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

Year 2010 *Paper* 265

A Targeted Maximum Likelihood Estimator of a Causal Effect on a Bounded Continuous Outcome

Susan Gruber* Mark J. van der Laan†

*University of California, Berkeley, sgruber@hsph.harvard.edu

†University of California, Berkeley, laan@berkeley.edu

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

http://biostats.bepress.com/ucbbiostat/paper265

Copyright ©2010 by the authors.

A Targeted Maximum Likelihood Estimator of a Causal Effect on a Bounded Continuous Outcome

Susan Gruber and Mark J. van der Laan

Abstract

Targeted maximum likelihood estimation of a parameter of a data generating distribution, known to be an element of a semiparametric model, involves constructing a parametric model through an initial density estimator with parameter epsilon representing an amount of fluctuation of the initial density estimator, where the score of this fluctuation model at epsilon=0 equals the efficient influence curve/canonical gradient. The latter constraint can be satisfied by many parametric fluctuation models, since it represents only a local constraint of its behavior at zero fluctuation. However, it is very important that the fluctuations stay within the semiparametric model for the observed data distribution, even if the parameter can be defined on fluctuations that fall outside the assumed observed data model. In particular, in the context of sparse data, a violation of this property can heavily affect the performance of the estimator. We demonstrate this in the context of estimation of a causal effect of a binary treatment on a continuous outcome that is bounded. It results in a targeted maximum likelihood estimator that inherently respects known bounds, and consequently is more robust in sparse data situations than the targeted MLE using a naive fluctuation model.

1 Introduction.

Targeted maximum likelihood estimation (TMLE) yields semiparametric efficient substitution estimators of parameters in semiparametric models (van der Laan and Rubin, 2006). In particular, it can be applied to estimating the statistical counterpart of a causal parameter. In this article a new targeted maximum likelihood estimator for estimating a causal effect of a binary treatment on a continuous outcome is introduced. This estimator is more robust than a previously presented TMLE procedure when there is sparsity in the data that decreases the identifiability of the parameter of interest.

Section 2 of the paper provides background on the application of TMLE methodology in the context of sparsity, and its power relative to other semiparametric efficient estimators by being a substitution estimator respecting global constraints of the semiparametric model. Even though an estimator can be asymptotically efficient without utilizing global constraints, the global constraints are instrumental in the context of sparsity with respect to the target parameter, motivating the need for semiparametric efficient *substitution* estimators, and for a careful choice of fluctuation function for the targeted MLE step that fully respects these global constraints. A rigorous demonstration of the proposed targeted MLE of the causal effect of a binary treatment on a bounded continuous outcome follows, and it is contrasted to a targeted MLE that makes use a fluctuation function that does not respect the bounds.

Simulation studies described in Section 3 compare the new TMLE estimator of the causal effect, which relies on a logistic fluctuation of an initial density estimate, with the traditional TMLE estimator, with and without sparsity in the data. Results for other commonly applied estimators, the inverse-probability-of-treatment weighted estimator (IPTW) (Hernan et al., 2000; Robins, 2000b), a double robust augmented IPTW estimator (aug-IPTW) (Robins and Rotnitzky, 2001; Robins et al., 2000; Robins, 2000a) that is efficient but not a substitution estimator, and the maximum likelihood substitution estimator according to a parametric model (MLE) (Robins, 1986) are also presented.

2 TMLE for causal effect estimation on a continuous outcome.

The targeted MLE is a semiparametric efficient substitution estimator of a target parameter $\Psi(P_0)$ of a true distribution $P_0 \in \mathcal{M}$, based on sampling n i.i.d. $O_1, \ldots, O_n$ from P_0. Here P_0 is known to be an element of a semiparametric statistical model $\mathcal{M}$. We will start with providing a succinct summary of how it works. For more

COBRA
Collection of Biostatistics
Research Archive

details we refer to our articles on this topic (van der Laan et al., 2009).

