Published in final edited form as:

J Am Stat Assoc. 2018 ; 113(524): 1541–1549. doi:10.1080/01621459.2017.1345743.

Interpretable Dynamic Treatment Regimes

Yichi Zhang^{*},

Department of Biostatistics, Harvard University

Eric B. Laber,

Department of Statistics, NC State University

Marie Davidian, and

Department of Statistics, NC State University

Anastasios A. Tsiatis

Department of Statistics, NC State University

Abstract

Precision medicine is currently a topic of great interest in clinical and intervention science. A key component of precision medicine is that it is evidence-based, i.e., data-driven, and consequently there has been tremendous interest in estimation of precision medicine strategies using observational or randomized study data. One way to formalize precision medicine is through a treatment regime, which is a sequence of decision rules, one per stage of clinical intervention, that map up-to-date patient information to a recommended treatment. An optimal treatment regime is defined as maximizing the mean of some cumulative clinical outcome if applied to a population of interest. It is well-known that even under simple generative models an optimal treatment regime can be a highly nonlinear function of patient information. Consequently, a focal point of recent methodological research has been the development of flexible models for estimating optimal treatment regimes. However, in many settings, estimation of an optimal treatment regime is an exploratory analysis intended to generate new hypotheses for subsequent research and not to directly dictate treatment to new patients. In such settings, an estimated treatment regime that is interpretable in a domain context may be of greater value than an unintelligible treatment regime built using ‘black-box’ estimation methods. We propose an estimator of an optimal treatment regime composed of a sequence of decision rules, each expressible as a list of “if-then” statements that can be presented as either a paragraph or as a simple flowchart that is immediately interpretable to domain experts. The discreteness of these lists precludes smooth, i.e., gradient-based, methods of estimation and leads to non-standard asymptotics. Nevertheless, we provide a computationally efficient estimation algorithm, prove consistency of the proposed estimator, and derive rates of convergence. We illustrate the proposed methods using a series of simulation examples and application to data from a sequential clinical trial on bipolar disorder.

^{*}The authors gratefully acknowledge funding from the National Science Foundation (DMS-1555141, DMS-1557733, DMS-1513579) and the National Institutes of Health (P01 CA142538).

Keywords

Precision medicine; treatment regimes; interpretability; decision lists; tree-based methods; research-practice gap

1 Introduction

Precision medicine is now almost universally recognized as a path to delivering the best possible healthcare (Collins and Varmus, 2015; Ashley, 2015; Jameson and Longo, 2015). Furthermore, technological advancements and investment in big-data infrastructure have made it possible to collect, store, and curate large amounts of patient-level data to inform the practice of precision medicine (Krumholz, 2014). Quantitative researchers have responded with a surge of methodological developments aimed at ‘mathematizing’ precision medicine in the form of treatment regimes, a sequence of decision rules, one per stage of clinical intervention, that map up-to-date patient information to a treatment recommendation; an optimal treatment regime is defined as maximizing the mean of some desirable clinical outcome if applied to a population of interest. It can be shown that even under the simplest generative models the optimal regime is a nonlinear function of patient information (Robins, 2004; Schulte et al., 2014; Laber et al., 2014); consequently, to avoid model misspecification, a recent trend is to apply flexible supervised learning methods to estimate optimal treatment regimes. These flexible methods include direct-search using large-margin classifiers (Zhao et al., 2012, 2015; Kang et al., 2014; Zhao et al., 2015; Xu et al., 2015); Q -learning with non-parametric regression models (Qian and Murphy, 2011; Zhao et al., 2011; Moodie et al., 2013; Zhou and Kosorok, 2016); and tree-based methods (Zhang et al., 2012; Laber and Zhao, 2015; Zhang et al., 2015; Doove et al., 2015). Further testament to the popularity of these methods is that the Journal of the American Statistical Association’s Theory and Methods Invited Paper and the Case Studies and Applications Invited Paper at the 2016 Joint Statistical Meetings will feature non-parametric methods for estimating treatment regimes (Zhou et al., 2015; Xu et al., 2015).

Flexible estimation methods mitigate the risk of model misspecification but potentially at the price of rendering the estimated regime unintelligible. This price is may be too high in settings where the primary role of an estimated optimal regime is to generate new scientific hypotheses or inform future research. For example, in the context of sequential multiple assignment clinical trials (SMARTs, Murphy, 2005; Lei et al., 2012) estimation of an optimal treatment regime is typically included as a secondary, exploratory analysis, as sizing the trial to ensure high-quality estimation of an optimal regime is complex (Laber et al., 2016). Tree-based regimes, like regression or classification trees, offer flexibility while retaining interpretability. Here, we propose a method for estimation of an optimal treatment regime that comprises a sequence of decision rules each of which is represented as a sequence if-then statements mapping logical clauses to treatment recommendations. Decision rules of this form are a special case of tree-based rules, known as decision lists (Rivest, 1987; Marchand and Sokolova, 2005; Letham et al., 2012; Wang and Rudin, 2015; Zhang et al., 2015), that are immediately interpretable in a domain context as they can be expressed in either flow-chart or paragraph form. Thus, regimes of this form are amenable to

critique and examination by clinicians and facilitate collaborative, iterative development of data-driven precision medicine. Furthermore, we shall show that despite the structure imposed by the decision lists, they are sufficiently expressive so as to provide high-quality regimes even under non-linear generative models previously used in the literature to illustrate the value of non-parametric estimation methods.

In addition to the clinical and scientific value of interpretable, list-based regimes, the proposed work provides a number of important methodological contributions. Unlike existing tree-based methods for estimating optimal treatment regimes, the proposed methodology applies to problems with an arbitrary number of treatment stages and treatments per stage. In principle, robust policy-search (Zhang et al., 2013) could be used with CART (Breiman et al., 1984) to estimate a multi-stage, tree-based treatment regime; however, this method relies on inverse probability of treatment weighting which rapidly becomes unstable as the number of treatment stages increases. A second contribution is that we prove that the proposed estimator is consistent for the optimal regime within the class of list-based regimes and derive rates of convergence for the proposed estimator. These theoretical results are non-trivial because the discreteness of the list precludes the use of standard asymptotic approaches; to our knowledge these are first results on convergence rates for decision lists and are therefore of independent interest. A third contribution is the proposed estimation algorithm used to construct the decision lists at each stage. This algorithm reduces computation time of naive recursive-splitting algorithm from $O(n^3)$ to $O(n \log n)$ where n is the number of subjects in the sample. Furthermore we modify the splitting criteria proposed by Zhang et al. (2015) to avoid (asymptotically) becoming stuck in a local mode.

In Section 2, we describe list-based treatment regimes and describe our estimation algorithm. In Section 3, we prove consistency of the proposed estimator and derive rates of convergence. In Section 4, we demonstrate the finite sample performance of the proposed method using simulation experiments. We illustrate the proposed method using data from a clinical trial in Section 5 and make concluding remarks in Section 6.

2 Methodology

2.1 Framework

Consider n *i.i.d.* observations collected from a sequential clinical trial with T stages; the proposed methodology also applies to observational data provided that standard causal assumptions required for Q -learning are satisfied (see Schulte et al., 2014, for a statement of these assumptions). In the assumed setup the observed data are

$\{(S_{it}, A_{it}, Y_{it}) : t = 1, \dots, T\}_{i=1}^n$, which comprise *i.i.d.* trajectories of the form $\{(S_t, A_t, Y_t) : t = 1, \dots, T\}$ where: $S_t \in \mathbb{R}^{p_t}$ is a vector of covariates measured at the beginning of the t -th stage; $A_t \in \mathcal{A}_t$ is the treatment actually received during the t -th stage; and $Y_t \in \mathbb{R}$ is a scalar outcome measured at the end of the t -th stage. Let $m_t = |\mathcal{A}_t|$ denote the number of available treatment options at the t -th stage. The final outcome of interest is the sum of immediate outcomes, $Y = \sum_{t=1}^T Y_t$. We assume that larger values of Y are better. Let X_t denote the

information available to the decision maker at stage t so that $\mathbf{X}_1 = \mathbf{S}_1$ and $\mathbf{X}_t = (\mathbf{X}_{t-1}, \mathbf{A}_t, \mathbf{S}_t)$ for $t > 1$. Let $\mathcal{X}_t \subset \mathbb{R}^{d_t}$ be the support of $\mathbf{X}_t$, where $d_t = \sum_{s=1}^t p_s + 2(t-1)$ is the dimension of $\mathbf{X}_t$.

A treatment regime $\boldsymbol{\pi} = (\pi_1, \dots, \pi_T)$ is a sequence of functions $\pi_t: \mathcal{X}_t \rightarrow \mathcal{A}_t$ so that under $\boldsymbol{\pi}$ a patient presenting with $\mathbf{X}_t = \mathbf{x}_t$ at stage t is recommended treatment $\pi_t(\mathbf{x}_t)$. For simplicity, we assume that all treatment are feasible for all patients; the extension to patient-specific sets of feasible treatments is straightforward (Schulte et al., 2014). For any regime $\boldsymbol{\pi}$, let $\mathbb{E}^{\boldsymbol{\pi}}$ denote expectation with respect to distribution induced by assigning treatments according to $\boldsymbol{\pi}$.

