

History-Adjusted Marginal Structural Models and Statically-Optimal Dynamic Treatment Regimes

Mark J. van der Laan*

Maya L. Petersen[†]

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

[†]Division of Epidemiology, School of Public Health, University of California, Berkeley, mayaliv@hotmail.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://www.bepress.com/ucbbiostat/paper158>

Copyright ©2004 by the authors.

History-Adjusted Marginal Structural Models and Statically-Optimal Dynamic Treatment Regimes

Mark J. van der Laan and Maya L. Petersen

Abstract

Marginal structural models (MSM) provide a powerful tool for estimating the causal effect of a treatment. These models, introduced by Robins, model the marginal distributions of treatment-specific counterfactual outcomes, possibly conditional on a subset of the baseline covariates. Marginal structural models are particularly useful in the context of longitudinal data structures, in which each subject's treatment and covariate history are measured over time, and an outcome is recorded at a final time point. However, the utility of these models for some applications has been limited by their inability to incorporate modification of the causal effect of treatment by time-varying covariates. Particularly in the context of clinical decision making, such time-varying effect modifiers are often of considerable or even primary interest, as they are used in practice to guide treatment decisions for an individual. In this article we propose a generalization of marginal structural models, which we call history-adjusted marginal structural models (HA-MSM). These models allow estimation of adjusted causal effects of treatment, given the observed past, and are therefore more suitable for making treatment decisions at the individual level and for identification of time-dependent effect modifiers. Specifically, a HA-MSM models the conditional distribution of treatment-specific counterfactual outcomes, conditional on the whole or a subset of the observed past up till a time-point, simultaneously for all time-points. Double robust inverse probability of treatment weighted estimators have been developed and studied in detail for standard MSM. We extend these results by proposing a class of double robust inverse probability of treatment weighted estimators for the unknown parameters of the HA-MSM. In addition, we show that HA-MSM provide a natural approach to identifying the dynamic treatment regime which follows, at each time-point, the history-adjusted (up till the most recent time point) optimal static treatment regime. We illustrate our results using an example drawn from the treatment of HIV infection.

1 Introduction.

The practice of much of modern medicine involves monitoring patient covariates over time and using these measurements to make treatment decisions. For example, patient risk factors and sequential measurements of blood pressure inform when antihypertensive medication is initiated, dose of antidepressant medication is often modified in response to changes in patient symptoms and side effects, and the decision to switch antiretroviral therapy regimen for an HIV-infected patient is based on measures of viral and patient response to treatment. Understanding how to use past covariates measured on a patient to make the best immediate treatment decision is central to clinical practice.

In the statistical literature, strategies or rules for changing treatment in response to changing covariates are termed "dynamic treatment regimens". Dynamic treatment regimens can be contrasted with static treatment regimens, which allow treatment to change over time but not in response to changing patient covariates. While dynamic treatment regimens are often of primary clinical interest, the focus of the majority of both randomized trials and observational studies has been estimation of the effects of static treatment regimens.

The discrepancy between the questions posed by clinical medicine and the questions currently being answered by epidemiologic research may be due in part to a lack of accessible analytic tools for estimating the causal effects of dynamic treatment regimens and for identifying dynamic treatment regimens that will optimize expected patient outcome. Two analytic methods currently available, Structural Nested Mean Models and G-computation (Murphy (2003), Robins (1997), Robins (2000)) both require substantial modeling and programming effort to implement.

Marginal Structural Models (MSM) are a well-established and powerful tool for causal inference. In longitudinal studies, MSM allow for control of confounders that are also affected by previous treatment (time-dependent confounders), and that are thus not amenable to traditional analytic approaches to confounding such as multivariable regression. In addition, MSM can be more robust to model misspecification, particularly when data are high dimensional, and can be implemented using standard software. However, until recently MSM could not identify or estimate modification of the causal effect of treatment by covariates that change over time. As a result, MSM have been able to identify optimal static treatment regimens, but have not been applied to identify optimal dynamic treatment regimens.

