Published in final edited form as:

J Am Stat Assoc. 2020 ; 115(530): 692–706. doi:10.1080/01621459.2018.1537919.

Estimating Dynamic Treatment Regimes in Mobile Health Using V-learning

Daniel J. Lockett*,

Department of Biostatistics, University of North Carolina at Chapel Hill

Eric B. Laber,

Department of Statistics, North Carolina State University

Anna R. Kahkoska,

Department of Nutrition, University of North Carolina at Chapel Hill

David M. Maahs,

Department of Pediatrics, Stanford University

Elizabeth Mayer-Davis,

Department of Nutrition, University of North Carolina at Chapel Hill

Michael R. Kosorok

Department of Biostatistics, University of North Carolina at Chapel Hill

Abstract

The vision for precision medicine is to use individual patient characteristics to inform a personalized treatment plan that leads to the best possible healthcare for each patient. Mobile technologies have an important role to play in this vision as they offer a means to monitor a patient's health status in real-time and subsequently to deliver interventions if, when, and in the dose that they are needed. Dynamic treatment regimes formalize individualized treatment plans as sequences of decision rules, one per stage of clinical intervention, that map current patient information to a recommended treatment. However, most existing methods for estimating optimal dynamic treatment regimes are designed for a small number of fixed decision points occurring on a coarse time-scale. We propose a new reinforcement learning method for estimating an optimal treatment regime that is applicable to data collected using mobile technologies in an outpatient setting. The proposed method accommodates an indefinite time horizon and minute-by-minute decision making that are common in mobile health applications. We show that the proposed estimators are consistent and asymptotically normal under mild conditions. The proposed methods are applied to estimate an optimal dynamic treatment regime for controlling blood glucose levels in patients with type 1 diabetes.

Keywords

Markov decision processes; Precision medicine; Reinforcement learning; Type 1 diabetes

*The authors gratefully acknowledge NIH P01 CA142538, NIH UL1 TR001111, NIH T32 CA201159, NSF DMS-1555141, NSF DMS-1513579, NSF DMS-1407732, NIH R01 AA023187

1 Introduction

The use of mobile devices in clinical care, called mobile health (mHealth), provides an effective and scalable platform to assist patients in managing their illness (Free et al., 2013; Steinhubl et al., 2013). Advantages of mHealth interventions include real-time communication between a patient and their health-care provider as well as systems for delivering training, teaching, and social support (Kumar et al., 2013). Mobile technologies can also be used to collect rich longitudinal data to estimate optimal dynamic treatment regimes and to deliver treatment that is deeply tailored to each individual patient. We propose a new estimator of an optimal treatment regime that is suitable for use with longitudinal data collected in mHealth applications.

A dynamic treatment regime provides a framework to administer individualized treatment over time through a series of decision rules. Dynamic treatment regimes have been well-studied in the statistical and biomedical literature (Murphy, 2003; Robins, 2004; Moodie et al., 2007; Kosorok and Moodie, 2015; Chakraborty and Moodie, 2013) and furthermore, statistical considerations in mHealth have been studied by, for example, Liao et al. (2016) and Klasnja et al. (2015). Although mobile technology has been successfully utilized in clinical areas such as diabetes (Quinn et al., 2011; Maahs et al., 2012), smoking cessation (Ali et al., 2012), and obesity (Bexelius et al., 2010), mHealth poses some unique challenges that preclude direct application of existing methodologies for dynamic treatment regimes. For example, mHealth applications typically have no definite time horizon in the sense that treatment decisions are made continually throughout the life of the patient with no fixed time point for the final treatment decision; estimation of an optimal treatment strategy must utilize data collected over a much shorter time period than that over which treatment would be applied in practice; the momentary signal may be weak and may not directly measure the outcome of interest; and estimation of optimal treatment strategies must be done online as data accumulate.

This work is motivated in part by our involvement in a study of mHealth as a management tool for type 1 diabetes. Type 1 diabetes is an autoimmune disease wherein the pancreas produces insufficient levels of insulin, a hormone needed to regulate blood glucose concentration. Patients with type 1 diabetes are continually engaged in management activities including monitoring glucose levels, timing and dosing insulin injections, and regulating diet and physical activity. Increased glucose monitoring and attention to self-management facilitate more frequent treatment adjustments and have been shown to improve patient outcomes (Levine et al., 2001; Haller et al., 2004; Ziegler et al., 2011). Thus, patient outcomes have the potential to be improved by diabetes management tools which are deeply tailored to the continually evolving health status of each patient. Mobile technologies can be used to collect data on physical activity, glucose, and insulin at a fine granularity in an outpatient setting (Maahs et al., 2012). There is great potential for using these data to create comprehensive and accessible mHealth interventions for clinical use. We envision application of this work for use before the artificial pancreas (Weinzimer et al., 2008; Kowalski, 2015; Bergenstal et al., 2016) becomes widely available.

In our motivating example as well as other mHealth applications, the goal is to treat a chronic disease over the long term. However, data will typically be collected over a short time period. Because the time frame of data collection is much shorter than the time frame of application, standard methods for longitudinal data analysis, such as generalized estimating equations or mixed models, cannot be used. We assume that the data collected in the field consist of a sample from a stationary Markov process, which allows us to estimate a dynamic treatment regime that will lead to good outcomes over the long term using data collected over a much shorter time period.

The sequential decision making process can be modeled as a Markov decision process (Puterman, 2014) and the optimal treatment regime can be estimated using reinforcement learning algorithms such as Q-learning (Murphy, 2005; Zhao et al., 2009; Tang and Kosorok, 2012; Schulte et al., 2014). Ertefaie (2014) proposed a variant of greedy gradient Q-learning (GGQ) to estimate optimal dynamic treatment regimes in infinite horizon settings (see also Maei et al., 2010). In GGQ, the form of the estimated Q-function dictates the form of the estimated optimal treatment regime. Thus, one must choose between a parsimonious model for the Q-function at the risk of model misspecification or a complex Q-function that yields unintelligible treatment regimes. Furthermore, GGQ requires modeling a non-smooth function of the data, which creates complications (Laber et al., 2014; Linn et al., 2017). Applications of mHealth require methods that can both estimate a policy from a fixed sample of retrospective data (offline estimation) and estimate a policy that is updated as data accumulate (online estimation). Online estimation has been given considerable attention in the field of reinforcement learning, particularly in engineering applications (Doya, 2000; Kober and Peters, 2012), with special consideration given to algorithms that provide fast updates in situations where new data accumulate multiple times per second. In our mHealth examples, treatment decisions may be made multiple times per day or even every hour or minute; however, rapid updating of the estimated policy at the scale of engineering applications is not needed. Therefore, we are able to focus on different aspects of the problem without needing to ensure very fast estimation. We propose an alternative estimation method for infinite horizon dynamic treatment regimes that is suited to mHealth applications. Our approach, which we call V-learning, involves estimating the optimal policy among a prespecified class of policies (Zhang et al., 2012, 2013). It requires minimal assumptions about the data-generating process and permits estimating a randomized decision rule that can be implemented online as data accumulate.

In Section 2, we describe the setup and present our method for offline estimation using data from a micro-randomized trial or observational study. In Section 3, we extend our method for application to online estimation with accumulating data. Theoretical results, including consistency and asymptotic normality of the proposed estimators, are presented in Section 4. We compare the proposed method to GGQ using simulated data in Section 5. A case study using data from patients with type 1 diabetes is presented in Section 6 and we conclude with a discussion in Section 7. Proofs of technical results are in the Appendix.

2 Offline estimation from observational data

We assume that the available data are $\left\{ \left(S_i^1, A_i^1, S_i^2, \dots, S_i^{T_i}, A_i^{T_i}, S_i^{T_i+1} \right) \right\}_{i=1}^n$, which comprise n independent, identically distributed trajectories $(S^1, A^1, S^2, \dots, S^T, A^T, S^{T+1})$, where: $S^t \in \mathbb{R}^p$ denotes a summary of patient information collected up to and including time t ; $A^t \in \mathcal{A}$ denotes the treatment assigned at time t ; and $T \in \mathbb{Z}_+$ denotes the (possibly random) patient follow-up time. In the motivating example of type 1 diabetes, S^t could contain a patient's blood glucose, dietary intake, and physical activity in the hour leading up to time t and A^t could denote an indicator that an insulin injection is taken at time t . We assume that the data-generating model is a time-homogeneous Markov process so that $S^{t+1} \perp\!\!\!\perp (A^{t-1}, S^{t+1}, \dots, A^1, S^1) | (A^t, S^t)$ and the conditional density $p(s^{t+1} | a^t, s^t)$ is the same for all $t \geq 1$. Let $L^t \in \{0, 1\}$ denote an indicator that the patient is still in follow-up at time t , i.e., $L^t = 1$ if the patient is being followed at time t and zero otherwise. We assume that L^t is contained in S^t so that $P(L^{t+1} = 1 | A^t, S^t, \dots, A^1, S^1) = P(L^{t+1} = 1 | A^t, S^t)$. It is not necessary for time points to be evenly spaced or homogeneous across patients. For example, we can define the decision times as those times at which data are observed and include time since the previous observation in S^t . Defining the decision times in this way allows us to effectively handle intermittent missing data, and thus we can assume that $L^t = 0$ implies $L^{t+1} = 0$ with probability one. Furthermore, we assume a known utility function $u : \mathbb{R}^p \times \mathcal{A} \times \mathbb{R}^p \rightarrow \mathbb{R}$ so that $U^t = u(S^{t+1}, A^t, S^t)$ measures the 'goodness' of choosing treatment A^t in state S^t and subsequently transitioning to state S^{t+1} . In our motivating example, the utility at time t could be a measure of how infrequently the patient's blood glucose concentration deviates from the optimal range over the hour preceding and following time t . The goal is to select treatments to maximize expected cumulative utility; treatment selection is formalized using a treatment regime (Schulte et al., 2014; Kosorok and Moodie, 2015) and the utility associated with any regime is defined using potential outcomes (Rubin, 1978).

Let $\mathcal{B}(\mathcal{A})$ denote the space of probability distributions over $\mathcal{A}$. A treatment regime in this context is a function $\pi : \text{dom } S^t \rightarrow \mathcal{B}(\mathcal{A})$ so that, under π , a decision maker presented with state $S^t = s^t$ at time t will select action $a^t \in \mathcal{A}$ with probability $\pi(a^t; s^t)$. Define $\vec{a}^t = (a^1, \dots, a^t) \in \mathcal{A}^t$, and $\vec{a}^\infty = (a^1, a^2, \dots) \in \mathcal{A}^\infty$. The set of potential outcomes is

$$\mathcal{W}^* = \left\{ S^1, S^{*2}(a^1), \dots, S^{*T^*}(\vec{a}^\infty) \left(\vec{a}^{T^*}(\vec{a}^\infty) - 1 \right) : T^*(\vec{a}^\infty) = \inf \left\{ t \geq 1 : L^{*t}(\vec{a}^t - 1) = 0 \right\}, \vec{a}^\infty \in \mathcal{A}^\infty \right\},$$

where $S^{*t}(\vec{a}^t - 1)$ is the potential state and $L^{*t}(\vec{a}^t - 1)$ is the potential follow-up status at time t under treatment sequence $\vec{a}^t - 1$. Thus, the potential utility at time t is $U^{*t}(\vec{a}^t) = u \left(S^{*(t+1)}(\vec{a}^t, a^t, S^{*t}(\vec{a}^t - 1)) \right)$. For any π , define $\{ \xi_\pi^t(\cdot) \}_{t \geq 1}$ to be a sequence of independent, $\mathcal{A}$ -valued stochastic processes indexed by $\text{dom } S^t$ such that $P \left\{ \xi_\pi^t(S^t) = a^t \right\} = \pi(a^t; S^t)$. The potential follow-up time under π is

$$T^*(\pi) = \sum_{t \geq 1} \sum_{\bar{a}^t \in \mathcal{A}^t} t \mathbb{1} \left\{ \sup_{\bar{a}^t} \bar{a}^t + 1 T^* \left(\bar{a}^t, \bar{a}^t + 1 \right) = t \right\} \prod_{v=1}^t \mathbb{1} \left[\xi_{\pi}^v \left\{ S^{*v} \left(\bar{a}^v - 1 \right) \right\} = a^v \right],$$

where $\bar{a}^t + 1 = (a^t + 1, a^t + 2, \dots)$. The potential utility under π at time t is

$$U^{*t}(\pi) = \begin{cases} \sum_{\bar{a}^t \in \mathcal{A}^t} U^{*t}(\bar{a}^t) \prod_{v=1}^t \mathbb{1} \left[\xi_{\pi}^v \left\{ S^{*v} \left(\bar{a}^v - 1 \right) \right\} = a^v \right], & \text{if } T^*(\pi) \geq t \\ 0, & \text{otherwise,} \end{cases}$$

where $S^{*1}(\bar{a}^0) = S^1$. Thus, utility is set to zero after a patient is lost to follow-up. However, in certain situations, utility may be constructed so as to take a negative value at the time point when the patient is lost to follow-up, e.g., if the patient discontinues treatment because of a negative effect associated with the intervention. Define the state-value function $V(\pi, s^t) = \mathbb{E} \left[\sum_{k \geq 0} \gamma^k U^{*(t+k)}(\pi) | S^t = s^t \right]$ (Sutton and Barto, 1998), where $\gamma \in (0, 1)$ is a fixed constant that captures the trade-off between short- and long-term outcomes. For any distribution $\mathcal{R}$ on $\text{dom } S^1$, define the value function with respect to reference distribution $\mathcal{R}$ as $V_{\mathcal{R}}(\pi) = \int V(\pi, s) d\mathcal{R}(s)$; throughout, we assume that this reference distribution is fixed. The reference distribution can be thought of as a distribution of initial states and we estimate it from the data in the implementation in Sections 5 and 6. For a prespecified class of regimes, Π , the optimal regime, $\pi_{\mathcal{R}}^{\text{opt}} \in \Pi$, satisfies $V_{\mathcal{R}}(\pi_{\mathcal{R}}^{\text{opt}}) \geq V_{\mathcal{R}}(\pi)$ for all $\pi \in \Pi$. The goal is to estimate $\pi_{\mathcal{R}}^{\text{opt}}$ using data collected from n patients, where patient i is followed for T_i time points, $i = 1, \dots, n$. Thus, T_i represents the number of treatment decisions made for patient i in the observed data; however, because the observed data are assumed to come from a time-homogeneous Markov chain, the estimated policy could be applied in a population indefinitely.

