The International Journal of Biostatistics

Volume 8, Issue 1

2012

Article 11

Targeted Minimum Loss Based Estimator that Outperforms a given Estimator

Susan Gruber, Harvard University
Mark J. van der Laan, University of California - Berkeley

Recommended Citation:

Gruber, Susan and van der Laan, Mark J. (2012) "Targeted Minimum Loss Based Estimator that Outperforms a given Estimator," *The International Journal of Biostatistics*: Vol. 8: Iss. 1, Article 11

DOI: 10.1515/1557-4679.1332

©2012 De Gruyter. All rights reserved.

Targeted Minimum Loss Based Estimator that Outperforms a given Estimator

Susan Gruber and Mark J. van der Laan

Abstract

Targeted minimum loss based estimation (TMLE) provides a template for the construction of semiparametric locally efficient double robust substitution estimators of the target parameter of the data generating distribution in a semiparametric censored data or causal inference model (van der Laan and Rubin (2006), van der Laan (2008), van der Laan and Rose (2011)). In this article we demonstrate how to construct a TMLE that also satisfies the property that it is at least as efficient as a user supplied asymptotically linear estimator. In particular it is shown that this type of TMLE can incorporate empirical efficiency maximization as in Rubin and van der Laan (2008), Tan (2008, 2010), Rotnitzky et al. (2012), and retain double robustness. For the sake of illustration we focus on estimation of the additive average causal effect of a point treatment on an outcome, adjusting for baseline covariates.

KEYWORDS: Asymptotic linearity of an estimator, causal effect, efficient influence curve, empirical efficiency maximization, confounding, G-computation formula, influence curve, loss function, nonparametric structural equation model, positivity assumption, randomization assumption, randomized trial, semiparametric statistical model, targeted maximum likelihood estimation, targeted minimum loss based estimation

Author Notes: This work was funded by NIH-grant R01 AI074345-04.

1 Introduction

Targeted minimum loss based estimation (TMLE) provides a template for the construction of semiparametric locally efficient double robust substitution estimators of the target parameter of the data generating distribution in a semiparametric censored data or causal inference model (van der Laan and Rubin (2006), van der Laan (2008), van der Laan and Rose (2011)). It is assumed that the data set is a realization of n independent and identically distributed random variables, the probability distribution of this random variable is known to be an element of a semiparametric statistical model, and the target parameter (mapping) is defined as a particular function of the possible probability distributions in this semiparametric model. A targeted minimum loss based estimator (TMLE) of the target parameter is defined by an initial estimator of a relevant part of the data generating distribution, a parametric submodel through an initial estimator, a loss function for this relevant part, minimizing the empirical risk of the loss function along the parametric submodel to iteratively update the initial estimator until convergence. This final estimator is the TMLE of the relevant part of the data generating distribution, and the evaluation of its target parameter value is the TMLE of the target parameter. By enforcing that the loss-based score of the submodel (at zero fluctuation of the initial estimator) spans the efficient influence curve of the target parameter (at the initial estimator), it follows that the TMLE of the relevant part of the data generating distribution solves the efficient score estimating equation, making the TMLE locally efficient and double robust, under regularity conditions. By choosing a parametric submodel with extra fluctuation parameters, the TMLE can be arranged to solve additional estimating equations, and thereby satisfy additional properties of interest (e.g., be an imputation estimator, see Gruber and van der Laan (2010a)). One particular example of such an iterative TMLE was presented in the original TMLE article, van der Laan and Rubin (2006), which involved also fluctuating the treatment/censoring mechanism, resulting in a TMLE that, assuming convergence, also equals an IPTW/IPCW estimator and is guaranteed to outperform the IPTW/IPCW estimator defined by the initial estimator of the treatment/censoring mechanism.

Another desirable property of an estimator is that it is guaranteed to be more efficient than a user-supplied class of estimators in the case that the censoring/treatment mechanism is correctly specified. This has been achieved with empirical efficiency maximization (Rubin and van der Laan (2008), Tan (2008, 2010), Cao, Tsiatis, and Davidian (2009), van der Laan and Rose (2011)). However, in general this technique as presented in Rubin and van der Laan (2008) may come at a cost of losing double robustness (e.g., see Robins and Rotnitzky (1992) and van der Laan and Robins (2003)). Tan (2008) demonstrates how in the context of estimating equation methodology for estimating a population mean under miss-

ingness the double robustness can be preserved. Recently, Rotnitzky, Lei, Sued, and Robins (2012) shows how to combine empirical efficiency maximization with double robust locally efficient substitution estimators, by fluctuating the treatment mechanism with a carefully chosen clever covariate derived from the empirical efficiency maximization procedure. Borrowing this fundamental idea, in this article we demonstrate that this enhanced efficiency property can be achieved with the above mentioned TMLE of van der Laan and Rubin (2006) (jointly updating treatment mechanism and outcome regression), by fluctuating the treatment mechanism with this additional clever covariate as suggested by Rotnitzky et al. (2012). TMLE with this property remains double robust, and is as efficient as any competing regular asymptotically linear estimator. For the sake of illustration we focus on estimation of the additive average causal effect of a point treatment on an outcome, adjusting for baseline covariates.

1.1 Organization

In Section 2 we present the statistical estimation problem. In Section 3 we present the TMLE, and the enhanced empirically efficient TMLE, and explain its properties. From the presentation in Section 3, for experts familiar with the theory of augmented IPCW estimating equations (Robins and Rotnitzky (1992), van der Laan and Robins (2003)) it will also be clear how this TMLE is generalized to all CARcensored data and causal inference models. In Section 4 we review the method for empirical efficiency maximization of Rubin and van der Laan (2008), and an adaptive version of it as presented in van der Laan and Rose (2011), used as an ingredient in the enhanced empirically efficient TMLE. In Section 5 we present simulations confirming the enhanced efficiency property of the TMLE presented in Section 3, and comparing it with the (non-double robust) empirical efficiency maximization estimator in Rubin and van der Laan (2008), and a regular TMLE. We end with some concluding remarks. We also provide an appendix with the R-code of the TMLEs implemented in the simulation study.

2 The statistical model, target parameter, and estimation problem

Let $O = (W, A, Y) \sim P_0$ be a random variable, where W represents a vector of baseline covariates, A a binary treatment, and Y a continuous or binary outcome with values in [0,1]. Let $g_0(A \mid W)$ be the conditional probability distribution of A, given W. Consider a statistical model \mathcal{M} that makes no assumptions about the marginal

distribution of W, and the conditional distribution of Y, given A, W, but might make assumptions about g_0 . In particular, it is assumed that $0 < g_0(1 \mid W) < 1$ so that the following target parameter is well defined. The statistical target parameter $\Psi: \mathcal{M} \to \mathbb{R}$ of interest is defined as

$$\Psi(P) = E_P(E_P(Y \mid A = 1, W) - E_P(Y \mid A = 0, W)).$$

If one assumes an underlying nonparametric structural equation model $W = f_W(U_W)$, $A = f_A(W, U_A)$, $Y = f_Y(W, A, U_Y)$ (Pearl (2000)), and the randomization assumption U_A is independent of U_Y , given W, then $\Psi(P_0)$ identifies the additive causal effect $E_0(Y(1) - Y(0))$, where $Y(a) = f_Y(W, a, U_Y)$ is the treatment-specific counterfactual. For the sake of estimation, we are only concerned with the statistical target parameter.

