



Published in final edited form as:

Biometrika. 2013 ; 100(3): . doi:10.1093/biomet/ast014.

Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions

Baqun Zhang,

Department of Preventive Medicine, 680 N. Lakeshore Drive, Suite 1400 Northwestern University, Chicago, Illinois, 60611 U.S.A

Anastasios A. Tsiatis,

Department of Statistics, North Carolina State University, Raleigh, North Carolina, 27695-8203, U.S.A

Eric B. Laber, and

Department of Statistics, North Carolina State University, Raleigh, North Carolina, 27695-8203, U.S.A

Marie Davidian

Department of Statistics, North Carolina State University, Raleigh, North Carolina, 27695-8203, U.S.A

Baqun Zhang: baqun.zhang@northwestern.edu; Anastasios A. Tsiatis: tsiatis@ncsu.edu; Eric B. Laber: eblaber@ncsu.edu; Marie Davidian: davidian@ncsu.edu

Summary

A dynamic treatment regime is a list of sequential decision rules for assigning treatment based on a patient's history. Q- and A-learning are two main approaches for estimating the optimal regime, i.e., that yielding the most beneficial outcome in the patient population, using data from a clinical trial or observational study. Q-learning requires postulated regression models for the outcome, while A-learning involves models for that part of the outcome regression representing treatment contrasts and for treatment assignment. We propose an alternative to Q- and A-learning that maximizes a doubly robust augmented inverse probability weighted estimator for population mean outcome over a restricted class of regimes. Simulations demonstrate the method's performance and robustness to model misspecification, which is a key concern.

Keywords

A-learning; Double robustness; Outcome regression; Propensity score; Q-learning

1. Introduction

Treatment of patients with chronic disease involves a series of decisions, where the clinician determines the next treatment to be administered based on all information available to that point. A dynamic treatment regime is a set of sequential decision rules, each corresponding to a decision point in the treatment process. Each rule inputs the available information and

outputs the treatment to be given from among the possible options. The optimal regime is that yielding the most favorable outcome on average if followed by the patient population.

Q- and A-learning are two main approaches for estimating the optimal dynamic treatment regime using data from a clinical trial or observational study. Q-learning (Watkins & Dayan, 1992) involves postulating at each decision point regression models for outcome as a function of patient information to that point. In A-learning (Robins, 2004; Murphy, 2003), models are posited only for the part of the regression involving contrasts among treatments and for treatment assignment at each decision point. Both are implemented through a backward recursive fitting procedure based on a dynamic programming algorithm (Bather, 2000). Under certain assumptions and correct specification of these models, Q- and A-learning lead to consistent estimation of the optimal regime. See Rosthøj et al. (2006), Murphy et al. (2007), Zhao et al. (2009) and Henderson et al. (2010) for applications; related methods are discussed by Robins (2004), Moodie et al. (2007), Robins et al. (2008), Almirall et al. (2010) and Orellana et al. (2010).

A concern with both Q- and A-learning is the effect of model misspecification on the quality of the estimated optimal regime. If one attempts to circumvent this difficulty by using flexible nonparametric regression techniques (Zhao et al., 2009), the estimated optimal rules may be complicated functions of possibly high-dimensional patient information that are difficult to interpret or implement and hence are unappealing to clinicians wary of black box approaches.

Given these drawbacks, we focus on a restricted class of treatment regimes indexed by a finite number of parameters, where the form of regimes in the class may be derived from posited regression models or prespecified on the grounds of interpretability or cost to depend on key subsets of patient information. Zhang et al. (2012) proposed an approach for estimating the optimal regime within such a restricted class for a single treatment decision based on maximizing directly a doubly robust augmented inverse probability weighted estimator for the population mean outcome over all regimes in the class, assuming that larger outcomes are preferred. Via the double robustness property, the estimated optimal regimes enjoy protection against model misspecification and comparable or superior performance than do competing methods. With judicious choice of the augmentation term, increased efficiency of estimation of the mean outcome is achieved, which translates into more precise estimators for the optimal regime.

We adapt this approach to two or more decision points. This is considerably more complex than for one decision and is based on casting the problem as one of monotone coarsening (Tsiatis, 2006, Chapter 7). We focus for simplicity on the case of two treatment options at each decision point, though the methods extend to a finite number of options. The methods lead to estimated optimal regimes achieving comparable performance to those derived via Q- or A-learning under correctly specified models and have the added benefit of protection against misspecification.

2. Framework

Assume there are K prespecified, ordered decision points and an outcome of interest, a function of information collected across all K decisions or ascertained after the K th decision, with larger values preferred. At each decision $k = 1, \dots, K$, there are two k -specific treatment options coded as 0,1 in the set of options \mathcal{A}_k ; write a_k to denote an element of \mathcal{A}_k . Denote a possible treatment history up to and including decision k as $\mathcal{A}_k = (a_1, \dots, a_k) \in \mathcal{A}_1 \times \dots \times \mathcal{A}_k = \mathcal{A}_k$.

We consider a potential outcomes framework. For a randomly chosen patient, let X_1 denote baseline covariates recorded prior to the first decision, and let $X_k^*(\bar{a}_{k-1})$ be the covariate information that would accrue between decisions $(k-1)$ and k were s/he to receive treatment history \bar{a}_{k-1} ($k=2, \dots, K$), taking values $x_k \in \mathcal{X}_k$. Let $Y^*(\bar{a}_K)$ be the outcome that would result were s/he to receive full treatment history \bar{a}_K . Then define the potential outcomes (Robins, 1986) as

$$W = \{X_1, X_2^*(a_1), \dots, X_K^*(\bar{a}_{K-1}), Y^*(\bar{a}_K) \text{ for all } \bar{a}_K \in \bar{\mathcal{A}}_K\}.$$

For convenience later, we include X_1 , which is always observed and hence is not strictly a potential outcome, in W , and write $\bar{X}_k^*(\bar{a}_{k-1}) = \{X_1, X_2^*(a_1), \dots, X_k^*(\bar{a}_{k-1})\}$ and $x_k = (x_1, \dots, x_k)$ for $k=1, \dots, K$, where then $\mathcal{X}_k = \mathcal{X}_1 \times \dots \times \mathcal{X}_k$.

A dynamic treatment regime $g = (g_1, \dots, g_K)$ is an ordered set of decision rules, where $g_k(x_k, \bar{a}_{k-1})$ corresponding to the k th decision takes as input a patient's realized covariate and treatment history up to decision k and outputs a treatment option $a_k \in \mathcal{A}_k(x_k, \bar{a}_{k-1}) \subseteq \mathcal{A}_k$. In general, $\mathcal{A}_k(x_k, \bar{a}_{k-1})$ is the set of feasible options at decision k for a patient with realized history (x_k, \bar{a}_{k-1}) , allowing that some options in \mathcal{A}_k may not be possible for patients with certain histories; here, $\mathcal{A}_k(x_k, \bar{a}_{k-1}) \subseteq \{0, 1\}$. Thus, a feasible treatment regime must satisfy $g_k(x_k, \bar{a}_{k-1}) \in \mathcal{A}_k(x_k, \bar{a}_{k-1})$ ($k=1, \dots, K$). Denote the class of all feasible regimes by \mathcal{G} .

For $g \in \mathcal{G}$ writing $\bar{g}_k = (g_1, \dots, g_k)$ for $k=1, \dots, K$ and $\bar{g}_K = g$, define the potential outcomes associated with g to be $W_g = \{X_1, X_2^*(g_1), \dots, X_K^*(\bar{g}_{K-1}), Y^*(g)\}$, where $X_k^*(\bar{g}_{k-1})$ is the covariate information that would be seen between decisions $k-1$ and k were a patient to receive the treatments dictated sequentially by the first $k-1$ rules in g , and $Y^*(g)$ is the outcome if s/he were to receive the K treatments determined by g . Thus, W_g is an element of W .