Firstly, one notes that $\Psi(P_0) = \Psi(Q_0)$ only depends on P_0 through a relevant part $Q_0 = Q(P_0)$ of P_0. Secondly, one proposes a loss function $L(Q)(O)$ so that $Q_0 = \arg\min_{Q \in \mathcal{Q}} E_0 L(Q)(O)$, where $\mathcal{Q} = \{Q(P) : P \in \mathcal{M}\}$. Thirdly, one uses minimum loss-based learning, such as super learning (van der Laan et al., 2007), fully utilizing the power and optimality results for loss-based cross-validation to select among candidate estimators, to obtain an initial estimator Q_n^0 of Q_0. Fourthly, one proposes a parametric fluctuation $Q_{ng}^0(\epsilon)$, possibly indexed by nuisance parameter $g_0 = g(P_0)$, so that

$$\frac{d}{d\epsilon} L(Q_{ng}^0(\epsilon))(O)\Big|_{\epsilon=0} = D^*(Q_n^0, g)(O), \tag{1}$$

where $D^*(Q_0, g_0)$ is the canonical gradient/efficient influence curve of $\Psi : \mathcal{M} \to \mathbb{R}$ at P_0. Fifthly, one computes the amount of fluctuation

$$\epsilon_n = \arg\min_{\epsilon} \sum_{i=1}^{n} L(Q_{ng_n}^0(\epsilon))(O_i),$$

where g_n is an estimator of the unknown nuisance parameter g_0. This yields an update $Q_n^1 = Q_{ng_n}^0(\epsilon_n)$. This updating of an initial estimator Q_n^0 into a next Q_n^1 is iterated till convergence resulting in a Q_n^*. Since at the last step the amount of fluctuation $\epsilon_n \approx= 0$, this final Q_n^* will solve the efficient influence curve estimating equation

$$0 = \sum_{i=1}^{n} D^*(Q_n^*, g_n)(O_i),$$

representing a fundamental ingredient for establishing asymptotic efficiency of $\Psi(Q_n^*)$: recall that an estimator is efficient if and only if it is asymptotically linear with influence curve equal to the efficient influence curve $D^*(Q_0, g_0)$. Finally, the targeted MLE of ψ_0 is the substitution estimator $\Psi(Q_n^*)$.

Thus we see that the targeted MLE involves constructing a parametric model $Q_n^0(\epsilon)$ through the initial estimator Q_n^0 with parameter ϵ representing an amount of fluctuation of the initial estimator, where the score of this fluctuation model at $\epsilon = 0$ equals the efficient influence curve. The latter constraint can be satisfied by many parametric models, since it represents only a local constraint of its behavior at zero fluctuation. However, it is very important that the fluctuations stay within the model for the observed data distribution, even if the parameter can be defined on fluctuations that fall outside the assumed observed data model. In particular, in the context of sparse data, a violation of this property can heavily affect the performance of the estimator.

COBRA
Collection of Biostatistics
Research Archive

One important strength of the semiparametric efficient targeted MLE relative to the alternative semiparametric efficient estimating equation methodology (van der Laan and Robins, 2003) is that it does respect the global constraints of the observed data model since it is a substitution estimator $\Psi(Q_n^*)$ with Q_n^* an estimator of a relevant part Q_0 of the true distribution of the data in the observed data model. The estimating equation methodology does not result in substitution estimators and thereby often ignore important global constraints of the observed data model, which comes at a price in the context of sparsity. Indeed, simulations have confirmed this gain of targeted MLE relative to the efficient estimating equation method in the context of sparsity (Stitelman and van der Laan, 2010), and it is again demonstrated in this article. However, if the targeted MLE starts violating this principle of being a substitution estimator by allowing Q_n^* to fall outside the assumed observed data model, this advantage is compromised. Therefore, it is crucial that a fluctuation model is used that is guaranteed to stay within the wished observed data model.

To demonstrate this important consideration of selecting a valid fluctuation model in the construction of targeted MLE, we consider the problem of estimating a causal effect of a binary treatment A on a continuous outcome Y, based on observing n i.i.d. copies of $O = (W, A, Y) \sim P_0$, where W is the set of confounders. Under nonparametric structural equation model (NPSEM) $W = f_W(U_W)$, $A = f_A(W, U_A)$, $Y = f_Y(W, A, U_Y)$ with a structure on the exogenous variables $U = (U_W, U_A, U_Y)$ satisfying the no unmeasured confounder assumption ($A \perp Y(a) \mid W$ for the counterfactuals $Y(a)$ defined by this NPSEM), the additive causal effect $E(Y(1) - Y(0))$ can be identified from the observed data distribution through the following statistical parameter of P_0:

$$\Psi(P_0) = E_0(E_0(Y \mid A = 1, W) - E_0(Y \mid A = 0, W)).$$

Suppose that it is known that $Y \in [a, b]$ for some $a < b$. Alternatively, one might have truncated the original data to fall in such an interval and focus on the causal effect of treatment on this truncated outcome, motivated by the fact that estimating conditional means of unbounded, or very heavy tailed, outcomes requires very large data sets.