Let $Q_W(P)$ be the marginal distribution of W under P, $\bar{Q}(P)(A,W) = E_P(Y \mid A,W)$, and we will denote corresponding parameter values with Q_W and \bar{Q} , respectively. Let $Q(P) = (Q_W(P), \bar{Q}(P))$. Note that $\Psi(P)$ only depends on P through $Q_W(P)$ and $\bar{Q}(P)$. Therefore, we will also use the notation

$$\Psi(Q) = E_{Q_W} \{ \bar{Q}(1, W) - \bar{Q}(0, W) \}.$$

Our goal is to estimate $\psi_0 = \Psi(Q_0)$ based on observing n i.i.d. copies O_1, \dots, O_n of $O \sim P_0 \in \mathcal{M}$.

The TMLE requires knowing the canonical gradient/efficient influence curve of the pathwise derivative of $\Psi: \mathcal{M} \to \mathbb{R}$. The efficient influence curve of $\Psi: \mathcal{M} \to \mathbb{R}$ at P is given by

$$D^{*}(P)(O) = \frac{2A-1}{g(A \mid W)} (Y - \bar{Q}(A, W)) + \{\bar{Q}(1, W) - \bar{Q}(0, W) - \Psi(Q)\}$$

$$\equiv D_{Y}^{*}(P)(O) + D_{W}^{*}(P)(W),$$

where the latter decomposition in a score $D_Y^*(P)$ of the conditional distribution of Y, given A, W, and score $D_W^*(P)$ of the marginal distribution of W will be utilized in TMLE. In order to establish the enhanced efficiency property of the proposed TMLE we will also utilize the augmented IPCW-representation of the efficient influence curve (Robins and Rotnitzky (1992), van der Laan and Robins (2003)):

$$D^{*}(P)(O) = \frac{2A-1}{g(A \mid W)}Y - \Psi(Q) - \left\{\frac{\bar{Q}(1,W)}{g(1 \mid W)} + \frac{\bar{Q}(0,W)}{g(0 \mid W)}\right\} (A - g(1 \mid W))$$

$$\equiv D_{IPTW}(Q,g)(O) + D_{CAR}(\bar{Q},g)(O),$$

where $D_{CAR}(\bar{Q}, g) = -H_{CAR}(Q, g)(W)(A - g(1 \mid W))$ with

$$H_{CAR}(Q,g)(W) \equiv \left\{ rac{ar{Q}(1,W)}{g(1\mid W)} + rac{ar{Q}(0,W)}{g(0\mid W)}
ight\}.$$

In cases where we want to stress the representation of the efficient influence curve $D^*(P)$ as an estimating function in ψ indexed by nuisance parameters \bar{Q}_0, g_0 , we will also use the notation $D_{IPTW}(\psi_0, g_0)$ for $D_{IPTW}(Q_0, g_0)$, and $D^*(\psi_0, \bar{Q}_0, g_0)$ for $D^*(P_0)$. We will use the subscript '0' to denote the truth, and 'n' to denote estimates from data.

Another ingredient of the TMLE presented in the next section is an influence curve D(P) of a competing regular asymptotically linear estimator of Ψ at P in the model \mathcal{M} . The TMLE ψ_n^* will be constructed so that it is at least as efficient as this competing estimator at P_0 in the case that we estimate g_0 consistently. By the representation theorem for the class of gradients in CAR-censored data models (van der Laan and Robins (2003), p. 65), it follows that

$$D(P) = D_{IPTW}(Q, g) + D_{CAR}(\bar{Q}^e, g)$$

for a particular function $\bar{Q}^e = \bar{Q}^e(P)$. Let \bar{Q}_0^e denote the true value of this parameter $P \to \bar{Q}^e(P)$.

The TMLE presented in the next section will use an estimator \bar{Q}_n^e of \bar{Q}_0^e in order to define a clever covariate $H_{CAR}(\bar{Q}_n^e,g_n^k)$ in the definition of the TMLE. As a consequence of this choice of clever covariate, the TMLE Q_n^*,g_n^* will solve

$$0 = P_n \left(D_{IPTW}(\psi_n^*, g_n^*) + D_{CAR}(\bar{Q}_n^e, g_n^*) \right),$$

and thereby have an influence curve at least as efficient as $D(P_0)$, if g_0 is estimated consistently.

A particular choice for \bar{Q}_0^e is defined by empirical efficiency maximization over a user-supplied working model as in Rubin and van der Laan (2008). That is, let $\{\bar{Q}_\beta:\beta\}$ be a parametric working model, and define

$$\bar{Q}^{e}(P_{0}) = \arg\min_{\bar{Q}_{\beta}} P_{0} \{ D_{IPTW}(\psi_{0}, g_{0}) + D_{CAR}(\bar{Q}_{\beta}, g_{0}) \}^{2}.$$
 (1)

Here we used the notation $Pf \equiv \int f(o)dP(o)$. With this choice, $D(P_0)$ represents the influence curve with minimal variance among the class of influence curves $\{D_{IPTW}(Q_0,g_0)+D_{CAR}(\bar{Q}_{\beta},g_0):\beta\}$ indexed by β .

3 The TMLE that is at least as efficient as competing estimator

The TMLE of $\Psi(Q_0)$ as presented in van der Laan and Rubin (2006) is defined by 1) a loss function $\mathcal{L}(Q,g) = \mathcal{L}(Q) + \mathcal{L}(g)$ for (Q_0,g_0) so that $Q_0 = \arg\min_Q P_0 \mathcal{L}(Q)$,

 $g_0 = \arg\min_g P_0 \mathcal{L}(g)$, 2) a submodel $\{Q(\varepsilon_1) : \varepsilon_1\}$ through Q at $\varepsilon_1 = 0$, a submodel $\{g(\varepsilon_2) : \varepsilon_2\}$ through g at $\varepsilon_2 = 0$, and 3) an initial estimator Q_n^0 , g_n^0 . The TMLE is defined by iterative minimization of the empirical risk, and updating:

$$\varepsilon_{1n} = \arg\min_{\varepsilon_1} P_n \mathcal{L}(Q_n^0(\varepsilon_1))$$

$$\varepsilon_{2n} = \arg\min_{\varepsilon_2} P_n \mathcal{L}(g_n^0(\varepsilon_2)),$$

 $Q_n^1 = Q_n^0(\varepsilon_{1n}), g_n^1 = g_n^0(\varepsilon_{2n}),$ and this updating process is iterated until $\varepsilon_n = (\varepsilon_{1n}, \varepsilon_{2n})$ ≈ 0 . The resulting Q_n^*, g_n^* solve the loss-based score equation:

$$P_n \frac{d}{d\varepsilon} \mathcal{L}(Q_n^*(\varepsilon), g_n^*(\varepsilon)) \bigg|_{\varepsilon=0} = 0.$$
 (2)

By defining the loss-function \mathcal{L} and submodel through (Q,g), one can control the estimating equation (2) solved by the TMLE. In particular, one wants the loss-based scores to span the efficient influence curve $D^*(Q_n^*, g_n^*)$ so that the resulting $\Psi(Q_n^*)$ will be double robust and locally efficient. Below we present a submodel $\{g(\mathcal{E}_2):\mathcal{E}_2\}$ so that the additional desired enhanced efficiency property is achieved as well. The goal is to guarantee that the estimator is at least as efficient as a given estimator, even when \bar{Q} is misspecified, and one application of this is to guarantee that it is at least as efficient as an empirically efficient estimator that maximizes efficiency over a parametric model. The TMLE presented in this article is improved locally efficient, and, as in Rotnitzky et al. (2012), is double robust and will be at least as efficient as the augmented IPTW estimator that uses \bar{Q}_n^e as an estimator of the true \bar{Q}_0 , and tailors \bar{Q}_n^e to maximize efficiency when g_n^* is consistent.

