University of California, Berkeley
U.C. Berkeley Division of Biostatistics Working Paper Series

Year 2007 *Paper* 226

A Note on Targeted Maximum Likelihood and Right Censored Data

Mark J. van der Laan* Daniel Rubin†

*Division of Biostatistics, School of Public Health, University of California, Berkeley, laan@berkeley.edu

†University of California, Berkeley, drubin@stat.berkeley.edu

A Note on Targeted Maximum Likelihood and Right Censored Data

Mark J. van der Laan and Daniel Rubin

Abstract

A popular way to estimate an unknown parameter is with substitution, or evaluating the parameter at a likelihood based fit of the data generating density. In many cases, such estimators have substantial bias and can fail to converge at the parametric rate. van der Laan and Rubin (2006) introduced targeted maximum likelihood learning, removing these shackles from substitution estimators, which were made in full agreement with the locally efficient estimating equation procedures as presented in Robins and Rotnitzsky (1992) and van der Laan and Robins (2003). This note illustrates how targeted maximum likelihood can be applied in right censored data structures. In particular, we show that when an initial substitution estimator is based on a Cox proportional hazards model, the targeted likelihood algorithm can be implemented by iteratively adding an appropriate time-dependent covariate.

1 Introduction

Suppose we observe a sample $\{O_i\}_{i=1}^n$ of independent and identically distributed observations, for

$$O = (W, \Delta = I(T \leq C), \tilde{T} = \min(T, C)) \sim P \in \mathcal{M}. \quad (1)$$

Here W is a vector of baseline covariates, T is a survival time, C is a censoring time, Δ is an indicator of censoring, P is the data generating distribution, and the statistical model $\mathcal{M}$ is a family of data generating distributions containing P. We will make the usual assumption that

$$\{T \perp C | W\}, \quad (2)$$

meaning survival and censoring times are conditionally independent given the baseline covariates. The log likelihood for a single observation can be written as

$$\begin{aligned} dP(w, \delta, \tilde{t}) &= dP(W = w) \\ &\times [dP(T = \tilde{t}|W = w)P(C \geq \tilde{t}|W = w)]^{\delta} \\ &\times [P(T > \tilde{t}|W = w)dP(C = \tilde{t}|W = w)]^{1-\delta}. \end{aligned} \quad (3)$$

The full data, which would have liked to observe, but could not be completely measured because of censoring, consists of the baseline covariates and survival times $\{X_i\}_{i=1}^n = \{W_i, T_i\} \sim F$. We can write $P = P_{F,G}$, for $G(\cdot|W)$ denoting the conditional cumulative distribution function of the censoring time. This note applies to general scenarios where the goal is to estimate a smooth (pathwise differentiable) Euclidean parameter $\mu(F) \in \mathbb{R}^k$, representing some feature of the full data distribution.

An example of such a parameter is simply the marginal survival probability at a fixed time t,

$$\mu(P_{F,G}) = \mu(F) = \bar{F}(t) = P(T > t). \quad (4)$$

Note that if the stronger unconditional independence assumption $\{T \perp C\}$ doesn't hold, the (1958) Kaplan-Meier estimator might not necessarily be consistent. Even if Kaplan-Meier assumptions aren't violated, the presence of informative baseline covariates make efficiency gains possible. When a randomized treatment $A \in \{0, 1\}$ is assigned at baseline, another important parameter could be the risk difference

$$\mu(F) = \bar{F}(t|A = 1) - \bar{F}(t|A = 0) = P(T > t|A = 1) - P(T > t|A = 0). \quad (5)$$

Additionally, interest might lie in regression parameters such as

$$\mu(F) = (\beta_0(F), \beta(F)) = \text{argmin}_{\beta_0,\beta} E_F|\log(T) - \beta_0 - \beta^T W|^2. \tag{6}$$

Note that $\mu(F)$ can here be defined without assuming an accelerated failure time model actually holds. It is simply a coefficient vector giving the best linear predictor of log survival from baseline covariates.