Given a class of regimes Π , an optimal regime satisfies, $\boldsymbol{\pi}^{\text{opt}} \in \Pi$ and $\mathbb{E}^{\boldsymbol{\pi}^{\text{opt}}} Y \geq \mathbb{E}^{\boldsymbol{\pi}} Y$ for all $\boldsymbol{\pi} \in \Pi$. Our goal is to construct an estimator of $\boldsymbol{\pi}^{\text{opt}}$ when Π is the class of list-based regimes. Each decision rule π_t in a list-based regime has the form:

$$\begin{aligned} &\text{If } \mathbf{x}_t \in R_{t1} \text{ then } a_{t1}; \\ &\text{else if } \mathbf{x}_t \in R_{t2} \text{ then } a_{t2}; \\ &\dots \\ &\text{else if } \mathbf{x}_t \in R_{tL_t} \text{ then } a_{tL_t}, \end{aligned} \quad (1)$$

where: each $R_{t\ell}$ is a subset of $\mathcal{X}_t$ with the restriction that $R_{tL_t} = \mathcal{X}_t$; $a_{t\ell} \in \mathcal{A}_t$; $\ell = 1, \dots, L_t$; and

L_t is the length of $\boldsymbol{\pi}_t$. Thus, a compact representation of $\boldsymbol{\pi}_t$ is $\{(R_{t\ell}, a_{t\ell})\}_{\ell=1}^{L_t}$. To increase interpretability, we restrict $R_{t\ell}$ to clauses involving thresholding with at most two covariates, hence $R_{t\ell}$ is an element of

$$\begin{aligned} \mathcal{R}_t = \{ &\mathcal{X}_t, \{ \mathbf{x} \in \mathcal{X}_t: x_{j1} \leq \tau_1 \}, \{ \mathbf{x} \in \mathcal{X}_t: x_{j1} > \tau_1 \}, \{ \mathbf{x} \in \mathcal{X}_t: x_{j1} \leq \tau_1 \text{ and } x_{j2} \leq \tau_2 \}, \\ &\{ \mathbf{x} \in \mathcal{X}_t: x_{j1} \leq \tau_1 \text{ and } x_{j2} > \tau_2 \}, \{ \mathbf{x} \in \mathcal{X}_t: x_{j1} > \tau_1 \text{ and } x_{j2} \leq \tau_2 \}, \\ &\{ \mathbf{x} \in \mathcal{X}_t: x_{j1} > \tau_1 \text{ and } x_{j2} > \tau_2 \}: 1 \leq j_1 < j_2 \leq d_t, \tau_1, \tau_2 \in \mathbb{R} \}, \end{aligned} \quad (2)$$

where j_1, j_2 are indices and τ_1, τ_2 are thresholds. We also impose an upper bound, $L_{\max}$, on list length L_t for all t . Hence, the class of regimes of interest is $\Pi = \bigotimes_{t=1}^T \Pi_t$, where

$$\Pi_t = \left\{ \{(R_{t\ell}, a_{t\ell})\}_{\ell=1}^{L_t} : R_{t\ell} \in \mathcal{R}_t, a_{t\ell} \in \mathcal{A}_t, L_t \leq L_{\max} \right\}.$$

Remark 1—We omit sets of the form $\{ \mathbf{x}_t \in \mathcal{X}_t: x_{j1} \leq \tau_1 \text{ or } x_{j2} \leq \tau_2 \}$ in the definition of $\mathcal{R}_t$ because such sets are expressible in terms of the sets already in $\mathcal{R}_t$. For example, the clause “if $\mathbf{x}_t \in R_{t1}$ then a_{t1} ” with $R_{t1} = \{ \mathbf{x}_t \in \mathcal{X}_t: x_{tj1} \leq \tau_1 \text{ or } x_{tj2} \leq \tau_2 \}$ can be written as “if $\mathbf{x}_t \in R'_{t1}$ then a_{t1} ; else if $\mathbf{x}_t \in R'_{t2}$ then a_{t1} ” with $R'_{t1} = \{ \mathbf{x}_t \in \mathcal{X}_t: x_{tj1} \leq \tau_1 \}$ and $R'_{t2} = \{ \mathbf{x}_t \in \mathcal{X}_t: x_{tj2} \leq \tau_2 \}$. Moreover, the latter form has the benefit of avoiding the measurement of x_{j2} for subjects satisfying $x_{j1} \leq \tau_1$, which may be an important consideration if x_{j2} refers to some biomarker

that is expensive to measure (see Zhang et al., 2015, for discussion of decision lists and measurement cost).

Remark 2—Under certain generative models, distinct sets in $\mathcal{R}_t$ may correspond to the same group of subjects with probability one. For example, if X_{t1} takes values in $\{0, 1\}$, the set $\{x \in \mathcal{X}_t: x_1 \leq 0\}$ and the set $\{x \in \mathcal{X}_t: x_1 \leq 0.5\}$ correspond to the same group of subjects.

To address this issue, it is tempting to require the threshold for x_1 to take values in the support of X_{t1} . Nevertheless, such requirement is not sufficient to ensure that different sets in $\mathcal{R}_t$ correspond to different groups of subjects. To see this, suppose $(X_{t1}, X_{t2})^T$ can take three possible values: $(0, 0)^T$, $(1, 0)^T$ and $(1, 1)^T$, e.g., if X_{t1} and X_{t2} are indicators of two symptoms where the second symptom can be present only when the first symptom is present. In this case, the set $\{x \in \mathcal{X}_t: x_1 \leq 0\}$ and the set $\{x \in \mathcal{X}_t: x_1 \leq 0 \text{ and } x_2 \leq 0\}$

correspond to the same group of subjects. Therefore, we allow the thresholds to take arbitrary values. In our theoretical analysis, we quantify dissimilarity of sets in $\mathcal{R}_t$ using a distance that accounts for the distribution of X_t .

To estimate π^{opt} we combine non-parametric Q -learning with policy-search (see Taylor et al., 2015, for a discussion of this idea in the context of single decision point). To develop our ideas, we first provide a high-level schematic for our algorithm, then we describe implementation and modeling details, and finally we discuss a computational insight that improves computation time. A complete description of our estimation algorithm is lengthy and technical and is therefore presented in the Supplemental Materials.

Define $Q_T(x_T, a_T) = \mathbb{E}(Y_T | X_T = x_T, A_T = a_T)$. Then it can be shown that

$\pi_T^{\text{opt}} = \arg \max_{\pi \in \Pi_T} \mathbb{E} Q_T \{X_T, \pi(X_T)\}$. Recursively, for $t = T-1, \dots, 1$ define

$Q_t(x_t, a_t) = \mathbb{E}[Y_t + Q_{t+1} \{X_{t+1}, \pi_{t+1}^{\text{opt}}(X_{t+1})\} | X_t = x_t, A_t = a_t]$ and subsequently it can be shown that $\pi_t^{\text{opt}} = \arg \max_{\pi_t \in \Pi_t} \mathbb{E} Q_t \{X_t, \pi_t(X_t)\}$ (Schulte et al., 2014). For each $t = 1, \dots, T$ let $\mathcal{Q}_t$ denote a postulated class of models for Q_t . Q -learning with policy-search follows directly from the foregoing definitions, a schematic is as follows.

(S1) Construct an estimator of Q_T in $\mathcal{Q}_T$, e.g., one could use penalized least squares

$$\hat{Q}_T = \arg \min_{Q_T \in \mathcal{Q}_T} \sum_{i=1}^n \{Y_{Ti} - Q_T(X_{Ti}, A_{Ti})\}^2 + \mathcal{P}_T(Q_T), \text{ where } \mathcal{P}_T(Q_T) \text{ is a penalty on the complexity of } Q_T. \text{ Define}$$

$$\hat{\pi}_T = \arg \max_{\pi \in \Pi_T} \sum_{i=1}^n \hat{Q}_T \{X_{Ti}, \pi(X_{Ti})\}.$$

(S2) Recursively, for $t = T-1, \dots, 1$ construct an estimator of Q_t in $\mathcal{Q}_t$, say $\hat{Q}_t$, e.g.,

$$\hat{Q}_t = \arg \min_{Q_t \in \mathcal{Q}_t} \sum_{i=1}^n \{Y_{ti} + \hat{Q}_{t+1} \{X_{(t+1)i}, \hat{\pi}_{t+1}(X_{(t+1)i})\} - Q_t(X_{ti}, A_{ti})\}^2 \text{ where}$$

$$+ \mathcal{P}_t(Q_t),$$

$\mathcal{P}_t(Q_t)$ is a penalty on the complexity of Q_t . Define

$$\hat{\pi}_t = \operatorname{argmax}_{\pi_t \in \Pi_t} \sum_{i=1}^n \hat{Q}_t\{X_{ti}, \pi_t(X_{ti})\}.$$

Implementation of the preceding schematic requires a choice of models for the Q -functions, a means of constructing an estimator within this class, and an algorithm for computing $\operatorname{argmax}_{\pi_t \in \Pi_t} \sum_{i=1}^n \hat{Q}_t\{X_{ti}, \pi_t(X_{ti})\}$. In our implementation, we use kernel ridge regression with an extended Gaussian kernel to construct estimators of the Q -functions and a greedy stepwise algorithm to approximate $\hat{\pi}_t$ from the estimated Q -functions.