3.1 Initial estimators

Let $Q_{W,n}^0$, \bar{Q}_n^0 , g_n^0 , and \bar{Q}_n^e be initial estimators of $Q_{W,0}$, \bar{Q}_0 , g_0 , and \bar{Q}_0^e , respectively. Let $Q_{W,n}^0 = Q_{W,n}$ be the empirical probability distribution of W_1, \ldots, W_n . The estimator of \bar{Q}_0 can be based on the least squares or (quasi-)log-likelihood loss function

$$\mathcal{L}(\bar{Q})(O) = -\left\{Y\log\bar{Q}(A, W) + (1 - Y)\log\{1 - \bar{Q}(A, W)\}\right\}. \tag{3}$$

This is the log-likelihood loss-function for \bar{Q}_0 if Y is binary. We refer to Gruber and van der Laan (2010b) in which this loss function is proposed for TMLE with a continuous bounded outcome $Y \in [0,1]$. By a simple linear transformation, this also provides a loss function for $Y \in [a,b]$ with bounded a,b. In particular, \bar{Q}_0 could be estimated with a loss-based super learner using this loss function for the cross-validation selector (van der Laan, Polley, and Hubbard (2007)).

The estimator of g_0 can be based on the log-likelihood loss function $\mathcal{L}(g) = -\log g$. The estimation method for \bar{Q}_0^e might depend on the type of parameter it represents. If \bar{Q}_0^e is defined by (1), then one could estimate it as

$$\bar{Q}_{n}^{e} = \arg\min_{\bar{Q}_{\beta}} P_{n} \{ D_{IPTW}(\psi_{n}^{0}, g_{n}^{0}) + D_{CAR}(\bar{Q}_{\beta}, g_{n}^{0}) \}^{2}, \tag{4}$$

where P_n denotes the empirical probability distribution of O_1, \ldots, O_n , and ψ_n^0 represents an estimator of ψ_0 that is consistent if g_n^0 is consistent. For example, ψ_n^0 could be any TMLE that takes \bar{Q}_n^0 and g_n^0 as initial estimator. We assume that \bar{Q}_n^e is consistent for \bar{Q}_0^e if g_n^0 is consistent.

3.2 Loss function

We select the log-likelihood loss functions $\mathcal{L}(g) = -\log g$, $\mathcal{L}(Q_W) = -\log Q_W$ for g_0 and $Q_{W,0}$, respectively, and we select $\mathcal{L}(\bar{Q})$ (3) as loss function for \bar{Q}_0 . Let $\mathcal{L}(Q,g) \equiv \mathcal{L}(\bar{Q}) + \mathcal{L}(Q_W) + \mathcal{L}(g)$ be the loss function for (Q_0,g_0) .

3.3 TMLE that is at least as efficient as competing estimator

Let $\bar{g}(W) \equiv g(1 \mid W)$. For a given \bar{Q}_n^k , g_n^k , define the submodels

$$\begin{array}{lcl} \operatorname{Logit} \bar{Q}_{n}^{k}(\varepsilon_{1}) & = & \operatorname{Logit} \bar{Q}_{n}^{k} + \varepsilon_{1} H^{*}(g_{n}^{k}) \\ \operatorname{Logit} \bar{g}_{n}^{k}(\varepsilon_{2}) & = & \operatorname{Logit} \bar{g}_{n} + \varepsilon_{21} H_{CAR}(\bar{Q}_{n}^{k}, g_{n}^{k}) + \varepsilon_{22} H_{CAR}(\bar{Q}_{n}^{e}, g_{n}^{k}), \end{array}$$

where $H^*(g_n^k) = (2A-1)/g(A \mid W)$. We also define a submodel $Q_{W,n}(\varepsilon_0) = (1+\varepsilon_0 D_W^*(Q_n^k))Q_{W,n}$ through the empirical probability distribution $Q_{W,n}$. Let $\varepsilon = (\varepsilon_0, \varepsilon_1, \varepsilon_2 = (\varepsilon_{21}, \varepsilon_{22}))$. This defines a submodel $(Q_n^k(\varepsilon), g_n^k(\varepsilon))$ through $(Q_n^k = (Q_{W,n}^k, \bar{Q}_n^k), g_n^k)$ at $\varepsilon = 0$. The scores $\frac{d}{d\varepsilon}L(Q_n^k(\varepsilon), g_n^k(\varepsilon))$ of $(\varepsilon_0, \varepsilon_1)$ at $\varepsilon = 0$ spans the efficient influence curve $D^*(Q_n^k, g_n^k)$. The score of ε_2 at $\varepsilon = 0$ spans any linear combination of $D_{CAR}(\bar{Q}_n^k, g_n^k)$ and $D_{CAR}(\bar{Q}_n^e, g_n^k)$.

Given a current estimator (Q_n^k, g_n^k) , we estimate ε with the MLE ε_n^k based on loss function $\mathcal{L}(Q, g)$:

$$\varepsilon_{0n}^{k} = \arg\min_{\varepsilon_{0}} -P_{n} \log Q_{W,n}^{k}(\varepsilon_{0})
\varepsilon_{1n}^{k} = \arg\min_{\varepsilon_{1}} P_{n} \mathcal{L}(\bar{Q}_{n}^{k}(\varepsilon_{1}))
\varepsilon_{2n}^{k} = \arg\min_{\varepsilon_{2}} -P_{n} \log g_{n}^{k}(\varepsilon_{2}).$$

Note that \mathcal{E}^k_{1n} and \mathcal{E}^k_{2n} can be fitted with standard univariate logistic regression incorporating an offset. We start with k=0. This defines the first step TMLE update $(Q^1_n=Q^0_n(\mathcal{E}^0_n),g^1_n=g^0_n(\mathcal{E}^0_n))$. We can iterate this updating algorithm until convergence so that $\mathcal{E}_n\approx 0$. Let (Q^*_n,g^*_n) be the final TMLE at convergence. Since $Q^0_{W,n}$ is the empirical probability distribution, we have $\mathcal{E}^k_{0n}=0$ for all k, so that this empirical probability distribution is not updated by the TMLE algorithm, i.e., $Q^*_n=(Q_{W,n},\bar{Q}^*_n)$. The TMLE of ψ_0 is the substitution estimator $\psi^*_n=\Psi(Q^*_n)$. An iterative TMLE algorithm involving updating both g^0_n and Q^0_n was pre-

An iterative TMLE algorithm involving updating both g_n^0 and Q_n^0 was presented and implemented in van der Laan and Rubin (2006)), that did not include the extra clever covariate $H(\bar{Q}_n^e, g_n^k)$. The important choice of extra clever covariate $H(\bar{Q}_n^e, g_n^k)$ in a model for g_0 in order to establish the enhanced efficiency property without losing double robustness was presented in Rotnitzky et al. (2012).

