Published in final edited form as:

J Am Stat Assoc. 2012 April 1; 107(498): 513–517. doi:10.1080/01621459.2012.665197.

Discussion of Evaluation of Viable Dynamic Treatment Regimes in a Sequentially Randomized Trial of Advanced Prostate Cancer, by Wang et al. 2012

Paul Chaffee and Mark van der Laan

1. INTRODUCTION

We would like to congratulate the authors of this article on a thoughtful and thorough design and data analysis. Their analysis represents a good example of how complex longitudinal data from observational studies or randomized trials should be analyzed, and reflects care and thought. The authors target questions of interest using sophisticated causal inference tools instead of allowing available software and common practice to dictate the analysis.

1.1 Background

The authors analyze data from a sequentially randomized controlled trial in which the longitudinal data structure on one patient can be coded as $O = (L(1), A(1), L(2), A(2), L(3), A(3), L(4), A(4), Y = L(5))$, where $L(1)$ denotes baseline covariates, $L(j)$, $j = 2, 3, 4$, encodes both efficacy and toxicity measurements, $A(j) = (A_1(j), A_2(j))$, $A_1(j)$ indicates a treatment category among specified options, $A_2(j)$ is a drop-out indicator, and Y is a particular final outcome such as a treatment efficacy score measured at 32 weeks. Let P_0 denote the true probability distribution of O . If the person drops-out, then the subsequent L -process, including Y , is defined deterministically in some arbitrary manner. Treatment decisions $A_1(j)$ were assigned either deterministically in response to the observed history of the patient, or randomly to three categories with probability $1/3$ in cases where this type of experimentation made sense from a medical point of view. Certain patients developed a history that resulted in a medically sensible deterministic switch of $A_1(j)$ to an alternative therapy (coded as OFF) prescribed by the treating physician, even though this deterministic switch was not a priori planned or foreseen in the protocol of the trial. The authors dealt with this by defining this switch as another natural treatment decision, as if planned by design, that can be incorporated in the targeted treatment rules of interest, instead of viewing it as right-censoring and thereby targeting parameters that cannot even be identified.

The authors define the quantity of interest as the mean of outcome Y (measured at 32 weeks after randomization) under a dynamic intervention d applied to the “intervention” nodes $A(j)$, denoted by $E_0 Y_d$. This intervention is chosen so that it fully respects the designed deterministic treatment switches as well as the unplanned deterministic switch to OFF, and the intervention enforces no drop-out till the end (i.e., $A_2(j) = 0$ for $j = 1, 2, 3, 4$). The dynamic intervention d is indexed by an assigned first line therapy for $A_1(1)$ among four possible choices, and a second line therapy, among the remaining three, to which the patient will be switched if the first line therapy fails and switching to this second line therapy does not violate the deterministic switches in the data generating experiment. If the patient experiences two consecutive successful responses on the same therapy, then the intervention stops. The intervention is also stopped if the patient experiences an event that enforces the treatment to be switched to OFF. Twelve such dynamic rules are considered, indexed by four choices of first line therapy and three choices for the second line therapy. The authors estimate the mean $E_0 Y_d$ for each of these twelve rules d and construct simultaneous confidence bands.

The dynamic rules are chosen so that this quantity of interest, $E_0 Y_d$, can actually be identified from the data generating experiment: i.e., the positivity assumption $P_0(A(j) = d_j | \bar{L}(j)) > 0$ a.e. holds, where we denote the history of all variables right before $A(j)$ —the “parents” of $A(j)$ —with $\bar{L}(j)$ and we use the notation $\bar{L}(j) = (L(1), L(2), \dots, L(j))$. A natural by-product of selecting such realistic rules is that these dynamic treatments are of practical and scientific interest—a point effectively argued by the authors. Therefore, we can represent $E_0 Y_d$ as $\Psi(P_0)$ for a specified mapping $\Psi: \mathcal{M} \rightarrow \mathbb{R}$ that maps a probability distribution of O in the statistical model $\mathcal{M}$ into a real number. We remind the reader that a statistical model is an assumed collection of possible probability distributions of the data structure O and thus represents all statistical knowledge about the experiment. The pure statistical estimation problem is now defined as estimation of $\Psi(P_0)$ based on n i.i.d. copies of $O \sim P_0$, with the knowledge that $P_0 \in \mathcal{M}$.