2.2 Kernel Ridge Regression

We use kernel ridge regression to estimate the Q -functions. Starting with the last stage, let $K_T(\cdot, \cdot)$ be a symmetric and positive definite function from $\mathbb{R}^{d_T} \times \mathbb{R}^{d_T}$ to $\mathbb{R}$, and let $\mathbb{H}_T$ be the corresponding reproducing kernel Hilbert space (RKHS). In our implementation, we employ an extension of the Gaussian kernel that employs different scaling factors in different variables: $K_T(\mathbf{x}, \mathbf{z}) = \exp\left\{-\sum_{j=1}^{d_T} \gamma_{Tj}(x_j - z_j)^2\right\}$, where $\gamma_T = (\gamma_{T1}, \dots, \gamma_{Td_T})^T$ is a tuning parameter and $\gamma_{Tj} > 0$ for all j . For each $a \in \mathcal{A}_T$, we estimate $Q_T(\cdot, a)$ via penalized least squares

$$\hat{Q}_T(\cdot, a) = \operatorname{argmin}_{f \in \mathbb{H}_T} \frac{1}{n_{Ta}} \sum_{i \in \mathcal{I}_{Ta}} \{Y_{iT} - f(X_{iT})\}^2 + \lambda_T \|f\|_{\mathbb{H}_T}^2,$$

where $\mathcal{I}_{Ta} = \{i: A_{iT} = a\}$, $n_{Ta} = |\mathcal{I}_{Ta}|$, and $\lambda_T > 0$ is a tuning parameter. Let $\mathbf{Y}_{Ta} = (Y_{iT})_{i \in \mathcal{I}_{Ta}}$ and $\mathbf{K}_{Ta} = \{K(X_{iT}, X_{jT})\}_{i, j \in \mathcal{I}_{Ta}}$. By the representer theorem (Kimeldorf and Wahba, 1971), $\hat{Q}_T(\mathbf{x}, a) = \sum_{i \in \mathcal{I}_{Ta}} K_T(\mathbf{x}, X_{iT}) \hat{\beta}_{iT a}$, where $\hat{\beta}_{Ta} = (\hat{\beta}_{iT a})_{i \in \mathcal{I}_{Ta}}$ satisfy $\hat{\beta}_{Ta} = \operatorname{argmin}_{\beta} \|\mathbf{Y}_{Ta} - \mathbf{K}_{Ta} \beta\|^2 + n_{Ta} \lambda_T \beta^T \mathbf{K}_{Ta} \beta$. Define $\hat{\pi}_T = \operatorname{argmax}_{\pi_T \in \Pi_T} \sum_{i=1}^n \hat{Q}_T\{X_{Ti}, \pi_T(X_{Ti})\}$.

Similarly, for each $t < T$ let $\mathbb{H}_t$ be the RKHS induced by the kernel $K_t(\mathbf{x}, \mathbf{z}) = \exp\left\{-\sum_{j=1}^{d_t} \gamma_{tj}(x_j - z_j)^2\right\}$, and $\gamma_t = (\gamma_{t1}, \dots, \gamma_{td_t})^T$ is a tuning parameter. Recursively, for each $t < T$, $a_t \in \mathcal{A}_t$, estimate $Q_t(\cdot, a)$ by

$$\hat{Q}_t(\cdot, a) = \operatorname{argmin}_{f \in \mathbb{H}_t} \frac{1}{n_{ta}} \sum_{i \in \mathcal{I}_{ta}} \left[Y_{it} + \hat{Q}_{t+1}\{X_{i, t+1}, \hat{\pi}_{t+1}(X_{i, t+1})\} - f(X_{it}) \right]^2 + \lambda_t \|f\|_{\mathbb{H}_t}^2,$$

where $\mathcal{J}_{ta} = \{i: A_{it} = a\}$, $n_{ta} = |\mathcal{J}_{ta}|$, $\mathbb{H}_t$ is an RKHS induced by the kernel

$$K_t(x, z) = \exp\left[-\sum_{j=1}^{d_t} \gamma_{tj}(x_j - z_j)^2\right], \text{ and } \lambda_t \gamma_t = (\gamma_{t1}, \dots, \gamma_{td_t})^T \text{ are tuning parameters.}$$

2.3 Construction of Decision Lists

In addition to a method for estimating the Q -functions, the proposed method requires a method for computing $\arg \max_{\pi_t \in \Pi_t} \sum_{i=1}^n \hat{Q}_t\{X_{it}, \pi_t(X_{it})\}$ where Π_t is the space of list-based

decision rules defined previously. Any element in Π_t can be expressed as $\{(R_{t\ell}, a_{t\ell})\}_{\ell=1}^{L_t}$,

however, simultaneous optimization over all regions and treatments is not computationally feasible except in very small problems. Instead, we propose an algorithm that constructs $\hat{\pi}_t$ using a greedy optimization procedure that optimizes one clause in $\hat{\pi}_t$ at a time; unlike many greed algorithms, the proposed method is consistent for the global maximizer. To provide intuition, we describe in detail the first two steps of this greedy algorithm before stating it in more general terms.

2.3.1 Estimation of the first clause—Define $\hat{\pi}_t^Q$ to be map $x_t \mapsto \arg \max_{a_t \in \mathcal{A}_t} \hat{Q}_t(x_t, a_t)$;

thus, $\hat{\pi}_t^Q$ is an optimal estimated decision rule at stage t using non-parametric Q -learning. To estimate the first clause (R_{t1}, a_{t1}) in π_t , we consider the following decision-list parameterized by R and a .

$$\begin{aligned} &\text{If } x_t \in R \text{ then } a; \\ &\text{else if } x_t \in \mathcal{X}_t \text{ then } \hat{\pi}_t^Q(x_t). \end{aligned} \quad (3)$$

If all subjects follow (3), the estimated mean outcome is

$$\frac{1}{n} \sum_{i=1}^n \left[I(X_{it} \in R) \hat{Q}_t(X_{it}, a) + I(X_{it} \notin R) \hat{Q}_t\{X_{it}, \hat{\pi}_t^Q(X_{it})\} \right]. \quad (4)$$

Hence, we can pick the maximizer of (4) as the estimator of (R_{t1}, a_{t1}) . Note that the difference between the estimated mean outcome under $\hat{\pi}_t^Q$ and that under (3) is

$n^{-1} \sum_{i=1}^n I(X_{it} \in R) [\hat{Q}_t\{X_{it}, \hat{\pi}_t^Q(X_{it})\} - \hat{Q}_t(X_{it}, a)]$, which measures the decrease in the estimated mean outcome when some part of $\hat{\pi}_t^Q$ is replaced with an if-then clause. This represents the price paid for interpretability, and by maximizing (4), we minimize this price.

To improve generalization performance, we add a complexity penalty to (3); in addition to encouraging parsimonious lists, we shall see that this penalty also ensures a unique

maximizer. Define $V(R) \in \{0, 1, 3\}$ to be the number of covariates needed to check inclusion in R . We define $\hat{R}_1$ and $\hat{a}_1$ as the maximizers over R and a in

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \left[I(X_{it} \in R) \hat{Q}_t(X_{it}, a) + I(X_{it} \notin R) \hat{Q}_t\{X_{it}, \hat{\pi}_t^Q(X_{it})\} \right] + \varsigma \left\{ \frac{1}{n} \sum_{i=1}^n I(X_{it} \in R) \right\} \\ & + \eta \{2 - V(R)\}, \end{aligned} \quad (5)$$

where $\varsigma, \eta > 0$ are tuning parameters. Thus, the first penalty term rewards regions R with large mass relative to the distribution of X_t whereas the second term rewards regions that involve fewer covariates. Moreover, we impose the constraint $n^{-1} \sum_{i=1}^n I(X_{it} \in R) > 0$ to avoid searching over vacuous clauses.