To construct an estimator of $\pi_{\mathcal{R}}^{\text{opt}}$, we make a series of assumptions that connect the potential outcomes in W^* with the data-generating model.

Assumption 1. Strong ignorability, $A^t \perp\!\!\!\perp W^* | S^t$ for all t .

Assumption 2. Consistency, $S^t = S^{*t}(\bar{A}^{t-1})$ for all t and $T = T^*(\bar{A}^{\infty})$.

Assumption 3. Positivity, there exists $c_0 > 0$ so that $P(A^t = a^t | S^t = s^t) \geq c_0$ for all $a^t \in \mathcal{A}, s^t \in \text{dom } S^t$, and all t .

In addition, we implicitly assume that there is no interference among the experimental units. These assumptions are common in the context of estimating dynamic treatment regimes (Robins, 2004; Hernan and Robins, 2010; Schulte et al., 2014). Assumption 1 implies that there are no unmeasured confounders and assumptions 1 and 3 hold by construction in a micro-randomized trial (Klasnja et al., 2015; Liao et al., 2016).

Let $\mu^t(a^t, s^t) = P(A^t = a^t | S^t = s^t)$ for each $t \geq 1$. In a micro-randomized trial, $\mu^t(a^t, s^t)$ is a known randomization probability; in an observational study, it must be estimated from the data. The following lemma characterizes $V_{\mathcal{R}}(\pi)$ for any regime, π , in terms of the data-generating model (see also Lemma 4.1 of Murphy et al., 2001). A proof is provided in the appendix.

Lemma 2.1. Let π denote an arbitrary regime and $\gamma \in (0, 1)$ a discount factor. Then, under assumptions 1–3 and provided interchange of the sum and integration is justified, the state-value function of π at s^t is

$$V(\pi, s^t) = \sum_{k \geq 0} \mathbb{E} \left[\gamma^k U^{t+k} \left\{ \prod_{v=0}^k \frac{\pi(A^{v+t}; S^{v+t})}{\mu^{v+t}(A^{v+t}; S^{v+t})} \right\} \mid S^t = s^t \right]. \quad (1)$$

The preceding result will form the basis for an estimating equation for $V_{\mathcal{R}}(\pi)$. Write the right hand side of (1) as

$$\begin{aligned} V(\pi, S^t) &= \mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} \left(U^t + \gamma \sum_{k \geq 0} \mathbb{E} \left[\gamma^k U^{t+k+1} \left\{ \prod_{v=0}^k \frac{\pi(A^{v+t+1}; S^{v+t+1})}{\mu^{v+t+1}(A^{v+t+1}; S^{v+t+1})} \right\} \mid S^{t+1} \right] \mid S^t \right) \right] \\ &= \mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} \left(U^t + \gamma V(\pi, S^{t+1}) \right) \mid S^t \right], \end{aligned}$$

from which it follows that

$$0 = \mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} \left\{ U^t + \gamma V(\pi, S^{t+1}) - V(\pi, S^t) \right\} \mid S^t \right].$$

Subsequently, for any function ψ defined on $\text{dom } S^t$, the state-value function satisfies

$$0 = \mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} \left\{ U^t + \gamma V(\pi, S^{t+1}) - V(\pi, S^t) \right\} \psi(S^t) \right], \quad (2)$$

which is an importance-weighted variant of the well-known Bellman optimality equation (Sutton and Barto, 1998).

Let $V(\pi, s; \theta^\pi)$ denote a model for $V(\pi, s)$ indexed by $\theta^\pi \in \Theta \subseteq \mathbb{R}^q$. We assume that the map $\theta^\pi \mapsto V(\pi, s; \theta^\pi)$ is differentiable everywhere for each fixed s and π . Let $\nabla_{\theta^\pi} V(\pi, s; \theta^\pi)$ denote the gradient of $V(\pi, s; \theta^\pi)$ and define

$$\begin{aligned} \Lambda_n(\pi, \theta^\pi) &= \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^{T_i} \frac{\pi(A_i^t; S_i^t)}{\mu^t(A_i^t; S_i^t)} \left\{ U_i^t + \gamma V(\pi, S_i^{t+1}; \theta^\pi) - V(\pi, S_i^t; \theta^\pi) \right\} \\ \nabla_{\theta^\pi} V(\pi, S_i^t; \theta^\pi). \end{aligned} \quad (3)$$

Given a positive definite matrix $\Omega \in \mathbb{R}^{q \times q}$ and penalty function $\mathcal{P} : \mathbb{R}^q \rightarrow \mathbb{R}_+$, define $\hat{\theta}_n^\pi = \arg \min_{\theta^\pi \in \Theta} \left\{ \Lambda_n(\pi, \theta^\pi)^\top \Omega \Lambda_n(\pi, \theta^\pi) + \lambda_n \mathcal{P}(\theta^\pi) \right\}$, where λ_n is a tuning parameter. Subsequently, $V(\pi, s; \hat{\theta}_n^\pi)$ is the estimated state-value function under π in state s . Thus, given a reference distribution, $\mathcal{R}$, the estimated value of a regime, π , is $\hat{V}_{n, \mathcal{R}}(\pi) = \int V(\pi, s; \hat{\theta}_n^\pi) d\mathcal{R}(S)$ and the estimated optimal regime is $\hat{\pi}_n = \arg \max_{\pi \in \Pi} \hat{V}_{n, \mathcal{R}}(\pi)$. The idea of V-learning is to use estimating equation (3) to estimate the value of any policy and maximize estimated value over a class of policies; we will discuss strategies for this maximization below.

V-learning requires a parametric class of policies. Assuming that there are K possible treatments, $a_1, \dots, a_K$, we can define a parametric class of policies as follows. Define $\pi(a_j; s, \beta) = \exp(s^\top \beta_j) / \left\{ 1 + \sum_{k=1}^{K-1} \exp(s^\top \beta_k) \right\}$ for $j = 1, \dots, K-1$, and $\pi(a_K; s) = 1 / \left\{ 1 + \sum_{k=1}^{K-1} \exp(s^\top \beta_k) \right\}$. This defines a class of randomized policies, Π , parametrized by $\beta = (\beta_1^\top, \dots, \beta_{K-1}^\top)^\top$, where β_k is a vector of parameters for the k -th treatment. Under a policy in this class defined by β , actions are selected stochastically according to the probabilities $\pi(a_j; s, \beta)$, $j = 1, \dots, K$. In the case of a binary treatment, a policy in this class reduces to $\pi(1; s, \beta) = \exp(s^\top \beta) / \{1 + \exp(s^\top \beta)\}$ and $\pi(0; s, \beta) = 1 / \{1 + \exp(s^\top \beta)\}$ for a $p \times 1$ vector β . This class of policies is used in the implementation in Sections 5 and 6.

V-learning also requires a class of models for the state value function indexed by a parameter, θ^π . We use a basis function approximation (Hastie et al., 2009; Long et al., 2010). Let $\Phi = (\phi_1, \dots, \phi_q)^\top$ be a vector of prespecified basis functions and let $\Phi(s_i^t) = \{\phi_1(s_i^t), \dots, \phi_q(s_i^t)\}^\top$. Let $V(\pi, s_i^t; \theta^\pi) = \Phi(s_i^t)^\top \theta^\pi$. Under this working model,

$$\begin{aligned} \Lambda_n(\pi, \theta^\pi) = & \left[n^{-1} \sum_{i=1}^n \sum_{t=1}^{T_i} \frac{\pi(A_i^t; S_i^t)}{\mu^t(A_i^t; S_i^t)} \left\{ \gamma \Phi(S_i^t) \Phi(S_i^{t+1})^\top - \Phi(S_i^t) \Phi(S_i^t)^\top \right\} \right] \theta^\pi \\ & + n^{-1} \sum_{i=1}^n \sum_{t=1}^{T_i} \left\{ \frac{\pi(A_i^t; S_i^t)}{\mu^t(A_i^t; S_i^t)} U_i^t \Phi(S_i^t) \right\}. \end{aligned} \quad (4)$$

Computational efficiency is gained from the linearity of $V(\pi, s_i^t; \theta^\pi)$ in θ^π ; flexibility can be achieved through the choice of Φ . We recommend Gaussian basis functions as they offer the highest degree of flexibility. However, along with Gaussian basis functions, we also examine the performance of V-learning using linear and polynomial basis functions in Sections 5 and 6 as these offer reasonable alternatives.

The algorithm for V-learning is given in Algorithm 1 below. The algorithm can be terminated when $\|\beta^k - \beta^{k-1}\|$ is below some small threshold. The update in step can be achieved with a variety of existing optimization methods. In our implementation, we use the BFGS algorithm (Dai, 2002) as implemented in the optim function in R software (R Core

Team, 2016). Because the objective function is not necessarily convex, care must be taken when selecting the starting point, β^1 (step 2). In our implementation, we use simulated annealing as implemented in the `optim` function in R to find an appropriate starting point.

Algorithm 1: V-learning.

1 Initialize a class of policies, $\Pi = \{\pi_\beta : \beta \in \mathcal{B}\}$, and a model, $V(\pi, s; \theta^\pi)$;
 2 Set $k = 1$ and initialize β^1 to a starting value in $\mathcal{B}$;
 3 **while** *Not converged* **do**
 4 Estimate $\hat{\theta}_n^{\pi\beta^k} = \arg \min_{\theta \in \Theta} \left\{ \Lambda_n(\pi\beta^k, \theta^{\pi\beta^k})^\top \Omega \Lambda_n(\pi\beta^k, \theta^{\pi\beta^k}) + \lambda_n \mathcal{P}(\theta^{\pi\beta^k}) \right\}$;
 5 Evaluate $\hat{V}_n, \mathcal{R}(\pi\beta^k) = \int V(\pi\beta^k, S; \hat{\theta}_n^{\pi\beta^k}) d\mathcal{R}(s)$;
 6 Set $\beta^{k+1} = \beta^k + \alpha^k \nabla_{\beta_k} \hat{V}_n, \mathcal{R}(\pi\beta^k)$ for some step size, α^k , where $\nabla_{\beta_k} \hat{V}_n, \mathcal{R}(\pi\beta^k)$ is the gradient of $\hat{V}_n, \mathcal{R}(\pi\beta^k)$;
 7 **end**

2.1 Greedy gradient Q-learning

Here we briefly discuss an existing method for infinite horizon dynamic treatment regimes, which will be used for comparison in the simulation studies in Section 5.

Ertefaie (2014) introduced greedy gradient Q-learning (GGQ) for estimating dynamic treatment regimes in infinite horizon settings (see also Maei et al., 2010; Murphy et al., 2016).

Define $Q^\pi(s^t, a^t) = \mathbb{E} \left[\sum_{k \geq 0} \gamma^k U^{t+k}(\pi) \mid S^t = s^t, A^t = a^t \right]$. The Bellman optimality equation (Sutton and Barto, 1998) is

$$Q^{\text{opt}}(s^t, a^t) = \mathbb{E} \left[U^t + \gamma \max_{a \in \mathcal{A}} Q^{\text{opt}}(S^{t+1}, a) \mid S^t = s^t, A^t = a^t \right]. \quad (5)$$

Let $Q(s, a; \eta^{\text{opt}})$ be a parametric model for $Q^{\text{opt}}(s, a)$ indexed by $\eta^{\text{opt}} \in H \subseteq \mathbb{R}^q$. In our implementation, we model $Q(s, a; \eta^{\text{opt}})$ as a linear function with interactions between all state variables and treatment. The Bellman optimality equation motivates the estimating equation

$$D_n(\eta^{\text{opt}}) = \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^{T_i} \left\{ U_i^t + \gamma \max_{a \in \mathcal{A}} Q(S_i^{t+1}, a; \eta^{\text{opt}}) - Q(S_i^t, A_i^t; \eta^{\text{opt}}) \right\} \nabla_{\eta^{\text{opt}}} Q(S_i^t, A_i^t; \eta^{\text{opt}}). \quad (6)$$

For a positive definite matrix, Ω , we estimate η^{opt} using $\hat{\eta}_n^{\text{opt}} = \arg \min_{\eta \in H} D_n(\eta)^\top \Omega D_n(\eta)$.

The estimated optimal policy in state s selects action $\hat{\pi}_n(s) = \max_{a \in \mathcal{A}} Q(s, a; \hat{\eta}_n^{\text{opt}})$. This optimization problem is non-convex and non-differentiable in η^{opt} . However, it can be solved with a generalization of the greedy gradient Q-learning algorithm of Maei et al. (2010), and hence is referred to as GGQ by Ertefaie (2014) and in the following.

The performance of GGQ has been demonstrated in the context of chronic diseases with large sample sizes and a moderate number of time points. However, in mHealth applications, it is common to have small sample sizes and a large number of time points, with decisions occurring at a fine granularity. In GGQ, the estimated policy depends directly $\hat{\eta}_n^{\text{opt}}$ and, therefore, depends on modeling the transition probabilities of the data-generating process. Furthermore, estimating equation (6) contains a non-smooth max operator, which makes estimation difficult without large amounts of data (Laber et al., 2014; Linn et al., 2017). V-learning only requires modeling the policy and the value function rather than the data-generating process and directly maximizes estimated value over a class of policies, thereby avoiding the non-smooth max operator in the estimating equation (compare equations (3) and (6)); these attributes may prove advantageous in mHealth settings.