For many parameters of interest, the prevailing estimation technique is to first fit the data generating distribution $P \in \mathcal{M}$ with some $\hat{P} \in \mathcal{M}$ according to maximizing likelihood over a submodel $\mathcal{M}_0 \subset \mathcal{M}$, and then forming the substitution estimator $\hat{\mu} = \mu(\hat{P})$. This note will focus on how to proceed when initially considering a substitution estimator based on the ubiquitous proportional hazards model introduced in Cox (1973). Unfortunately, substitution estimators often have poor performance. As discussed in Robins and Ritov (1997), they can be heavily biased because the choice of $\hat{P} \in \mathcal{M}$ was made without regard to the parameter of interest. Such estimators can be inconsistent, or lead to arbitrarily bad rates of convergence, while simpler schemes can sometimes guarantee the parametric $n^{-1/2}$ rate.

For example, Robins and Rotnitzky (2005) review inverse probability of censoring weighted (IPCW) estimators in survival analysis, which can lead to $\sqrt{n}$-consistent, asymptotically linear estimators if the censoring mechanism $\bar{G}(\cdot|W)$ can be well approximated. Frequently censoring is caused by study termination, and the censoring time is independent of the survival time and baseline covariates, in which case $\bar{G}(\cdot|W) = \bar{G}(\cdot)$ can be efficiently estimated with the Kaplan-Meier curve.

Suppose $D_{\text{Full}}(W,T|F) = D_{\text{Full}}(W,T|\mu(F),\eta(F)) : (W,T) \to \mathbb{R}^k$ is an estimating function for μ we could use with access to the full data $\{X_i = (W_i, T_i)\}_{i=1}^n$. That is, suppose $E_F[D_{\text{Full}}(W,T|\mu,\eta(F))] = 0$ at $\mu = \mu(F)$, and that with no censoring we could reliably estimate $\mu(F)$ with the solution to

$$0 = \frac{1}{n}\sum_{i=1}^{n} D_{\text{Full}}(W_i, T_i|\mu, \eta_n). \tag{7}$$

Here the left side is the zero vector in $\mathbb{R}^k$, and η_n is an estimator of the nuisance parameter $\eta(F)$. For the three parameters given by (4), (5), and (6), respective full data estimating equations could be

$$D_{\text{Full}}(W,T|F) = I(T > t) - \mu(F) \tag{8}$$

$$D_{\text{Full}}(W,T|F) = I(T > t)\left(\frac{A}{P(A=1)} - \frac{1-A}{P(A=0)}\right) - \mu(F) \tag{9}$$

$$D_{\text{Full}}(W,T|F) = W(\log(T) - \beta_0 - \beta^T W) \text{ recalling } \mu(F) = [\beta_0, \beta]^T. \tag{10}$$

Inverse probability of censoring weighted estimation maps the full data estimating function into

$$D_{\mathrm{IPCW}}(O|P) = D_{\mathrm{IPCW}}(O|\mu, \eta(F), G) = \frac{D_{\mathrm{Full}}(O|\mu, \eta(F))\Delta}{\bar{G}(\tilde{T}_{-}|W)}, \tag{11}$$

which is a function of the observed data $O = (W, \Delta, \tilde{T})$. It is easy to verify $E_P[D_{\mathrm{IPCW}}(O|P)] = E_F[D_{\mathrm{Full}}(W, T|F)] = 0$. Hence, we can use it as an estimating equation for $\mu(F)$, after fitting nuisance parameters $\eta(F)$ and $G(\cdot|W)$. While simple, IPCW estimating equations are suboptimal in terms of both efficiency and robustness. We refer to van der Laan and Robins (2003) for a survey of estimating function methodology in survival analysis.