Define an optimal treatment regime $g^{\text{opt}} = (g_1^{\text{opt}}, \dots, g_K^{\text{opt}}) \in \mathcal{G}$ as satisfying

$$E\{Y^*(g^{\text{opt}})\} \geq E\{Y^*(g)\}, g \in \mathcal{G}. \quad (1)$$

That is, g^{opt} is a regime that maximizes expected outcome were all patients in the population to follow it. The optimal regime g^{opt} may be determined via dynamic programming, also referred to as backward induction. At the K th decision point, for any $x_K \in \mathcal{X}_K, \bar{a}_{K-1} \in \bar{\mathcal{A}}_{K-1}$, define

$$g_K^{\text{opt}}(\bar{x}_K, \bar{a}_{K-1}) = \arg \max_{a_K \in \Phi_K(\bar{x}_K, \bar{a}_{K-1})} E\{Y^*(\bar{a}_{K-1}, a_K) | \bar{X}_K^*(\bar{a}_{K-1}) = \bar{x}_K\}, \quad (2)$$

$$V_K(\bar{x}_K, \bar{a}_{K-1}) = \max_{a_K \in \Phi_K(\bar{x}_K, \bar{a}_{K-1})} E\{Y^*(\bar{a}_{K-1}, a_K) | \bar{X}_K^*(\bar{a}_{K-1}) = \bar{x}_K\}. \quad (3)$$

For $k=K-1, \dots, 2$ and any $x_k \in \mathcal{X}_k, \bar{a}_{k-1} \in \bar{\mathcal{A}}_{k-1}$, define

$$g_k^{\text{opt}}(\bar{x}_k, \bar{a}_{k-1}) = \arg \max_{a_k \in \Phi_k(\bar{x}_k, \bar{a}_{k-1})} E[V_{k+1}(\bar{x}_k, X_{k+1}^*(\bar{a}_{k-1}, a_k), \bar{a}_{k-1}, a_k) | \bar{X}_k^*(\bar{a}_{k-1}) = \bar{x}_k],$$

$$V_k(\bar{x}_k, \bar{a}_{k-1}) = \max_{a_k \in \Phi_k(\bar{x}_k, \bar{a}_{k-1})} E[V_{k+1}(\bar{x}_k, X_{k+1}^*(\bar{a}_{k-1}, a_k), \bar{a}_{k-1}, a_k) | \bar{X}_k^*(\bar{a}_{k-1}) = \bar{x}_k].$$

For $k = 1$, $x_1 \in \mathcal{X}_1$, $g_1^{\text{opt}}(x_1) = \arg \max_{a_1 \in \Phi_1(x_1)} E[V_2\{x_1, X_2^*(a_1), a_1\} | X_1 = x_1]$ and $V_1(x_1) = \max_{a_1 \in \Phi_1(x_1)} E[V_2\{x_1, X_2^*(a_1), a_1\} | X_1 = x_1]$. Thus, g_k^{opt} yields the treatment option at decision k that maximizes the expected potential outcome given prior covariate and treatment history. At decisions $k = K - 1, \dots, 1$, g_k^{opt} dictates the option that maximizes the expected potential outcome that would be achieved if the optimal rules were followed in the future. An argument that g^{opt} , so defined, satisfies (1) is given in an unpublished report by Schulte et al. (2013) available from the last author.

This definition of an optimal regime is intuitively given in terms of potential outcomes. In practice, with the exception of X_1 , W cannot be observed for any patient. Rather, a patient is observed to experience only a single treatment history. Let A_k be the observed treatment received at decision k and let $\bar{a}_k = (A_1, \dots, A_k)$ be observed treatment history up to decision k . Let X_k be the covariate information observed between decisions $k - 1$ and k under the observed treatment history \bar{a}_{k-1} ($k = 2, \dots, K$), with history $X_k = (X_1, \dots, X_k)$ for $k = 1, \dots, K$ to decision k . Let Y be the observed outcome under \bar{a}_K . The observed data on a patient are (X_K, \bar{a}_K, Y) , and the data available from a clinical trial or observational study involving n subjects are independent and identically distributed $(X_{K_i}, \bar{a}_{K_i}, Y_i)$ for $i = 1, \dots, n$.

Under the following standard assumptions, g^{opt} may equivalently be expressed in terms of the observed data. The consistency assumption states that

$X_k = X_k^*(\bar{A}_{k-1}) = \sum_{\bar{a}_{k-1} \in \bar{\mathcal{A}}_{k-1}} X_{k-1}^*(\bar{a}_{k-1}) I(\bar{A}_{k-1} = \bar{a}_{k-1})$ for $k = 2, \dots, K$, and $Y = Y^*(\bar{a}_K) = \sum_{\bar{a}_K \in \bar{\mathcal{A}}_K} Y^*(\bar{a}_K) I(\bar{a}_K = \bar{a}_K)$; that is, a patient's observed covariates and outcome are the same as the potential ones s/he would exhibit under the treatment history actually received. The stable unit treatment value assumption (Rubin, 1978), implies that a patient's covariates and outcome are not influenced by treatments received by other patients. A version of the sequential randomization assumption (Robins, 2004) states that W is independent of A_k conditional on (X_k, \bar{a}_{k-1}) . This is satisfied by default for data from a sequentially randomized clinical trial (Murphy, 2005), but is not verifiable from data from an observational study. It is reasonable to believe that decisions made in an observational study are based on a patient's covariate and treatment history; however, all such information associated with treatment assignment and outcome must be recorded in the X_k to validate the assumption.

Under these assumptions, from §1 of the Supplementary Material,

$p_{Y^*(\bar{a}_K) | \bar{X}_K^*(\bar{a}_{K-1})}(y | \bar{x}_K) = p_{Y | \bar{X}_K, \bar{A}_K}(y | \bar{x}_K, \bar{a}_K)$, so that $E\{Y^*(\bar{a}_K) | \bar{X}_K^*(\bar{a}_{K-1}) = \bar{x}_K\} = E(Y | \bar{X}_K = \bar{x}_K, \bar{A}_K = \bar{a}_K)$. Thus, letting $Q_K(x_K, \bar{a}_K) = E(Y | X_K = x_K, \bar{a}_K = \bar{a}_K)$, (2) and (3) become

$$g_K^{\text{opt}}(\bar{x}_K, \bar{a}_{K-1}) = \arg \max_{a_K \in \Phi_K(\bar{x}_K, \bar{a}_{K-1})} Q_K(\bar{x}_K, \bar{a}_{K-1}, a_K), \quad V_K(\bar{x}_K, \bar{a}_{K-1}) = \max_{a_K \in \Phi_K(\bar{x}_K, \bar{a}_{K-1})} Q_K(\bar{x}_K, \bar{a}_{K-1}, a_K).$$

Using $p_{X_k^*(\bar{a}_{k-1}) | \bar{X}_{k-1}^*(\bar{a}_{k-2})}(\bar{x}_k | \bar{x}_{k-1}) = p_{X_k | \bar{X}_{k-1}, \bar{A}_{k-1}}(x_k | \bar{x}_{k-1}, \bar{a}_{k-1})$, for $k = K, \dots, 2$,

$$\begin{aligned} Q_k(\bar{x}_k, \bar{a}_k) &= E\{V_{k+1}(\bar{x}_k, X_{k+1}, \bar{a}_k) | \bar{X}_k = \bar{x}_k, \bar{A}_k = \bar{a}_k\} \quad (k = K - 1, \dots, 1), \\ g_k^{\text{opt}}(\bar{x}_k, \bar{a}_{k-1}) &= \arg \max_{a_k \in \Phi_k(\bar{x}_k, \bar{a}_{k-1})} Q_k(\bar{x}_k, \bar{a}_{k-1}, a_k) \quad (k = K - 1, \dots, 2), \\ V_k(\bar{x}_k, \bar{a}_{k-1}) &= \max_{a_k \in \Phi_k(\bar{x}_k, \bar{a}_{k-1})} Q_k(\bar{x}_k, \bar{a}_{k-1}, a_k) \quad (k = K - 1, \dots, 2), \end{aligned}$$

and $g_1^{\text{opt}}(x_1) = \arg \max_{a_1 \in \Phi_1(x_1)} Q_1(x_1, a_1)$, $V_1(x_1) = \max_{a_1 \in \Phi_1(x_1)} Q_1(x_1, a_1)$. The $Q_k(x_k, a_k)$ and $V_k(x_k, a_{k-1})$ are referred to as Q-functions and value functions and are derived from the distribution of the observed data.

3. Q- and A- learning

Q-learning is based on the developments in §2. Linear or nonlinear models $Q_k(x_k, a_k)$ in a finite-dimensional parameter β_k may be posited and estimators \hat{g}_k obtained via a backward iterative process for $k = K, \dots, 1$ by solving least squares estimating equations; see §2 of the

Supplementary Material. The estimated optimal regime is $\hat{g}_Q^{\text{opt}} = (\hat{g}_{Q,1}^{\text{opt}}, \dots, \hat{g}_{Q,K}^{\text{opt}})$, where

$$\hat{g}_{Q,1}^{\text{opt}}(x_1) = g_{Q,1}^{\text{opt}}(x_1; \hat{\beta}_1) = \arg \max_{a_1 \in \Phi_1(x_1)} Q_1(x_1, a_1; \hat{\beta}_1), \text{ and}$$

$$\hat{g}_{Q,k}^{\text{opt}}(\bar{x}_k, \bar{a}_{k-1}) = g_{Q,k}^{\text{opt}}(\bar{x}_k, \bar{a}_{k-1}; \hat{\beta}_k) = \arg \max_{a_k \in \Phi_k(\bar{x}_k, \bar{a}_{k-1})} Q_k(\bar{x}_k, \bar{a}_{k-1}, a_k; \hat{\beta}_k) \text{ for } k = 2,$$

\dots, K . Unless all models are correctly specified, \hat{g}_Q^{opt} may not be a good estimator for g^{opt} .