2.3.2 Estimation of the second clause—To estimate the second clause we consider the following decision list parameterized by R and a

$$\begin{aligned} & \text{If } \mathbf{x} \in \hat{R}_{t1} \text{ then } \hat{a}_{t1}; \\ & \text{else if } \mathbf{x} \in R \text{ then } a; \\ & \text{else if } \mathbf{x} \in \mathcal{X}_t \text{ then } \hat{\pi}_t^Q(\mathbf{x}). \end{aligned} \quad (6)$$

If all the subjects follow the regime (6), the estimated mean outcome is

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n I(X_{it} \in \hat{R}_{t1}) \hat{Q}_t(X_{it}, \hat{a}_{t1}) + \frac{1}{n} \sum_{i=1}^n I(X_{it} \notin \hat{R}_{t1}, X_{it} \in R) \hat{Q}_t(X_{it}, a) \\ & + \frac{1}{n} \sum_{i=1}^n I(X_{it} \notin \hat{R}_{t1}, X_{it} \notin R) \hat{Q}_t\{X_{it}, \hat{\pi}_t^Q(X_{it})\}. \end{aligned} \quad (7)$$

Note that the first term in (7) can be dropped during the optimization as it is independent of R and a . As in (5), we maximize the penalized criterion

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n I(X_{it} \notin \hat{R}_{t1}, X_{it} \in R) \hat{Q}_t(X_{it}, a) + \frac{1}{n} \sum_{i=1}^n I(X_{it} \notin \hat{R}_{t1}, X_{it} \notin R) \hat{Q}_t\{X_{it}, \hat{\pi}_t^Q(X_{it})\} \\ & + \varsigma \left\{ \frac{1}{n} \sum_{i=1}^n I(X_{it} \notin \hat{R}_{t1}, X_{it} \in R) \right\} + \eta \{2 - V(R)\} \end{aligned} \quad (8)$$

with respect to $R \in \mathcal{R}_b$, $a \in \mathcal{A}_t$ and subject to the constraint

$n^{-1} \sum_{i=1}^n I(X_{it} \notin \hat{R}_{t1}, X_{it} \in R) > 0$. We continue this procedure until either every subject gets a recommended treatment, namely $R_{t\ell} = \mathcal{X}_t$ for some ℓ or the maximum length is reached, ℓ

$= L_{\max}$. If the maximum list length is reached, we set $R_{tL_{\max}} = \mathcal{X}_t$ to ensure that the regime applies to every subject and choose $\hat{a}_{tL_{\max}}$ be the estimated best single treatment for all remaining subjects.

2.3.3 Estimation of all clauses—An algorithmic description of the proposed algorithm is given below. Additional computational details, including the time complexity, are given in the next section.

Step 1 Initialize $\ell = 1$.

Step 2 If $\ell < L_{\max}$, compute

$$\begin{aligned} (\hat{R}_{t\ell}, \hat{a}_{t\ell}) = \arg \max_{R \in \mathcal{R}_t, a \in \mathcal{A}_t} \frac{1}{n} \sum_{i=1}^n & \left[I(X_{it} \in \hat{G}_{t\ell}, X_{it} \in R) \hat{Q}_t(X_{it}, a) \right. \\ & + I(X_{it} \in \hat{G}_{t\ell}, X_{it} \notin R) \hat{Q}_t(X_{it}, \hat{\pi}_t^Q(X_{it})) \left. \right] + \varsigma \left\{ \frac{1}{n} \sum_{i=1}^n I(X_{it} \in \hat{G}_{t\ell}, X_{it} \in R) \right\} \\ & + \eta \{2 - V(R)\} \end{aligned} \quad (9)$$

subject to $n^{-1} I(X_{it} \in \hat{G}_{t\ell}, X_{it} \in R) > 0$, where $\hat{G}_{t-1} = \mathcal{X}_t$,

$\hat{G}_{t\ell} = \mathcal{X}_t \setminus \left(\bigcup_{k < \ell} \hat{R}_{tk} \right)$ for $\ell \geq 2$, and $V(R) \in \{0, 1, 2\}$ is the number of variables used to define R . It is easy to verify that the objective function above reduces to (5) when $\ell = 1$ and to (8) when $\ell = 2$. If $\ell = L_{\max}$, set

$$\left(\hat{R}_{t\ell}, \hat{a}_{t\ell} \right) = \arg \max_{R \in \mathcal{R}_t, a \in \mathcal{A}_t} \frac{1}{n} \sum_{i=1}^n I(X_{it} \in \hat{G}_{t\ell}) \hat{Q}_t(X_{it}, a) + \eta \{2 - V(R)\}. \quad (10)$$

The solution of (10) must satisfy $V(R) = 0$ and hence $\hat{R}_{t\ell} = \mathcal{X}_t$. Consequently the last clause does apply to all the rest subjects.

Step 3 If $\hat{R}_{t\ell} = \mathcal{X}_t$ then go to Step 4; otherwise, increase ℓ by 1 and repeat Steps 2 and 3.

Step 4 Output $\hat{\pi}_t = \{(\hat{R}_{tk}, \hat{a}_{tk})\}_{k=1}^{\ell}$.

2.3.4 Implementation details and time complexity—Computation of $(\hat{R}_{t\ell}, \hat{a}_{t\ell})$ in (9) requires special attention because the objective function is non-differentiable and non-convex. We first argue that brute-force search can be used to obtain $(\hat{R}_{t\ell}, \hat{a}_{t\ell})$. Although $\mathcal{R}_t$ contains infinitely many elements, because the objective function in (9) is piecewise linear, for each covariate it suffices to consider n thresholds located at the order statistics of that

covariate. Hence, the number of thresholds to enumerate is of order n^2 . In addition, there are $d_t(d_t + 1)/2$ choices for variables in $R_{t\ell}$, m_t choices for $a_{t\ell}$ and each evaluation of (9) takes $O(n)$ operations. Therefore, the time complexity for finding $(\hat{R}_{t\ell}, \hat{a}_{t\ell})$ via brute-force search is $O(n^3 d_t^2 m_t)$. Unfortunately, the factor n^3 is overwhelming even when the sample size n is moderate.

Instead of brute-force search, we propose a novel algorithm to compute $(\hat{R}_{t\ell}, \hat{a}_{t\ell})$, that substantially reduces the time complexity. Note that the n^3 factor is due to the enumeration of thresholds and the evaluation of the objective function in (9). By reorganizing the enumeration and evaluation, the proposed algorithm reduces the n^3 factor to $n \log n$. Thus, with this implementation, the proposed algorithm can be applied to large datasets; this is appealing in an era of ‘big-data’ where large data-bases are being mined to generate hypotheses about precision medicine.

Proposition 1: For each t and ℓ the estimator $(\hat{R}_{t\ell}, \hat{a}_{t\ell})$ in (9) can be computed within $O(n \log n d_t^2 m_t)$ operations.

The proof of this result is constructive but technical so we provide a sketch of the main idea here and relegate the remaining details to the Supplemental Materials. Suppose R involves only one covariate: $R = \{x : x_j \leq \tau\}$. For fixed t, j and a , we observe that, up to a constant independent of τ , the objective function in (9) is of the form

$F(\tau) = n^{-1} \sum_{i=1}^n I(X_{ijt} \leq \tau) U_i + I(X_{ijt} > \tau) V_i$, where U_i and V_i are constants. As discussed previously, we need only to compute $F(\tau)$ for τ equal to observed covariate values, X_{ijt} . Let $i_1 < \dots < i_n$ be a permutation of $1, \dots, n$ such that $X_{i_1 jt} \leq \dots \leq X_{i_n jt}$. Then, it can be shown

that $F(X_{i_s jt}) = F(X_{i_{s-1} jt}) + U_{i_s} - V_{i_s}$, $s = 2, \dots, n$. Hence, one can enumerate all possible values for τ and evaluate $F(\tau)$ in $O(n)$ time, in contrast to $O(n^2)$ time for brute-force search. A similar recursive relationship can be established if R is of the form $\{x : x_j > \tau\}$. When R involves two covariates, we combine this sorting technique with binary search tree (Cormen et al., 2009), which enables us to find the thresholds in $O(n \log n)$ time.

Remark 3: The proposed algorithm differs from that in Zhang et al. (2015) in two important ways. First, the two algorithms maximize different objective functions. In Zhang et al. (2015), regime (3) is replaced by “if $x \in R$ then a ; else if $x \in \mathcal{X}_t$ then a' ”, where R , a and a' are obtained by maximizing the estimated mean outcome under such a regime. However, this criterion fails to account for subsequent splits in the decision lists and can thereby get stuck in a local mode. In contrast, the proposed algorithm approximates the remaining list with the estimated optimal regime using non-parametric Q -learning. To illustrate the difference between the two objective functions, consider a scenario with $T = 1$ stage, a single covariate $S_1 \sim \text{Uniform}(-2, 2)$ and suppose that $\hat{Q}_1(x, a) = Q_1(x, a) = ax(x - 1)$, $a \in \{-1, 1\}$. Assume ζ and η are small but positive. Then the solution of (9) is $\hat{R}_{11} = \{x : x \leq \tau\}$ and $\hat{a}_{11} = 1$ with $\tau \approx$

0. Nevertheless, if the term $\hat{\pi}_t^Q(X_{it})$ were replaced by a fixed treatment $a' = a$, the solution would be $\hat{R}_{11} = \mathcal{X}_1$ and $\hat{a}_{11} = 1$, leading to a suboptimal regime. A second difference between the proposed algorithm and the one proposed in Zhang et al. (2015) is that the latter requires a pre-specified set of candidate thresholds for each predictor, and its time complexity is the same as brute-force search if we use all the unique values as candidate thresholds.

