. and Gerds, T. A. and Andersen, P. K. (1988). Misspecified poisson regression models for large-scale registry data: inference for ‘large n and small p ’. *Statistics in medicine* 35(7)1117–1129.
- van der Laan, M. J. and Gruber, S. (2016). One-step targeted minimum loss-based estimation based on universal least favorable one-dimensional submodels. *The international journal of biostatistics* 12(1)351–378.
- Moore, K. L. and van der Laan, M. J. (2009). Increasing power in randomized trials with right censored outcomes through covariate adjustment. *Journal of biopharmaceutical statistics* 19(6)1099–1131.
- Moore, K. L. and van der Laan, M. J. (2009). Covariate adjustment in randomized trials with binary outcomes: targeted maximum likelihood estimation. *Statistics in medicine* 28(1)39–64.
- Moore, K. L. and van der Laan (2009). Application of time-to-event methods in the assessment of safety in clinical trials. *Design and Analysis of Clinical Trials with Time-to-Event Endpoints*. Taylor & Francis 455–482.
- Hubbard, A. E., van der Laan, M. J. and Robins, J. M. (2000). Nonparametric locally efficient estimation of the treatment specific survival distribution with right censored data and covariates in observational studies. *Statistical Models in Epidemiology, the Environment, and Clinical Trials*, 135–177.
- M. J. van der Laan, S. Dudoit, and van der Vaart, A. W. (2006). The cross-validated adaptive epsilon-net estimator. *Statistics & Decisions* 24(3), 373–395.
- Bang, H. and J. M. Robins (2005). Doubly robust estimation in missing data and causal inference models. *Biometrics* 61(4), 962–973.
- Bibaut, A. F. and M. J. van der Laan (2019, July). Fast rates for empirical risk minimization over càdlàg functions with bounded sectional variation norm. *arXiv e-prints*, arXiv:1907.09244.
- Bickel, P. J., C. A. J. Klaassen, Y. Ritov, and J. A. Wellner (1993). Efficient and adaptive inference in semiparametric models.
- Breiman, L. (1996). Stacked regressions. *Machine learning* 24(1), 49–64.
- Chakraborty, B. and E. E. Moodie (2013). *Statistical methods for dynamic treatment regimes*. Springer.
- Dawid, A. P. and V. Didelez (2010). Identifying the consequences of dynamic treatment strategies: A decision-theoretic overview. *Statistics Surveys* 4, 184–231.
- Fleming, T. R. and D. P. Harrington (2011). *Counting processes and survival analysis*, Volume 169. John Wiley & Sons.

- Gill, R. D. (1994). Lectures on survival analysis. In *Lectures on Probability Theory*, pp. 115–241. Springer.
- Gill, R. D. and S. Johansen (1990). A survey of product-integration with a view toward application in survival analysis. *The annals of statistics* 18(4), 1501–1555.
- Gill, R. D. and J. M. Robins (2001). Causal inference for complex longitudinal data: the continuous case. *Annals of Statistics*, 1785–1811.
- Gill, R. D., M. J. van der Laan, and J. A. Wellner (1995). *Inefficient estimators of the bivariate survival function for three models*, Volume 31. Annales de l’Institut Henri Poincaré.
- Gneiting, T. and A. E. Raftery (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association* 102(477), 359–378.
- Hernán, M. A. (2010). The hazards of hazard ratios. *Epidemiology (Cambridge, Mass.)* 21(1), 13.
- Hernán, M. Á., B. Brumback, and J. M. Robins (2000). Marginal structural models to estimate the causal effect of zidovudine on the survival of hiv-positive men. *Epidemiology*, 561–570.
- Hernán, M. A., B. Brumback, and J. M. Robins (2001). Marginal structural models to estimate the joint causal effect of nonrandomized treatments. *Journal of the American Statistical Association* 96(454), 440–448.
- Hernán, M. A., B. A. Brumback, and J. M. Robins (2002). Estimating the causal effect of zidovudine on cd4 count with a marginal structural model for repeated measures. *Statistics in medicine* 21(12), 1689–1709.
- Hernán, M. A., E. Lanoy, D. Costagliola, and J. M. Robins (2006). Comparison of dynamic treatment regimes via inverse probability weighting. *Basic & clinical pharmacology & toxicology* 98(3), 237–242.
- Hernan, M. A. and J. M. Robins (2020). *Causal Inference*. Boca Raton, FL: Chapman & Hall/CRC.
- Karim, M. E., P. Gustafson, J. Petkau, H. Tremlett, the Long-Term Benefits, and A. E. of Beta-Interferon for Multiple Sclerosis (BeAMS) Study Group (2016). Comparison of statistical approaches for dealing with immortal time bias in drug effectiveness studies. *American journal of epidemiology* 184(4), 325–335.
- Kessing, L. V., H. C. Rytgaard, T. A. Gerds, M. Berk, C. T. Ekstrøm, and P. K. Andersen (2019). New drug candidates for depression – a nationwide population-based study. *Acta Psychiatrica Scandinavica* 139(1), 68–77.
- Lendle, S. D., J. Schwab, M. L. Petersen, and M. J. van der Laan (2017). ltmle: An R package implementing targeted minimum loss-based estimation for longitudinal data. *Journal of Statistical Software* 81(1), 1–21.
- Lok, J. J. (2008). Statistical modeling of causal effects in continuous time. *The Annals of Statistics* 36(3), 1464–1507.
- Martinussen, T., S. Vansteelandt, and P. K. Andersen (2018). Subtleties in the interpretation of hazard ratios. *arXiv preprint arXiv:1810.09192*.

