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

Year 2009 *Paper* 246

Collaborative Targeted Maximum Likelihood Estimation

Mark J. van der Laan* Susan Gruber†

*University of California - Berkeley, laan@berkeley.edu

†UC Berkeley, sgruber65@yahoo.com

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/paper246

Copyright ©2009 by the authors.

Collaborative Targeted Maximum Likelihood Estimation

Mark J. van der Laan and Susan Gruber

Abstract

Collaborative double robust targeted maximum likelihood estimators represent a fundamental further advance over standard targeted maximum likelihood estimators of causal inference and variable importance parameters. The targeted maximum likelihood approach involves fluctuating an initial density estimate, (Q), in order to make a bias/variance tradeoff targeted towards a specific parameter in a semi-parametric model. The fluctuation involves estimation of a nuisance parameter portion of the likelihood, g. TMLE and other double robust estimators have been shown to be consistent and asymptotically normally distributed (CAN) under regularity conditions, when either one of these two factors of the likelihood of the data is correctly specified.

In this article we provide a template for applying collaborative targeted maximum likelihood estimation (C-TMLE) to the estimation of pathwise differentiable parameters in semi-parametric models. The procedure creates a sequence of candidate targeted maximum likelihood estimators based on an initial estimate for Q coupled with a succession of increasingly non-parametric estimates for g. In a departure from current state of the art nuisance parameter estimation, C-TMLE estimates of g are constructed based on a loss function for the relevant factor Q_0, instead of a loss function for the nuisance parameter itself. Likelihood-based cross-validation is used to select the best estimator among all candidate TMLE estimators in this sequence. A penalized-likelihood loss function for Q_0 is suggested when the parameter of interest is borderline-identifiable.

We present theoretical results for "collaborative double robustness," demonstrating that the collaborative targeted maximum likelihood estimator is CAN when Q and g are both mis-specified, providing that g solves a specified score equation

implied by the difference between the Q and the true Q_0.

This marks an improvement over the current definition of double robustness in the estimating equation literature.

We also establish an asymptotic linearity theorem for the C-DR-TMLE of the target parameter, showing that the C-DR-TMLE is more adaptive to the truth, and, as a consequence, can even be super efficient if the first stage density estimator does an excellent job itself with respect to the target parameter.

This research provides a template for targeted efficient and robust loss-based learning of a particular target feature of the probability distribution of the data within large (infinite dimensional) semi-parametric models, while still providing statistical inference in terms of confidence intervals and p-values. This research also breaks with a taboo (e.g., in the propensity score literature in the field of causal inference) on using the relevant part of likelihood to fine-tune the fitting of the nuisance parameter/censoring mechanism/treatment mechanism.

1 Introduction

Researchers acknowledge that questions about our infinite-dimensional, semiparametric world are not well-addressed by parametric models. More sophisticated tools are needed to wrest meaning from data. We can and should develop and utilize methods specifically designed to estimate a relatively small-dimensional precisely specified parameter within such a semiparametric model that is identifiable from the data. The ideal method would be entirely a priori specified, have desirable statistical properties, avoid reliance on ad hoc or arbitrary specifications, and be computationally feasible.

Suppose one observes a sample of independent and identically distributed observations from a particular data generating distribution P_0 in a semiparametric model, and that one is concerned with estimation of a particular pathwise differentiable parameter of the data generating distribution. A parameter should be viewed as a mapping from the semiparametric model to the parameter space (e.g., real line). A parameter mapping is pathwise differentiable at P_0 if it is differentiable along all smooth parametric sub-models through P_0, and its derivative is uniformly bounded as a linear mapping on the Hilbert space of all scores of these parametric submodels. Intuitively, a pathwise differentiable parameter is a parameter which has a finite generalized Cramer-Rao information lower bound, so that in principle, under enough regularity conditions, it is possible to construct an estimator which behaves like a sample mean of i.i.d. random variables. Due to the curse of dimensionality implied by the infinite dimension of semi-parametric models, standard (non-parametric) maximum likelihood estimation is often ill defined or breaks down due to overfitting, while, on the other hand, regularized sieve-based maximum likelihood estimation results in overly biased plug-in estimators of the target parameter of interest.

The latter is due to the fact that such likelihood based estimators seek and achieve a bias-variance trade-off that is optimal for the density of the distribution of the data itself. Since the variance of an optimally smoothed density estimator is typically much larger than the variance of a smooth (pathwise-differentiable) parameter of the density estimator, the substitution estimators are often too biased relative to their variance. That is, substitution estimators based on density estimators involving optimal (e.g., likelihood-based) bias-variance trade-off (for the whole density) are not *targeted* towards the parameter of interest.

Motivated by this problem with the bias-variance trade-off of maximum likelihood estimation in semiparametric models, while still wanting to preserve the log-likelihood as the principle criterion in estimation, in van der Laan and

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

Rubin (2006) we introduced and developed a targeted maximum likelihood estimator of the parameter of interest.

The targeted maximum likelihood estimator of the distribution of the data is obtained by fluctuating an initial estimator of the relevant part of the data generating distribution with a parametric fluctuation model whose score at the initial estimator (i.e., at zero fluctuation) equals or includes the efficient influence curve of the parameter of interest, and estimating the fluctuation parameter (i.e., amount of fluctuation) with standard parametric maximum likelihood, treating the initial estimator as offset. Iteration of this targeted maximum likelihood modification step results in a so called k-th step targeted maximum likelihood estimator, and its limit in k solves the actual efficient influence curve equation defined by setting the empirical mean of the efficient influence curve equal to zero. The latter estimator we called the targeted maximum likelihood estimator, which also results in a corresponding plug-in targeted maximum likelihood estimator of the parameter of interest by applying the parameter mapping to the targeted maximum likelihood estimator.

This targeted maximum likelihood step using the fluctuation model removes bias of the initial estimator with respect to (w.r.t.) the target parameter, while increasing the variance of the estimator till the level of the semi-parametric information bound, thereby resulting in a consistent, asymptotically linear, and semi-parametric (locally) efficient estimator.

Although in a variety of applications the fluctuation model is known, e.g., randomized controlled trials with known treatment assignment and missingness mechanism, the fluctuation model typically depends on an unknown nuisance parameter, which then needs to be estimated as well. In censored data models satisfying the so called coarsening at random (CAR) assumption this nuisance parameter typically represents the censoring mechanism, and the density of the data factors in the relevant part of the density and the censoring mechanism density (e.g., Heitjan and Rubin (1991), Jacobsen and Keiding (1995), Gill et al. (1997)).

In this case, the bias reduction obtained at the targeted maximum likelihood step depends on how and how well we estimate the nuisance parameter. Specifically, the targeted maximum likelihood estimator is a so called double robust locally efficient estimator in censored data models (including causal inference models with the full data representing a collection of treatment regimen specific counterfactuals) in which the censoring mechanism satisfies the coarsening at random assumption. This means that, under regularity conditions, it is consistent and asymptotically linear if either the initial estimator is consistent or the nuisance parameter is consistent, and it is efficient in the semiparametric model if the initial estimator is consistent. Another ap-

Collection of Biostatistics
Research Archive

proach for double robust locally efficient estimation is the estimating equation methodology (see, van der Laan and Robins (2003), and review below).

An outstanding open problem that obstructs the robust practical application of double robust estimators (in particular, in nonparametric censored data or causal inference models) is the selection of a sensible model or estimator of the nuisance parameter: this is particularly true when the efficient influence curve estimating equation involves inverse probability of censoring or treatment weighting, due to the enormous sensitivity of the estimator of the parameter of interest to the estimator of the nuisance parameter. A relevant recent discussion of these issues is found in Kang and Schafer (2007a), Ridgeway and McCaffrey (2007), Robins et al. (2007), Tan (2007), Tsiatis and Davidian (2007), Kang and Schafer (2007b).

Given an initial estimator, we are concerned with constructing an estimator of the nuisance parameter, that results in a better bias-variance trade-off (i.e. better MSE) for the resulting targeted maximum likelihood estimator of the target parameter than current practice. In this article we introduce a new strategy for nuisance parameter estimator selection for targeted maximum likelihood estimators that addresses this challenge by using the log-likelihood of the targeted maximum likelihood estimator (of the relevant density) indexed by the nuisance parameter estimator as the principal selection criterion. The nuisance parameter estimators needed for the targeting step are selected based on the relevant log-likelihood loss function of the resulting targeted maximum likelihood estimator, not on a loss function for the nuisance parameter itself. This approach takes into account the established fit of the initial estimator, and that the resulting estimator of the target parameter is indeed based on the relevant part of the likelihood.

Recognizing that the selected estimator of the nuisance parameter is very much a function of the goodness of fit of the initial estimator led to the development of a new theory of collaborative double robust estimation. The asymptotic linearity theory presented below involves characterizing a true minimal nuisance parameter indexed by the initial estimator limit, $g_0(Q)$, that results in an efficient influence curve that is unbiased for the target parameter. This defines the collaborative double robustness of the efficient influence curve. Given a nested sequence of increasingly non-parametric estimators, g_δ, there is a $g_{\delta_{min}}$ corresponding to $g_0(Q)$ which makes the efficient influence curve unbiased for the target parameter. In addition, all estimates of g in the sequence that are more nonparametric than the estimator indexed by δ_{min}, i.e. $\delta > \delta_{min}$, also make the efficient influence curve unbiased for the target parameter. These results allow us to establish asymptotic linearity of the collaborative double

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

robust targeted maximum likelihood estimator, under appropriate regularity conditions.

The theory is fascinating, and results in potentially super efficient estimators, the main intuition, in the context of CAR-censored data models, is that the covariates that enter the treatment mechanism and censoring mechanism estimator (i.e., nuisance parameter estimator) used to define the fluctuation model in the targeted maximum likelihood step should explain the difference between the initial estimator, Q_n, and the true relevant density, Q_0.

Such collaborative double robust estimators involve a variety of choices, including the choice of initial estimator and the choice of collaborative nuisance parameter estimator, but all solve the efficient influence curve equation and all rely on the collaborative nuisance parameter estimator being correctly specified so that the wished unbiasedness of the efficient influence curve is achieved in the limit. We propose using cross-validation w.r.t. a targeted loss function to select among these different collaborative targeted maximum likelihood estimators of the relevant density. In addition, we suggest the square of their influence curve or square of the efficient influence curve as a particularly suitable loss function, corresponding with selection of the estimator with minimal asymptotic variance.

An overview of relevant literature

The construction of efficient estimators of pathwise differentiable parameters in semi-parametric models requires utilizing the so called efficient influence curve, defined as the canonical gradient of the pathwise derivative of the parameter. A fundamental result of the efficiency theory is that a regular estimator is efficient if and only if it is asymptotically linear with influence curve equal to the efficient influence curve. We refer to Bickel et al. (1997), and Andersen et al. (1993). There are two distinct approaches for construction of efficient (or locally efficient) estimators: the estimating equation approach that uses the efficient influence curve as an estimating equation (e.g., one-step estimators based on the Newton-Raphson algorithm in Bickel et al. (1997)), and the targeted MLE that uses the efficient influence curve to define a targeted fluctuation function, and maximizes the likelihood in that targeted direction.

The construction of locally efficient estimators in censored data models in which the censoring mechanism satisfies the so called coarsening at random assumption (Heitjan and Rubin (1991), Jacobsen and Keiding (1995), Gill et al. (1997)) has been a particular focus area. This also includes the theory for locally efficient estimation of causal effects under the sequential randomization assumption (SRA), since the causal inference data structure can be viewed as

Collection of Biostatistics
Research Archive

a missing data structure on the intervention-specific counterfactuals, and SRA implies the coarsening at random assumption on the missingness mechanism, while not implying any restriction on the data generating distribution.

Gill and Robins (2001) present an implicit construction of counterfactuals as a mapping from the observed data distribution, such that the observed data structure augmented with the counterfactuals satisfies the consistency assumption and the SRA. Yu and van der Laan (2002) provide a particular explicit construction of counterfactuals from the observed data structure in terms of quantile-quantile functions, satisfying the consistency assumption and SRA. These results show that, without loss of generality, one can view causal inference as a missing data structure estimation problem. Causal graphs make explicit the real assumptions needed to claim that these counterfactuals are actually the counterfactuals of interest.

Inverse probability of censoring weighted (IPCW) estimators were originally developed to correct for confounding-induced bias in causal effect estimation. Theory for IPCW estimation and augmented locally efficient IPCW-estimator based on estimating functions defined in terms of the orthogonal complement of the nuisance tangent space in CAR-censored data models (including the optimal estimating function implied by efficient influence curve) was originally developed in Robins (1993), Robins and Rotnitzky (1992). Many papers build on this framework (see van der Laan and Robins (2003) for a unified treatment of this estimating equation methodology, and references). In particular, double robust locally efficient augmented IPCW-estimators have been developed (Robins and Rotnitzky (2001b), Robins and Rotnitzky (2001), Robins et al. (2000), Robins (2000a), van der Laan and Robins (2003),Neugebauer and van der Laan (2005), Yu and van der Laan (2003)).

Causal inference for multiple time-point interventions under sequential randomization was first addressed by Robins in the eighties: e.g. Robins (1986), Robins (1989).

The popular propensity score methods to assess causal effects of single time point interventions (e.g., Rosenbaum and Rubin (1983), Sekhon (2008), Rubin (2006)) have no natural generalization to multiple time-point interventions and may be inefficient (and less robust) estimators for single time point interventions, relative to the locally efficient double robust estimators such as the augmented IPCW and the targeted MLE. One crucial ingredient of these proposed methods is propensity score estimation in the absence of any knowledge of the outcomes.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

Structural nested models and marginal structural models for single and multiple time point static treatment regimens were proposed by Robins as well: Robins (1997b), Robins (1997a), Robins (2000b). Many application papers on marginal structural models exist, involving the application of estimating equation methodology (IPCW and DR-IPCW): e.g., Hernan et al. (2000), Robins et al. (2000a), Bryan et al. (2003), Yu and van der Laan (2003). In van der Laan et al. (2005) history adjusted marginal structural models were proposed as a natural extension of marginal structural models, and it was shown that the latter also imply an individualized treatment rule of interest (a so called history adjusted statically optimal treatment regimen): see Petersen et al. (2005) for an application to the "when to switch" question in HIV research.

Murphy et al. (2001) present a nonparametric estimator for a mean under a dynamic treatment in an observational study. Structural nested models for modeling and estimating an optimal dynamic treatment were proposed by Murphy (2003), Robins (2003), Robins (2005a), Robins (2005b). Marginal structural models for user supplied set of dynamic treatment regimens were developed and proposed in van der Laan (2006), van der Laan and Petersen (2007) and, simultaneously and independently, in a technical report authored by Rotnizky and co-workers (2006), and Robins et al. (2008). van der Laan and Petersen (2007) also includes a data analysis application of these models to assess the mean outcome under a rule that switches treatment when CD4-count drops below a cut-off, and the optimal cut-off is estimated as well. Another practical illustration in sequentially randomized trials of these marginal structural models for realistic individualized treatment rules is presented in Bembom and van der Laan (2007).

Unified loss based learning based on cross-validation was developed invan der Laan and Dudoit (2003), including construction of adaptive minimax estimators for infinite dimensional parameters of the full data distribution in CAR-censored data and causal inference models: see also van der Laan et al. (2006), van der Vaart et al. (2006), van der Laan et al. (2004), Dudoit and van der Laan (2005), Keleş et al. (2002), Sinisi and van der Laan (2004).

The oracle results for the cross-validation selector inspired a unified super learning methodology mapping a library of candidate estimators into a weighted combination with optimal cross-validated risk, thereby resulting in an estimator which either achieves the best possible parametric model rate of convergence up till a log-n-factor, or it is asymptotically equivalent with the oracle selected estimator that selects the best set of weights for the given data set. These results rely on the assumption that the loss function is uniformly bounded and that the number of candidates in the library is polynomial in sample size (van der Laan et al. (2007), Polley and van der Laan (2009)).

COBRA
Collection of Biostatistics
Research Archive

The super learning methodology applied to a loss function for the G-computation formula factor, Q_0, in causal inference, or the full-data distribution factor, Q_0, of the observed data distribution in CAR-censored data models, provides substitution estimators of the target parameter ψ_0. However, although these super learners of Q_0 are optimal w.r.t. the dissimilarity with Q_0 implied by the loss function, the corresponding substitution estimators will be overly biased for a smooth parameter mapping Ψ. This is due to the fact that cross-validation makes optimal choices w.r.t. the (global) loss-function specific dissimilarity, but the variance of $\Psi(\hat{Q})$ is of smaller order than the variance of $\hat{Q}$ itself.

van der Laan and Rubin (2006) integrates the loss-based learning of Q_0 into the locally efficient estimation of pathwise differentiable parameters, by enforcing the restriction in the loss-based learning that each candidate estimator of Q_0 needs to be a targeted maximum likelihood estimator (thereby, in particular, enforcing each candidate estimator of Q_0 to solve the efficient influence curve estimating equation). Another way to think about this is that each loss function $L(Q)$ for Q_0 has a corresponding targeted loss function $L(Q^*)$, with Q^* the targeted MLE algorithm applied to initial Q, and we apply the loss-based learning to the latter targeted version of the loss function $L(Q)$. Rubin and van der Laan (2008) propose the square of efficient influence curve as a valid and sensible loss function $L(Q)$ for selection and estimation of Q_0 in models in which g_0 can be estimated consistently, such as in randomized controlled trials.

The implications of this targeted loss based learning are that Q_0 is estimated optimally (maximally adaptive to the true Q_0) w.r.t. the targeted loss function $L(Q^*)$ using the super learning methodology, and due to the targeted MLE step the resulting substitution estimator of ψ_0 is now asymptotically linear as well if the targeted fluctuation function is estimated at a good enough rate (and only requiring adjustment by confounders not yet accounted for by initial estimator: see collaborative targeted MLE): either way, bias reduction will occur as long as the censoring/treatment mechanism is estimated consistently. Targeted MLE have been applied in a variety of estimation problems: Bembom et al. (2008), Bembom et al. (2009) (physical activity), Tuglus and van der Laan (2008) (biomarker analysis), Rosenblum et al. (2009) (AIDS), van der Laan (2008a) (case control studies), Rose and van der Laan (2008) (case control studies), Rose and van der Laan (2009) (matched case control studies), Moore and van der Laan (2009) (causal effect on time till event, allowing for right-censoring), van der Laan (2008b) (adaptive designs, and multiple time point interventions), Moore and van der Laan (2007) (randomized trials

UC BERKELEY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

with binary outcome). We refer to van der Laan et al. (September, 2009) for collective readings on targeted maximum likelihood estimation.

1.1 Advantages of TMLE relative to augmented IPCW estimating function methodology

Even though the augmented IPCW-estimator is also double robust, targeted maximum likelihood estimation has the following important advantages relative to estimating equation methods such as the augmented-IPCW estimator: 1) the TMLE is a substitution estimator and thereby respects global constraints of model such as that one might be estimating a probability in $[0, 1]$ or a (monotone) survival function at a finite set of points, 2) since, given an initial estimator, the targeted MLE step involves maximizing the likelihood along a smooth parametric targeted fluctuation model, it does *not* suffer from *multiple solutions* of a (possibly *non-smooth* in the parameter) estimating equation, 3) the TMLE does *not* require that the efficient influence curve can be represented as an estimating function in the target parameter, and thereby applies to *all* pathwise differentiable parameters 4) it can use the cross-validated log-likelihood (of the targeted maximum likelihood estimator), or any other cross-validated risk of an appropriate loss function for the relevant factor Q_0 of the density of the data, as principle criterion to select among different targeted maximum likelihood estimators indexed by different initial estimators or targeted maximum likelihood steps.

The latter allows fine tuning of initial estimator of Q_0 as well as the fine tuning of the estimation of the unknowns (e.g., censoring/treatment mechanism g_0) of the fluctuation function applied in the targeted maximum likelihood step, thereby utilizing the excellent theoretical and practical properties of the loss-function specific cross-validation selector. In particular, this property results in a collaborative double robust, and possibly super efficient, TMLE, as introduced and studied in this article, thereby adding theoretical and practical properties that go beyond the double robustness and efficiency. In contrast, the augmented-IPCW estimator cannot be evaluated based on a loss function for Q_0 alone: the augmented-IPCW estimator is not a substitution estimator $\Psi(Q_n^*)$ for some Q_n^* of Q_0, as is the TMLE. Instead the augmented-IPCW estimator ψ_n is a certain function of an initial Q_n and g_n, where the performance of g_n is scored based on the orthogonal loglikelihood of g_0, for which a good fit can result in bad fit of ψ_0. In trying to address these shortcomings of the augmented IPCW-estimators we converged to the targeted MLE and, subsequent refinement, the collaborative targeted MLE.

Collection of Biostatistics
Research Archive

1.2 Organization of article

In Section 2 we present a description of the two stage collaborative targeted maximum likelihood methodology, the first stage representing the initial estimator, and the second stage representing the construction of a sequence of targeted maximum likelihood estimators indexed by increasingly nonparametric nuisance parameter estimators, and log-likelihood based cross-validation to select among the TMLEs and thereby select the nuisance parameter estimator. The second stage of the C-DR-TMLE can be viewed as a mapping from an initial estimator of the relevant density into a particular estimator of the nuisance parameter needed in the fluctuation function, and corresponding targeted maximum likelihood estimator using this nuisance parameter estimator in the targeted maximum likelihood step. We also provide the rational for the consistency of this C-DR-TML estimator under the collaborative double robustness assumption, relying on the earlier established oracle property of the log-likelihood-based cross-validation selector, which itself relies on the assumption that the log-likelihood loss function is uniformly bounded.

In Section 3 we define and study collaborative double robustness of the efficient influence curve. In particular, we define true nuisance parameters depending on a choice of relevant density (i.e., limit of initial estimator), which make the efficient influence curve an unbiased function for the target parameter. A collaborative targeted maximum likelihood estimator solves the efficient influence curve equation and relies on the nuisance parameter estimator to consistently estimate this true initial estimator-specific nuisance parameter or more nonparametric nuisance parameter. We also discuss alternative collaborative nuisance parameter estimators that can be used in the targeted MLE or in estimating equation methodology.

In Section 4 we prove an asymptotic linearity theorem for such collaborative double robust estimators, such as the collaborative double robust targeted maximum likelihood estimator, and discuss the conditions and implications of this theorem. In particular, this theorem provides us with influence curve based confidence intervals and tests of null hypotheses. A study of the influence curve teaches us that the C-DR-TMLE can be super efficient.

In Section 5 we consider targeted loss functions that can be used to select among different C-DR-TMLEs indexed by different initial estimators and choices of nuisance parameter estimator. These targeted loss functions can also be used to build the candidate nuisance parameter estimators within a C-DR-TMLE estimator, and thereby to construct the sequence of corresponding candidate targeted maximum likelihood estimators in the collaborative targeted maximum likelihood algorithm. Even though we enforce the use of a

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

log-likelihood-based cross-validation selector to select among these candidate targeted maximum likelihood estimators in the C-DR-TMLE algorithm, we propose a penalized log-likelihood loss function that is more targeted towards the target parameter in the case the target parameter is borderline identifiable. This penalty is particularly important to robustify the estimation procedure in situations in which the variance of the efficient influence curve easily blows up to infinity for certain realization of the nuisance parameter estimator (e.g., close to zero inverse weights).

In section 6 we consider estimation of a causal effect in a marginal structural model, and define the collaborative double robust targeted penalized maximum likelihood estimator of the unknown parameters of the marginal structural model. In section 7 we present a simulation study and data analysis for the C-DR-TMLE of the causal effect $EY(1) - Y(0)$ of a binary treatment A, adjusting for baseline confounders W, based on observing n i.i.d. copies of a time-ordered data structure $(W, A, Y = Y(A))$. A discussion in Section 8 provides a global overview. TMLE as an imputation estimator is described in an appendix.

1.3 An example to keep in mind

Although the methodology is completely general, throughout the paper we ground the discussion by referring to the following example, estimation of the additive causal effect of a binary treatment on an outcome. This example is rich enough to illustrate the ideas and methods, and has been used intensively in the causal inference literature. In this subsection we provide the notation, and objects required to define the C-DR-TMLE.

Let $O = (W, A, Y = Y(A)) \sim P_0$ be an observed missing data structure on full data structure $X = (W, Y(0), Y(1))$ with missingness binary variable $A \in \{0, 1\}$. For concreteness, we consider the case that Y is binary. Suppose the model for P_0 is nonparametric, that the missingness mechanism $g_0(1 \mid X) = P_0(A = 1 \mid X) = P_0(A = 1 \mid W)$ satisfies the coarsening at random assumption, and that our target parameter is the causal additive risk

$$\begin{aligned} \Psi(P_0) = \Psi^F(Q_0) &= E_0 Y(1) - Y(0) \\ &= E_0\{E_0(Y \mid A = 1, W) - E_0(Y \mid A = 0, W)\}, \end{aligned}$$

where $Q_0 = (Q_{01}, Q_{02})$ denotes the marginal distribution of W and conditional distribution of Y, given A, W, respectively. For notational convenience, we will suppress the F from "Full Data Parameter" in Ψ^F. We note that $dP_0(O) = Q_0(O)g_0(A \mid X) = Q_{01}(W)Q_{02}(Y \mid A, W)g_0(A \mid X)$.

Collection of Biostatistics
Research Archive

The efficient influence curve of Ψ at $dP_0 = Q_0 g_0$ is given by

$$D^*(Q_0, g_0) = h_{g_0}(A, W)(Y - Q_0(A, W)) + Q_0(1, W) - Q_0(0, W) - \Psi(Q_0),$$

where $Q_0(A, W) = E_{Q_0}(Y \mid A, W)$, $h_{g_0}(A, W) = A/g_0(1 \mid W) - (1 - A)/g_0(0 \mid X)$. We note that h_{g_0} also plays the role of the clever covariate in the targeted maximum likelihood fluctuation of the conditional distribution of Y, given A, W: $\log Q(\epsilon)/(1 - Q(\epsilon)) = \log Q/(1 - Q) + \epsilon h_{g_0}$.

We also note that an alternative representation of the efficient influence curve is given by the augmented IPCW-representation:

$$\begin{aligned} D^*(Q_0, g_0) &= \left(\frac{A}{g_0(1 \mid X)} - \frac{1 - A}{g_0(0 \mid X)}\right) Y - \Psi(Q_0) \\ &\quad - \left(\frac{A}{g_0(1 \mid X)} - 1\right) Q_0(1, W) + \left(\frac{1 - A}{g_0(0 \mid X)} - 1\right) Q_0(0, W) \\ &= D_{IPCW}(g_0, \psi_0) - D_{CAR}(Q_0, g_0), \end{aligned}$$

where $D_{IPCW}(g_0, \psi_0) = (A/g_0(1) - (1 - A)/g_0(0))Y - \Psi(Q_0)$ is the IPCW-estimating function, and $D_{CAR}(Q_0, g_0)$ is its projection onto T_{CAR} defined as the sub-Hilbert space of $L_0^2(P_0)$ consisting of all functions of (A, W) with conditional mean zero, given W. Here $L_0^2(P_0)$ is the Hilbert space of functions of O endowed with inner product $\langle h_1, h_2 \rangle_{P_0} = E_{P_0} h_1(O) h_2(O)$.

2 Collaborative double robust targeted maximum likelihood estimators

We will describe the proposed collaborative double robust targeted maximum likelihood estimators in the context of censored data models, but the generalization to general semi-parametric models is immediate. We first review targeted maximum likelihood estimation and loss-based cross-validation in order to provide a foundation for the explanation of C-DR-TMLE.

2.1 Targeted MLE in CAR-censored data model

Let $O = \Phi(C, X)$ be a censored data structure on a full data random variable X, where C denotes the censoring variable. We assume coarsening at random so that the observed data structure $O \sim P_0$ has a probability distribution whose density w.r.t an appropriate dominating measure factors as $dP_0(O) = Q_0(O) g_0(O \mid X)$, where Q_0 is the part of the distribution of X that

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

is identifiable, and g_0 denotes the conditional probability distribution of O, given X, which we often refer to as the censoring mechanism. By CAR, we have $g_0(O \mid X) = h(O)$ for some measurable function h. If C is observed itself, then g_0 denotes the conditional distribution of C, given X.

A semiparametric model $\mathcal{M}$ for the probability distribution P_0 of the observed data structure O is implied by a model $\mathcal{Q}$ for the full-data distribution factor Q_0, and a model $\mathcal{G}$ for the censoring mechanism g_0. The conditional distribution of O, given X, is identified by the conditional distribution of C, given X. For notational convenience, we will denote both with g_0. Let $O_1, \ldots, O_n$ be n independent and identically distributed (i.i.d.) observations of the experimental unit O with probability distribution $P_0 \in \mathcal{M}$. Let P_n be the empirical probability distribution of $O_1, \ldots, O_n$ which puts mass $1/n$ on each of the n observations.