The A-learning method we consider is a version of g-estimation (Robins, 2004); see §2 of the Supplementary Material. Write $Q_k(x_k, a_k)$ as $h_k(x_k, a_{k-1}) + a_k C_k(x_k, a_{k-1})$, $h_k(x_k, a_{k-1}) = Q_k(x_k, a_{k-1}, 0)$ and $C_k(x_k, a_{k-1}) = Q_k(x_k, a_{k-1}, 1) - Q_k(x_k, a_{k-1}, 0)$. We refer to $C_k(x_k, a_{k-1})$ as the Q-contrast function; with two treatment options, $A_k C_k(x_k, a_{k-1})$ is the optimal-blip-to-zero function of Robins (2004). Posit models $C_k(x_k, a_{k-1}; \beta_k)$ and $C_1(x_1; \beta_1)$, depending on parameters β_k ; and models $h_k(x_k, a_{k-1}; \beta_k)$ and $h_1(x_1; \beta_1)$, with parameters β_k for $k = K, \dots, 2$. Let $\pi_k(x_k, a_{k-1}) = \text{pr}(A_k = 1 \mid X_k = x_k, A_{k-1} = a_{k-1})$ and $\pi_1(x_1) = \text{pr}(A_1 = 1 \mid X_1 = x_1)$ be the propensities for treatment, which are unknown unless the data are from a sequentially randomized trial, and specify models $\pi_k(x_k, a_{k-1}; \psi_k)$, $k = K, \dots, 2$, and $\pi_1(x_1; \psi_1)$, e.g., logistic regression models. Estimators $\hat{\pi}_k$ may be found iteratively for $k = K, \dots, 1$ by solving for $\hat{\pi}_k$ and $\hat{\pi}_k$ estimating equations given in §2 of the Supplementary Material, substituting the maximum likelihood estimators $\hat{\pi}_k$. As $Q_k(x_k, a_k)$ is maximized by $a_k =$

$I\{C_k(x_k, a_{k-1}) > 0\}$, the estimated optimal regime is $\hat{g}_A^{\text{opt}} = (\hat{g}_{A,1}^{\text{opt}}, \dots, \hat{g}_{A,K}^{\text{opt}})$, where

$$\hat{g}_{A,1}^{\text{opt}}(x_1) = g_{A,1}^{\text{opt}}(x_1; \hat{\psi}_1) = I\{C_1(x_1; \hat{\psi}_1) > 0\} \text{ and}$$

$$\hat{g}_{A,k}^{\text{opt}}(\bar{x}_k, \bar{a}_{k-1}) = g_{A,k}^{\text{opt}}(\bar{x}_k, \bar{a}_{k-1}; \hat{\psi}_k) = I\{C_k(\bar{x}_k, \bar{a}_{k-1}; \hat{\psi}_k) > 0\}, \text{ for } k = 2, \dots, K. \text{ If the contrast}$$

and propensity models are correctly specified, then \hat{g}_A^{opt} will be consistent for g^{opt} even if $h_k(x_k, a_{k-1}; \beta_k)$ for $k = K, \dots, 2$, and $h_1(x_1; \beta_1)$ are misspecified, and \hat{g}_A^{opt} will consistently estimate g^{opt} . Thus, the quality of \hat{g}_A^{opt} depends on how close the $C_k(x_k, a_{k-1}; \beta_k)$ are to the true contrast functions.

As discussed in §2 of the Supplementary Material, the efficient version of A-learning is so complex as to be infeasible to implement. The implementation of A-learning we use in the empirical studies of §5 is likely as close to efficient as could be hoped in practice.

See the unpublished report of Schulte et al. (2013) for a detailed account of both methods.

4. Proposed robust method

Q- and A-learning are predicated on the postulated models for the Q-functions and Q-contrast functions, respectively, so the resulting estimated regime may be far from g^{opt} if these models are misspecified. We propose an alternative approach that may be robust to such misspecification, based on directly estimating the optimal regime in a specified class of regimes.

Models $Q_k(x_k, k; k)$ or $C_k(x_k, k-1; k)$, whether correct or not, define classes of regimes \mathcal{G} , $\beta = (\beta_1^T, \dots, \beta_K^T)^T$, or \mathcal{G} , indexed analogously by β , whose elements may often be simplified. For example, with $K = 2$, if $C_2(x_2, a_1; 2) = \beta_2 + \beta_2 x_2$ and $C_1(x_1; 1) = \beta_1 + \beta_1 x_1$, the corresponding regimes $g = (g_1, g_2)$ take $g_1(x_1) = \mathbb{I}(\beta_1 + \beta_1 x_1 > 0)$ and $g_2(x_2, a_1) = \mathbb{I}(\beta_2 + \beta_2 x_2 > 0)$. If prior knowledge suggests that treatment 1 would benefit patients with smaller values of X_1 or X_2 , then all reasonable regimes should have $\beta_1 < 0$ and $\beta_2 < 0$, and elements of \mathcal{G} may be expressed in terms of $\beta_1 = -\beta_1 / \beta_1$ and $\beta_2 = -\beta_2 / \beta_2$ as $g = (g_1, g_2)$, $g_1(x_1) = \mathbb{I}(\beta_1 > x_1)$ and $g_2(x_2, a_1) = \mathbb{I}(\beta_2 > x_2)$, $\beta = (\beta_1, \beta_2)^T$.

This suggests considering a class \mathcal{G} , with elements $g = (g_1, \dots, g_K)$, indexed by $\eta = (\eta_1^T, \dots, \eta_K^T)^T$, of form $\{g_1(x_1), \dots, g_K(x_K, K-1)\}$. If \mathcal{G} is derived from models $Q_k(x_k, k; k)$ or $C_k(x_k, k-1; k)$, then $\eta = (\eta_1, \dots, \eta_K)$ or $\eta = (\eta_1, \dots, \eta_K)$ is a many-to-one function of β or β , and $g^{\text{opt}} \in \mathcal{G}$ if these models are correct. Here, estimating $g^{\text{opt}} = \arg \max_{g \in \mathcal{G}} E\{Y^*(g)\}$ defining the regime g_η^{opt} , say, will yield an estimator for g^{opt} . If these models are misspecified, $\eta = (\eta_1, \dots, \eta_K)$ may not converge in probability to η^{opt} , and resulting regimes may be far from optimal. If instead the form of elements of \mathcal{G} is chosen directly based on interpretability or cost, independently of such models, \mathcal{G} may or may not contain g^{opt} , but g_η^{opt} is still of interest as the optimal regime among those deemed realistic in practice.