- Murphy, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 65(2), 331–355.
- Murphy, S. A. (2005). An experimental design for the development of adaptive treatment strategies. *Statistics in medicine* 24(10), 1455–1481.
- Murphy, S. A., M. J. van der Laan, J. M. Robins, and C. P. P. R. Group (2001). Marginal mean models for dynamic regimes. *Journal of the American Statistical Association* 96(456), 1410–1423.
- Petersen, M., J. Schwab, S. Gruber, N. Blaser, M. Schomaker, and M. van der Laan (2014). Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. *Journal of causal inference* 2(2), 147–185.
- Robins, J. (1986). A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical modelling* 7(9-12), 1393–1512.
- Robins, J. (1989a). The control of confounding by intermediate variables. *Statistics in medicine* 8(6), 679–701.
- Robins, J. (1992). Estimation of the time-dependent accelerated failure time model in the presence of confounding factors. *Biometrika* 79(2), 321–334.
- Robins, J., L. Orellana, and A. Rotnitzky (2008). Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in medicine* 27(23), 4678–4721.
- Robins, J. M. (1987). Addendum to “a new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect”. *Computers & Mathematics with Applications* 14(9-12), 923–945.
- Robins, J. M. (1989b). The analysis of randomized and non-randomized aids treatment trials using a new approach to causal inference in longitudinal studies. *Health service research methodology: a focus on AIDS*, 113–159.
- Robins, J. M. (1998). Marginal structural models. 1997 proceedings of the american statistical association, section on bayesian statistical science (pp. 1–10). Retrieved from.
- Robins, J. M. (2000a). Marginal structural models versus structural nested models as tools for causal inference. In *Statistical models in epidemiology, the environment, and clinical trials*, pp. 95–133. Springer.
- Robins, J. M. (2000b). Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association*, Volume 1999, pp. 6–10.
- Robins, J. M. (2002). Analytic methods for estimating hiv-treatment and cofactor effects. In *Methodological Issues in AIDS Behavioral Research*, pp. 213–288. Springer.
- Robins, J. M., M. A. Hernan, and B. Brumback (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology* 11(5), 550–560.
- Robins, J. M., M. A. Hernán, and U. Siebert (2004). Effects of multiple interventions. *Comparative quantification of health risks: global and regional burden of disease attributable to selected major risk factors* 1, 2191–2230.

- Røysland, K. (2011). A martingale approach to continuous-time marginal structural models. *Bernoulli* 17(3), 895–915.
- Røysland, K. (2012). Counterfactual analyses with graphical models based on local independence. *The Annals of Statistics* 40(4), 2162–2194.
- Stitelman, O. M., V. De Gruttola, and M. J. van der Laan (2012). A general implementation of tmle for longitudinal data applied to causal inference in survival analysis. *The international journal of biostatistics* 8(1).
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* 58(1), 267–288.
- Tsiatis, A. (2007). *Semiparametric theory and missing data*. Springer Science & Business Media.
- van der Laan, M. J. (2010a). Targeted maximum likelihood based causal inference: Part I. *The International Journal of Biostatistics* 6(2).
- van der Laan, M. J. (2010b). Targeted maximum likelihood based causal inference: Part II. *The international journal of biostatistics* 6(2).
- van der Laan, M. J. (2017). A generally efficient targeted minimum loss based estimator based on the highly adaptive lasso. *The International Journal of Biostatistics* 13(2).
- van der Laan, M. J. and S. Dudoit (2003). Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples.
- van der Laan, M. J. and S. Gruber (2012). Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The international journal of biostatistics* 8(1).
- van der Laan, M. J. and M. L. Petersen (2007). Causal effect models for realistic individualized treatment and intention to treat rules. *The international journal of biostatistics* 3(1).
- van der Laan, M. J., E. C. Polley, and A. E. Hubbard (2007). Super learner. *Statistical applications in genetics and molecular biology* 6(1).
- van der Laan, M. J. and J. M. Robins (2003). *Unified methods for censored longitudinal data and causality*. Springer Science & Business Media.
- van der Laan, M. J. and S. Rose (2011). *Targeted learning: causal inference for observational and experimental data*. Springer Science & Business Media.
- van der Laan, M. J. and S. Rose (2018). *Targeted learning in data science: causal inference for complex longitudinal studies*. Springer.
- van der Laan, M. J. and D. Rubin (2006). Targeted maximum likelihood learning. *The International Journal of Biostatistics* 2(1).
- van der Vaart, A. W. (2000). *Asymptotic statistics*, Volume 3. Cambridge university press.
- van der Vaart, A. W., S. Dudoit, and M. J. van der Laan (2006). Oracle inequalities for multi-fold cross validation. *Statistics & Decisions* 24(3), 351–371.

- van der Vaart, A. W. and J. A. Wellner (1996). Weak convergence. In *Weak convergence and empirical processes*, pp. 16–28. Springer.
- Wolpert, D. H. (1992). Stacked generalization. *Neural networks* 5(2), 241–259.
- Zhang, B., A. A. Tsiatis, E. B. Laber, and M. Davidian (2012). A robust method for estimating optimal treatment regimes. *Biometrics* 68(4), 1010–1018.
- Zhang, B., A. A. Tsiatis, E. B. Laber, and M. Davidian (2013). Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika* 100(3), 681–694.