Let $Y^* = (Y - a)/(b - a)$ be the linearly transformed outcome within $[0, 1]$, and define

$$\Psi^*(P_0) = E_0(E_0(Y^* \mid A = 1, W) - E_0(Y^* \mid A = 0, W)).$$

We note that

$$\Psi(P_0) = (b - a)\Psi^*(P_0).$$

An estimate, limit distribution, and confidence interval for $\Psi^*(P_0)$ is now immediately mapped into an estimate, limit distribution, and confidence interval for $\Psi(P_0)$,

COBRA
Collection of Biostatistics
Research Archive

by simple multiplication by $(b-a)$. As a consequence, without loss of generality, we can assume $a=0$ and $b=1$ so that $Y \in [0,1]$.

The efficient influence curve of the statistical parameter $\Psi : \mathcal{M} \to \mathbb{R}$, defined on a nonparametric statistical model $\mathcal{M}$ for P_0, at the true distribution P_0, is given by

$$D^*(P_0) = \frac{2A-1}{g_0(A \mid W)}(Y - \bar{Q}_0(W,A)) + \bar{Q}_0(1,W) - \bar{Q}_0(0,W) - \Psi(Q_0),$$

where $\bar{Q}_0(W,A) = E_0(Y \mid W,A)$, and $Q_0 = (Q_W, \bar{Q}_0)$ denotes both this conditional mean $\bar{Q}_0$ as well as the marginal distribution Q_W of W. Note that indeed $\Psi(P_0)$ only depends on P_0 through $\bar{Q}_0$ and the marginal distribution of W. We will use the notation $\Psi(P_0)$ and $\Psi(Q_0)$ interchangeably.

We will now define a targeted MLE of $\Psi(Q_0)$ as follows. Let $\bar{Q}_n^0$ be an initial estimator of $\bar{Q}_0(W,A) = E(Y \mid A,W)$ with predicted values in $(0,1)$. In addition, we estimate P_W with the empirical distribution of W. Let Q_n^0 denote the resulting initial estimator of Q_0. The targeted MLE step will also require an estimator g_n of $g_0 = P_{A|W}$. Only the conditional mean $\bar{Q}_n^0$ will be modified by the targeted MLE procedure defined below: this makes sense since the empirical distribution of W is already a nonparametric maximum likelihood estimator so that no bias gain with respect to the target parameter will be obtained by modifying it.

We can represent the estimator $\bar{Q}_n^0$ as $\bar{Q}_n^0 = \frac{1}{1+\exp(-f_n^0)}$ with $f_n^0 = \log(\bar{Q}_n^0/(1-\bar{Q}_n^0))$. Consider now the fluctuation model

$$\bar{Q}_n^0(\epsilon) = \frac{1}{1+\exp(-\{f_n^0 + \epsilon h\})},$$

with parameter ϵ, indexed by a function

$$h(g_n)(W,A) = \frac{2A-1}{g_n(A \mid W)}.$$

Equivalently, we can write this as $\text{logit}\bar{Q}_n^0(\epsilon) = \text{logit}\bar{Q}_n^0 + \epsilon h(g_n)$.

Consider now the following loss function for $\bar{Q}_0$:

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

Note that this is the log-likelihood of the conditional distribuiton of a binary outcome Y, but now extended to continuous outcomes in $[0,1]$. It is thus known that this loss function is a valid loss function for the conditional distribution of a binary Y, but we need that it is a valid loss function for a conditional mean of a continuous $Y \in [0,1]$. We have the following lemma establishing this result about this loss function.