We propose an approach to estimation of g_η^{opt} in a given class \mathcal{G} by developing an estimator for $E\{Y^*(g)\}$ that is robust to model misspecification and maximizing it in \mathcal{G} . We cast the problem as one of monotone coarsening. Following Tsiatis (2006, §7.1), for fixed g , let $\bar{g}_k = (g_1, \dots, g_k)$, for $k = 1, \dots, K$, and let $\bar{g}_K = g$. Identify full data as the potential outcomes $W_{g_\eta} = \{X_1, X_2^*(g_{\eta_1}), \dots, X_K^*(\bar{g}_{\eta_{K-1}}), Y^*(g_\eta)\}$, and let $\bar{X}_k^*(\bar{g}_{\eta_{k-1}}) = \{X_1, X_2^*(g_{\eta_1}), \dots, X_k^*(\bar{g}_{\eta_{k-1}})\}$. Let \mathcal{C} be a discrete coarsening variable taking values $1, \dots, K$, corresponding to $K + 1$ levels of coarsening, reflecting the extent to which the observed treatments received are consistent with those dictated by g . In the general coarsened data set up, when $\mathcal{C} = k$, we observe $G_k(W_g)$, a many-to-one function of W_g ; when $\mathcal{C} = K$, we observe $G_K(W_g) = W_g$, the full data. Here, under the consistency assumption, this is as follows. If $A_1 = g_1(X_1)$, then $\mathcal{C} = 1$; that is, $\mathbb{I}(\mathcal{C} = 1) = \mathbb{I}\{A_1 = g_1(X_1)\}$, and we observe $G_{\mathcal{C}}(W_g) = G_1(W_g) = X_1$. None of the observed treatments are consistent with following g , so X_2, \dots, X_K, Y are not consistent with g . If $A_1 = g_1(X_1)$ and $A_2 = g_2(X_2, g_1(X_1))$, then $\mathcal{C} = 2$, $\mathbb{I}(\mathcal{C} = 2) = \mathbb{I}\{A_1 = g_1(X_1)\} \mathbb{I}\{A_2 = g_2(X_2, g_1(X_1))\}$, and $G_{\mathcal{C}}(W_{g_\eta}) = G_2(W_{g_\eta}) = \bar{X}_2^*(g_{\eta_1}) = \bar{X}_2$. Only the treatment at decision 1 and the ensuing X_2 are consistent with g . Likewise, $\mathbb{I}(\mathcal{C} = 3) = \mathbb{I}\{A_1 = g_1(X_1)\} \mathbb{I}\{A_2 = g_2(X_2, g_1(X_1))\} \mathbb{I}\{A_3 = g_3(X_3)\}$, where $g_3(X_3)$ is shorthand for $g_3[X_3, g_1(X_1), g_2(X_2, g_1(X_1))] = g_3\{X_3, \bar{g}_2(X_2)\}$ and $\bar{g}_2(X_2) = [g_1(X_1), g_2(X_2, g_1(X_1))]$ and similarly for general k ; and $G_{\mathcal{C}}(W_{g_\eta}) = G_3(W_{g_\eta}) = \bar{X}_3^*(\bar{g}_{\eta_2}) = \bar{X}_3$. Continuing in this fashion, $\mathbb{I}(\mathcal{C} = K) = \mathbb{I}\{A_{K-1} = \bar{g}_{K-1}\{X_{K-1}, \bar{g}_{K-2}(X_{K-2})\}\} \mathbb{I}\{A_K = g_K\{X_K, \bar{g}_{K-1}(X_{K-1})\}\}$, and $G_{\mathcal{C}}(W_{g_\eta}) = G_K(W_{g_\eta}) = \bar{X}_K^*(\bar{g}_{\eta_{K-1}}) = \bar{X}_K$. Finally, if $\mathcal{C} = \bar{g}_K\{X_K, \bar{g}_{K-1}(X_{K-1})\}$, $G_{\mathcal{C}}(W_g) = G_K(W_g) = W_g = (X_1, \dots, X_K, Y)$. Here, the observed data are consistent with having followed all K rules in g . The coarsening is monotone in that $G_k(W_g)$ is a coarsened version of $G_k(W_g)$, $k > k$, and $G_k(W_g)$ is a many-to-one function of $G_{k+1}(W_g)$.

Coarsened data are said to be coarsened at random if, for each k , the probability that the data are coarsened at level k , given the full data, depends only on the coarsened data, so only on data that are observed at level k (Tsiatis, 2006, §7.1). Under the consistency and sequential randomization assumptions, it may be shown using results in §3 of the Supplementary

Material that the coarsening here is at random. Define the coarsening discrete hazard $\text{pr}(\mathcal{C} = k | \mathcal{C} \leq k, W_g)$ to be the probability that the observed treatments cease to be consistent with g at decision k , given they are consistent prior to k and all potential outcomes. Under coarsening at random, this hazard is a function only of the coarsened data, that is, the data observed through decision k , which we write as $\text{pr}(\mathcal{C} = k | \mathcal{C} \leq k, W_g) = \lambda_{\eta,k}(W_g)$. Then, from above, for $k = 1, \dots, K$, $\lambda_{\eta,k}(W_g) = \text{pr}\{A_k \neq g_k(X_k) | X_1, \dots, X_k\}$, which can be expressed in terms of the propensity for treatment at decision 1 as $\lambda_{\eta,k}(W_g) = \text{pr}\{A_1 \neq g_1(X_1) | X_1\} \{1 - \text{pr}\{A_1 = g_1(X_1)\}\}^{g_1(X_1)}$. Similarly, for $k = 2, \dots, K$,

$$\begin{aligned} \lambda_{\eta,k}(W_{g_{\eta}}) &= \lambda_{\eta,k}(\bar{X}_k) = \text{pr}\{A_k \neq g_{\eta,k}(\bar{X}_k, \bar{A}_{k-1}) | \bar{X}_k, \bar{A}_{k-1}\} \\ &= \bar{g}_{\eta,k-1}(\bar{X}_{k-1}) \\ &= \pi_k\{\bar{X}_k, \bar{g}_{\eta,k-1}(\bar{X}_{k-1})\}^{1-g_{\eta,k}(\bar{X}_k, \bar{g}_{\eta,k-1}(\bar{X}_{k-1}))} \\ &\quad \times [1 - \pi_k\{\bar{X}_k, \bar{g}_{\eta,k-1}(\bar{X}_{k-1})\}]^{g_{\eta,k}(\bar{X}_k, \bar{g}_{\eta,k-1}(\bar{X}_{k-1}))}. \end{aligned}$$

We may then express the probabilities of being consistent with g through at least the k th decision, so having $\mathcal{C} > k$, given all potential outcomes, in terms of the discrete hazards. Under coarsening at random, these probabilities depend only on the observed data through decision k . That is, $\text{pr}(\mathcal{C} > k | W_g) = K - \sum_{k=1}^k \lambda_{\eta,k}(W_g) = K - \sum_{k=1}^k \lambda_{\eta,k}(X_k)$, where

$$K_{\eta,k}(\bar{X}_k) = \prod_{k'=1}^k \{1 - \lambda_{\eta,k'}(\bar{X}_{k'})\} \text{ (Tsiatis, 2006, §8.1).}$$

We now use these developments to deduce the form of estimators for $E\{Y^*(g)\}$. From the theory of Robins et al. (1994) for general monotonely coarsened data, under coarsening at random, if the coarsening mechanism is correctly specified, which corresponds here to correct specification of the $\lambda_{\eta,k}(X_k)$, and hence of the propensity models, all regular, asymptotically linear, consistent estimators (Tsiatis, 2006, Chapter 3) for $E\{Y^*(g)\}$ for fixed g have the form

$$\sum_{i=1}^n \frac{I(\mathcal{C}_{\eta,i} = \infty)}{K_{\eta,K}(\bar{X}_{K,i})} Y_i + \sum_{k=1}^K \frac{I(\mathcal{C}_{\eta,i} = k) - \lambda_{\eta,k}(\bar{X}_{k,i}) I(\mathcal{C}_{\eta,i} \geq k)}{K_{\eta,k}(\bar{X}_{k,i})} L_k(\bar{X}_{k,i}), \quad (4)$$

where $L_k(X_k)$ are arbitrary functions of X_k . The optimal choice leading to (4) with smallest asymptotic variance is $L_{\eta,k}^{\text{opt}}(\bar{x}_k) = E\{Y^*(g_{\eta}) | \bar{X}_k^*(\bar{g}_{\eta,k-1}) = \bar{x}_k\}$. The right hand term in (4) augments the first, itself a consistent estimator for $E\{Y^*(g)\}$ when the $\lambda_{\eta,k}(X_k)$ are correctly specified, to gain efficiency. As in Tsiatis (2006, §10.3), (4) is doubly robust in that it is a consistent estimator for $E\{Y^*(g)\}$ if either the $\lambda_{\eta,k}(X_k)$ are correctly specified or if the $L_k(X_k)$ are equal to $L_{\eta,k}^{\text{opt}}(\bar{X}_{k,i}) (k=1, \dots, K)$; see §4 of the Supplementary Material.

To implement (4), one must specify $\lambda_{\eta,k}(X_k)$ and $L_k(X_k)$. The first follow from specifying $\text{pr}(A_1 = 1 | X_1 = x_1)$, $\text{pr}(A_k = 1 | X_k = x_k, X_{k-1} = x_{k-1})$ for $k = 2, \dots, K$. If these are unknown, as in A-learning, posit models $\text{pr}(A_1 = 1 | X_1; \gamma_1)$, $\text{pr}(A_k = 1 | X_k; \gamma_k)$ for $k = 2, \dots, K$, and estimate $\lambda_{\eta,k}$ by $\hat{\lambda}_{\eta,k}$ ($k = 1, \dots, K$). With $\gamma = (\gamma_1^T, \dots, \gamma_K^T)^T$ and $\hat{\gamma}^T = (\hat{\gamma}_1^T, \dots, \hat{\gamma}_K^T)^T$, this implies that $\lambda_{\eta,k}(X_k; \gamma) = \text{pr}(A_1 = 1 | X_1; \gamma_1)^{1-g_1(X_1)} \{1 - \text{pr}(A_1 = 1 | X_1; \gamma_1)\}^{g_1(X_1)}$,

$$\lambda_{\eta,k}(\bar{X}_k; \gamma_k) = \pi_k\{\bar{X}_k, \bar{g}_{\eta,k-1}(\bar{X}_{k-1}); \gamma_k\}^{1-g_{\eta,k}(\bar{X}_k, \bar{g}_{\eta,k-1}(\bar{X}_{k-1}))} \times [1 - \pi_k\{\bar{X}_k, \bar{g}_{\eta,k-1}(\bar{X}_{k-1}); \gamma_k\}]^{g_{\eta,k}(\bar{X}_k, \bar{g}_{\eta,k-1}(\bar{X}_{k-1}))}$$

and $K_{\eta,k}(\bar{X}_k; \gamma) = \prod_{k'=1}^k \{1 - \lambda_{\eta,k'}(\bar{X}_{k'}; \gamma_{k'})\}$, and suggests substituting $L_k(X_k)$ and $K_{\eta,k}(X_k)$ in (4).