3.4 Estimating equations solved by TMLE, and resulting alternative representations of the TMLE

We assume that the algorithm converges. In that case, (Q_n^*, g_n^*) solves the score equations for the sub-model $\{Q_n^*(\varepsilon), g_n^*(\varepsilon) : \varepsilon\}$ at $\varepsilon = (\varepsilon_0, \varepsilon_1, \varepsilon_{21}, \varepsilon_{22}) = 0$. As a consequence, the TMLE solves the following equations:

$$P_{n}D^{*}(\psi_{n}^{*}, \bar{Q}_{n}^{*}, g_{n}^{*}) = 0$$

$$P_{n}D_{IPTW}(\psi_{n}^{*}, g_{n}^{*}) = 0$$

$$P_{n}D_{CAR}(\bar{Q}_{n}^{e}, g_{n}^{*}) = 0$$

$$P_{n}D^{*}(\psi_{n}^{*}, \bar{Q}_{n}^{e}, g_{n}^{*}) = 0.$$

This allows for a variety of representations of the TMLE. It is a plug in estimator

$$\psi_n^* = \Psi(Q_n^*);$$

it is an IPTW estimator

$$\psi_n^* = \frac{1}{n} \sum_{i=1}^n \frac{2A_i - 1}{g_n^*(A_i \mid W_i)} Y_i;$$

it is an augmented IPCW-estimating equation based estimator

$$\psi_n^* = \frac{1}{n} \sum_{i=1}^n \frac{2A_i - 1}{g_n^*(A_i \mid W_i)} Y_i - H_{CAR}(Q_n^*, g_n^*)(W_i) (A_i - \bar{g}_n^*(W_i)),$$

corresponding with the implicit estimator \bar{Q}_n^*, g_n^* of the nuisance parameters (\bar{Q}_0, g_0) of the estimating function $D^*(\psi, \bar{Q}_0, g_0)$ in ψ ; and, finally, it is also an augmented

IPCW-estimating equation based estimator

$$\psi_n^* = \frac{1}{n} \sum_{i=1}^n \frac{2A_i - 1}{g_n^*(A_i \mid W_i)} Y_i - H_{CAR}(Q_n^e, g_n^*)(W_i) (A_i - \bar{g}_n^*(W_i)),$$

corresponding with estimating Q_0 with Q_n^e .

Even though the TMLE allows these representations, the actual construction of the TMLE is fundamentally different from these estimators since the nuisance parameters g_0, Q_0 are estimated with the implicitly defined TMLE g_n^*, Q_n^* itself. So from that perspective it is misleading to refer to the TMLE as also being an IPCW and augmented-IPCW estimator, since the latter estimators are defined by solving the corresponding estimating equation at an initial explicit estimator of the nuisance parameters g_0, Q_0 . Nonetheless, these representations help us understand that the proposed TMLE will inherit the same asymptotic properties as these estimating equation based estimators, while retaining fundamental advantages by also being a well defined substitution estimator respecting the global constraints of the statistical model and target parameter TMLE does not require a closed form estimating equation, but indirectly solves the efficient influence curve equation (and avoids the problem of multiple solutions) as a by-product of minimizing a loss function for fitting the fluctuation parameter. In addition, by simultaneously fitting the additional parameter derived by Rotnitzky et al. (2012), the proposed algorithm ensures that additional score equations associated with empirical efficiency maximization are solved. The iterative procedure can be carried out to arbitrary precision.

In the case that Q_n^e is defined by empirical efficiency maximization (1), then the latter estimator is the estimator of Rubin and van der Laan (2008), obtained by maximizing empirical efficiency of the class of estimating functions $D(\psi, \bar{Q}_\beta, g_0)$ (or equivalently, $D(\psi, f_\beta, g_0)$, as reviewed in next section) over the working model $\{\bar{Q}_\beta:\beta\}$ at the (implicit) estimator g_n^* of g_0 , and defining the estimator of ψ_0 as the solution of the corresponding estimating equation. Note that in this particular situation \bar{Q}_n^e could be used as an initial estimator of \bar{Q}_0 in the above algorithm. However, in general, substituting \bar{Q}_n^e for \bar{Q}_n^0 has drawbacks. In the above procedure \bar{Q}_n^0 plays the role of the best initial estimate of the relevant portion of the likelihood, while \bar{Q}_n^e comes from the influence curve of an estimator that provides an upper bound on the efficiency of the parameter estimate. This estimator may be less than ideal, for example, \bar{Q}_n^e might be a parametric regression-based estimator, when in fact it is not necessary to place such a restriction on \bar{Q}_n^0 , thus it is important to distinguish between these two objects, and obtaining a separate estimate of \bar{Q}_n^0 is preferred.

3.5 Properties of TMLE

The TMLE presented above satisfies both the definition of the TMLE as well as the definition of the empirical efficient maximization (estimating equation based) estimator of Rubin and van der Laan (2008), using the implicit estimator g_n^* for g_0 . As a consequence, it inherits the properties of both the TMLE, as a locally efficient double robust substitution estimator, as well as the empirically efficient maximization estimator of Rubin and van der Laan (2008), as an estimator that is maximally efficient among a user supplied class of asymptotically linear estimators in the case that g_0 is estimated consistently. For the sake of being self-contained we present here the rationale resulting in these properties. Formal proofs of these properties would require regularity conditions, and is beyond the scope of this article. Therefore, below we present the general statements, and refer to the general theorems that would have to be applied to formally establish the claimed asymptotic properties. For a completely worked out proof of a TMLE for the additive treatment effect, we refer to Zheng and van der Laan (2010) and van der Laan and Rose (2011).

By the fact that it is a TMLE that solves the efficient influence curve estimating equation $P_nD^*(\psi_n^*, \bar{Q}_n^*, g_n^*) = 0$ it follows that ψ_n^* will be consistent if either \bar{Q}_n^* or g_n^* is consistent. In addition, under regularity conditions (e.g., van der Laan and Robins (2003), van der Laan and Rubin (2006)), ψ_n^* will be an asymptotically linear estimator if either \bar{Q}_n^* or g_n^* is consistent, and it will be efficient if both are consistent. This corresponds with stating that ψ_n^* is a double robust locally efficient estimator.