Let $\Psi : \mathcal{M} \to \mathbb{R}^d$ be a d-dimensional parameter that is path-wise differentiable at each $P \in \mathcal{M}$ (w.r.t. a class of finite dimensional paths through P) with canonical gradient $D^*(P)$: i.e., for a rich class of parametric submodels $\{P(\delta) : \delta\} \subset \mathcal{M}$ through P at $\delta = 0$ with score $S \in L_0^2(P)$, $L_0^2(P)$ being the Hilbert space of mean zero functions of O endowed with inner product $\langle h_1, h_2 \rangle_P = E h_1 h_2(O)$ (i.e., the covariance operator), we have

$$\frac{d}{d\delta} \Psi(P(\delta))|_{\delta=0} = E_P D^*(P) S.$$

Because $D^*(P)$ is an element of the Hilbert space in $L_0^2(P)$ generated by all scores S of these parametric submodels (the so called tangent space), it is the canonical gradient $D^*(P)$, also called the efficient influence curve at P. Any $D(P)$ such that $E_P D^*(P) S = E_P D(P) S$ for all scores S in the tangent space is called a gradient of the path-wise derivative. Thus the canonical gradient is the unique gradient that is an element of the tangent space. For the sake of illustration, it is assumed that $\Psi(P_{Q,g}) = \Psi^F(Q)$ for some Ψ^F: i..e, the parameter of interest is a parameter of the full data distribution of X. The efficient influence curve $D^*(P)$ at P with $dP = Qg$ will also be denoted with $D^*(Q, g)$.

The Targeted Maximum Likelihood estimator indexed by initial (Q, g): Given any $P \in \mathcal{M}$ with $dP = Qg$, let $\{P(\epsilon) : \epsilon\} \subset \mathcal{M}$ be a submodel with finite dimensional parameter ϵ, dominated by P, through P at $\epsilon = 0$, and whose scores at $\epsilon = 0$ span a finite dimensional space within $L_0^2(P)$ that includes the (components of the) efficient influence curve $D^*(P) = D^*(Q, g)$. Because our parameter of interest is a parameter of Q_0 and the factorization $dP_0 = Q_0 g_0$, it follows that such a fluctuation model can be chosen to only fluctuate Q with a submodel $Q_g(\epsilon) \subset \mathcal{Q}$, where this fluctuation model will be

Collection of Biostatistics
Research Archive

indexed by g. Let $dP(\epsilon) = Q_g(\epsilon)g$ be such a fluctuation model with fluctuation parameter ϵ. In van der Laan and Rubin (2006) we also consider fluctuation models that vary both Q and g.

At a given (Q, g), one can now define a k-th step targeted maximum likelihood version $Q_g^k(P_n)$ of Q_0 as follows. Let $L(Q) = -\log Q$ be the log-likelihood loss. Firstly, let $Q_g^1(P_n) = Q_g(\epsilon_n^1)$, where

$$\epsilon_n^1 = \arg\min_\epsilon P_n L(Q_g(\epsilon)).$$

Here we use the notation $Pf = \int f(o)dP(o)$. In general, $Q_{gn}^k = Q_g^k(P_n) = Q_g^{k-1}(P_n)(\epsilon_n^k)$, where

$$\epsilon_n^k = \arg\min_\epsilon P_n L(Q_g^{k-1}(P_n)(\epsilon)),\ k = 1, \ldots.$$

One iterates this updating till ϵ_n^k equals zero within a user supplied precision. The final update is refered to as the (iterative) targeted maximum likelihood estimator $Q_{gn}^* = Q_g^*(P_n)$, indexed by the initial starting point (Q, g).

The Targeted Maximum Likelihood estimator indexed by initial estimator and estimator of nuisance parameter: The above procedure, applied to an initial estimator Q_n^0, and an estimator g_n of g_0, defines the k-th step targeted maximum likelihood estimator and its limit in k, Q_n^*, as introduced and analyzed in van der Laan and Rubin (2006). By definition, the targeted maximum likelihood estimator (Q_n^*, g_n) solves the efficient influence curve equation:

$$0 = P_n D^*(Q_n^*, g_n).$$

Remark: Cross-validated initial estimator in the targeted MLE. If the initial estimator is an over-fit, then the bias reduction of the targeted MLE algorithm is not as effective. To protect against such cases one can use a cross-validated initial estimator. Specifically, let $B_n \in \{0,1\}^n$ be a random variable that splits the sample in a training sample $\{i : B_n(i) = 0\}$ and validation sample $\{i : B_n(i) = 1\}$, and, let P_{n,B_n}^0, P_{n,B_n}^1, denote the empirical distribution of the training and validation sample, respectively. The above targeted MLE iterative algorithm is now given by: $Q_{gn}^k = Q_g^k(P_n) = Q_g^{k-1}(P_n)(\epsilon_n^k)$, where

$$\epsilon_n^k = \arg\min_\epsilon E_{B_n} P_{n,B_n}^1 L(Q_g^{k-1}(P_{nB_n}^0)(\epsilon)),\ k = 1, \ldots.$$

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

2.2 Loss-based cross-validation to select among (collaborative) targeted maximum likelihood estimators

Consider a loss function $L^*(Q)$ for Q_0 that satisfies

$$Q_0 = \arg\min_Q P_0 L^*(Q).$$

Or, more precisely, we only require that $\Psi\left(\arg\min_Q P_0 L^*(Q)\right) = \Psi(Q_0)$. An example of such a loss function is the the log-likelihood $L^*(Q)(O) = L(Q) = -\log Q(O)$. Each loss function has a corresponding dissimilarity $d(Q, Q_0) = P_0\{L^*(Q) - L^*(Q_0)\}$.

Given different targeted maximum likelihood estimators, $P_n \to \hat{Q}_k^*(P_n)$, of Q_0, for example, indexed by different initial estimators, we can use a preferred loss-function based cross-validation to select among them. Specifically, let $B_n \in \{0,1\}^n$ be a random variable that splits the sample in a training sample $\{i : B_n(i) = 0\}$ and validation sample $\{i : B_n(i) = 1\}$, and, let P^0_{n,B_n}, P^1_{n,B_n}, denote the empirical distribution of the training and validation sample, respectively. The loss-function based cross-validation selector is now defined by

$$\hat{k}(P_n) = \arg\min_k E_{B_n} P^1_{n,B_n} L^*(\hat{Q}_k^*(P^0_{n,B_n})).$$

The resulting targeted maximum likelihood estimator is then given by

$$\hat{Q}_n^* = \hat{Q}^*_{\hat{k}(P_n)}(P_n).$$

Cross-validation selector: Consider a preferred loss function that satisfies

$$\sup_Q \frac{\mathrm{VAR}_{P_0}\{L^*(Q) - L^*(Q_0)\}}{P_0\{L^*(Q) - L^*(Q_0)\}} \leq M_2, \tag{1}$$

and that is uniformly bounded

$$\sup_{O,Q} \mid L^*(Q) - L^*(Q_0) \mid (O) < M_1 < \infty,$$

where the supremum is over the support of P_0, and over all possible candidate estimators of Q_0 that will ever be considered. The first property (1) applies to the log-likelihood loss function and any weighted squared residual loss function, among others. The property (1) is essentially equivalent with the assumption that the loss-function based dissimilarity $d(Q, Q_0) = P_0 L^*(Q) - L^*(Q_0)$ is quadratic in a distance between Q and Q_0. The property (1) has been proven for log-likelihood loss functions and weighted L^2-loss functions, and is

Collection of Biostatistics
Research Archive

in essence equivalent with stating that the loss function implies a quadratic dissimilarity $d(Q, Q_0)$ (see van der Laan and Dudoit (2003)). If this property does not hold for the loss function, the rates $1/n$ for second order terms in the below stated oracle inequality reduce to the rate $1/\sqrt{n}$.

For such loss functions, the cross-validation selector satisfies the following (so called) oracle inequality: for any $\delta > 0$,

$$E_{B_n}\{P_0 L(\hat{Q}_{\hat{k}}(P^0_{n,B_n}) - L(Q_0)\} \leq (1+2\delta)E_{B_n} \min_k P_0\{L(\hat{Q}_k(P^0_{n,B_n})) - L(Q_0)\} + 2C(M_1, M_2, \delta)\frac{1+\log K(n)}{np},$$

where the constant $C(M_1, M_2, \delta) = 2(1+\delta)^2(M_1/3 + M_2/\delta)$ (see page 25 of van der Laan and Dudoit (2003)). This result proves (see van der Laan and Dudoit (2003) for the precise statement of these implications) that, if the number of candidates $K(n)$ is polynomial in sample size, then the cross-validation selector is either asymptotically equivalent with the oracle selector (based on sample of size of training samples, as defined on right-hand side of above inequality), or it achieves the parametric rate $\log n/n$ for convergence w.r.t. $d(Q, Q_0) \equiv P_0\{L(Q) - L(Q_0)\}$. So in most realistic scenarios, in which none of the candidate estimators achieve the rate of convergence one would have with an a priori correctly specified parametric model, the cross-validated selected estimator selector performs asymptotically exactly as well (up till constant!) as the oracle selected estimator. These oracle results are generalized for estimated loss functions $L^*_n(Q)$ that approximate a fixed loss function $L^*(Q)$. If $\arg\min_Q P_0 L^*_n(Q) \neq Q_0$, then the oracle inequality also presents second order terms due to the estimation of the loss function.

This preferred loss function based cross-validation can now be used to select among different candidate targeted maximum likelihood estimators indexed by different initial estimators, and possibly different censoring mechanism estimators. Specifically, we will use a preferred targeted loss function to select among different collaborative targeted maximum likelihood estimators, which are just special targeted maximum likelihood estimators in the sense that g_n is estimated in collaboration with the initial Q_n.

For a given loss function $L(Q)$, and an estimator $\hat{Q}(P_n)$, we will refer to $P_n L(\hat{Q}(P_n))$ as the entropy of the fit $\hat{Q}(P_n)$. Similarly, for a loss function $L_1(g)$ of g_0, and an estimator $\hat{g}(P_n)$, we will refer to $P_n L_1(\hat{g}(P_n))$ as the entropy of $\hat{g}(P_n)$. Both the preferred loss function for Q_0, as well as this loss function L_1 for g_0 represent important choices. For example, one likes to select the loss function L_1 so that the dissimilarity $P_0\{L_1(g) - L_1(g_0)\}$ measures strongly how well g approaches the optimal fluctuation function implied by g_0.

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

In other words, we need to keep in mind how g is used, namely that it is used to fit the wished fluctuation function implied by g_0. For example, if the clever covariate defining the fluctuation function is given by $A - E(A/\sigma^2(A,W) \mid W)/E(1/\sigma^2(A,W) \mid W)$, as in the semiparametric regression model $E(Y \mid A,W) - E(Y \mid A=0,W) = \beta A$, one might want to define as loss function $L(\theta_1(g),\theta_2(g)) = w_1(A/\sigma^2(A,W) - \theta_1(W))^2 + w_2(1/\sigma^2(A,W) - \theta_2(W))^2$, for weight-functions w_1, w_2 (functions of W), and θ_1, θ_2 representing the numerator and denominator of the conditional expectations in the clever covariate. Similarly, the preferred loss function for Q_0 can be tuned to represent a dissimilarity $d(Q,Q_0)$ that measures strongly how well $\Psi(Q)$ approximates $\Psi(Q_0)$. We discuss such choices in more detail in a later section.

2.3 Building a collaborative estimator of censoring mechanism/nuisance parameter

A C-TMLE estimator is constructed by building a family of candidate estimators, then choosing the best among them, using cross-validation to drive the choice to Q_0. However, we also rely upon a loss function when building each candidate nuisance parameter (e.g. censoring mechanism) estimator, and it is not necessary that these two loss functions be the same. In fact, as part of building a collaborative nuisance parameter estimator in the collaborative T-MLE procedure, we couple an increase in the log-likelihood entropy of the targeted maximum likelihood estimator with an increase in the g_0-loss function specific entropy of the corresponding nuisance parameter estimator. In this manner, we arrange that, for increasing sample size, the cross-validation selector will be driven towards the selection of targeted maximum likelihood estimator with an initial estimator closer to Q_0 *and simultaneously* a more and more nonparametric estimator of g_0 (thereby achieving the full wished bias reduction in the limit).

That is, given a collection of candidate estimators of g_0, ordered by empirical fit w.r.t. a loss function for g_0 such as the log-likelihood, we will build a sequence of targeted maximum likelihood estimators of Q_0 ordered by log-likelihood entropy and indexed by increasingly nonparametric estimators of g_0, where the extend of being nonparamatric is measured by the L_1-entropy. Subsequently, we use the cross-validated log-likelihood for Q_0 to choose among these candidate targeted maximum likelihood estimators.

There are many possible approaches that construct such an ordered sequence of targeted maximum likelihood estimators in which a next element in the sequence has both a higher entropy for the Q_0-loss as well as a higher

Collection of Biostatistics
Research Archive

g_0-loss entropy for its corresponding censoring estimator. Of course, the strict ordering is not what drives the properties of the resulting estimator, but the sequence should represent an approximately monotone function in the log-likelihood entropy of Q_0 and L_1-entropy of g_0.

This procedure represents one particular approach for constructing a targeted maximum likelihood estimator that uses a collaboratively estimated nuisance parameter. We refer to any algorithm that maps into a targeted maximum likelihood estimator that uses a collaborative nuisance parameter estimator (relative to the Q-estimator), as a collaborative targeted maximum likelihood estimator.

2.4 A template for collaborative targeted MLEs

We present the following template providing a class of collaborative targeted maximum likelihood estimators.

Initial estimator of Q_0: Build an estimator Q_n of Q_0, such as a super learner based on the log-likelihood loss function $L(Q)$, or any other loss function.

Preferred loss function for Q_0: Let $L^*(Q)$ be a (targeted) loss function for Q_0. We note that the loss function can also be data dependent, and, in particular, the choice of loss function can depend on an initial estimator Q_n of Q_0, and corresponding collaborative estimator g_n (see DR-IPCW loss functions in van der Laan and Dudoit (2003), and our section on targeted loss functions).

Loss function for g_0: Let $L_1(g)$ be a loss function for g_0.

Candidate estimators of censoring mechanism/nuisance parameter: For each δ in an index set, let $g_{n\delta}$ be a candidate estimator of g_0. Let $d(\delta) = P_n L_1(g_{n\delta})$ denote the entropy of $g_{n\delta}$, thereby measuring how data adaptive $g_{n\delta}$ is, and for a maximal value $d(\delta)$ or for $d(\delta)$ approximating a maximum value we have that $g_{n\delta}$ is actually a consistent estimator of g_0.

Select ordered sequence of entropies for censoring mechanism (nuisance parameter) estimator
Select a sequence $d^0 > d^1 > \ldots > d^K$.

Select initial targeted maximum likelihood estimator: We start out with a g_n^0 with entropy larger than d^0 and a corresponding targeted maximum likelihood estimator $Q_n^{*0} = Q_{ng_n^0}^*$ applied to initial estimator Q_n.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

We refer to the pair (g_n^0, Q_n^{*0}) as the initial targeted maximum likelihood estimator in the sequence of targeted maximum likelihood estimators that will be constructed below.

Construct next targeted maximum likelihood estimator in sequence

We are given an current initial estimator Q_n^k, a current targeted maximum likelihood estimator (g_n^k, Q_n^{k*}) in our sequence of targeted maximum likelihood estimators, with Q_n^{k*} being the targeted maximum likelihood estimator applied to current initial estimator Q_n^k and nuisance parameter estimator g_n^k. The current nuisance parameter estimator g_n^k has entropy larger than d^k. We are also given k, and thereby two corresponding entropy values $d^k > d^{k+1}$. (we note that the initial estimator does not get updated at each step k, but it corresponds with one of the elements in current sequence of targeted maximum likelihood estimators)

Consider an algorithm that searches among a specified set of candidate estimators $g_{n\delta}$ with $\{\delta : d^k > d(\delta) > d^{k+1})\}$ with the goal of minimizing the preferred loss L^*-fit of the targeted maximum likelihood estimator, applied to initial Q_n^k:

$$\delta \rightarrow P_n L^*(Q_{n\delta}^{k*}). \tag{2}$$

Recall that $Q_{n\delta}^{k*}$ denotes the targeted maximum likelihood estimator that uses the optimal fluctuation model identified by censoring mechanism $g_{n\delta}$ applied to initial estimator Q_n^k. Let $g_{n\delta_n}$ be the selected estimator. If either the fit is improved relative to current T-MLE Q_n^{k*},

$$P_n L^*(Q_{n\delta_n}^{k*}) < P_n L^*(Q_n^{k*}),$$

or the above holds for the log-likelihood loss function $L(Q) = -\log Q$ on which the targeted maximum likelihood algorithm operates, then we accept δ_n, and thereby the next targeted maximum likelihood estimator, $g_n^{k+1} = g_{n\delta_n}, Q_n^{k+1*} = Q_{n\delta_n}^{k*}$, in the sequence we are constructing. The algorithm now delivered its next $k+1$-th targeted maximum likelihood estimator. We set $k = k+1$, keep the initial estimator Q_n^k unchanged, and the current targeted maximum likelihood estimator (g_n^k, Q_n^{k*}) is now updated.

If this monotonicity condition fails to hold for both the log-likelihood fit as well as the preferred loss function fit, then we reject this δ_n, and update the initial estimator Q_n^k by setting it equal to the current targeted maximum likelihood estimator Q_n^{k*}. We now, rerun the above procedure with initial $Q_n^k = Q_n^{k*}$, and same $d^k > d^{k+1}$. This time the resulting δ_n

Collection of Biostatistics
Research Archive

will always be accepted since the log-likelihood fit of a targeted maximum likelihood estimator (a maximum likelihood fluctuation of an initial estimator) is larger than the log-likelihood of initial estimator. So the algorithm now delivers the next $k+1$-th targeted maximum likelihood estimator $g_n^{k+1} = g_{n\delta_n}, Q_n^{k+1*} = Q_{n\delta_n}^{k*}$ in its sequence. We set $k = k+1$, the initial estimator is still set at Q_n^k, and the current targeted maximum likelihood estimator (g_n^k, Q_n^{*k}) is now updated (the last one in sequence so far).

k-th step collaborative targeted maximum likelihood estimator:
The above algorithm maps a running current initial estimator, a current targeted MLE (g_n^k, Q_n^{*k}) (the lastly constructed in current sequence), into a new targeted MLE (g_n^{k+1}, Q_n^{*k+1}), and possible updated current initial estimator. We start this algorithm with $k = 0$, and iterate it. This now defines the k-th step collaborative targeted maximum likelihood estimator (g_n^k, Q_n^{*k}), $k = 0, 1, 2, \ldots, K$.

We are guaranteed that the fit of Q_n^{*k} is either increasing w.r.t. the preferred loss function (most likely, since that is the loss we minimize at each step), or it is increasing w.r.t the log-likelihood loss used to define the targeted maximum likelihood step, relative to previous targeted maximum likelihood estimator Q_n^{*k-1}. In addition, the corresponding g_n^k has a L_1-fit that is larger than the L_1-fit of g_n^{k-1}. At every step in which the initial estimator is updated, we also know that the log-likelihood fit is increasing.

Cross-validation to select number of iterations k in k-th step C-TMLE:
Given this sequence of k-th step collaborative targeted maximum likelihood estimators $P_n \to (Q_n^{k*} =)\hat{Q}^{k*}(P_n)$, using estimator g_n^k, it remains to select k, $k = 0, 1, \ldots, K$.

We select k based on the cross-validated log-likelihood:

$$k_n = \underset{k}{\text{argmax}}\ E_{B_n} P_{n,B_n}^1 L(\hat{Q}^{k*}(P_{n,B_n}^0)),$$

where the random vector $B_n \in \{0,1\}^n$ denotes a cross-validation scheme such as V-fold cross-validation, and P_{n,B_n}^0, P_{n,B_n}^1 are the empirical probability distributions of the training sample $\{i : B_n(i) = 0\}$ and validation sample $\{i : B_n(i) = 1\}$, respectively, as identified by the split vector B_n.

This finalizes the mapping from the initial estimator Q_n, and the data, into a collaborative estimator of the censoring mechanism, $g_n = g_n^{k_n}$.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

We refer to $Q_n^* = Q_n^{k_n*}$, paired with collaborative estimator g_n, as the collaborative targeted maximum likelihood estimator of Q_0.

The Collaborative (Double Robust) Targeted Maximum Likelihood Estimator: The corresponding targeted maximum likelihood estimator of $\psi_0 = \Psi^F(Q_0)$ is given by the substitution estimator

$$\Psi(Q_n^*) = \Psi(Q_n^{k_n*}) = \Psi(\hat{Q}^{k_n*}(P_n)).$$

We refer to this estimator as the collaborative (double robust) targeted maximum likelihood estimator (C-DR-TMLE or C-TMLE) of ψ_0, and we recall that it is paired with a collaborative estimator g_n.

C-TMLE solves an efficient influence curve equation: Since the C-TMLE is a targeted maximum likelihood estimator $Q_n^{k_n*}$, applying the fluctuation function with censoring mechanism estimator $g_n = g_n^{k_n}$ to the estimator $Q_n^{k_n}$, it solves the efficient influence curve equation:

$$0 = P_n D^*(Q_n^*, g_n).$$

This is a fundamental property of the collaborative targeted MLEs driving the targeted bias reduction w.r.t. the target parameter of interest,ψ_0.

Selection among candidate C-TMLEs: The collaborative targeted maximum likelihood estimator depends on a choice of initial estimator Q_n^0, and choices that concern the second stage. As a consequence, one might have a set of collaborative targeted maximum likelihood estimators (Q_{nj}^*, g_{nj}) indexed by such choices, $j = 1, \ldots, J$. We can now select among these estimators Q_{nj}^* based on loss-based cross-validation using the preferred loss function L^* for Q_0.

Selection based on empirical efficiency maximization: Since, under regularity conditions of our asymptotic linearity theorem, each j-specific C-TMLE is asymptotically linear with influence curve $IC_j(Q_j^*, g_{0j})$ (equal to $D^*(Q_j^*, g_{0j})$ plus a contribution from g_{nj}), we can select j as the minimizer of a (cross-validated) estimate of the variance of $IC_j(Q_j^*, g_{0j})$, or, if ψ_0 has dimension larger than 1, then we can minimize an estimate of the variance of a function of ψ_0. One could here ignore the contribution from g_{nj} and thus use the cross-validated or empirical variance of the efficient influence curve at the collaborative targeted maximum likelihood estimator:

$$j_n = \arg\min_j E_{B_n} P_{n,B_n}^1 \left\{ D^*(\hat{Q}_j^*(P_{n,B_n}^0), \hat{g}_j(P_{n,B_n}^0)) \right\}^2.$$

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Generalization. The above C-TMLE can also be called the collaborative minimum loss estimator. The loss function $L(Q)$ needs to satisfy that the derivative of $\epsilon \to L(Q_g(\epsilon))$ at $\epsilon = 0$ for a suitably constructed path $\{Q_g(\epsilon) : \epsilon\}$ equals the efficient influence curve $D^*(Q, g)$, where the efficient influence curve at P only depends on $Q(P)$ and $g(P)$, while the target parameter $\Psi(P) = \Psi^F(Q(P))$ depends on P only through $Q(P)$. No further structure is needed for the above template (such as $dP = Q * g$, or CAR-censored data structure).

2.5 The rationale of the consistency of the collaborative-TMLE

The C-TMLE procedure starts with an initial estimator Q_n of Q_0. Suppose that the sequence constructed in the C-TMLE template consists of a finite number K of targeted maximum likelihood estimators Q_n^{k*}. By construction, the last targeted maximum likelihood estimator in this sequence uses a censoring mechanism estimator that is nonparametric (maximal g_0-entropy): i.e, the nuisance parameter estimator g_n^K as selected by the K-th step C-TMLE converges to the true g_0. We also know that g_n^k is increasingly nonparametric in k, $k = 1, \ldots, K$.

For simplicity, we also assume that the k-th targeted maximum likelihood estimator in the sequence is obtained by applying the targeted maximum likelihood algorithm to the previous targeted maximum likelihood estimator in sequence. This is not necessary, since we can apply the argument to the subsequence for which that is true (the elements in the sequence at which the targeted maximum likelihood update is actually carried out), but it simplifies the presentation.

Consider the limits $Q^*_{kg_k}$ of the targeted maximum likelihood estimators Q_n^{*k} in our sequence, where g_k is the limit of g_n^k, $k = 1, \ldots, K$, and thus $g_K = g_0$. We also know that $P_n \log Q^*_{nkg_{nk}}$ is increasing in k, by the fact that each element in the sequence is a targeted maximum likelihood estimator applied to previous element in sequence (as initial estimator in the T-MLE algorithm). Therefore, $P_0 \log Q^*_{kg_k}$ is non-decreasing in k. As discussed in introduction, if the log Q is uniformly bounded in all its candidates Q, then the cross-validation selector of k is asymptotically equivalent with the oracle selector $\tilde{k}_n = \arg\max P_0 \log Q^*_{kg_{nk}}$. For n large enough, this oracle selector behaves as $\tilde{k} = \arg\max_k P_0 \log Q^*_{kg_k}$, where this maximum might be non-unique. One maximum is obtained at $k = K$, giving $P_0 \log Q^*_{Kg_0}$ and, we know that $\Psi(Q^*_{Kg_0}) = \psi_0$. So if $\tilde{k} = K$, then the c-tmle will be consistent for ψ_0. Suppose that $\tilde{k}$ is actually smaller than K. Then we have, suppressing the g's in the

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

notation,

$$P_0 \log Q^*_{\tilde{k}} = P_0 \log Q^*_{\tilde{k}+1} = \ldots = P_0 \log Q^*_K.$$

We know that Q^{*k+1} is a T-MLE with Q^{*k} as initial. So the above equalities are only possible if $Q^{*k+1} = Q^{*k}$ for $k = \tilde{k}, \ldots, K-1$. Thus $Q^*_K = Q^*_{\tilde{k}}$. Since Q^*_K is a targeted MLE at nuisance parameter g_0, it follows that $\epsilon \to P_0 \log Q^*_{K,g_0}(\epsilon)$ is maximized at $\epsilon = 0$: compare with $P_n \log Q^*_{nK}(\epsilon)$ is maximized at $\epsilon = 0$ by definition of the T-MLE algorithm. Since we just showed that $Q^*_K = Q^*_{\tilde{k}}$, it also follows now

$$\epsilon \to P_0 \log Q^*_{\tilde{k},g_0}(\epsilon)$$

is maximized at $\epsilon = 0$. In particular, this means that the derivative at $\epsilon = 0$ equals zero, giving us:

$$0 = P_0 D^*(Q^*_{\tilde{k}}, g_0).$$

However, the efficient influence curve typically satisfies that $P_0 D^*(Q, g_0) = 0$ implies $\Psi(Q) = \psi_0$, which then implies $\Psi(Q^*_{\tilde{k}}) = \psi_0$. Thus $\Psi(Q^*_{n\tilde{k}})$ is consistent, and thereby $\Psi(Q^*_{nk_n})$ is consistent.

Figure 1 illustrates the collaborative nature of the construction of a sequence of increasingly data-adaptive nuisance parameter estimators, $\{g^1_n, \ldots, g^K_n\}$, and its relation to the performance of the initial estimator. We generated 5000 observations of $O = (W, A, Y)$ from data generating distribution $dP_0 = Q_0 g_0$ defined as:

$$\begin{aligned} logit\left(g_0(A \mid W)\right) &= .15W_1 + .1W_2 + W_3 - W_4 \\ Q_0(A, W) &= A + 3W_1 - 6W_2 + 4W_3 - 5W_5 + 3W_4 \end{aligned}$$

where W_1 through W_5 are independent random variables $\sim N(0,1)$, $Y = Q_0(A,W) + \epsilon, \epsilon \sim N(0,1)$, and g_0 is the conditional density of A given confounding variables $W = \{W_1, W_2, W_3, W_4, W_5\}$. We applied the C-TMLE to estimate the effect of binary treatment A on outcome Y, adjusting for W, defined as $\psi_0 = E_W\left(E(Y \mid A = 1, W) - E(Y \mid A = 0, W)\right)$.

A kernel density estimator was applied to Y and to the predicted values of two initial estimators of $Q_0 = E_0(Y \mid A, W)$, which we denote with $\hat{Q}^0_{n,poor}$, and $\hat{Q}^0_{n,good}$, respectively. These estimators were obtained with the D/S/A algorithm (Sinisi and van der Laan, 2004), a data-adaptive machine learning approach to model selection that was set to search over all second degree polynomials of size six. The kernel density estimates are displayed in plots on the left hand side of the figure.

In addition, we plotted the kernel density estimates of the predicted values of each set of the collaboratively-constructed candidate $\hat{g}$ estimators, and we can compare them with the density of the true predictions $g_0(1 \mid W) = P_0(A =$

Collection of Biostatistics
Research Archive

1|W). These are plotted on the right hand side of the figure, overlaid with the density estimator applied to the true values $g_0(1 \mid W)$. When the initial fit of Q_0 is poor, the nuisance parameter estimator g_n^k converges quickly to g_0 in k, and the selected candidate estimator closely approximates g_0. Plots in the bottom half of the figure shows the behavior of the C-TMLE procedure when Q_n^0 is a good estimate of Q_0. When the initial fit of Q_0 is good, the nuisance parameter estimator grows slowly towards g_0, and a candidate estimator that estimates a true treatment mechanism that adjusts for fewer covariates than the true treatment mechanism g_0 that was used to generate the data.