Despite advantages of the estimating function methodology outlined in this survey, likelihood based substitution estimators remain more prevalent in many applications. This could be for a variety of reasons, among them outlier concerns due to inverse weighting, computational considerations, unfamiliarity, and inertia. To remedy the situation, van der Laan and Rubin (2006) introduced targeted maximum likelihood. Given an initial fit $\hat{P}$ of the data generating distribution, the procedure iteratively updates the fit by maximizing likelihood along submodels chosen to best target the parameter of interest $\mu(F)$. The algorithm maps an initial $\hat{P}$ into a $\hat{P}^{\star} \in \mathcal{M}$, at which the substitution estimator $\mu(\hat{P}^{\star})$ is also the solution to a well-chosen estimating equation. Hence, the resulting estimator is a familiar type of likelihood based substitution estimator, inheriting the benefits of $\sqrt{n}$-convergence, asymptotic linearity, and local efficiency implied by estimating function theory. Targeted maximum likelihood works as follows:

1. Form an initial fit $\hat{P} \in \mathcal{M}$ of the data generating distribution.

2. Create a smooth (regular) parametric submodel of $\mathcal{M}$, parametrized by an ϵ, passing through $\hat{P}$ at $\epsilon = 0$. Ensure the linear span of the score vector at $\hat{P}$ includes the efficient influence curve for parameter $\mu(P)$ at $\hat{P}$. The efficient influence curve will be discussed in the sequel, and is formally discussed in Bickel et al. (1998) and Chapter 1.4 of van der Laan and Robins (2003).

3. Estimate ϵ with maximum likelihood.

4. Define a new density estimator as the corresponding update to the original estimator $\hat{P}$.

5. Iterate steps 2-4 until convergence. Of course, the procedure can be applied without iteration, and van der Laan and Rubin (2006) argued that most bias reduction should occur in the first step.

The efficient influence curve $D(O|P) = D(O|\mu, \eta, G, F)$, or a scaled version thereof, is in a strong sense the optimal estimating equation for the parameter of interest. If the nuisance parameters on which it depends are estimated accurately, and regularity conditions are met, the estimating equation gives rise to the regular asymptotically linear estimator with the smallest possible asymptotic variance. Further, as discussed in van der Laan and Robins (2003), it has desirable robustness properties. If either the initial full data fit $\hat{F}$ is a good approximation to F, or the censoring mechanism estimate $\hat{G}$ is a good approximation to G, using the estimating function $D(O|\mu, \eta, \hat{G}, \hat{F})$ can ensure asymptotic linearity.

Suppose no special parametric or semiparametric assumptions are made on the full data model, and the $D_{\text{Full}}(W, T|F)$ given earlier would be a valid estimating equation with uncensored data. Robins and Rotnitzsky (1992) show the efficient influence curve at $P_{F,G}$ is given by a scaled version of,

$$\begin{aligned} D(O|P) &= D(O|\mu(P), \eta(P), F(P), G(P)) \qquad (12)\\ &= D_{\text{IPCW}}(O|\mu, \eta, G) + \int_t \frac{E_F[D_{\text{Full}}(W, T|\mu, \eta)|W, T > t]}{\bar{G}(\tilde{T}|W)} dM_G(t). \end{aligned}$$

Here the last term is an integral with respect to the martingale,

$$M_G(t) = I(\tilde{T} \le t, \Delta = 0) - \int_{-\infty}^{t} I(\tilde{T} \ge s) \frac{dG(s|W)}{\bar{G}(s_-|W)}. \qquad (13)$$

Examine the targeted likelihood algorithm. Upon convergence to $\hat{P}^\star$, the relevant submodel's likelihood is maximized at $\epsilon = 0$. Hence, the score at $\epsilon = 0$ will have empirical mean zero. But from the choice of submodel, this means the efficient influence curve $D(O|\hat{P}^\star) = D(O|\mu(\hat{P}^\star), \eta(\hat{P}^\star), F(\hat{P}^\star), G(\hat{P}^\star))$ will have empirical mean zero. In other words, the substitution estimator $\mu(\hat{P}^\star)$ will solve the efficient influence curve estimating equation, based on plug-in estimators for the curve's nuisance parameters.