Before we proceed with explaining the enhanced efficiency property we first provide background on estimating equation based estimators in CAR censored data models (Robins and Rotnitzky (1992), van der Laan and Robins (2003)). Suppose that ψ_n is an estimator that solves the estimating equation $0 = P_n D^*(\psi, \bar{Q}, g_0)$ for some Q. Then it follows that ψ_n is asymptotically linear with influence curve $D^*(\psi_0, \bar{Q}, g_0)$. In addition, if \bar{Q}_n converges to \bar{Q} , then under weak regularity conditions, we have that the solution ψ_n of $P_nD^*(\psi,\bar{Q}_n,g_0)$ is also asymptotically linear with influence curve $D^*(\psi_0, \bar{Q}, g_0)$. By Theorem 2.3 in van der Laan and Robins (2003), if the estimator g_n^* of g_0 is such that a particular specified smooth function $\Phi(g_n^*)$ is an efficient estimator of $\Phi(g_0)$ so that its influence curve is an element of the tangent space $T_{CAR}(P_0) = \{S(A \mid W) : E_{g_0}(S \mid W) = 0\}$ of g at P_0 under CAR, then, under regularity conditions, the solution ψ_n of $P_nD^*(\psi,\bar{Q}_n,g_n^*)=0$ is asymptotically linear with an influence curve that has a variance smaller than or equal to the variance of $D^*(\psi, \bar{Q}, g_0)$. (Here $T_{CAR}(P_0)$ consists of all functions, S, of A, W with conditional mean zero, given W, which corresponds with the definition of $T_{CAR}(P_0)$ as the tangent space of the censoring mechanism g as in van der Laan and Robins (2003).) That is, consistent (and efficient) estimation of the orthogonal nuisance parameter g_0 only improves the efficiency of the estimating equation based estimator of ψ_0 . Since g_n^* is a pure MLE-based estimator, under regularity conditions, and under the assumption that g_n^0 is consistent for g_0 , one can show that $\Phi(g_n^*)$ is an asymptotically linear estimator of $\Phi(g_0)$ with influence curve in $T_{CAR}(P_0)$.

Given that we know that $P_nD^*(\psi_n^*, \bar{Q}_n^e, g_n^*) = 0$, if g_n^* is consistent for g_0 and \bar{Q}_n^e converges to \bar{Q}_0^e , it follows that ψ_n^* will be asymptotically linear with an influence curve with variance smaller than or equal to the variance of $D^*(\psi_0, \bar{Q}_0^e, g_0)$. That is, in the case that g_n^* is consistent, the TMLE ψ_n^* is at least as efficient as the competing estimator whose influence curve equals $D^*(\psi_0, \bar{Q}_0^e, g_0)$.

4 Empirical efficiency maximization

This section concerns the estimation of Q_0^e that forms an ingredient of the TMLE presented above. We first review empirical efficiency maximization as presented in Rubin and van der Laan (2008), and then we demonstrate how empirical efficiency maximization can be embedded in loss-based learning of Q_0 by using as loss function the square of the efficient influence curve (van der Laan and Rose (2011)).

4.1 Empirical Efficiency Maximization as in Rubin, van der Laan (2008)

In order to determine a solution that optimizes the variance of the influence curve among a class of influence curves the following method was proposed in Rubin and van der Laan (2008). Firstly, it is noted that

$$D^{*}(g,Q) = D_{IPTW}(Q,g) - H_{CAR}(Q,g)(A - g(1 \mid W))$$

= $H_{g}^{*}(A,W)(Y - f(Q,g)(W)) - \Psi(Q)$
= $D^{*}(g,f(Q,g),\Psi(Q)),$

where $H_g^*(A, W) = (2A - 1)/g(A \mid W)$, and

$$f(Q,g) = g(1 \mid W)\bar{Q}(0,W) + g(0 \mid W)\bar{Q}(1,W).$$

Note that $D^*(g, f, \psi) = D_{IPTW}(g, \psi) - H_{CAR}(f, g)(A - g(1 \mid W))$, where

$$H_{CAR}(f,g) = \frac{f(W)}{g(1 \mid W)g(0 \mid W)}.$$

Therefore, minimizing the variance of $D^*(g, Q, \psi_0)$ over Q is equivalent with minimizing the variance of $D^*(g, f, \psi_0)$ over a corresponding class of f's. We have

$$VAR_{P_0}D^*(g_0, f, \psi_0) = E_0 \left\{ H_{g_0}^{*2}(A, W)(Y - f(W))^2 \right\}.$$

As a consequence, minimizing the variance of $D^*(g,Q,\psi_0)$ over f's corresponds with minimum weighted least squares, whose solution is often unique, and for which numerical routines are often available. Thus, given a working model $\{f_\beta:\beta\}$ for f, one can now define an optimal choice

$$f^{e}(P_{0}) = \arg\min_{f_{\beta}} P_{0} \left\{ H_{g_{0}}^{*2}(A, W) (Y - f_{\beta}(W))^{2} \right\}.$$
 (5)

The choice (5) can be estimated with weighted least squares by regressing Y on W using weights $H_{g_0}^{*2}$. An estimator f_n^e of f_0^e results in a clever covariate $H_{CAR}(f_n^e, g_n^k)$ in the k-th step of the TMLE-algorithm presented in the previous section.

4.2 Adaptive empirical efficiency maximization

The choice $\bar{Q}_e(P_0)$ (1) corresponds with minimizing the empirical risk of the loss function $\mathcal{L}_{g_0}(\bar{Q}) = \{D^*(\psi_0, \bar{Q}, g_0)\}^2$ over a working model $\{\bar{Q}_\beta : \beta\}$. Note that \mathcal{L}_{g_0} is indeed a valid loss function since $\bar{Q}_0 = \arg\min_{\bar{Q}} P_0 \mathcal{L}_{g_0}(\bar{Q})$ (van der Laan and Robins (2003)). The strength of this loss function is that its loss-based dissimilarity is given by

$$P_0\{\mathscr{L}_{g_0}(\bar{Q}) - \mathscr{L}_{g_0}(\bar{Q}_0)\} = P_0\{D^*(\psi_0, \bar{Q}, g_0) - D^*(\psi_0, \bar{Q}_0, g_0)\}^2,$$

which follows by the Pythagorean theorem (van der Laan and Rose (2011)). In some cases (as in the previous subsection), one can define another (e.g., squared error) loss function that has the same loss-based dissimilarity, but with an empirical risk that might be easier to minimize. The validity of the loss function relies on g_0 being known or consistently estimated. At the known g_0 , this loss-based dissimilarity is targeted towards ψ_0 since it concerns approximating the true efficient influence curve, and it also corresponds with minimizing the variance of the influence curves $D^*(\psi_0, \bar{Q}, g_0)$ over \bar{Q} .

Instead of working with a single working model, we can alternatively use loss-based learning, using cross-validation based on this loss function \mathcal{L}_{g_0} (van der Laan and Dudoit (2003), van der Laan et al. (2007)). For example, suppose one considers a collection of K working models $\{\bar{Q}_{\beta^k}^k:\beta^k\}$, $k=1,\ldots,K$. Each working model results in an estimator \bar{Q}_n^k defined by the minimizer of the empirical risk $P_n\mathcal{L}_{g_0}(\bar{Q}_{\beta^k}^k)$ over the working model indexed by parameter vector β^k . One can now select the choice k of working model with the V-fold cross-validation selector

$$k_n = \arg\min_{k} \sum_{v=1}^{V} \sum_{i \in \text{Val}_v} \mathcal{L}_{g_0}(\bar{\mathcal{Q}}_{n,v}^k)(O_i),$$

where Val_{v} is the validation sample for the v-th sample split, and $\overline{Q}_{n,v}$ is the fit of the k-th working model based on the training sample Train_{v} (i.e., the complement of

 Val_{v}) for the v-th sample split, $v = 1, \dots, V$ The estimator would now be $\bar{Q}_{n}^{e} = \bar{Q}_{\beta_{n}^{k_{n}}}^{k_{n}}$, which plays the role of an estimator of \bar{Q}_{0}^{e} .