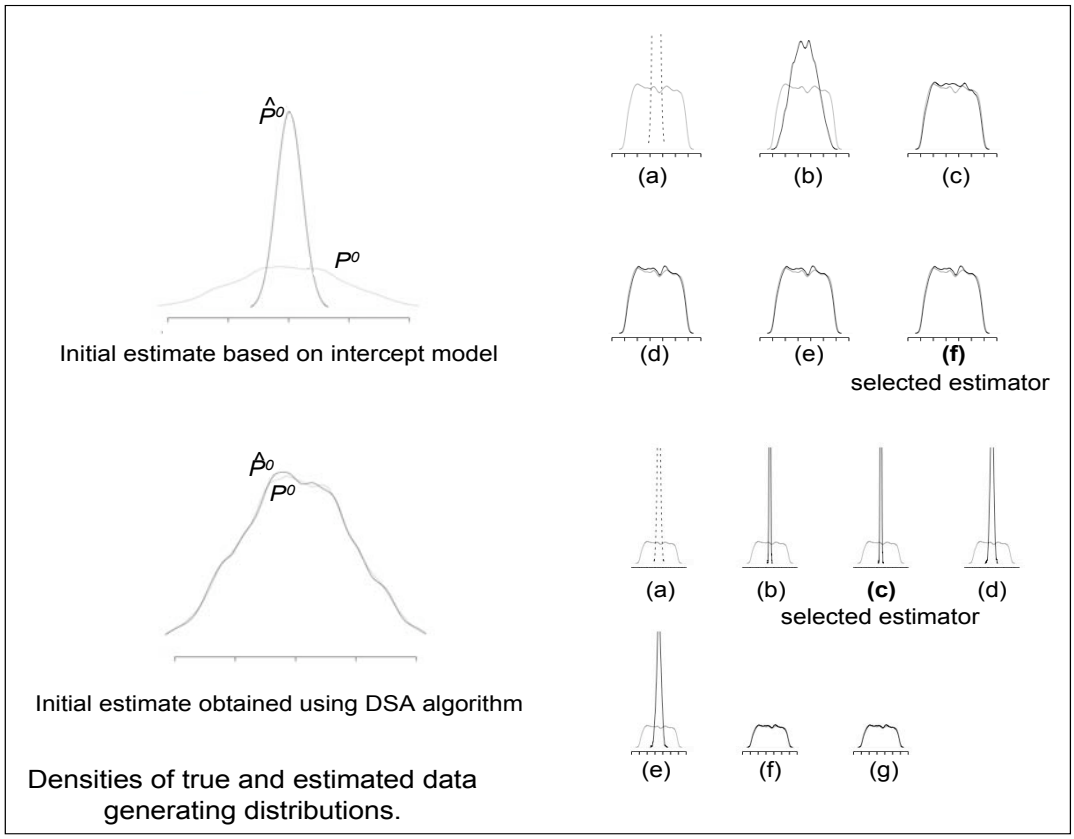

Figure 1: Construction of a sequence of nuisance parameter estimators based on a poor initial fit of the density (top) and a good initial fit for the density (bottom). Kernel estimates of true densities Q_0 and g_0 are shown in gray.

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

2.6 Revisiting the additive causal effect example

Recall that the targeted maximum likelihood estimator applied to an estimator, Q_n, of $P(Y = 1 \mid A, W)$ is obtained by running a univariate logistic regression of Y, with offset the initial estimator, on an estimate of the univariate clever covariate $h_{g_0}(A, W) = A/g_0(1 \mid W) - (1 - A)/g_0(0 \mid W)$, implied by the treatment mechanism estimator, using an estimator g_n of treatment mechanism $P(A = 1 \mid W)$.

The collaborative targeted maximum likelihood estimation procedure starts with computing Q_n^0, an initial estimator of $P(Y = 1 \mid A, W)$ using super learning, and then collaboratively generating a sequence of targeted maximum likelihood estimators. These use increasingly nonparametric estimators of g_0, applied to subsequent targeted maximum likelihood updates of the initial estimator (as needed to guarantee the monotonicity in fit). In this way the sequence of constructed targeted maximum likelihood estimators has increasing log-likelihood fit. The selection of the sequence of increasingly nonparametric treatment mechanism estimators was based on maximizing the fit of the corresponding targeted maximum likelihood estimators of $P(Y = 1 \mid A, W)$, as outlined in our template, thus very much driven by the outcome data. Likelihood based cross-validation selects the wished targeted maximum likelihood estimator, with its paired treatment mechanism estimator, from this sequence. It is assumed that the resulting selection of the estimator of g_0 is nonparametric enough so that the collaborative double robustness of the efficient influence curve as presented in next section is utilized, and, thereby, that our asymptotic linearity theorem in later section can indeed be applied.

A collaborative targeted maximum likelihood estimator constructed in this manner has made every effort to make the estimator of the additive causal effect as unbiased as possible. If we now construct a set of such collaborative targeted maximum likelihood estimators, possibly indexed by different initial estimators, and different ways of constructing the sequence of targeted maximum likelihood estimators, we can then select among these estimators the estimator with minimal estimated variance (based on the influence curve). To obtain an honest estimate of the variance of the resulting estimator, just as one obtained honest cross-validated risk of an estimator that internally uses cross-validation, one uses the honest cross-validated variance of the influence curve of the complete estimator, including cross-validating this final selection step that involves minimizing the variance.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

3 Collaborative double robustness of estimating functions in CAR censored data models

In this section we establish a new kind of collaborative robustness of the class of estimating functions in CAR-censored data models, where, as in van der Laan and Robins (2003), the class of estimating functions is implied by the orthogonal complement of the nuisance tangent space of the target parameter $\Psi : \mathcal{M} \to \mathbb{R}^d$. This orthogonal complement of the nuisance tangent space equals the space spanned by the gradients of the pathwise derivative of Ψ, and thus includes the canonical gradient/efficient influence curve. The collaborative robustness result teaches us that the censoring mechanism required to obtain an unbiased estimating function at a mis-specified Q for the parameter of interest need not always condition on the whole full data structure. In fact, it teaches us that the better Q approximates Q_0 the less of an adjustment by full data random variables is necessary for the censoring mechanism to still obtain an unbiased estimating function for the parameter of interest. The precise collaborative property of $(Q, g_0(Q))$ such that $P_0 D(\psi_0, g_0(Q), Q) = 0$ will be explicitly specified, where D represents the estimating function, such as the one implied by the canonical gradient.

3.1 The formal collaborative robustness result

The new form of double robustness we wish to establish is understood as follows. Consider an estimating function $D(\Psi(Q), G, Q)$ for the parameter of interest ψ_0 that is indexed by nuisance parameters (G_0, Q_0), and which is already known to satisfy the classical double robustness property: for any G under which ψ_0 is identifiable from $P_{Q_0,G}$, we have $E_0 D(\psi_0, G, Q) = 0$ if either $Q = Q_0$ or $G = G_0$ (van der Laan and Robins (2003)). Given a Q, we are interested in the question under what conditional distribution $G_{0\delta}$ of censoring variable C, given a reduction $X(\delta)$ of X, will we still have $P_0 D(\psi_0, G_{0\delta}, Q) = 0$ and thereby that D is an unbiased estimating function for ψ at this mis-specified Q.

Firstly, we note that $P_0 D(\psi_0, G, Q) = P_0\{D(\psi_0, G, Q) - D(\psi_0, G, Q_0)\} + P_0 D(\psi_0, G, Q_0)$, and the latter term is zero under any G that allows identifiability of ψ_0. Thus, it remains to determine for what $G_{0\delta}$ we will have $P_0\{D(\psi_0, G_{0\delta}, Q) - D(\psi_0, G_{0\delta}, Q_0)\} = 0$. This choice of $G_{0\delta}$ (e.g., it includes G_0 itself) is not unique but will be dependent on a difference $Q - Q_0$ in the sense that $X(\delta)$ has to be rich enough so that it contains a difference $Q - Q_0$.

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

By the general representation theorem for estimating functions that are orthogonal to the nuisance tangent space of the target parameter (Theorem 1.6, van der Laan and Robins (2003)), one can typically represent an estimating function $D(\psi_0, G, Q)$ as an Inverse Probability of Censoring Weighted Estimating function $D_{IPCW}(G, \psi_0)$ plus a function $D_{CAR}(Q, G)$ in the tangent space $T_{CAR}(G)$ of the censoring mechanism at G. The function $D_{CAR}(Q, G)$ is defined as the projection of $-D_{IPCW}(G, \psi_0)$ on the tangent space $T_{CAR}(G) = \{h(O) : E_G(h(O) \mid X) = 0\}$ of the censoring mechanism when only assuming coarsening at random, where this projection is carried out in the Hilbert space of all functions of O with mean zero and finite variance endowed with inner product the covariance operator $\langle f_1, f_2 \rangle = E_{Q,g} f_1(O) f_2(O)$. In other instances, the D_{IPCW} might depend on Q_0 through another parameter beyond ψ_0, in which case it will need to be assumed that this parameter is correctly specified.

This teaches us that $P_0\{D(\psi_0, G, Q) - D(\psi_0, G, Q_0)\} = P_0\{D_{CAR}(Q, G) - D_{CAR}(Q_0, G)\}$, since the IPCW-difference equals zero. This representation theorem also teaches us that for all Q we have that $D_{CAR}(Q, G)$ has conditional mean zero under G, given X. In addition, this same theorem also shows that $Q \to D_{CAR}(Q, G)$ is linear in Q. Therefore, it remains to show that $P_0 D_{CAR}(Q - Q_0, G) = 0$. Now, inspection of the proof that the conditional mean of $D_{CAR}(Q', G)$ under G equals zero for a Q' involves typically conditioning on a rich enough reduction of X so that a particular function indexed by Q' is fixed under the conditioning. Thus, the censoring mechanism only needs to condition on a particular function of $Q - Q_0$.

This is best illustrated with a concrete censored data structure. For example, consider the right censored data structures $O = (C, \bar{X}(C))$, where $X(t)$ is a time dependent process, $X = (X(t) : t)$ represents the full data structure, and $\bar{X}(t) = \{X(s) : s \leq t\}$ represents the sample path up till time t. For this censored data structure, one can represent the projection of D_{IPCW} onto T_{CAR} as $D_{CAR}(Q, G) = \int H_{Q,G}(u, \bar{X}(u-)) dM_G(u)$, where

$$\begin{aligned} H_{Q,G}(u, \bar{X}(u-)) &= E_{Q,G}(D_{IPCW,G} \mid C = u, \bar{X}(u)) - E(D_{IPCW,G} \mid C \geq u, \bar{X}(u)) \\ dM_G(u) &= I(C = u) - I(C \geq u) d\Lambda_{C|X}(u \mid X), \end{aligned}$$

and $\Lambda_{C|X}$ is the cumulative hazard of C, given X. For details, we refer to chapter 3 in van der Laan and Robins (2003). Here $dM_G(u)$ is a Martingale satisfying $E(dM_G(u) \mid \bar{X}(u), C \geq u) = 0$. Due to the linearity of the conditional expectation operator, we have $D_{CAR}(Q - Q_0, G) = \int H_{Q-Q_0,G}(u, \bar{X}(u)) dM_G(u)$. By conditioning on $H_{Q-Q_0,G}(u, \bar{X}(u))$ within the integral, and using $E(dM_G(u) \mid \bar{X}(u), C \geq u) = 0$, it follows that $D_{CAR}(Q - Q_0, G)$ also has mean zero

Collection of Biostatistics
Research Archive

under a censoring mechanism s.t. $\lambda_C(u \mid X)$ only depends on $\bar{X}(u)$ (it only depends on $\bar{X}(u)$ by CAR) through $H_{Q-Q_0,G}(u, \bar{X}(u))$. One can factorize $H_{Q-Q_0,G} = H_1(G)H_2(Q - Q_0)$, so that adjustment in $\lambda_C(u \mid X)$ by the time-dependent covariate $H_2(Q - Q_0)(u, \bar{X}(u))$ suffices. Alternatively, it also suffices if the censoring mechanism uses a self-iterated adjustment by $H_{Q-Q_0,G}$ as described later in this section. If Q approximates Q_0, this function $H_{Q-Q_0,G}(u, \bar{X}(u))$ will be shrunk to zero, so that less conditioning becomes necessary.

The following much simpler (but in essence making the same point) example helps to further illustrate the general collaborative double robustness property of the efficient influence curve. Suppose the observed censored data structure is $O = (W, \Delta, \Delta Y)$ and $X = (W, Y)$ is the full data random variable, where Δ is the censoring variable. Suppose one wishes to estimate $\psi_0 = E_0 Y$. The efficient influence curve is given by

$$D(\psi_0, \Pi_0, Q_0) = D_{IPCW}(\psi_0, \Pi_0) + D_{CAR}(Q_0, \Pi_0),$$

where

$$\begin{aligned} D_{IPCW}(\psi_0, \Pi_0) &= Y \frac{\Delta}{\Pi_0(W)} - \psi_0 \\ D_{CAR}(Q_0, \Pi_0) &= -E(Y \mid \Delta = 1, W) \left(\frac{\Delta}{\Pi_0(W)} - 1 \right), \end{aligned}$$

$\Pi_0(W) = P_0(\Delta = 1 \mid W)$ and $Q_0(W) = E_0(Y \mid W, \Delta = 1)$. Consider a Q. We are interested in the question under what conditional distribution $\Pi_{0\delta}$ of Δ, given a reduction $W(\delta)$ of W, will we still have $P_0 D(\psi_0, \Pi_{0\delta}, Q) = 0$ and thereby that D is an unbiased estimating function for ψ at this mis-specified Q. Firstly, we note that $P_0 D(\psi_0, \Pi, Q) = P_0\{D(\psi_0, \Pi, Q) - D(\psi_0, \Pi, Q_0)\} + P_0 D(\psi_0, \Pi, Q_0)$, and the latter term is zero under any Π for which $P_0(\Pi(W) > 0) = 1$. Thus, it remains to determine for what $\Pi_{0\delta}$ $P_0\{D(\psi_0, \Pi_{0\delta}, Q) - D(\psi_0, \Pi_{0\delta}, Q_0)\} = 0$.

This teaches us that $P_0\{D(\psi_0, \Pi, Q) - D(\psi_0, \Pi, Q_0)\} = P_0\{D_{CAR}(Q, \Pi) - D_{CAR}(Q_0, \Pi)\}$, since the IPCW-difference equals zero:

$$P_0\{D(\psi_0, \Pi, Q) - D(\psi_0, \Pi, Q_0)\} = \frac{(Q - Q_0)(W)}{\Pi_0(W)} (\Delta - \Pi_0(W)).$$

Note that we used here that $Q \to D_{CAR}(Q, \Pi)$ is linear in Q. Therefore, it remains to show that $P_0 D_{CAR}(Q - Q_0, \Pi) = 0$. This can be represented as $H(Q - Q_0, \Pi_0)(W)(\Delta - \Pi_0(W))$ as above, with $H(Q - Q_0, \Pi_0)(W) = (Q - Q_0)(W)/\Pi_0(W)$.

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

The proof that the conditional mean of $D_{CAR}(Q - Q_0, \Pi)$ under Π equals zero involves conditioning on a rich enough reduction of W so that $Q - Q_0$ is captured by the conditioning: if $(Q - Q_0)(W)$ only depends on W through $W(\delta)$, then

$$E\frac{(Q - Q_0)(W)}{\Pi_0(W(\delta))}(\Delta - \Pi_0(W(\delta))) = 0.$$

In particular, we have that the conditional mean of $D_{CAR}(Q - Q_0, \Pi_0)$, given $(Q - Q_0)(W)$, equals zero if $\Pi_0(W) = P(\Delta = 1 \mid Q - Q_0(W))$. This shows that if, for example, $(Q - Q_0)(W)$ only depends on one component W_1, then $P_0 D(\psi_0, \Pi_0, Q) = 0$ for $\Pi_0(W_1) = P_0(\Delta = 1 \mid W_1)$, and, more general, for $\Pi_0(W')$ with $W_1 \subset W'$. That is, the better job Q does in approximating Q_0 the less inverse probability of missingness weighting is required to still obtain an unbiased estimating function for ψ_0.

Summary: Consider the efficient influence curve $D(Q, G)$. Suppose we already know that for Q with $\Psi(Q) = \psi_0$ $P_0 D(Q_0, G) = 0$ for all G. Given a Q with $\Psi(Q) = \psi_0$, characterize the set of $G_0(Q)$s for which $P_0 D(Q_0, G_0(Q)) - D(Q, G_0(Q)) = 0$. For such Q and corresponding $G_0(Q)$'s we have $P_0 D(Q, G_0(Q)) = 0$. Given the representation theorem for estimating functions derived from the orthogonal complement of the nuisance tangent space, it appears that we need to determine the conditional distributions $G_{0\delta}$ of C, given a reduction $X(\delta)$ of X, for which $E_0 D_{CAR}(Q - Q_0, G_{0\delta}(Q)) = 0$. Thus we need to determine the conditional distributions $G_0(Q)$ of C that solves the score equation $E_0 D_{CAR}(Q - Q_0, G_0) = 0$ of score $D_{CAR}(Q - Q_0, G_0)$. In particular, if $G_0(Q)$ is a MLE of a finite dimensional parameter (e.g., same dimension as ψ_0), whose score spans $D_{CAR}(Q - Q_0, G_0)$, then $E_0 D_{CAR}(Q - Q_0, G_0) = 0$. More generally, if G_0 is a limit of an efficient (e.g. NPMLE) estimator in a model for G_0 that has a tangent space at G_0 that contains $D_{CAR}(Q - Q_0, G_0)$, then this G_0 also satisfies $E_0 D_{CAR}(Q - Q_0, G_0) = 0$. In addition, a self-iterated iterative MLE, starting with arbitrary offset G, for a parameter with score $D_{CAR}(Q - Q_0, G)$, at G, can be employed as well, as presented below, resulting in an updated $G_0(Q)$ of G so that $E_0 D_{CAR}(Q - Q_0, G_0(Q)) = 0$.

We will now present the general result which can be applied to any CAR-censored data model as defined and studied in van der Laan and Robins (2003).

Theorem 1 *(***Collaborative Double Robustness of Efficient Influence Curve/Estimating Functions***)*

CAR-censored data model: *Let $O = \Phi(C, X) \sim P_0$ be a censored data structure with full data random variable $X \sim P_{X0}$, and censoring variable C with conditional probability distribution G_0 of C, given X. Assume G_0 satisfies the coarsening at random assumption. Let $g_0(C \mid X) = dG_0(C \mid X)$*

COBRA
Collection of Biostatistics
Research Archive

a probability density of G_0 w.r.t. an appropriate dominating measure that satisfies coarsening at random itself. Let $\mathcal{M}$ denote the observed data model for P_0. Due to CAR, we have w.r.t. an appropriate dominating measure $dP_0(O) = Q_0(O)g_0(O \mid X)$, where $g_0(O \mid X)$ is only a function of O (by CAR), and Q_0 denotes the identifiable part of the full data distribution P_{X0} (Gill et al. (1997)). (Here we abused notation to indicate that the conditional density of O, given X, is a deterministic function of the conditional density of C, given X, and, in fact, represents the identifiable part of the censoring mechanism G_0.) Let $\mathcal{Q}$ and $\mathcal{G}$ be models for Q_0 and G_0 which imply a model $\mathcal{M} = \{dP = Qg : Q \in \mathcal{Q}, G \in \mathcal{G}\}$ for P_0.

Parameter of interest: *Let $\Psi : \mathcal{M} \to \mathbb{R}^d$ be pathwise differentiable parameter of interest and it is assumed that $\Psi(P_0) = \Psi^F(Q_0)$ is only a function of Q_0. Let $D^*(Q, G)$ be the efficient influence curve/canonical gradient of Ψ at $dP = Qg$.*

We make the following assumptions:

Augmented "PCW"-representation of efficient influence curve:
(PCW stands for Probability of Censoring Weighted) For each $Q \in \mathcal{Q}$, $G \in \mathcal{G}$,

$$D^*(G, Q) = D_{h(G,Q)}(G, Q) = D_{h(G,Q),PCW}(G, \Gamma(Q)) + D_{h(G,Q),CAR}(G, Q'),$$

for mappings $(G, Q) \to h(G, Q)$, $(h, G, Q) \to D_{h,PCW}(G, \Gamma(Q))$, $(h, G, Q') \to D_{h,CAR}(G, Q'(Q, G))$, both defined on $\mathcal{H} \times \mathcal{G} \times \mathcal{Q}$, a parameter mapping Γ on $\mathcal{Q}$, and $(G, Q) \to Q'(G, Q)$.

(We refer to Theorem 1.3 in van der Laan and Robins (2003) for such a general representation of the efficient influence curve and, more generally, the orthogonal complement of the nuisance tangent space, where the CAR-components are elements of the tangent space T_{CAR} of G consisting of all functions of O with conditional mean zero, given X, under G. Under that representation, we have that $E_0 D_{h,PCW}(G_0, \gamma_0) = 0$ and $D_{h,CAR}(G_0, Q')$ has conditional mean zero, given X, for all Q'.)

Linearity of CAR-component: *$Q' \to D_{h,CAR}(G, Q')$ is linear on a set $\overline{\mathcal{Q}}'$ containing $\{Q'(G, Q) : G, Q\}$ in the sense that for all $h \in \mathcal{H}$, and all $Q_1, Q_2 \in \overline{\mathcal{Q}}'$*

$$D_{h,CAR}(G, Q_1') - D_{h,CAR}(G, Q_2') = D_{h,CAR}(G, Q_1' - Q_2').$$

Robustness for mis-specified censoring mechanism: *For all $Q_0 \in \mathcal{Q}_0$ and $G \in \mathcal{G}(Q_0) \subset \mathcal{G}$, where (e.g.,) $\mathcal{G}(Q_0)$ is defined as all censoring*

COBRA
Collection of Biostatistics
Research Archive

mechanisms G for which ψ_0 can be identified from $dP = dQ_0 g$, we have

$$E_0 D_h(G, Q_0) = 0 \text{ for all } h \in \mathcal{H}.$$

Robustness of CAR-component: *For a reduction $X(\delta)$ of X (i.e., $X(\delta) = f(X, \delta)$ for some function f), let $G_{0\delta}$ be the conditional distribution of C, given $X(\delta)$.*

Let $\overline{\mathcal{Q}_\delta}'$ be a set within $\overline{\mathcal{Q}}'$ for which for each $\bar{Q}' \in \overline{\mathcal{Q}_\delta}'$

$$E_0 D_{h,CAR}(G_{0\delta}, \bar{Q}') = 0.$$

(Typically, one can select $\overline{\mathcal{Q}_\delta}'$ as all functions in $\overline{\mathcal{Q}}'$ that are only functions of X through $X(\delta)$.)

Let $\Gamma(Q) = \Gamma(Q_0)$ (typically implying $\Psi(Q) = \psi_0$), $G_{0\delta} \in \mathcal{G}(Q_0)$, and assume $Q' - Q'_0 \in \overline{\mathcal{Q}_\delta}'$, where $Q' = Q'(G_{0\delta}, Q)$ and $Q'_0 = Q'(G_{0\delta}, Q_0)$. Then

$$E_0 D^*(G_{0\delta}, Q) = 0.$$

We also have for all $G \in \mathcal{G}(Q_0)$

$$E_0 D^*(G, Q_0) = 0.$$

Proof. Suppose $\Gamma(Q) = \Gamma(Q_0)$ and $Q' - Q'_0 \in \overline{\mathcal{Q}_\delta}'$. Let $G_0^* = G_{0\delta}$ be the conditional distribution of C, given $X(\delta)$, and assume it is an element of $\mathcal{G}(Q_0)$.

By the "Augmented 'PCW'-representation of efficient influence curve" assumption, we have

$$E_0 D^*(G_0^*, Q) = E_0 D_h(G_0^*, Q)$$

for some $h \in \mathcal{H}$. Thus,

$$\begin{aligned} E_0 D^*(G_0^*, Q) &= E_0 D_h(G_0^*, Q) \\ &= E_0\{D_h(G_0^*, Q) - D_h(G_0^*, Q_0)\} + E_0 D_h(G_0^*, Q_0). \end{aligned}$$

By the assumption that $G_0^* \in \mathcal{G}(Q_0)$, it follows that the last term $E_0 D_h(G_0^*, Q_0) = 0$.

By the "PCW-representation" assumption we have

$$\begin{aligned} E_0\{D_h(G_0^*, Q) - D_h(G_0^*, Q_0)\} = {} & E_0\{D_{h,PCW}(G_0^*, \Gamma(Q)) - D_{h,PCW}(G_0^*, \Gamma(Q_0))\} \\ & + E_0\{D_{h,CAR}(G_0^*, Q'(Q, G_0^*)) - D_{h,CAR}(G_0^*, Q'(Q_0, G_0^*))\}. \end{aligned}$$

COBRA Preprint Series
Collection of Biostatistics Research Archive

By the assumption that $\Gamma(Q) = \Gamma(Q_0)$, the first term equals zero. By the "linearity of CAR-component"-assumption we have that the last term equals:

$$E_0\{D_{h,CAR}(G_0^*, Q') - D_{h,CAR}(G_0^*, Q_0')\} = E_0 D_{h,CAR}(G_0^*, Q' - Q_0'),$$

where $Q' = Q'(G_0^*, Q)$ and $Q_0' = Q'(G_0^*, Q_0)$.

We assumed that $Q' - Q_0' \in \overline{\mathcal{Q}_\delta}'$. Thus, by the "Robustness of CAR-component"-assumption we have that

$$E_0 D_{h,CAR}(G_0^*, Q' - Q_0') = 0.$$

This proves $E_0 D^*(G_0^*, Q) = 0$. □

3.2 Examples illustrating the collaborative double robustness in censored data models

For the sake of illustration, we will now explicitly establish the collaborative double robustness of the efficient influence curve estimating function in two additional examples. These results are also corollaries of the above general Theorem 1.

3.2.1 Example I: Marginal additive causal effect in nonparametric model

We have the following double robustness result for our additive causal effect example.

Theorem 2 *Let $dP_0 = Q_0 dG_0$ be the distribution of $O = (W, A, Y)$ and let the model for P_0 be nonparametric.*

Let $\Psi(Q_0) = E_{Q_{01}}\{E_{Q_{02}}(Y \mid A = 1, W) - E_{Q_{02}}(Y \mid A = 0, W)\}$ be the parameter on this model, where it is assumed that it is identifiable from P_0. Here Q_{01} denotes marginal distribution of W and Q_{02} the conditional distribution of Y, given A, W. The efficient influence curve of Ψ at $P = (Q, G)$ is given by

$$D^*(Q, G)(O) = h(G)(A, W)(Y - Q_2(A, W)) + Q_2(1, W) - Q_2(0, W) - \Psi(Q),$$

where $Q_2(A, W) = E_Q(Y \mid A, W)$ denotes the conditional mean of Y, given A, W, under $Q = (Q_1, Q_2)$.

Assume

$$(Q_{02} - Q_2)(A, W) = E_{Q_0}(Y - Q_2(A, W) \mid A, W) = f_0(A, W(Q))$$

CUBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

is only a function of $A, W(Q)$ for a $W(Q) = \Phi(Q_2, W)$ for some mapping Φ: i.e., $W(Q)$ denotes a reduction or subset of the full vector random variable W indexed by Q.

Let $dG_0(Q)$ be the conditional distribution of A, given $W(Q)$. If $\Psi(Q) = \Psi(Q_0)$, then

$$E_{P_0} D^*(Q, G_0(Q)) = 0.$$

Or, equivalently, if we represent $D^*(Q, G)$ as $D^*(\Psi(Q), Q, G)$, then

$$E_{P_0} D^*(\psi_0, Q, G_0(Q)) = 0.$$

We also have: If $Pr(P_G(A = 0 \mid W) * P_G(A = 1 \mid W) > 0) = 1$, then

$$E_{P_0} D^*(Q_0, G) = 0,$$

or equivalently,

$$E_{P_0} D^*(\psi_0, Q_0, G) = 0.$$

Proof. The last statement is easy and well known (e.g., van der Laan and Robins (2003)). The first statement needs to be proved, or can be derived as a corollary of Theorem 1. Note, if $\Psi(Q) = \psi_0$, then

$$E_0 D^*(Q, G_0(Q)) = E_0 h(G_0)(A, W(Q))(Y - Q(A, W)) + Q(1, W) - Q(0, W) - \psi_0.$$

If $E_0(Y - Q(A, W) \mid A, W) = f_0(A, W(Q))$ is only a function of $A, W(Q)$, then it follows by first taking the conditional mean, given A, W, and then taking the mean of A, given $W(Q)$,

$$\begin{aligned} E_0 D^*(Q, G_0(Q)) &= E_0 h(G_0)(A, W(Q)) f_0(A, W(Q)) \\ &\quad + Q(1, W) - Q(0, W) - \psi_0 \\ &= E_0 f_0(1, W(Q)) - f_0(0, W(Q)) + Q(1, W) - Q(0, W) \\ &\quad - \psi_0. \end{aligned}$$

Now, note that $f_0(A, W(Q)) = Q_0(A, W) - Q(A, W)$, which proves that the latter quantity equals zero.