UC BERKELEY
Collection of Biostatistics
Research Archive

Lemma 1 *We have that*

$$\bar{Q}_0 = \underset{\bar{Q}}{\text{argmin}}\ E_0 L(\bar{Q}),$$

where the minimum is taken over all functions of (W, A) *which map into* $(0, 1)$*. In addition, define the fluctuation function*

$$logit\bar{Q}(\epsilon) = logitQ + \epsilon h.$$

For any function h *we have*

$$\frac{d}{d\epsilon} L(\bar{Q}(\epsilon))\Big|_{\epsilon=0} = h(W, A)(Y - \bar{Q}(W, A)).$$

Proof: Let Q_1 be a local minimum and consider the fluctuation $Q_1(\epsilon)$ defined above. Then the derivative of $E_0 L(Q_1(\epsilon))$ at $\epsilon = 0$ equals zero. However,

$$-\frac{d}{d\epsilon} L(Q_1(\epsilon))\Big|_{\epsilon=0} = h(W, A)(Y - Q_1(W, A)).$$

Thus, it follows that

$$E_0 h(W, A)(Y - Q_1(W, A)) = E_0 h(W, A)(Q_0 - Q_1)(W, A).$$

But this needs to hold for any function $h(W, A)$, which proves that $Q_1 = Q_0$ a.e. □

This proves that $L(\bar{Q})$ is a valid loss function for the conditional mean $\bar{Q}_0$. Indeed, we can use $L(\bar{Q})$ as loss function to construct an initial estimator of $\bar{Q}_0$, and or use cross-validation to select among candidate targeted maximum likelihood estimators, such as in the collaborative targeted MLE procedure. For the purpose of construction of an initial estimator one could also use a minimum loss-based super learner based on the squared error loss function $L_2(\bar{Q}) = (Y - \bar{Q}(W, A))^2$, possibly with weights.

Given an initial estimator $\bar{Q}_n^0$, and our proposed fluctuation function $\bar{Q}_n^0(\epsilon)$, we have

$$\frac{d}{d\epsilon} L(\bar{Q}_n^0(\epsilon))\Big|_{\epsilon=0} = h(g)(W, A)(Y - \bar{Q}_n^0(W, A)),$$

giving us the wished first component D_1^* of the efficient influence curve $D^* = D_1^* + D_2^*$.

Let's use the log-likelihood loss function, $-logQ_W$, as loss function for the marginal distribution of W, so that our combined loss function is given by $L(Q) = -logQ_W + L(\bar{Q})$. In addition, we use as fluctuation of the empirical distribution Q_{Wn}, $Q_{Wn}(\epsilon_1) = (1 + \epsilon_1 D_2^*(Q))Q_{Wn}$, where $D_2^*(Q) = \bar{Q}(W, 1) - \bar{Q}(W, 0) - \Psi(Q)$

Collection of Biostatistics
Research Archive

is the remaining component of the efficient influence curve. With these choices we indeed now have that

$$\frac{d}{d\epsilon} L(Q(\epsilon))\bigg|_{\epsilon=0} = D^*(Q, g).$$

This shows that we succeeded in defining a loss function for $Q_0 = (Q_W, \bar{Q}_0)$ and fluctuation function so that the wished derivative (1) indeed yields the efficient influence curve.

The MLE of ϵ_1 equals zero, so that the update of Q_{Wn} equals Q_{Wn} itself. The empirical mean of the component $D_2^* = \bar{Q}(W, 1) - \bar{Q}(W, 0) - \Psi(Q)$ of the efficient influence curve is always equal to zero, due to the fact that we estimate the marginal distribution of W with the empirical distribution of W.

The amount of fluctuation of ϵ for fluctuating $\bar{Q}_n^0$ is given by

$$\epsilon_n^0 = \underset{\epsilon}{\operatorname{argmin}}\ P_n L(\bar{Q}_n^0(\epsilon)).$$

This "maximum likelihood" estimator of ϵ can be computed with generalized linear regression using the binomial link, i.e. the logistic regression MLE procedure, simply ignoring that the outcome is not binary, which also corresponds with iterative reweighted least squares estimation using weights $1/Q(1-Q)$.

This provides us with the targeted MLE update $Q_n^1 = Q_n^0(\epsilon_n^0)$, where the empirical distribution of W did not get updated, and $\bar{Q}_n^0$ did get updated as $\bar{Q}_n^0(\epsilon_n^0)$. Iterating this procedure now defines the targeted MLE Q_n^*, but as in the binary outcome case, we have that $Q_n^2 = Q_n^1(\epsilon_n^1) = Q_n^1$ since the next MLE $\epsilon_n^1 = 0$. Thus convergence occurs in one step, so that $Q_n^* = Q_n^1$. The targeted MLE of ψ_0 is thus given by $\Psi(Q_n^*) = \Psi(Q_n^1)$. As predicted, we have that the targeted MLE Q_n^* solves the efficient influence curve estimating equation $P_n D^*(Q_n^*, g_n, \Psi(Q_n^*)) = 0$.