As shown in van der Laan and Rose (2011), the general oracle results of the cross-validation selector k_n apply to this loss function $\mathcal{L}_{g_0}(\bar{Q})$ (van der Laan and Dudoit (2003)), under the assumption that the efficient influence curve is uniformly bounded in supremum norm (i.e., $\delta < g_0(1 \mid W) < 1 - \delta$ for some $\delta > 0$). As a consequence, under this boundedness condition, if none of the working models are correctly specified, the cross-validation selector will asymptotically make the optimal choice, even if the number K of working models grows polynomial in sample size, while, if one of the working models is correctly specified, then the resulting \bar{Q}_n^e will converge at rate $1/\sqrt{n}$ to \bar{Q}_0 . These oracle results will also apply if g_0 is estimated at a rate faster than that at which \bar{Q}_0 is estimated (van der Laan and Dudoit (2003)).

As a consequence, if g_0 is estimated consistently, the augmented IPTW estimator that uses \bar{Q}_n^e as estimator of \bar{Q}_0 will now have an influence curve that is more efficient than $D^*(\psi_0, \bar{Q}_{\beta^k}^k, g_0)$ for any β^k , and any $k=1,\ldots,K$. We note that this estimator \bar{Q}_n^e based on using loss-based (super) learning can now be used in the clever covariate $H_{CAR}(\bar{Q}_n^e, g_n^k)$ in the TMLE proposed in the previous section. The TMLE presented in the previous section using this estimator \bar{Q}_n^e as estimator of \bar{Q}_0^e will now not only be a double robust locally efficient substitution estimator, but, if g_0 is estimated consistently, it will also be at least as efficient as the augmented IPTW estimator that uses \bar{Q}_n^e as estimator of \bar{Q}_0 .

If one implements an \mathcal{L}_{g_0} -based super learner with a library of candidate estimators of \bar{Q}_0 that includes nonparametric estimators, so that at least one candidate in the library will be asymptotically consistent for \bar{Q}_0 , then $\bar{Q}_0^e = \bar{Q}_0$ and \bar{Q}_n^e is now an estimator of the globally optimal \bar{Q}_0 . However, if one estimates g_0 consistently, then \bar{Q}_n^e is also a fully targeted estimator of \bar{Q}_0 in the sense that it is tailored to result in a best estimate of the efficient influence curve itself. Again, we can use this estimator \bar{Q}_n^e in the clever covariate $H_{CAR}(\bar{Q}_n^e, g_n^k)$ in the TMLE proposed in the previous section. The resulting TMLE is not only a double robust locally efficient substitution estimator, but also an augmented IPTW estimator that estimates \bar{Q}_0 with an estimator \bar{Q}_n^e that is tailored to maximize efficiency in the case that g_n is consistent.

5 Simulations illustrating empirical efficiency property of TMLE

We illustrate the additional enhanced efficiency property of the TMLE proposed in the previous section by comparing the performance of the enhanced TMLE with the standard TMLE and the empirical efficiency maximization estimator proposed in Rubin and van der Laan (2008). Two data generating distributions were defined, and the additive treatment effect parameter was estimated for one thousand samples of size n = 500 drawn from each. Results in Table 1 below verify that when there are no violations of the positivity assumption, at a misspecified working model for \bar{Q}_0 and correctly specified working model for g_0 , the performance of the enhanced TMLE is on a par with the empirical efficiency maximization estimator, and both outperform the TMLE that does not aim for maximal efficiency.

Data were generated according to the following two mechanisms:

$$\begin{split} W_1, W_2 \sim N(0,1) \\ g_{0,1}(1 \mid W) &= 0.5 \\ g_{0,2}(1 \mid W) &= \operatorname{Expit}(-0.3 - 0.1W_1 - 0.3W_2) \\ P_{0,1}(Y = 1 \mid A, W) &= \operatorname{Expit}(-1 + A + W_1 + 2.5W_1^2) \\ P_{0,2}(Y = 1 \mid A, W) &= \operatorname{Expit}(-1 + A + W_1 + 2.5W_1^2 - 0.2W_2). \end{split}$$

The first simulation study mimics a randomized controlled trial in which treatment assignment is independent of baseline covariates $W = (W_1, W_2)$. The probability of being assigned to the treatment group is 0.5 for all subjects. In the second study W_1 and W_2 confound the effect of treatment on the outcome. For this simulation true treatment assignment probabilities ranged between 0.26 and 0.60. The true values of the target parameter for these two simulations are $\psi_{0,1} = 0.1579$ and $\psi_{0,2} = 0.1570$. These true values were obtained as an average of the additive effect $(Y_1 - Y_0)$ calculated from the full data for ten samples of size $n = 10^7$. For both studies, (misspecified) logistic linear regression of Y on (A, W_1) was used to obtain the initial estimate of \bar{Q}_0 , and the correctly specified logistic regression model was used to obtain the initial estimate of g_0 . The estimators of ψ_0 are of the following form:

$$\begin{split} & \psi_n^{TMLE} = \frac{1}{n} \sum_{i=1}^n \left\{ \bar{Q}_n^{1*}(1, W_i) - \bar{Q}_n^{1*}(0, W_i) \right\}, \\ & \psi_n^{Empeff} = \frac{1}{n} \sum_{i=1}^n \frac{2A_i - 1}{g(A_i \mid W_i)} \left(Y_i - f_n^e(W_i) \right), \\ & \psi_n^{TMLE_{en}} = \frac{1}{n} \sum_{i=1}^n \left\{ \bar{Q}_n^{k*}(1, W_i) - \bar{Q}_n^{k*}(0, W_i) \right\}. \end{split}$$

Here $\operatorname{Logit} \bar{Q}_n^{1*} = \operatorname{Logit} \bar{Q}_n^0 + \varepsilon_n H_{g_n}^*$, with ε_n fit by maximum likelihood estimation. The target function $f_0^e(W) = f_{c_0,\alpha_0,\beta_0}(W)$, which defines the empirical efficiency maximization estimator, is defined in terms of a working model $f_{c_0,\alpha_0,\beta_0}(W) = c + \operatorname{Expit}(\alpha + \beta W_1)$ (see section 4 above). The true values (c_0,α_0,β_0) of the coefficients are estimated with weighted least squares using the nlm function in R (Team, 2010) and weights $\{H_{g_n}^*\}^2$. Finally, $\bar{Q}_n^{k^*}(A,W)$ is a targeted estimate of \bar{Q}_0 obtained by applying the iterative TMLE procedure described in Section 3 to initial estimates \bar{Q}_n^0, g_n, f_n^e , where k denotes the final step. Convergence was defined as $abs(\varepsilon_1) < 0.00001$ and $abs(\varepsilon_2) < 0.00001$, and typically occurred after two to three iterations (i.e., k typically equals 2 or 3). Table 1 also reports unadjusted estimates

	Simulation 1			Simulation 2			
	Bias	Var	MSE	Bias	Var	MSE	
Unadj	0.0014	0.0017	0.0017	-0.0023	0.0017	0.0017	
TMLE	0.0015	0.0017	0.0017	-0.0006	0.0017	0.0017	
Emp eff	0.0011	0.0015	0.0015	-0.0001	0.0015	0.0015	
TMLE	0.0012	0.0015	0.0015	-0.0002	0.0015	0.0015	

Table 1: Additive treatment effect estimates, 1000 samples (n = 500).