□

The implication of this result is that, given an estimate Q of Q_0, we only need to estimate $G_0(Q)$, conditioning on $W(Q)$, or any conditional distribution that conditions on more than $W(Q)$. Thus, if Q already succeeds in explaining most of the true regression $E_0(Y \mid A, W)$, then only little inverse weighting with $G_0(Q) = P(A = \cdot \mid W(Q))$ remains to be done. That is, the amount and manner of inverse weighting required to obtain a consistent estimator of the causal effect ψ_0 can be adapted to the approximation error of Q relative to the true regression.

Collection of Biostatistics
Research Archive

3.2.2 Example II: Semiparametric regression

Let $O = (W, A, Y) \sim P_0$. Assume the model $E_0(Y \mid A, W) - E_0(Y \mid A = 0, W) = A\beta_0 V$ for some $V \subset W$. If the variance of Y, given A, W, only depends on W, then the efficient score of β_0 at P_0 can be represented as

$$D^*(\Pi_0, \theta_0, \beta_0)(O) = (A - \Pi_0(W))(Y - A\beta_0 V - \theta_0(W)),$$

where $\Pi_0(W) = E_0(A \mid W)$, and $\theta_0(W) = E_0(Y \mid A = 0, W)$. For the sake of illustration we will use this simpler representation, but the same double robustness applies to the general efficient influence curve representation as (e.g.) presented in van der Laan and Robins (2003).

Theorem 3 *Suppose $E_0(Y - A\beta_0 V - \theta(W) \mid A, W) = f_0(W(\theta))$ for some function f_0 of $W(\theta)$ where $W(\theta) = \Phi(W, \theta)$ is function of W and θ. Note that this states that $\theta_0(W) - \theta(W) = f_0(W(\theta))$ is only a function of a reduction $W(\theta)$ of W. Let $\Pi_0(\theta)(W) = E_0(A \mid W(\theta))$. Then*

$$E_0 D^*(\Pi_0(\theta), \theta, \beta_0) = 0$$

We also have

$$E_0 D^*(\Pi, \theta_0, \beta_0) = 0$$

Proof. Only the first robustness result needs to be proved. First take the conditional mean, given A, W, which results in the term $E_0(A - \Pi_0(\theta)(W(\theta)))$ $f_0(W(\theta))$. Subsequently, we take the conditional mean, given $W(\theta)$, which proves it equals zero. □

3.3 Construction of collaborative double robust estimators

By using a collaborative estimator $g_n(Q)$ of a $g_0(Q)$ in the set of conditional distributions that conditions on the required function of $Q - Q_0$ (and g_0 itself), one can construct collaborative double robust estimators. For example, one could use the targeted maximum likelihood estimator applied to initial estimator Q_n and using the resulting collaborative estimator $g_n(Q_n)$. One can also use estimating equation methodology, solving for ψ_n in $0 = P_n D^*(\psi, Q_n, g_n(Q_n))$. The formal asymptotic linearity (and thereby asymptotic normality) of such estimators is studied in the next section. Our proposed collaborative targeted maximum likelihood procedure is one particular collaborative double robust targeted maximum likelihood estimator, which also involves updating the initial estimator Q_n beyond the construction of an appropriate $g_n(Q_n)$. However,

UC BERKELEY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

we could also simply have taken $g_n(Q_n)$ from our proposed collaborative targeted maximum likelihood procedure, and still use the targeted maximum likelihood estimator with initial estimator Q_n. In addition, we could also have used our proposed collaborative estimator $g_n(Q_n)$ to solve an estimating equation $0 = P_n D^*(Q_n, g_n(Q_n), \psi) = 0$ in ψ.

Other methods for construction of collaborative estimators $g_n(Q_n)$ are of interest as well. For example, one could consider a collection of one-dimensional fluctuations of Q_n and use maximum likelihood to test these fluctuations. In this manner one can select a dimension reduction involving the X-components that still significantly increase the log-likelihood (or other loss function) beyond the initial fit Q_n. One could then fit g_n by running a machine learning algorithm that only conditions on the selected components. This procedure only uses the initial estimator to obtain a dimension reduction, but from then on it uses an external procedure based on the loss function for g_0.

Given an initial estimator Q_n, another idea of interest for construction of a collaborative estimator $g_n(Q_n)$ is the following. One first constructs a sequence of increasingly nonparametric estimators $\hat{g}_j$ of g_0, $j = 1, \ldots, J$. These estimators could already be based on a dimension reduction based on offset by initial estimator. Given an initial estimator $\hat{Q}$, we select the following estimator of g_0:

$$j_n = \arg\min_j \parallel E_{B_n} P^1_{n,B_n} D^*(\hat{Q}(P^0_{n,B_n}), \hat{g}_j(P_n)) \parallel^2,$$

where B_n denotes a random variable in $\{0,1\}^n$ defining a random split in training sample $\{i : B_n(i) = 0\}$ and validation sample $\{i : B_n(i) = 1\}$, and P^0_{n,B_n}, P^1_{n,B_n} denote the empirical probability distributions of the training and validation sample, respectively. Thus, one selects the estimator that minimizes the Euclidean norm of the cross-validated mean of efficient influence curve at the estimator $\hat{Q}$. If j is too small, then $P_0 D^*(\hat{Q}, g_j)$ will be non zero, so that j_n will always select a large enough j for n tending to infinity. If, on the other hand, j is large enough so that $P_0 D^*(\hat{Q}, g_j) = 0$, then the expectation of $\{P_n D^*(\hat{Q}, g_j)\}^2$ will be equal to its variance which will be increasing in j, so that smaller j's, but larger than the critical one, will be preferred. One now defines as collaborative estimator the estimator $g_n(Q_n) \equiv \hat{g}_{j_n}(P_n)$ indexed by this choice j_n.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

3.4 Estimating the sufficient minimal adjustment covariate from the data

Let $H(g_0, Q - Q_0)$ be the component that needs to be adjusted for in $g_0 = g_0(Q)$. One could estimate this component from the data using appropriate methodology. If, given an arbitrary initial fit g, one would add $H(g, Q - Q_0)$ as main term in a fluctuation model of g, and the fluctuation function is chosen so that the score of the coefficient of this main term at zero equals $D_{CAR}(Q - Q_0, g_0)$, then the MLE-update of g will solve the wished score equation $P_0 D_{CAR}(Q - Q_0, g_0) = 0$. We can refer to $H(g_0, Q - Q_0)$ as the minimal adjustment covariate, needed to obtain the wished collaborative robustness.

The sufficient covariate $H(g, Q - Q_0)$, that is needed to update g, depends on g itself as well, so that, even given an estimate of $Q - Q_0$, the above maximum likelihood update of an initial g does not work. There are two approaches that can be used to deal with this self-dependence of the minimal adjustment covariate for g.

Firstly, one can extract the few components of only $Q - Q_0$, and enforce nonparametric adjustment by these covariates in a fit of $g_0(Q)$ of the censoring mechanism. In this manner, the resulting censoring mechanism estimator will estimate a true conditional distribution that conditions on covariates that imply the value of $H(g_0(Q), Q - Q_0)$. Secondly, one can also only adjust for $H(g^0, Q - Q_0)$ as a main term, given an estimate g^0, and iterate this updating process of g till convergence. In the latter case, as mentioned above, it is assumed that the score of the fluctuation of g implied by this main term extension $H(g, Q - Q_0)$ at zero equals $D_{CAR}(g, Q - Q_0)$. Let's illustrate these two approaches with an example.

For example, in the additive causal effect example with $O = (W, A, Y)$, $\psi_0 = EY(1) - Y(0)$ and Y continuous, we have $H(g_0, Q - Q_0) = \frac{1}{g_0(1|W)} E(Y - Q \mid A = 1, W) + \frac{1}{g_0(0|W)} E(Y - Q \mid A = 0, W)$, and the score of ϵ, at $\epsilon = 0$, of logistic regression model $g_\epsilon(1 \mid W) = 1/(1 + \exp(-C^0(W) - \epsilon H(g, Q - Q_0))$, using an offset $C^0(W)$, is given by $D_{CAR}(g, Q - Q_0) = H(g, Q - Q_0)(A - g(1 \mid W))$.

One can estimate $E(Y - Q \mid A = 1, W)$ and $E(Y - Q \mid A = 0, W)$ with a machine learning algorithm, treating an initial estimate Q as offset (possibly cross-validated to make the offset independent of Y_i). If Y is binary, we would use Q as offset and one could estimate these $Q - Q_0$-components by running a logistic regression with Q as offset. Given this estimate of the two $Q - Q_0$-components that span $H(g_0, Q - Q_0)$, one could now force nonparametric adjustment by these two estimated covariates in the estimate of g_0.

UC BERKELEY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

Alternatively, given an initial estimator g_n^0, one could obtain an estimate H_n^0 by plugging in an estimate of $Q_0 - Q$, and this initial g_n^0. One could now force in this H_n^0 as a main term in g_n^0, resulting in an updated g_n^1. This process is iterated till convergence. In the limit we have that g_n solves $0 = P_n H(g_n, \widehat{Q - Q_0})(A - g_n(1 \mid W))$. Here g_n^0 would already be a collaborative estimator of $g_0(Q)$, such as the one proposed in our collaborative targeted maximum likelihood estimator, so that the collaborative double robustness is preserved (we do not want to only rely on correct specification of $\widehat{Q - Q_0}$, and thereby of this sufficient minimal covariate)

We do not advise starting the above iterative algorithm at a purposely misspecified estimator g_n^0. Instead we want to apply the above iterative algorithm at a collaborative estimator g_n^0, such as the one presented in our template of the C-TMLE in Section 2. For example, after having run the C-TMLE in Section 2, we would carry out a subsequent update of the resulting collaborative estimator g_n by applying the above iterative updating algorithm, starting at $g_n^0 = g_n$, and using an estimator of $Q_0 - Q$. If we would only include this estimate of the sufficient covariate $H(g_0, Q - Q_0)$ in g_n, then the consistency of the estimator ψ_n fully relies on correct estimation of $Q - Q_0$, and thereby on correct estimation of Q_0, and therefore would not utilize the collaborative double robustness of the efficient influence curve. Instead of carrying out a subsequent update of a collaborative estimator g_n using the iterative algorithm, we could incorporate an estimate H_n (or its $(Q - Q_0)$-components) in our proposed template for the collaborative targeted MLE by forcing it in our candidate censoring mechanism estimators.

In the above additive causal risk example, we estimate $H(g, Q - Q_0)$ by plugging in an initial estimate g^0 and $Q - Q_0$, and the iterative adjustment succeeds in its goal as long as the estimate of $Q - Q_0$ is correct, even if (the initial) g is misspecified. One could also use a representation such as $H(g_0, Q - Q_0) = E_{g_0,Q_0}(D_{IPCW}(g_0, Q) \mid A = 1, W) - E_0(D_{IPCW}(g_0, Q) \mid A = 0, W)$, and estimate the two regressions by regressing an $IPCW$-function indexed by g_0, Q on A, W, and evaluate it at $A = 1$ and $A = 0$, respectively. Here $D_{IPCW} = (Y - Q)\{A/g_0(A \mid W) - (1 - A)/g_0(A \mid W)\}$. One could now apply the above iterative updating algorithm to this (non-substitution based) manner of estimating $H(g_0, Q - Q_0)$. As shown in (e.g.) van der Laan and Robins (2003) for monotone censored data structures and causal inference data structures, involving censoring and treatment actions over time, $D_{CAR}(Q - Q_0, g_0)$ does allow such a representation $\sum_j H_j(g_0, Q - Q_0)(A(j) - g_{0j}(1 \mid \mathcal{F}(j))$, where $\mathcal{F}(j)$ represents the history before censoring or treatment $A(j)$, and $H_j(g_0, Q - Q_0) = E_0(D_{IPCW} \mid A(j) = 1, \mathcal{F}(j)) - E_0(D_{IPCW} \mid A(j) = 0, \mathcal{F}(j))$ for some $IPCW$-function. The disadvantage of this approach is that it relies on g_0 representing a

Collection of Biostatistics
Research Archive

true conditional distribution, while in the iterative substitution based approach the main term adjustment at a possibly misspecified g still yields the wished collaborative double robustness.

4 Asymptotic linearity of collaborative double robust TMLE

The collaborative targeted maximum likelihood estimator Q_n^* equals a k_n-th step collaborative targeted maximum likelihood estimator, and thereby equals a targeted maximum likelihood estimator with a starting estimator Q_n^k (e.g., the $k_n - 1$-th collaborative targeted maximum likelihood estimator), and the censoring mechanism estimator $g_n = g_{n\delta_n}$ as selected in the k_n-step, given the collection of candidate estimators $g_{n\delta}$ indexed by δ ranging over an index set.

Thus, just like the targeted maximum likelihood estimator, the collaborative targeted maximum likelihood estimator $\psi_n = \Psi(Q_n^*)$ of ψ_0 solves the efficient influence curve estimating equation

$$0 = P_n D^*(Q_n^*, g_n, \psi_n).$$

For simplicity, we will make the assumption that the efficient influence curve at a $P_{Q,g}$ can be represented as an estimating function in ψ: i.e., the efficient influence curve at P can be represented as $D^*(Q(P), g(P), \psi(Q(P)))$ for some mapping $(Q, g, \psi) \to D^*(Q, g, \psi)$. However, the theorem in this section can be generalized to any efficient influence curve $D^*(Q, g)$ at a data generating distribution $P_{Q,g}$.

It is a reasonable assumption that Q_n^* converges to some element Q^* in the model for Q_0, where Q^* is not necessarily equal to the true Q_0. In addition, let's assume that, for each δ, the δ-specific censoring mechanism estimator $g_{n\delta}$ converges to some $g_{0\delta}$. For example, if δ indicates an adjustment set, then it might be assumed that $g_{n\delta}$ converges to the true conditional distribution, given this δ-specific adjustment set.

For a given Q, we define $\delta(Q)$ as the index δ with entropy $d(\delta)$ minimal and so that

$$P_0 D^*(Q, g_{0\delta(Q)}, \psi_0) = 0.$$

In other words, given the family of adjustments indexed by δ, $\delta(Q)$ represents the minimal adjustment necessary in the censoring mechanism to obtain the collaborative double robustness/unbiased estimating function for ψ_0. It is then a natural assumption that

$$P_0 D^*(Q, g_{0\delta}, \psi_0) = 0 \text{ for each } \delta \text{ with } d(\delta) \geq d(\delta(Q)).$$

bepress
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

In other words, if one uses a more nonparametric estimator of the censoring mechanism than needed (i.e.., than $\delta(Q)$), then one certainly obtains the wished unbiasedness.

We will assume that, as n converges to infinity, then the selected censoring mechanism estimator $g_n = g_{n\delta_n}$ converges to a fixed $g_{0\delta_0}$ representing the limit of a $g_{n\delta_0}$, not necessarily equal to the conditional distribution, given the full X. For notational convenience, we will also denote this limit with g_0.

It is assumed that $d(\delta_0) \geq d(Q^*)$ so that

$$0 = P_0 D^*(Q^*, g_0, \psi_0),$$

which will be the fundamental assumption for asymptotic normality of the C-TMLE. In other words, it is assumed that our collaborative C-TMLE procedure selects a nonparametric enough estimator g_n for the censoring mechanism (in collaboration with Q_n^*) so that the required unbiasedness of the efficient influence curve estimating function is achieved.

To derive the influence curve of $\Psi(Q_n^*)$, the asymptotic linearity theorem below assumes also that the limit of the selected censoring mechanism estimator satisfies

$$P_0 D^*(Q_n^*, g_0, \psi_0) = o_P(1/\sqrt{n}). \tag{3}$$

As a consequence of this assumption (3), the influence curve does not involve a contribution requiring the analysis of a function of Q_n^*. This important simplification of the influence curve allows straightforward calculation of standard errors for the C-TMLE. The assumption (3) requires the limit g_0 to be nonparametric enough w.r.t. the actual estimator Q_n^* so that enough orthogonality is achieved to make the contribution $P_0 D^*(Q_n^*, g_0, \psi_0)$ second order.

Why assumption (3) holds for C-TMLE: We now explain why this assumption is reasonable for the C-TMLE.

Define $g_0(Q_n^*)$ as $g_{0\delta_n^*}$ with $\delta_n^* = \min\{\delta : d(\delta) \geq d(\delta_0), P_0 D^*(Q_n^*, g_{0\delta}, \psi_0) = 0\}$. In other words, $g_0(Q_n^*)$ corresponds with the limit of the least nonparametric estimator (among all estimators more nonparametric than the one identified by δ_0) that still yields the wished unbiasedness of the estimating function at Q_n^*, and it as close as possible to $g_0 = g_{0\delta_0}$.

We note that

$$\begin{aligned} P_0 D^*(Q_n^*, g_0, \psi_0) - D^*(Q^*, g_0, \psi_0) = {} & P_0 D^*(Q_n^*, g_0(Q_n^*), \psi_0) - D^*(Q^*, g_0(Q_n^*), \psi_0) \\ & + R_n, \end{aligned}$$

where R_n is a second order term (like R_{n1} below) involving the difference $Q_n^* - Q^*$ and $g_0(Q_n^*) - g_0$. By definition of $g_0(Q_n^*)$ and the fact that Q_n^*

COBRA
Collection of Biostatistics
Research Archive

converges to Q^*, it is reasonable to assume $g_0(Q_n^*) \to g_0$ as $n \to \infty$. So R_n is a second order term, so that it is reasonable to assume $R_n = o_P(1/\sqrt{n})$.

By definition of $g_0(Q_n^*)$, we do not only have

$$P_0 D^*(Q_n^*, g_0(Q_n^*), \psi_0) = 0,$$

but also that $g_0(Q_n^*)$ is equally or more nonparametric than $g_0(Q^*)$ so that

$$P_0 D^*(Q^*, g_0(Q_n^*), \psi_0) = 0.$$

This implies now that indeed

$$P_0 D^*(Q_n^*, g_0, \psi_0) = o_P(1/\sqrt{n}).$$

Finally, we note that the next theorem can be applied to any collaborative double robust estimator, as discussed in previous section, not only the collaborative double robust targeted maximum likelihood estimator.

Theorem 4 *Let $(Q, g, \psi) \to D^*(Q, g, \psi)$ be a well defined function that maps any possible $(Q, g, \Psi(Q))$ into a function of O. Let $O_1, \ldots, O_n \sim P_0$ be i.i.d, and let P_n be the empirical probability distribution. Let $Q \to \Psi(Q)$ be a d-dimensional parameter, where $\psi_0 = \Psi(Q_0)$ is the parameter value of interest. In the following template for proving asymptotic linearity of $\Psi(Q_n^*)$ as an estimator of $\Psi(Q_0)$, Q_n^* represents the collaborative targeted maximum likelihood estimator, but it can be any estimator.*

Let Q^* denote the limit of Q_n^*. Let g_n be an estimator and g_0 denote its limit.

Assume

Efficient Influence Curve Estimating Equation: $0 = P_n D^*(Q_n^*, g_n, \psi_n)$, *where $\psi_n = \Psi(Q_n^*)$.*

Censoring Mechanism Estimator is Nonparametric Enough:

$$P_0 D^*(Q^*, g_0, \psi_0) = 0.$$

$$P_0 D^*(Q_n^*, g_0, \psi_0) = o_P(1/\sqrt{n}).$$

(Above we show why the latter is indeed a second order term for the C-TMLE.)

Consistency:

$$P_0(D^*(Q_n^*, g_n, \psi_n) - D^*(Q^*, g_0, \psi_0))^2 \to 0 \text{ in probability,}$$

as $n \to \infty$. And the same is assumed if one or two of the triplets (Q_n^*, g_n, ψ_n) is replaced by its limit (Q^*, g_0, ψ_0).

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Identifiability/Invertibility: *$c_0 = -d/d\psi_0 P_0 D^*(Q^*, g_0, \psi_0)$ exists and is invertible.*

Donsker Class: *$\{D^*(Q, g, \Psi(Q)) : Q, g\}$ is P_0-Donsker, where (Q, g) vary over sets that contain (Q_n^*, g_n), (Q^*, g_n), (Q_n^*, g) with probability tending to 1.*

Contribution due to Censoring Mechanism Estimation: *Define the mapping $g \to \Phi(g) \equiv P_0 D^*(Q^*, g, \psi_0)$. Assume $\Phi(g_n) - \Phi(g_0) = (P_n - P_0)IC_{g_0} + o_P(1/\sqrt{n})$ for some mean zero function $IC_{g_0} \in L_0^2(P_0)$.*

Second order terms: *Define second order term*

$$\begin{aligned} R_{n1} &= P_0\{D^*(Q_n^*, g_n, \psi_n) - D^*(Q^*, g_n, \psi_n)\} \\ &\quad - P_0\{D^*(Q_n^*, g_0, \psi_0) - D^*(Q^*, g_0, \psi_0)\}, \end{aligned}$$

and assume $R_{n1} = o_P(1/\sqrt{n})$. Note R_{n1} is a second order term involving difference between $Q_n^* - Q$ and $g_n - g_0$.

Define second order term

$$\begin{aligned} R_{n2} &= P_0\{D^*(Q^*, g_n, \psi_n) - D^*(Q^*, g_0, \psi_n)\} \\ &\quad - P_0\{D^*(Q^*, g_n, \psi_0) - D^*(Q^*, g_0, \psi_0)\}, \end{aligned}$$

and assume $R_{n2} = o_P(1/\sqrt{n})$. Note R_{n2} is a second order term involving differences $g_n - g_0$ and $\psi_n - \psi_0$.

Then, ψ_n is asymptotically linear estimator of ψ_0 at P_0 with influence curve

$$IC(P_0) = c_0^{-1} \left\{ D^*(Q^*, g_0, \psi_0) + IC_{g_0} \right\}.$$

That is,

$$\psi_n - \psi_0 = (P_n - P_0)IC(P_0) + o_P(1/\sqrt{n}).$$

In particular, $\sqrt{n}(\psi_n - \psi_0)$ converges in distribution to a multivariate normal distribution with mean zero and covariance matrix $\Sigma_0 = E_0 IC(P_0) IC(P_0)^\top$.

Proof: The principal equations are $0 = P_n D^*(Q_n^*, g_n, \psi_n)$ and $P_0 D^*(Q^*, g_0, \psi_0) = 0$. So, we have

$$P_0 D^*(Q^*, g_0, \psi_n) - D^*(Q^*, g_0, \psi_0) = -\left\{ P_n D^*(Q_n^*, g_n, \psi_n) - P_0 D^*(Q^*, g_0, \psi_n) \right\}.$$

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Let $c_0 = -\frac{d}{d\psi_0} P_0 D^*(Q^*, g_0, \psi_0)$. Then,

$$\begin{aligned} c_0(\psi_n - \psi_0) + o(\mid \psi_n - \psi_0 \mid) &= (P_n - P_0)D^*(Q^*, g_0, \psi_n) \\ &\quad + P_n\{D^*(Q_n^*, g_n, \psi_n) - D^*(Q^*, g_n, \psi_n)\} \\ &\quad + P_n\{D^*(Q^*, g_n, \psi_n) - D^*(Q^*, g_0, \psi_n)\}. \end{aligned}$$

We denote the three terms on the right with I,II and III, and deal with them separately below.

I: By the Donsker condition, and consistency condition, we have

$$(P_n - P_0)\{D^*(Q^*, g_0, \psi_n) - D^*(Q^*, g_0, \psi_0)\} = o_P(1/\sqrt{n}).$$

Thus, we obtain $(P_n - P_0)D^*(Q^*, g_0, \psi_0) + o_P(1/\sqrt{n})$ as first term approximation. We refer to van der Vaart and Wellner (1996) for this empirical process theorem.

II: We have

$$\begin{aligned} P_n\{D^*(Q_n^*, g_n, \psi_n) - D^*(Q^*, g_n, \psi_n)\} &= \\ & P_n - P_0)\{D^*(Q_n^*, g_n, \psi_n) - D^*(Q^*, g_n, \psi_n)\} \\ & + P_0\{D^*(Q_n^*, g_n, \psi_n) - D^*(Q^*, g_n, \psi_n)\}. \end{aligned}$$

The first term is $o_P(1/\sqrt{n})$ by our Donsker class condition, and consistency condition at Q_n^*, g_n, ψ_n. We also have

$$P_0\{D^*(Q_n^*, g_n, \psi_n) - D^*(Q^*, g_n, \psi_n)\} = P_0\{D^*(Q_n^*, g_0, \psi_0) - D^*(Q^*, g_0, \psi_0) + R_{n1},$$

where

$$\begin{aligned} R_{n1} &= P_0\{D^*(Q_n^*, g_n, \psi_n) - D^*(Q^*, g_n, \psi_n) - D^*(Q_n^*, g_0, \psi_0) - D^*(Q^*, g_0, \psi_0)\} \\ &= o_P(1/\sqrt{n}), \end{aligned}$$

by assumption.

R_{n1} is a second order term involving $Q_n^* - Q^*$ and $(g_n, \psi_n) - (g_0, \psi_0)$. It remains to consider the term $P_0\{D^*(Q_n^*, g_0, \psi_0) - D^*(Q^*, g_0, \psi_0)\}$, which is $o_P(1/\sqrt{n})$ by "Censoring Mechanism is Nonparametric Enough"-assumption.

III: We have

$$\begin{aligned} P_n\{D^*(Q^*, g_n, \psi_n) - D^*(Q^*, g_0, \psi_n) &= \\ & (P_n - P_0)\{D^*(Q^*, g_n, \psi_n) - D^*(Q^*, g_0, \psi_n)\} \\ & + P_0\{D^*(Q^*, g_n, \psi_n) - D^*(Q^*, g_0, \psi_n)\}. \end{aligned}$$

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

The first term is $o_P(1/\sqrt{n})$ by Donsker class condition, and consistency condition at Q_n^*, g_n, ψ_n. We also have

$$P_0\{D^*(Q^*, g_n, \psi_n) - D^*(Q^*, g_0, \psi_n)\} = P_0\{D^*(Q^*, g_n, \psi_0) - D^*(Q^*, g_0, \psi_0)\} + R_{n2},$$

where

$$\begin{aligned} R_{n2} &= P_0\{D^*(Q^*, g_n, \psi_n) - D^*(Q^*, g_0, \psi_n) - D^*(Q^*, g_n, \psi_0) - D^*(Q^*, g_0, \psi_0)\} \\ &= o_P(1/\sqrt{n}), \end{aligned}$$

by assumption. Thus the third term equals $P_0 D^*(Q^*, g_n, \psi_0) - D^*(Q^*, g_0, \psi_0)$, which, by definition, equals $\Phi(g_n) - \Phi(g_0)$. We assumed that $\Phi(g_n) - \Phi(g_0) = (P_n - P_0)IC_{g_0} + o_P(1/\sqrt{n})$. Thus, the third term equals $(P_n - P_0)IC_{g_0} + o_P(1/\sqrt{n})$.

We can thus conclude that

$$\psi_n - \psi_0 = (P_n - P_0)c_0^{-1}\{D^*(Q^*, g_0, \psi_0) + IC_{g_0}\} + o_P(\mid \psi_n - \psi_0 \mid) + o_P(1/\sqrt{n}).$$

This implies $\mid \psi_n - \psi_0 \mid = O_P(1/\sqrt{n})$, and thereby the stated asymptotic linearity. □

4.1 Statistical Inference

If $Q^* = Q_0$, then $IC_{g_0} = 0$, so that the influence curve reduces to the efficient influence curve $D^*(Q_0, g_0, \psi_0)$ at a possibly weakly adjusted g_0. If g_n converges to the fully adjusted conditional distribution, given X, then we know that IC_{g_0} equals minus the projection of $D^*(Q^*, g_0, \psi_0)$ onto the tangent space of the model used by g_n (van der Laan and Robins (2003), Section 2.3.7). We suggest that, even if g_0 is not the fully adjusted censoring mechanism, we will typically still have that $D^*(Q^*, g_0, \psi_0)$ is a conservative influence curve. In other words, if Q_n starts approximating the true Q_0, then the IC_{g_0} contribution gets smaller and smaller, while if Q_n stays away from Q_0, then g_n starts approximating the fully adjusted g_0, in which case, inference based on D^* is conservative. This might explain why we see good coverage in our simulations based on "influence curve" $D^*(Q_n^*, g_n, \psi_n^*)$. If g_n corresponds with a parametric MLE estimator (for a data adaptively selected parametric model), then we propose to use the parametric delta-method to compute the analytic formula for the influence curve IC_{g_0} in order to obtain an accurate influence curve.

One can estimate the covariance matrix $\Sigma = E_0 IC IC^\top$ of the influence curve with the empirical covariance matrix $\Sigma_n = 1/n \sum_{i=1}^n \hat{IC}(O_i)\hat{IC}(O_i)^\top$, and statistical inference can be based on the corresponding mean zero multivariate normal distribution, as usual.

Collection of Biostatistics
Research Archive

4.2 Selection among difference collaborative targeted maximum likelihood estimators

Suppose that we have a set of candidate collaborative targeted maximum likelihood estimators $(\hat{Q}_k^*(P_n), \hat{g}_k(P_n))$, $k = 1, \ldots, K$. Suppose that each of these estimators satisfy the conditions of the theorem. For example, these might be collaborative targeted maximum likelihood estimators as defined in our template, using different initial estimators indexed by k, but the same collaborative estimator for the censoring mechanism *as a function of the data and the initial estimator* (thus still resulting in different realizations if the initial estimators are different). Then $\Psi(\hat{Q}_k^*(P_n))$ is asymptotically linear with influence curve $D^*(Q_k^*, g_{0k}, \psi_0)$, $k = 1, \ldots, K$. We can now select among these candidate C-DR-TMLEs by maximizing the estimated efficiency, as in Rubin and van der Laan (2008).