We note that, even if there is strong confounding causing some large values of $h_{g_n^0}$, the resulting targeted MLE $\bar{Q}_n^*$ remains bounded in $(0, 1)$, so that the targeted MLE $\Psi(Q_n^*)$ fully respects the global constraints of the observed data model. On the other hand, the augmented IPTW estimator obtained by solving $P_n D^*(Q_n^0, g_n, \psi) = 0$ in ψ yields the estimator

$$\psi_n = \frac{1}{n} \sum_{i=1}^{n} h_{g_n}^0(W_i, A_i)(Y_i - \bar{Q}_n^0(W_i, A_i)) + \bar{Q}_n^0(W, 1) - \bar{Q}_n^0(W, 0),$$

which can easily fall outside $[0, 1]$ if for some observations W_i, $g_n(1 \mid W_i)$ is close to 1 or 0. This represents the price of not being a substitution estimator.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Contrasting with targeted MLE using linear fluctuation function. Alternatively, we would employ the targeted MLE using the $L_2(\bar{Q}) = (Y - \bar{Q}(W, A))^2$ loss function, and fluctuation function $\bar{Q}^0(\epsilon) = \bar{Q}^0 + \epsilon h(g)$, so that (1) is still satisfied. In this case, large values of $h(g)$ will result in predicted values of $\bar{Q}^0(\epsilon_n)$ that are out of the bounds $[a, b]$. Therefore, this version of targeted MLE is not respecting the global constraints of the model, i.e., the knowledge that $Y \in [a, b]$. A comparison based on simulated data of the targeted MLE using the logistic fluctuation function and the targeted MLE using this linear fluctuation function is provided in the next section.

3 Simulation studies for the additive effect of a binary point treatment on a continuous outcome.

Two simulation studies illustrate the effects of employing a logistic vs. linear fluctuation on TMLE estimator performance with and without sparsity in the data, where a high degree of sparsity corresponds to a target parameter that is borderline-identifiable. As above, the parameter of interest is defined as the marginal effect of a binary point treatment on the outcome, $\psi_0 = E_W[E[Y \mid A = 1, W] - E[Y \mid A = 0, W]]$.

The "traditional" targeted maximum likelihood approach to estimating an additive treatment effect when the outcome is continuous is to fluctuate the initial density estimate on a linear scale. Given $\bar{Q}_n^0(A, W)$, an initial estimate of the conditional mean of Y given (A, W), the fluctuation function is defined as $\bar{Q}_n^0(\epsilon) = \bar{Q}_n^0 + \epsilon(h_{g_n})$ and the loss function $L(\bar{Q})$ is chosen to be the squared error loss function, so that we still have the required constraint (1). The estimate ϵ_n can be obtained by estimating ϵ with a linear regression of Y on h_{g_n}, using the initial fit, $\bar{Q}_n^0(A, W)$, as offset.

A second TMLE estimate using the logistic fluctuation method described in Section 2 is also obtained. Y is transformed into $Y^* \in [0, 1]$ by shifting and scaling the values. In the simulation setting, Y is not bounded, so that we do not have an a priori a and b bound on Y. Instead of truncating Y and redefining the target parameter as the causal effect on the truncated Y, we still aim to estimate the causal effect on the original Y. Therefore, we set $a = \min(Y), b = \max(Y)$, and

$$Y^* = \frac{Y - a}{b - a}.$$

An initial estimate, $\bar{Q}_n^{0,Y^*}(A, W) = E(Y^*|A, W)$, is obtained, and then represented as a logistic function of its logit-transformation. Note that $\text{logit}(x)$ is not defined