 $\psi_n^{unadj} = E_n(Y=1 \mid A=1) - E_n(Y=1 \mid A=0)$, where $E_0(Y=1 \mid A)$ is estimated with univariate logistic regression of Y on A. The unadjusted estimator is unbiased in simulation 1, but biased in simulation 2.

Results in Table 2 verify the claim made in section 3.4 that in addition to being a double robust locally efficient substitution estimator, the enhanced TMLE is also an IPTW estimator, an augmented IPTW estimating-equation based estimator

Table 2: Alternative re	presentations of	additive tre	eatment effect	estimates.

	Simulation 1				Simulation 2			
	Bias	Var	MSE		Bias	Var	MSE	
TMLE _{en}	0.0012	0.0015	0.0015	_	-0.0002	0.0015	0.0015	
IPTW	0.0012	0.0015	0.0015	-	-0.0002	0.0015	0.0015	
$AIPTW_a$	0.0013	0.0015	0.0015	-	-0.0002	0.0015	0.0015	
$AIPTW_b$	0.0012	0.0015	0.0015	-	-0.0002	0.0015	0.0015	

in which the nuisance parameters Q_0, g_0 are estimated with the TMLE Q_n^*, g_n^* , and an augmented IPTW estimating-equation based estimating in which the nuisance parameters Q_0, g_0 are estimated with $(Q_{W,n}, \bar{Q}_n^e), g_n^*$. Recall that the latter is the empirical efficiency maximization estimator of Rubin and van der Laan (2008), except that g_0 is estimated with g_n^* instead of the initial estimator g_n .

We next investigate estimator performance under increasing levels of confounding. In simulation 3a the treatment assignment mechanism is held fixed and confounding is made stronger by increasing the association between W_2 and the outcome, Y. In simulation 3b the conditional distribution of Y given (A, W) is held fixed while the association between W_2 and A increases, leading to violations of the positivity assumption as confounding grows stronger. For each simulation estimates were obtained for 1000 samples of size n = 500 with $g_n(1 \mid W)$ bounded away from 0 and 1 at level (p, 1-p), with $p = (10^{-9}, 0.01, 0.025, 0.05, 0.1)$.

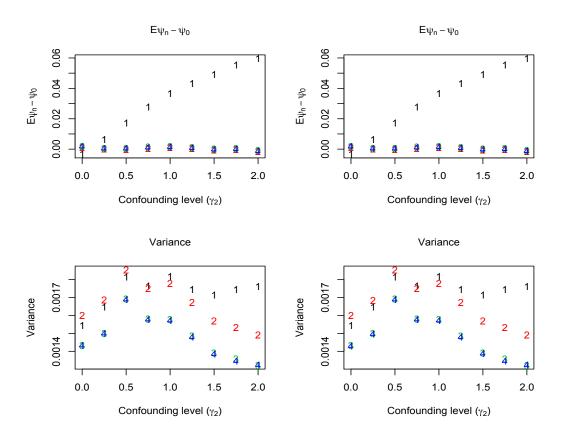


Figure 1: Simulation 3a: Estimator bias and variance at each value of γ_2 , two truncation levels for $g_n(1 \mid W)$, $(10^{-9}, 1 - 10^{-9})$ (left), and (0.1, 0.9) (right). Symbols on the plot refer to 1: Unadjusted, 2: TMLE, 3: Emp Eff, 4: TMLE_{en}.

Data for simulation 3 were generated as

$$W_1, W_2 \sim N(0, 1)$$

 $g_{0,3}(1 \mid W) = \text{Expit}(-0.3 - 0.1W_1 - \gamma_1 W_2)$
 $P_{0,3}(Y = 1 \mid A, W) = \text{Expit}(-1 + A + W_1 + 2.5W_1^2 - \gamma_2 W_2)$

with γ_1 fixed at 0.3 and γ_2 set to $(0,0.1,0.2,\ldots,2)$ for simulation 3a, and γ_1 set to $(0,0.2,\ldots,1)$ while γ_2 was fixed at 1, for simulation 3b.

Figure 1 summarizes results for simulation 3a with bounds on $g_n(1 \mid W)$ set to either $(10^{-9}, 1-10^{-9})$ or (0.1, 0.9). The bias of the unadjusted estimator (1) increases with γ_2 , while the TMLE (2), Emp Eff (3), and TMLE_{en} (4) estimators remain unbiased. When confounding is strong, the unadjusted estimator has the highest variance, followed by TMLE, while as predicted by theory, the variance of the TMLE_{en} estimator closely matches that of the empirical efficiency estimator, designed to minimize variance. Because the treatment assignment mechanism does not lead to a violation of the positivity assumption $(0.14 < g_0(1 \mid W) < 0.77)$, results are the same regardless of the choice of truncation level for $g_n(1 \mid W)$.

Estimator performance under increasing practical violations of the positivity assumption is illustrated in Figure 2, which shows results at three truncation levels of $g_n(1 \mid W)$: $(10^{-9}, 1-10^{-9})$, (0.025, 0.975), and (0.1, 0.9). Increasing truncation introduces a small amount of bias into TMLE, the empirical efficiency maximization estimator, and TMLE_{en}, but this amount is dwarfed by the bias of the unadjusted estimator increases with increased confounding, and is slightly ameliorated by increased truncation of $g_n(1 \mid W)$. At extreme violations of the positivity assumption (see Table 3) the variance of TMLE_{en}(4) is slightly larger than that of the empirical efficiency maximization estimator (3), but overall these two estimators are very close to one another.

Table 3: True conditional treatment assignment probabilities as a function of γ_1 .

γ_1	Range of $g_0(A \mid W)$		γ_1	Range of $g_0(A \mid W)$		
0	0.305	0.551	0.6	0.035	0.926	
0.2	0.212	0.676	0.8	0.013	0.969	
0.4	0.090	0.837	1	0.005	0.987	

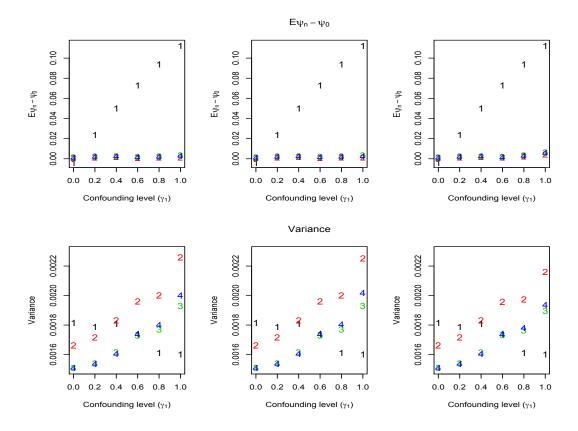


Figure 2: Simulation 3b: Estimator bias and variance at each value of γ_1 . Columns correspond to truncation level for $g_n(1 \mid W)$, $(10^{-9}, 1 - 10^{-9})$ (left), (0.025, 0.05) (center), and (0.1, 0.9) (right). Symbols on the plot refer to 1: Unadjusted, 2: TMLE, 3: Emp Eff, 4: TMLE_{en}.