This note is devoted to showing how targeted likelihood can be implemented when the initial fit is based on Cox's proportional hazards model. The initial fit to the data generating distribution can be decomposed into fits of the baseline covariate distribution P_W, the censoring mechanism $G(|W)$ representing the conditional distribution $\mathcal{L}(C|W)$, and the conditional survival distribution $\mathcal{L}(T|W)$. As we'll mention in Section 3, it will be convenient to use the empirical distribution placing mass $\frac{1}{n}$ on $W_1, ..., W_n$ to fit the baseline covariate distribution. As previously observed, the censoring mechanism can be fit with the Kaplan-Meier product-limit estimator if we believe censoring is independent of survival, but arbitrary initial fits can be used. Neither the baseline covariate distribution fit not the censoring mechanism fit will be updated at any step of the targeted likelihood algorithm. The Cox model is meant for estimating the conditional survival distribution. Note that the methodology can applied without necessarily believing the model holds, and targeted likelihood can allow us to consistently estimate parameters such as (4), (5), and (6) using a misspecified model. We'll consider a variant of the model assuming,

$$\Lambda(t|W) = \int_{-\infty}^{t} \frac{dF(s|W)}{\bar{F}(s_-|W)} = \Lambda_0(t)\exp(\beta^T L(W,t)). \tag{14}$$

Here $L(\cdot, W)$ is a specified function allowing multiplicative effect on conditional hazard to change with time. Coefficient vector β can be estimated by $\hat{\beta}$ through maximizing Cox's (1973) partial likelihood, while the Breslow (1974) estimator $\hat{\Lambda}_0(\cdot)$ is commonly used to fit the baseline cumulative hazard function $\Lambda_0(\cdot)$. Together, these fits determine a fit $\hat{\Lambda}(\cdot|W)$ of the conditional cumulative hazard $\Lambda(\cdot|W)$, and consequently the conditional survival distribution $\mathcal{L}(T|W)$. Taken together, $\hat{P}_W$, $\hat{G}(\cdot|W)$ and $\hat{\Lambda}(\cdot|W)$ determine the initial fit $\hat{P}$ of the data generating distribution. This is step 1 of the targeted likelihood algorithm, and it remains to be seen how $\hat{P}$ can be mapped into the $\hat{P}^\star$ providing an accurate substitution estimator for the parameter of interest.

2 Statement of Main Result

The targeted likelihood algorithm can be implemented by iteratively adding an appropriate time-dependent covariate to the Cox proportional hazards model. Letting $\hat{P}$ denote the initial data generating fit just mentioned, and $\bar{G}_n(\cdot|W) = 1 - \hat{G}(\cdot|W)$ the corresponding censoring mechanism fit, define the function

$$h(w,t|\hat{P}) = \frac{D_{\text{Full}}(w,t|\hat{P}) - E_{\hat{P}}[D_{\text{Full}}(w,T|\hat{P})|W=w, T>t]}{\bar{G}_n(t_-|w)}. \tag{15}$$

For fixed baseline cumulative hazard fit $\hat{\Lambda}(\cdot)$ and coefficient vector fit $\hat{\beta}$, consider the submodel

$$\Lambda_\epsilon(t|W) = \hat{\Lambda}_0(t)\exp(\hat{\beta}^T L(W,t) + \epsilon^T h(W,t|\hat{P})), \tag{16}$$

parametrized by $\epsilon \in \mathbb{R}^k$. Here ϵ has the same dimension as the parameter $\mu(F)$ and efficient influence curve $D(O|P)$.

Choosing $\hat{\epsilon}$ to maximize the likelihood of observed data $\{O_i\}_{i=1}^n$ corresponds to carrying out an iteration of the targeted maximum likelihood algorithm. The remainder of this note sketches the argument, without attempting to be overly formal.

When the data generating distribution fit $\hat{P}$ is updated based on the fit to this model, the procedure can be iterated. Hence, iteration corresponds to repeatedly adding a time-dependent covariate vector to an existing proportional hazards model, and using maximum likelihood to fit the associated coefficient vector while keeping everything else in the model fixed. Standard software can be used to fit ϵ via maximum likelihood, but this will require being able to evaluate covariate $h(W, t|\hat{P})$, which could be cumbersome due to the conditional expectation in its second term.