3 Theoretical Results

For each $\ell = 1, 2, \dots$, define the population analogs of (9) and (10) as follows,

$$(R_{t\ell}^*, a_{t\ell}^*) = \operatorname{argmax}_{R \in \mathcal{R}_t, a \in \mathcal{A}_t} \Psi_{t\ell}(R, a), \text{ where}$$

$$\begin{aligned} \Psi_{t\ell}(R, a) = & E[I(X_t \in G_{t\ell}^*, X_t \in R)Q(X_t, a) + I(X_t \in G_{t\ell}^*, X_t \notin R)Q\{X_t, \pi_t^Q(X_t)\}] \\ & + \varsigma \Pr(X_t \in G_{t\ell}^*, X_t \in R) + \eta\{2 - V(R)\}, \end{aligned} \quad (11)$$

and $G_{t\ell}^* = \mathcal{X}_t$ if $\ell = 1$ and $G_{t\ell}^* = \mathcal{X}_t \setminus (\cup_{k < \ell} R_{tk}^*)$ otherwise, until either $R_{t\ell}^* = \mathcal{X}_t$ or $\ell = L_{\max}$. In the latter case, instead of (11) we define

$$\Psi_{t\ell}(R, a) = E\{I(X_t \in G_{t\ell}^*)Q(X_t, a)\} + \eta\{2 - V(R)\}. \quad (12)$$

Let $L_t^* = \min\{\ell : R_{t\ell}^* = \mathcal{X}_t\}$ and $\pi_t^* = \{(R_{t\ell}^*, a_{t\ell}^*)\}_{\ell=1}^{L_t^*}$. In (11) and (12), the Q -functions are defined as $Q_T(\mathbf{x}, a) = E(Y_T | \mathbf{X}_T = \mathbf{x}, A_T = a)$, $Q_t(\mathbf{x}, a) = E[Y_t + Q_{t+1}\{X_{t+1}, \pi_{t+1}^*(X_{t+1})\} | \mathbf{X}_t = \mathbf{x}, A_t = a]$ for $t = T-1, \dots, 1$. Furthermore, let $\pi_t^Q(\mathbf{x}) = \operatorname{argmax}_{a \in \mathcal{A}_t} Q_t(\mathbf{x}, a)$ for all t .

We assume that all the covariates and outcomes are bounded. This is a common assumption in the context of nonparametric regression; the extension to include unbounded covariates is possible but at the expense of additional complexity.

Assumption 1

There exists $b > 0$ such that $\|\mathbf{X}_t\|_\infty \leq b$ and $|Y_t| \leq b$ with probability one for all $t = 1, \dots, T$.

We also assume positivity (Robins, 2004), which ensures that $Q_t(\mathbf{x}, a)$ is well-defined for all $a \in \mathcal{A}_t$.

Assumption 2

For each t and $a \in \mathcal{A}_t$, $\Pr(A_t = a | \mathbf{X}_t) \geq \varpi$ almost surely for some positive constant ϖ .

A crucial intermediate step in deriving the asymptotic behavior of $\hat{\pi}_t$'s is establishing convergence of $\hat{Q}_t$ to Q_b to facilitate this step we require a certain degree of smoothness in Q_t . A common means of imposing smoothness is to assume differentiability (see, e.g., Stone, 1982). However, the non-differentiable maximization operator that is implicit in the definition of the Q -functions forces us to consider a weaker notion of smoothness. Denote $B_t = [-b, b]^{d_t} \subset \mathbb{R}^{d_t}$. For any function $f: B_t \rightarrow \mathbb{R}$, define the r -th difference $\Delta_h^r(f)$ by $\Delta_h^r(f)(x) = \sum_{i=0}^r \binom{r}{i} (-1)^{r-i} f(x + ih)$ if $x \in B_{t,r,h}$ and 0 otherwise, where r is a positive integer, $h = (h_1, \dots, h_{d_t})^T$, $h_j \geq 0$ for all $j = 1, \dots, d_t$ and $B_{t,r,h} = \{x \in B_t: x + ih \in B_t \text{ for all } i \in \{0, \dots, r\}\}$. Define the r -th modulus of smoothness of f by $\omega_r(f, s) = \sup_{\|h\|_2 \leq s} \sup_{x \in B_t} |\Delta_h^r(f)(x)|$.

The definition above is similar to Eberts and Steinwart (2013, Definition 2.1), but replaces the L_p norm with the supremum norm. This modification allows us to drop the requirement that X_t have a density with respect to Lebesgue measure. Thus, our analysis applies when X_t contains discrete covariates.

The concept of modulus of smoothness generalizes the concept of differentiability. To see this, consider an example where $d = 1$. We observe that $\lim_{h \rightarrow 0} h^{-1} \Delta_h^1(f)(x) = f'(x)$. Suppose $|f'(x)|$ is bounded, then for sufficiently small h , there exists a constant C_f such that $|\Delta_h^1(f)(x)| \leq C_f |h|$. Hence, any continuously differentiable function f defined on a finite interval, satisfies $\omega_1(f, s) = O(s)$, as $s \rightarrow 0$. Generally, if f is r -times continuously differentiable, then $\omega_r(f, s) = O(s^r)$ as $s \rightarrow 0$. In addition, some non-differentiable functions also satisfy this condition. Consider $f(x) = |x|$ and $f(x) = \max(x, 0)$. It is easy to verify that $|\Delta_h^1(f)(x)| \leq |h|$ for any x . Thus $f(x) = |x|$ and $f(x) = \max(x, 0)$ also satisfy $\omega_1(f, s) = O(s)$ though f is not differentiable at 0. We make the following assumption regarding the smoothness of the Q -functions.

Assumption 3

For each t , there exists a positive integer r_t such that $\omega_{r_t}\{Q_t^*(\cdot, a), s\} = O(s^{r_t})$ as $s \rightarrow 0$, for any $a \in \mathcal{A}_t$.

In order to study the probabilistic convergence of $\hat{\pi}_t$ to π_t^* , it is necessary to define an appropriate distance between $\hat{R}_{t\ell}$ and $R_{t\ell}^*$. In view of Remark 2, the distance should incorporate the distribution of X_b , thus, we define $\rho_\lambda(R_1, R_2) = \Pr\{X_t \in (R_1 \ominus R_2)\}$, where $C \ominus D$ denotes the symmetric set difference between sets C and D . It can be verified that $\rho_\lambda(R_1, R_2)$ is non-negative, symmetric, and satisfies the triangle inequality. Note that $\rho_\lambda(R_1, R_2) = 0$ indicates only that R_1 and R_2 refer to the same group of subjects with probability one with respect to X_t but does not imply $R_1 = R_2$. For example, suppose X_t takes values in $\{0, 1\}$, $R_1 = \{x: x = 0\}$ and $R_2 = \{x: x = 0.5\}$. Then $\rho_\lambda(R_1, R_2) = 0$, as expected. Furthermore, the use of ρ_λ helps to avoid the issue of non-unique representations of R when some covariates can

be expressed using others. For example, if $X_{t1} = -X_{t2}$ and both are continuous, then $\rho(\{x : x_1 \leq \tau\}, \{x : x_2 > \tau\}) = 0$. Thus, our goal is to identify an equivalence class of clauses that each describe the same subset of patients. We require the following identifiability assumption on the equivalence class of optimal clauses.

Assumption 4

For each t and ℓ the following inequalities hold:

- i. There exists a constant $\kappa > 0$ such that $\Psi_{t\ell}(R, a_{t\ell}^*) \leq \Psi_{t\ell}(R_{t\ell}^*, a_{t\ell}^*) - \kappa \rho_t^2(R, R_{t\ell}^*)$ as $\rho_t(R, R_{t\ell}^*) \rightarrow 0$;
- ii. For any $\delta > 0$, there exists a constant $\epsilon > 0$ such that $\Psi_{t\ell}(R, a_{t\ell}^*) \leq \Psi_{t\ell}(R_{t\ell}^*, a_{t\ell}^*) - \epsilon$ for all $R \in \mathcal{R}_t$ with $\rho_t(R, R_{t\ell}^*) > \delta$;
- iii. There exists a constant $\varsigma > 0$ such that $\Psi_{t\ell}(R, a) \leq \Psi_{t\ell}(R_{t\ell}^*, a_{t\ell}^*) - \varsigma$ for all $R \in \mathcal{R}_t$ and $a \in \mathcal{A}_t \setminus \{a_{t\ell}^*\}$.

Assumption 4 guarantees the uniqueness of $(R_{t\ell}^*, a_{t\ell}^*)$ in the sense that if $(\tilde{R}_{t\ell}^*, \tilde{a}_{t\ell}^*)$ is another maximizer of $\Psi_{t\ell}(R, a)$, then $\rho(R_{t\ell}^*, \tilde{R}_{t\ell}^*) = 0$ and $a_{t\ell}^* = \tilde{a}_{t\ell}^*$. Moreover, condition (i) assumes that $\Psi_{t\ell}(R, a_{t\ell}^*)$ behaves like a quadratic function in a neighborhood of $R_{t\ell}^*$. When $\mathbf{X}_t$ has bounded density and $R, R_{t\ell}^*$ use the same covariates, $\Psi_{t\ell}$ can be viewed as a function of the threshold values and condition (i) implies that $\Psi_{t\ell}$ behaves like a quadratic function near the optimal threshold values, which is a common condition in parametric models.