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

when $x = 0$ or 1. Therefore in practice $\bar{Q}_n^{0,Y^*}(A, W)$ is bounded away from 0 and 1 by truncating it at $(\alpha, (1-\alpha))$. We used $\alpha = 0.005$ in these simulation studies, which did not yield appreciably different results than setting $\alpha = 0.001$ or $\alpha = 0.01$. The function $\bar{Q}_n^{0,Y*}$ is fluctuated on the logit scale with logit$\bar{Q}_n^{0,Y^*}(\epsilon) =$ logit$\bar{Q}_n^{0,Y^*} + \epsilon h(g_n)$, using the same clever covariate, $h_{g_n}(A, W)$, employed in the linear fluctuation described above. Fitting ϵ is again carried out using standard software, but this time using logistic regression of Y^* on $h_{g_n}(A, W)$ with offset $logit(Q_n^{0,Y^*}(A, W))$. This results in the updated $\bar{Q}_n^{1,Y^*}$. Fitted values for $\bar{Q}_n^{1,Y^*}(A, W)$ are mapped back to the original scale: $\bar{Q}_n^{1,Y} = \bar{Q}_n^{1,Y^*}(A, W)*(b-a)+a$. The marginal distribution is estimated with the empirical distribution of W, giving the $Q_n^* = Q_n^1 = (Q_{W,n}, \bar{Q}_n^{1,Y})$ of $(Q_W, \bar{Q}_0)$. The estimate

$$\psi_n = \Psi(Q_n^*) = \frac{1}{n}\sum_{i=1}^{n} \bar{Q}_n^{1,Y})(1, W_i) - \bar{Q}_n^{1,Y}(0, W_i)$$

is the targeted MLE of the wished additive causal effect ψ_0.

Parameter estimates were also obtained using the augmented inverse probability of treatment weighed estimator (aug-IPTW)

$$\begin{aligned}\psi_n^{aug-IPTW} &= \frac{1}{n}\sum_{i=1}^{n} \frac{2A-1}{g_n(A_i \mid W_i)}(Y_i - \bar{Q}_n^0(W_i, A_i)) \\ &\quad + \frac{1}{n}\sum_{i=1}^{n}(\bar{Q}_n^0(1, W_i) - \bar{Q}_n^0(0, W_i)).\end{aligned}$$

Both the targeted MLE and the augmented IPTW estimator are double robust so that these estimators will be consistent for ψ_0 if either g_n or $\bar{Q}_n^0$ is consistent for g_0 and $\bar{Q}_0$, respectively. Both the targeted MLE and the augmented IPTW estimator are asymptotically efficient if both g_n and $\bar{Q}_n^0$ are consistent.

In this simulation study we will use simple parametric MLE's as initial estimators $\bar{Q}_n^0$ and g_n, even though we recommend the utilization of super learning in practice. The purpose of this simulation is to investigate the performance of the updating step under misspecified and correctly specified $\bar{Q}_n^0$, and for that purpose we can work with parametric MLE fits.

Results from two estimation methods that are not double robust and semiparametric efficient are included as well. The maximum likelihood estimator according to a parametric model for $\bar{Q}_0$ (MLE), used as initial estimator in the targeted MLE and augmented IPTW, is included for the sake of evaluating the bias reduction step carried out by these two semiparametric efficient procedures. Inverse probability of treatment weighted (IPTW) estimators are consistent when $g_n(A, W)$ is a consistent

Collection of Biostatistics
Research Archive

estimator of the treatment mechanism $g_0(A, W) = P(A = 1|W)$, but are known to be inefficient. These two estimators are defined as

$$\begin{aligned}
\psi_n^{MLE} &= \frac{1}{n}\sum_{i-1}^{n}(\bar{Q}_n^0(1, W_i) - \bar{Q}_n^0(0, W_i), \\
\psi_n^{IPTW} &= \frac{1}{n}\sum_{i=1}^{n}(2A-1)\frac{Y_i}{g_n(A_i, W_i)}.
\end{aligned}$$

3.1 Data generation

Covariates W_1, W_2, W_3 were generated as independent binary random variables,

$$W_1, W_2, W_3 \sim Bernoulli(0.5).$$

Two treatment mechanisms were defined that differ only in the values of the coefficients for each covariate:

$$g_0(1 \mid W) = P(A = 1 \mid W) \quad = \quad \text{logit}^{-1}(aW_1 + bW_2 + cW_3).$$

We consider two settings:

$$a_1 = 0.5, b_1 = 1.5, c_1 = -1 \text{ and } a_2 = 1.5, b_2 = 4.5, c_2 = -3.$$

We refer to these two treatment mechanisms as $g_{0,1}$ and $g_{0,2}$, respectively. The observed outcome Y was generated as

$$\begin{aligned}
Y &= \bar{Q}_0(A, W) + e,\ e \sim N(0, 1), \\
\bar{Q}_0(A, W) &= A_j + 2W_1 + 3W_2 - 4W_3.
\end{aligned}$$

For both simulations the true additive causal effect equals one: $\psi_0 = 1$. Treatment assignment probabilities based on mechanism $g_{0,1}$ range from 0.269 to 0.881, indicating no sparsity in the data for simulation 1. In contrast, treatment assignment probabilities based on mechanism $g_{0,2}$ range from (0.047 to 0.998). Simulation 2 poses a more challenging estimation problem in the context of sparse data. In both simulations predicted values for $g_n(A \mid W)$ are bounded away from 0 and 1 by truncating at $(p, (1-p))$, with $p = 0.01$.