3 Online estimation from accumulating data

Suppose we have accumulating data $\{(S_i^1, A_i^1, S_i^2, \dots)\}_{i=1}^n$, where S_i^t and A_i^t represent the state and action for patient $i = 1, \dots, n$ at time $t - 1$. At each time t , we estimate an optimal policy in a class, Π , using data collected up to time t , take actions according to the estimated optimal policy, and estimate a new policy using the resulting states. Let $\hat{\pi}_n^t$ be the estimated policy at time t , i.e., $\hat{\pi}_n^t$ is estimated after observing state S^{t+1} and before taking action A^{t+1} . If Π is a class of randomized policies, we can select an action for a patient presenting with $S^{t+1} = s^{t+1}$ according to $\hat{\pi}_n^t(\cdot; s^{t+1})$, i.e., we draw A^{t+1} according to the distribution $P(A^{t+1} = a) = \hat{\pi}_n^t(a; s^{t+1})$. If a class of deterministic policies is of interest, we can inject some randomness into $\hat{\pi}_n^t$ to facilitate exploration. One way to do this is an ϵ -greedy strategy (Sutton and Barto, 1998), which selects the estimated optimal action with probability $1 - \epsilon$ and otherwise samples equally from all other actions. Because an ϵ -greedy strategy can be used to introduce randomness into a deterministic policy, we can assume a class of randomized policies.

At each time $t - 1$, let $\hat{\theta}_{n,t}^\pi = \arg \min_{\theta^\pi \in \Theta} \left\{ \Lambda_{n,t}(\pi, \theta^\pi)^\top \Omega \Lambda_{n,t}(\pi, \theta^\pi) + \lambda_n \mathcal{P}(\theta^\pi) \right\}$ where Ω , λ_n , and $\mathcal{P}$ are as defined in Section 2 and

$$\Lambda_{n,t}(\pi, \theta^\pi) = \frac{1}{n} \sum_{i=1}^n \sum_{v=1}^t \frac{\pi(A_i^v; S_i^v)}{\hat{\pi}_n^{v-1}(A_i^v; S_i^v)} \left\{ U_i^v + \gamma V(\pi, S_i^{v+1}; \theta^\pi) - V(\pi, S_i^v; \theta^\pi) \right\} \nabla_{\theta^\pi} V(\pi, S_i^v; \theta^\pi) \quad (7)$$

with $\hat{\pi}_n^0$ some initial randomized policy. We note that estimating equation (7) is similar to (3), except that $\hat{\pi}_n^{v-1}$ replaces μ^v as the data-generating policy. Given the estimator of the value of π at time t , $\hat{V}_{n,\mathcal{R},t}(\pi) = \int V(\pi, s; \hat{\theta}_{n,t}^\pi) d\mathcal{R}(s)$, the estimated optimal policy at time t is $\hat{\pi}_n^t = \arg \max_{\pi \in \Pi} \hat{V}_{n,\mathcal{R},t}(\pi)$. In practice, we may choose to update the policy in batches rather than at every time point. An alternative way to encourage exploration through the action space is to choose $\hat{\pi}_n^t = \arg \max_{\pi \in \Pi} \{ \hat{V}_{n,\mathcal{R},t}(\pi) + \alpha^t \hat{\psi}^t(\pi) \}$ for some sequence α^t 0, where $\hat{\psi}^t(\pi)$ is a measure of uncertainty in $\hat{V}_{n,\mathcal{R},t}(\pi)$. An example of this is upper confidence bound sampling, or UCB (Lai and Robbins, 1985).

In some settings, when the data-generating process may vary across patients, it may be desirable to allow each patient to follow an individualized policy that is estimated using only that patient's data. Suppose that n patients are followed for an initial T_1 time points after which the policy $\hat{\pi}_n^1$ is estimated. Then, suppose that patient i follows $\hat{\pi}_n^1$ until time T_2 , when a policy $\hat{\pi}_i^2$ is estimated using only the states and actions observed for patient i . This procedure is then carried out until time T_K for some fixed K with each patient following their own individual policy which is adapted to match the individual over time. We may also choose to adapt the randomness of the policy at each estimation. For example, we could select $\epsilon_1 > \epsilon_2 > \dots > \epsilon_K$ and, following estimation k , have patient i follow policy $\hat{\pi}_i^k$ with probability $1 - \epsilon_k$ and policy $\hat{\pi}_n^1$ with probability ϵ_k . In this way, patients become more likely to follow their own individualized policy and less likely to follow the initial policy over time, reflecting increasing confidence in the individualized policy as more data become available. The same class of policies and model for the state value function can be used as in Section 2.

4 Theoretical results

In this section, we establish asymptotic properties of $\hat{\theta}_n^\pi$ and $\hat{\pi}_n$ for offline estimation.

Because the proposed online estimation procedure involves performing offline estimation repeatedly in smaller batches, our theoretical results apply to each smaller batch as the number of observations increases. Developing more general theory for online estimation is an interesting topic for future research. Throughout, we assume assumptions 1–3 from Section 2.

Let $\hat{\theta}_n^\pi = \arg \min_{\theta^\pi \in \Theta} \{ \Lambda_n(\pi, \theta^\pi)^\top \Omega \Lambda_n(\pi, \theta^\pi) + \lambda_n(\theta^\pi)^\top \theta^\pi \}$. Thus, we consider the special case where the penalty function is the squared Euclidean norm of Θ^π . We will assume that $\lambda_n = o_P(n^{-1/2})$. All of our results hold for any positive definite matrix, Ω . Assume the working model for the state value function introduced in Section 2, i.e.,

$V(\pi, s; \theta^\pi) = \Phi(s)^\top \theta^\pi$. For fixed π , denote the true θ^π by θ_0^π , i.e., $V(\pi, s) = \Phi(s)^\top \theta_0^\pi$. Let $v = \int \Phi(s) d\mathcal{R}(s)$ so that $V_{\mathcal{R}}(\pi) = v^\top \theta_0^\pi$. Define $\hat{V}_{n,\mathcal{R}}(\pi) = \{ \mathbb{E}_n \Phi(S) \}^\top \hat{\theta}_n^\pi$, where $\mathbb{E}_n$ denotes

the empirical measure of the observed data. Let $\Pi = \{\pi_\beta : \beta \in \mathcal{B}\}$ be a parametric class of policies and let $\hat{\pi}_n = \pi_{\hat{\beta}_n}$ where $\hat{\beta}_n = \arg \max_{\beta \in \mathcal{B}} \hat{V}_n(\pi_\beta)$.

Our main results are summarized in Theorems 4.2 and 4.3 below. Because each patient trajectory is a stationary Markov chain, we need to use asymptotic theory based on stationary processes; consequently, some of the required technical conditions are more difficult to verify than those for i.i.d. data. Define the bracketing integral for a class of functions, $\mathcal{F}$, by $J_{[]}(\delta, \mathcal{F}, L_r(P)) = \int_0^\delta \sqrt{\log N_{[]}(\epsilon, \mathcal{F}, L_r(P))} d\epsilon$, where the bracketing number for $\mathcal{F}$, $N_{[]}(\epsilon, \mathcal{F}, L_r(P))$, is the number of $L_r(P)$ ϵ -brackets needed such that each element of $\mathcal{F}$ is contained in at least one bracket (see Chapter 2 of Kosorok, 2008). For any stationary sequence of possibly dependent random variables, $\{X^t\}_{t=1}^\infty$; let $\mathcal{M}_b^c$ be the σ -field generated by $X^b, \dots, X^c$ and define $\zeta(k) = \mathbb{E}[\sup_{m \geq 1} \{ |P(B|\mathcal{M}_1^m) - P(B)| : B \in \mathcal{M}_{m+k}^\infty \}]$. We say that the chain $\{X^t\}_{t=1}^\infty$ is absolutely regular if $\zeta(k) \rightarrow 0$ as $k \rightarrow \infty$ (also called β -mixing in Chapter 11 of Kosorok, 2008). We make the following assumptions.

Assumption 4. There exists a $2 < \rho < \infty$ such that

1. $\mathbb{E}[|U^t|^{3\rho}] < \infty$, $\mathbb{E}[\|\Phi(S^t)\|^{3\rho}] < \infty$, and $\mathbb{E}\|S^t\|^{3\rho} < \infty$.
2. The sequence $\{(S^t, A^t)\}_{t=1}^\infty$ is absolutely regular with $\sum_{k=1}^\infty k^{2/(\rho-2)} \zeta(k) < \infty$.
3. The bracketing integral of the class of policies, $J_{[]}(\infty, \Pi, L_{3\rho}(P)) < \infty$.

Assumption 5. There exists some $c_1 > 0$ such that

$$\inf_{\pi \in \Pi} c_1 \mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} \left\{ \Phi(S^t) \Phi(S^t)^\top - \gamma^2 \Phi(S^{t+1}) \Phi(S^{t+1})^\top \right\} \right] c \geq c_1 \|c\|^2$$

for all $c \in \mathbb{R}^q$

Assumption 6. The map $\beta \mapsto V_{\mathcal{R}}(\pi_\beta)$ has a unique and well separated maximum over β in the interior of $\mathcal{B}$; let β_0 denote the maximizer.

Assumption 7. The following condition holds:

$$\sup_{\|\beta_1 - \beta_2\| \leq \delta} \|\pi_{\beta_1}(A; S) - \pi_{\beta_2}(A; S)\| \rightarrow 0 \text{ as } \delta \downarrow 0.$$

Remark 4.1. Assumption 4 requires certain finite moments and that the dependence between observations on the same patient vanishes as observations become further apart. In Lemma 8.2 in the appendix, we verify part 3 of assumption 4 and assumption 7 for the class of policies introduced in Section 2. However, note that the theory holds for any class of policies satisfying the given assumptions, not just the class considered here. Assumption 5 is needed to show the existence of a unique θ_0^π uniformly over Π and can be verified empirically by checking that certain data-dependent matrices are invertible. Assumption requires that the true optimal decision in each state is unique (see assumption A. 8 of

Ertefaie, 2014) and is a standard assumption in M-estimation (see chapter 14 of Kosorok, 2008). Assumption 7 requires smoothness on the class of policies.

The main results of this section are stated below. Theorem 4.2 states that there exists a unique solution to $0 = \mathbb{E} \Lambda_n(\pi, \theta^\pi)$ uniformly over Π and that the estimator $\hat{\theta}_n$ converges weakly to a mean zero Gaussian process in $\ell^\infty(\Pi)$.

Theorem 4.2. Under the given assumptions, the following hold.

1. For all $\pi \in \Pi$, there exists a $\theta_0^\pi \in \mathbb{R}^q$ such that $\mathbb{E} \Lambda_n(\pi, \theta^\pi)$ has a zero at $\theta^\pi = \theta_0^\pi$.

Moreover, $\sup_{\pi \in \Pi} \|\theta_0^\pi\| < \infty$ and $\sup_{\|\beta_1 - \beta_2\| \leq \delta} \|\theta_0^{\pi\beta_1} - \theta_0^{\pi\beta_2}\| \rightarrow 0$ as $\delta \downarrow 0$.

2. Let $\mathbb{G}(\pi)$ be a tight, mean zero Gaussian process indexed by Π with covariance $\mathbb{E}\{\mathbb{G}(\pi_1) \mathbb{G}(\pi_2)\} = w_1(\pi_1)^{-1} w_0(\pi_1, \pi_2) w_1(\pi_2)^{-1}$ where

$$w_0(\pi_1, \pi_2) = \mathbb{E} \left[\frac{\pi_1(A^t; S^t) \pi_2(A^t; S^t)}{\mu^t(A^t; S^t)^2} \left\{ U^t + \gamma \Phi(S^t + 1) \theta_0^{\pi_1} - \Phi(S^t) \theta_0^{\pi_1} \right\} \right. \\ \left. \left\{ U^t + \gamma \Phi(S^t + 1) \theta_0^{\pi_2} - \Phi(S^t) \theta_0^{\pi_2} \right\} \Phi(S^t) \Phi(S^t)^\top \right]$$

and

$$w_1(\pi) = \mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} \Phi(S^t) \left\{ \Phi(S^t) - \gamma \Phi(S^t + 1)^\top \right\} \right].$$

Then, $\sqrt{n}(\hat{\theta}_n^\pi - \theta_0^\pi) \rightsquigarrow \mathbb{G}(\pi)$ in $\ell^\infty(\Pi)$.

3. Let $\mathbb{G}(\pi)$ be as defined in part 2. Then, $\sqrt{n}\{\hat{V}_n, \hat{\mathcal{R}}(\pi) - V_{\mathcal{R}}(\pi)\} \rightsquigarrow \sqrt{\mathbb{I}} \mathbb{G}(\pi)$ in $\ell^\infty(\Pi)$.

Theorem 4.3 below gives us that the estimated optimal policy converges in probability to the true optimal policy over Π and that the estimated value of the estimated optimal policy converges to the true value of the estimated optimal policy.

Theorem 4.3. Under the given assumptions, the following hold.