Define the value of a decision rule at time t , say π_t , as $V(\pi_t) = E[Q_t(\mathbf{X}_t, \pi_t(\mathbf{X}_t))]$. Our analysis focuses on how close $\hat{\pi}_t$ is to π_t^* , and how well $\hat{\pi}_t$ performs compared to π_t^* in terms of value. For each t , let $\mathcal{S}_t$ be the indices of the signal variables defining the function Q_t , and let $\mathcal{N}_t = \{1, \dots, d_t\} \setminus \mathcal{S}_t$ be the indices of the noise variables. Write $d_t^{\mathcal{S}} = |\mathcal{S}_t|$, $d_t^{\mathcal{N}} = |\mathcal{N}_t|$ and hence $d_t = d_t^{\mathcal{S}} + d_t^{\mathcal{N}}$. Recall that L_T^* is the length of π_T^* . Define $\phi_T = (2/3)^{L_T^*} (2r_T)/(2r_T + d_T)$, and $\phi_t = (2/3)^{L_t^*} \min\left\{(2r_t)/(2r_t + d_t), \phi_{t+1}\right\}$ for $t = T-1, \dots, 1$. Define $\phi_t^{\mathcal{S}}$ in the same way but with d_t replaced by $d_t^{\mathcal{S}}$. As $\gamma_t, \lambda_t, \zeta_t$ and η_t may depend on n , we may write $\gamma_{n,t}, \lambda_{n,t}, \zeta_n$ and η_n to emphasize such dependence. The following theorem establishes finite sample bounds. A proof is given in the Supplemental Materials.

Theorem 1

For each t , assume $\max_j \gamma_{n,t,j} = O\left(n^{2/(2r_t + d_t)}\right)$, $\lambda_t = O(n^{-1})$, $\sup_n \zeta_n < \infty$, and $\sup_n \eta_n < \infty$. Under Assumptions 1–4, for any $\xi > 0$,

$$\Pr\{\hat{\pi}_t(X_t) \neq \pi_t^*(X_t)\} \leq c_1 n^{-\phi_t + \xi}, \Pr\{V_t(\pi_t^*) - V_t(\hat{\pi}_t) \geq c_2 n^{-\phi_t + \xi} + c_3 n^{-1/2\tau}\} \leq e^{-\tau},$$

where c_i 's are constants independent of n and τ .

Moreover, if $\max_{j \in \mathcal{S}_t} \gamma_{n,t,j} = O\left\{n^{2/(2r_t + d_t)}\right\}$ and $\max_{j \in \mathcal{N}_t} \gamma_{n,t,j} = O(1)$, then the inequalities above holds with ϕ_t replaced by $\phi_t^{\mathcal{S}}$.

The minimax convergence rate for a nonparametric regression estimator of an r_T -times continuously differentiable function is $O\left\{n^{-(2r_T)/(2r_T + d_T)}\right\}$ (Stone, 1982). By extending the technique in Eberts and Steinwart (2013), we show in the Supplemental Material that the estimated Q -function $\hat{Q}_T$ converges to its true value Q_T at a nearly optimal rate

$O\left\{n^{-(2r_T)/(2r_T + d_T) + \xi'}\right\}$, where $\xi' > 0$ can be arbitrarily small. The construction of $\hat{\pi}_T$ involves estimating L_T^* pairs of parameters (R_{it}, a_{it}) , one pair for each if-then clause. The estimation of each pair (R_{it}, a_{it}) reduces the convergence rate by an additional factor of $2/3$. The underlying idea for this phenomenon is analogous to the problems analyzed in Kim and Pollard (1990). In earlier stages, the estimation of Q -functions is further complicated by the fact that $Q_{t+1}\{X_{t+1}, \pi_{t+1}(X_{t+1})\}$ is not observed but estimated via $\hat{Q}_{t+1}\{X_{t+1}, \hat{\pi}_{t+1}(X_{t+1})\}$.

When all the covariates are discrete, it can be shown that the convergence rate of the estimated regime $\hat{\pi}_t$ does not inherit the slow convergence rate from the underlying nonparametric regressions. The following result is proved in the Supplemental Material.

Theorem 2

For each t , assume $\max_{j \in \mathcal{S}_t} \gamma_{n,t,j} = O\left\{n^{2/(2r_t + d_t)}\right\}$, $\lambda_T = O(n^{-1})$, $\sup_n \zeta_n < \infty$ and $\sup_n \eta_n < \infty$.

Furthermore, assume that the distribution of X_t is discrete. Namely, for each t there exists a finite set $\tilde{\mathcal{X}}_t$ such that $\Pr(X_t \in \tilde{\mathcal{X}}_t) = 1$. Under Assumptions 1–4,

$$\Pr\{\hat{\pi}_t(X_t) \neq \pi_t^*(X_t)\} \leq c_1 e^{-c_2 n}, \Pr\{V_t(\pi_t^*) - V_t(\hat{\pi}_t) \geq c_3 n^{-1/2\tau}\} \leq e^{-\tau},$$

where c_i 's are constants independent of n and τ .

In both theorems, the convergence rates are independent of ζ_n and η_n . However, the choice of ζ_n and η_n has an impact on the limiting treatment regime π_t^* . In practice, we suggest to tune γ_n and λ_n by minimizing the cross validated mean squared error in the kernel ridge regression, and tune ζ_n and η_n by maximizing the cross validated value of the regime.

4 Simulation Studies

We conducted a series of simulation experiments to examine the empirical performance of the proposed method. Five scenarios were considered. The first four came from Zhao et al. (2015) and the fifth was adapted from Murphy (2003). Scenario I consists of two stages, two treatment options at each stage, and the covariates exhibit nonlinear effects: $S_1 = (S_{1,1}, \dots, S_{1,50})^T$, are independent standard normal random variables; A_1 is Uniform $\{-1, 1\}$; Y_1 is Normal $(\mu_1, 1)$, where $\mu_1 = 0.5S_{1,3}A_1$; S_2 is empty; A_2 is Uniform $\{-1, 1\}$; Y_2 is Normal $(\mu_2, 1)$ where $\mu_2 = \{(S_{1,1}^2 + S_{1,2}^2 - 0.2)(0.5 - S_{1,1}^2 - S_{1,2}^2) + Y_1\}A_2$. Scenario II consists of time-varying covariates. In this scenario: S_1 , A_1 and A_2 were generated in the same way as scenario I; Y_1 is Normal $(\mu_1, 1)$, where $\mu_1 = (1 + 1.5S_{1,3})A_1$; $S_2 = (S_{2,1}, S_{2,2})^T$; $S_{2,1}$ is Bernoulli with success probability $1 - \Phi(1.25S_{1,1}A_1)$; $S_{2,2}$ is Bernoulli with success probability $1 - \Phi(-1.75S_{1,2}A_1)$; and Y_2 is Normal $(\mu_2, 1)$, where $\mu_2 = (0.5 + Y_1 + 0.5A_1 + 0.5S_{2,1} - 0.5S_{2,2})A_2$. In Scenario III, A_1 , A_2 , A_3 are Uniform $\{-1, 1\}^3$; $S_{1,1}$, $S_{1,2}$, $S_{1,3}$ are *i.i.d.* Normal $(45, 15^2)$; S_2 is Normal $(1.5S_{1,1}, 10^2)$; S_3 is Normal $(0.5S_2, 10^2)$; $Y_1 = Y_2 = 0$, and Y_3 is Normal $(\mu_3, 1)$, where $\mu_3 = 20 - 0.6S_{1,1} - 40\{I(A_1 > 0) - I(S_{1,1} > 30)\}^2 - 0.8S_2 - 60\{I(A_2 > 0) - I(S_2 > 40)\}^2 - 1.4S_3 - 40\{I(A_3 > 0) - I(S_3 > 40)\}^2$. Scenario IV is the same as Scenario III except that many noise variables were added. In addition to $S_{1,1}$, $S_{1,2}$ and $S_{1,3}$, we generated $S_{1,4}, \dots, S_{1,50}$ *i.i.d.* from Normal $(45, 15^2)$. Scenario V involves ten stages and multiple treatment options at each stage. See Murphy (2003) for background and motivation for this scenario. For $t \in \{1, \dots, 10\}$, treatments were coded as a pair of values $A_t = (A_{t1}, A_{t2})^T$, generated as follows. First, A_{t1} is drawn from Uniform $\{0, 1\}$. Second, if $A_{t1} = 0$ then A_{t2} is drawn from Uniform $\{0, 1, 2, 3\}$, and otherwise A_{t2} is drawn uniformly from $\{1, 2, 3\}$. Thus, there are $m_t = 7$ treatment candidates at each decision point. In addition, $U_1, \dots, U_{10}$ are *i.i.d.* Normal $(0, 0.01)$; $S_1 = 0.5 + U_1$; $S_t = 0.5 + 0.2S_{t-1} - 0.07A_{t-1,1}A_{t-1,2} - 0.01(1 - A_{t-1,1})A_{t-1,2} + U_t$ for $t = 2$; Y_t is Normal $(\mu_t, 0.64)$, where $\mu_t = 30I(t = 1) - 5U_t - 6\{A_{t1} - I(S_t > 5/9)\}^2 - 1.5A_{t1}(A_t - 2S_t)^2 - 1.5(1 - A_{t1})(A_{t2} - 5.5S_t)^2$ for each t .