In a companion paper, we introduce a generalization of Marginal Structural Models, called History-Adjusted Marginal Structural Models (HA-MSM), that allow for the identification and estimation of causal effect modification by time-varying covariates in longitudinal data settings with time-dependent confounding. These models can be implemented using weighted least squares regression and standard software. In this article, we illustrate how HA-MSM can be used to identify a rule for making treatment decisions over time, based on current history of time-varying covariates, that represents a specific type of optimal dynamic treatment regimen. This dynamic treatment regimen corresponds closely to the needs of clinical practitioners, in that it allows a practitioner to use a patient's most recent measured covariates to update, at each patient visit, the future (static) treatment plan that will maximize the patient's expected long-term outcome. We use an example drawn from the treatment of human immunodeficiency virus (HIV) infection to illustrate how HA-MSM can be used to identify and estimate this type of optimal dynamic treatment regimen, and how this dynamic treatment regimen can be interpreted.

2 Antiretroviral therapy for the treatment of HIV infection: When to switch treatment regimens?

In this article, we address the issue of how long an HIV-infected patient who is virologically failing his/her antiretroviral therapy (ART) regimen should wait before switching to a new regimen. Switching to a new suppressive regimen immediately will prevent the potential decline in CD4 T-cells associated with ongoing viral replication. However, delay also runs the risk of prematurely depleting future treatment options. In addition, some patients may experience continued immunologic benefit from regimens that do not achieve complete virologic suppression (Piketty et al. (2001), Deeks et al. (2002), Deeks et al. (2000)). It has been hypothesized that this phenomenon is due to a decrease in viral fitness that occurs as a result of viral mutations which confer resistance to some antiretroviral therapies (Deeks et al. (2002)). Switching too early for these patients may represent a switch from a moderately effective but potentially long-lived regimen to a more effective but shorter-lived regimen. Finally, immunologic benefits conferred by

regimens despite incomplete virologic suppression may attenuate over time, with the accumulation of additional resistance mutations that restore viral fitness (Deeks (2003)). The optimal time to switch may change as a result of time-varying covariates. For example, the emergence over time of certain viral mutations associated with increased resistance or improved viral fitness may increase the depletion of CD4 T-cells that results from ongoing viral replication, and thus alter the optimal switch time (Deeks (2001)).

To illustrate the dynamic treatment regimen identified by HA-MSM, we use the following data structure. HIV-infected subjects become eligible for our analysis when they virologically rebound on an ART regimen. On each randomly sampled subject, we measure the following data at each time point during the study:

1. Treatment status, a binary variable indicating whether or not the subject has switched ART regimens. We denote treatment status at time point t as $A(t)$, where $A(t) = 1$ as long as a subject remains on his/her original non-suppressive therapy, and $A(t) = 0$ from the time point when a subject switches to a new therapy until the end of the study. In order to simplify presentation, we assume in this paper that subjects can switch therapy only once.
2. Covariates, denoted $L(t)$. These may include covariates measured only at baseline, $L(0)$, as well as time-dependent covariates. We focus here on one time-dependent covariate of particular interest, the presence of specific viral mutations thought to improve viral fitness, which we denote using the binary variable $S(t)$. Covariates measured on a subject also include CD4 T-cell count. Our outcome of interest for a given time point t is CD4 T-cell count m months in the future, which we denote $Y(t + m)$.

The observed data for each randomly sampled subject can be represented chronologically as:

$$O = (L(0), A(0), \dots, L(K), A(K), L(K + 1)) \quad (1)$$

where we assume that $L(t)$ occurs before $A(t)$ and $K + 1$ denotes the end of the study

3 Review of History-Adjusted Marginal Structural Models.

In this section, we briefly review the assumptions and statistical framework for History-Adjusted Marginal Structural Models. For a thorough introduction to the methodology, we refer the reader to our companion papers.