Several options exist for specification of the $L_k(X_k)$. The simplest is to take $L_k(X_k) = 0$, yielding the inverse probability weighted estimator

$$IPWE(\eta) = \sum_{i=1}^n \frac{I(\mathcal{C}_{\eta,i} = \infty)}{K_{\eta,K}(\bar{X}_{K,i}; \hat{\gamma})} Y_i, \quad (5)$$

which is consistent for $E\{Y^*(g)\}$ if $L_1(X_1; \gamma)$ and $L_k(X_k, \dots, X_{k-1}; \gamma)$ ($k = 2, \dots, K$), and hence $K_{\eta,k}(X_k; \gamma)$, are correctly specified, but otherwise may be inconsistent. The corresponding estimator for g_{η}^{opt} is found by estimating g_{η}^{opt} by \hat{g}_{IPWE}^{opt} , say, maximizing (5) in η . As (5) is based on data only from subjects whose entire treatment history is consistent with g , it is relatively less efficient than estimators that use all the data, discussed next.

To take greatest advantage of the potential for improved efficiency through the augmentation term in (4), an obvious approach is to posit and fit parametric models

approximating the conditional expectations $L_{\eta,k}^{opt}(\bar{x}_k) = E\{Y^*(g_{\eta}) | \bar{X}_k^*(\bar{g}_{\eta_{k-1}}) = \bar{x}_k\}$, and substitute these into (4) along with $L_k(X_k; \gamma)$ and $K_{\eta,k}(X_k; \gamma)$. To this end, let $\mu_K(X_K, \dots, X_1) = E\{Y | X_K = x_K, \dots, X_1 = x_1\}$ and $f_K(x_K, \dots, x_1) = \mu_K\{x_K, \dots, x_1, g_K(x_K, \dots, x_1)\}$. Then define iteratively, for $k = K - 1, \dots, 2$, the quantities $\mu_k(x_k, \dots, x_1) = E\{f_{k+1}(x_k, X_{k+1}, \dots, X_K) | X_k = x_k, \dots, X_1 = x_1\}$ and $f_k(x_k, \dots, x_1) = \mu_k\{x_k, \dots, x_1, g_k(x_k, \dots, x_1)\}$; for $k = 1$, $\mu_1(x_1, a_1) = E\{f_2(x_1, X_2, a_1) | X_1 = x_1, A_1 = a_1\}$, $f_1(x_1) = \mu_1\{x_1, g_1(x_1)\}$. In §5 of the Supplementary Material, we demonstrate that $L_{\eta,k}^{opt}(\bar{X}_k) = \mu_{\eta,k}\{\bar{X}_k, \bar{g}_{\eta_k}(\bar{X}_k)\}$.

This suggests specifying η -dependent models $\mu_k(x_k, \dots, x_1; \xi_k)$ depending on parameters ξ_k , $k = 1, \dots, K$. For fixed η , estimators $\hat{\mu}_k$ for μ_k may be found iteratively by solving in ξ_k

$$\sum_{i=1}^n \frac{\partial \mu_{\eta,k}(\bar{X}_{ki}, \bar{A}_{ki}; \xi_k)}{\partial \xi_k} \{\hat{f}_{(k+1)i} - \mu_{\eta,k}(\bar{X}_{ki}, \bar{A}_{ki}; \xi_k)\} = 0 \quad (k=1, \dots, K),$$

where $\hat{f}_{(k+1)i} = \mu_k(X_{ki}, \dots, X_{ki}; \hat{g}_{\eta_k})$ is the vector of partial derivatives of $\mu_k(X_{ki}, \dots, X_{ki}; \hat{g}_{\eta_k})$ with respect to elements of ξ_k , $f_{(k+1)i} = Y_i$ and $f_{ki} = \mu_k[X_{ki}, \dots, X_{ki}; g_k\{X_{ki}, \dots, X_{ki}\}]$; $k = K, \dots, 2$. The fitted $\hat{\mu}_k\{X_k, \bar{g}_k(X_k); \xi_k\}$ may then be used to approximate $L_{\eta,k}^{opt}(\bar{x}_k)$ in (4). While these models almost certainly are not correct, as specification of a compatible sequence of models for $k = 1, \dots, K$ is a significant challenge, they may be reasonable approximations to the true conditional expectations. Thus, define

$$DR(\eta) = \sum_{i=1}^n \frac{I(\mathcal{C}_{\eta,i} = \infty)}{K_{\eta,K}(\bar{X}_{K,i}; \hat{\gamma})} Y_i + \sum_{k=1}^K \frac{I(\mathcal{C}_{\eta,i} = k) - \lambda_{\eta,k}(\bar{X}_{k,i}; \hat{\gamma}_k) I(\mathcal{C}_{\eta,i} \geq k)}{K_{\eta,k}(\bar{X}_{k,i}; \hat{\gamma}_k)} \mu_{\eta,k}\{\bar{X}_{k,i}, \bar{g}_{\eta_k}(\bar{X}_{k,i}); \hat{\xi}_k\}, \quad (6)$$

which, by virtue of the double robustness property, will be consistent for $E\{Y^*(g)\}$ if either $L_1(X_1; \gamma)$ and $L_k(X_k, \dots, X_{k-1}; \gamma)$ ($k = K, \dots, 2$), are correctly specified, or the the $\mu_k(x_k, \dots, x_1; \xi_k)$ are. If all of these models were correct, then (6) would achieve optimal efficiency. As for (5), estimation of g_{η}^{opt} follows by maximizing (6) in η to obtain $\hat{\eta}_{DR}^{opt}$.

A computational challenge is that the models $\mu_k(x_k, \dots, x_{k-1}; \xi_k)$ must be refitted for each value of η encountered in the optimization algorithm used to carry out the maximization. A practical alternative when regimes in \mathcal{G} are derived from models is to substitute for $L_k(X_{k,i})$

in (4) fitted Q-functions $Q_k\{X_k, \bar{g}_k(X_k); k\}$ for $k = K, \dots, 1$ obtained from Q-learning; holding k fixed, these depend on k only through $\bar{g}_k(X_k)$. While these are not strictly models for $E\{Y^*(g_\eta)|\bar{X}_k(\bar{g}_{\eta_k})\}$, the hope is that they will be close enough to achieve near optimal efficiency gains over (5). Thus, estimate g_η^{opt} by maximizing in η to obtain $\hat{\eta}_{\text{AIPWE}}^{\text{opt}}$

$$\text{AIPWE}(\eta) = \sum_{i=1}^n \frac{I(\mathcal{C}_{\eta,i} = \infty)}{K_{\eta,K}(\bar{X}_{K,i}; \hat{\gamma})} Y_i + \sum_{k=1}^K \frac{I(\mathcal{C}_{\eta,i} = k) - \lambda_{\eta,k}(\bar{X}_{k,i}; \hat{\gamma}_k) I(\mathcal{C}_{\eta,i} \geq k)}{K_{\eta,k}(\bar{X}_{k,i}; \hat{\gamma})} Q_k\{\bar{X}_{k,i}, \bar{g}_{\eta_k}(\bar{X}_{k,i}); \hat{\beta}_k\}. \quad (7)$$

See §6 of the Supplementary Material for a similar proposal when \mathcal{G} is determined directly.

Standard errors for these estimators for $E\{Y^*(g_\eta^{\text{opt}})\}$ may be obtained via the sandwich technique (Stefanski & Boos, 2002) based on the argument in Zhang et al. (2012, Equation (4)).

5. Simulation studies

We have carried out several simulation studies to evaluate the performance of the proposed methods, each involving 1000 Monte Carlo data sets.