Estimates were obtained for 1000 samples of size $n = 1000$ from each data generating distribution. Treatment assignment probabilities, $g_0(A \mid W)$, were estimated using a correctly specified logistic regression model. A correctly specified main terms regression model was used to obtain $\bar{Q}_{cor}^0(A, W)$. In addition, a misspecified initial estimate, $\bar{Q}_{mis}^0(A, W)$, was obtained by regressing Y on A.

COBRA
Collection of Biostatistics
Research Archive

We expect MLE estimates based on $\bar{Q}^0_{cor}$ to be unbiased and efficient, while those based on $\bar{Q}^0_{mis}$ will be biased. IPTW estimates only depend on consistent estimation of g_0, thus are identical regardless of how $\bar{Q}_0$ is estimated. For both simulations g_n is a consistent estimator, thus it is reasonable to expect unbiased IPTW estimates, with more variation in simulation 2 estimates. The targeted MLE and the augmented IPTW are known to be unbiased if g_n is consistent, and asymptotically efficient when both $\bar{Q}_0$ and g_0 are consistently estimated. Though correctly estimating g_0 will asymptotically correct for any bias due to mis-specification of $\bar{Q}^0_n$, this is not guaranteed in finite samples, especially when there is sparsity. For simulation 2 we expect TMLE$_{log}$, using the logistic fluctuation, to outperform TMLE$_{lin}$, using the linear fluctuation.

3.2 Results

Table 1 reports the average estimate, bias, empirical variance, and mean squared error (MSE) for each estimator, under different specifications of the initial estimator $\bar{Q}^0_n$. In all cases g_n is consistent, and bounded at (0.01, 0.99).

Table 1: Estimator performance for simulations 1 and 2 when the initial estimator of $\bar{Q}_0$ is correct and misspecified. Results are based on 1000 samples of size $n = 1000$, $g_n(A, W)$ bounded at (.01,.99) for all estimators.

	$\bar{Q}_0$ **correctly estimated**				$\bar{Q}_0$ **incorrectly estimated**			
	ave	bias	var	MSE	ave	bias	var	MSE
Simulation 1								
MLE	1.003	0.003	0.005	0.005	3.075	2.075	0.030	4.336
IPTW	1.006	0.006	0.009	0.009	1.006	0.006	0.009	0.009
aug-IPTW	1.003	0.003	0.005	0.005	1.005	0.005	0.010	0.010
TMLE$_{log}$	0.993	−0.007	0.005	0.005	0.993	−0.007	0.006	0.006
TMLE$_{lin}$	0.993	−0.007	0.005	0.005	0.993	−0.007	0.006	0.006
Simulation 2								
MLE	1.001	0.001	0.009	0.009	4.653	3.653	0.025	13.370
IPTW	1.554	0.554	0.179	0.485	1.554	0.554	0.179	0.485
aug-IPTW	0.999	−0.001	0.023	0.023	1.708	0.708	0.298	0.798
TMLE$_{log}$	0.989	−0.011	0.037	0.037	0.722	−0.278	0.214	0.291
TMLE$_{lin}$	0.986	−0.014	0.042	0.042	−0.263	−1.263	2.581	4.173

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

In simulation 1, when $\bar{Q}_0$ is correctly estimated all estimators perform quite well, though as expected, IPTW is the least efficient. However, when $\bar{Q}_0$ is incorrectly estimated, the MLE estimator is biased and has high variance relative to the other estimators. Because $g_n(A \mid W)$ is correctly specified, IPTW and aug-IPTW provide unbiased estimates, as do both TMLEs. TMLE$_{log}$ is on a par with TMLE$_{lin}$, as there is no sparsity in the data, and both are more efficient than any of the other estimators.