1. Let $\hat{\beta}_n = \arg \max_{\beta \in \mathcal{B}} \hat{V}_n, \hat{\mathcal{R}}(\pi_\beta)$ and $\beta_0 = \arg \max_{\beta \in \mathcal{B}} V_{\mathcal{R}}(\pi_\beta)$. We have that $\|\hat{\beta}_n - \beta_0\| \xrightarrow{P} 0$.
2. Let $\hat{\beta}_n$ and β_0 be defined as in part 1. Then, $|V_{\mathcal{R}}(\pi_{\hat{\beta}_n}) - V_{\mathcal{R}}(\pi_{\beta_0})| \xrightarrow{P} 0$.
3. Let $\sigma_0^2 = \sqrt{\mathbb{I}} w_1(\pi_{\beta_0})^{-1} w_0(\pi_{\beta_0}, \pi_{\beta_0}) w_1(\pi_{\beta_0})^{-1} \sqrt{\mathbb{I}}$. Then, $\sqrt{n} \{\hat{V}_n, \hat{\mathcal{R}}(\pi_{\hat{\beta}_n}) - V_{\mathcal{R}}(\pi_{\beta_0})\} \rightsquigarrow N(0, \sigma_0^2)$.

4. A consistent estimator for σ_0^2 is

$$\hat{\sigma}_n^2 = \left\{ \mathbb{E} \Phi(S^t) \right\}^\top \hat{w}_1(\pi_{\hat{\beta}_n})^{-1} \hat{w}_0(\pi_{\hat{\beta}_n}, \pi_{\hat{\beta}_n}) \hat{w}_1(\pi_{\hat{\beta}_n})^{-\top} \left\{ \mathbb{E}_n \Phi(S^t) \right\},$$

where

$$\begin{aligned} \hat{w}_0(\pi_1, \pi_2) = \mathbb{E}_n \left[\frac{\pi_1(A^t; S^t) \pi_2(A^t; S^t)}{\mu^t(A^t; S^t)^2} \left\{ U^t + \gamma \Phi(S^t + 1) \hat{\theta}_n^{\pi_1} - \Phi(S^t) \hat{\theta}_0^{\pi_1} \right\} \right. \\ \left. \left\{ U^t + \gamma \Phi(S^t + 1) \hat{\theta}_n^{\pi_2} - \Phi(S^t) \hat{\theta}_n^{\pi_2} \right\} \Phi(S^t) \Phi(S^t)^\top \right] \end{aligned}$$

and

$$\hat{w}_1(\pi) = \mathbb{E}_n \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} \Phi(S^t) \left\{ \Phi(S^t) - \gamma \Phi(S^t + 1) \right\}^\top \right].$$

Proofs of the above results are in the Appendix along with a result on bracketing entropy that is needed for the proof of Theorem 4.2 and a proof that the class of policies introduced above satisfies the necessary bracketing integral assumption.

5 Simulation experiments

In this section, we examine the performance of V-learning on simulated data. Section 5.1 contains results for offline estimation and Section 5.2 contains results for online estimation. All simulation results are averaged across 50 replications.

5.1 Offline simulations

Our implementation of V-learning follows the setup in Section 2. Maximizing $\hat{V}_{n, \mathcal{R}}(\pi)$ is done using a combination of simulated annealing and the BFGS algorithm as implemented in the `optim` function in R software (R Core Team, 2016). We note that $\hat{V}_{n, \mathcal{R}}(\pi)$ is differentiable in π , thereby avoiding some of the computational complexity of GGQ. However, the objective is not necessarily convex. In order to avoid local maxima, simulated annealing with 1000 function evaluations is used to find a neighborhood of the maximum; this solution is then used as the starting value for the BFGS algorithm.

We use the class of policies introduced in Section 2. Although we maximize the value over a class of randomized policies, the true optimal policy is deterministic. To prevent the coefficients of $\hat{\beta}_n$ from diverging to infinity, we add an L2 penalty when maximizing over β .

To prevent overfitting, we use an L2 penalty when computing $\hat{\theta}_n^\pi$, i.e., $\mathcal{P}(\theta^\pi) = (\theta^\pi)^\top \theta^\pi$. We let Ω be the identity matrix. Simulation results with alternate choices of $\mathcal{P}$ and Ω are given in the appendix. Tuning parameters can be used to control the amount of randomness in the estimated policy. For example, increasing the penalty when computing $\hat{\beta}_n$ is one way to

encourage exploration through the action space because $\beta = 0$ defines a policy where each action is selected with equal probability.

We consider three different models for the state-value function: (i) linear; (ii) second degree polynomial; and (iii) Gaussian radial basis functions (RBF). The Gaussian RBF is $\phi(x; \kappa, \tau^2) = \exp \left\{ - (x - \kappa)^2 / 2\tau^2 \right\}$. We use $\tau = 0.25$ and $\kappa = 0, 0.25, 0.5, 0.75, 1$ to create a basis of functions and apply this basis to the state variables after scaling them to be between 0 and 1. Each model also implicitly contains an intercept.

We begin with the following simple generative model. Let the two-dimensional state vector be $S_i^t = (S_{i,1}^t, S_{i,2}^t)^T$, $i = 1, \dots, n$, $t = 1, \dots, T$. We initiate the state variables as independent standard normal random variables and let them evolve according to

$S_{i,1}^t = (3/4) (2A_i^{t-1} - 1) S_{i,1}^{t-1} + (1/4) S_{i,1}^{t-1} S_{i,2}^{t-1} + \epsilon_1^t$ and $S_{i,2}^t = (3/4) (1 - 2A_i^{t-1}) S_{i,2}^{t-1} + (1/4) S_{i,1}^{t-1} S_{i,2}^{t-1} + \epsilon_2^t$, where A_i^t takes values in $\{0, 1\}$ and ϵ_1^t and ϵ_2^t are independent $N(0, 1/4)$ random variables. Define the utility function by

$U_i^t = u(S_i^t, A_i^t) = 2S_{i,1}^{t+1} + S_{i,2}^{t+1} - (1/4) (2A_i^t - 1)$. At each time t , we must make a decision to treat or not with the goal of maximizing the components of S while treating as few times as possible. Treatment has a positive effect on S_1 and a negative effect on S_2 . We generate A_i^t from a Bernoulli distribution with mean $1/2$. In estimation, we assume that the generating model for treatment is known, as would be the case in a micro-randomized trial.

We generate samples of n patients with T time points per patient from the given generative model after an initial burn-in period of 50 time points. The burn-in period ensures that our simulated data are sampled from an approximately stationary distribution. We estimate policies using V-learning with three different types of basis functions and GGQ. After estimating optimal policies, we simulate 100 patients following each estimated policy for 100 time points and take the mean utility under each policy as an estimate of the value of that policy. Estimated values are found in Table 1 with Monte Carlo standard errors along with the observed value. Recall that larger values are better. The policies estimated using V-learning produce better outcomes than the observational policy and the policy estimated using GGQ. V-learning produces the best outcomes using Gaussian basis functions. In the Appendix, we give results for the same simulation settings given in Table 1 for alternate choices of $\mathcal{P}$ and Ω . Table 8 presents results with the LASSO penalty when Ω is the identity matrix, Table 9 presents results with the L2 penalty when $\Omega = \{E_n S^t (S^t)^T\}^{-1}$, and Table 10 presents results with the LASSO penalty when $\Omega = \{E_n S^t (S^t)^T\}^{-1}$.

Next, we simulate cohorts of patients with type 1 diabetes to mimic the mHealth study of Maahs et al. (2012). Maahs et al. (2012) followed a small sample of youths with type 1 diabetes and recorded data at a fine granularity using mobile devices. Blood glucose levels were tracked in real time using continuous glucose monitoring, physical activity was measured continuously using accelerometers, and insulin injections were logged by an insulin pump. Dietary data were recorded by 24-hour recall over phone interviews.

In our simulation study, we divide each day of follow-up into 60 minute intervals. Thus, for one day of follow-up, we observe $T = 24$ time points per simulated patient and a treatment decision is made every hour. Our hypothetical mHealth study is designed to estimate an optimal dynamic treatment regime for the timing of insulin injections based on patient blood glucose, physical activity, and dietary intake with the goal of controlling future blood glucose as close as possible to the optimal range. To this end, we define the utility at time t as a weighted sum of hypo- and hyperglycemic episodes in the 60 minutes preceding and following time t . Weights are -3 when glucose < 70 (hypoglycemic), -2 when glucose > 150 (hyperglycemic), -1 when $70 < \text{glucose} \leq 80$ or $120 < \text{glucose} \leq 150$ (borderline hypo- and hyperglycemic), and 0 when $80 < \text{glucose} \leq 120$ (normal glycemia). Utility at each time point ranges from -6 to 0 with larger utilities (closer to 0) being more preferable. For example, a patient who presents with an average blood glucose of 155 mg/dL over time interval $t-1$, takes an action to correct their hyperglycemia, and presents with an average blood glucose of 145 mg/dL over time interval t would receive a utility of $U^t = -3$. Weights were chosen to reflect the relative clinical consequences of high and low blood glucose. For example, acute hypoglycemia, characterized by blood glucose levels below 70 mg/dL, is an emergency situation that can result in coma or death.

Simulated data are generated as follows. At each time point, patients are randomly chosen to receive an insulin injection with probability 0.3 , consume food with probability 0.2 , partake in mild physical activity with probability 0.4 , and partake in moderate physical activity with probability 0.2 . Grams of food intake and counts of physical activity are generated from normal distributions with parameters estimated from the data of Maahs et al. (2012). Initial blood glucose level for each patient is drawn from a normal distribution with mean 100 and standard deviation 25 . Define the covariates for patient i collected at time t by $(\text{Gl}_i^t, \text{Di}_i^t, \text{Ex}_i^t)^T$, where Gl_i^t is average blood glucose level, Di_i^t is total dietary intake, and Ex_i^t is total counts of physical activity as would be measured by an accelerometer. Glucose levels evolve according to

$$\begin{aligned} \text{Gl}^t = & \mu(1 - \alpha_1) + \alpha_1 \text{Gl}^{t-1} + \alpha_2 \text{Di}^{t-1} + \alpha_3 \text{Di}^{t-2} + \alpha_4 \text{Ex}^{t-1} + \alpha_5 \text{Ex}^{t-2} \\ & + \alpha_6 \text{In}^{t-1} + \alpha_7 \text{In}^{t-2} + e, \end{aligned} \quad (8)$$

where In^t is an indicator of an insulin injection received at time t and $e \sim \mathcal{N}(0, \sigma^2)$. We use the parameter vector $\alpha = (\alpha_1, \dots, \alpha_7)^T = (0.9, 0.1, 0.1, -0.01, -0.01, -2, -4)^T$, $\mu = 100$, and $\sigma = 5.5$ based on a linear model fit to the data of Maahs et al. (2012). The known lag-time in the effect of insulin is reflected by $\alpha_6 = -2$ and $\alpha_7 = -4$. Selecting $\alpha_1 < 1$ ensures the existence of a stationary distribution.

We define the state vector for patient i at time t to contain average blood glucose, total dietary intake, and total physical activity measured over previous time intervals; we include blood glucose and physical activity for the previous two time intervals and dietary intake for the previous four time intervals. Let n denote number of patients and T denote number of time points per patient. Our choices for n and T are based on what is feasible for an mHealth outpatient study (dietary data was collected on two days by Maahs et al., 2012). For each replication, the optimal treatment regime is estimated with V-learning using three different

types of basis functions and GGQ. The generative model for insulin treatment is not assumed to be known and we estimate it using logistic regression. We record mean outcomes in an independent sample of 100 patients followed for 100 time points with treatments generated according to each estimated optimal regime. Simulation results (estimated values under each regime and Monte Carlo standard errors along with observed values) are found in Table 2. Again, V-learning with Gaussian basis functions performs the best out of all methods, generally producing large values and small standard errors. V-learning with the linear model underperforms and GGQ underperforms somewhat more so.

5.2 Online simulations

In practice, it may be useful for patients to follow a dynamic treatment regime that is updated as new data are collected. Here we consider a hypothetical study wherein n patients are followed for an initial period of T' time points, an optimal policy is estimated, and patients are followed for an additional $T - T'$ time points with the estimated optimal policy being continuously updated. At each time point, $t \leq T$, actions are taken according to the most recently estimated policy. Recall that V-learning produces a randomized decision rule from which to sample actions at each time point. When selecting an action based on a GGQ policy, we incorporate an ϵ -greedy strategy by selecting the action recommended by the estimated policy with probability $1 - \epsilon$ and otherwise randomly selecting one of the other actions. At the t th estimation, we use $\epsilon = 0.5^t$, allowing ϵ to decrease over time to reflect increasing confidence in the estimated policy. A burn-in period of 50 time points is discarded to ensure that we are sampling from a stationary distribution. We estimate the first policy after 12 time points and a new policy is estimated every 6 time points thereafter. After T time points, we estimate the value as the average utility over all patients and all time points after the initial period.

Table 3 presents mean outcomes under policies estimated online using data generated according to the simple two covariate generative model introduced at the beginning of Section 5.1. There is some variability across n and T regarding which type of basis function is best, but V-learning with a polynomial basis generally produces the best outcomes. GGQ performs well in large samples.

Next, we study the performance of online V-learning in simulated mHealth studies of type 1 diabetes by following the generative model described in (8). Mean outcomes are found in Table 4. Gaussian V-learning performs the best out of all methods. Across all variants of V-learning, outcomes improve with increased follow-up time.

Finally, we consider online simulations using individualized policies as outlined at the end of Section 3. Consider the simple two covariate generative model introduced above but let state variables evolve according to $S_{i,1}^t = \mu_i(2A_i^{t-1} - 1) - S_{i,1}^{t-1} + (1/4)S_{i,1}^{t-1} - S_{i,2}^{t-1} + \epsilon_1^t$ and $S_{i,2}^t = \mu_i(1 - 2A_i^{t-1}) - S_{i,2}^{t-1} + (1/4)S_{i,1}^{t-1} - S_{i,2}^{t-1} + \epsilon_2^t$ where μ_i is a subject-specific term drawn uniformly between 0.4 and 0.9. Including μ_i ensures that the optimal policy differs across patients. Table 5 contains mean outcomes for online simulation where a universal policy is estimated using data from all patients and where individualized policies are estimated using only a single patient's data. Because data are generated in a such a way that the optimal

policy varies across patients, individualized policies achieve better outcomes than universal policies.