3 Sketch of Argument that Adding Covariate Implements Targeted Likelihood Algorithm

Following Bickel et. al. (1998), we can define the tangent space $T(P)$ as the closure in $L_0^2(P)$ of the linear span of all scores of regular parametric submodels of $\mathcal{M}$ through P. It is well known that if the model is nonparametric, the tangent space is saturated, meaning that $T(P) = L_0^2(P)$. It is also easy to see the tangent space can be decomposed into the three tangent spaces corresponding to scores through P fluctuating the baseline covariate distribution $\mathcal{L}(W)$, conditional survival distribution $\mathcal{L}(T|W)$, and censoring mechanism $\mathcal{L}(C|W)$. These three tangent spaces

$$T_{\mathrm{W}}(P) = \{r(W) \in L_0^2(P) : \; E[r(W)] = 0\} \quad (17)$$
$$T_{\mathrm{F}}(P) = \{v(O) \in L_0^2(P) : \; E[v(O)|C, W] = 0\} \quad (18)$$
$$T_{\mathrm{CAR}}(P) = \{v(O) \in L_0^2(P) : \; E[v(O)|T, W] = 0\} \quad (19)$$

are orthogonal, giving us the direct sum

$$T(P) = L_0^2(P) = T_{\mathrm{W}}(P) \oplus T_{\mathrm{F}}(P) \oplus T_{\mathrm{CAR}}(P) \quad (20)$$

and the decomposition of any $v(O) \in L_0^2(P)$ into

$$\begin{aligned} v(O) &= \Pi(v(\cdot)|T(P))(O) \quad (21) \\ &= \Pi(v(\cdot)|T_{\mathrm{W}}(P))(O) + \Pi(v(\cdot)|T_{\mathrm{F}}(P))(O) + \Pi(v(\cdot)|T_{\mathrm{CAR}}(P))(O). \end{aligned}$$

This decomposition can be applied to the efficient influence curve $D(O|P)$. To find a submodel through P with score equal to this influence curve, it is thus only necessary to find submodels varying $\mathcal{L}(W)$, $\mathcal{L}(T|W)$, $\mathcal{L}(C|W)$ that give $\Pi(D(\cdot|P)|T_{\mathrm{W}}(P))(O)$, $\Pi(D(\cdot|P)|T_{\mathrm{F}}(P))(O)$, and $\Pi(D(\cdot|P)|T_{\mathrm{CAR}}(P))(O)$ as their respective scores.

3.1 Baseline Covariate Distribution

Letting $\hat{P}_W$ denote the empirical distribution on the baseline covariates $\{W_i\}_{i=1}^n$ given as the initial fit in the previous section, and $\hat{P}$ the initial fit for the entire data generating distribution P, we can trivially define the submodel

$$dP_W^{(\delta)} = \frac{\exp(\delta\Pi(D(O|\hat{P})|T_W(\hat{P})))}{\int \exp(\delta\Pi(D(O|\hat{P})|T_W(\hat{P}))) d\hat{P}_W^{(\delta)}} d\hat{P}_W. \tag{22}$$

The projection operator is given by $\Pi(v(O)|T_W(\hat{P})) = E_{\hat{P}}[v(O)|W]$, but this will not be relevant for our purposes. It can be verified that this submodel gives the desired score of $\Pi(D(O|\hat{P})|T_{\mathrm{W}}(\hat{P}))$.

In fact, the exponential family technique can always be used to define a submodel of a nonparametric model having a desired score. We could have simply used the exponential family $dP^{(\delta)}(O) \propto \exp(\delta D(O|\hat{P}) d\hat{P}(O)$ for the entire data generating distribution, but targeted likelihood becomes more difficult to implement than in our Cox model formulation.

The specific choice of submodel through P_W is not at all important for the targeted likelihood procedure, so long as it gives rise to the correct score. This is because $\hat{P}_W$ is never updated from its initial empirical distribution fit, as this is the nonparametric maximum likelihood estimate (NPMLE) for P_W. Consequently, in each iteration of the targeted likelihood algorithm, the $\hat{P}_k$ to be used as a substitution estimator corresponds to using the empirical distribution baseline covariate fit.

We mean to focus attention on when the survival distribution, meaning the marginal $\mathcal{L}(T)$ or conditional $\mathcal{L}(T|W)$ law, is of primary interest, rather than the baseline covariate distribution $\mathcal{L}(W)$. If there is concern substitution estimation of $\mu(F)$ based on the empirical $\hat{P}_W$ might lead us astray, the problem would have to be reconsidered.