The authors estimate $E_0 Y_d$ with a stabilized inverse probability of censoring weighted (IPCW) estimator, defined as the inverse probability of treatment and censoring weighted empirical mean of the outcome over all subjects who followed the rule d . Specifically, the stabilized IPCW-estimator of $\psi_0 = \Psi(P_0)$ is defined as

$$\psi_n^{IPCW} \equiv \frac{\sum_i Y_i I(A_i = \bar{d}(L_i)) / g_0(A_i | L_i)}{\sum_i I(A_i = \bar{d}(L_i)) / g_0(A_i | L_i)},$$

where $A_i = (A_i(1), \dots, A_i(4))$, $L_i = (L_i(1), \dots, L_i(5))$, $g_0(A_i | L_i) \equiv \prod_j P(A_i(j) | \bar{L}_i(j-1))$, and for notational convenience we use $A = \bar{L}(L)$ to denote the set of relations ($A(1) = d_1(L(1))$, $A(2) = d_2(Pa(A(2)))$, $\dots$, $A(4) = d_4(Pa(A(4)))$). Thus g_0 denotes the conditional probability, given covariates, of following a specified treatment through the entire trial and not being censored. We use subscript n to denote sample-based estimators.

Since, ignoring drop-out for now, the treatment mechanism is known in a sequentially randomized trial, such estimators are guaranteed to be consistent and asymptotically linear (and thereby asymptotically normally distributed), and they are also quite stable since the positivity assumption is practically satisfied. By estimating the treatment mechanism and incorporating covariates that are predictive of the outcome, thus ignoring the knowledge that the treatment mechanism is known, such an IPCW-estimator can be made highly efficient, or even fully efficient. The authors indeed use an estimator of g_0 for this reason. If drop-out would have been quite severe and informative, then a typical IPCW estimator may be quite inefficient, biased, and unstable (see e.g., Stitelman et al. (2011) and van der Laan and Rose (2011)). Given that drop-out was quite minimal in this particular trial, the selected estimation procedure, which is easy to understand and implement, appears to represent a reasonable choice.

In our previous discussion of the first analysis of this trial (Bembom and van der Laan 2007) we were given some room to provide constructive suggestions and demonstrated an analysis of realistic dynamic treatment rules with the available data, but we find the current analysis thorough. At that time, some of the data compiled for analysis in the current paper was not available. The new, augmented data set led the authors to re-define the set of outcomes, and to propose target parameters indexed by a set of treatment rules different from those based on the earlier data.

1.2 Substitution Estimators

We define the parameter of interest as a mapping $\Psi: \mathcal{M} \rightarrow \mathbb{R}$ from the statistical model $\mathcal{M}$ to the parameter space, where for simplicity we consider scalar parameters. A substitution

estimator of Ψ is a plug-in estimator $\Psi(\mathbf{P}_n)$ with $\mathbf{P}_n \in \mathcal{M}$ being an estimator of the true P_0 . In our research over the last decade we have focused on the development of double robust locally efficient *substitution* estimators based on the targeted minimum loss-based learning framework proposed in van der Laan and Rubin (2006), building on a rich literature in semiparametric models. As substitution estimators, these estimators not only are asymptotically optimized, but also fully respect the global constraints implied by the statistical model and the target parameter, which makes them robust under sparsity. We refer to van der Laan and Rose (2011) for a comprehensive book on targeted learning, including contributions of many authors, documenting the motivation, developments, and many of the applications of targeted minimum loss-based estimation (TMLE) to complex data structures. We would like to use the opportunity of this discussion to present a double robust, locally efficient substitution estimator for causal effects of multiple time point interventions, recently presented in van der Laan and Gruber (2011), which represents a particular TMLE inspired by innovative advances in general double robust estimation in Bang and Robins (2005). TMLE provides an alternative to IPCW estimation that is generally more efficient, less biased (TMLE is double robust), more finite sample-robust under sparsity, and inherits the virtues of substitution estimators mentioned above. Without a re-analysis of the data set, it is difficult to judge the potential benefit of this procedure for this particular study, but certainly if bias due to informative drop-out had been a serious concern it would have been worthwhile to apply this procedure.