HA-MSM are a generalization of standard Marginal Structural Models, as introduced by Robins (Robins (2000), Robins (1999)). Both methods rely on the identical counterfactual framework for causal inference, which assumes that each individual has counterfactual outcomes and other covariate processes corresponding to each possible treatment history. Applied to the HIV example, each individual has a set of counterfactuals corresponding to the CD4 T-cell count, viral mutations, etc., over time that would have been observed for that individual under each possible time until switching therapy.

This set of counterfactuals is the Full Data for a given individual, denoted:

$$X^{Full} = (\bar{L}_{\bar{a}}(K + 1), \bar{a} \in \mathcal{A}), \quad (2)$$

where \mathcal{A} denotes the set of possible treatment histories (in our example, possible times until switching to a new antiretroviral therapy), and $\bar{L}_{\bar{a}}(t)$ denotes the counterfactual covariate processes that would have been observed through time t under treatment history \bar{a} . We use the notation $\bar{X}(t) = (X(0), X(1), \dots, X(t))$ to denote the history of a variable up through time t .

We make the following assumptions:

- Consistency assumption (CA) The observed data for an individual are equal to the counterfactual covariate processes corresponding to the individual's observed treatment history

$$O = (\bar{L}(K + 1), \bar{A}(K)) = (\bar{L}_{\bar{A}}(K + 1), \bar{A}(K)) \quad (3)$$

- Temporal ordering assumption (TA) Covariates are not affected by treatment that occurs after they are measured

$$L_{\bar{a}}(t) = L_{\bar{a}(t-1)}(t) \quad (4)$$

- Sequential randomization assumption (SRA) Treatment at each time point is independent of the full data, given the observed past.

$$A(t) \perp X^{Full} | \bar{A}(t - 1), \bar{L}(t), t = 0, \dots, K \quad (5)$$

- Experimental treatment assignment assumption (ETA) For each time point, for every subpopulation defined by treatment and covariate history up to that time point, every possible treatment which is compatible with treatment history up till that time point must have positive probability of occurring.

$$\text{For all possible } \bar{A}(t-1), \bar{L}(t), P(A(t) = a(t) | \bar{A}(t-1), \bar{L}(t)) > 0 \quad (6)$$

$$\text{for all } a(t) \in a^*(t) : (\bar{A}(t-1), a^*(t)) \in \mathcal{A}(t)$$

Under these assumptions, Marginal Structural Models model the treatment-specific counterfactual outcome, possibly conditional on baseline covariates. History-Adjusted Marginal Structural Models extend this methodology by recognizing that, at each time point during a study, the covariate and treatment history up till that time-point can be considered baseline covariates. In other words, one could imagine fitting, at each time point, a separate model of the counterfactual outcome indexed by treatment after that time point, conditional on some subset of interest of treatment and covariate history up till that time point. However, in many settings fitting a separate model for each time point may be neither practical nor desirable; HA-MSM also allow one to assume a common model across time points.

Specifically, let $\underline{a}(j, j+m-1) \equiv (a(j), a(j+1), \dots, a(j+m-1))$ denote a future treatment regimen from time j until the outcome is measured. We denote the counterfactual outcome m months later if a subject were to follow his/her observed treatment history up till time $j-1$ ($\bar{A}(j-1)$) and a specified counterfactual future treatment regimen beginning at time j ($\underline{a}(j, j+m-1)$) as $Y_{\bar{A}(j-1)\underline{a}(j+m-1)}(j+m)$. These counterfactuals are defined for each time point $j \in \{0, \dots, K+1-m\}$ and for each possible treatment regimen, beginning at each time point j , that is consistent with a subject's treatment history up till time j .