3.2 Censoring Mechanism

As discussed in Chapter 1.4.4 of van der Laan and Robins (2003), the efficient influence curve $D(O|P)$ is orthogonal to the tangent space T_{CAR} generated from scores of submodels varying the censoring mechanism $\mathcal{L}(C|W)$. Hence, $\Pi(D(\cdot|P)|T_{\mathrm{CAR}}(P)) = 0$, and we do not need to perturb the censoring mechanism from its initial fit in the targeted maximum likelihood algorithm.

3.3 Conditional Survival Time Distribution

Note from Chapter 1.4 of van der Laan and Robins (2003) that the efficient influence curve at P can be written as

$$D(O|P) = D_{\mathrm{IPCW}}(O|P) - \Pi(D(\cdot|P)|T_{\mathrm{CAR}}(P))(O), \tag{23}$$

and that $T_{\mathrm{F}}(P)$ is orthogonal to $T_{\mathrm{CAR}}(P)$. Together these facts clearly imply

$$\Pi(D(\cdot|P)|T_{\mathrm{F}}(P)) = \Pi(D_{\mathrm{IPCW}}(\cdot|P)|T_{\mathrm{F}}(P)). \tag{24}$$

Thus, we only need to show the submodel through the $\mathcal{L}(T|W)$ fit in the previous section gives rise to a score equal to the IPCW estimating function's projection on tangent space $T_{\mathrm{F}}(P)$.

Define the counting process $N(t) = I(\tilde{T} \leq t, \Delta = 1)$ jumping at an observed failure time. Recalling $\Lambda(\cdot|W)$ represents the conditional cumulative hazard function for $\mathcal{L}(T|W)$, the associated Doob-Meyer martingale is

$$M(t) = N(t) - \int_{-\infty}^{t} I(\tilde{T} \geq s) d\Lambda(s|W). \tag{25}$$

From Theorem 1.1 of van der Laan and Robins (2003), interchanging the completely symmetric $T_{\mathrm{CAR}}(P)$ and $T_{\mathrm{F}}(P)$, the projection operator is given by

$$\Pi(v|T_{\mathrm{F}}(P)) = \int (E_P[v(O)|W, T = t, C \geq t] - E_P[v(O)|W, T > t, C \geq t]) dM(t). \tag{26}$$

We can apply this result with $v(O) = D_{\mathrm{IPCW}}(O) = \frac{D_{\mathrm{Full}}(W,T|P)\Delta}{\bar{G}(\tilde{T}_-|W)}$. Given that $\{T = t, C \geq t\}$ implies $\Delta = 1$, it is clear $E_P[v(O)|W, T = t, C \geq t] = \frac{D_{\mathrm{Full}}(W,t|P)}{\bar{G}(t_-|W)}$. Further, it is an elementary calculation to show $E_P[v(O)|W, T > t, C \geq t]$ is equal to $E_P[D_{\mathrm{Full}}(W, T|P)|W, T > t]/\bar{G}(t_-|W)$. Hence, the efficient influence curve $D(O|P)$ has projection on tangent space $T_{\mathrm{F}}(P)$ of

$$\Pi(D(\cdot|P)|T_{\mathrm{F}}(P)) = \int h(W, t|P) dM(t), \tag{27}$$

for the $h(W, t|P)$ defined in (15). However, as reviewed in Lemma 3.2 of van der Laan and Robins (2003), $\int g(W, t) dM(t)$ is simply the score at $\epsilon = 0$

of a submodel through P varying conditional cumulative hazard of $\mathcal{L}(T|W)$ through

$$\Lambda_\epsilon(t|W) = \Lambda(t|W)\exp(\epsilon^T g(W,t|P)). \quad (28)$$

Thus, the projection $\Pi(D(\cdot|\hat{P})|T_{\mathrm{F}}(\hat{P}))$ in $L_0^2(\hat{P})$ is exactly the score at $\epsilon = 0$ of the submodel (16). Recall that this was the desired result, from our decomposition of $D(O|P)$ into projections on $T_{\mathrm{W}}(P)$, $T_{\mathrm{CAR}}(P)$ and $T_{\mathrm{F}}(P)$. By adding $h(W,t|\hat{P})$ as a time-dependent covariate to a Cox model, fixing the censoring mechanism fit, and placing a submodel through the baseline covariate empirical distribution fit, we can obtain the efficient influence curve as a score. Because the baseline covariate fit will never be perturbed, targeted likelihood proceeds by iteratively updating the initial Cox model fit.