In this discussion we will also demonstrate a particular TMLE, proposed in van der Laan and Rubin (2006), chosen to guarantee to outperform the IPCW estimator, even under severe misspecification of the time-dependent covariate and outcome regressions. Though not discussed here, we refer to recent original work by Rotnitzky et al. (2012) who go beyond this by presenting a general method for constructing a double robust, locally efficient estimator that is guaranteed to beat any user-supplied estimator (not just this IPCW estimator). This has nice applications in the sequentially randomized trial analyzed by the authors here. Subsequently, we applied Rotnitzky's idea to the TMLE framework to construct TMLEs that are guaranteed to outperform any user supplied estimator (van der Laan and Gruber 2012). We refrain from discussing these important relevant advances in this discussion, but wish to highlight this important work by Rotnitzky and co-authors, and the resulting enhanced property of TMLE. In this discussion we will also present a small scale simulation aiming to imitate a longitudinal data structure of a two time-point sequentially randomized trial without right-censoring, and we conclude with a few remarks. We will start with some background that motivates the development of TMLE.

2. WHY NOT NPMLE?

The density P_0 of the data-generating distribution factorizes as $P_0 = Q_0 g_0$, where $g_0 = \prod_j g_0(A(j) | Pa(A(j)))$ represents the known treatment and unknown right-censoring mechanism, while $Q_0 = \prod_{j=1}^5 P_0(L(j) | Pa(L(j)))$ represents the unspecified relevant factor of the likelihood. Note that $g_0 = g(P_0)$ and $Q_0 = Q(P_0)$ are themselves parameters of the data generating distribution. The statistical model $\mathcal{M}$ for P_0 consists of all such distributions of O . As a result of sequential randomization, and assuming that right-censoring also satisfies the sequential randomization and positivity assumptions, the distribution of the counterfactual $L_d = (L(1), L_d(2), L_d(3), L_d(4), L_d(5) = Y_d)$ under dynamic intervention d is given by the so-called G-computation formula

$$Q_{0,d}(l) = \prod_j P_0(L(j)=l(j) | \bar{L}(j-1)=\bar{l}(j-1), \bar{A}(j-1)=d(\bar{L}(j-1))).$$

In particular, $E_0 Y_d$ is defined as the mean of the marginal distribution under $Q_{0,d}$ of component Y_d . This defines now EY_d as a statistical target parameter $\Psi(P)$ of $P \in \mathcal{M}$ which depends on P only through $Q = Q(P)$. If $L(j)$ are discrete-valued, then we can estimate each conditional distribution of $L(j)$, given its parents $Pa(L(j))$, with empirical proportions, resulting in a pure empirical estimator, Q_n , of the true $Q_0 = Q(P_0)$. The corresponding plug-in estimator $\Psi(Q_n)$ would be the NPMLE. The problem with this estimator is that it is often not defined due to empty strata in finite samples, and, even when defined, it is often too variable to be useful in finite samples.

3. WHY NOT SMOOTHED NPMLE?

In general, one will need to apply smoothing to estimate Q_0 . For example, one could use a k -th nearest neighbor regression estimator for each conditional distribution of a binary variable coding $L(j)$, given its parents. However, in order to guarantee that these nearest neighbor regression estimators are consistent, one needs the number of neighbors in the neighborhood over which one smoothes to go to infinity. In a covariate space of dimension k this means that nh^k has to go to infinity, and thereby that the smoothing width h converges to zero slower than $n^{-1/k}$. If one averages outcomes over a neighborhood of width h , then the bias will be $O(h)$. As a consequence, if $k \geq 2$, this smoothed NPMLE will have a bias $O(n^{-1/k})$ that cannot be neglected for statistical inference: confidence intervals ignoring this bias will have asymptotically zero coverage, and p -values cannot be trusted either. Using parametric regression models represents an extreme form of smoothing that is guaranteed to result in bias that will not even converge to zero as sample size goes to infinity, making confidence intervals particularly meaningless.