The first simulation adopts the scenario in Moodie et al. (2007) of a study in which HIV-infected patients are randomized to initiate antiretroviral therapy or not, coded as 1 or 0, at baseline and again at six months to determine the optimal regime for therapy initiation. We generated baseline CD4 count $X_1 \sim N(450, 150)$, where $N(\mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2 ; baseline treatment A_1 as Bernoulli with success probability $\text{pr}(A_1 = 1 | X_1) = \text{expit}(2 - 0.006X_1)$, where $\text{expit}(u) = e^u / (1 + e^u)$; six-month CD4 count X_2 , conditional on (X_1, A_1) , as $N(1.25X_1, 60)$; and treatment at six months A_2 as Bernoulli with $\text{pr}(A_2 = 1 | X_2, A_1) = A_1 + (1 - A_1)\text{expit}(0.8 - 0.004X_2)$. Here, patients with $A_1 = 1$ continue on therapy with certainty. The outcome Y , one-year CD4 count, conditional on (X_2, A_2) , was normal with mean $400 + 1.6X_1 - |250 - X_1\{A_1 - I(250 - X_1 > 0)\}^2 - (1 - A_1)|720 - 2X_2|A_2 - I(720 - 2X_2 > 0)|^2$ and variance 60^2 . The true Q-contrast functions are thus $C_2(x_1, x_2, a_1) = (1 - a_1)(720 - 2x_2)$, $C_1(x_1) = 250 - x_1$, the optimal treatment regime $g^{\text{opt}} = (g_1^{\text{opt}}, g_2^{\text{opt}})$ has $g_1^{\text{opt}}(x_1) = I(250 - x_1 > 0)$, $g_2^{\text{opt}}(\bar{x}_2, a_1) = I\{a_1 + (1 - a_1)(720 - 2x_2) > 0\} = I\{a_1 + (1 - a_1)(360 - x_2) > 0\}$ and $E\{Y^*(g^{\text{opt}})\} = 1120$.

For A-learning, we took

$$h_2(\bar{x}_2, a_1; \alpha_2) = \alpha_{20} + \alpha_{21}x_1 + \alpha_{22}a_1 + \alpha_{23}a_1x_1 + \alpha_{24}(1 - a_1)x_2, \quad C_2(\bar{x}_2, a_1; \psi_2) = (1 - a_1)(\psi_{20} + \psi_{21}x_2),$$

$h_1(x_1, a_1) = \alpha_{10} + \alpha_{11}x_1$, and $C_1(x_1; \beta_1) = \beta_{10} + \beta_{11}x_1$; and, analogously, for Q-learning,

$$Q_2(\bar{x}_2, \bar{a}_2; \beta_2) = \beta_{20} + \beta_{21}x_1 + a_1(\beta_{22} + \beta_{23}x_1) + \beta_{24}(1 - a_1)x_2 + a_2(1 - a_1)(\beta_{25} + \beta_{26}x_2), \\ Q_1(x_1, a_1; \beta_1) = \beta_{10} + \beta_{11}x_1 + a_1(\beta_{12} + \beta_{13}x_1),$$

so the Q-contrast functions are correct, but the Q-functions are misspecified. Here, $C_2(x_2, 1; \psi_2) = 0$, respecting that $\psi_2(x_2, 1) = \{1\}$. We used correct propensity models $\psi_2(x_2, a_1 = 0; \psi_2) = \text{expit}(\alpha_{20} + \alpha_{21}x_2)$, $\psi_1(x_1; \psi_1) = \text{expit}(\alpha_{10} + \alpha_{11}x_1)$ and incorrect models $\psi_2(x_2, a_1 = 0; \psi_2) = \psi_2$, $\psi_1(x_1; \psi_1) = 1$.

For maximizing $ipwe(\cdot)$ in (5), $dr(\cdot)$ in (6), and $aipwe(\cdot)$ in (7) to obtain $\hat{\eta}_{IPWE}^{opt}$, $\hat{\eta}_{DR}^{opt}$, and $\hat{\eta}_{AIPWE}^{opt}$ we considered the class of regimes \mathcal{G} with elements $g = (g_1, g_2)$,

$$g_{\eta_2}(\bar{x}_2, a_1) = I\{a_1 + (1 - a_1)(\eta_{20} + \eta_{21}x_2) > 0\}, g_{\eta_1}(x_1) = I(\eta_{10} + \eta_{11}x_1 > 0),$$

so that $\beta_2 = (\beta_{20}, \beta_{21})^T$, $\beta_1 = (\beta_{10}, \beta_{11})^T$, $\eta = (\eta_1^T, \eta_2^T)^T$ and $\eta^{opt} = (250, -1, 360, -1)^T$. Clearly, $g^{opt} \in \mathcal{G}$. We used the same propensity models, and, for (7), Q-function models as above; for (6), we posited $\mu_2(x_2, a_1; \beta_2) = \beta_{20} + \beta_{21}x_1 + a_1(\beta_{22} + \beta_{23}x_1) + \beta_{24}(1 - a_1)x_2 + a_2(1 - a_1)(\beta_{25} + \beta_{26}x_2)$ and $\mu_1(x_1, a_1; \beta_1) = \beta_{10} + \beta_{11}x_1 + a_1(\beta_{12} + \beta_{13}x_1)$ for each β . To achieve a unique representation, we fixed $(\beta_{21}, \beta_{11}) = (-1, -1)$ and determined β_{20}, β_{10} via a grid search; because $ipwe(\cdot)$, $dr(\cdot)$ and $aipwe(\cdot)$ are step functions of β with jumps at (x_{1j}, x_{2j}) ($i, j = 1, \dots, n$), we maximized in β over all (x_{1j}, x_{2j}) .

The second scenario is the same as the first except that the models for the Q-contrast functions are misspecified. Specifically, the generative distribution of Y given (X_2, a_2) is now normal with mean $400 + 1.6X_1 - |250 - 0.6X_1| \{A_1 - I(250 - X_1 > 0)\}^2 - (1 - A_1)|720 - 1.4X_2| \{A_2 - I(720 - 2X_2 > 0)\}^2$ and variance 60^2 , so that, from the discussion below (2) of Moodie et al. (2007), the implied true contrast functions are no longer of the form above, but all posited models were taken to be the same as in the first simulation.

Tables 1 and 2 show the results. For Q- and A-learning, we report $\hat{\eta}(\cdot)$ and $\hat{\eta}^{opt}(\cdot)$. The column $\hat{E}(\cdot^{opt})$ shows for each estimator the Monte Carlo average and standard deviation of the estimated values of $E\{Y^*(g_{\eta}^{opt})\}$ reflecting performance for estimating the true achievable mean outcome under the true optimal regime, while $E(\cdot^{opt})$ reflects performance of the estimated optimal regime itself. For each Monte Carlo data set, this is the true mean outcome that would be achieved if the estimated optimal regime were followed by the population was determined by simulation, and the values reported are the Monte Carlo average and standard deviation of these simulated quantities. When compared to the true $E\{Y^*(g_{\eta}^{opt})\} = 1120$, these measure the extent to which the estimated optimal regimes approach the performance of the true optimal regime.

For the first simulation, from Table 1, because the Q-functions are misspecified, the Q-learning estimators for β_{10} and β_{20} are biased, while those from A-learning based on postulated Q-contrast functions that include the truth are consistent when the propensity model is correct. When the propensity model is incorrect, Q-learning is unaffected; however, A-learning yields biased estimators for β_{10} and β_{20} identical to those from Q-learning, as linear models are used for $C_2(x_2, a_1; \beta_2)$, $C_1(x_1; \beta_1)$, $h_2(x_2, a_1; \beta_2)$ and $h_1(x_1; \beta_1)$ (Chakraborty et al., 2010). Although Q-learning results in poor estimation of β_{10} and β_{20} , efficiency loss for estimating the optimal regime is negligible, as the proportion of benefit the estimated regime achieves if used in the entire population relative to the true optimal regime is virtually one. A possible explanation is that patients near the true decision boundary have $C_2(X_2, a_1)$, $C_1(X_1)$ close to zero, and few patients would receive treatment 1 according to the true decision rule for the first time point. This also follows from the fact that for regime $g = (0, g_2^{opt})$, the corresponding expectation is 1114. When the propensity model is correct, the estimators based on $dr(\cdot)$ and $aipwe(\cdot)$ yield estimated regimes comparable to those found by A-learning in terms of true mean outcome achieved, despite yielding relatively inefficient estimators for β_{10} and β_{20} A-learning, perhaps for the same reason as above. When the propensity model is incorrect, the $dr(\cdot)$ and $aipwe(\cdot)$ estimators yield estimated regimes that are still close to the optimal. The $ipwe(\cdot)$ estimator show relatively poorer performance, especially when the propensity score model is incorrect,

which is not unexpected; this estimator only uses information from patients whose treatment histories are consistent with following g and hence is inefficient.