6 Case study: Type 1 diabetes

Machine learning is currently under consideration in type 1 diabetes through studies to build and test a “closed loop” system that joins continuous blood glucose monitoring and subcutaneous insulin infusion through an underlying algorithm. Known as the artificial pancreas, this technology has been shown to be safe in preliminary studies and is making headway from small hospital-based safety studies to large-scale outpatient effectiveness studies (Ly et al., 2014, 2015). Despite the success of the artificial pancreas, the rate of uptake may be limited and widespread use may not occur for many years (Kowalski, 2015). The proposed method may be useful for implementing mHealth interventions for use alongside the artificial pancreas or before it is widely available.

Studies have shown that data on food intake and physical activity to inform optimal decision making can be collected in an inpatient setting (see, e.g., Cobry et al., 2010; Wolever and Mullan, 2011). However, Maahs et al. (2012) demonstrated that rich data on the effect of food intake and physical activity can be collected in an outpatient setting using mobile technology. Here, we apply the proposed methodology to the observational data collected by Maahs et al. (2012).

The full data consist of $N = 31$ patients with type 1 diabetes, aged 12–18. Glucose levels were monitored using continuous glucose monitoring and physical activity tracked using accelerometers for five days. Dietary data were self-reported by the patient in telephone-based interviews for two days. Patients were treated using either an insulin pump or multiple daily insulin injections. We use data on a subset of $n = 14$ patients treated with an insulin pump for whom full follow-up is available on days when dietary information was recorded. This represents 28 patient-days of data, with which we use V-learning to estimate an optimal treatment policy. An advantage of mHealth is the ability to collect data passively, limiting the amount of missing data. There is no intermittent missingness in this data set.

The setup closely follows the simulation experiments in Section 5.1. Patient state at each time, t , is taken to be average glucose level and total counts of physical activity over the two previous 60 minute intervals and total food intake in grams over the four previous 60 minute intervals. The goal is to learn a policy to determine when to administer insulin injections based on prior blood glucose, dietary intake, and physical activity. The utility at time t is a weighted sum of glycemic events over the 60 minutes preceding and following time t with weights defined in Section 5.1. A treatment regime with large value will minimize the number of hypo- and hyperglycemic episodes weighted to reflect the clinical importance of each. We note that because $V(\pi, s; \theta^\pi)$ is linear in θ^π , we can evaluate $\hat{V}_{n, \mathcal{R}}(\pi)$ with only the mean of $\Phi(S)$ under $\mathcal{R}$. These were estimated from the data. Because we cannot simulate data following a given policy to estimate its value, we report the parametric value estimate $\hat{V}_{n, \mathcal{R}}(\hat{\pi}_n)$. Interpreting the parametric value estimate is difficult because of the effect the discount factor has on estimated value. We cannot compare parametric value

estimates to mean outcomes observed in the data. Instead, we use $\mathbb{E}_n \sum_{t \geq 0} \gamma^t U^t$ as an estimate of value under the observational policy.

We estimate optimal treatment strategies for two different action spaces. In the first, the only decision made at each time is whether or not to administer an insulin injection, i.e., the action space contains a single binary action. In the second, the action space contains all possible combinations of insulin injection, physical activity, and food intake. This corresponds to a hypothetical mHealth intervention where insulin injections are administered via an insulin pump and suggestions for physical activity and food intake are administered via a mobile app.

Table 6 contains parametric value estimates for policies estimated using V-learning for the two action spaces outlined above with different basis functions and discount factors. These results indicate that improvements in glycemic control can come from personalized and dynamic treatment strategies that account for food intake and physical activity. Improvement results from a dynamic insulin regimen (binary action space), and, in most cases, further improvement results from a comprehensive mHealth intervention including suggestions for diet and exercise delivered via mobile app in addition to insulin therapy (multiple action space). When considering multiple actions, the policy estimated using a polynomial basis and $\gamma = 0.7$ achieves a 64% increase in value and the policy estimated using a Gaussian basis and $\gamma = 0.8$ achieves a 68% increase in value over the observational policy. Although the small sample size is a weakness of this study, these results represent a significant improvement in value despite the sample size.

Finally, we use an example hyperglycemic patient to illustrate how an estimated policy would be applied in practice. One patient in the data presented at a specific time with an average blood glucose of 229 m g/dL over the previous hour and an average blood glucose of 283 m g/dL over the hour before that. The policy estimated with $\gamma = 0.7$ and a polynomial basis recommends each action according to the probabilities in Table 7. Because this patient presented with blood glucose levels that are higher than the optimal range, the policy recommends actions that would lower the patient's blood glucose levels, assigning a probability of 0.79 to insulin and a probability of 0.21 to insulin combined with activity.

7 Conclusion

The emergence of mHealth has provided great potential for the estimation and implementation of dynamic treatment regimes. Mobile technologies can be used both in the collection of rich longitudinal data to inform decision making and the delivery of deeply tailored interventions. The proposed method, V-learning, addresses a number of challenges associated with estimating dynamic treatment regimes in mHealth applications. V-learning directly estimates a policy which maximizes the value over a class of policies and requires minimal assumptions on the data-generating process. Furthermore, V-learning permits estimation of a randomized decision rule which can be used in place of existing strategies (e.g., ϵ -greedy) to encourage exploration in online estimation. A randomized decision rule can also provide patients with multiple treatment options. Estimation of an optimal policy for different populations can be handled through the use of different reference distributions.

V-learning and mobile technologies have the potential to improve patient outcomes in a variety of clinical areas. We have demonstrated, for example, that the proposed method can be used to estimate treatment regimes to reduce the number of hypoand hyperglycemic episodes in patients with type 1 diabetes. The proposed method could also be useful for other mHealth applications as well as applications outside of mHealth. For example, V-learning could be used to estimate dynamic treatment regimes for chronic illnesses using electronic health records data. Future research in this area may include increasing flexibility through use of a semiparametric model for the state-value function. Alternatively, nonlinear models for the state-value function may be informed by underlying theory or mathematical models of the system of interest. Data-driven selection of tuning parameters for the proposed method may help to improve performance. Developing theory for alternative penalty functions, such as the LASSO penalty, is another important step. Accounting for patient availability and feasibility of a sequence of treatments can be done by setting constraints on the class of policies. This will ensure that the resulting mHealth intervention is able to be implemented and that the recommended decisions are consistent with domain knowledge.

It would also be worthwhile to generalize our asymptotic results to permit nonstationarity. While we believe that stationarity is generally a reasonable assumption for moderate stretches of time—including when we do online estimation using moderately large batches of observations wherein both the patient dynamics and treatment policy remain approximately constant over each batch—stationarity would not in general hold when either the patient dynamics or treatment policy change more rapidly. For example, online estimation with treatment policy changes after each observation could induce nonstationarity. We conjecture that our asymptotic results will continue to hold in this setting, as our simulation studies in Section 5.2 seem to indicate.

Acknowledgments

We thank the editor, associate editor, and reviewers for helpful comments which led to a significantly improved paper.

Appendix

Proofs

Proof of Lemma 2.1. Let π be an arbitrary policy and $\gamma \in (0, 1)$ a fixed constant. Suppose we observe a state $S^t = s^t$ at time t and let $\vec{a}^{t-1} = (a^1, \dots, a^{t-1})$ be the sequence of actions resulting in $S^t = s^t$, i.e., $S^{*t}(\vec{a}^{t-1}) = s^t$. Let $\vec{a}^{k+1} = (a^t, \dots, a^{t+k}) \in \mathcal{A}^{k+1}$ be a potential sequence of actions taken from time t to time $t+k$. We have that

$$\begin{aligned}
V(\pi, s^t) &= \sum_{k \geq 0} \gamma^k \mathbb{E} \left[U^{*(t+k)}(\pi) \middle| S^t = s^t \right] \\
&= \sum_{k \geq 0} \gamma^k \mathbb{E} \left[\sum_{\bar{a}^t+k \in \mathcal{A}^{t+k}} U^{*(t+k)}(\bar{a}^t+k) \prod_{v=t}^{t+k} 1 \left[\xi^v \pi \left\{ S^{*v}(\bar{a}^v-1) \right\} = a^v \right] \middle| S^t = s^t \right] \\
&= \sum_{k \geq 0} \gamma^k \sum_{\bar{a}^k+1 \in \mathcal{A}^{k+1}} U^{*(t+k)}(\bar{a}^t-1, a^{-k}+1) \left[\prod_{v=1}^{t+k} \mathbb{E}(1 \left[\xi^v \pi \left\{ S^{*v}(\bar{a}^v-1) \right\} = a^v \right]) \middle| S^t = s^t \right] \\
&= \sum_{k \geq 0} \gamma^k \sum_{\bar{a}^k+1 \in \mathcal{A}^{k+1}} U^{*(t+k)}(\bar{a}^t-1, a^{-k}+1) \prod_{v=t}^{t+k} \pi \left\{ a^v; S^{*v}(\bar{a}^v-1) \right\} \prod_{v=t}^{t+k} \frac{\mu^v \left\{ a^v; S^{*v}(\bar{a}^v-1) \right\}}{\mu^v \left\{ a^v; S^{*v}(\bar{a}^v-1) \right\}} \\
&= \sum_{k \geq 0} \gamma^k \mathbb{E} \left[U^{t+k} \left[\prod_{v=0}^k \frac{\pi(a^t+v; s^t+v)}{\mu^t(a^t+v; s^t+v)} \right] \middle| S^t = s^t \right],
\end{aligned}$$

where we let $\pi(\bar{a}^t; s^t) = 0$ for all $\bar{a}^t$ and s^t whenever $t > T^*(\pi)$. The last equality uses the consistency and strong ignorability assumptions.

Proof of Theorem 4.2. Proof of part 1: We first note that θ_0^π must solve

$$0 = \mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} \left[U^t + \left\{ \gamma \Phi(S^t+1) - \Phi(S^t) \right\}^T \theta^\pi \right] \Phi(S^t) \right],$$

or

$$\mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} \left\{ \Phi(S^t) - \gamma \Phi(S^t+1) \right\}^T \right] \theta^\pi = \mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} U^t \Phi(S^t) \right],$$

which is equivalent to $w_1(\pi) \theta^\pi = w_2(\pi)$ where

$$w_2(\pi) = \mathbb{E} \left[\pi(A^t; S^t) \mu^t(A^t; S^t)^{-1} U^t \Phi(S^t) \right]. \text{ We have that}$$

$$\left\| \mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} U^t \Phi(S^t) \right] \right\| \leq c_0^{-1} \left(\mathbb{E} |U^t|^2 \right)^{1/2} \left(\mathbb{E} \|\Phi(S^t)\|^2 \right)^{1/2} < \infty,$$

by assumption 3, part 1 of assumption 4 and the Cauchy-Schwarz inequality. Let $c \in \mathbb{R}^q$ be arbitrary and note that

$$\begin{aligned}
&\mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} c^T \Phi(S^t) \Phi(S^t+1)^T c \right] \\
&\leq \left[\mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} c^T \Phi(S^t) \otimes^2 c \right] \cdot \mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} c^T \Phi(S^t+1) \otimes^2 c \right] \right]^{1/2},
\end{aligned}$$

by the Cauchy-Schwarz inequality, where $u^{\otimes 2} = uu^\top$. This implies that

$$\begin{aligned} c^\top w_1(\pi)c &\geq \mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} c^\top \Phi(S^t) \otimes^2 c \right] - \mathbb{E} \left[\frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} c^\top \Phi(S^t) \otimes^2 c \right]^{1/2} \mathbb{E} \left[\gamma^2 \frac{\pi(A^t; S^t)}{\mu^t(A^t; S^t)} c^\top \Phi(S^t + 1) \otimes^2 c \right]^{1/2} \\ &= A - A^{1/2} B^{1/2} \\ &= A^{1/2} (A^{1/2} - B^{1/2}) \\ &= \frac{A^{1/2}(A - B)}{A^{1/2} + B^{1/2}}, \end{aligned}$$

where we simplify notation by defining $A = \mathbb{E} \left[\pi(A^t; S^t) \mu^t(A^t; S^t)^{-1} c^\top \Phi(S^t) \otimes^2 c \right]$ and $B = \mathbb{E} \left[\gamma^2 \pi(A^t; S^t) \mu^t(A^t; S^t)^{-1} c^\top \Phi(S^t + 1) \otimes^2 c \right]$. We have that

$$\begin{aligned} A^{1/2} + B^{1/2} &\leq c_0^{-1/2} \|c\| \left\{ \mathbb{E} \left\| \Phi(S^t) \right\|^2 \right\}^{1/2} + c_0^{-1/2} \|c\| \left\{ \mathbb{E} \left\| \Phi(S^t + 1) \right\|^2 \right\}^{1/2} \\ &= 2c_0^{-1/2} \|c\| \left\{ \mathbb{E} \left\| \Phi(S^t) \right\|^2 \right\}^{1/2} \\ &< \infty, \end{aligned}$$

by the Cauchy-Schwarz inequality, the fact that $\mathbb{E} \left\| \Phi(S^t) \right\|^2 = \mathbb{E} \left\| \Phi(S^t + 1) \right\|^2$ by time-homogeneity, and part 1 of assumption 4. Also, $A - A^{1/2} B^{1/2}$ and $A - B - c_1 \|c\|^2$ by assumption 5. Thus,

$$\begin{aligned} A - A^{1/2} B^{1/2} &\geq \frac{c_1^{3/2} \|c\|^3}{2c_0^{-1/2} \|c\| \left\{ \mathbb{E} \left\| \Phi(S^t) \right\|^2 \right\}^{1/2}} \\ &= \frac{c_0^{1/2} c_1^{3/2} \|c\|^2}{2 \left\{ \mathbb{E} \left\| \Phi(S^t) \right\|^2 \right\}^{1/2}}, \end{aligned}$$

which finally implies that $w_1(\pi)$ is invertible and thus $\theta_0^\pi = w_1(\pi)^{-1} w_2(\pi)$ is well-defined uniformly over $\pi \in \Pi$. Using the fact that $c^\top w_1(\pi)c \leq k_0 \|c\|^2$ for a constant $k_0 > 0$, we can show that $\|w_1(\pi)^{-1}\| \leq k_1^{-1}$ for some constant $k_1 > 0$, where $\|\cdot\|$ is the usual matrix norm when applied to a matrix. Therefore,

$$\|\theta_0^\pi\| \leq k_1^{-1} \|w_2(\pi)\| \leq c_0^{-1} k_1^{-1} \left\{ \mathbb{E}(U^t)^2 \right\}^{1/2} \left\{ \mathbb{E} \left\| \Phi(S^t) \right\|^2 \right\}^{1/2} < \infty. \text{ Finally, it follows from assumptions 5 and 7 that } \sup \|\beta_1 - \beta_2\| \leq \delta \|\theta^{\pi\beta_1} - \theta^{\pi\beta_2}\| \rightarrow 0 \text{ as } \delta \downarrow 0.$$