4. SUPER LEARNING: BETTER, BUT NOT ENOUGH

The typical bias of an estimator of Q_0 that converges at a particular rate to Q_0 will be of the same order as that rate. By constructing an adaptive estimator of Q_0 that is able to adapt to underlying smoothness or structure of Q_0 , one can construct an estimator that might achieve a better rate of convergence than the typical minimax rate of convergence implied by the statistical model $\{Q(P): P \in \mathcal{M}\}$ for Q_0 . Such an estimator will therefore also have less bias. For that purpose, one can use super learning (van der Laan and Dudoit 2003; van der Laan et al. 2007; van der Laan and Rose 2011), which involves using cross-validation to select among a family of candidate estimators of Q_0 that are indexed by different fine-tuning parameters and approximation strategies. For example, the library of the super learner could include kernel regression estimators indexed by different bandwidths and different degrees of orthogonality of the kernels, so that the best choice among all these kernel estimators will achieve the minimax rate for the model that assumes that the true underlying smoothness of Q_0 is known. Due to the oracle inequality for the cross-validation selector, the super learner will also achieve this adaptive minimax rate of convergence. So super learning results in less biased estimators of $\psi_0 \equiv \Psi(Q_0)$ relative to using a nearest neighbor type of estimator or standard kernel regression estimator. However, the best rate of convergence one can realistically hope for is still worse than $1/\sqrt{n}$ (i.e., the rate achieved for a correctly specified parametric model), so that the bias for ψ_0 implied by the super learner fit will still be non-negligible and result in a plug-in estimator that is *not* asymptotically linear.

Additional bias reduction is thus essential, but this bias reduction should be fully targeted towards the target parameter. That is, the estimator Q_n does not need to solve all score equations as does the NPMLE, but it needs to solve certain target parameter-specific score equations, such as the efficient influence curve/efficient score for the target parameter. For that purpose, we combine super learning with targeted minimum loss-based estimation, resulting in the targeted learning approach (van der Laan and Rubin 2006; van der Laan and Rose 2011). This step will need to involve an estimator of the true treatment and censoring mechanism g_0 , and, if g_0 is unknown, super learning is again recommended to obtain maximal bias reduction for this step. For this purpose it is useful to note that g_0 can be factorized as a product of treatment probabilities and censoring probabilities. If treatment $A_1(j)$ at time j has more than two possible values, then one could code $A_1(j)$ in terms of a few binaries, and still estimate the conditional probability of $A(j)$ with corresponding logistic regressions, and corresponding adaptive regression estimators for binary outcomes.

5. TARGETED MINIMUM LOSS BASED ESTIMATION

5.1 Efficient Influence Curve

We remind the reader that an efficient influence curve $D^*(P)$ at P is the canonical gradient of the pathwise derivative of $\Psi: \mathcal{M} \rightarrow \mathbb{R}$ along a richly selected class of parametric submodels through P (see, e.g., Bickel et al. 1997; van der Vaart 1998; van der Laan and Robins 2003; van der Laan and Rose 2011). A regular asymptotically linear estimator has an influence curve that equals one of the gradients of the pathwise derivative, which explains why the canonical gradient is also called the efficient influence curve. In particular, the asymptotic variance of a regular asymptotically linear estimator is larger than or equal to the variance of the efficient influence curve. An estimator ψ_n of ψ_0 is asymptotically efficient at P_0 among the class of regular estimators if and only if it is asymptotically linear at P_0 with influence curve equal to the efficient influence curve: i.e. $\psi_n - \psi_0 = 1/n \sum_{i=1}^n D^*(P_0)(O_i) + o_p(1/\sqrt{n})$. A minimal condition for a substitution estimator $\Psi(\mathbf{P}_n)$ to be asymptotically efficient is that it solves $0 = \sum_{i=1}^n D^*(\mathbf{P}_n)(O_i)$ (or at least up till an $o_p(1/\sqrt{n})$ -term). This equation for $\mathbf{P}_n$ does not define $\mathbf{P}_n$ (it is just one equation for an infinite dimensional $\mathbf{P}_n$), but provides an important and necessary characteristic of $\mathbf{P}_n$ and makes the bias of $\Psi(\mathbf{P}_n)$ second order. The efficient influence curve at P depends on a relevant part of P , $Q(P)$, and a purely nuisance part, $g(P)$, and so we also use the notation $D^*(P) = D^*(Q(P), g(P))$.