HA-MSM estimate the expectation (or some other parameter) of these treatment-specific counterfactuals, indexed by future treatment regimens beginning at time j , conditional on a subset of covariate processes observed up till time j ($V(j) \subset (\bar{A}(j-1)\bar{L}(j))$)

$$E_0(Y_{\bar{A}(j-1)\underline{a}(j+m-1)}(j+m) | V(j)) \equiv \theta_0(j, \underline{a}(j, j+m-1) | V(j)), j = 0, \dots, K+1-m \quad (7)$$

We note that if the set of possible future treatment regimens depends on a subject's treatment history, treatment history must be included in $V(j)$.

For example, if only a single change in treatment regimen is allowed during the course of the study, whether or not such a change has already occurred will restrict whether or not it can occur in the future, and thus past history of treatment change must be included in $V(j)$. In our example, HA-MSM are concerned with estimating the conditional expectation of counterfactual CD4 T-cell count m months later among subjects who have not yet switched therapy, given observed treatment history and the current presence of viral mutations ($V(j) = (\bar{A}(j-1), S(j))$), if a subject were to follow his/her observed treatment up till time j (not switching) and then switch treatment at a specified time after j .

4 The HA-MSM dynamic treatment regimen: Updating the best future treatment plan based current covariate values.

A dynamic treatment regimen is a rule or function that gives a recommended treatment decision at each time point, based on patient characteristics measured up till the time point. For example, we are interested here in decision rules for when to switch to a new antiretroviral treatment, given time spent on current non-suppressive treatment and current viral mutation profile. We can represent a dynamic treatment regimen with a vector $d = (d_0, \dots, d_K)$, where d_t represents the treatment decision at time t (switch treatment or not). In our current example, say we wish to find the dynamic treatment regimen, beginning at time $j = 0$ that maximizes the mean outcome (CD4 T-cell count) at time m . In other words, we wish to maximize $E[Y_d(m)]$, for $t = 0, \dots, m-1$, and where $Y_d(m)$ is the CD4 T-cell count we would have observed at time m if the whole population had followed the dynamic treatment regimen d .

For each time point, HA-MSM allow us to identify the future static treatment regimen that, given treatment history and covariates of interest up till that time point, will maximize the expectation of the outcome. Applied to our example, History-Adjusted Marginal Structural Models can be used to estimate how much longer subjects should remain on their current non-suppressive therapy, given how long they have already been on it and their viral mutation profile, in order to maximize the expected CD4 T-cell count

m months later. Formally,

$$\arg \max_{\underline{a}(j, j+m-1) \in \underline{a}^*(j, j+m-1): (\bar{A}(j-1), \underline{a}^*(j, j+m-1)) \in \mathcal{A}} \underline{a}_0(j, j+m-1 | V(j)) \equiv \theta_0(j, \underline{a}(j, j+m-1) | V(j)) \quad (8)$$

Note that $\underline{a}_0(j, j+m-1 | V(j))$ is the optimal future *static* treatment regimen at time j ; it represents a fixed point in the future, within subgroups defined by treatment and covariate history at that time point, at which treatment should be switched to maximize CD4 T-cell count m months later.

The optimal future static treatment regimen estimated by HA-MSM (8) immediately suggests an interesting *dynamic* treatment regimen. At each time point $j, \dots, K+1-m$, HA-MSM identify the best future static treatment regimen given a subset of an individual's treatment history and covariates up till that time point (where "best" is defined as the future static regimen that is expected to maximize outcome at time $j+m$). The dynamic treatment regimen consists of following, at each time point, the first action of an individual's optimal static future treatment regimen at that time point. At subsequent time points, the optimal static future treatment regimen can then be updated in response to changes in covariates and treatment history. Formally, the dynamic treatment regimen given by HA-MSM consists of

$$d_0(V(j)) \equiv a_0(j, | V(j)), j = 0, \dots, K+1-m. \quad (9)$$

That is, $d_0(V(j))$ equals the first treatment action of the history- (at time j) adjusted optimal static treatment regimen $\underline{a}_0(j, j+m-1 | V(j))$.