Specifically, let Ψ be a one-dimensional parameter. We now select the k that minimizes the cross-validated variance of the influence curve:

$$k_n = \arg\min_k E_{B_n} P^1_{n,B_n} D^{*2}(\hat{Q}_k^*(P^0_{n,B_n}), \hat{g}_k(P_{n,B_n}), \psi_n).$$

Thus, we would use the estimator $\psi_n = \Psi(\hat{Q}^*_{k_n}(P_n))$. If Ψ is multidimensional, then one needs to agree on a real valued criterion applied to the covariance matrix of the influence curve, such as the sum of the variances along the diagonal, and minimize over k the criterion of the cross-validated covariance matrix of the k-specific influence curve.

4.3 Irregular C-TMLE and super efficiency

If g_n converges to the fully adjusted $g_0(\cdot \mid X)$ (fully adjusting for X, under CAR) and Q_n^* converges to Q_0, then it follows that ψ_n is asymptotically linear with influence curve equal to the efficient influence curve $D^*(Q_0, g_0, \psi_0)$. So in that case, ψ_n is an asymptotically efficient estimator and thereby also a regular estimator.

Due to the particular way g_n is constructed in response to Q_n, it is easily argued that the collaborative targeted MLE can be an irregular estimator and can be super efficient by achieving an asymptotic variance that is smaller than the variance of the efficient influence curve. In particular, our previous arguments showed that if the initial estimator is a maximum likelihood estimator according to a correctly specified parametric model, then g_n will avoid nonparametric fits, thereby staying away from estimating the fully adjusted g_0 that would result in an efficient estimator in first order. In this case, by

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

the above theorem, the influence curve of ψ_n will be equal to $D^*(Q_0, g_0, \psi_0)$, using a non-fully adjusted g_0, so that the variance of the influence curve will be smaller than the variance of the efficient influence curve that involves a fully adjusted g_0.

The super efficiency may have very attractive features in practice. For example, there might be a covariate that is very predictive of censoring/treatment, but have no relation to the outcome. The C-TMLE will now decide to not adjust for this covariate at all in the selected censoring mechanism, and as a consequence, it might achieve the efficiency bound for the data structure excluding this covariate, but still assuming CAR, so that the C-TMLE will have smaller asymptotic variance than the efficiency bound. The resulting super efficient estimator not only shows improved precision, but also yields more reliable confidence intervals, by avoiding heavily non-robust (and harmful) operations. In most practical scenarios, such a covariate will still have a weak link with the outcome. In this case, for very large sample sizes, the C-TMLE will adjust for this covariate and thereby only be asymptotically efficient, but it will still behave as a super efficient estimator for practical sample sizes, by not adjusting for this covariate. That is, it invests in effective bias reduction focussing on covariates that are still predictive of the outcome, taking into account the already included initial estimator. This behavior is completely compatible with an estimator that aims to minimize mean squared error of the estimator of the target parameter, and certainly avoids steps that both increase bias as well as variance.

Finally, we remark that in simulations in which Q_n converges fast to the true Q_0, g_n seems to have a temptation to converge to a random choice g_0 that is beyond the required minimal censoring mechanism with probability 1. That is, likelihood based cross-validation might over-select the adjustment in the censoring mechanism relative to the minimal adjustment, and the amount of over-selection remains random (but small) for large sample sizes (this is a known property of cross-validation). This naturally results in an irregularity of the estimator. Simulations have not shown practical problems for statistical inference, but this remains an area of study.

5 Targeted loss functions implied by efficient influence curve

The template of the collaborative targeted maximum likelihood estimators is based on 1) a log-likelihood loss function (i.e., same loss function that is maximized at targeted maximum likelihood step) to select among candidate tar-

COBRA
Collection of Biostatistics
Research Archive

geted maximum likelihood estimators indexed by increasingly nonparametric estimators of censoring mechanism, and 2) a preferred loss function to compare targeted maximum likelihood estimators using different censoring mechanism estimators, in order to build these candidate censoring mechanism estimators. One can also use a preferred loss function to select among different candidate collaborative targeted maximum likelihood estimators (e.g., indexed by different initial estimators).

In this section we propose targeted loss functions implied by the efficient influence curve of Ψ. Firstly, the log-likelihood can be replaced by a penalized log-likelihood that is sensitive to sparse data bias w.r.t. target parameter, as defined in next subsection This penalized log-likelihood can also play the role of the preferred loss function. In the second subsection we propose as preferred loss function the cross-validated variance of the efficient influence curve, relying on an overall collaborative estimator of censoring mechanism w.r.t. an initial estimator, or a candidate specific collaborative estimator. In the last subsection, we utilize the mean of the efficient influence curve as a criterion to generate a targeted loss function for Q that incorporates a sequence of increasingly nonparametric estimators of the censoring mechanism.

5.1 The MSE-penalized cross-validated log-likelihood

In the C-DR-TMLE we applied loglikelihood based cross-validation to select among different targeted maximum likelihood estimators, indexed by different censoring mechanism estimators. We propose here a penalized log-likelihood criterion that results in robust estimators in the context of sparse data w.r.t. the parameter of interest.

Consider candidate (e.g., collaborative) targeted maximum likelihood estimators $P_n \to \hat{Q}^*_\delta(P_n)$ of the true $Q_0 \in \mathcal{M}$, targeting a parameter $\psi_0 = \Psi(Q_0)$, indexed by δ. Our proposed criterion for selecting δ is

$$\delta_n = \underset{\delta}{\operatorname{argmax}}\, E_{B_n} P^1_{n,B_n} \log \hat{Q}^*_\delta(P^0_{n,B_n}) - MSE(P_n)(\delta),$$

where the first term is the cross-validated log-likelihood for the candidate estimator $\hat{Q}^*_\delta(P_n)$, and $MSE(P_n)(\delta)$ is an estimator of the mean squared error (variance plus bias-squared) of the substitution estimator $\hat{\Psi}(\hat{Q}^*_\delta(P_n))$ as an estimator of its δ-specific limit (thus ignoring asymptotic bias). The $MSE(P_n)(\delta)$ is possibly appropriately scaled relative to the log-likelihood term. The sole motivation for the proposed additional penalty term is to make the criterion more targeted towards ψ_0, while still preserving the log-likelihood as the dominant term in regular situations: i..e, asymptotically, the penalty is negligible (in regular situations the MSE behaves as $1/n$).

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

5.1.1 Variance of targeted maximum likelihood estimator relative to its δ-limit

If the target parameter cannot be reasonably identified from the data the log-likelihood of the targeted maximum likelihood estimator is not sensitive enough to such a singularity: in fact, on many occasions this just means that the targeted maximum likelihood algorithm will be ineffective (i.e., the maximum likelihood fluctuations get too noisy) so that in essence the log-likelihood of the initial estimator drives the selection.

Therefore it is crucial that the log-likelihood terms are penalized by a term that blows up (in the negative direction) for δ-values for which the variance (or bias, addressed in next subsection) of the targeted maximum likelihood estimator $\Psi(\hat{Q}^*_\delta(P_n))$ relative to its limit $\psi_0(\delta) = \Psi(\hat{Q}^*_\delta(P_0))$ gets large. Since we can derive the influence curve of the targeted maximum likelihood estimator $\Psi(\hat{Q}^*_\delta(P_n))$ as an estimator of $\psi_0(\delta)$, this variance can be estimated with the variance of this influence curve at this targeted maximum likelihood estimator $\hat{Q}^*_\delta(P_n)$. As follows from the study of TMLE in van der Laan and Rubin (2006) one can often use as approximate influence curve the efficient influence curve $D^*(Q^*_\delta, g_\delta)$ at the limit of the targeted maximum likelihood estimator $(\hat{Q}^*_\delta, \hat{g}_\delta)$, which simplifies the penalty while it remains equally effective.

We first define the cross-validated covariance matrix

$$\frac{\Sigma(P_n)(\delta)}{n} = \frac{1}{n} E_{B_n} P^1_{n,B_n} \left\{ D^*(\hat{P}^*_\delta(P^0_{n,B_n})) D^*(\hat{P}^*_\delta(P^0_{n,B_n}))^\top \right\}.$$

For example, if the target parameter is 1-dimensional (i.e., $d = 1$), then we have

$$\frac{\sigma^2(P_n)(\delta)}{n} = \frac{1}{n} E_{B_n} P^1_{n,B_n} \left\{ D^*(\hat{P}^*_\delta(P^0_{n,B_n})) \right\}^2.$$

For example, one can define the variance term of the MSE as

$$\sigma^2(P_n)(\delta) = a \Sigma(P_n)(\delta) a^\top,$$

for a user supplied vector a, so that $\sigma^2(P_n)/n$ represents the variance estimate of the estimator of $a^\top \psi_0(\delta)$.

Our proposal presented in the MSE-subsection below will actually have the form

$$\sigma^2(P_n)(\delta) = \sum_{j=1}^{d} a_j^\top \Sigma(P_n)(\delta) a_j, \tag{4}$$

where a_j are the row vectors of the square root of a user supplied matrix such as the inverse of the the correlation matrix of $\Sigma(P_n)(\delta)$.

COBRA
Collection of Biostatistics
Research Archive

5.1.2 Bias of targeted maximum likelihood estimator relative to its δ-limit

We might also wish to estimate the bias of the targeted maximum likelihood estimator $\Psi(\hat{Q}^*_\delta(P_n))$ relative to its limit $\psi_0(\delta)$ (even though in most applications the variance appears to drive the MSE term). For example, this could be done with the bootstrap:

$$E_{P_n}\left\{\Psi(\hat{Q}^*_\delta(P_n^\#)) - \Psi(\hat{Q}^*_\delta(P_n))\right\},$$

where $P_n^\#$ represents the empirical distribution of a bootstrap sample $O_1^\#,\ldots,O_n^\#$ from the empirical distribution P_n. However, this would be much too computer intensive in many applications in which the targeted maximum likelihood estimator involves data adaptive model or algorithm selection. By noting that a bootstrap sample corresponds on average with 2/3 of the n observations, the following analogue bias estimate can be viewed as an approximation of this bootstrap bias that only requires 3 times applying the targeted maximum likelihood estimator to a sample of size $n * 2/3$:

$$B(P_n)(\delta) = E_{B_{n3}}\left\{\Psi(\hat{Q}^*_\delta(P^0_{n,B_n})) - \Psi(\hat{Q}^*_\delta(P_n)\right\},$$

where B_{n3} denotes the 3-fold cross-validation scheme.

If $d = 1$, then we will add to the variance term in the previous section the squared bias $B(P_n)^2$ to create a MSE-term. If $d > 1$, then in our proposal below we will construct an appropriate function of $B(P_n)$ representing the analogue of the variance term (4):

$$b(P_n)^2(\delta) \equiv \sum_j (a_j^\top B(P_n)(\delta))^2.$$

Additional rationale behind bias term: To provide further understanding of this kind of bias estimate $B(P_n)$, we note the following. Let $\hat{\Psi}(P_n)$ be an estimator of its target $\hat{\Psi}(P_0)$, where it plays the role of the δ-specific targeted maximum likelihood estimator $\Psi(\hat{Q}^*_\delta(P_n))$. The fundamental assumption allowing statistical inference for $\hat{\Psi}(P_0)$ is the assumption of asymptotic linearity:

$$\hat{\Psi}(P_n) - \hat{\Psi}(P_0) = (P_n - P_0)D(P_0) + R(P_n), \tag{5}$$

where $D(P_0)$ is the influence curve of the estimator, and $R(P_n)$ is the remainder. The asymptotic linearity assumption now assumes that $R(P_n) = o_P(1/\sqrt{n})$.

CORA
Access to CURRENT
Collection of Biostatistics
Research Archive

The representation (5) of the mapping $P_n \to \hat{\Psi}(P_n)$ implies for any cross-validation scheme B_n

$$\begin{aligned} B(P_n) &= E_{B_n}\hat{\Psi}(P^0_{nB_n}) - \hat{\Psi}(P_n) \\ &= E_{B_n}\left\{\hat{\Psi}(P^0_{nB_n}) - \hat{\Psi}(P_0)\right\} - \left\{\hat{\Psi}(P_n) - \hat{\Psi}(P_0)\right\} \\ &= E_{B_n}\left\{(P^0_{nB_n} - P_0)D(P_0) + R(P^0_{n,B_n})\right\} \\ &\quad - \left\{(P_n - P_0)D(P_0) + R(P_n)\right\} \\ &= E_{B_n}R(P^0_{n,B_n}) - R(P_n), \end{aligned}$$

where we use that $E_{B_n}P^0_{n,B_n}D(P_0) = P_nD(P_0)$. Thus, our proposed bias estimate $B(P_n)$ *equals*, for any cross-validation scheme, an average difference of the remainder applied to a subsample of size $n(1-p)$ and the full sample of size n. Therefore, one can conclude that this term will be very sensitive to a large remainder (e.g., second order terms) in the asymptotic linearity expansion (5).

5.1.3 MSE of targeted maximum likelihood estimator relative to its δ-limit

If $d = 1$, then we define the MSE term as

$$MSE(P_n)(\delta) = \frac{\sigma^2(P_n)(\delta)}{n} + B(P_n)^2.$$

If $d > 1$, then we assume that we are provided with a user-specified $d \times d$ symmetric positive definite matrix ρ, so that the square root of this matrix $\rho^{1/2}$ exists. Our MSE term will represent the expectation of the Euclidean norm of $\rho^{1/2}(\hat{\Psi}-\psi)$, or equivalently, the expectation of $(\hat{\Psi}-\psi)^\top\rho(\hat{\Psi}-\psi)$. One concrete proposal is to set $\rho^{1/2}$ equal to the square root of the inverse of an estimate of the correlation matrix of the asymptotic covariance matrix of $\sqrt{n}(\hat{\Psi} - \psi)$, so that the linearly transformed vector has uncorrelated components.

Let a_j be the j-th row of the matrix $\rho^{1/2}$, $j = 1, \ldots, d$. The wished MSE term is now the sum of the MSEs of the linear combination $a_j^\top\hat{\Psi}$. Therefore, the MSE term is represented as

$$MSE(P_n)(\delta) = \frac{1}{n}\sum_j a_j^\top\Sigma(P_n)(\delta)a_j + n\left\{a_j^\top B(P_n)(\delta)\right\}.$$

This is equivalent to defining a variance term

$$\frac{\sigma^2(P_n)(\delta)}{n} = \frac{1}{n}\sum_j a_j^\top\Sigma(P_n)(\delta)a_j,$$

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

a bias term

$$b(P_n)(\delta) = \sum_j \left\{ a_j^\top B(P_n)(\delta) \right\},$$

and defining

$$MSE(P_n)(\delta) = \frac{\sigma^2(P_n)(\delta)}{n} + \{b(P_n)(\delta)\}^2.$$

5.2 Targeted loss functions relying on a collaborative estimate of censoring mechanism

Let $D^*(Q_0, g_0, \psi_0)$ be the efficient influence curve at $dP_0 = Q_0 g_0$ for the parameter $\Psi : \mathcal{M} \to \mathbb{R}$. Consider a set of estimators $\hat{Q}^k(P_n)$ that are all more nonparametric than an initial estimator $\hat{Q}^0(P_n)$. Let $\hat{g}^0$ be a collaborative estimator, relative to the initial estimator $\hat{Q}^0$, so that $P_0 D^*(Q^{0*}, g^0, \psi_0) = 0$. Since, $\hat{Q}^k$ is more nonparametric than $\hat{Q}^0$, it is reasonable to assume that $\hat{g}^0$ is also a collaborative estimator for $\hat{Q}^k$, so that $P_0 D^*(Q^{k*}_{g^0}, g^0, \psi_0) = 0$, where now $Q^{k*}_{g^0}$ denotes the limit of the targeted maximum likelihood estimator of $\hat{Q}^k$ using $\hat{g}^0$.

We can use as criterion for selection among $\hat{Q}^k$, the cross-validated estimated variance of $D^*(Q^{k*}_{g^0}, g^0, \Psi(Q^{k*}_{g^0}))$, where $Q^{k*}_{g^0}$ denotes the limit of the targeted maximum likelihood estimator of $\hat{Q}^k$ using $\hat{g}^0$. Of course, this is also a selection among the collaborative targeted maximum likelihood estimators $\hat{Q}^{k*}_{\hat{g}^0}$.

By our asymptotic linearity theorem, this selection among $\hat{Q}^k$ corresponds with minimizing the variance of the influence curve of the (collaborative) targeted maximum likelihood estimators $\Psi(\hat{Q}^{k*}_{\hat{g}^0})$ based on initial estimator $\hat{Q}^k$ using the collaborative estimator $\hat{g}^0$. The crucial assumption is that $\hat{g}^0$ is indeed a collaborative estimator estimating a true g_0 that involves enough adjustment, so that $P_0 D^*(Q^{k*}_{g^0}, g^0, \psi_0) = 0$. However, this assumption is needed for construction of estimators of ψ_0, and is already relying on weaker assumption than double robustness.

This criterion can now be used as the preferred loss function in our template for the collaborative targeted maximum likelihood estimator to build the candidate censoring mechanism estimators. It will require obtaining a single collaborative estimator $\hat{g}^0$ after having obtained the initial estimator: for example, one could carry out a dimension reduction based on improved fits relative to initial estimator, and estimate it with a super learner only adjusting for the selected variables.

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

We already discussed that the same criterion can be used to select among any set of candidate collaborative targeted maximum likelihood estimators $(\hat{Q}^{k*}_{\hat{g}^k}, \hat{g}^k)$, relying on collaborative estimators $\hat{g}^k$, so that $P_0 D^*(\hat{Q}^{k*}, g^k, \psi_0) = 0$:

$$k_n = \arg\min_k E_{B_n} P^1_{n,B_n} D^{*2}(\hat{Q}^{k*}(P^0_{n,B_n}), \hat{g}^k(P^0_{n,B_n}), \Psi(\hat{Q}^{k*}(P_n)).$$

However, the above only relies on a single collaborative estimator $\hat{g}^0$, and is therefore particularly suitable for building the second stage within the collaborative targeted MLE template.

5.3 Targeted loss functions incorporating a sequence of increasingly nonparametric estimators of censoring mechanism

Consider an initial estimator $\hat{Q}$. We wish to develop a criterion to select among candidate estimators of Q_0 that are more nonparametric than $\hat{Q}$, such as estimators using $\hat{Q}$ as offset. A special application of the loss function presented in this subsection is that $\hat{Q}$ is empty. Consider a sequence of increasingly nonparametric estimators $\hat{g}_j$, $j = 1, \ldots, J$, of g_0, which could be based on $\hat{Q}$: For example, they might be based on a dimension reduction which has as goal to estimate $Q_0 - \hat{Q}$. Consider the following hypothetical criterion for a candidate Q:

$$Q \to \sum_{j=1}^{J} w(j) \parallel P_0 D^*(Q, \hat{g}_j(P_0)) \parallel^2,$$

for a list of weights/scalars $(w(j) : j = 1, \ldots, J)$, where $\| \cdot \|$ denotes a possibly weighted Euclidean norm applied to the vector of the form $P_0 D^*(Q, g)$.

Firstly, if $Q = Q_0$, then the criterion is minimized. In addition, let $j_0(Q)$ be such that $P_0 D^*(Q, \hat{g}_j(P_0)) = 0$ for all $j \geq j_0$. Since this mean zero property holds if $\hat{g}_j$ solves a score implied by $Q - Q_0$, it follows that the closer Q is to Q_0 the more j-specific terms within this sum are close to zero. Finally, by the fact that D^* is a gradient of the path-wise derivative of the target parameter Ψ, for $j \geq j_0$, we have that $P_0 D^*(Q, \hat{g}_j(P_0)) \approx \Psi(Q) - \psi_0$ (either exact, or in first order), which shows that this criterion also targets ψ_0 (see van der Laan and Robins (2003), Section 1.4).

The cross-validated analogue of this criterion is given by

$$D_n(\hat{Q}^k) \equiv E_{B_n} \sum_{j=1}^{J} w(j) \parallel P^1_{n,B_n} D^*(\hat{Q}^k(P^0_{n,B_n}), \hat{g}_j(P_n)) \parallel^2 .$$

COBRA
A BEPRESS
Collection of Biostatistics
Research Archive

If ψ_0 is one dimensional, a sensible choice of weights are given by the inverse of the variance of the empirical means, so that it downweights the noisy j-specific signals:

$$w(j)^{-1} = P^1_{n,B_n}\{D^*(\hat{Q}^k(P^0_{n,B_n}), \hat{g}_j(P_n))\}^2,$$

and makes the criterion $D_n()$ unit free. Analogues of such weighting can be obtained for multi-dimensional ψ_0 as well, and are recommended.

One can add this criterion to a valid loss function such as the log-likelihood criterion $L(\hat{Q}^k)$, giving a more targeted loss function

$$L^*(\hat{Q}^k) \equiv L(\hat{Q}^k) + D_n(\hat{Q}^k).$$

The additional term D_n preserves the validity of the loss function L (i.e., its minimum still identifies Q_0), while it makes the selection targeted towards ψ_0. One can add both D_n as well as the above presented MSE-penalty, where the latter is asymptotically negligible but important in sparse data (w.r.t. ψ_0) situations.

6 Example: Targeted maximum likelihood estimation of the marginal structural model

Suppose we observe $O = (W, A, Y = Y(A))$, where W are baseline covariates, A is a discrete treatment, and Y is a subsequently measured outcome. It is assumed that A is realized in response to the realization of W, and Y is realized in response to both W and A. The full data structure on the experimental unit is $X = (W, (Y(a) : a))$, so that A represents the missingness variable for the missing data structure O on X. We assume the randomization assumption: $g_0(a \mid X) = P_0(A = a \mid X) = P_0(A = a \mid W)$.

Consider a marginal structural model for the full data distribution

$$E_0(Y(a) \mid V) = m(a, V \mid \beta_0)$$

that models the causal effect of a treatment intervention $A = a$ on the outcome Y. For example, one might assume a simple linear model $m(a, V \mid \beta_0) = \beta_0(a, V, aV)$.

Since it is often unreasonable to assume such a parametric form, but such parametric forms can still provide very meaningful projections of the true causal curve, we consider the nonparametric extensions of the parameter β_0:

$$\Psi_h(P_0) = \underset{\beta}{\operatorname{argmin}}\, E_{P_0} \sum_a h(a, V)(Q_0(a, W) - m(a, V \mid \beta))^2,$$

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

where $Q_0(a,w) = E_0(Y \mid A = a, W)$. We have that $\Psi_h(P_0)$ represents a projection of $E_0(Y(a) \mid V)$ onto the working model $m(\mid \beta_0)$. Specifically,

$$\Psi_h(P_0) = \underset{\beta}{\operatorname{argmin}}\ E_{P_0} \sum_a h(a,V) \left(E(Y(a) \mid V) - m(a, V \mid \beta)\right)^2.$$

In particular, if $E_0(Y(a) \mid V) = m(a, V \mid \beta_0)$, then for each h we have $\Psi_h(P_0) = \beta_0$. Without the randomization assumption and consistency assumption, we can interpret $\Psi_h(P_0)$ as the same projection, but $E(Y(a) \mid V) = E(E(Y \mid A = a, W) \mid V)$ now represents a (non-causal) dose response curve of the effect of A on Y that controls for the measured confounders W.

We note that this nonparametric extension only depends on P_0 through the conditional mean of Y, given A, W, and the marginal distribution of W. For simplicity, we will also use the notation $\Psi_h(Q_0)$, where Q_0 now denotes both the marginal distribution of W and the conditional distribution of Y, given A, W.

The efficient estimating function for this nonparametric extension Ψ_h of β_0 is given by:

$$\begin{aligned} D_h(P_0)(O) &= \frac{h_1(A,V)}{g_0(A \mid W)}(Y - Q_0(A,W)) \\ &\quad + \sum_a h_1(a,V)(Q_0(a,W) - m(a, V \mid \Psi_h(P_0))), \end{aligned}$$

where $h_1(a,V) = h(a,V)\frac{d}{d\psi} m(a, V \mid \psi)$. We will assume that $h_1(A,V) = d/d\psi_0 m(A, V \mid \psi_0) h(A,V)$ is chosen so that h_1 does not depend on ψ_0, which is easily arranged for the case that m is linear in ψ and that m is logistic linear. For example, if we use the linear form $(1, a, V, aV)$, then, if m is linear, then we can choose $h_1 = d/d\psi_0 m(A, V \mid \psi_0) g(A \mid V) = ((1, A, V, AV) g(A \mid V)$, and, if m is logistic linear, then we can choose $h_1 = d/d\psi_0 m(A, V \mid \psi_0)/(m(1 - m)(A, V \mid \psi_0)) g(A \mid V)$, which equals $(1, A, V, AV) g(A \mid V)$, again.

Let $D_h^*(P_0) = -c_0^{-1} D_h(P_0)$ be the corresponding efficient influence curve obtained by standardizing the efficient estimating function by the negative of the inverse of the derivative matrix $c_0 = d/d\psi_0 E_0 D_h(P_0)$ (noting that $D_h(P_0)$ can indeed be viewed as a function in ψ_0).

If Y is continuous and we use a normal error regression model as a working model, then a targeted maximum likelihood estimator of ψ_{h0} can be obtained by adding to an initial estimator $Q^0(A,W)$ of $E_0(Y \mid A, W)$ the d-dimensional ϵ-extension $\epsilon C_h(g)(A,W)$, where

$$C_h(g)(A,W) = \frac{h_1(A,V)}{g(A \mid W)},$$

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

for some fit g of g_0, and fitting ϵ with maximum likelihood (i.e., least squares estimation) using Q^0 as offset. The resulting update $Q^1(A,W)$ is now a first step targeted maximum likelihood estimator. It is also the actual targeted maximum likelihood estimator, since iteration is not resulting in further updates (the clever covariate does not change at a targeted MLE update). One estimates the distribution of W with the empirical distribution. The estimate Q^1 and the empirical distribution of W now yields a substitution estimate of the target parameter ψ_{h0}. If Y is binary, the same $\epsilon C_h(g)(A,W)$ is added on the logit scale, and ϵ is fitted with maximum likelihood estimation. Again, the targeted maximum likelihood estimator converges in one step.

6.1 Penalized log-likelihood for candidate treatment mechanism fits

Let $\hat{Q}(P_n)$ be an initial regression estimator of $Q_0 = E_0(Y \mid A,W)$. For a given $P_n \to \hat{g}(P_n)$, let $\hat{Q}^*_{\hat{g}}(P_n)$ be the targeted maximum likelihood estimator corresponding with the covariate $C_h(\hat{g}(P_n))$. Let B_n be a cross-validation scheme, and let P^1_{n,B_n} and P^0_{n,B_n} be the empirical distributions of the validation and training sample, respectively, as identified by $B_n \in \{0,1\}^n$. Let

$$\hat{\Sigma}_{CV}(P_n)(\hat{g}) = E_{B_n} P^1_{n,B_n} D^*_h(\hat{Q}^*_{\hat{g}}(P^0_{n,B_n}), \hat{g}(P^0_{n,B_n}))^2$$

be the cross-validated estimate of the covariance matrix of the efficient influence curve at the estimator $\hat{Q}$ and $\hat{g}$. We also consider the empirical estimate of this covariance matrix

$$\hat{\Sigma}(P_n)(\hat{g}) = P_n D^*_h(\hat{Q}^*_{\hat{g}}(P_n), \hat{g}(P_n))^2.$$

Let

$$\hat{B}(P_n)(\hat{g}) = E_{B_n} \hat{\Psi}_{\hat{g}}(P^0_{n,B_n}) - \hat{\Psi}_{\hat{g}}(P_n)$$

be the bias estimator for the targeted maximum likelihood estimator $\hat{\Psi}_{\hat{g}}(P_n) = \Psi_h(\hat{Q}^*_{\hat{g}}(P_n))$ obtained by plugging in $\hat{Q}^*_{\hat{g}}(P_n)$ in the parameter mapping Ψ. Here we can use three-fold cross-validation as choice for B_n.

We will penalize the log-likelihood with the an estimate of the following average mean squared error

$$\frac{1}{n}\sum_{i=1}^{n} E\left(m(a_i, v_i \mid \hat{\Psi}_{\hat{g}}(P_n))) - m(a_i, v_i \mid \hat{\Psi}_{\hat{g}}(P_0))\right)^2.$$

This mean squared error can be decomposed as $1/n\sum_{i=1}^n \mathrm{Var}(m(a_i, v_i \mid \hat{\Psi}_{\hat{g}}(P_n)))$ and $1/n\sum_{i=1}^n \mathrm{Bias}^2(m(a_i, v_i \mid \hat{\Psi}_{\hat{g}}(P_n)))$. The variance terms of this mean

COBRA
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

squared error can be estimated by

$$\frac{\sigma_i^2(\hat{g})}{n} \equiv \frac{z(a_i, v_i)^\top \hat{\Sigma}(P_n)(\hat{g}) z(a_i, v_i)}{n},$$

where

$$z(a_i, v_i) = \left. \frac{d}{d\beta} m(a_i, v_i \mid \beta) \right|_{\beta = \hat{\Psi}_{\hat{g}}(P_n)}.$$

We keep open the option that one uses either the cross-validated covariance matrix $\hat{\Sigma}_{CV}(P_n)$ or the empirical covariance matrix $\hat{\Sigma}(P_n)$.