5.2 TMLE Algorithm

The first step in TMLE is to represent $\Psi(P)$ as a mapping of a part of P , say $\Psi(Q)$ with $Q = Q(P)$. One then obtains an initial estimate Q_n of Q_0 and g_n of g_0 , and subsequently one constructs a targeted estimator Q_n^*, g_n^* , obtained by iteratively minimizing an empirical risk w.r.t an appropriate loss function along a cleverly chosen submodel through Q_n and g_n so that the final update (Q_n^*, g_n^*) of the initial estimator, (Q_n, g_n) , solves the efficient influence curve score equation $P_n D^*(Q_n^*, g_n^*) = 0$. We use super-script * to indicate that these estimators are targeted. Here we used the notation $P_n f = 1/n \sum_{i=1}^n f(O_i)$. The iterative updating of g_n is not necessary (in which case $g_n^* = g_n$), but can provide additional gains in the TMLE as discussed below. The corresponding TMLE of ψ_0 is simply the plug-in estimator $\psi_n^* = \Psi(Q_n^*)$. Under regularity conditions, ψ_n^* is a double robust, locally efficient asymptotically linear estimator: that is, $\Psi(Q_n^*)$ is consistent and asymptotically linear if either g_n^* or Q_n^* is consistent, and it is efficient if both are consistent.

One could construct a targeted estimator of the relevant factor of the likelihood as in van der Laan (2010); Stitelman et al. (2011); van der Laan and Rose (2011), involving coding $L(j)$ in

terms of binaries, fitting the conditional distributions of these binaries with data adaptive logistic regression, and targeting this fit with a subsequent targeted maximum likelihood update step involving univariate logistic regression with a clever covariate using the initial fit as offset.

Alternatively, we can note that

$$\begin{aligned}\bar{Q}_4^d(\bar{l}(4)) &= E(Y|\bar{A}(4)=d(\bar{l}(4)), \bar{L}(4)=\bar{l}(4)) \\ &= E(Y_d|\bar{L}_d(4)=\bar{l}(4)) \\ \bar{Q}_3^d(\bar{l}(3)) &= E(\bar{Q}_4^d(\bar{L}(4))|\bar{A}(3)=d(\bar{l}(3)), \bar{L}(3)=\bar{l}(3)) \\ &= E(Y_d|\bar{L}_d(3)=\bar{l}(3)),\end{aligned}$$

and so on, so that

$$\bar{Q}_0^d = E(\bar{Q}_1^d(L(1))) = EY_d.$$

In this manner, it follows that EY_d is a function Ψ of a vector of iteratively defined conditional means $\bar{Q}^d = (\bar{Q}_4^d, \dots, \bar{Q}_0^d)$. Suppose that Y is binary. The initial regression $\bar{Q}_4^d$ can be estimated with logistic regression. The targeting step involves adding a clever covariate $I(\bar{A} = d(\bar{L}))/g$ whose coefficient can be fitted with univariate logistic regression using the initial estimator as offset. This targeted estimator $\bar{Q}_{4,n}^{d,*}$ of $\bar{Q}_4^d$ is now used as an outcome in $(0,1)$ in the next *logistic* regression of $\bar{Q}_4^d$ onto $(3) = d(\bar{L}(3)), \bar{L}(3)$, and again, the resulting fit is targeted by fluctuating the fit with a univariate logistic regression model using a clever covariate $I(\bar{A}(3)=d(\bar{L}(3)))/\prod_{j=1}^3 g_j$. This sequential regression procedure finally ends by taking the empirical mean of the estimator of $\bar{Q}_1^d$ over the baseline covariates $L(1)$, which is the final targeted estimate of EY_d . The same estimator applies to continuous $Y \in (0, 1)$, and thereby for bounded continuous outcomes. For details about this TMLE $\psi_n^* = \Psi(\bar{Q}_n^{d,*})$ obtained by plugging in a targeted estimator $\bar{Q}_n^{d,*}$ of these iteratively defined conditional means $\bar{Q}_0^d$, we refer to van der Laan and Gruber (2011). In the latter article we also present some simulations demonstrating its practical performance.