Applied to the HIV example, $d_0(V(j))$ represents a dynamic treatment regimen in which subjects either switch treatment or not at each time point, based on the history-adjusted optimal static future treatment regimen at that time point. This approach yields for each subject a vector of recommended treatment actions (switch or not at each time point), $d_0 = (d_{00}, \dots, d_{0(K+1-m)})$ that can depend on the subject's current measured covariates. In practice, applying this dynamic treatment regimen consists of following a patient on non-suppressive therapy over time, and at each follow-up visit, using the patient's current treatment history and viral mutations to choose the future treatment plan (switch time) that will maximize the patient's expected CD4 T-cell count m months later. This treatment plan is then followed until the next visit, at which point the treatment plan can be revised based on new covariate values. Note that this does indeed constitute a dynamic treatment

regimen, in that the treatment decision made at each time point can change in response to changes in patient covariates.

The dynamic treatment regimen (9) identified by HA-MSM can be compared to two common types of optimal treatment regimens discussed in the literature. The *optimal static treatment regimen* (as can be estimated with a standard MSM), as applied to the HIV example is the total duration of non-suppressive therapy (measured from a fixed time point, such as the beginning of the study, $j = 0$) within subgroups defined by *baseline* covariates that would result in the highest mean CD4 T-cell count in the population m months later. In other words, the optimal static treatment regimen indicates how long subjects should remain on non-suppressive therapy in total, given only their characteristics at the beginning of the study. Unlike the history-adjusted dynamic treatment regimen, the decision about when to switch therapy cannot be modified over the course of follow-up in response to changing patient characteristics.

In contrast the *optimal dynamic treatment regimen* as defined by Murphy (2003), Robins (2003) optimizes a dynamic treatment regimen-specific marginal outcome (at a fixed time point) over a collection of dynamic treatment regimens. Applied to the HIV example, the optimal dynamic treatment regimen is the decision rule for when to switch therapy, based on how long subjects have already been on current therapy and viral mutation profile, that, if applied to all members of the population, would result in the highest mean CD4 T-cell count m months after the start of the study. This differs from the HA-MSM dynamic treatment regimen, which identifies at each time point the first action of the optimal static treatment regimen. Certainly, this is an interesting dynamic treatment regimen, and has the advantage of being easier to estimate than the truly optimal dynamic treatment regimen. Finally, it is interesting to note that in the current example we identify a dynamic treatment regimen aimed at optimizing a continuously changing outcome (CD4 T-cell count m months in the future). In many applications, optimizing such a "moving" outcome, rather than an outcome at a fixed time point, may indeed be the researchers goal.

5 Using HA-MSM to identify the best time to switch therapy

Implementation of HA-MSM is presented in detail in our companion papers. Here we illustrate how the parameters estimated by HA-MSM can be used to identify the optimal history-adjusted treatment regimen (9), and how that regimen can be interpreted. We are interested in the question “Given a subject’s current viral mutation profile and the length of time he/she has already been on the current non-suppressive ART regimen, when should a subject switch therapy?” We remind the reader that, because in our example a subject is only allowed to switch therapy once, the decision of when to switch will only apply to those individuals who have not switched already.

We begin by proposing a simple HA-MSM model for the effect of total exposure time on final CD4 T-cell count with effect modification by current mutation profile. We summarize the $\underline{a}(j, j + m - 1)$ as the duration from time j until a subject switches therapy: $c(j) \equiv \sum_{l=j}^{j+m-1} a(l)$. Consider the model

$$E_0(Y_{\bar{A}(j-1), \underline{a}(j, j+m-1)}(j+m) \mid A(j-1) = 1, S(j)) \\ = I(A(j-1) = 1)(\beta_0 + \beta_1 c(j) + \beta_2 S(j) + \beta_3 j + \beta_4 c(j) S(j) + \beta_5 c(j) j),$$

$j = 0, \dots, K + 1 - m$. Note that we condition on not already having switched therapy $I(A(j-1) = 1)$, and that, among this population, j is the elapsed time a subject has already spent on non-suppressive therapy.