6 Discussion

The TMLE represents a template for construction of a loss-based substitution estimator of a target parameter defined on a semiparametric model, defined by a choice of loss function for a relevant part of the data generating distribution, a parametric submodel, and a strategy for iteratively minimizing the empirical risk over the parametric submodel. The choice of submodel and loss function defines the score equations the TMLE will solve. In this manner it can be arranged that the TMLE solves not only the efficient score equation, but also an estimating equation corresponding with the influence curve of a competing estimator. By solving this estimating equation the TMLE is at least as efficient as the competing estimator in the case this competing estimator is asymptotically linear and g_n^0 is consistent.

In this article we demonstrated this type of TMLE for the simple point treatment data structure (W,A,Y) and the additive effect parameter. Our presentation is straightforwardly generalized to general CAR-censored data models, and target parameters, since we only relied on a general representation of the efficient influence curve as an augmented IPCW-function as presented in Robins and Rotnitzky (1992), van der Laan and Robins (2003). Suppose now that the target parameter is multivariate. One needs to define the collection of real valued parameters, and one needs to define a competing estimator for each of these real valued parameters. For example, one might define one single real valued parameter as a function of the multivariate parameter, or one might define each component of the target parameter as a real valued parameter. Each of the real valued parameters now implies an influence curve of the corresponding competing estimator. Each of these influence curves implies a clever covariate for the treatment mechanism playing the role of $H(Q_n^e, g_n^k)$ in the above TMLE algorithm. The resulting TMLE will not only be a double robust locally efficient substitution estimator of the target parameter, but it will also estimate each of the real valued parameters in a more efficient way than the competing estimators, in the case that g_0 is estimated consistently.

Appendix: R Implementation

The R function below calculates the enhanced TMLE for binary outcomes. Required arguments are Y (binary outcome vector), A (binary treatment indicator vector), and initial estimates $\bar{Q}_n^0(A,W)$, $g_n^0(A\mid W)$, and $f_n^e(W)$. $\bar{Q}_n^0(A,W)$ is an $n\times 3$ matrix containing values for $\bar{Q}_n^0(A,W)$, $\bar{Q}_n^0(0,W)$, and $\bar{Q}_n^0(1,W)$ on the logit scale. Predicted values for $g_n(A\mid W)$ are bounded away from 0 and 1.

```
bound <- function(x, bounds){</pre>
  x[x<min(bounds)] <- min(bounds)</pre>
  x[x>max(bounds)] <- max(bounds)</pre>
  return(x)
}
tmle_en \leftarrow function(Y, A, g1W, Q, f, gbds = c(10^-9, 1-10^-9)){
  g1w <- bound(g1W, gbds)
  eps1 <- eps2 <- Inf
  epsilon <- .00001
  maxIter <- 30
  iterations <- 0
  while((any(abs(c(eps1, eps2)) > epsilon)) & iterations <= maxIter){</pre>
    iterations <- iterations + 1
    h \leftarrow cbind(A/g1W - (1-A)/(1-g1W), 1/g1W, -1/(1-g1W))
    m \leftarrow glm(Y \sim -1 + offset(Q[,"QAW"]) + h[,1], family=binomial)
    eps1 <- coef(m)
```

```
Q \leftarrow Q + eps1*h
    h2 \leftarrow plogis(Q[,"Q1W"])/g1W + plogis(Q[,"Q0W"])/(1-g1W)
    h3 \leftarrow f/(g1W * (1-g1W))
    g <- glm(A ~ -1 + offset(qlogis(g1W)) + h2 + h3, family=binomial)
    g1W <- bound(predict(g, type = "response"), gbds)
    eps2 <- coef(g)
  Q <- plogis(Q)
  psi.en
          <- mean(Q[,"Q1W"] - Q[,"Q0W"])
  psi.IPTW \leftarrow mean((A/g1W - (1-A)/(1-g1W)) * Y)
 psi.AIPTWQstargstar <- mean((A/g1W - (1-A)/(1-g1W)) * Y
                          - (Q[,"Q1W"]/g1W - Q[,"Q0W"]/(1-g1W))*(A-g1W))
 psi.AIPTWQegstar <- mean((A/g1W - (1-A)/(1-g1W)) * Y
                          - f/(g1W * (1-g1W)) * (A-g1W))
  return(c(psi.en, psi.IPTW, psi.AIPTWQstargstar, psi.AIPTWQegstar))
}
```

References

- Cao, W., A. Tsiatis, and M. Davidian (2009): "Improving efficiency and robustness of the doubly robust estimator for a population mean with incomplete data," *Biometrika*, 96,3, 723–734.
- Gruber, S. and M. van der Laan (2010a): "An application of collaborative targeted maximum likelihood estimation in causal inference and genomics," *The International Journal of Biostatistics*.
- Gruber, S. and M. van der Laan (2010b): "A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome," *The International Journal of Biostatistics*, 6:1.
- Pearl, J. (2000): *Causality: Models, Reasoning, and Inference*, Cambridge University Press, Cambridge, 2nd edition.
- Robins, J. and A. Rotnitzky (1992): "Recovery of information and adjustment for dependent censoring using surrogate markers," in *AIDS Epidemiology*, Methodological issues, Bikhäuser.
- Rotnitzky, A., Q. Lei, M. Sued, and J. Robins (2012): "Improved double-robust estimation missing data and causal inference models," *Biometrika*, 99.
- Rubin, D. and M. van der Laan (2008): "Empirical efficiency maximization: Improved locally efficient covariate adjustment in randomized experiments and survival analysis," *The International Journal of Biostatistics*, Vol. 4, Iss. 1, Article 5.
- Tan, Z. (2008): "Comment: Improved local efficiency and double robustness," *International Journal of Biostatistics*.

- Tan, Z. (2010): "Bounded, efficient, and doubly robust estimation with inverse weighting," *Biometrika*, 97, 661–82.
- Team, R. D. C. (2010): *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
- van der Laan, M. and S. Dudoit (2003): "Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilonnet estimator: Finite sample oracle inequalities and examples," Technical report, Division of Biostatistics, University of California, Berkeley.
- van der Laan, M., E. Polley, and A. Hubbard (2007): "Super learner," *Statistical Applications in Genetics and Molecular Biology*, 6.
- van der Laan, M. and J. Robins (2003): *Unified methods for censored longitudinal data and causality*, Springer, New York.
- van der Laan, M. and S. Rose (2011): *Targeted Learning: Prediction and Causal Inference for Observational and Experimental Data*, Springer, New York.
- van der Laan, M. and D. Rubin (2006): "Targeted maximum likelihood learning," *The International Journal of Biostatistics*, 2.
- van der Laan, M. J. (2008): "The construction and analysis of adaptive group sequential designs," Technical Report 232, www.bepress.com/ucbbiostat/paper232, University of California, Berkeley.
- Zheng, W. and M. J. van der Laan (2010): "Asymptotic theory for cross-validated targeted maximum likelihood estimation," Technical Report 273, www.bepress.com/ucbbiostat/paper273, University of California, Berkeley.