In the second simulation, the values of $|C_2(X_2, A_1)|$ and $|C_1(X_1)|$ for patients near the true decision boundary are larger than in the first simulation, and the posited Q-contrast functions are no longer correct. From Table 2, the A- and Q- learning estimators perform similarly, both yielding estimated regimes far from optimal. Those based on $\text{dr}(\cdot)$ and $\text{aipwe}(\cdot)$ are almost identical to g^{opt} on average and perform almost identically to the true optimal regime, regardless of whether or not the propensity model is correct. Again, the estimator based on $\text{ipwe}(\cdot)$ in (5) performs poorly. Evidently, augmentation even using incorrect models leads to considerable gains over $\text{ipwe}(\cdot)$ regardless of whether or not the propensity model is correct.

The third scenario involved $K = 3$ decision points. To achieve average numbers of patients consistent with the regime comparable to those in the $K = 2$ cases, we took $n = 1000$. We generated X_1, A_1, X_2 as in the previous two scenarios; A_2 as Bernoulli with $\text{pr}(A_2 = 1 | X_2, A_1) = \text{expit}(0.8 - 0.004X_2)$; twelve-month CD4 count X_3 , conditional on (X_2, A_2) , as $N(0.8X_2, 60)$; treatment at twelve months A_3 as Bernoulli with $\text{pr}(A_3 = 1 | X_3, A_2) = \text{expit}(1 - 0.004X_3)$; and the outcome Y , 18-month CD4 count, conditional on (X_3, A_3) , as normal with mean $400 + 1.6X_1 - |500 - 1.4X_1| \{A_1 - I(500 - 2X_1 > 0)\}^2 - |720 - 1.4X_2| \{A_2 - I(720 - 2X_2 > 0)\}^2 - |600 - 1.4X_3| \{A_3 - I(600 - 2X_3 > 0)\}^2$ and variance 60^2 . The optimal treatment regime $g^{\text{opt}} = (g_1^{\text{opt}}, g_2^{\text{opt}}, g_3^{\text{opt}})$ has $g_1^{\text{opt}}(x_1) = I(250 - x_1 > 0)$, $g_2^{\text{opt}}(\bar{x}_2, a_1) = I(360 - x_2 > 0)$, $g_3^{\text{opt}}(\bar{x}_3, \bar{a}_2) = I(300 - x_3 > 0)$ and $E\{Y^*(g^{\text{opt}})\} = 1120$.

For A-learning, we took

$$\begin{aligned} h_3(\bar{x}_3, \bar{a}_2; \alpha_3) &= \alpha_{30} + \alpha_{31}x_1 + a_1(\alpha_{32} + \alpha_{33}x_1) + \alpha_{34}x_2 + a_2(\alpha_{35} + \alpha_{36}x_2) + \alpha_{37}x_3, \\ C_3(\bar{x}_3, \bar{a}_2; \psi_2) &= \psi_{30} + \psi_{31}x_3, \quad h_2(\bar{x}_2, a_1; \alpha_2) = \alpha_{20} + \alpha_{21}x_1 + a_1(\alpha_{22} + \alpha_{23}x_1) + \alpha_{24}x_2, \\ C_2(\bar{x}_2, a_1; \psi_2) &= \psi_{20} + \psi_{21}x_2, \quad h_1(x_1; \alpha_1) = \alpha_{10} + \alpha_{11}x_1, \quad C_1(x_1; \psi_1) = \psi_{10} + \psi_{11}x_1, \end{aligned}$$

and for Q-learning

$$\begin{aligned} Q_3(\bar{x}_3, \bar{a}_2; \beta_3) &= \beta_{30} + \beta_{31}x_1 + a_1(\beta_{32} + \beta_{33}x_1) + \beta_{34}x_2 + a_2(\beta_{35} + \beta_{36}x_2) + \alpha_{37}x_3 + a_3(\beta_{38} + \beta_{39}x_3), \\ Q_2(\bar{x}_2, \bar{a}_2; \beta_2) &= \beta_{20} + \beta_{21}x_1 + a_1(\beta_{22} + \beta_{23}x_1) + \beta_{24}x_2 + a_2(\beta_{25} + \beta_{26}x_2), \\ Q_1(x_1, a_1; \beta_1) &= \beta_{10} + \beta_{11}x_1 + a_1(\beta_{12} + \beta_{13}x_1); \end{aligned}$$

thus, both Q- and Q-contrast functions are misspecified. We used correct propensity models $\psi_3(x_3, a_2; \beta_3) = \text{expit}(\beta_{30} + \beta_{31}x_3)$, $\psi_2(x_2, a_1; \beta_2) = \text{expit}(\beta_{20} + \beta_{21}x_2)$, $\psi_1(x_1; \beta_1) = \text{expit}(\beta_{10} + \beta_{11}x_1)$ and incorrect models $\psi_3(x_3, a_2; \beta_3) = \beta_3$, $\psi_2(x_2, a_1; \beta_2) = \beta_2$, $\psi_1(x_1; \beta_1) = \beta_1$.

For the three proposed estimators, we took the class of regimes \mathcal{G} to have elements $g = (g_1, g_2, g_3)$, $g_3(x_3, a_2) = I(\beta_{30} + \beta_{31}x_3 > 0)$, $g_2(x_2, a_1) = I(\beta_{20} + \beta_{21}x_2 > 0)$, $g_1(x_1) = I(\beta_{10} + \beta_{11}x_1 > 0)$, so $\beta_3 = (\beta_{30}, \beta_{31})^T$, $\beta_2 = (\beta_{20}, \beta_{21})^T$, $\beta_1 = (\beta_{10}, \beta_{11})^T$, $\eta = (\eta_1^T, \eta_2^T, \eta_3^T)^T$ and $g^{\text{opt}} = (250, -1, 360, -1, 300, -1)^T$, so $g^{\text{opt}} \in \mathcal{G}$. We used the same propensity models, and, for (7), Q-function models as above, and fixed $(\beta_{31}, \beta_{21}, \beta_{11}) = (-1, -1, -1)$. To carry out the maximizations, we used a genetic algorithm discussed by Goldberg (1989), implemented in the `rgenoud` package in R (Mebane & Sekhon, 2011); see §7 of the Supplementary Material for details.

Table 3 show the results. Q-learning performs poorly, as expected. When the propensity model is correctly specified, results for A-learning and the proposed methods are similar to those in the second scenario, with the estimated regimes based on $dr(\cdot)$ and $aipwe(\cdot)$ achieving near-optimal performance and associated reliable inference on the true achievable mean outcome $E\{Y^*(g)\}$. When the propensity models are misspecified, the situation is similar for these estimators in terms of performance; however, inference on $E\{Y^*(g)\}$ is markedly degraded. In both cases, performance of the estimator based on $ipwe(\cdot)$ is quite poor. Intuitively, as the number of decisions K increases, it is not unexpected that all methods can suffer from diminished performance. Research is needed on the design of sequentially randomized trials to ensure adequate sample size for reliable inference on multi-decision regimes.

In §8 of the Supplementary Material, we present results of a more complex scenario; the qualitative conclusions are similar.

All simulations here, and others we have conducted, suggest that Q- and A-learning can yield biased estimators for parameters defining the optimal regime if the Q-functions or Q-contrast functions are misspecified. Under these conditions, the resulting estimated optimal regimes can perform poorly in terms of achieving the expected potential outcome of the true optimal regime. In contrast, the proposed approach using (6) or (7) exhibits robustness to misspecification of either one of the outcome regression or propensity score models. Under these circumstances, the estimators of regime parameters are relatively unbiased, and the expected potential outcome under the estimated optimal regime approaches that of the true optimal regime. Moreover, the proposed methods lead to reliable estimation of the expected potential outcome under the true regime, with coverage probabilities close to the nominal level. Even when both outcome regression and propensity models are misspecified, the proposed methods can yield estimated optimal regimes that do not show substantial degradation of performance in terms of achieved expected potential outcome relative to the true optimal regime. In this case, inference on the expected outcome under the true optimal regime can be compromised, although, interestingly, the methods performed well in this regard under these conditions in the second simulation scenario. Collectively, our results suggest that the proposed methods are attractive alternatives to Q- and A-learning owing to their robustness to such model misspecification. As the estimator based on $aipwe(\cdot)$ is much less computationally intensive than $dr(\cdot)$ and performs similarly, we recommend it for practical use.

In §9 of the Supplementary Material, we report on application of the methods to a study to compare treatment options in patients with nonpsychotic major depressive disorder.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This research was supported by grants from the US National Institutes of Health.

References

- Almirall D, Ten Have T, Murphy SA. Structural nested mean models for assessing time-varying effect moderation. *Biometrics*. 2010; 66:131–139. [PubMed: 19397586]
- Bather, J. *Decision Theory: an Introduction to Dynamic Programming and Sequential Decisions*. Chichester: Wiley; 2000.