In this linear model, the history-adjusted optimal static treatment regimen is relatively straightforward. To choose the future treatment regimen that maximizes the expected CD4 T-cell count, we simply look at our estimates of β . Among subjects with key viral mutations ($S(j) = 1$), the estimated effect of each additional month waiting to switch therapy is: $\beta_1 + \beta_4 + \beta_5(j)$; similarly, among subjects without viral mutations ($S(j) = 0$), the estimated effect of each additional month waiting to switch is $\beta_1 + \beta_5(j)$. If the goal is to maximize CD4 count m months later, when this sum is positive the best statically optimal treatment plan is to wait to switch, and when this sum becomes negative, the best plan is to switch immediately. The resulting dynamic treatment regimen suggested is thus the following:

- For $\beta_5 < 0$, when a subjects has key viral mutations, switch therapy immediately if the subject has already been on non-suppressive therapy

for more than $(\beta_1 + \beta_4)/\beta_5$ months. If a subject does not have key viral mutations, switch therapy immediately if a subject has been on non-suppressive therapy for more than β_1/β_5 months. Note that this rule implies that, for $\beta_1, \beta_4, \beta_5$ all negative, never wait to switch regardless of a patient's current viral mutations or duration on non-suppressive therapy.

- For $\beta_5 > 0$ when a subject has key viral mutations, switch therapy if a subject has already been on non-suppressive therapy for less than $-(\beta_1 + \beta_4)/\beta_5$ months, If a subject does not have key viral mutations, switch therapy immediate if a subject has already been on non-suppressive therapy for less than $-\beta_1/\beta_5$ months. Note that this rule implies that, for $\beta_1, \beta_4, \beta_5$ all positive, always wait to switch regardless of a patients current viral mutations or duration on non-suppressive therapy.

6 Discussion


In our previous paper, we introduced history-adjusted marginal structural models, a generalization of marginal structural models that make possible the estimation of time-dependent effect modification. Here, we have illustrated how history-adjusted marginal structural models also make possible the identification of a specific type of dynamic treatment regimen. HA-MSM allow us to identify treatment decision rules that can be based on time-dependent covariates, a crucial application in the medical sciences.

7 Acknowledgements

Maya Petersen is supported by a Predoctoral Fellowship from the Howard Hughes Medical Institute.

References

- S. Deeks, J. Barbour, Martin J., Swanson M., and Grant R. Sustained CD4+ T cell response after virologic failure of protease inhibitor-based regimens



- in patients with human immunodeficiency virus infection. *J Infect Dis*, 181(3):946–53, 2000.
- S. G. Deeks. Virologic and immunologic consequences of discontinuing combination antiretroviral drug therapy in HIV-infected patients with detectable viremia. *New England Journal of Medicine*, 344(7):472–480, 2001.
- S. G. Deeks. Treatment of antiretroviral resistant HIV-1 infection. *Lancet*, 362(9400):2002–2011, 2003.
- S. G. Deeks, J. D. Barbour, R. M. Grant, and J. M. Martin. Duration and predictors of CD4 T-cell gains in patients who continue combination therapy despite detectable plasma viremia. *AIDS*, 16(2):201–207, 2002.
- S.A. Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B*, 65(2):331–584, 2003.
- C. Piketty, L. Weiss, F. Thomas, A. S. Mohamed, L. Belc, and M. D. Kazatchkine. Long-term clinical outcome of Human Immunodeficiency Virus-infected patients with discordant immunologic and virologic responses to a protease inhibitor-containing regimen. *Journal of Infectious Diseases*, 183(9):1328–1335, 2001.
- 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, 1997.
- J.M. Robins. Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association: Section on Bayesian Statistical Science*, pages 6–10, 1999.
- 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, 2000.
- 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.