6. STATISTICAL PROPERTIES

If g_0 is estimated consistently, then, under appropriate regularity conditions, the latter TMLE of ψ_0 will be consistent and asymptotically linear, even if all the conditional regressions for $\bar{Q}_j^d, j = 4, 3, 2, 1$ are inconsistent. In addition, if $\bar{Q}^d$ is estimated consistently, then this TMLE is asymptotically efficient. That is, as with all TMLEs, this TMLE is a double robust and locally efficient substitution estimator. It respects the global constraints of the statistical model and target parameter by using logistic regressions (see Gruber and van der Laan 2010; van der Laan and Gruber 2011). The asymptotic variance of ψ_n^* can be estimated with $1/n^2 \sum_{i=1}^n \{D^*(\bar{Q}_n^*, g_n)\}^2(O_i)$, and statistical inference proceeds accordingly.

7. TMLE THAT IS GUARANTEED TO BEAT IPCW

The efficient influence curve can also be represented as $D^*(P) = D_{IPCW}(P) - D_{CAR}(P)$, where $D_{IPCW}(P) = \mathbb{I}(A = d(L))(Y - \psi)/g$ is the estimating function solved by the stabilized IPCW estimator, while D_{CAR} is a score of g (see, e.g. Robins and Rotnitzky 1992; van der Laan and Robins 2003). As proposed in van der Laan and Rubin (2006), given a current estimator $\bar{Q}_n^k, g_n^k$ we can construct a submodel through g_n^k with score at zero-fluctuation equal to $D_{CAR}(\bar{Q}_n^k, g_n^k)$. By coding $A(j)$ in terms of binaries, this submodel can be chosen to correspond with fluctuating each of these binary conditional distributions with a univariate logistic regression in a cleverly chosen covariate, using the current estimator as offset. By also iteratively updating g_n^k in the TMLE algorithm described above along this clever submodel, we obtain a TMLE Q_n^*, g_n^* that solves both the efficient influence curve equation $P_n D^*(Q_n^*, g_n^*) = 0$, and the score $P_n D_{CAR}(Q_n^*, g_n^*) = 0$ for the censoring mechanism. As a consequence, it also solves $P_n D_{IPCW}(Q_n^*, g_n^*) = 0$ so that ψ_n^* is now also a stabilized IPCW estimator, and is guaranteed to be at least as efficient as the stabilized IPCW estimator used in the study.

8. A SIMULATED EXAMPLE

We simulated data that mimics a simplified two time-point sequential randomized controlled trial with structure for a single observation $O = (L(1), A(1), L(2), A(2), L(3) = Y)$. Each $A(j)$ was binary and completely randomized, and each $L(j)$ (including $L(3)$) was a single, continuous random variable in the interval $[0, 1]$. We did not incorporate drop-out in this data-generating experiment.

We chose as parameters of interest for this data the four possible treatment history-specific counterfactual parameters EY where $Y \in \{(0, 0), (0, 1), (1, 0), (1, 1)\}$. In words, these are the mean population outcomes under the specified sequential interventions on the variables $(A(1), A(2))$. We report results for one of them ($EY_{0,0}$) under two different data-generating distributions. In the first distribution $L(j)$ is a particular function of $(j-1)$ and $\bar{L}(j-1)$. The second incorporates similar relationships, but the influence of $Pa(L(j))$ on $L(j)$ for each j is very weak (see Appendix).

The true values of these parameters of the simulated data distribution are, of course, not of interest, since they are entirely fabricated. The results of interest are the relative performances of the estimators in estimating these values.