In each scenario, we considered sample sizes $n = 100, 200$, and 400 . We generated 1000 data sets for each sample size and estimated the optimal treatment regime using the proposed method. In each stage, we tuned the scaling vector in the Gaussian kernel, γ_b , as well as the amount of penalty, λ_b , via leave-one-out cross validation. The cross validated error was minimized via a Quasi-Newton type algorithm (Kim et al., 2010) with a random starting value. During the construction of decision lists, at each decision point we tuned ζ and η via five-fold cross validation over a pre-specified grid. We picked the combination that led to the largest cross validated outcome.

To form a basis for comparison, we also implemented Q -learning with linear models, nonparametric Q -learning with random forests, backward outcome weighted learning (BOWL), and simultaneous outcome weighted learning (SOWL; Zhao et al., 2015). In Q -learning with linear Q -functions, we fit the working $Q_t(X_t, A_t) = \sum_{a \in \mathcal{A}_t} I(A_t = a)X_t^T \beta_{t,a}$.

Motivated by Qian and Murphy (2011), we imposed an ℓ_1 penalty to reduce overfitting. The Q -functions were estimated by ℓ_1 regularized least squares, implemented in the R package glmnet (Friedman et al., 2010). The covariates were standardized to have mean zero and

variance one before entering the model, and the tuning parameter was selected by five-fold cross validation. Our implementation of non-parametric Q -learning used the R package `randomForest` with default parameters settings (Liaw and Wiener, 2002). We implemented BOWL and SOWL according to the descriptions in Zhao et al. (2015). Linear kernels were used, and the amount of regularization was chosen by five-fold cross validation. Note that BOWL and SOWL assume binary treatment options and thus are not applicable in Scenario V.

We measure the quality of an estimated treatment regime by the mean outcome under that treatment regime; we approximate this mean outcome using an independent test set of size 10^5 . The results are displayed in Table 1. In Scenario I, the second stage Q -function is highly nonlinear, and most methods tended to assign a single treatment to all patients in the second stage, leading to a mean outcome of 6.70. In contrast, the proposed method is able to correctly individualize treatment as the sample size increased and thus produce a higher mean outcome. In Scenario II, both Q -functions at the first and the second stages are linear. Hence, as expected, Q -learning with linear models performs best. Nevertheless, the proposed method and non-parametric Q -learning perform well and shows marked improvement over BOWL and SOWL. In Scenario III, Q -learning with linear models suffers from model misspecification whereas the proposed method and non-parametric Q -learning both perform well. Furthermore, although the Q -functions are complicated, the optimal treatment regime consists of linear functions of covariates. Hence, both BOWL and SOWL perform well in this scenario. Recall that scenario IV is the same as scenario III except for the addition of many noise variables. Thus, the results for scenario IV demonstrate a sensitivity to noise variables in BOWL and SOWL. One possible reason for this is that both BOWL and SOWL utilizes ℓ_2 penalties, which fails to exclude noise variables. In Scenario V, the proposed method outperforms competing methods, especially when the sample size is small. The reason might be due to the nonparametric estimation of Q -functions and the simple form of decision list compared to a random forest, as simpler treatment regimes tends to have better generalizability.

5 Data Analysis

As an illustration of the proposed method, we use data from the Systematic Treatment Enhancement Program for Bipolar Disorder (STEP-BD) to estimate an interpretable treatment regime for treating bipolar disorder (Sachs et al., 2003). We focus on the randomized acute depression (RAD) pathway in STEP-BD, which is a Sequential Multiple Assignment Randomized Trial (SMART) and provides the data needed to build treatment regimes.

One purpose of STEP-BD is to assess the effectiveness of adding antidepressants to mood stabilizers in treating patients with bipolar disorder. Although antidepressants were often assigned to supplement mood stabilizers in practice, it was found that the adjunctive antidepressant medication did not show much improvement over the use of mood stabilizers alone (Sachs et al., 2007). Thus, it is of scientific interest to tailor the use of antidepressants based on individual and time-dependent characteristics.

The RAD pathway in STEP-BD is a randomized trials with two stages. At both stages, patients always received one or more mood stabilizers chosen by their psychiatrists. In addition, they might receive one antidepressant in the form of bupropion or paroxetine. At week 0, patients were randomized to receive bupropion, paroxetine or placebo with probability 0.25, 0.25 and 0.5, respectively. After 6 weeks, patients returned to their psychiatrists for evaluation on response status. In another 6 weeks, responders continued their initial treatments, non-responders who received either bupropion or paroxetine initially were offered an increased dose, and non-responders who received placebo initially were randomized to received bupropion or paroxetine with equal probability. At week 12, patients returned to their psychiatrists for final measurements.

In this clinical trial, the covariate X_1 of a patient consists of his/her age, gender, marital status, education level, employment status, bipolar type, nature of the episode prior to the current depressive episode, summary score for depression (SUM-D) at baseline, and summary score for mood elevation (SUM-ME) at baseline. The treatment A_1 takes three values: bupropion, paroxetine and placebo. The covariate X_2 consists of SUM-D at week 6, SUM-ME at week 6, and indicators for nine different adverse events at week 6. The treatment A_2 is either bupropion or paroxetine for non-responders who received placebo in the first stage. For other patients, A_2 is the same as A_1 . The outcomes are $Y_1 = 0$ and $Y_2 = \text{SUM-D at week 12}$. Note that smaller values of SUM-D and SUM-ME indicates better clinical status. A complete description of these variables is provided in the Supplemental Materials.

We apply the propose method to estimate an interpretable treatment regime. For simplicity, we only include patients with complete baseline and stage 1 information. And we use the last-value-carry-forward strategy if the SUM-D at week 12 is missing. The estimated optimal decision rule at the first stage is:

If SUM-D at week 0 > 8.625 then bupropion;
 else if SUM-D at week 0 ≤ 4.875 and race is not white then paroxetine;
 else placebo.

The estimated regime suggests that the baseline SUM-D is informative in treatment selection. Recall that smaller values of SUM-D indicate lower symptoms. Hence, an interpretation of the estimated regime is: patients with severe depression symptoms should receive bupropion, while non-white patients with minor depression symptoms should receive paroxetine. Although applying an antidepressant medication to all patients did not lead to a better mean outcome relative to not applying antidepressants to any of the patients (Sachs et al., 2007), the estimated regime indicates that personalizing the use of antidepressants based on SUM-D may improve the overall mean outcome. The estimated optimal decision rule at the second stage is:

If SUM-ME at week 6 ≤ 0.875 then bupropion;
 else if SUM-D at week 6 > 8.5 then bupropion;
 else paroxetine.

From this rule it can be seen that patients with large SUM-D or low SUM-ME are assigned to bupropion.

6 Discussion

The current trend in methodological research for estimation of optimal treatment regimes seems to be the development of increasingly flexible models to mitigate risk of model misspecification. This trend is aligned with the notion that an estimated optimal regime will be used to make treatment decisions for future patients. However, in many settings an estimated optimal regime is not used to make treatment decisions but rather is used to generate hypotheses and inform future research. Indeed, our view is that the development of a precision medicine strategy should be the culmination of an iterative process of hypothesis generation and validation. With this perspective, the ability to interpret and estimated optimal regime in a domain context is paramount.

We used list-based regimes to ensure interpretability of the estimated regimes. Our proposed estimation algorithm combines non-parametric Q -learning with policy-search and consistently estimates the optimal regime under mild assumptions. In principle, the proposed estimation framework could be used to estimate interpretable optimal regimes of other forms, e.g., more general tree structures or rule-based systems. Nevertheless, the simplicity of list-based regimes that ensures parsimony and interpretability also appears to have regularizing effect that improves generalization performance.

The recognition that estimated optimal regimes are often not used directly to select treatments for patients but instead are part of an iterative, collaborative process opens many new lines of research beyond estimation of interpretable regimes. These include methods for visualization, models for shared-decision making, models for patient preference and utility construction, and methods for constructing prediction sets for outcome trajectories in multistage decision problems. We are currently pursuing several of these research areas.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

- Ashley EA. 2015; The precision medicine initiative: a new national effort. *Journal of the American Statistical Association*. 313(21):2119–2120.
- Breiman, L, Friedman, JH, Olshen, RA, Stone, CJ. *Classification and Regression Trees*. New York: CRC Press; 1984.
- Collins FS, Varmus H. 2015; A new initiative on precision medicine. *New England Journal of Medicine*. 372(9):793–795. [PubMed: 25635347]
- Cormen, TH, Leiserson, CE, Rivest, RL, Stein, C. *Introduction to algorithms*. 3. MIT Press and McGraw-Hill; 2009.
- Doove L, Dusseldorp E, Van Deun K, Van Mechelen I. 2015A novel method for estimating optimal tree-based treatment regimes in randomized clinical trials. Technical report.
- Eberts M, Steinwart I. 2013; Optimal regression rates for SVMs using Gaussian kernels. *Electronic Journal of Statistics*. 7:1–42.