In simulation 2 all estimators except IPTW are unbiased when $\bar{Q}_0$ is correctly estimated. In this case, both TMLE estimators have higher variance than aug-IPTW, and all three are more efficient than IPTW, but less efficient than the parametric MLE estimator. Though asymptotically the IPTW estimator is expected to be unbiased in this simulation, since g_n is a consistent estimator of g_{0_2}, these results demonstrate that in finite samples, heavily weighting a subset of observations not only increases variance, but can also bias the estimate.

When the model for $\bar{Q}_0$ is misspecified in simulation 2, The MLE estimator is even more biased than it was in simulation 1. The efficiency of all three double-robust efficient estimators suffers in comparison with simulation 1 as well. Nevertheless, TMLE$_{log}$, using the logistic fluctuation, has the lowest MSE of all estimators. Its superiority over TMLE$_{lin}$ in terms of bias and variance is clear. TMLE$_{log}$ also outperforms aug-IPTW with respect to both bias and variance, and performs much better than IPTW or MLE.

4 Discussion.

When an estimation procedure incorporates weights, observations with large weights can heavily influence the point estimate and inflate the variance. Truncating these weights is a common approach to reducing the variance, but it generally introduces bias. The presented TMLE of an additive causal effect of a point treatment intervention, incorporating a logistic fluctuation of the initial conditional mean estimate, dampens the effect of these heavily weighted observations, thereby heavily reducing the reliance on truncation. As a substitution estimator, the proposed TMLE of the additive causal effect respects the global constraints of the observed data model. Simulation study results indicate that this approach is on a par with, and in the context of sparsity often superior to, fluctuating on the linear scale. In particular it is more robust when there is sparsity in the data, outperforming MLE, IPTW, and aug-IPTW.

For the sake of demonstration we considered estimation of the additive causal effect. However, the same targeted MLE, using the logistic fluctuation, can be used to estimate other point-treatment causal effects, including parameters of a marginal

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

structural model. The newly proposed loss function has also applications in prediction of a bounded outcome, and for targeted MLE of the causal effect of a multiple time point intervention in which the final outcome is bounded and continuous. We also pointed out that the proposed fluctuation function and loss function, and corresponding targeted MLE, should also be used for continuous outcomes for which no a priori bounds are known, by simply using the minimal and maximal observed outcome values. In this way, these choices naturally robustify the targeted MLE by enforcing that the updated initial estimator will not predict outcomes outside the observed range.

The TMLE approach presented here using a logistic fluctuation of an initial estimate of the conditional mean of the continuous outcome retains all properties of targeted maximum likelihood estimators, including influence curve-based inference. The method presented here extends to collaborative targeted maximum likelihood estimation without modification.

References

M. A. Hernan, B. Brumback, and J. M. Robins. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology*, 11(5):561–570, 2000.

J. M. Robins and A. Rotnitzky. Comment on the Bickel and Kwon article, "Inference for semiparametric models: Some questions and an answer". *Statistica Sinica*, 11(4):920–936, 2001.

J. M. Robins, A. Rotnitzky, and M.J. van der Laan. Comment on "On Profile Likelihood" by S.A. Murphy and A.W. van der Vaart. *Journal of the American Statistical Association – Theory and Methods*, 450:431–435, 2000.

J.M. Robins. Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association*, 2000a.

J.M. Robins. A new approach to causal inference in mortality studies with sustained exposure periods - application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7:1393–1512, 1986.

J.M. Robins. Marginal structural models versus structural nested models as tools for causal inference. In *Statistical models in epidemiology, the environment, and clinical trials (Minneapolis, MN, 1997)*, pages 95–133. Springer, New York, 2000b.

eScholarship
Access Repository
Collection of Biostatistics
Research Archive

O.M. Stitelman and M.J. van der Laan. Collaborative targeted maximum likelihood for time to event data. Technical Report 260, Division of Biostatistics, University of California, Berkeley, 2010.

M.J. van der Laan and J.M. Robins. *Unified methods for censored longitudinal data and causality*. Springer, New York, 2003.

M.J. van der Laan and D. Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1), 2006.

M.J. van der Laan, E. Polley, and A. Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(25), 2007. ISSN 1.

M.J. van der Laan, S. Rose, and S. Gruber. Readings in targeted maximum likelihood estimation. Technical report 254, Division of Biostatistics, University of California, Berkeley, Sept 2009.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press