The sequential regressions required for the TMLE were estimated with main terms logistic regressions, and thus possibly misspecified for $\bar{Q}_j, j < 3$. $P(A(j) = a | Pa(A(j)))$ for use in the IPCW estimator was also estimated using main terms logistic regression, and included all elements of $Pa(A(j))$, which is a consistent estimator of the true conditional treatment probabilities. We found the TMLE to be quite insensitive to the choice of g incorporated (i.e., g_n vs. g_0), and we thus only include results using g_0 .

We simulated data at three samples sizes, 50, 100, 500, but include results only for $n = 100$ since the relative performances of the three estimators were quantitatively similar at all samples sizes. We report estimated bias, variance and relative MSE, the latter defined as $MSE(\cdot)/MSE(\psi_{n,g_0}^{IPCW})$.

A more thorough simulation study would first evaluate the relative variance of the efficient influence curve and the influence curve of the stabilized IPCW estimator under a variety of settings, so that one knows how much information has been neglected by the IPCW

estimator. One can then also investigate the relative gain of an estimator such as TMLE, which takes full advantage of the asymptotic efficiency bound.

9. CONCLUDING REMARKS

Though randomized trials make the IPCW estimator a valid estimator that is relatively efficient and unbiased in the absence of censoring, knowledge of the treatment mechanism can be utilized by constructing locally efficient substitution estimators of the target parameter that are guaranteed asymptotically linear and more efficient than (e.g.) the standard IPCW estimator, even under severe misspecification of the relevant factor of the likelihood. The potential gain in efficiency and bias of these estimators relative to current practice in actual clinical trials that are subject to informative censoring remains an important area of research. In addition, the utilization of machine (super) learning to maximize efficiency and minimize bias due to informative drop-out represents another potential improvement in the analysis of randomized controlled trials. Obviously, such estimators still need to be specified a priori. Asymptotic theory teaches us that adaptive estimation is essential for obtaining an asymptotically linear and thereby consistent and normally distributed estimator of the target parameter in semiparametric models where we lack knowledge of the parametric models involved (see main Theorem in Zheng and van der Laan 2010 and chap. 27 in van der Laan and Rose 2011 by the same authors). There is an evident need for adaptive estimation in typical randomized controlled trials that are subject to drop-out or other forms of missingness that are only partially understood.

References

- Bang H, Robins J. Doubly Robust Estimation in Missing Data and Causal Inference Models. *Biometrics*. 2005; 61:962–972. [PubMed: 16401269]
- Bembom O, van der Laan M. Comment: Statistical Methods for Analyzing Sequentially Randomized Trials. *Journal of the National Cancer Institute*. 2007; 99 (21):1577–1582. [PubMed: 17971533]
- Bickel, P.; Klaassen, C.; Ritov, Y.; Wellner, J. Efficient and Adaptive Estimation for Semiparametric Models. Springer-Verlag; 1997.
- Gruber S, van der Laan M. A Targeted Maximum Likelihood Estimator of a Causal Effect on a Bounded Continuous Outcome. *International Journal of Biostatistics*. 2010; 6:article 26. www.bepress.com/ijb/vol6/iss1/26. [PubMed: 21731529]
- Robins, J.; Rotnitzky, A. AIDS epidemiology. Basel: Birkhäuser; 1992. Recovery of Information and Adjustment for Dependent Censoring Using Surrogate Markers.
- Rotnitzky A, Lei Q, Sued M, Robins J. Enhanced Double-Robust Locally Efficient Estimation in Regression Models for Longitudinal Studies. 2012 To appear in *Biometrika*.
- Stitelman, O.; Gruttola, VD.; van der Laan, M. Technical Report. Vol. 281. Division of Biostatistics, University of California; Berkeley: 2011. A General Implementation of TMLE for Longitudinal Data Applied to Causal Inference in Survival Analysis. <http://www.bepress.com/ucbbiostat/paper281>
- van der Laan M. Targeted Maximum Likelihood Based Causal Inference: Part I. *International Journal of Biostatistics*. 2010; 6(2):Article 2. [PubMed: 21969992]
- van der Laan, M.; Dudoit, S. Technical Report. Vol. 130. Division of Biostatistics, University of California; Berkeley: 2003. Unified Cross-Validation Methodology for Selection Among Estimators and a General Cross-Validated Adaptive Epsilon-net Estimator: Finite Sample Oracle Inequalities and Examples.
- van der Laan, M.; Gruber, S. Technical Report. Division of Biostatistics, University of California; Berkeley: 2011. Targeted Minimum Loss Based Estimation of an Intervention Specific Mean Outcome; p. 290 To appear in *International Journal of Biostatistics*
- van der Laan M, Polley E, Hubbard A. Super Learner. *Statistical Applications in Genetics and Molecular Biology*. 2007; 6(1):Article 25.