- Chakraborty B, Murphy SA, Strecher V. Inference for non-regular parameters in optimal dynamic treatment regimes. *Statist. Meth. Med. Res.* 2010; 19:317–343.
- Goldberg, DE. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, MA: Addison-Wesley; 1989.
- Henderson R, Ansell P, Alshibani D. Regret-regression for optimal dynamic treatment regimes. *Biometrics*. 2010; 66:1192–1201. [PubMed: 20002404]
- Mebane WR, Sekhon JS. Genetic optimization using derivatives: the rgenoud package for R. *J. Statist. Soft.* 2011; 42:1–26.
- Moodie EEM, Richardson TS, Stephens DA. Demystifying optimal dynamic treatment regimes. *Biometrics*. 2007; 63:447–455. [PubMed: 17688497]
- Murphy SA. Optimal dynamic treatment regimes (with discussion). *J. Royal Statist. Soc., Ser. B.* 2003; 58:331–366.
- Murphy SA. An experimental design for the development of adaptive treatment strategies. *Statist. Med.* 2005; 24:1455–1481.
- Murphy SA, Oslin DW, Rush AJ, Zhu J. Methodological challenges in constructing effective treatment sequences for chronic psychiatric disorders. *Neuropsychopharmacology*. 2007; 32:257–262. [PubMed: 17091129]
- Orellana L, Rotnitzky A, Robins J. Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, part I: Main content. *Int. J. Biostatist.* 2010; 6(Issue 2) Article 8.
- Robins JM. A new approach to causal inference in mortality studies with sustained exposure periods: Applications to control of the healthy worker survivor effect. *Math. Model.* 1986; 7:1393–1512.
- Robins, JM. Optimal structured nested models for optimal sequential decisions. In: Lin, DY.; Heagerty, PJ., editors. *Proceedings of the Second Seattle Symposium on Biostatistics*. New York: Springer; 2004. p. 189-326.
- Robins J, Orellana L, Rotnitzky A. Estimation and extrapolation of optimal treatment and testing strategies. *Statist. Med.* 2008; 27:4678–4721.
- Robins JM, Rotnitzky A, Zhao LP. Estimation of regression coefficients when some regressors are not always observed. *J. Am. Statist. Assoc.* 1994; 89:846–866.
- Rosthøj S, Fullwood C, Henderson R, Stewart S. Estimation of optimal dynamic anticoagulation regimes from observational data: A regret-based approach. *Statist. Med.* 2006; 25:4197–4215.
- Rubin DB. Bayesian inference for causal effects: The role of randomization. *Ann. Statist.* 1978; 6:34–58.
- Stefanski LA, Boos DD. The calculus of M-estimation. *Amer. Statist.* 2002; 56:29–38.
- Tsiatis, AA. *Semiparametric Theory and Missing Data*. New York: Springer; 2006.
- Watkins CJCH, Dayan P. Q-learning. *Mach. Learn.* 1992; 8:279–292.
- Zhang B, Tsiatis AA, Laber EB, Davidian M. A robust method for estimating optimal treatment regimes. *Biometrics*. 2012; 68:1010–1018. [PubMed: 22550953]
- Zhao Y, Kosorok MR, Zeng D. Reinforcement learning design for cancer clinical trials. *Statist. Med.* 2009; 28:3294–3315.

Table 1

Results for the first simulation scenario, Q-contrast functions correct, 1000 Monte Carlo data sets, $n = 500$. For the true optimal regime $g^{\text{opt}} = g_{\eta}^{\text{opt}} \in \mathcal{G}_{\eta}$, $\text{opt} = (250, -1, 360, -1)^T$ and $E\{Y^*(g_{\eta}^{\text{opt}})\} = 1120$

Estimator	10	20	$\hat{E}(\text{opt})$	SE	Cov.	$E(\text{opt})$
Q-learning	228 (17)	322 (25)	1117 (12)	-	-	1119 (1)
Propensity score correct						
A-learning	245 (18)	359 (20)	1121 (11)	-	-	1120 (1)
AIPWE (7)	210 (73)	363 (33)	1125 (12)	12	93.1	1118 (2)
DR (6)	211 (73)	363 (34)	1125 (12)	12	93.1	1118 (2)
IPWE (5)	268 (72)	397 (83)	1183 (24)	34	59.2	1105 (18)
Propensity score correct						
A-learning	228 (17)	322 (25)	1117 (12)	-	-	1119 (1)
AIPWE (7)	259 (51)	390 (47)	1123 (12)	12	93.8	1116 (4)
DR (6)	262 (48)	386 (45)	1123 (12)	12	94.3	1116 (4)
IPWE (5)	349 (49)	471 (63)	1554 (56)	64	0.0	1075 (22)

AIPWE, DR, and IPWE, estimators based on maximizing $\text{aipwe}(\cdot, \text{dr}(\cdot))$, and $\text{ipwe}(\cdot)$, respectively; 10, 20, Monte Carlo average estimates (standard deviation); $\hat{E}(\text{opt})$, Monte Carlo average (standard deviation) of estimated $E\{Y^*(g_{\eta}^{\text{opt}})\}$; SE, Monte Carlo average of sandwich standard errors; Cov., coverage of associated 95% Wald-type confidence intervals for $E(\text{opt})$; $E(\text{opt})$, Monte Carlo average (standard deviation) of values $E\{Y^*(\hat{g}_{\eta}^{\text{opt}})\}$ obtained using 10^6 Monte Carlo simulations for each data set.

Table 2

Results for the second simulation scenario, Q-contrast functions incorrect, 1000 Monte Carlo data sets, $n = 500$. For the true optimal regime $g^{\text{opt}} = g_{\eta^{\text{opt}}} \in \mathcal{G}_{\eta^{\text{opt}}} = (250, -1, 360, -1)^T$ and $E\{Y^*(g_{\eta^{\text{opt}}})\} = 1120$. All quantities are as in Table 1

Estimator	10	20	$\hat{E}(\text{opt})$	SE	Cov.	$E(\text{opt})$
Q-learning	381 (33)	386 (45)	1104 (12)	-	-	1088 (6)
Propensity score correct						
A-learning	364 (29)	453 (26)	1115 (12)	-	-	1087 (3)
AIPWE (7)	250 (21)	359 (9)	1120 (12)	12	94.7	1118 (3)
DR (6)	250 (23)	360 (13)	1121 (11)	12	96.3	1118 (3)
IPWE (5)	305 (67)	432 (86)	1182 (27)	38	70.1	1096 (12)
Propensity score incorrect						
A-learning	381 (33)	386 (45)	1104 (12)	-	-	1088 (6)
AIPWE (7)	255 (24)	363 (28)	1116 (12)	12	93.5	1118 (6)
DR (6)	255 (25)	364 (28)	1116 (12)	12	93.3	1117 (7)
IPWE (5)	361 (47)	480 (69)	1571 (59)	67	0.0	1086 (5)

Table 3

Results for the third simulation scenario, $K = 3$, Q -contrast functions incorrect, 1000 Monte Carlo data sets, $n = 1000$. For the true optimal regime $g^{\text{opt}} = g_{\eta}^{\text{opt}} \in \mathcal{G}_{\eta}$, $g^{\text{opt}} = (250, -1, 360, -1, 300, -1, 300, -1)^T$ and $E\{Y^*(g_{\eta}^{\text{opt}})\} = 1120$. All quantities are as in Table 1

Estimator	10	20	30	$\hat{E}(\text{opt})$	SE	Cov.	$E(\text{opt})$
Q-learning	179 (58)	412.9 (28)	341 (33)	1058 (13)	-	-	1086 (9)
Propensity score correct							
A-learning	319 (12)	462 (11)	387 (12)	1108 (12)	-	-	1071 (3)
AIPWE (7)	263 (41)	362 (14)	300 (7)	1121 (10)	10	94.6	1116 (5)
DR (6)	263 (37)	361 (11)	300 (8)	1121 (10)	10	94.2	1117 (5)
IPWE (5)	399 (132)	618 (138)	450 (132)	1297 (63)	103	56.2	1008 (75)
Propensity score incorrect							
A-learning	179 (58)	413 (28)	341 (33)	1041 (12)	-	-	1086 (9)
AIPWE (7)	360 (48)	371 (39)	310 (30)	1200 (26)	27	9.0	1104 (10)
DR (6)	386 (35)	362 (26)	314 (39)	1208 (26)	27	4.5	1102 (9)
IPWE (5)	412 (42)	521 (60)	415 (58)	2459 (148)	167	0.0	1055 (14)