The bias terms of this mean squared error can be estimated as

$$B_i(\hat{g}) \equiv E_{B_n} m(a_i, v_i \mid \hat{\Psi}_{\hat{g}}(P^0_{n,B_n})) - m(a_i, v_i \mid \hat{\Psi}_{\hat{g}}(P_n)).$$

If m is linear in β, then this reduces to

$$B_i(\hat{g}) = m(a_i, v_i \mid B(P_n)).$$

Thus, we obtain the following mean squared error estimate for the targeted maximum likelihood estimator $\hat{\Psi}_{\hat{g}}(P_n)$ for a given g-estimator:

$$\widehat{MSE}(P_n)(\hat{g}) \equiv \frac{1}{n} \sum_{i=1}^{n} \left\{ \frac{\sigma_i^2(\hat{g})}{n} + B_i(\hat{g})^2 \right\}.$$

We suggest that the penalized log-likelihood could also only be penalized by the empirical variance component of the MSE. Therefore, we also define

$$\sigma^2(P_n)(\hat{g}) \equiv \frac{1}{n} \sum_{i=1}^{n} \frac{\sigma_i^2(\hat{g})}{n}.$$

Consider now the following two penalized log-likelihood criterions for $\hat{g}$, given the initial estimator $\hat{Q}^0$:

$$L(\hat{g} \mid \hat{Q}^0) = \frac{1}{n} \sum_{i=1}^{n} (Y_i - \hat{Q}^*_{\hat{g}}(P_n)(W_i, A_i))^2 + \widehat{MSE}(P_n)(\hat{g}),$$

or

$$L(\hat{g} \mid \hat{Q}^0) = \frac{1}{n} \sum_{i=1}^{n} (Y_i - \hat{Q}^*_{\hat{g}}(P_n)(W_i, A_i)))^2 + \sigma^2(P_n)(\hat{g}).$$

For Y binary, the RSS is replaced by the log-likelihood of Y, given A, W.

We will use the penalized log-likelihood as loss function to build the candidate treatment mechanism estimators, according to the template for the collaborative targeted MLE in Section 2.

UC Berkeley Collection of Biostatistics Research Archive

6.2 Algorithm for estimating the treatment mechanism based on penalized log-likelihood

Given any candidate adjustment set $W^* \subset W$, let an estimator $\hat{g}(P_n)(W^*)$ of $g_0(A \mid W^*)$ be specified. This allows us to define a criterion in candidate adjustment sets W^*, given the current estimator $\hat{Q}$:

$$L(W^* \mid \hat{Q}) \rightarrow L(\hat{g}(P_n)(W^*) \mid \hat{Q}).$$

Thus, one can evaluate/score any given adjustment set W^* with $L(W^* \mid \hat{Q})$.

Given $\hat{Q}$, one can now use this empirical criterion in adjustment sets to construct an estimator of $g_0(\hat{Q})$ with a greedy type algorithm maximizing over a set of candidate adjustment sets. One starts with the empty adjustment set and selects the best addition move among a set of candidate addition moves based on the criterion. One iterates this process until there does not exist an addition move that improves the criterion. More aggressive greedy algorithms can be implemented as well, as with any machine learning algorithm that is based on iterative local maximization of an empirical criterion. One could apply this algorithm to candidate adjustments sets that have a certain size or entropy for the corresponding $\hat{g}(P_n)(W^*)$.

Alternatively, one creates a sequence of nested (increasing in size) adjustment sets W_j^*, $j = 1, \ldots, J$, for each W_j^* one obtains a particular estimator $\hat{g}_j(P_n)$ of $g_0(A \mid W_j^*)$ (e.g., using super learning), and maximizes the penalized log-likelihood criterion over all these J adjustment sets.

In our algorithm in the next subsection defining the sequence of C-TMLEs we apply this greedy algorithm to candidate estimators that are more non-parametric than the selected estimator of g_0 in the previous step.

6.3 Iteration to obtain sequence of targeted maximum likelihood estimators indexed by increasingly non-parametric estimators of treatment mechanism

Given an initial estimator $\hat{Q}$ of $E(Y \mid A, W)$ and a corresponding estimator $\hat{g}(\hat{Q})$ defined above, sometimes denoted with $\hat{g}$, we define a resulting targeted maximum likelihood estimator

$$\hat{Q}_{\hat{g}}^*(P_n) = \hat{Q}(P_n) + \epsilon_n h(\hat{g}(\hat{Q})(P_n)),$$

where ϵ_n is the least squares estimator of the regression coefficient ϵ treating $\hat{Q}(P_n)$ as offset and $h(\hat{g}(\hat{Q})(P_n))$ as covariate. We can define this as a first step

COBRA
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

targeted maximum likelihood estimator based on an initial $\hat{Q}(P_n)$, and corresponding censoring mechanism estimator $\hat{g}(\hat{Q})$). Let's denote this operation as:

$$\hat{Q}^1(P_n) = \hat{Q}(P_n) + \epsilon_n^1 h(\hat{g}(\hat{Q})(P_n)).$$

This process can now be iterated by replacing $\hat{Q}(P_n)$ by this update $\hat{Q}^1(P_n)$:

$$\hat{Q}^2(P_n) = \hat{Q}^1(P_n) + \epsilon_n^2 h(\hat{g}(\hat{Q}^1)(P_n)),$$

where we require that the next censoring mechanism estimator $\hat{g}(\hat{Q}^1)(P_n)$ is obtained with the same algorithm as presented in above subsection, but now maximizing over candidate estimators that are more nonparametric than the previously obtained $\hat{g}(\hat{Q})(P_n)$.

In general, we define the k-th step of this targeted maximum likelihood estimator as

$$\hat{Q}^k(P_n) = \hat{Q}^{k-1}(P_n) + \epsilon_n^k h(\hat{g}(\hat{Q}^{k-1})(P_n)),$$

where $\hat{g}(\hat{Q}^{k-1})(P_n)$ involves maximizing over more nonparametric candidate estimators than $\hat{g}(\hat{Q}^{k-2})(P_n)$.

This algorithm results in a sequence of k-th step collaborative targeted maximum likelihood estimators $\Psi(\hat{Q}^k(P_n))$ of ψ_0, and corresponding increasingly nonparametric censoring mechanism estimators $\hat{g}^k(P_n)$ (i.e., $\hat{g}(\hat{Q}^{k-1})(P_n)$ in above notation), $k = 1, \ldots, K$.

We could also have defined candidate targeted maximum likelihood estimators using a forward selection algorithm each time finding the best next term to add in the treatment mechanism, so that k denotes the number of terms included in $\hat{g}^k$. In that case, k corresponds with the size of the model for $\hat{g}^k$, and the targeted MLE step would be carried out when needed in order to guarantee an increase in either the penalized or non-penalized log-likelihood fit of $\hat{Q}^k$, as described in our general template for collaborative targeted maximum likelihood estimation in Section 2.

6.4 Collaborative TMLEs

If the initial estimator $\hat{Q}$ is indexed by a choice δ_1 and the choice of algorithm $\hat{g}(\hat{Q})$ is indexed by a δ_2, then, for each δ_1, δ_2, this results in candidate k-th step collaborative targeted maximum likelihood estimators $P_n \to \hat{Q}^k_{\delta_1,\delta_2}(P_n)$, corresponding treatment mechanism estimators $P_n \to \hat{g}^k_{\delta_2}(P_n)$, and corresponding $P_n \to \Psi(\hat{Q}^k_{\delta_1,\delta_2}(P_n))$ targeted maximum likelihood estimators of ψ_0, indexed by k.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

For each δ_1, δ_2, in order to select among these candidate targeted maximum likelihood estimators indexed by k, we use the cross-validated penalized log-likelihood defined as:

$$\begin{aligned} L(k, \delta_1, \delta_2) &= E_{B_n} P^1_{n,B_n} \left(Y - \hat{Q}^k_{\delta_1,\delta_2}(P^0_{n,B_n})(W, A) \right)^2 \\ &\quad + \widehat{MSE}_{CV}(P_n)(\hat{Q}^k_{\delta_1,\delta_2}, \hat{g}^k_{\delta_2}). \end{aligned}$$

This results now in candidate (δ_1, δ_2)-specific collaborative targeted maximum likelihood estimators $\hat{Q}^*_{\delta_1,\delta_2}$, with corresponding initial estimator $\hat{Q}_{\delta_1}$ and collaborative treatment mechanism estimator $\hat{g}_{\delta_1,\delta_2}$ (note the choice of k is now a function of δ_1, δ_2 so that also the collaborative estimator $\hat{g}$ is indexed by these choices).

6.5 Selection among candidate collaborative targeted maximum likelihood estimators

We could select δ_1, δ_2 by minimizing the same cross-validated penalized log-likelihood, e.g., by simply simultaneously minimizing the above criterion over the triplets (k, δ_1, δ_2). Alternatively, we could employ empirical efficiency maximization for all these candidate collaborative targeted maximum likelihood estimators that are assumed to be asymptotically linear with influence curve the efficient influence curve plus a contribution from the collaborative estimator of g_0, as stated in our asymptotic linearity theorem. Thus, by ignoring this latter contribution to the influence curve, we could also select δ_1, δ_2 as the minimizer of the sum of the variances of the components of the efficient influence curve of ψ_0: (with $\delta = (\delta_1, \delta_2)$)

$$\delta_n = \arg\min_{\delta} \sum_{j=1}^{d} P_n D^*_j(\hat{Q}^*_\delta, \hat{g}_\delta)^2.$$

Other criteria based on the vector-efficient influence curve could be considered as well.

6.6 Statistical inference based on CLT

The resulting collaborative targeted maximum likelihood estimator $Q_n = \hat{Q}^*(P_n)$ and corresponding $g_n = \hat{g}(P_n)$ solve the efficient influence curve equation $0 = P_n D^*(\Psi(Q_n), g_n, Q_n)$, so that $\psi_n = \Psi(Q_n)$ can be analyzed with our

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

asymptotics theorem, and inference can be based on the influence curve. So we could estimate the covariance matrix as

$$\Sigma_n = E_{B_n} P^1_{n,B_n} D^*(\hat{Q}^*(P^0_{n,B_n}), \hat{g}^*(P^0_{n,B_n}))^2,$$

where one should include the g_n-component to obtain more accurate inference. Statistical inference would be based on the normal working model $\psi_n \sim N(\psi_0, \Sigma_n)$ to construct confidence intervals, confidence bands, p-values, and possible multiple testing adjusted p-values.

7 Simulation

In this section we first describe an implementation of the C-TMLE algorithm, then review other estimators in the literature before presenting the results of three simulations designed to offer a performance comparison across a variety of situations commonly found in the analysis of real-world data. Though each of these estimators described below is capable of providing an unbiased estimate of the parameter of interest under ideal conditions, results indicate that the C-TMLE estimator consistently performs as well or better than the others across all simulations. We end by comparing performance of the new C-TMLE estimator with the standard TMLE.

7.1 C-TMLE implementation

The general C-TMLE procedure is to create several stage 1 (non-targeted) density estimators and carry out stage 2 procedures for each of these. Penalized cross-validation is used to choose among the final candidate estimators that are indexed by stage 1 and stage 2 candidates. The implementation presented here is based on only one initial stage 1 estimate for simplicitly. We describe specific choices that were employed for the simulations and data analysis presented in the following sections, occasionally noting other implementation options.

Step 1: Obtain an estimate Q^0_n of $Q_0(A, W)$. A data-adaptive machine learning approach to obtaining this initial estimate is recommended. The super learner (SL) is a prediction algorithm that creates a weighted combination of predictions of many individual prediction algorithms, with weights selected using V-fold cross-validation (van der Laan et al., 2007). In practice, it is important to include algorithms in the SL library of predictors that cover different model spaces, e.g. support vector machines, splines, neural nets, etc., since the true best estimation method is unknown. If SL is not used, any particular data adaptive machine learning

Collection of Biostatistics
Research Archive

algorithm providing a consistent estimate is acceptable. For the simulations described below, the DSA algorithm was used to provide the initial estimate of the the true regression of Y on treatment A and confounders W.

Step 2: Generate candidate second stage estimators Q_n^k. In the simulations forward selection was used to build a sequence of updates for g_0 that are increasing in size.

Though not required, a sensible approach is to use the intercept model for g to construct a first clever covariate, h_1, used to create the first targeted maximum likelihood candidate, Q_n^1.

$$g_n^1(1 \mid W) = P(A = 1),\ g_n^1(0 \mid W) = P(A = 0)$$

$$h_1 = \left(\frac{I[A = 1]}{g_n^1(1 \mid W)} - \frac{I[A = 0]}{g_n^1(0 \mid W)} \right).$$

This results in the first candidate second stage estimator $Q_n^1 = Q_n^0 + \epsilon_1 h_1$, where ϵ_1 is fitted by least-squares regression of Y on h_1 with offset Q_n^0. Next we create an updated model for g by adding a main term to the intercept, and the resulting targeted MLE using the corresponding clever covariate is evaluated. The best main term is selected based on a penalized log-likelihood criterion for the targeted MLE fit. Additional terms are incorporated in the g-fit as long as they increase the overall penalized log-likelihood for the resulting Q_0-targeted MLE fit. Thus the penalized likelihood is defined as the empirical sum of squared residuals at the resulting Q_0-fit plus a penalty term proportional to the estimated variance of the target parameter, the empirical variance of D^*, the main component of the efficient influence curve (see below), at the resulting Q_0-fit and the candidate g-fit. In the event that no terms in the model for g increase the penalized likelihood of the resulting Q_0-fit, the targeted MLE update is carried out with the clever covariate that provided the best penalized log-likelihood, and the above process is iterated with this new initial estimator and next clever covariate indexed by g fits that are still building on last g-fit.

As an example suppose that in addition to the intercept term, m terms, ordered $1, \ldots, m$, are incorporated into the model for g, at which point no further increase of the penalized log likelihood is possible. We define

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

candidate estimators Q_n^2 through Q_n^{m+1} as:

$$\begin{aligned} Q_n^2 &= Q_n^1 + \epsilon_2 h_2 \\ Q_n^3 &= Q_n^1 + \epsilon_3 h_3 \\ &\vdots \\ Q_n^{m+1} &= Q_n^1 + \epsilon_{m+1} h_{m+1} \end{aligned}$$

where the corresponding models g_n^{i+1} contains all the terms in the model for g_n^i plus one additional term, $i = 2, \ldots, m$. At this point Q_n^{m+1} is considered as a new "initial" estimate of the true regression, and the entire process starts over in order to build a second clever covariate augmenting the previous fit g_n^{m+1} used in h_{m+1}. To continue the example, $Q_n^{m+2} = Q_n^{m+1} + \epsilon_{m+2} h_{m+2}$. This process is iterated until all terms are incorporated into the final model for g. If the maximal number of terms that can be added is given by K, then this results in K candidate estimators Q_n^k, $k = 1, \ldots, K$, corresponding with treatment mechanism estimators g_n^k, $k = 1, \ldots, K$. Note that the number of clever covariates in Q_n^k that are added to the initial estimator Q_n^0 cannot be predicted, and depends on how many covariates can be added to the treatment mechanism estimator in each iteration before reaching the local maximum (not allowing a further increase of the penalized log-likelihood).

Note that the model for g is not restricted to main terms only. For example, variables can be created that correspond to higher-order terms. In addition, a categorical or continuous covariate can be split into many binary covariates, thereby allowing for more nonparametric modeling of the effect of a single covariate. When there are many covariates it might be desirable in practice to terminate the procedure before all covariates have been incorporated into the model for g, though care must be taken to ensure that none of the candidates thereby excluded from the subsequent selection process potentially maximize the penalized log-likelihood criterion. SL can be integrated into the second stage as well. A series of increasingly non-parametric propensity score SL estimates can be obtained based on different adjustment sets. These SL fits are used as the main terms for the stage 2 forward selection to build candidate $\hat{g}$ estimators.

The presented algorithm illustrates that the number of clever covariates used to update the initial estimator Q_n^0 depends entirely on the likelihood and cannot be pre-determined. Terms are incorporated into the model for g for a single clever covariate until there is a decrease in the likelihood.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

At that point the estimate is updated from $Q_n^m \rightarrow Q_n^{(m+1)}$ and the process iterates until all candidate TMLEs have been constructed.

We also note that we can represent these estimators Q_n^k and corresponding treatment mechanism estimators g_n^k as mappings $\hat{Q}^k$ and $\hat{g}^k$ applied to the empirical distribution P_n: $Q_n^k = \hat{Q}^k(P_n)$, $g_n^k = \hat{g}^k(P_n)$, $k = 1, \ldots, K$. These mappings $P_n \rightarrow \hat{Q}^k(P_n)$ represent our candidate estimators of the true regression Q_0, and in the next step we use cross-validation to select among these candidate algorithms.

Step 3: Select the estimator that maximizes the V-fold cross-validated penalized likelihood, where V was set to 5. Maximizing the penalized likelihood is equivalent to minimizing the residual sum of squares (RSS) plus a penalty term corresponding to the mean squared error (MSE), which can be decomposed into variance and bias terms:

$$k^* = \underset{k}{\operatorname{argmin}}\, cvRSS_k + cvVar_k + n * cvBias_k^2.$$

These terms are defined as follows:

$$\begin{aligned}
cvRSS_k &= \sum_{v=1}^{V} \sum_{i \in Val(v)} (Y_i - \hat{Q}^k(P_{nv}^0)(W_i, A_i))^2 \\
cvVar_k &= \sum_{v=1}^{V} \sum_{i \in Val(v)} D^{*2}(\hat{Q}^k(P_{nv}^0), \hat{g}_k(P_n), \Psi(\hat{Q}^k(P_{nv}^0)))(O_i) \\
cvBias_k &= \frac{1}{V} \sum_{v=1}^{V} \Psi(\hat{Q}^k(P_{nv}^0)) - \Psi(\hat{Q}^k(P_n)) \\
D^*(Q, g, \Psi(Q))(O) &= \frac{I[A=1] - I[A=0]}{g(A \mid W)}(Y - Q(A, W)) \\
&\quad + \frac{1}{n} \sum_{i=1}^{n} Q(1, W) - Q(0, W) - \Psi(Q)
\end{aligned}$$

where v ranging from 1 to V indexes the validation set $Val(v)$ for the vth fold, $\Psi(Q)$ is a mapping from Q to the parameter of interest, and $\hat{Q}^k(P_{nv}^0)$ denotes the k-th C-TMLE applied to the corresponding training sample P_{nv}^0, containing $n(1-p)$ observations, with $p = 1/V$.

There are many variations for obtaining ψ_n^{C-TMLE}. For example, given an a priori set of candidate nuisance parameter estimators, $\hat{g}_j$, that includes

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

highly non-parametric candidates we could construct clever covariates $h_j(g)$, and then use forward selection with this set of clever covariates, using the initial estimator as off-set, to build (second stage) model-fits for Q_0 of increasing size, where each term in the model corresponds to one of the clever covariates. The number of clever covariates that are added in this forward-selection algorithm can be selected using likelihood-based cross-validation.

Note that in contrast with the algorithm described above, in which previous coefficients are used as fixed offsets in the regression, coefficients in front of each term are estimated by least squares, thereby solving the efficient influence equation corresponding to each $\hat{g}_j$, in particular the most non-parametric of these. Because these covariates are highly correlated, refitting all coefficients in front of clever covariates at each step in the forward selection algorithm is likely to result in highly variable coefficient estimates, and therefore less stability in the estimate of the parameter of interest.

Another alternative approach is to define $\psi_n^{C-TMLE} = \psi(Q_n^1)$, where $Q_n^1 = Q_n^0 + \epsilon_n h(g_n^{k^*})$ is the targeted MLE updating the initial estimator with the final selected clever covariate defined by carrying out the k^* moves in the above forward selection algorithm to obtain a g-fit, where k^* is the optimal number of moves selected by likelihood-based cross-validation (exactly as above). This variation did not improve performance in simulation studies not presented in this article. We mention these alternatives only to underscore the fact that C-TMLE methodology can be implemented in a variety of ways, and is not limited to the specific implementation presented here.

7.1.1 Inference

The variance of the influence curve (IC) of the C-TMLE provides suitable inference, under certain regularity conditions, and assuming that the collaborative estimator g_n converges to a $g_0 = g_0(Q)$, where $g_0(Q)$ represents a true conditional distribution of A given $W(Q)$ for a subset or reduction $W(Q)$ of all covariates W, so that $P_0 D^*(Q, g_0(Q), \psi_0) = 0$. For example, it suffices that the limit $g_0(Q)$ is a true conditional distribution of A, given $W(Q)$, for a $W(Q)$ such that $(Q_0 - Q)(1, W), (Q_0 - Q)(0, W)$ only depend on W through $W(Q)$. The asymptotics theorem presented above states that ψ_n is an asymptotically linear estimator of ψ_0 with influence curve

$$IC(P_0) = D^*(Q, g_0, \psi_0) + IC_{g_0},$$

where IC_{g_0} denotes the influence curve of the linearization of $P_0 D^*(Q, g_n, \psi_0)$ viewed as an estimator of $P_0 D^*(Q, g_0, \psi_0)$. This additional term IC_{g_0} represents the contribution to the influence curve from the estimator g_n. The

COBRA
Collection of Biostatistics
Research Archive

formula for the efficient influence curve/canonical gradient $D^*(Q, g_0, \psi_0)$ is given in the previous section for the particular causal effect parameter, $\psi_0 = E_W[E[Y \mid A = 1, W] - E[Y \mid A = 0, W]]$.

In our application of C-TMLE, g_n is a data adaptively selected logistic regression model fitted with maximum likelihood estimation. Thus, if we define $\{g_\alpha : \alpha\}$ as the logistic regression model selected, and α_n is the MLE, then $g_n = g_{\alpha_n}$. We will approximate IC_{g_0} with the influence curve of the asymptotic linearization of $P_0 D^*(Q, g_{\alpha_n}, \psi_0) - D^*(Q, g_\alpha, \psi_0)$. This IC_{g_0} can now be determined with a straightforward application of the delta method. The formula for IC_{g_0} we derived is given by:

$$IC_{g_0}(O) = -a_0 \cdot IC_\alpha(O)$$

where

$$\begin{aligned} a_0 &= P_0(Y - Q(A,W))\overrightarrow{W} h_\alpha(A,W), \\ h_\alpha(A,W) &= \left[\frac{A g_\alpha(0 \mid W)}{g_\alpha(1 \mid W)} + \frac{(1-A) g_\alpha(1 \mid W)}{g_\alpha(0 \mid W)}\right], \\ IC_\alpha(O) &= P_0\left[\overrightarrow{W}\overrightarrow{W}^T g_\alpha(1 \mid W) g_\alpha(0 \mid W)\right]^{-1} (A - g_\alpha(1 \mid W))\overrightarrow{W}. \end{aligned}$$

The notation $\overrightarrow{W}$ is used to denote the vector of main terms that is included in the logistic regression model g_{α_n}. Note that a_0 is a vector of the same dimension as $\overrightarrow{W}$.

This influence curve is estimated by its empirical analog, given by:

$$\widehat{IC}_{g_0}(O) = -a_n \cdot \widehat{IC}_\alpha(O)$$

where

$$\begin{aligned} a_n &= \frac{1}{n}\sum_{i=1}^{n} (Y_i - \hat{Q}(A_i, W_i))\overrightarrow{W}_i h_{\alpha_n}(A_i, W_i), \\ h_{\alpha_n}(A_i, W_i) &= \left[\frac{A_i g_{\alpha_n}(0 \mid W_i)}{g_{\alpha_n}(1 \mid W_i)} + \frac{(1-A_i) g_{\alpha_n}(1 \mid W_i)}{g_{\alpha_n}(0 \mid W_i)}\right], \\ \widehat{IC}_\alpha(O) &= \left[\frac{1}{n}\sum_{i=1}^{n} \overrightarrow{W}_i \overrightarrow{W}_i^T g_{\alpha_n}(1 \mid W_i) g_{\alpha_n}(0 \mid W_i)\right]^{-1} (A - g_{\alpha_n}(1 \mid W))\overrightarrow{W}. \end{aligned}$$

The standard error of the C-TMLE is now estimated as $\mathrm{SE}(\psi_n) = \sqrt{var(IC)/n}$, where $\mathrm{var}(IC) = 1/n \sum_i \hat{IC}_i^2$ is the sample variance of the estimated influence

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

curve. A 95% confidence interval (CI) is constructed as $\psi_n \pm 1.96\text{SE}(\psi_n)$. The bootstrap is an alternative valid method for asymptotically valid inference, but it is much more computationally intensive.

We remark that it is good practice to incorporate the additional term IC_{g_0} in the influence curve, thereby targeting the true influence curve of the estimator. We can provide the following qualitative understanding of the contribution of IC_{g_0} to the influence curve of the estimator. If Q_n^0 converges to the true Q_0, then the term IC_{g_0} equals zero, and if Q_n^0 is inconsistent, and g_n converges to the fully adjusted g_0, then IC_{g_0} is known to reduce the variance of the influence curve (section 2.3.4 van der Laan and Robins (2003)). Based on these two facts, we suggest that ignoring the contribution IC_{g_0} will typically result in asymptotically conservative confidence intervals. Empirical evidence presented in Section 6 using finite samples ($n = 1000$) supports this. However, from a theoretical point of view, there seems to be no guarantee that IC_{g_0} always reduces the variance.

7.2 Current methods for estimating marginal causal treatment effects

Current methods for estimating the marginal causal effect of a treatment A on outcome Y are compared with C-TMLE on simulated data below. The estimators under consideration are the G-computation estimator (Robins, 1986), the IPTW estimator (Hernan et al. (2000), Robins (2000b)), a double robust IPTW estimator (DR-IPTW), (Robins and Rotnitzky (2001); Robins et al. (2000); Robins (2000a)), a propensity score estimator (Rosenbaum and Rubin, 1983) that calculates the marginal treatment effect as the mean across strata defined by the conditional probability of receiving treatment, and an extension to propensity score estimators implemented in *Matching*, a publicly available R package (Sekhon (2008)).

Recall that our parameter of interest is given by: $\psi_0 = E_W[E[Y \mid A = 1, W] - E[Y \mid A = 0, W]]$. Each of the estimators we are considering rely on estimates of one or both of the following: $Q_0(A, W) \equiv E[Y \mid A, W]$ and $g_0(A, W) \equiv P(A \mid W)$. The first conditional distribution can be estimated by, for example, a regression of Y on A and W. The second, which we refer to as the treatment mechanism, is sometimes known, for example in a randomized trial. When the treatment mechanism is unknown it can be estimated by a logistic regression of A on W. Each estimator is defined below.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

$$
\begin{aligned}
\psi_n^{Gcomp} &= \frac{1}{n}\sum_{i=1}^{n}(Q_n^0(1,W_i) - Q_n^0(0,W_i)) \\
\psi_n^{IPTW} &= \frac{1}{n}\sum_{i=1}^{n}\left[I(A_i=1) - I(A_i=0)\right]\frac{Y_i}{g_n(A_i,W_i)} \\
\psi_n^{DR-IPTW} &= \frac{1}{n}\sum_{i=1}^{n}\frac{[I(A_i=1) - I(A_i=0)]}{g_n(A_i \mid W_i)}(Y_i - Q_n^0(W_i,A_i)) \\
&\quad + \frac{1}{n}\sum_{i=1}^{n}(Q_n^0(1,W_i) - Q_n^0(0,W_i)) \\
\psi_n^{C-TMLE} &= \frac{1}{n}\sum_{i=1}^{n}(Q_n^*(1,W_i) - Q_n^*(0,W_i)) \\
\psi_n^{PropScore} &= \frac{1}{n}\sum_{i=1}^{n}(Q_n^0(1,s_i) - Q_n^0(0,s_i)) \\
\psi_n^{Matching} &= \frac{1}{n}\sum_{i=1}^{n}(Q_n^0(1,m_i) - Q_n^0(0,m_i))
\end{aligned}
$$

where Q_n^0 refers to an initial estimate of $Q_0(A,W)$, Q_n^* refers to an updated targeted estimate of $Q_0(A,W)$, described in detail in the next section. For the propensity score method, s_i indicates a stratum of the propensity score of covariate vector W_i, and $Q_n^0(a,s)$ denotes an estimator of the true conditional mean $E(Y \mid A=a, S=s)$ given treatment and propensity score. In the last equation m_i indicates a set of matched observations to which subject i is assigned, where matches are based on minimizing a distance between the user supplied covariates W. Each set of matched observations indexed by m results in a corresponding mean regression $Q_n^0(a,m)$ representing an estimate of $E(Y \mid A=a, M=m)$. The creation of the partitioning in sets of matched observations is only a function of the data (W_i, A_i), $i=1,\ldots,n$, thus ignoring the outcome data.

Regarding asymptotic properties of the estimators, the G-computation estimator relies on consistent estimation of Q_0, the IPTW estimator relies on consistent estimation of g_0, while the DR-IPTW estimator yields consistent estimates if one or both nuisance parameters are estimated consistently.