4 Discussion

In this note, we've shown how a Cox-based substitution estimator can be made to solve a locally efficient estimating equation, if appropriate covariates are added to an initial fit. Estimating equation approaches are often avoided in favor of more familiar substitution estimators, despite their theoretical advantages outlined in van der Laan and Robins (2003). By representing estimating function procedures as fits to commonplace Cox models, we hope to make the methodology more amenable. This parallels results given in van der Laan and Rubin (2006) and Moore and van der Laan (2007) demonstrating the targeted likelihood algorithm can be implemented in causal inference problems by adding covariates to linear and logistic regression models, although in those cases the algorithm was shown to converge in a single iteration.

Several serious caveats are in order. Primarily, while we've suggested how to perform targeted maximum likelihood, our exposition was hardly a formal proof. Further, van der Laan and Rubin (2006) listed several criteria to ensure convergence of the iterative algorithm, which have not been checked in this work, although we expect them to hold. Finally, while it sounds straightforward to iteratively add a time-dependent covariate to a Cox model, we have glossed over the specific details of how to implement our procedure.

Bembom et al. (2007) showed targeted likelihood estimates of variable importance measures could enhance biomarker discovery procedures. We have here introduced similar locally efficient doubly robust estimators suitable for right censored data structures, and also expect benefits to become apparent in real world applications.

References

[1] Bembom, B., Petersen, M.L., Rhee, S., Fessel, J.W., Sinisi, S.E., Shafer, R.W., and van der Laan, M.J. (2007). Biomarker Discovery Using Targeted Maximum Likelihood Estimation: Application to the Treatment of Antiretroviral Resistant HIV Infection. *U.C. Berkeley Division of Biostatistics Working Paper Series*, Working Paper 221.

[2] Bickel, P.J., Klaassen, C.A.J., Ritov, Y., and Wellner, J.A (1998). *Efficient and Adaptive Estimation for Semiparametric Models.* Springer-Verlag, New York.

[3] Breslow, N.E. (1974). Covariance analysis of censored survival data. *Biometrics.* 30, 89-99.

[4] Cox, D.R. (1972). Regression models and life tables (with discussion). *Journal of the Royal Statistical Society B,* 34, 187-220.

[5] Kaplan, E.L., and P. Meier. (1958). Nonparametric estimator from incomplete observations. *Journal of the American Statistical Association.* 53, 457-481.

[6] Moore, K.L. and van der Laan M.J. (2007). Covariate Adjustment in Randomized Trials with Binary Outcomes: Targeted Maximum Likelihood Estimation. *U.C. Berkeley Division of Biostatistics Working Paper Series*, Working Paper 215.

[7] Robins, J.M. and Ritov, Y. (1997). Toward a curse of dimensionality appropriate (CODA) asymptotic theory for semi-parametric models. *Statistics in Medicine,* 16, 285-319.

[8] Robins, J.M. and Rotnitzky, A. (1992). Recovery of information and adjustment for dependent censoring using surrogate markers. *AIDS Methodology - Methodological Issues.* Jewell, N., Dietz, K, and Farewell, W., eds. Birkhäuser, Boston, 297-331.

[9] Robins J.M. and Rotnitzky, A. (2005). Inverse probability weighted estimation in survival analysis. *Encyclopedia of Biostatistics*, Second Edition, Editors: Armitage, P. and Colton, T., Wiley & Sons, New York.

[10] van der Laan, M.J. and Robins, J.M. (2003). *Unified Methods for Censored Longitudinal Data and Causality.* Springer-Verlag, New York.

[11] van der Laan, M.J. and Rubin, D.B. (2006). Targeted Maximum Likelihood Learning. *The International Journal of Biostatistics,* Vol. 2, Iss. 1, Article 11.