Proof of part 2: Define

$$\mathcal{Z} = \left\{ \Phi(s^t) \Phi(s^t)^\top / \mu^t(a^t; s^t), \gamma \Phi(s^t) \Phi(s^t + 1)^\top / \mu^t(a^t; s^t), u^t \Phi(s^t) / \mu^t(a^t; s^t) \right\}.$$

Let G be an envelope for $\mathcal{G}$, for example $G(s^{t+1}, a^t, s^t) = \max_{g \in \mathcal{G}} g(s^{t+1}, a^t, s^t)$. By part 1 of assumption 4, $\mathbb{E}G^{3\rho} < \infty$. Part 4 of Lemma 8.1 below gives us that $\mathcal{G}$ is Donsker. Since Π satisfies $J_{\Pi} \{ \infty, \Pi, L_{3\rho}(P) \} < \infty$ to, we have that

$$\mathcal{F}_1 = \left\{ \Omega^{1/2} \frac{\pi(a^t; s^t)}{\mu^t(a^t; s^t)} \Phi(s^t) \{ \Phi(s^t) - \gamma \Phi(s^{t+1}) \}^\top : \pi \in \Pi \right\}$$

satisfies $J_{\Pi} \{ \infty, \mathcal{F}_1, L_{3\rho}(P) \} < \infty$ by parts 1 and 2 of Lemma 8.1 below. Moreover, $R(a^t, s^t) = \|\Phi(s^t)\| \cdot \|\Phi(s^t) - \gamma \Phi(s^{t+1})\| / \mu^t(a^t; s^t)$ is an envelope for $\mathcal{F}_1$ with $\mathbb{E}R^{3\rho} < \infty$ by assumption 3 and part 1 of assumption 4. Thus, $\mathcal{F}_1$ is Donsker. Let

$$\mathcal{F}_2 = \left\{ \Omega^{1/2} \frac{\pi(a^t; s^t)}{\mu^t(a^t; s^t)} u^t \Phi(s^t) : \pi \in \Pi \right\}.$$

Similar arguments yield that $\mathcal{F}_2$ is Donsker.

Now, let $\hat{A}(\pi) = \{\mathbb{E}_n f_{1\pi} : f_{1\pi} \in \mathcal{F}_1\}$ and $\hat{B}(\pi) = \{\mathbb{E}_n f_{2\pi} : f_{2\pi} \in \mathcal{F}_2\}$. Let $\hat{A}'(\pi) = \hat{A}(\pi) + \lambda_n \hat{A}(\pi)^{-\top}$. We have that $\hat{\theta}_n^\pi = \hat{A}'(\pi)^{-1} \hat{B}(\pi)$. Thus,

$$\begin{aligned} \sqrt{n}(\hat{\theta}_n^\pi - \theta_0^\pi) &= \sqrt{n} \left\{ \hat{A}'(\pi)^{-1} \hat{B}(\pi) - \hat{A}'(\pi)^{-1} \hat{A}'(\pi) \theta_0^\pi \right\} + o_P(1) \\ &= \hat{A}'(\pi)^{-1} \sqrt{n} \left\{ \hat{B}(\pi) - \hat{A}'(\pi) \theta_0^\pi \right\} + o_P(1) \\ &= \hat{A}'(\pi)^{-1} \sqrt{n} \left\{ \hat{B}(\pi) - \hat{A}(\pi) \theta_0^\pi \right\} + \hat{A}'(\pi)^{-1} \sqrt{n} \left\{ \hat{A}(\pi) - \hat{A}'(\pi) \right\} \theta_0^\pi + o_P(1) \\ &= \hat{A}'(\pi)^{-1} \sqrt{n} \left\{ \hat{B}(\pi) - \hat{A}(\pi) \theta_0^\pi \right\} + o_P(1) \end{aligned}$$

where $o_P(1)$ doesn't depend on π , because $\hat{A}'(\pi)^{-1} \xrightarrow{P} w_1(\pi)^{-1} < \infty$ uniformly over $\pi \in \Pi$ by assumption 3 and part 1 of assumption 4, $\sup_{\pi \in \Pi} \|\theta_0^\pi\| < \infty$ by part 1 of this theorem, and $\sqrt{n} \left\{ \hat{A}(\pi) - \hat{A}'(\pi) \right\} = -\sqrt{n} \lambda_n \hat{A}(\pi)^{-\top} = o_P(1)$ because $\lambda_n = o_P(n^{-1/2})$. Using arguments similar to those in the previous paragraph, one can show that $\mathcal{F}_3 = \{f_{2\pi} - f_{1\pi}\theta : f_{1\pi} \in \mathcal{F}_1, f_{2\pi} \in \mathcal{F}_2, \pi \in \Pi, \theta \in B_*\}$ is Donsker, where B_* is any finite collection of elements of $\mathbb{R}^q$. By part 1 of this theorem, there exists a bounded, closed set B_0 such that $\theta_0^\pi \in B_0$ for all $\pi \in \Pi$. Let $\mathbb{G}_n(\pi, \theta) = \sqrt{n}(\mathbb{E}_n - \mathbb{E})(f_{2\pi} - f_{1\pi}\theta)$. Note that

$$\begin{aligned} \sup_{\pi \in \Pi} \|\mathbb{G}_n(\pi, \theta_1) - \mathbb{G}_n(\pi, \theta_2)\| &\leq \sup_{\pi \in \Pi} \|\sqrt{n}(\mathbb{E}_n - \mathbb{E})f_{1\pi}\| \cdot \|\theta_1 - \theta_2\| \\ &\leq R^* \|\theta_1 - \theta_2\|, \end{aligned}$$

where $R^* = O_P(1)$ by the Donsker property of $\mathcal{F}_1$ and R^* doesn't depend on π . Thus, $\mathbb{G}_n(\pi, \theta)$ is stochastically equicontinuous on B_0 . Combined with the Donsker property of $\mathcal{F}_3$ for arbitrary B_* , we have that the class

$\mathcal{F}_4 = \{f_{2\pi} - f_{1\pi}\theta : f_{1\pi} \in \mathcal{F}_1, f_{2\pi} \in \mathcal{F}_2, \pi \in \Pi, \theta \in B_0\}$ is Donsker. Using Slutsky's Theorem, Theorem 11.24 of Kosorok (2008), the fact that $\mathcal{F}_1$ is Glivenko-Cantelli, and the fact that $\theta_0^\pi = (\mathbb{E}f_{1\pi})^{-1} \mathbb{E}f_{2\pi}$, we have that

$$\sqrt{n} (\hat{\theta}_0^\pi - \theta_0^\pi) = \hat{A}'(\pi)^{-1} \mathbb{G}_n(\pi, \theta_0^\pi) \rightsquigarrow w_1(\pi)^{-1} \mathbb{G}_0(\pi), \text{ in } \ell^\infty(\Pi) \text{ where } \mathbb{G}_0(\pi) \text{ is a mean zero Gaussian process indexed by } \Pi \text{ with covariance } \mathbb{E}\{\mathbb{G}_0(\pi_1) \mathbb{G}_0(\pi_2)\} = w_0(\pi_1, \pi_2).$$

Proof of part 3: We have that

$$\begin{aligned} \sqrt{n} \{\hat{V}_{n, \mathcal{R}}(\pi) - V_{\mathcal{R}}(\pi)\} &= \sqrt{n} \mathbb{E}_n \Phi(S^t) (\hat{\theta}_n^\pi - \theta_0^\pi) \\ &\rightsquigarrow v^T w_1(\pi)^{-1} \mathbb{G}_0(\pi) \end{aligned}$$

in $\ell^\infty(\Pi)$ by Slutsky's Theorem.

Proof of Theorem 4.3. Proof of part 1: Following part 3 of Theorem 4.2, we have that $\sup_{\beta \in \mathcal{B}} |\hat{V}_{n, \mathcal{R}}(\pi_\beta) - V_{\mathcal{R}}(\pi_\beta)| \xrightarrow{P} 0$. Combining this with the unique and well separated maximum condition (assumption 6), continuity of $V_{\mathcal{R}}(\pi_\beta)$ in β , and Theorem 2.12 of Kosorok (2008), yields the result in part 1. Part 2 follows from part 1 of this theorem, part 1 of Theorem 4.2, and the continuous mapping theorem. Part 3 follows from parts 2 and 3 of Theorem 4.2. The proof of part 4 follows standard arguments. $\square$

Lemma 8.1. *Let $\mathcal{F}$ and $\mathcal{G}$ be function classes with respective envelopes F and G . Let*

$\|F\|_u = (\mathbb{E}|F|^u)^{1/u}$. For any $1 \leq r, s_1, s_2 < \infty$ with $s_1^{-1} + s_2^{-1} = 1$,

1. $J_{[]}(\infty, \mathcal{F} \cdot \mathcal{G}, L_r(P)) \leq 2 (\|F\|_{rs_1} + \|G\|_{rs_2}).$
 $[J_{[]}(\infty, \mathcal{F}, L_{rs_1}(P)) + J_{[]}(\infty, \mathcal{G}, L_{rs_2}(P))]$
2. $J_{[]}(\infty, \mathcal{F} + \mathcal{G}, L_r(P)) \leq 2 [J_{[]}(\infty, \mathcal{F}, L_r(P)) + J_{[]}(\infty, \mathcal{G}, L_r(P))].$
3. *For any*
 $0 \leq r \leq \infty, J_{[]}(\infty, \mathcal{F} \cup \mathcal{G}, L_r(P)) \leq \sqrt{\log 2} (\|F\|_r + \|G\|_r) + J_{[]}(\infty, \mathcal{F}, L_r(P))$
 $+ J_{[]}(\infty, \mathcal{G}, L_r(P))$
4. *If $\mathcal{G}$ is a finite class, $J_{[]}(\infty, \mathcal{G}, L_r(P)) \leq 2\|G\|_r \sqrt{\log |\mathcal{G}|}$, where $|\mathcal{G}|$ denotes the cardinality of $\mathcal{G}$.*

Proof of Lemma 8.1. Proof of part 1: Let $1 \leq r, s_1, s_2 < \infty$ with $s_1^{-1} + s_2^{-1} = 1$ and let $(\mathcal{L}_F, u_F)$ and $(\mathcal{L}_G, u_G)$ be $L_{rs_1}(P)$ and $L_{rs_2}(P)$ ϵ -brackets, respectively. Choose $\mathcal{L}_F = \{f_1, f_2, \dots, u_F\}$ and $\mathcal{L}_G = \{g_1, g_2, \dots, u_G\}$ and consider the bracket for any $f_2 g_2$ defined by $f_1 g_1 \pm (F|u_G - \mathcal{L}_G| + G|u_F - \mathcal{L}_F|)$. Note that $f_1 g_1 + F|u_G - \mathcal{L}_G| + G|u_F - \mathcal{L}_F| - f_2 g_2 = F|u_G - \mathcal{L}_G| + G|u_F - \mathcal{L}_F| - F|g_1 - g_2| - G|f_1 - f_2| \geq 0$, because $f_2 g_2 - f_1 g_1 = f_2 g_2 + f_2 g_1 - f_2 g_1 - f_1 g_1 = F|g_1 - g_2| + G|f_1 - f_2|$. Similarly, $f_2 g_2 + F|u_G - \mathcal{L}_G| + G|u_F - \mathcal{L}_F| - f_1 g_1 \geq 0$. Thus, these brackets hold all $f_2 g_2$ for $f_2 \in (\mathcal{L}_F, u_F)$ and $g_2 \in (\mathcal{L}_G, u_G)$. Now, $\|F|u_G - \mathcal{L}_G| + G|u_F - \mathcal{L}_F|\|_r = \|F\|_{rs_1} \epsilon + \|G\|_{rs_2} \epsilon$ by Minkowski's inequality and Hölder's inequality, and it follows that

$$N_{[]} \left\{ 2\epsilon(\|F\|_{rs1} + \|G\|_{rs2}), \mathcal{F} \cup \mathcal{G}, L_r(P) \right\} \leq N_{[]} \left\{ \epsilon, \mathcal{F}, L_{rs1}(P) \right\} N_{[]} \left\{ \epsilon, \mathcal{G}, L_{rs2}(P) \right\}.$$