Notice that ψ_n^{C-TMLE} is a G-computation estimate. However, unlike G-compu-tation, which is consistent only when Q_n is a consistent estimator for Q_0, C-TMLE estimates are consistent if either Q_0 or g_0 is estimated consis-

UC Berkeley Collection of Biostatistics Research Archive

Hosted by The Berkeley Electronic Press

tently. ψ_n^{C-TMLE} can equivalently be formulated as a double-robust IPTW estimator:

$$\begin{aligned}\psi_n^{C-TMLE} &= \frac{1}{n}\sum_{i=1}^{n}\frac{[I(A_i=1)-I(A_i=0)]}{g_n^*(A_i \mid W_i)}(Y_i - Q_n^*(W_i, A_i)) \\ &\quad + \frac{1}{n}\sum_{i=1}^{n}(Q_n^*(1, W_i) - Q_n^*(0, W_i))\end{aligned}$$

The propensity score method implemented uses the *Deletion/Substitution/ Addition* (DSA) algorithm (Sinisi and van der Laan, 2004) to model conditional treatment probabilities given covariates W. This data-adaptive algorithm searches over a large space of polynomial models by adding, subtracting, or substituting terms, starting with a base user-specified regression model. The final model, selected by cross-validation with the L2 loss function, was used to estimate a propensity score for each observation. Observations were then divided into five strata based on the quantiles of these propensity scores. Regression of Y on A and strata indicator variables using the full model enabled the calculation of stratum-specific treatment effects, which were averaged to obtain the marginal effect. The Matching estimator generalizes the propensity score approach by carefully matching observations in the treatment and control groups in such a way that potential confounders are evenly distributed, across the matches.

The Matching procedure relies on the genetic algorithm (Holland and Reitman (1977)) to achieve this goal. This is a non-parametric approach for selecting weights on covariates that are in turn are used to determine which observations are matched. Candidate sets of matches are evaluated based on a loss function and a distance metric specified at run-time, and are used to generate successive sets of candidates that achieve good balance Sekhon (2008). The marginal treatment effect is the average effect across strata defined by the matches.

Propensity score methods are especially effective when overall match quality is a function of true confounders. Estimates can suffer even when overall match quality is high if a small subset of covariates responsible for introducing the most bias into the estimate is unevenly distributed between treatment and control groups. Because matches are made without regard to the outcome variable, these methods do not exploit all information available in the data and are known to be less than fully efficient (Abadie and Imbens, 2006). A violation of the experimental treatment assignment assumption, also called the positivity assumption, is known to reduce the quality of the match and

Collection of Biostatistics
Research Archive

introduce bias into the estimate, and can be detected once the matches have been specified. The lack of identifiability as measured by such an assumption results in potential bias for each method, but the augmented IPTW, targeted MLE, and G-computation method allow reliance on extrapolation.

7.3 Comparison of estimators

For each simulation we have a data structure $O = (W, A, Y)$, where W=($W_1, \ldots,$ W_6) is a set of potential confounders of the relationship between binary treatment variable A and continuous outcome Y. Our parameter of interest is the marginal causal effect of treatment on the outcome: $\psi = E_W[E[Y \mid A = 1, W] - E[Y \mid A = 0, W]]$. The simulations are designed to demonstrate estimator performance in the face of confounding of the relationship between treatment and outcome, complex underlying data-generating distributions, and practical violations of the Experimental Treatment Assumption (ETA), i.e., $P(A = a \mid W) < \alpha$, for some small α, implying that there is very little possibility of observing both treated and untreated subjects for some combination of covariates present in the data.

These simulations are designed specifically to illustrate features of the C-TMLE estimator, and there are other simulations for which the relative performance of the estimators would differ. For example, when a correct model for the underlying data generating distribution is known, a parametric regression approach would be optimal. When the outcome is rare, as is often the case in safety analysis, we would not expect the initial fit, Q_n^0, to have much predictive power. In this case, the fully adjusted g_0 is very likely needed for full bias reduction, so creating and evaluating intermediate candidates with C-TMLE may be needlessly computationally expensive. Standard TMLE might be a better approach. Adjusting for many confounders may lead to violations of the ETA assumption when n is small relative to the number of confounders or if the confounders are very strongly predictive of treatment. There are two ways to deal with this. First, one could extrapolate based on model assumptions to arrive at an estimate of the desired parameter. Secondly, one could acknowledge that the parameter of interest is not identifiable from the data, and choose an adjustment set that provides bias reduction without yielding an estimate with variance so large that it is essentially meaningless. This latter approach is taken by second stage of the C-TMLE procedure, by basing the selection of confounders based on the penalized log-likelihood, while the extrapolation approach is still present through the initial first stage estimator.

A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

Data generation Covariates $W_1, \ldots W_5$ were generated as independent normal random variables. W_6 is a binary variable.

$$\begin{aligned} W_1, W_2, W_3, W_4, W_5 &\sim N(0,1) \\ logit(P(W_6 = 1 \mid W_1, W_2, W_3, W_4, W_5)) &= .3W_1 + .2W_2 - 3W_3 \end{aligned}$$

Two treatment mechanisms were defined:

$$logit(g_{1,0}) = logit(P(A{=}1 \mid W_1,\ W_2,\ W_3,\ W_4,\ W_5,\ W_6)) = .3W_1 + .2W_2 - 3W_3$$
$$logit(g_{2,0}) = logit(P(A{=}1 \mid W_1,\ W_2,\ W_3,\ W_4,\ W_5,\ W_6)) = .15(.3W_1 + .2W_2 - 3W_3)$$

The observed outcome Y was generated as

$$Y = Q_{i,0}(A, W) + \epsilon,\ \epsilon \sim N(0,1)$$

with corresponding regression equations:

$$\begin{aligned} Q_{1,0}(A, W) &= A + .5W_1 - 8W_2 + W_3 + 8W_3 - 2W_5 \\ Q_{2,0}(A, W) &= A + .5W_1 - 8W_2 + W_3 + 8W_3^2 - 2W_5 \end{aligned}$$

We consider three different data-generating distributions, $(Q_{1,0}, g_{1,0})$ in simulation 1, $(Q_{2,0}, g_{1,0})$ in simulation 2, and $(Q_{2,0}, g_{2,0})$ in simulation 3. Note that W_6 is strongly correlated with treatment mechanism A in simulations 1 and 2 (corr=0.54), but is not an actual confounder of the relationship between A and Y. W_1, W_2, and W_3 are confounders. The linear nature of the confounding due to W_3 in simulation 1 differs from that in simulations 2 and 3, where the true functional form is quadratic. In this way simulations 2 and 3 mimic realistic data analysis scenarios in which the unknown underlying functional form is seldom entirely captured by the regression model used in the analysis. Finally, the treatment mechanism in simulations 1 and 2 leads to ETA violations ($p(A = a \mid W)$ ranges between 9×10^{-7} and 0.9999978, approximately one-third of the probabilites are outside the range (0.05, 0.95)). In simulation 3 there are no ETA violations ($0.11 < p(A = a \mid W) < 0.88$). In each simulation the true value of the parameter of interest is 1.

7.3.1 Simulation

1000 samples of size $n = 1000$ were drawn from each data generating distribution. Marginal treatment effect estimates were calculated based on the unadjusted regression of Y on A, Gcomp, IPTW, DR-IPTW, propensity score and C-TMLE methods.

Collection of Biostatistics
Research Archive

A main-effects model for Gcomp and DR-IPTW, $\hat{Q}$, was obtained using the DSA algorithm with the maximum model size set to seven. A model for the treatment mechanism $\hat{g}$ used in IPTW, DR-IPTW, propensity score, and Matching estimation was also selected by DSA, again restricted to main terms. The Matching function considered this treatment mechanism model as merely one additional covariate, indistinguishable from the other potential confounders, W. The procedure was run using default settings, except population size for each generation was increased to 200. In contrast, the C-TMLE algorithm includes an aggressive search through a larger space of models to obtain an initial estimate of the density. As a proxy for the super-learner algorithm we used the DSA algorithm to select a model for $\hat{Q}$ containing at most six terms, allowing quadratic terms and two-way interactions.

We expect to see that the estimators that rely on consistent estimation of Q_0 are unbiased in simulation 1, (Gcomp, DR-IPTW, C-TMLE), while estimators relying on consistent estimation of g_0 are unbiased in simulation 3 (IPTW, DR-IPTW, propScore, Matching, C-TMLE).

7.3.2 Results

Table 1: Mean estimate and standard error (SE) for each estimator based on 1000 iterations with sample size $n = 1000$. $\psi_0 = 1$.

	Simulation 1		**Simulation 2**		**Simulation 3**	
	$\overline{\psi}_n$	SE	$\overline{\psi}_n$	SE	$\overline{\psi}_n$	SE
Unadj	−11.97	0.64	−0.98	0.91	0.29	0.86
Gcomp	0.99	0.09	0.76	1.22	0.95	0.68
IPTW	−4.36	0.72	0.03	0.76	0.83	0.90
DR-IPTW	0.99	0.09	0.94	0.62	1.03	0.80
C-TMLE	0.99	0.09	1.00	0.10	1.00	0.07
PropScore	-1.09	1.27	0.42	1.38	0.93	0.59
Matching	−1.22	0.82	0.54	0.73	0.96	0.25

Mean estimates of the treatment effect and standard errors are shown in Table 1 for each simulation. Mean estimates and (0.025, 0.975) quantiles of the probability distribution of each estimator are plotted in Figures 2 and 3.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

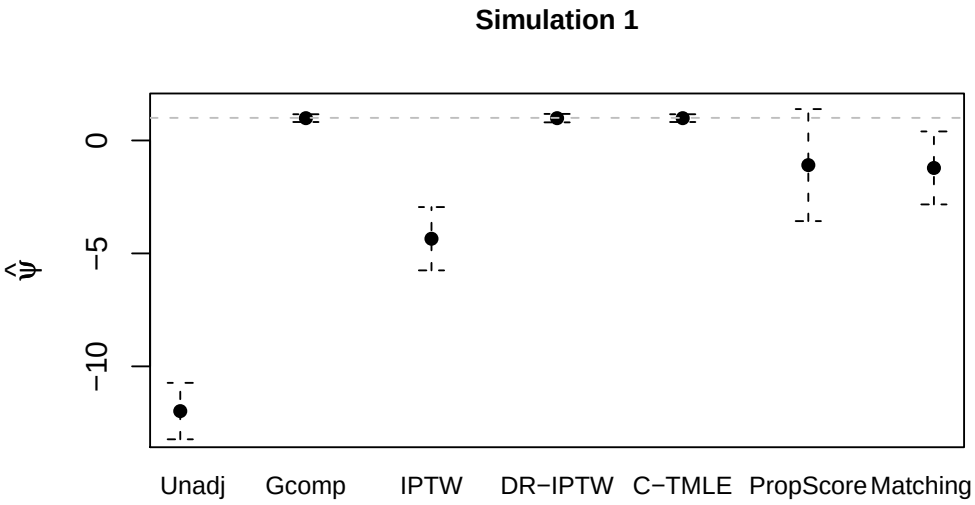

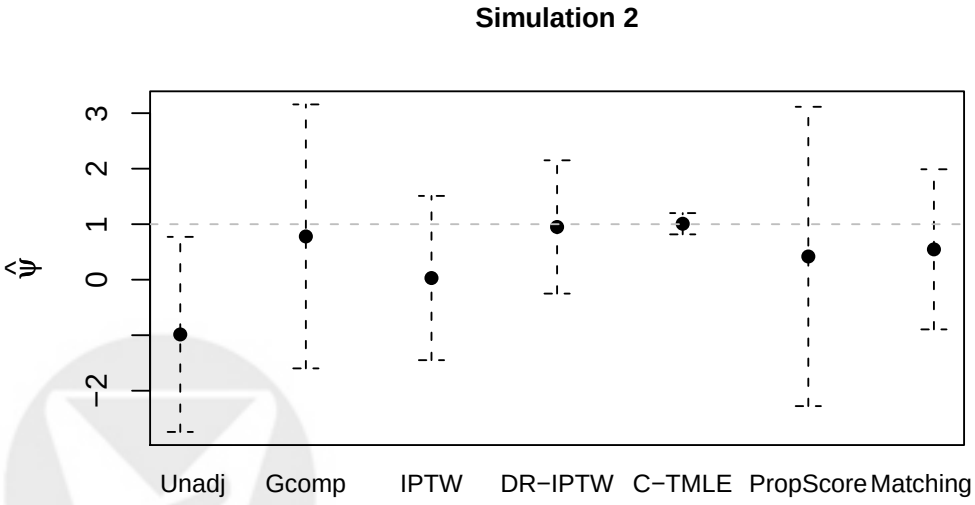

Figure 2: Mean estimates and (0.025,0.975) quantiles for each estimation method, simulations 1 and 2. Dashed line is at true parameter value.

SCRIA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

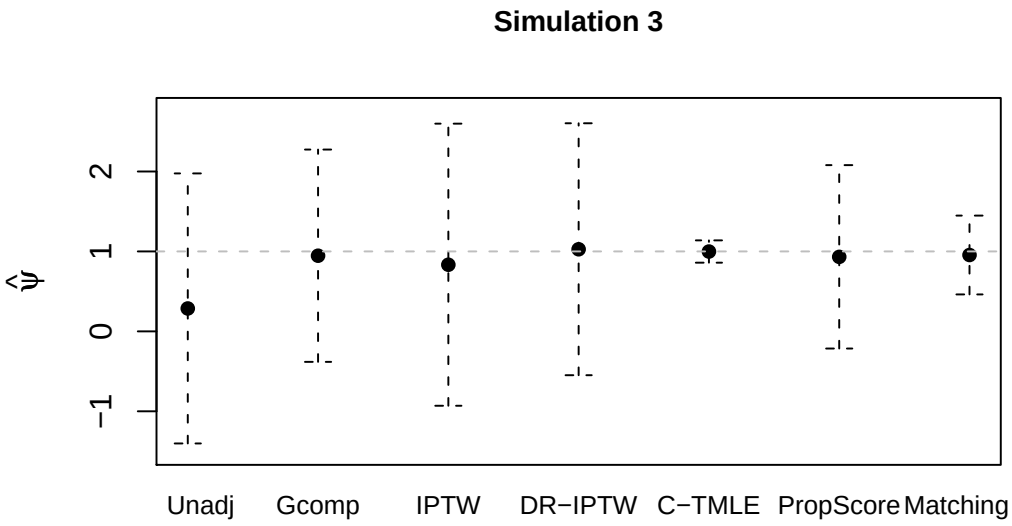

Figure 3: Mean estimates and (0.025,0.975) quantiles for each estimation method, simulation 3. Dashed line is at true parameter value.

Figures 2 and 3 illustrate each estimator's behavior. As expected, estimators relying on consistent estimation of Q_0 are unbiased in simulation 1, estimators relying on consistent estimation of g_0 are unbiased in simulation 3.

- The unadjusted estimator yields biased results in all three simulations due to its failure to adjust for confounders.
- The G-computation estimator performs well in simulation 1 when the model is correctly specified. We understand that misspecification (simulations 2 and 3) will often, though not always, lead to bias in the estimates. However the plots highlight another phenomenon that is easy to overlook. the inability of the misspecified model to adequately account for the variance in the outcome often leads to large residual variance of the estimator, and in practice would have low power to reject a null hypothesis.
- Truncation bias due to ETA violations causes the IPTW estimator using truncated weights to fail in simulations 1 and 2. The estimate is not

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

biased in simulation 3, but the variance is so large that even in this setting where we'd expect IPTW to be reliable it would fail to produce a significant result.

- DR-IPTW estimates are unbiased and have low variance when the functional form is correctly modeled by the regression equation (simulation 1). Though we see little bias in the other two simulations, the variance is large due to misspecification of the treatment mechanism. Because W_6 is a strong predictor of A and is indistinguishable from a true confounder of the relationship between Y and A it is always included in the treatment mechanism, behavior that does not help achieve an accurate estimate of the true treatment effect.

- Propensity score estimators are known to perform poorly when there are ETA violations, e.g. simulations 1 and 2 (Sekhon (2008)). Researchers constructing the propensity score could observe this and choose an alternate propensity score model, but without using information about the outcome this choice would likely be made based on the predictive power of the model, not the potential bias reduction. Both propensity score-based methods do a reasonable job in simulation 3. Abadie and Imbens (2006) shows that matching estimators will not obtain the semiparametric efficiency bound. This theory is borne out in simulation 3, where neither matching methodology confidence interval is as tight as that of the collaborative targeted maximum likelihood estimator.

7.4 Comparison of C-TMLE and TMLE

The double robust property of the targeted maximum likelihood estimator obviates the need for accurate estimation of both Q_0 and g_0 since correct specification of either one leads to consistent estimates of the parameter of interest. However, accurate estimates of both are needed to achieve the Cramer-Rao efficiency bound. Implementations of the standard targeted maximum likelihood estimator (TMLE) therefore strive for ideal estimates of both Q_0 and g_0.

In contrast, the collaborative nature of the second stage of the C-TMLE estimation algorithm leads to selection of an estimator, g_n, that targets that portion of the treatment mechanism needed to reduce bias not already adequately addressed by the first stage estimator for Q_0. For example, covariates included in the model for Q_n^0 might not be selected into the model for g_n because they do not increase the penalized log-likelihood. At the same time, confounders that are not adequately adjusted for in the initial density estimate

Collection of Biostatistics
Research Archive

are quickly added to model for g_n unless the gain in bias reduction is offset by too great an increase in variance. When the initial estimate of the density is a very good fit for the true underlying density, TMLE and C-TMLE have similar performance w.r.t bias, but the C-TMLE will have smaller variance by selecting a g_n that targets non fully adjusted g_0, resulting in a possibly super efficient estimator. When the initial fit is less good, C-TMLE makes judicious choices regarding inclusion of covariates in the treatment mechanism. As predicted by theory, again, this might lead to lower variances when no covariates cause ETA violations. When inclusion of all confounding covariates does violate the ETA assumption, the C-TMLE estimator, in essence, targets a less ambitious data adaptively selected parameter that is identifiable. Data were simulated to illustrate these phenomena.

7.4.1 Data generation

Covariates $W_1, W2$, and W_3 were generated as independent random uniform variables over the interval $[0, 1]$. W_4 and W_5 are independent normally distributed random variables.

$$\begin{aligned} W_1, W_2, W_3 &\sim U(0,1) \\ W_4, W_5 &\sim N(0,1) \end{aligned}$$

Treatment mechanism g_0 was designed so that W_3 is highly predictive of treatment:

$$logit(g_0) = logit(P(A = 1 \mid W)) = 2W_1 + W_2 - 5W_3 + W_5$$

The observed outcome Y was generated as

$$Y = Q_0(A, W) + \epsilon,\ \epsilon \sim N(0, 1)$$

with corresponding regression equation:

$$Q_0(A, W) = A + 4W_1 - 5W_2 + 5W_4W_5$$

7.4.2 Simulation

C-TMLE and TMLE estimates of the parameter of interest, again defined as $\psi = E_W[E[Y \mid A = 1, W] - E[Y \mid A = 0, W]]$, were obtained for 1000 samples of size $n = 1000$ drawn from data generating distribution (Q_0, g_0). For this study we deliberately select a misspecified main-terms only model for Q_0 by running the DSA algorithm on 100,000 observations drawn from that same

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

distribution. $P(A = a \mid W)$ for these observations ranges from 0.004 to 0.996. Approximately 17% of the observations have covariates indicating that the probability of receiving treatment is less than 0.05, indicating that practical ETA violations in finite samples will cause unstable TMLE estimates.

For each iteration an initial regression, Q_n^0, was obtained by fitting the DSA-selected model, $Y = A + W_1 + W_2$, on n observations in the sample. We expect that any estimate of ψ based solely on this model is likely to be incorrect because the model fails to take into account the effect on the outcome of the missing interaction term, and also fails to adjust for the confounding effect of W_5. The targeting step for both targeted maximum likelihood estimators reduces this bias.

In order to construct the covariate used to target the parameter of interest in the updating step of the TMLE algorithm we obtain an estimate g_n of g_0 by running the DSA algorithm, allowing quadratic terms and two-way interaction terms to enter the model. This model was not fixed over the 1000 iterations; the model selection process was carried out each time a sample was drawn from the population. Similarly, covariates that were candidates for inclusion in the model for g_n in the second stage of the C-TMLE estimation algorithm include $(W_1, \ldots, W_5, W_1^2, \ldots, W_5^2)$, and all two-way interaction terms $(W_i W_j)$, where $i \neq j$.

7.4.3 Results

Results of the simulation are shown in Table 2. A small number of aberrant TMLE estimates were major contributers to the variance of that estimator. The three highest TMLE estimates of the treatment effect were (771.91, 37.22, 9.52). It is likely that these high values arise from atypical samples containing observations that presented unusually strong ETA issues. In contrast, all C-TMLE estimates calculated from those same samples range between 0.307 and 1.698. Both estimators' average treatment effect estimates are not far from the true value, $\psi_0 = 1$. As expected, the variance of the TMLE estimator is many times larger than that of the C-TMLE estimator.

Not surprisingly, W_3, the strong predictor of treatment that is not a true confounder of the relationship between treatment and outcome, is included in every one of the 1000 models for g_n selected by the DSA algorithm, but it is included in only 35 of the models constructed in the second stage of the C-TMLE algorithm. At the same time, the interaction term $W_4 W_5$ is included in only two out of 1000 models for g_0 selected by DSA, but is present in 576, more than half, of the collaborative models.

This clearly demonstrates the differences between TMLE's reliance on an

Collection of Biostatistics
Research Archive

external estimate of g_0 and the collaborative approach to estimating the treatment mechanism used by C-TMLE. However, we note that the degradation of TMLE performance under sparsity is due to the unboundedness of the fluctuation function, and can be mitigated by employing an alternative fluctuation function that respects known bounds on the data model. Though a full discussion is beyond the scope of this paper, details may be found in Gruber and van der Laan (2010).

Table 2: Comparison of C-TMLE and TMLE estimators at different levels of truncation. Mean estimate and variance based on 1000 iterations.

	truncation level	# obs truncated	$\overline{\psi}_n$	variance
C-TMLE	∞	0	0.98	0.04
TMLE	∞	0	1.73	597.52
	40	1	1.36	162.38
	10	2	0.94	1.99
	5	9	0.92	1.68

7.4.4 Confidence Intervals

The variance of the influence curve provides the basis for calculation of a 95% confidence interval for the C-TMLE estimate.

$$95\%CI = \psi^{C-TMLE} \pm 1.96\sqrt{(var(IC)/n)}$$

Two sets of confidence intervals were constructed for each of the 1000 iterations in simulation 4, with Q_n^0 misspecified by a main-terms only regression model. As described above, one set of CIs is based on $D^*(Q, g)$, the first term of the IC. The second set is based on the variance of $D^*(Q, g) + IC_g$, which includes the contribution from the estimation of g_n. Table 3 shows that CIs based on D^* alone are conservative when the model for Q_n^0 is misspecified, as expected. In contrast, observed coverage closely approximates the nominal 95% coverage rate when the contribution from the IC_g term is taken into account.

Confidence intervals were also created for an additional 1000 samples from the same data generating distribution that were analyzed using a correct model

UCBRA
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

for Q_n^0. Coverage rates for these confidence intervals are given in Table 3. When Q_n^0 is correctly specified we observe little difference in the coverage rate whether or not we take the contribution from IC_g into account, indicating zero contribution to the variance from the estimate of g_n. Attaining the nominal rate indicates that inference is reliable even when the estimator is super efficient.

Table 3: Empirical coverage of 1000 confidence intervals constructed at a nominal 95% level. SE calculated as $\sqrt{var(IC)/n}$, where the IC was estimated with and without IC_g.

	Coverage	
	$D^*(Q, g_0)$	$D^*(Q, g_0) + IC_g$
Q_n^0 misspecified	.979	.943
Q_n^0 correct	.932	.933

7.5 Data Analysis

We apply the C-TMLE estimator to an observational dataset previously analyzed with the goal of identifying HIV mutations that affect response to the antiretroviral drug lopinavir. (Bembom et al., 2009, 2008) The data includes observations on $O = (W, A, Y)$, where the outcome, Y, is the change in $\log_{10}$ viral load measured at baseline and at follow-up after treatment has been initiated. If follow-up viral load was beneath the limit of detection Y was set to the maximal change seen in the population. $A \in \{0, 1\}$ is an indicator of the presence or absence of a mutation of interest, taking on the appropriate value for each of the 26 candidate mutations in 26 separate analyses. W consists of 51 covariates including treatment history, baseline characteristics, and indicators of the presence of additional HIV mutations. Practical ETA violations stemming from high correlations among some of the covariates and/or low probability of observing a given mutation of interest make it difficult to obtain stable low variance estimates of the association between A and Y. Bembom used a targeted maximum likelihood estimation approach incorporating data-adaptive selection of an adjustment set that relies on setting a limit on the maximum allowable truncation bias introduced by truncating treatment probabilities less than α to some specified lower limit. Covariates whose inclusion in the adjustment set introduces an unacceptable amount of bias are not

COBRA
Collection of Biostatistics
Research Archive

selected. That study's findings showed good greement with Stanford HIVdb mutation scores, values on a scale of 0 to 20 (http://hivdb.stanford.edu, as of September, 2007, subsequently modified), where 20 indicates evidence exists that the mutation strongly inhibits response to drug treatment and 0 signifies that the mutation confers no resistance. Because the C-TMLE method includes covariates in the treatment mechanism only if they improve the targeting of the parameter of interest without having too adverse an effect on the MSE, we expect similar performance without having to specify truncation levels or an acceptable maximum amount of bias.

7.5.1 Analysis description

The dataset consists of 401 observations on 372 subjects. Correlations due to the few subjects contributing more than one observation were ignored. Separate analyses was carried out for each mutation. In each, an initial density estimate, Q_n^0, was obtained using DSA restricted to addition moves only to select a main-terms model containing at most 20 terms, where candidate terms in W include pre-computed interactions detailed in Bembom et al. A was forced into the model. An estimate of the effect on change in viral load was recorded for each mutation. Influence curve-based variance estimates incorporating the contribution from estimating g given by the IC_g term, was used to construct 95% confidence intervals.

7.5.2 Results

Table 4 lists the Stanford mutation score associated with each of the HIV mutations under consideration, as well as the C-TMLE estimate of the adjusted effect of mutation on lopinavir resistance. 95% confidence intervals were constructed based on the variance of the IC. Confidence intervals entirely above zero indicate a mutation increases resistance to lopinavir. Eight of the twelve mutations having a mutation score of 10 or greater fall into this category. Point estimates for the remaining four mutations were positive, but the variance was too large to produce a significant result. Only one of the six mutations thought to confer slight resistance to lopinavir was flagged by the procedure, though with the exception of p10FIRVY point estimates were positive. Stanford mutation scores of 0 for four of the five mutations found to have a significantly negative effect on drug resistance support the conclusion that these mutations do not increase resistance, but are not designed to offer confirmation that a mutation can decrease drug resistance. However, Bembom et al. report that there is some clinical evidence that two of these mutations, 30N and 88S, do

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

Table 4: Stanford score (2007), C-TMLE estimate and 95% confidence interval for each mutation. Starred confidence intervals do not include 0.

mutation	score	estimate	95% CI
p50V	20	1.703	(0.760, 2.645)*
p82AFST	20	0.389	(0.091, 0.688)*
p54VA	11	0.505	(0.241, 0.770)*
p54LMST	11	0.369	(0.002, 0.735)*
p84AV	11	0.099	(-0.139, 0.337)
p46ILV	11	0.046	(-0.222, 0.315)
p82MLC	10	1.610	(1.377, 1.843)*
p47V	10	0.805	(0.282, 1.328)*
p84C	10	0.602	(0.471, 0.734)*
p32I	10	0.544	(0.325, 0.763)*
p48VM	10	0.306	(-0.162, 0.774)
p90M	10	0.209	(-0.063, 0.481)
p33F	5	0.300	(-0.070, 0.669)
p53LY	3	0.214	(-0.266, 0.695)
p73CSTA	2	0.635	(0.278, 0.992)*
p24IF	2	0.229	(-0.215, 0.674)
p10FIRVY	2	−0.266	(-0.545, 0.012)
p71TVI	2	0.019	(-0.243, 0.281)
p23I	0	0.822	(-0.014, 1.658)
p36ILVTA	0	0.272	(-0.001, 0.544)
p16E	0	0.239	(-0.156, 0.633)
p20IMRTVL	0	0.178	(-0.111, 0.467)
p63P	0	−0.131	(-0.417, 0.156)
p88DTG	0	−0.426	(-0.842,-0.010)*
p30N	0	−0.440	(-0.853,-0.028)*
p88S	0	−0.474	(-0.781,-0.167)*

indeed decrease lopinavir resistance.

Our findings are quite consistent with the Stanford mutation scores and with the results from the previous analysis using the data-adaptively selected adjustment set targeted maximum likelihood estimation approach. The C-TMLE method was able to achieve these results without relying on ad hoc or user-specified tuning parameters.