- Friedman J, Hastie T, Tibshirani R. 2010; Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*. 33(1):1–22. [PubMed: 20808728]
- Jameson JL, Longo DL. 2015; Precision medicine personalized, problematic, and promising. *New England Journal of Medicine*. 372(23):2229–2234. [PubMed: 26014593]
- Kang C, Janes H, Huang Y. 2014; Combining biomarkers to optimize patient treatment recommendations. *Biometrics*. 70(3):695–707. [PubMed: 24889663]
- Kim D, Sra S, Dhillon IS. 2010; Tackling box-constrained optimization via a new projected quasi-newton approach. *SIAM Journal on Scientific Computing*. 32(6):3548–3563.
- Kim J, Pollard D. 1990; Cube root asymptotics. *The Annals of Statistics*. 18(1):191–219.
- Kimeldorf G, Wahba G. 1971; Some results on Tchebycheffian spline functions. *Journal of Mathematical Analysis and Applications*. 33(1):82–95.
- Krumholz HM. 2014; Big data and new knowledge in medicine: the thinking, training, and tools needed for a learning health system. *Health Affairs*. 33(7):1163–1170. [PubMed: 25006142]
- Laber E, Zhao Y, Regh T, Davidian M, Tsiatis AA, Stanford JB, Zeng D, Kosorok MR. 2016; Sizing a phase ii trial to find a nearly optimal personalized treatment strategy. *Statistics in Medicine*.
- Laber EB, Linn KA, Stefanski LA. 2014; Interactive model building for Q-learning. *Biometrika*. 101(4):831–847. [PubMed: 25541562]
- Laber EB, Zhao YQ. 2015; Tree-based methods for individualized treatment regimes. *Biometrika*. 102(3):501–514. [PubMed: 26893526]
- Lei H, Nahum-Shani I, Lynch K, Oslin D, Murphy SA. 2012; A “SMART” design for building individualized treatment sequences. *Annual Review of Clinical Psychology*. 8(1):21–48.
- Letham, B, Rudin, C, McCormick, TH, Madigan, D. Technical Report TR609. Department of Statistics, University of Washington; 2012. Building interpretable classifiers with rules using Bayesian analysis.
- Liaw A, Wiener M. 2002; Classification and regression by randomForest. *R news*. 2(3):18–22.
- Marchand M, Sokolova M. 2005; Learning with decision lists of data-dependent features. *Journal of Machine Learning Research*. 6:427–451.
- Moodie EEM, Dean N, Sun YR. 2013; Q-learning: Flexible learning about useful utilities. *Statistics in Biosciences*. 6:1–21.
- Murphy SA. 2003; Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B*. 65(2):331–355.
- Murphy SA. 2005; An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine*. 24(10):1455–1481. [PubMed: 15586395]
- Qian M, Murphy SA. 2011; Performance guarantees for individualized treatment rules. *Annals of Statistics*. 39(2):1180–1210. [PubMed: 21666835]
- Rivest RL. 1987; Learning decision lists. *Machine Learning*. 2(3):229–246.
- Robins, JM. Optimal structural nested models for optimal sequential decisions. In: Lin, DY, Heagerty, PJ, editors *Proceedings of the Second Seattle Symposium in Biostatistics*. Vol. 179. New York: Springer; 2004. 189–326. *Lecture Notes in Statistics*
- Sachs GS, Nierenberg AA, Calabrese JR, Marangell LB, Wisniewski SR, Gyulai L, Friedman ES, Bowden CL, Fossey MD, Ostacher MJ, Ketter TA, Patel J, Hauser P, Rapport D, Martinez JM, Allen MH, Miklowitz DJ, Otto MW, Dennehy EB, Thase ME. 2007; Effectiveness of adjunctive antidepressant treatment for bipolar depression. *New England Journal of Medicine*. 356(17):1711–1722. [PubMed: 17392295]
- Sachs GS, Thase ME, Otto MW, Bauer M, Miklowitz D, Wisniewski SR, Lavori P, Lebowitz B, Rudorfer M, Frank E, Nierenberg AA, Fava M, Bowden C, Ketter T, Marangell L, Calabrese J, Kupfer D, Rosenbaum JF. 2003; Rationale, design, and methods of the systematic treatment enhancement program for bipolar disorder (STEP-BD). *Biological Psychiatry*. 53(11):1028–1042. [PubMed: 12788248]
- Schulte PJ, Tsiatis AA, Laber EB, Davidian M. 2014; Q- and A-learning methods for estimating optimal dynamic treatment regimes. *Statistical Science*. 29(4):640–661. [PubMed: 25620840]
- Stone CJ. 1982; Optimal global rates of convergence for nonparametric regression. *The Annals of Statistics*. 10(4):1040–1053.

- Taylor JMG, Cheng W, Foster JC. 2015; Reader reaction to “a robust method for estimating optimal treatment regimes” by Zhang et al. (2012). *Biometrics*. 71(1):267–273.
- Wang F, Rudin C. 2015 Falling rule lists. *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*. :1013–1022.
- Xu Y, Müller P, Wahed AS, Thall PF. 2015 Bayesian nonparametric estimation for dynamic treatment regimes with sequential transition times. *Journal of the American Statistical Association*.
- Xu Y, Yu M, Zhao YQ, Li Q, Wang S, Shao J. 2015; Regularized outcome weighted subgroup identification for differential treatment effects. *Biometrics*. 71(3):645–653. [PubMed: 25962845]
- Zhang B, Tsiatis AA, Davidian M, Zhang M, Laber E. 2012; Estimating optimal treatment regimes from a classification perspective. *Stat*. 1(1):103–114. [PubMed: 23645940]
- Zhang B, Tsiatis AA, Laber EB, Davidian M. 2013; Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika*. 100(3):681–694.
- Zhang Y, Laber EB, Tsiatis A, Davidian M. 2015; Using decision lists to construct interpretable and parsimonious treatment regimes. *Biometrics*. 71(4):895–904. [PubMed: 26193819]
- Zhao Y, Zeng D, Laber EB, Kosorok MR. 2015; New statistical learning methods for estimating optimal dynamic treatment regimes. *Journal of the American Statistical Association*. 110(510): 583–598. [PubMed: 26236062]
- Zhao Y, Zeng D, Rush AJ, Kosorok MR. 2012; Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*. 107(499):1106–1118. [PubMed: 23630406]
- Zhao Y, Zeng D, Socinski MA, Kosorok MR. 2011; Reinforcement learning strategies for clinical trials in nonsmall cell lung cancer. *Biometrics*. 67(4):1422–1433. [PubMed: 21385164]
- Zhao YQ, Zeng D, Laber EB, Song R, Yuan M, Kosorok MR. 2015; Doubly robust learning for estimating individualized treatment with censored data. *Biometrika*. 102(1):151–168. [PubMed: 25937641]
- Zhou X, Kosorok MR. 2016 Nearest neighbor rules for optimal treatment regimes.
- Zhou X, Mayer-Hamblett N, Khan U, Kosorok MR. 2015 Residual weighted learning for estimating individualized treatment rules. *Journal of the American Statistical Association*.

Table 1

Simulation results. Given a scenario and a sample size, each method constructed 1000 treatment regimes, one per each simulated dataset. The number in each cell is the outcome under the estimated treatment regime, averaged over 1000 replications, with standard deviation in parentheses. In the header, n is the sample size, DL refers to the proposed decision list based approach, Q -lasso refers to the Q -learning approach with linear model and lasso penalty, Q -RF refers to the Q -learning approach using random forest

Scenario	n	DL	Q -lasso	Q -RF	BOWL	SOWL
I	100	6.63 (0.24)	6.55 (0.58)	6.70 (0.05)	6.70 (0.05)	6.70 (0.05)
I	200	6.73 (0.24)	6.64 (0.33)	6.70 (0.05)	6.70 (0.05)	6.70 (0.05)
I	400	6.94 (0.16)	6.66 (0.26)	6.70 (0.05)	6.70 (0.05)	6.70 (0.05)
II	100	3.66 (0.10)	3.68 (0.08)	3.41 (0.17)	3.15 (0.05)	2.77 (0.52)
II	200	3.71 (0.04)	3.73 (0.04)	3.62 (0.12)	3.22 (0.08)	2.84 (0.33)
II	400	3.73 (0.03)	3.75 (0.02)	3.71 (0.04)	3.37 (0.14)	2.91 (0.28)
III	100	14.49 (2.77)	5.42 (4.54)	12.94 (2.07)	10.65 (2.40)	10.27 (2.33)
III	200	17.42 (1.42)	7.88 (1.63)	15.79 (1.59)	13.09 (2.20)	12.98 (1.88)
III	400	18.60 (0.71)	8.41 (0.65)	18.02 (0.73)	15.33 (1.56)	16.22 (1.58)
IV	100	13.38 (3.14)	4.54 (5.17)	11.47 (2.31)	6.72 (1.71)	6.04 (2.18)
IV	200	17.33 (1.87)	7.69 (2.33)	14.82 (1.75)	8.90 (1.13)	8.34 (1.99)
IV	400	18.84 (0.70)	8.61 (0.96)	17.04 (1.02)	10.75 (0.68)	9.38 (2.30)
V	100	23.68 (1.09)	12.97 (3.40)	17.83 (1.63)	–	–
V	200	25.94 (0.51)	13.80 (2.57)	21.60 (1.28)	–	–
V	400	26.80 (0.29)	16.65 (1.71)	24.73 (0.65)	–	–