Next we note that

$$N_{[]} \left\{ \epsilon, \mathcal{F} \cup \mathcal{G}, L_r(P) \right\} \leq N_{[]} \left\{ \frac{\epsilon}{2(\|F\|_{rs1} + \|G\|_{rs2})}, \mathcal{F}, L_{rs1}(P) \right\} N_{[]} \left\{ \frac{\epsilon}{2(\|F\|_{rs1} + \|G\|_{rs2})}, \mathcal{G}, L_{rs2}(P) \right\}$$

and thus

$$\begin{aligned} J_{[]} \left\{ \infty, \mathcal{F} \cup \mathcal{G}, L_r(P) \right\} &\leq \int_0^2 \|F\|_{rs1} \|G\|_{rs2} \sqrt{\log N_{[]} \left\{ \frac{\epsilon}{2(\|F\|_{rs1} + \|G\|_{rs2})}, \mathcal{F}, L_{rs1}(P) \right\}} d\epsilon + \int_0^2 \|F\|_{rs1} \|G\|_{rs2} \sqrt{\log N_{[]} \left\{ \frac{\epsilon}{2(\|F\|_{rs1} + \|G\|_{rs2})}, \mathcal{G}, L_{rs2}(P) \right\}} d\epsilon \\ &\leq 2(\|F\|_{rs1} + \|G\|_{rs2}) \left[J_{[]} \left\{ \infty, \mathcal{F}, L_{rs1}(P) \right\} + J_{[]} \left\{ \infty, \mathcal{G}, L_{rs2}(P) \right\} \right]. \end{aligned}$$

The proof of part 2 follows from Lemma 9.25 part (i) of Kosorok (2008) after a change of variables. Proof of part 3: First note that

$$N_{[]} \left\{ \epsilon, \mathcal{F} \cup \mathcal{G}, L_r(P) \right\} \leq N_{[]} \left\{ \epsilon, \mathcal{F}, L_r(P) \right\} + N_{[]} \left\{ \epsilon, \mathcal{G}, L_r(P) \right\},$$

whence it follows that

$$\begin{aligned} J_{[]} \left\{ \infty, \mathcal{F} \cup \mathcal{G}, L_r(P) \right\} &= \int_0^{2(\|F\|_r + \|G\|_r)} \sqrt{\log N_{[]} \left\{ \epsilon, \mathcal{F} \cup \mathcal{G}, L_r(P) \right\}} d\epsilon \\ &\leq \int_0^{2(\|F\|_r + \|G\|_r)} \sqrt{\log \left[N_{[]} \left\{ \epsilon, \mathcal{F}, L_r(P) \right\} + N_{[]} \left\{ \epsilon, \mathcal{G}, L_r(P) \right\} \right]} d\epsilon \\ &\leq \int_0^{2(\|F\|_r + \|G\|_r)} \sqrt{\log 2 + \log N_{[]} \left\{ \epsilon, \mathcal{F}, L_r(P) \right\} + \log N_{[]} \left\{ \epsilon, \mathcal{G}, L_r(P) \right\}} d\epsilon \\ &\leq \int_0^{2(\|F\|_r + \|G\|_r)} \sqrt{\log 2} d\epsilon + J_{[]} \left\{ \infty, \mathcal{F}, L_r(P) \right\} + J_{[]} \left\{ \infty, \mathcal{G}, L_r(P) \right\}, \end{aligned}$$

where the second inequality uses the fact that $a + b \leq 2ab$ for all $a, b \geq 1$.

Proof of part 4: If $\mathcal{G}$ is finite, then $N_{[]} \left\{ \epsilon, \mathcal{G}, L_r(P) \right\} \leq |\mathcal{G}|$. Thus,

$$\begin{aligned} J_{[]} \left\{ \infty, \mathcal{G}, L_r(P) \right\} &= \int_0^{2\|G\|_r} \sqrt{\log N_{[]} \left\{ \epsilon, \mathcal{G}, L_r(P) \right\}} d\epsilon \\ &\leq \int_0^{2\|G\|_r} \sqrt{\log |\mathcal{G}|} d\epsilon, \end{aligned}$$

which completes the proof.

Lemma 8.2. Define the class of functions

$$\Pi = \left\{ \pi_{\tilde{\beta}}(a; s) = \frac{a_J + \sum_{j=1}^{J-1} a_j \exp(s^T \tilde{\beta}_j)}{1 + \sum_{j=1}^{J-1} \exp(s^T \tilde{\beta}_j)} : \tilde{\beta} = (\tilde{\beta}_1^T, \dots, \tilde{\beta}_{J-1}^T)^T, \tilde{\beta} \in \mathcal{B} \subset \mathbb{R}^{p(J-1)} \right\}$$

for a compact set $\mathcal{B}$ and $2 \leq J < \infty$ where $a = (a_1, \dots, a_J)^T$. Then, there exists a $b_0 < \infty$ such that for any $1 \leq r < \infty$, $J_{\Pi}(\epsilon, \Pi, L_r(P)) \leq b_0 \|S\|_r \sqrt{p(J-1)} \pi$, which is finite whenever $\|S\|_r < \infty$. Furthermore, $\sup \|\tilde{\beta}_1 - \tilde{\beta}_2\| \leq \delta \mathbb{E} \|\pi_{\tilde{\beta}_1}(A; S) - \pi_{\tilde{\beta}_2}(A; S)\| \rightarrow 0$ as $\delta \downarrow 0$.

Proof of Lemma 8.2. For $\tilde{\beta}_1, \tilde{\beta}_2 \in \mathcal{B}$, define $d(\tilde{\beta}_1, \tilde{\beta}_2) = \max_{1 \leq j \leq J-1} \|\tilde{\beta}_{1j} - \tilde{\beta}_{2j}\|$ and $b_0 = \sup_{\tilde{\beta}_1, \tilde{\beta}_2 \in \mathcal{B}} \|\tilde{\beta}_1 - \tilde{\beta}_2\| < \infty$ because $\mathcal{B}$ is compact. By the mean value theorem, for any $\tilde{\beta}_1, \tilde{\beta}_2 \in \mathcal{B}$, there exists a point $\tilde{\beta}_*$ on the line segment between $\tilde{\beta}_1$ and $\tilde{\beta}_2$ such that

$$\pi_{\tilde{\beta}_1}(a; s) - \pi_{\tilde{\beta}_2}(a; s) = \frac{1}{1 + \sum_{j=1}^{J-1} \exp(s^T \tilde{\beta}_{*j})} \left[\sum_{j=1}^{J-1} \{a_j - \pi_{\tilde{\beta}_*}(a; s)\} \exp(s^T \tilde{\beta}_{*j}) s^T (\tilde{\beta}_{1j} - \tilde{\beta}_{2j}) \right],$$

which implies that

$$|\pi_{\tilde{\beta}_1}(a; s) - \pi_{\tilde{\beta}_2}(a; s)| \leq \|s\| d(\tilde{\beta}_1, \tilde{\beta}_2). \quad (9)$$

It follows from equation (9) that assumption 7 holds for this particular class of policies.

Now, $N_{\Pi}(\epsilon, \Pi, L_r(P)) \leq N(\epsilon, \mathcal{B}, d)$ by Theorem 9.23 of Kosorok (2008).

Furthermore, $N(\epsilon, \mathcal{B}, d) \leq \max \left\{ (b_0/\epsilon)^{p(J-1)}, 1 \right\}$, and thus

$$\begin{aligned} J_{\Pi}(\epsilon, \Pi, L_r(P)) &\leq 2 \|S\|_r \int_0^{b_0} \sqrt{p(J-1)} \{\log b_0 + \log(1/\epsilon)\} d\epsilon \\ &\leq 2 \|S\|_r b_0 \sqrt{p(J-1)} \int_0^1 \sqrt{\log(1/\epsilon)} d\epsilon \\ &= 2 \|S\|_r b_0 \sqrt{p(J-1)} \int_0^\infty u^{1/2} \exp(-u) du \\ &= \|S\|_r b_0 \sqrt{p(J-1)} \pi, \end{aligned}$$

which proves the result.

Additional simulation results

Table 8:

Monte Carlo value estimates for offline simulations with $\gamma = 0.9$, $\mathcal{P}$ is the LASSO penalty, and Ω is the identity.

n	T	Linear VL	Polynomial VL	Gaussian VL	GGQ	Observed
25	24	0.123 (0.0773)	0.117 (0.1067)	0.128 (0.0960)	0.025 (0.0330)	-0.005
	36	0.117 (0.0900)	0.120 (0.0933)	0.138 (0.0992)	0.030 (0.0341)	-0.004

n	T	Linear VL	Polynomial VL	Gaussian VL	GGQ	Observed
50	48	0.122 (0.0782)	0.103 (0.1002)	0.141 (0.0878)	0.028 (0.0301)	0.000
	24	0.109 (0.0727)	0.122 (0.0954)	0.153 (0.0692)	0.028 (0.0321)	-0.005
	36	0.137 (0.0782)	0.127 (0.1061)	0.141 (0.0816)	0.024 (0.0285)	0.003
100	48	0.110 (0.0761)	0.127 (0.0860)	0.147 (0.0778)	0.029 (0.0347)	0.000
	24	0.125 (0.0802)	0.129 (0.0854)	0.164 (0.0609)	0.027 (0.0289)	-0.001
	36	0.151 (0.0739)	0.148 (0.0822)	0.131 (0.0897)	0.025 (0.0356)	-0.002
	48	0.131 (0.0726)	0.132 (0.0814)	0.169 (0.0666)	0.030 (0.0325)	-0.001

Table 9:

Monte Carlo value estimates for offline simulations with $\gamma = 0.9$, $\mathcal{P}$ is Euclidean norm, and Ω is the inverse Fisher information.

n	T	Linear VL	Polynomial VL	Gaussian VL	GGQ	Observed
25	48	0.136 (0.0862)	0.118 (0.0925)	0.153 (0.0785)	0.022 (0.0272)	-0.005
	36	0.147 (0.0768)	0.132 (0.1057)	0.124 (0.0909)	0.026 (0.0355)	-0.004
	24	0.128 (0.0897)	0.146 (0.0826)	0.116 (0.1067)	0.020 (0.0317)	0.000
50	48	0.113 (0.0954)	0.129 (0.0964)	0.123 (0.1052)	0.027 (0.0275)	-0.005
	36	0.116 (0.0973)	0.149 (0.0940)	0.152 (0.0798)	0.029 (0.0289)	0.003
	24	0.109 (0.0899)	0.132 (0.0998)	0.124 (0.0932)	0.024 (0.0285)	0.000
100	48	0.167 (0.0652)	0.155 (0.0743)	0.144 (0.0897)	0.025 (0.0291)	-0.002
	36	0.167 (0.0731)	0.153 (0.0988)	0.155 (0.0851)	0.027 (0.0311)	-0.002
	24	0.137 (0.0868)	0.175 (0.0615)	0.148 (0.0978)	0.026 (0.0332)	-0.001

Table 10:

Monte Carlo value estimates for offline simulations with $\gamma = 0.9$, $\mathcal{P}$ is the LASSO penalty, and Ω is the inverse Fisher information.

n	T	Linear VL	Polynomial VL	Gaussian VL	GGQ	Observed
25	48	0.123 (0.0750)	0.123 (0.0912)	0.140 (0.0930)	0.025 (0.0301)	-0.005
	36	0.139 (0.0780)	0.110 (0.1013)	0.138 (0.0813)	0.024 (0.0344)	-0.004
	24	0.135 (0.0690)	0.110 (0.1177)	0.143 (0.0718)	0.023 (0.0282)	0.000
50	48	0.118 (0.0705)	0.124 (0.0994)	0.137 (0.0802)	0.030 (0.0287)	-0.006
	36	0.117 (0.0827)	0.123 (0.0972)	0.121 (0.0804)	0.030 (0.0292)	0.003
	24	0.128 (0.0807)	0.113 (0.1085)	0.137 (0.0921)	0.023 (0.0282)	0.000
100	48	0.131 (0.0563)	0.123 (0.1015)	0.167 (0.0472)	0.029 (0.0295)	-0.001
	36	0.132 (0.0735)	0.148 (0.0851)	0.161 (0.0670)	0.029 (0.0334)	-0.002
	24	0.149 (0.0612)	0.137 (0.1003)	0.156 (0.0687)	0.023 (0.0267)	-0.001

References

- Ali AA, Hossain SM, Hovsepian K, Rahman MM, Plarre K, and Kumar S (2012). mpuff: Automated detection of cigarette smoking puffs from respiration measurements. In Proceedings of the 11th International Conference on Information Processing in Sensor Networks, pp. 269–280. ACM.
- Bergental RM, Garg S, Weinzimer SA, Buckingham BA, Bode BW, Tamborlane WV, and Kaufman FR (2016). Safety of a hybrid closed-loop insulin delivery system in patients with type 1 diabetes. *Journal of the American Medical Association* 316 (13), 1407–1408. [PubMed: 27629148]
- Bexelius C, Löf M, Sandin S, Lagerros YT, Forsum E, and Litton J-E (2010). Measures of physical activity using cell phones: Validation using criterion methods. *Journal of Medical Internet Research* 12 (1), e2. [PubMed: 20118036]
- Chakraborty B and Moodie EE (2013). *Statistical Methods for Dynamic Treatment Regimes*. Springer.
- Cobry E, McFann K, Messer L, Gage V, VanderWel B, Horton L, and Chase HP (2010). Timing of meal insulin boluses to achieve optimal postprandial glycemic control in patients with type 1 diabetes. *Diabetes Technology & Therapeutics* 12 (3), 173–177. [PubMed: 20151766]
- Dai Y-H (2002). Convergence properties of the BFGS algorithm. *SIAM Journal on Optimization* 13 (3), 693–701.
- Doya K (2000). Reinforcement learning in continuous time and space. *Neural computation* 12 (1), 219–245. [PubMed: 10636940]
- Ertefaie A (2014). Constructing dynamic treatment regimes in infinite-horizon settings. arXiv preprint arXiv:1406.0764.
- Free C, Phillips G, Watson L, Galli L, Felix L, Edwards P, Patel V, and Haines A (2013). The effectiveness of mobile-health technologies to improve health care service delivery processes: A systematic review and meta-analysis. *PLoS Med* 10 (1), e1001363. [PubMed: 23458994]
- Haller MJ, Stalvey MS, and Silverstein JH (2004). Predictors of control of diabetes: Monitoring may be the key. *The Journal of Pediatrics* 144 (5), 660–661. [PubMed: 15127007]
- Hastie T, Tibshirani R, and Friedman JH (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (2 ed.). New York: Springer.
- Hernan MA and Robins JM (2010). *Causal Inference*. CRC Boca Raton, FL.
- Klasnja P, Hekler EB, Shiffman S, Boruvka A, Almirall D, Tewari A, and Murphy SA (2015). Microrandomized trials: An experimental design for developing just-in-time adaptive interventions. *Health Psychology* 34 (S), 1220.
- Kober J and Peters J (2012). Reinforcement learning in robotics: A survey In *Reinforcement Learning*, pp. 579–610. Springer.
- Kosorok MR (2008). *Introduction to Empirical Processes and Semiparametric Inference*. New York: Springer.
- Kosorok MR and Moodie EE (2015). *Adaptive Treatment Strategies in Practice: Planning Trials and Analyzing Data for Personalized Medicine*, Volume 21 SIAM.
- Kowalski A (2015). Pathway to artificial pancreas systems revisited: Moving downstream. *Diabetes Care* 38 (6), 1036–1043. [PubMed: 25998296]
- Kumar S, Nilsen WJ, Abernethy A, Atienza A, Patrick K, Pavel M, Riley WT, Shar A, Spring B, Spruijt-Metz D, et al. (2013). Mobile health technology evaluation: The mhealth evidence workshop. *American Journal of Preventive Medicine* 45 (2), 228–236. [PubMed: 23867031]
- Laber EB, Linn KA, and Stefanski LA (2014). Interactive model building for Q-learning. *Biometrika* 101 (4), 831. [PubMed: 25541562]
- Lai TL and Robbins H (1985). Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics* 6 (1), 4–22.
- Levine B-S, Anderson BJ, Butler DA, Antisdel JE, Brackett J, and Laffel LM (2001). Predictors of glycemic control and short-term adverse outcomes in youth with type 1 diabetes. *The Journal of Pediatrics* 139 (2), 197–203. [PubMed: 11487743]
- Liao P, Klasnja P, Tewari A, and Murphy SA (2016). Sample size calculations for micro-randomized trials in mhealth. *Statistics in Medicine* 35 (12), 1944–1971. [PubMed: 26707831]

- Linn KA, Laber EB, and Stefanski LA (2017). Interactive Q-learning for quantiles. *Journal of the American Statistical Association* 112 (518), 638–649. [PubMed: 28890584]
- Long N, Gianola D, Rosa GJ, Weigel KA, Kranis A, and Gonzalez-Recio O (2010). Radial basis function regression methods for predicting quantitative traits using SNP markers. *Genetics Research* 92 (3), 209–225. [PubMed: 20667165]
- Ly TT, Breton MD, Keith-Hynes P, De Salvo D, Clinton P, Benassi K, Mize B, Chernavsky D, Place J, Wilson DM, et al. (2014). Overnight glucose control with an automated, unified safety system in children and adolescents with type 1 diabetes at diabetes camp. *Diabetes Care* 37 (8), 2310–2316. [PubMed: 24879841]
- Ly TT, Roy A, Grosman B, Shin J, Campbell A, Monirabbasi S, Liang B, von Eyben R, Shanmugham S, Clinton P, et al. (2015). Day and night closed-loop control using the integrated medtronic hybrid closed-loop system in type 1 diabetes at diabetes camp. *Diabetes Care* 38 (7), 1205–1211. [PubMed: 26049550]
- Maahs DM, Mayer-Davis E, Bishop FK, Wang L, Mangan M, and McMurray RG (2012). Outpatient assessment of determinants of glucose excursions in adolescents with type 1 diabetes: Proof of concept. *Diabetes Technology & Therapeutics* 14 (8), 658–664. [PubMed: 22853720]
- Maei HR, Szepesvári C, Bhatnagar S, and Sutton RS (2010). Toward off-policy learning control with function approximation. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 719–726.
- Moodie EE, Richardson TS, and Stephens DA (2007). Demystifying optimal dynamic treatment regimes. *Biometrics* 63 (2), 447–455. [PubMed: 17688497]
- Murphy S, Deng Y, Laber E, Maei H, Sutton R, and Witkiewitz K (2016). A batch, off-policy, actor-critic algorithm for optimizing the average reward. *arXiv preprint arXiv:1607.05047*.
- Murphy SA (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B* 65 (2), 331–355.
- Murphy SA (2005). A generalization error for Q-learning. *Journal of Machine Learning Research* 6 (Jul), 1073–1097. [PubMed: 16763665]
- Murphy SA, van der Laan MJ, and Robins JM (2001). Marginal mean models for dynamic regimes. *Journal of the American Statistical Association* 96 (456), 1410–1423. [PubMed: 20019887]
- Puterman ML (2014). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.
- Quinn CC, Shardell MD, Terrin ML, Barr EA, Ballew SH, and Gruber-Baldini AL (2011). Cluster-randomized trial of a mobile phone personalized behavioral intervention for blood glucose control. *Diabetes Care* 34 (9), 1934–1942. [PubMed: 21788632]
- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Robins JM (2004). Optimal structural nested models for optimal sequential decisions. In *Proceedings of the Second Seattle Symposium on Biostatistics*, pp. 189–326. Springer.
- Rubin D (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics* 6 (1), 34–58.
- Schulte PJ, Tsiatis AA, Laber EB, and Davidian M (2014). Q- and A-learning methods for estimating optimal dynamic treatment regimes. *Statistical Science* 29 (4), 640–661. [PubMed: 25620840]
- Steinhubl SR, Muse ED, and Topol EJ (2013). Can mobile health technologies transform health care? *Journal of the American Medical Association* 310 (22), 2395–2396. [PubMed: 24158428]
- Sutton R and Barto A (1998). *Reinforcement Learning: An Introduction*. The MIT Press.
- Tang Y and Kosorok MR (2012). Developing adaptive personalized therapy for cystic fibrosis using reinforcement learning. The University of North Carolina at Chapel Hill Department of Biostatistics Technical Report Series Working paper 30.
- Weinzimer SA, Steil GM, Swan KL, Dziura J, Kurtz N, and Tamborlane WV (2008). Fully automated closed-loop insulin delivery versus semiautomated hybrid control in pediatric patients with type 1 diabetes using an artificial pancreas. *Diabetes Care* 31 (5), 934–939. [PubMed: 18252903]
- Wolever T and Mullan Y (2011). Sugars and fat have different effects on postprandial glucose responses in normal and type 1 diabetic subjects. *Nutrition, Metabolism and Cardiovascular Diseases* 21 (9), 719–725.

- Zhang B, Tsiatis AA, Laber EB, and Davidian M (2012). A robust method for estimating optimal treatment regimes. *Biometrics* 68 (4), 1010–1018. [PubMed: 22550953]
- Zhang B, Tsiatis AA, Laber EB, and Davidian M (2013). Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika* 100 (3), 681–694.
- Zhao Y, Kosorok MR, and Zeng D (2009). Reinforcement learning design for cancer clinical trials. *Statistics in Medicine* 28 (26), 3294–3315. [PubMed: 19750510]
- Ziegler R, Heidtmann B, Hilgard D, Hofer S, Rosenbauer J, and Holl R (2011). Frequency of SMBG correlates with hba1c and acute complications in children and adolescents with type 1 diabetes. *Pediatric Diabetes* 12 (1), 11–17. [PubMed: 20337978]

Table 1:

Monte Carlo value estimates for offline simulations with $\gamma = 0.9$.

n	T	Linear VL	Polynomial VL	Gaussian VL	GGQ	Observed
25	24	0.118 (0.0892)	0.091 (0.0825)	0.110 (0.0979)	0.014 (0.0311)	-0.005
	36	0.108 (0.0914)	0.115 (0.0911)	0.112 (0.0919)	0.029 (0.0280)	-0.004
	48	0.106 (0.0705)	0.071 (0.0974)	0.103 (0.0757)	0.031 (0.0350)	0.000
50	24	0.124 (0.0813)	0.109 (0.1045)	0.118 (0.0879)	0.016 (0.0355)	-0.005
	36	0.126 (0.0818)	0.134 (0.0878)	0.136 (0.0704)	0.027 (0.0276)	0.003
	48	0.101 (0.0732)	0.109 (0.0767)	0.115 (0.0763)	0.020 (0.0245)	0.000
100	24	0.117 (0.0895)	0.135 (0.0973)	0.140 (0.0866)	0.019 (0.0257)	0.011
	36	0.113 (0.0853)	0.105 (0.1033)	0.139 (0.0828)	0.021 (0.0312)	0.012
	48	0.111 (0.0762)	0.143 (0.0853)	0.114 (0.0699)	0.031 (0.0306)	-0.001

Table 2:

Monte Carlo value estimates for simulated T1D cohorts with $\gamma = 0.9$.

<i>n</i>	<i>T</i>	Linear VL	Polynomial VL	Gaussian VL	GGQ	Observed
25	24	-2.716 (1.2015)	-2.335 (0.9818)	-2.018 (1.2011)	-3.870 (0.9225)	-2.316
	36	-2.700 (1.2395)	-2.077 (1.0481)	-1.760 (0.8468)	-3.644 (0.8745)	-2.261
	48	-2.496 (1.1986)	-2.236 (1.1978)	-1.751 (0.9887)	-2.405 (1.1025)	-2.365
50	24	-2.545 (1.1865)	-2.069 (1.0395)	-1.605 (0.8064)	-3.368 (1.0186)	-2.263
	36	-2.644 (1.1719)	-2.004 (0.9074)	-1.778 (0.8496)	-3.099 (0.9722)	-2.336
	48	-2.469 (1.1635)	-2.073 (0.9870)	-2.102 (1.2078)	-2.528 (0.9571)	-2.308
100	24	-2.350 (1.1171)	-2.128 (1.0520)	-1.612 (0.7203)	-3.272 (0.8636)	-2.299
	36	-2.547 (1.1852)	-2.116 (0.8518)	-1.672 (0.8643)	-3.232 (0.7951)	-2.321
	48	-2.401 (1.0643)	-2.204 (1.0400)	-1.494 (0.5413)	-2.820 (0.8442)	-2.351

Table 3:Value estimates for online simulations with $\gamma = 0.9$.

n	T	Linear VL	Polynomial VL	Gaussian VL	GGQ
25	24	0.0053	0.0149	-0.0100	-0.0081
	36	0.0525	0.0665	0.0310	0.0160
	48	0.0649	0.0722	0.0416	0.0493
50	24	0.0164	0.0117	0.0037	0.0058
	36	0.0926	0.0791	0.0666	0.0227
	48	0.1014	0.0894	0.0512	0.0434
100	24	0.0036	-0.0157	0.0200	0.0239
	36	0.0766	0.0626	0.0907	0.0540
	48	0.0728	0.0781	0.0608	0.0818

Table 4:Value estimates for online estimation of simulated T1D cohorts with $\gamma = 0.9$

n	T	Linear VL	Polynomial VL	Gaussian VL	GGQ
25	24	-2.3887	-1.9713	-1.8860	-3.2027
	36	-2.3784	-2.1535	-1.7857	-3.5127
	48	-2.2190	-2.0679	-1.6999	-3.2280
50	24	-2.3405	-2.2313	-1.7761	-2.8976
	36	-2.2829	-2.0922	-1.6016	-3.1589
	48	-2.1587	-1.9669	-1.5948	-2.8729
100	24	-2.3229	-2.2295	-1.9138	-3.0865
	36	-2.2927	-2.1608	-1.9030	-3.3483
	48	-2.2096	-2.0454	-1.8252	-2.9428

Table 5:

Value estimates for online V-learning simulations with universal and patient-specific policies when $\gamma = 0.9$.

<i>n</i>	<i>T</i>	Universal policy	Patient-specific policy
25	24	0.0282	0.1813
	36	0.1025	0.1700
	48	0.0977	0.1944
50	24	0.0164	0.2771
	36	0.0768	0.2617
	48	0.0752	0.3038
100	24	0.0160	0.4230
	36	0.0960	0.2970
	48	0.1140	0.3197

Table 6:

Parametric value estimates for V-learning applied to type 1 diabetes data.

Action space	Basis	$\gamma = 0.7$	$\gamma = 0.8$	$\gamma = 0.9$
Binary	Linear	-6.20	-9.35	-15.99
	Polynomial	-3.91	-9.03	-17.50
	Gaussian	-3.44	-13.09	-25.52
Multiple	Linear	-6.47	-9.92	-0.49
	Polynomial	-2.44	-6.80	-14.48
	Gaussian	-8.45	-3.58	-21.18
Observational policy		-6.77	-11.28	-21.79

Table 7:

Probabilities for each action as recommended by estimated policy for one example patient.

Action	Probability
No action	< 0.0001
Physical activity	< 0.0001
Food intake	< 0.0001
Food and activity	< 0.0001
Insulin	0.7856
Insulin and activity	0.2143
Insulin and food	0.0002
Insulin, food, and activity	< 0.0001