- van der Laan, M.; Robins, J. Unified Methods for Censored Longitudinal Data and Causality. Berlin Heidelberg New York: Springer; 2003.
- van der Laan, M.; Rose, S. Targeted Learning: Causal Inference for Observational and Experimental Data. Berlin Heidelberg New York: Springer; 2011.
- van der Laan M, Rubin DB. Targeted Maximum Likelihood Learning. International Journal of Biostatistics. 2006; 2(1):Article 11.
- van der Vaart, A. Asymptotic Statistics. New York: Cambridge; 1998.
- Zheng, W.; van der Laan, M. Technical Report. Vol. 273. Division of Biostatistics, University of California; Berkeley; 2010. Asymptotic Theory for Cross-validated Targeted Maximum Likelihood Estimation.

APPENDIX: DATA-GENERATION FUNCTIONS FOR SIMULATIONS

As mentioned, we ran simulations under two different data-generation functions. The influence of $Pa(L(j))$ on $L(j)$ for $j = 2, 3$ was much stronger under the first distribution.

- Data generation under “ P_0 ”:

$$L(1) = \text{expit}(3 * Z_1) \text{ where } Z_1 \sim N(0, 1)$$

$$A(1) \sim \text{Ber}(0.5)$$

$$L(2) = \text{expit}(\text{logit}(L(1)) + 2 * A(1) + Z_2) \text{ where } Z_2 \sim N(0, 0.16)$$

$$A(2) \sim \text{Ber}(0.5)$$

$$Y = \text{expit}(\text{logit}(L(1)) + 1.5 * A(1) + \text{logit}(L(2)) - 0.5 * A(2) + Z_3) \text{ with } Z_3 \sim N(0, 0.16)$$

- Data generation under “ $P_0^\dagger$ ”:

$$L(1) = \text{expit}(3 * Z_4) \text{ where } Z_4 \sim N(0, 1)$$

$$A(1) \sim \text{Ber}(0.5)$$

$$L(2) = \text{expit}(0.01 * \text{logit}(L(1)) + 0.05 * A(1) + Z_5) \text{ where } Z_5 \sim N(0, 16)$$

$$A(2) \sim \text{Ber}(0.5)$$

$$Y = \text{expit}(0.01 * \text{logit}(L(1)) + 0.1 * A(1) + 0.01 * \text{logit}(L(2)) - 0.05 * A(2) + Z_6) \text{ with } Z_6 \sim N(0, 12.25)$$

$$Z_i \perp Z_j \text{ for all } i \neq j.$$

Table 1

Comparison of TMLE and IPCW in estimation of a single treatment history-specific parameter EY at $n = 100$ in simulated data under two different data-generation mechanisms, P_0 and $P_0^\dagger$. Results based on 2000 simulations. IPCW (g_0) refers to the IPCW estimator using the true g_0 for the treatment mechanism; IPCW (g_n) is the IPCW estimator with g_0 estimated from data.

$P_0 (n = 100)$	Bias	Var	Rel MSE
TMLE	-0.0068	0.0018	0.25
IPCW (g_n)	*	0.0040	0.52
IPCW (g_0)	*	0.0077	1.00

$P_0^\dagger (n = 100)$	Bias	Var	Rel MSE
TMLE	*	0.0036	0.6
IPCW (g_n)	*	0.0064	1.0
IPCW (g_0)	*	0.0061	1.0

* indicates $< 10^{-3}$