Collection of Biostatistics
Research Archive

7.6 Summary

Simulation studies demonstrate the collaborative double robustness and efficiency of C-TMLE methodology, which allows for consistent efficient estimation in situations when other estimators can fail to perform adequately. In practice these failures may lead to biased estimates and to confidence intervals that fail to attain the correct coverage, as suggested by the IPTW results in simulations 1 and 2, where weights depend on a variable highly predictive of treatment that is not a true confounder of the relationship between Y and A. It is worth noting that the unadjusted estimator applied to data from a randomized controlled trial in which randomization fails to evenly distribute confounders across treatment arms will also yield (finite sample) biased results, as we saw in simulations 1,2, and 3.

As simulations 2 and 3 demonstrate, a misspecied parametric model not only results in biased estimates, but can also easily fail to adequately explain the variance in the outcome. Therefore estimates of the parameter of interest will have a larger variance than the semiparametric information bound achieved by an efficient estimator, such as C-TMLE. Such misspecified parametric models can easily result in the construction of a confidence interval that contains 0, and therefore a failure to reject a false null hypothesis, even when the point estimate is close to the true value of the parameter of interest. Since misspecied parametric models are the rule rather than the exception, in the analysis of data from an unknown data-generating distribution, using C-TMLE combined with super learning for the initial estimator, is a prudent course of action.

Estimators that rely on nuisance parameter estimation (IPTW, DR-IPTW, TMLE, propensity score-based estimation) break down when there are ETA violations, failing to reduce bias, or even increasing bias, while incurring high variance that renders estimates meaningless (no statistical signicance). An effort to reduce variance through truncation introduces bias into the estimate, and requires a careful trade-off. C-TMLE addresses these issues, in the sense that it is able to utilize the covariates for effective bias reduction, avoiding harmful bias reduction efforts. As a targeted-MLE, the bias-variance tradeoff is targeted towards the estimation of the parameter of interest, not the estimate of the entire density. The collaborative nature of the estimation of the treatment mechanism in the C-TMLE confers three advantages:

1. The treatment mechanism model will exclude covariates that are highly predictive of treatment but do not truly confound the relationship between treatment and the outcome.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

2. The treatment mechanism model will include only covariates that help adjust for residual bias remaining after stage 1 adjustment.

3. Cross-validation based on a penalized log-likelihood will not select a treatment mechanism model that includes a term that leads to violations of the ETA assumption and thereby large variance of the corresponding targeted MLE without the benefit of a meaningful bias reduction.

Influence-curve based inference is theoretically sound, and achieves the desired coverage rate across a wide range of simulations, in addition to the ones presented.

8 Discussion

For most data sets little to no knowledge is available about the data generating distribution. Consequently, the true model is a large infinite dimensional semi-parametric model. In such models many data adaptive approaches can be considered for fitting the true distribution of the data, based on different approximation function spaces, different searching strategies for maximizing an empirical criterion (such as the empirical log-likelihood) over these spaces, and different methods for selecting the fine tuning parameters indexing the function spaces and search strategies. Depending on the true data generating distribution, these algorithms will have very different levels of performance in approximating the true data generating distribution. As a consequence, cross-validation based super learning should be employed to find the best weighted combination among a large user supplied set of candidate estimators of the true data generating distribution. The user has an option to select the wished loss function. The oracle property of the cross-validation selector (van der Vaart et al. (2006), van der Laan et al. (2006)) teaches us that the super learner will asymptotically perform exactly as well, w.r.t. the loss-based (e.g., Kullback-Leibler) dissimilarity measure, as the best weighted combination of the candidate algorithms optimized for each data set. A crucial assumption is that the loss function used in this super learning methodology is uniformly bounded in all the candidate estimators. This could be viewed as a semi-parametric model assumption, and it is essential for robustness of the resulting estimator against outliers. It can be arranged by bounding the candidate estimators so that the loss function remains bounded.

The super learner fit represents a best fit for the purpose of estimation of the whole distribution of the data w.r.t. the loss-function specific dissimilarity,

Collection of Biostatistics
Research Archive

so that the bias-variance trade-off is not targeted (enough) w.r.t. the parameter of interest. Overall, the resulting estimate it overly biased for the smooth target parameter.

Therefore, our methodology involves a second targeted modification of the first stage super learner fit that aims to reduce the bias w.r.t the target parameter, while simultaneously increasing the likelihood fit. A single fluctuation function that would yield asymptotic optimal bias reduction as defined by the efficient influence curve of the target parameter is determined. This fluctuation function needs to have a score-vector at zero fluctuation whose linear span includes the efficient influence curve of the target parameter. This fluctuation function depends on an unknown nuisance parameter of the data generating distribution, such as a censoring mechanism.

In one particular embodiment of our C-TMLE, we define an iterative sequence of subsequent fluctuations, starting with the initial super learner fit, where the subsequent fluctuation functions are estimated with increasingly nonparametric estimates of the nuisance parameter, including a final fully non-parametrically estimated fluctuation function. In addition, by construction, we make sure that for each fluctuation function the nuisance parameter estimator that results in maximal increase in likelihood (or preferred loss function) fit is selected, among the candidate nuisance parameter estimators that are more nonparametric than the one selected at previous fluctuation function. In this way, we arrange that most of the targeted bias reduction occurs in the first few fluctuations. The actual number of times we carry out the subsequent update is selected with likelihood based cross-validation.

Essentially, we try to move towards the asymptotically optimal bias reduction (identified by the true nuisance parameter/censoring mechanism used in the data generating distribution) along a sequence of targeted bias reduction steps, but we stop moving towards this asymptotically optimal bias reduction when it results in a loss of likelihood fit as measured by the cross-validated log-likelihood. In general, our template for C-TMLE allow a fine sequence of nested targeted bias reduction steps (i.e., a fine sequence of candidate second stage estimators indexed by increasingly nonparametric nuisance parameter estimators) whose fits contain this set of candidate-fits as a subsequence, thereby potentially providing an additional improvement in practical performance of the resulting C-TMLE.

Theoretical results teach us that this push towards the asymptotically optimal bias reduction takes into account how well the initial estimator approximates the true distribution, by giving preference to targeted bias reduction steps that improve the log-likelihood fit using the initial estimator as off-set. As a consequence, the C-TMLE is able to avoid selecting irrelevant or harmful

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

(w.r.t. relevant factor of density) fits of the nuisance parameter, even though such fits might improve the overall fit of the nuisance parameter. That is, the fit of the nuisance parameter is targeted towards our primary goal, the parameter of interest.

Specifically, we establish a general asymptotic linearity theorem for collaborative targeted maximum likelihood estimators, which provides us with the influence curve of our estimator, and thereby statistical inference. Fortunately, the influence curve happens to be equal to the efficient influence curve at the collaborative nuisance parameter estimator (its limit) plus a contribution of the nuisance parameter estimator as an estimator of its limit, providing immediate variance estimates. Gains in efficiency, resulting in possible super efficiency, are obtained by estimating the nuisance parameter collaboratively, in relation to the initial estimator. An inspection of the efficient influence curve allows us to precisely define the sufficient and minimal term $H(g_0, Q - Q_0)$ one needs to adjust for in g_0 to achieve the wished full bias reduction. We propose to estimate this term and force adjustment by this term in the candidate censoring mechanism estimators within the C-TMLE procedure, without relying on its correct specification, thereby preserving and only enhancing the double robustness of the C-TMLE.

The targeted maximum likelihood estimator relies itself on maximizing an empirical mean of a loss function over a fluctuation function, and the derivative of this empirical criterion at zero fluctuation needs to be the empirical mean of the efficient influence curve at the estimator to be fluctuated. We refer to this loss as the log-likelihood loss, but it needs to be understood that this choice can already be targeted (e.g,. van der Laan (2008b)) in the sense that it is typically implied by the efficient influence curve of the target parameter (e.g, one would not use a likelihood based on factors that are not needed to evaluate the target parameter). In addition, we propose to replace this log-likelihood by a penalized log-likelihood, where the penalty is scaled appropriately, has negligible contribution for nicely identifiable target parameters, but blows up for fits that result in extremely variable or biased estimators of the parameter of interest. Even though the penalty's effect on the Kullback-Leibler dissimilarity is asymptotically negligible for identifiable parameters, for parameters that are borderline identifiable, this penalty can yield dramatic additional finite sample improvements to the C-TMLE. In essence, it builds in a sensible robustness of the resulting C-TMLE as an estimator of the target parameter.

Given the initial estimator, the candidate censoring mechanism estimators, and the corresponding sequence of targeted MLE indexed by these increasingly nonparametric estimators of the censoring mechanism, the log-likelihood or penalized log-likelihood is used for cross-validation selection of the depth of the

COBRA
Collection of Biostatistics
Research Archive

bias reduction (i.e, for selection among these candidate targeted MLEs) in the C-TMLE algorithm. However, a preferred loss function can be used to build the initial estimator, to build the candidate censoring mechanism estimators, and to select among different collaborative C-TMLEs. In particular, one could use the penalized log-likelihood as preferred loss function.

We propose as a possibly more targeted selection criterion the cross-validated variance of the square of the efficient influence curve (e.g., if target parameter is one-dimensional), where a collaborative estimator is used for the nuisance parameter (censoring mechanism) in the efficient influence curve: i.e. we propose as loss function for a candidate Q the square of the efficient influence curve at its targeted version, using a collaborative estimator of g_0. Indeed, $E_0 D^{*2}(Q^*_{g_0}, g_0, \Psi(Q^*_{g_0}) = E_0 D^{*2}(Q^*_{g_0}, g_0, \psi_0)$ is a variance of a collaborative double robust estimator of ψ_0, indexed by initial estimator Q, using known g_0, and is thereby a valid and sensible loss function. In order to also take into account that g_0 is estimated, one could also minimize the variance of the actual influence curve of the collaborative double robust targeted MLE using Q as initial and using g_n as collaborative estimator. This would imply the same square of influence curve, but now the influence curve equals D^* plus an contribution from estimation of g_0.

Finally, we are also able to select a targeted loss function for g_0 in the C-TMLE template, by making its loss-based dissimilarity a measure of goodness of fit of the g_0-specific fluctuation function in which the estimator of g_0 is used.

The collaborative double robust maximum likelihood estimator utilizes 1) loss-based super learning to obtain the best initial estimator of Q_0 (grounded by theory for cross-validation selector), 2) loss-based super learning to generate best estimators of candidate censoring mechanisms adjusting for increasingly large adjustment sets (grounded by theory for cross-validation selector), 3) specification of loss functions that target the needed Q_0, g_0 for identification of the efficient influence curve, 4) targeted maximum likelihood estimation along a fluctuation function using such a censoring mechanism estimator to remove bias for target parameter (grounded by theory for targeted maximum likelihood estimation), 4) Q_0-(penalized)-likelihood based cross-validation to select among such candidate censoring mechanism estimators, and thereby obtain the collaborative estimator of the censoring mechanism for the fluctuation function that removes the bias, while avoids unnecessary loss in variance (grounded by oracle results for cross-validation, collaborative double robustness, and our asymptotic linearity theorem), 5) efficient influence curve based dimension reductions (the minimal sufficient covariate) to be included in the censoring mechanism estimators that allow for effective bias reduction in the T-MLE, if correctly specified, (and will not harm the collaborative double

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

robustness if not) (grounded by theory on efficient influence curve representation and our collaborative double robustness theorem), 6) efficient influence curve based loss function for Q_0 to build more targeted candidate censoring mechanism estimators 7) efficient influence curve based loss functions for Q_0 to make the initial estimator more targeted, to make the selection among candidate C-TMLE more targeted, resulting in smaller asymptotic variance of the resulting C-TML estimator of the target parameter (grounded by empirical efficiency maximization theory, and our asymptotic linearity theorem for C-TMLE).

To summarize, this article provides a template for loss-based adaptive (super) efficient estimators in semiparametric models that are targeted towards a particular target feature of the distribution of the data, and for which statistical inference is available.

References

P.K. Andersen, O. Borgan, R.D. Gill, and N. Keiding. *Statistical Models Based on Counting Processes*. Springer-Verlag New York, 1993.

O. Bembom, M.L. Petersen, S.-Y. Rhee, W. J. Fessel, S.E. Sinisi, R.W. Shafer, and M.J. van der Laan. Biomarker discovery using targeted maximum likelihood estimation: Application to the treatment of antiretroviral resistant hiv infection. *Statistics in Medicine*, page http://www3.interscience.wiley.com/journal/121422393/abstract, 2008.

O. Bembom, M.J. van der Laan, T. Haight, and I.B. Tager. Lifetime and current leisure time physical activity and all-cause mortality in an elderly cohort. *Epidemiology*, 2009.

Oliver Bembom and Mark van der Laan. Statistical methods for analyzing sequentially randomized trials, commentary on JNCI article Adaptive therapy for androgen independent prostate cancer: A randomized selection trial including four regimens, by Peter F. Thall, C. Logothetis, C. Pagliaro, S. Wen, M.A. Brown, D. Williams, R. Millikan (2007). *Journal of the National Cancer Institute*, 99(21):1577–1582, 2007.

P.J. Bickel, C.A.J. Klaassen, Y. Ritov, and J. Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. Springer-Verlag, 1997.

J. Bryan, Z. Yu, and M.J. van der Laan. Analysis of longitudinal marginal structural models. *Biostatistics*, 5(3):361–380, 2003.

COBRA
Collection of Biostatistics
Research Archive

S. Dudoit and M.J. van der Laan. Asymptotics of cross-validated risk estimation in estimator selection and performance assessment. *Statistical Methodology*, 2(2):131–154, 2005.

R. Gill and J.M. Robins. Causal inference in complex longitudinal studies: continuous case. *Ann. Stat.*, 29(6), 2001.

R.D. Gill, M.J. van der Laan, and J.M. Robins. Coarsening at random: characterizations, conjectures and counter-examples. In D.Y. Lin and T.R. Fleming, editors, *Proceedings of the First Seattle Symposium in Biostatistics*, pages 255–94, New York, 1997. Springer Verlag.

D.F. Heitjan and D.B. Rubin. Ignorability and coarse data. *Annals of statistics*, 19(4):2244–2253, December 1991.

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.

M. Jacobsen and N. Keiding. Coarsening at random in general sample spaces and random censoring in continuous time. *Annals of Statistics*, 23:774–86, 1995.

J. Kang and J. Schafer. Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data (with discussion). *Statistical Science*, 22:523–39, 2007a.

J. Kang and J. Schafer. Rejoinder: Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical Science*, 22:574–80, 2007b.

S. Keleş, M. van der Laan, and S. Dudoit. Asymptotically optimal model selection method for regression on censored outcomes. *Technical Report, Division of Biostatistics, UC Berkeley*, 2002.

K.L. Moore and M.J. van der Laan. Covariate adjustment in randomized trials with binary outcomes. Technical report 215, Division of Biostatistics, University of California, Berkeley, April 2007.

K.L. Moore and M.J. van der Laan. Application of time-to-event methods in the assessment of safety in clinical trials. In Karl E. Peace, editor, *in Design, Summarization, Analysis & Interpretation of Clinical Trials with Time-to-Event Endpoints*. Chapman and Hall, 2009.

UC Berkeley Division of Biostatistics Working Paper Series
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

S.A. Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B*, 65(2), 2003.

S.A. Murphy, M.J. van der Laan, and J.M. Robins. Marginal mean models for dynamic treatment regimens. *Journal of the American Statistical Association*, 96:1410–1424, 2001.

R. Neugebauer and M.J. van der Laan. Why prefer double robust estimates. *Journal of Statistical Planning and Inference*, 129(1-2):405–426, 2005.

Maya L. Petersen, Steven G. Deeks, Jeffrey N. Martin, and Mark J. van der Laan. History-adjusted marginal structural models: Time-varying effect modification and dynamic treatment regimens. Technical report 199, Division of Biostatistics, University of California, Berkeley, December 2005.

E.C. Polley and M.J. van der Laan. Predicting optimal treatment assignment based on prognostic factors in cancer patients. In Karl E. Peace, editor, *in Design, Summarization, Analysis & Interpretation of Clinical Trials with Time-to-Event Endpoints*. Chapman and Hall, 2009.

G. Ridgeway and D. McCaffrey. Comment: Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data (with discussion). *Statistical Science*, 22:540–43, 2007.

J. Robins, L. Orallana, and A. Rotnitzky. Estimaton and extrapolation of optimal treatment and testing strategies. *Statistics in Medicine*, 27(23): 4678–4721, 2008.

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, 2001a.

J. M. Robins, M. A. Hernan, and B. Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000a.

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, 2000b.

J. M. Robins, M. Sued, Q. Lei-Gomez, and A. Rotnitzky. Comment: Performance of double-robust estimators when "inverse probability" weights are highly variable. *Statistical Science*, 22:544–559, 2007.

Collection of Biostatistics
Research Archive

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. Discussion of "optimal dynamic treatment regimes" by susan a. murphy. *Journal of the Royal Statistical Society: Series B*, 65(2):355–366, 2003.

J.M. Robins. Optimal structural nested models for optimal sequential decisions. In Heagerty P.J Lin, D.Y, editor, *Proceedings of the 2nd Seattle symposium in biostatistics*, volume 179, pages 189–326, 2005a.

J.M. Robins. Optimal structural nested models for optimal sequential decisions. Technical report, Department of Biostatistics, Havard University, 2005b.

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. The analysis of randomized and non-randomized aids treatment trials using a new approach in causal inference in longitudinal studies. In L. Sechrest, H. Freeman, and A. Mulley, editors, *Health Service Methodology: A Focus on AIDS*, pages 113–159. U.S. Public Health Service, National Center for Health SErvices Research, Washington D.C., 1989.

J.M. Robins. Information recovery and bias adjustment in proportional hazards regression analysis of randomized trials using surrogate markers. In *Proceeding of the Biopharmaceutical section*, pages 24–33. American Statistical Association, 1993.

J.M. Robins. Causal inference from complex longitudinal data. In Editor M. Berkane, editor, *Latent Variable Modeling and Applications to Causality*, pages 69–117. Springer Verlag, New York, 1997a.

J.M. Robins. Structural nested failure time models. In P. Armitage, T. Colton, P.K. Andersen, and N. Keiding, editors, *The Encyclopedia of Biostatistics*. John Wiley and Sons, Chichester, UK, 1997b.

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.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

J.M Robins and A. Rotnitzky. Comment on Inference for semiparametric models: some questions and an answer, by Bickel, P.J. and Kwon, J. *Statistica Sinica*, 11:920–935, 2001b.

J.M. Robins and A. Rotnitzky. Recovery of information and adjustment for dependent censoring using surrogate markers. In *AIDS Epidemiology*, Methodological issues. Bikhäuser, 1992.

S. Rose and M.J. van der Laan. Simple optimal weighting of cases and controls in case-control studies. *The International Journal of Biostatistics*, page http://www.bepress.com/ijb/vol4/iss1/19/., 2008.

S. Rose and M.J. van der Laan. Why match? investigating matched case-control study designs with causal effect estimation. *The International Journal of Biostatistics*, page http://www.bepress.com/ijb/vol5/iss1/1/., 2009.

P.R. Rosenbaum and D.B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70:41–55, 1983.

M. Rosenblum, S.G. Deeks, M.J. van der Laan, and D.R. Bangsberg. The risk of virologic failure decreases with duration of hiv suppression, at greater than 50% adherence to antiretroviral therapy. *PLoS ONE*, 4(9): e7196.doi:10.1371/journal.pone.0007196, 2009.

D.B. Rubin. *Matched Sampling for Causal Effects.* Cambridge University Press, Cambridge, MA, 2006.

D.B. Rubin and M.J. van der Laan. 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, 2008.

J.S. Sekhon. Multivariate and propensity score matching software with automated balance optimization: The matching package for R. *Journal of Statistical Sotware, Forthcoming*, 2008.

S. Sinisi and M.J. van der Laan. The deletion/substitution/addition algorithm in loss function based estimation: Applications in genomics. *Journal of Statistical Methods in Molecular Biology*, 3(1), 2004.

Z. Tan. Comment: Understanding or, ps and dr. *Statistical Science*, 22: 560–568, 2007.

CORRA
A SERVICE OF THE REPOSITORY
Collection of Biostatistics
Research Archive

A. Tsiatis and M. Davidian. Comment: Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data (with discussion). *Statistical Science*, 22:569–73, 2007.

C. Tuglus and M.J. van der Laan. Targeted methods for biomarker discovery, the search for a standard. *UC Berkeley Working Paper Series*, page http://www.bepress.com/ucbbiostat/paper233/., 2008.

M.J. van der Laan. Causal effect models for intention to treat and realistic individualized treatment rules. Technical report 203, Division of Biostatistics, University of California, Berkeley, 2006.

M.J. van der Laan. Estimation based on case-control designs with known prevalance probability. *The International Journal of Biostatistics*, page http://www.bepress.com/ijb/vol4/iss1/17/, 2008a.

M.J. van der Laan. The construction and analysis of adaptive group sequential designs. Technical report 232, Division of Biostatistics, University of California, Berkeley, March 2008b.

M.J. van der Laan and S. Dudoit. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples. Technical report, Division of Biostatistics, University of California, Berkeley, November 2003.

M.J. van der Laan and M.L. Petersen. Causal effect models for realistic individualized treatment and intention to treat rules. *International Journal of Biostatistics*, 3(1), 2007.

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, S. Dudoit, and S. Keles. Asymptotic optimality of likelihood-based cross-validation. *Statistical Applications in Genetics and Molecular Biology*, 3, 2004.

M.J. van der Laan, M.L. Petersen, and M.M. Joffe. History-adjusted marginal structural models and statically-optimal dynamic treatment regimens. *The International Journal of Biostatistics*, 1(1):10–20, 2005.

Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

M.J. van der Laan, S. Dudoit, and A.W. van der Vaart. The cross-validated adaptive epsilon-net estimator. *Statistics and Decisions*, 24(3):373–395, 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 on targeted maximum likelihood estimation. *Technical report, working paper series http://www.bepress.com/ucbbiostat/paper254/*, September, 2009.

A. van der Vaart and J. Wellner. *Weak Convergence and Empirical Processes*. Springer-Verlag, New York, 1996.

A.W. van der Vaart, S. Dudoit, and M.J. van der Laan. Oracle inequalities for multi-fold cross-validation. *Statistics and Decisions*, 24(3):351–371, 2006.

Z. Yu and M.J. van der Laan. Construction of counterfactuals and the G-computation formula. Technical report, Division of Biostatistics, University of California, Berkeley, 2002.

Z. Yu and M.J. van der Laan. Double robust estimation in longitudinal marginal structural models. Technical report, Division of Biostatistics, University of California, Berkeley, 2003.

References

A. Abadie and G.W. Imbens. Large sample properties of matching estimators for average treatment effects. *Econometrica*, 74:235–67, 2006.

O. Bembom, J.W. Fessel, R.W. Shafer, and M.J. van der Laan. Data-adaptive selection of the adjustment set in variable importance estimation. Technical report, DigitalCommons@Florida Atlantic University [http://digitalcommons.fau.edu/cgi/oai2.cgi] (United States), 2008. URL `http://www.bepress.com/ucbbiostat/paper231`.

O. Bembom, M.L. Petersen, S.-Y. Rhee, W. J. Fessel, S.E. Sinisi, R.W. Shafer, and M.J. van der Laan. Biomarker discovery using targeted maximum likelihood estimation: Application to the treatment of antiretroviral resistant HIV infection. *Statistics in Medicine*, 28:152–72, 2009.

Collection of Biostatistics
Research Archive

S. Gruber and M.J. van der Laan. A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. Technical report 265, Division of Biostatistics, University of California, Berkeley, May 2010.

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.H. Holland and J.S. Reitman. Cognitive systems based on adaptive algorithms. *SIGART Bull.*, 63:49–49, 1977. ISSN 0163-5719. doi: http://doi.acm.org/10.1145/1045343.1045373.

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.

P.R. Rosenbaum and D.B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70:41–55, 1983.

D.B. Rubin and M.J. van der Laan. Doubly robust ecological inference. Technical report 236, Division of Biostatistics, University of California, Berkeley, May 2008.

J.S. Sekhon. Multivariate and propensity score matching software with automated balance optimization: The matching package for R. *Journal of Statistical Software, Forthcoming*, 2008.

COBRA
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

S. Sinisi and M.J. van der Laan. The Deletion/Substitution/Addition algorithm in loss function based estimation: Applications in genomics. *Journal of Statistical Methods in Molecular Biology*, 3(1), 2004.

M.J. van der Laan and S. Gruber. Collaborative double robust penalized targeted maximum likelihood estimation. *The International Journal of Biostatistics*, 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.

Appendix

TMLE as an imputation estimator

Consider the goal of estimating EY_1. It is desirable for the estimator Q_n to satisfy $0 = \sum_i \{E_{Q_n}(Y_1|O_i) - P_{Q_n}(Y_1 = 1)\}$, which makes $\Psi(Q_n)$ an imputation estimator. In other words, averaging the imputed values $E_{Q_n}(Y_1|O_i)$ for Y_1 under Q_n for $i = 1, .., n$, gives the same estimator as the mean of Y_1 under Q_n. Below, we show that this holds if Q_n solves the score equation $\sum_i A_i(Y_i - E_{Q_n}(Y|A_i, W_i))$. With this insight, we can now compute a targeted maximum likelihood estimator that is also an imputation estimator, by using a bivariate ϵ fluctuation function involving both the clever covariate and the term A. Similarly, we can construct a targeted maximum likelihood estimator that is also an imputation estimator for the additive causal effect parameter $EY(1) - Y(0)$. This problem was presented in Rubin and van der Laan (2008).

Details of this derivation follow. Let $O = (W, A, Y)$, with Y binary. Consider the score $D(Q)(O) = E_Q(Y_1 \mid O) - E_Q Y_1$. We note that

$$E_Q(Y_1 \mid O) = I(A = 1)Y + I(A = 0)E_Q(Y \mid A = 1, W).$$

We wish to find the component of $D(Q)$ that is in the tangent space of the conditional distribution of Y, given A, W. We have

$$E(D(Q) \mid A, W) = E(Y \mid A = 1, W).$$

Collection of Biostatistics Research Archive

We decompose

$$\begin{aligned} D(Q) &= D_1(Q) + D_2(Q) \\ &= \{D(Q) - E(Y \mid A=1, W)\} + \{E(Y \mid A=1, W) - E_Q Y_1\}. \end{aligned}$$

The first component can be written as

$$\{E(D_1(Q) \mid Y=1, A, W) - E(D_1(Q) \mid Y=0, A, W)\}(Y - E_Q(Y \mid A, W)),$$

which reduces to $A(Y - E_Q(Y \mid A, W))$.

Consider now the score $D(Q)(O) = E_Q(Y_0 \mid O) - E_Q Y_0$. We note that

$$E_Q(Y_0 \mid O) = I(A=1)E_Q(Y \mid A=0, W) + I(A=0)Y.$$

We wish to find the component of $D(Q)$ that is in the tangent space of the conditional distribution of Y, given A, W. We have

$$E(D(Q) \mid A, W) = E(Y \mid A=0, W).$$

We decompose

$$D(Q) = \{D(Q) - E(Y \mid A=0, W)\} + \{E(Y \mid A=0, W) - E_Q Y_0\}.$$

We can write this first component as

$$\{E(D_1(Q) \mid Y=1, A, W) - E(D_1(Q) \mid Y=0, A, W)\}(Y - E_Q(Y \mid A, W)),$$

which reduces to $(1-A)(Y - E_Q(Y \mid A, W))$.

Suppose now that one wishes to solve the score equation of score $D(Q) = E_Q(Y_1 - Y_0 \mid O_i) - E_Q(Y_1 - Y_0)$. We follow the same proof as above. Firstly,

$$\begin{aligned} E_Q(Y_1 - Y_0 \mid O) &= I(A=1)\{Y - E_Q(Y \mid A=0, W)\} \\ &\quad + I(A=0)\{E_Q(Y \mid A=1, W) - Y\}. \end{aligned}$$

We wish to find the component of $D(Q)$ that is in the tangent space of the conditional distribution of Y, given A, W. We have $E(D(Q) \mid A, W) = E_Q(Y \mid A=1, W) - E_Q(Y \mid A=0, W))$. We decompose

$$D(Q) = D_1(Q) + D_2(Q) \equiv \{D(Q) - E(D(Q) \mid A, W)\} + \{E(D(Q) \mid A, W) - E_Q Y_1\}.$$

We can write this first component as

$$\{E(D_1(Q) \mid Y=1, A, W) - E(D_1(Q) \mid Y=0, A, W)\}(Y - E_Q(Y \mid A, W)),$$

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

which reduces to $(2A-1)(Y - E_Q(Y \mid A, W)$.

Therefore, if we want Q_n^* to be an imputation estimator for both EY_0, EY_1, then we wish to have an estimator Q_n^* that solves the bivariate score equation

$$0 = \sum_i (1, A_i)(Y_i - E_{Q_n^*}(Y \mid W_i, A_i)).$$

This can be arranged by applying the targeted MLE update with the fluctuation function corresponding with covariate-extension $\epsilon(1, A, h_g(A, W))$, where h_g denotes the clever covariate for the target parameter (say) $EY_1 - EY_0$. If one wishes to solve the score equation of score $D(Q) = E_Q(Y_1 - Y_0 \mid O_i) - E_Q(Y_1 - Y_0)$, then, one uses the bivariate extension $\epsilon_1(2A-1) + \epsilon_2 h_g$.

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive