Published in final edited form as:

J Am Stat Assoc. 2016 ; 111(514): 875–887. doi:10.1080/01621459.2015.1066682.

Active Clinical Trials for Personalized Medicine

Stanislav Minsker, Ying-Qi Zhao, and Guang Cheng

Abstract

Individualized treatment rules (ITRs) tailor treatments according to individual patient characteristics. They can significantly improve patient care and are thus becoming increasingly popular. The data collected during randomized clinical trials are often used to estimate the optimal ITRs. However, these trials are generally expensive to run, and, moreover, they are not designed to efficiently estimate ITRs. In this article, we propose a cost-effective estimation method from an active learning perspective. In particular, our method recruits only the “most informative” patients (in terms of learning the optimal ITRs) from an ongoing clinical trial. Simulation studies and real-data examples show that our active clinical trial method significantly improves on competing methods. We derive risk bounds and show that they support these observed empirical advantages. Supplementary materials for this article are available online.

Keywords

Active learning; Clinical trial; Individualized treatment rule; Personalized medicine; Risk bound

1. Introduction

It is widely recognized that different patients respond differently to the same treatment. Recent advances in personalized medicine have the potential to improve treatment decisions in clinical practice by tailoring the clinical interventions to the patient characteristics. These characteristics include demographics, medical histories, and genetic or genomic information (Hamburg and Collins 2010). It is anticipated that these new developments in personalized medicine may salvage some failed medications, which is especially important given the low overall success rate recently observed in clinical trials (DiMasi et al. 2010).

The success of personalized medicine is contingent on the correct identification of the best treatments for each individual. One research direction is subgroup analysis, where the patients are grouped based on the estimated individual-level treatment differences (Cai et al. 2011; Foster, Taylor, and Ruberg 2011). Alternatively, vigorous research has focused on

CONTACT Ying-Qi Zhao, yqzhao@biostat.wisc.edu, Department of Biostatistics & Medical Informatics, University of Wisconsin, Madison, WI, 53792.

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/r/JASA.

Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JASA.

Supplementary Materials

Supplementary materials available online include: additional simulation results with normal biomarkers; sensitivity analysis regarding different prespecified parameters; performance of a doubly robust augmented inverse probability weighted estimator; and proof of Lemma A.1.

finding optimal treatment regimens, which yield the greatest benefit overall for the whole population. Some methods involve fitting a regression model for the response, and recommending the treatment achieving the best prediction (Qian and Murphy 2011). Instead, Zhao et al. (2012) explored the optimal individualized therapies from a classification perspective (see also Zhang et al. 2012). All these methods are implemented using data from randomized clinical trials (RCTs). However, traditional RCTs are primarily designed to confirm the efficacy of new treatments; they do not generate comprehensive personalized therapy rules in an *efficient* manner. Consequently, post-mining data from RCTs are not ideal for finding optimal treatment strategies (Cui et al. 2002; Lagakos 2006). RCTs generally require a large sample size to demonstrate the efficacy of a candidate treatment, and they can be expensive to run because of the need to treat and monitor a large number of subjects. In addition, RCTs waste trial resources on subjects who experience treatment effects that are relatively large. Therefore, it would be desirable to design cost-effective clinical trials for personalized medicine. Such trials would highlight individual differences in responses and take advantage of continuing advances revealed in the trial (Singer 2005).

We propose an active learning framework for conducting clinical trials, called *active clinical trials*. In these trials, patients are judiciously recruited so that the optimal individualized treatment rules (ITRs) can be learned with fewer patients being randomized; this is a cost-effective method. Unlike traditional RCTs, we will exclude the patients for whom the benefit from one of the treatments is clearly observed, thus concentrating on those for whom the difference is less pronounced. Specifically, within the classification framework (Zhao et al. 2012), we first construct “confidence intervals” for the optimal decision boundary using the data accumulated so far based on either a frequentist or Bayesian approach. We then selectively enroll the patients whose optimal treatments are hard to determine, that is, their benefit differences from the different treatments are “small,” based on the above confidence intervals. Those patients are viewed as the most informative for the purpose of learning the optimal ITRs, and thus they are recruited for randomizations at the next stage. As will be seen in the empirical and theoretical analysis, the real-time selection of the right patients indeed improves the chance of discovering optimal ITRs with a drastically reduced sample size and cost.

Our work is related to the topic of budgeted learning (Madani, Lizotte, and Greiner 2004; Raghavan, Madani, and Jones 2006; Deng et al. 2007) in computer science. These works aim to find the most accurate classifier subject to a fixed budget. The main difference is that the responses for the subjects are assumed to be *known*, and the covariate information is required. Our work in the clinical trial setting sequentially applies treatments to selected subjects and then observes their responses. Another related topic in computer science is the “multi-armed bandit problem” (Robbins 1952), where resources are allocated among multiple arms given a fixed budget. One tries to maximize the cumulative rewards over all allocations. This perspective is also taken in response-adaptive trials (Rosenberger and Lachin 1993; Hu and Rosenberger 2006), which allocate more patients to the better treatment based on the available patient responses. Thus, such trials can be viewed as exploitative, since most of them adopt myopic strategies by exploiting the current best arm. Our method is more exploratory, exploring patients close to the decision boundary to learn better ITRs for *future* practice. Some adaptive enrichment designs (Wang, James Hung, and

O'Neill 2009; Simon and Simon 2013) are similar in that they allow the eligibility criteria to be updated during the trial. However, our primary goal is to construct informative and favorable ITRs rather than to confirm the efficacy of one treatment over another. Our method is also different from adaptive randomization procedures, which primarily aim to place more subjects onto the treatment arm that is more promising (Rosenberger, Sverdlov, and Hu 2012). Deng, Pineau, and Murphy (2011b) proposed a minimax bandit model for clinical trials that carefully distributes the trial resources, with the same objective as our method. However, they specify subpopulations in advance, and the optimal treatments are learned for each subpopulation. The resulting ITRs may not be ideal if the subpopulations are not formed correctly.

In Section 2, we introduce a general methodology for conducting active clinical trials, along with methods for constructing optimal ITRs using the accumulated data. In Section 3, we discuss the theoretical properties of our approach by providing a finite sample upper bound on the difference in the expected outcome under the estimated ITR and the optimal ITR. In Section 4, we conduct extensive simulation studies to examine the empirical performance and also compare the results with those of Deng, Pineau, and Murphy (2011b). Two real-data examples are presented in Section 5, and Section 7 provides a discussion. The technical details are presented in the Appendix and the online supplementary materials.

2. Methodology

2.1 Optimal Individual Treatment Rule

In this section, we discuss a probabilistic framework for studying the optimal ITRs that is similar to that of Qian and Murphy (2011) and Zhao et al. (2012). Let (X, A, R) be a random triple with a joint distribution P . Here, $X \in \mathbb{R}^p$ denotes the patient's baseline covariates with marginal distribution Π , A is a binary treatment assignment taking values in $\{-1, 1\}$, and R stands for the treatment outcome (a larger value of R corresponds to a better outcome). An ITR $D(\cdot)$ is defined as a function from the covariate space $\mathbb{R}^p$ into the treatment space $\{-1, 1\}$.

We use the value function, denoted as $V(D)$, to measure the quality of D , which is a marginal mean outcome representing the overall population mean were all patients to receive treatment according to D . Our goal is to identify the optimal ITR that yields the maximum $V(D)$. For any ITR D , let P^D be the distribution of (X, A, R) when $A = D(X)$, and let $\mathbb{E}^D$ be the corresponding expectation. The value function $V(D) = \mathbb{E}^D(R)$. Qian and Murphy (2011) showed that $V(D)$ can be expressed as

$$V(D) = \mathbb{E} \left[\frac{RI(A=D(X))}{\pi(A; X)} \right], \quad (2.1)$$

where $\mathbb{E}$ denotes the expectation w.r.t. the joint distribution P , $I(\cdot)$ is an indicator function, and $\pi(a; X)$ is the conditional probability $P(A = a|X)$ for $a \in \{-1, 1\}$. For simplicity, we assume a pure randomization scheme with equal probability for different assignments, that is, $\pi(a; X) = 1/2$, throughout the article.

Let D^* denote the optimal treatment rule that maximizes $V(D)$. By rewriting $V(D)$ as $V(D) = \mathbb{E}[\mathbb{E}[R\mathbb{I}(D(X) = 1)|A = 1, X] + \mathbb{E}[R\mathbb{I}(D(X) = -1)|A = -1, X]]$, we obtain

$$D^*(x) = \text{sign}\{f^*(x)\}, \quad (2.2)$$

where $f^*(x) = \mathbb{E}[R|A = 1, X = x] - \mathbb{E}[R|A = -1, X = x]$ is called the *contrast function*. The *optimal decision boundary* is just the level set $\{x \in \mathbb{R}^p : f^*(x) = 0\}$. As can be seen, the overall benefit will be maximized if the patients satisfying $f^*(x) \geq 0$ receive the alternative treatment (treatment 1), and the others receive standard care (treatment -1). Hence, the decision is tailored according to the patient's characteristics represented by X .

To estimate the optimal ITR D^* from the data $\{X^{(i)}, A^{(i)}, R^{(i)}\}_{i=1}^n$, usually collected from a clinical trial, one can fit a parametric or nonparametric regression model for $\mathbb{E}[R|A = a, X = x]$ (equivalently, $f^*(x)$), and then estimate the optimal ITR by plugging the fitted model into (2.2). Alternatively, we can replace the problem of maximizing (2.1) by minimizing a weighted classification error $\mathbb{E}[R\mathbb{I}(A \neq D(X))/\pi(A; X)]$. In this case, existing classification techniques, for example, support vector machines, can be adapted to estimate $D^*(x)$; see Zhao et al. (2012) for more details. All of the above methods take the whole (randomized) clinical trial data as an input, and thus they have no influence on the data-collection process. They belong to the class of *passive learning* methods, also called *batch learning*. In clinical trials, the patient-recruiting process is usually long, and the treatment and monitoring process can be extremely expensive. With a limited budget and a fixed sample size, we should wisely allocate the resources, that is, decide who to recruit, to learn the optimal ITRs at a lower cost. This motivates us to propose active clinical trials that can identify the optimal ITRs with a significantly reduced cost.

2.2 Active Clinical Trials

The active learning (AL) approach in the classification literature is shown to produce accurate classifiers with a significantly reduced number of label requests; see Dasgupta, Monteleoni, and Hsu (2007), Balcan, Hanneke, and Wortman (2008), Castro and Nowak (2008), Koltchinskii (2010), Hanneke (2011), and Minsker (2012). Recall that the estimation of ITRs can be thought of as a classification problem. For example, a patient with a small outcome given the assigned treatment is potentially misclassified. We use AL techniques to select the “most informative” patients based on a given pool of prognostic variables, and then we randomize only the patients with these “informative” characteristics. The intuition behind this patient-selection process is simple: if we have confidence that certain patients will benefit from a particular treatment, we should not recruit them into the study since they are not likely to contribute to the estimation of ITRs. However, it is hard to determine the optimal treatment, that is, the sign of $f^*(x)$, for the patients whose baseline variables are “close” to the optimal decision boundary. As will be seen from the empirical and theoretical analysis, patients with such features are more “informative” in the sense that their outcomes after randomization provide more insight into the optimal ITRs (to be applied to *future* patients). Moreover, since patients will not be randomized if we have confidence in the treatment they should receive, they are less likely to be exposed to ineffective medications.

We assume that there is no delay in observing the outcome. For $S \subset \text{supp}(\Pi)$ and a function $f: \mathbb{R}^p \mapsto \mathbb{R}$, let $f|_S: S \mapsto \mathbb{R}$ be the restriction of f onto S , and define

$$\|f\|_{\infty, S} = \|f|_S\|_{\infty} := \sup_{x \in S} |f(x)|. \text{ Given } \delta > 0, \text{ set}$$

$$\mathcal{F}_{\infty, S}(f, \delta) := \{g: S \mapsto \mathbb{R}: \|g - f\|_{\infty, S} \leq \delta\}$$

as a δ -band around f on S . A precise description is given in Algorithm 1. Here, we introduce the important notion of an *active* set (Minsker 2012), which underlies the majority of active learning algorithms. At step k of our Algorithm 1, the active set AS_k is defined to be the set of baseline variables for which the best treatment is not yet known. The active set is characterized by the confidence interval, that is, x belongs to the active set if and only if the confidence interval for $f^*(x)$ contains both positive and negative elements; Figures 1 and 2 illustrate this. The approximation of AS_k using a regular set act_k is discussed in Appendix A. 1.

In practice, it is not necessary to evaluate the active sets explicitly at each step. Instead, at step k , we repeat the following two steps until we reach the number of patients that can be randomized:

1. Recruit a group of new patients with baseline variables $X^{(i,k)}$, $i = 1, \dots, N_k$.
2. Let $\hat{f}_{i,k}(X^{(i,k)})$ be the estimator of $f^*(X^{(i,k)})$ based on S_k and $I_{i,k} = [\hat{f}_k(X^{(i,k)}) - \delta(X^{(i,k)}), \hat{f}_k(X^{(i,k)}) + \delta(X^{(i,k)})]$ be the confidence interval, where S_k is defined in Algorithm 1 and represents the available sample in the k th step.
 - a. If $0 \notin I_{i,k}$: drop the patient from the study.
 - b. If $0 \in I_{i,k}$: randomize the patient, record the outcome R , and add $(X^{(i,k)}, A^{(i,k)}, R^{(i,k)})$ to the current sample.

The above procedure gives a dataset containing “the most informative” observations for predicting the optimal treatment. For every new patient, this dataset will be used to predict the treatment rule based on that individual’s baseline variables. At each iteration, we must specify an estimator of the contrast function and the corresponding confidence interval. In the next section, we will propose two construction methods based on kernel estimation and Gaussian process regression, respectively. We remark that the empirical performance of our method is essentially controlled by the confidence level that we set. This level also determines the rejection ratio, for example, $N = 100$ out of 140 total iterations in simulation Scenario 1 of Section 4.

2.3 Confidence Interval Construction

2.3.1 Kernel Smoothing Approach—Let K be a smooth kernel function (with bandwidth h_n) satisfying Assumption (A2) in Section 3. Define

$K_{h_n}(x_0 - X) := K((x_0 - X)/h_n)/h_n^p$. Let $\eta_j(x) = \mathbb{E}[R|A = j, X = x]$, $j = \pm 1$, be the conditional expectation of R given A and X . Given n observations from P , we propose the

estimators $\hat{\eta}_j(\cdot; h_n)$ for $j = \pm 1$, and the plug-in estimator $\hat{f}(\cdot; h_n)$ of the contrast function $f^*(\cdot)$. At any fixed point x_0 ,

$$\hat{\eta}_j(x_0; h_n) = \frac{\sum_{i=1}^n K_{h_n}(x_0 - X^{(i)}) I(A^{(i)} = j) R^{(i)}}{\sum_{i=1}^n K_{h_n}(x_0 - X^{(i)}) I(A^{(i)} = j)}, j = \pm 1,$$

$$\hat{f}(x_0; h_n) = \hat{\eta}_1(x_0; h_n) - \hat{\eta}_{-1}(x_0; h_n). \quad (2.3)$$

Here, $h_n = h_n(x_0)$ is an adaptive bandwidth parameter varying with x_0 . We need to choose a proper $h_n(x_0)$, which controls the local amount of data near x_0 , to optimally balance the estimation bias and variance; see Appendix A.2 for the technical details.

Assuming that both η_1 and η_{-1} are Lipschitz continuous with Lipschitz constants bounded by L , we define

$$h_{n,j}(x_0) = \inf \left\{ h > 0 : L^2 h^2 \geq \frac{C_1(K, P)}{\sum_{i=1}^n I\{\|x_0 - X^{(i)}\|_2 \leq h\} I\{A^{(i)} = j\}} \right\}, j = \pm 1,$$

Algorithm 1

Active Clinical Trials for Personalized Medicine

input : Sample size limit N , confidence α

output : $\hat{D} = \text{sign}(\hat{f})$

1 $k = 0$, $\text{act}_0 = \text{supp}(\Pi)$;

Set the initial $N_0 := 2 \left\lfloor \sqrt{N} \right\rfloor$ and $LB = N - N_0$;

for $i = 1$ **to** N_0 **do**

Recruit $X^{(i,0)}$ from Π ;

Randomize $X^{(i,0)}$ to treatment $A^{(i,0)} = 1$ or -1 with equal probability;

Observe $R^{(i,0)}$;

end

2

Construct the estimator $\hat{f}_0(x)$ of f^* from $S_0 = \left\{ \left(X^{(i,0)}, A^{(i,0)}, R^{(i,0)} \right) \right\}_{i=1}^{N_0}$

while $LB > 0$ **do**

```

3    $\hat{\mathcal{F}}_k := \left\{ f: f|_{\text{act}_k} \in \mathcal{F}_{\infty, \text{act}_k}(\hat{f}_k; \frac{3}{2}\delta_k), f|_{\text{supp}(\Pi) \setminus \text{act}_k} \equiv \hat{f}_{k-1}|_{\text{supp}(\Pi) \setminus \text{act}_k} \right\}$ 
   confidence band size */;
4    $k := k + 1;$ 
    $AS_k := \{x \in \text{supp}(\Pi) : \exists f_1, f_2 \in \hat{\mathcal{F}}_{k-1}, \text{sign}(f_1(x)) \neq \text{sign}(f_2(x))\}$  /* active set */;
   Approximate  $AS_k$  with a regular set  $\text{act}_k$ ;
5   if  $\text{act}_k \cap \text{supp}(\Pi) = \emptyset$  then
       break
   end
   else
        $N_k = 2N_{k-1};$ 
       for  $i = 1$  to  $\lfloor N_k \cdot \Pi(\text{act}_k) \rfloor$  do
           Recruit  $X^{(i,k)}$  from the active set  $\hat{\Pi}_k := \Pi_{\text{act}_k}(dx)$ ;
           Randomize  $X^{(i,k)}$  to treatment  $A^{(i,k)} = 1$  or  $-1$  with equal probability;
           Observe  $R^{(i,k)}, S_k := \left\{ \left( X^{(i,k)}, A^{(i,k)}, R^{(i,k)} \right), i \leq \lfloor N_k \cdot \Pi(\text{act}_k) \rfloor \right\};$ 
       end
       Construct the estimator  $\hat{f}_k(\cdot)$  of  $f^*$  based on  $S_k$ ;
        $LB := LB - \lfloor N_k \cdot \Pi(\text{act}_k) \rfloor$ ;
        $\hat{f} := \hat{f}_k$  /* keeping track of the most recent estimator */;
   end
end

```

where $C_1(K, P)$ is a constant depending on the kernel K and distribution P , and $\|\cdot\|_2$ is the usual Euclidean norm. Set

$$h_n(x_0) := \max (h_{n,1}(x_0), h_{n,-1}(x_0)). \quad (2.4)$$

This choice mimics the usual “bias-variance tradeoff”: indeed, the bias of $\hat{\eta}(x_0; h)$ is bounded by the order of Lh , while $(\sum_{i=1}^n I\{\|x_0 - X^{(i)}\|_2 \leq h\} I\{A^{(i)} = j\})^{-1}$ plays the role of the variance parameter. Based on the above choice of $h_n(x_0)$, that is, (2.4), we define the radius of the confidence interval to be

$$\delta(x_0) := t \cdot Lh_n(x_0),$$

where t controls the coverage probability. The display above depends on the unknown constants $C_1(K, P)$ that must be chosen before running the algorithm. Given a certain confidence level, we recommend selecting the “confidence parameter” t by reverting the coverage probability error that decays exponentially fast with t ; see Corollary A.1.

Remark 1. We are aware that the proposed kernel methods are unfortunately affected by the “curse of dimensionality.” However, our theoretical analysis reveals that the empirical

performance of our kernel estimate essentially depends on the “intrinsic dimension” d of the support of the marginal distribution Π (see Assumption (A3) in Section 3), which might be much lower than the ambient dimension p . The concept of the intrinsic dimension originated in geometry and was later introduced to statistics, for example, Allard, Chen, and Maggioni (2012). It characterizes a low-dimensional representation embedded in a high-dimensional space. For example, when the data are linear, we may use principal component analysis to identify the subspace that contains the data. The number of important components is the intrinsic dimension d , which can be significantly smaller than the total number of covariates p . In practice, we follow the local singular value decomposition method, suggested in Little, Jung, and Maggioni (2009), to determine the intrinsic dimension.

2.3.2 Gaussian Process Regression Approach—The new method presented in this section is particularly well suited for applications since it is completely data-driven and does not require us to specify any unknown parameters in advance; its performance is demonstrated in Section 4. However, the price is that we must assume that the conditional distribution of the outcome R given (A, X) is Gaussian. In other words, if an individual with baseline variables X receives treatment $A = j$, we assume that the outcome satisfies $R = \eta_j(X) + \varepsilon$ for $j = \pm 1$, where $\varepsilon \sim N(0, \sigma^2)$. In this situation, we take a Bayesian approach by imposing a Gaussian process prior (defined by the mean $m(\cdot)$ and covariance $k(\cdot, \cdot)$) on functions $\eta_1(\cdot)$ and $\eta_{-1}(\cdot)$; see chap. 2 of Rasmussen and Williams (2006). For simplicity, the mean function $m(\cdot)$ is set to be identically zero throughout this article.

As before, the confidence interval of the contrast function builds upon those of $\eta_1(\cdot)$ and $\eta_{-1}(\cdot)$. Hence, we start from the inference procedure for $\eta_1(\cdot)$. Let $\{(X^{(i)}, R^{(i)})\}_{i=1}^{n_1}$ be the observations corresponding to the individuals who received treatment 1. The covariance function of the Gaussian process prior is set as a slight variant of the squared exponential kernel: given $x = (x_1, \dots, x_p)$ and $x' = (x'_1, \dots, x'_p) \in \mathbb{R}^p$,

$$k_\gamma(x, x') = \gamma_0 \exp \left(- \sum_{l=1}^p \frac{(x_l - x'_l)^2}{2\gamma_l^2} \right)$$

for some positive $\gamma_0, \dots, \gamma_p$. Let K_γ be a $p \times p$ matrix with entries $(K_\gamma)_{l_1, l_2} = k_\gamma(X_{l_1}, X_{l_2})$, $l_1, l_2 = 1, \dots, p$. Under the Gaussian error assumption, the marginal distribution of the vector $\mathbf{R} = (R^{(1)}, \dots, R^{(n_1)})^T$ given $\mathbf{X} = (X^{(1)}, \dots, X^{(n_1)})$ is a multivariate Gaussian with mean 0 and covariance matrix $K_{\mathbf{R}} := K_\gamma + \sigma^2 I_p$, where I_p is a $p \times p$ identity matrix. The “optimal” value $\Gamma_* := (\gamma_0^*, \dots, \gamma_p^*, \sigma_*^2)$ of the parameter $\Gamma = (\gamma_0, \dots, \gamma_p, \sigma^2)$ is then “learned” from the data by finding a local maximum of the marginal log-likelihood

$$\log p(\mathbf{R}|\mathbf{X}, \Gamma) := -\frac{1}{2} \mathbf{R}^T K_{\mathbf{R}}^{-1} \mathbf{R} - \frac{1}{2} \log \det K_{\mathbf{R}} - \frac{n_1}{2} \log 2\pi$$

with respect to $(\gamma_0, \dots, \gamma_p, \sigma^2)$; see sec. 5.4.1 in Rasmussen and Williams (2006) for more details. Therefore, the Gaussian process regression is a *global* method that can automatically

select the bandwidth by maximizing the data likelihood. This is in contrast with the previous kernel method where the bandwidth is selected *locally*; see Equation (2.4) above.

We next construct the confidence interval based on the posterior distribution of $\eta_1(\cdot)$ with the optimal Γ^* . Given a new observation with the baseline variable x_0 , the value of $\eta_1(x_0)$ is estimated by the posterior mean, denoted $\hat{\eta}_1(x_0)$, that is,

$$\hat{\eta}_1(x_0) := k_{\Gamma^*}(x_0, \mathbf{X})K_{\mathbf{R}}^{-1}\mathbf{R},$$

where $k_{\Gamma^*}(x, x') := \gamma_0^* \exp(-\sum_{l=1}^p (x_l - x'_l)^2 / 2(\gamma_l^*)^2)$ for $x, x' \in \mathbb{R}^p$, and

$$k_{\Gamma^*}(x_0, \mathbf{X}) = (k_{\Gamma^*}(x_0, X^{(1)}), \dots, k_{\Gamma^*}(x_0, X^{(n_1)})).$$

The variance of the posterior distribution is given by

$$\hat{\sigma}_1^2(x_0) := k_{\Gamma^*}(x_0, x_0) - k_{\Gamma^*}(x_0, \mathbf{X})K_{\mathbf{R}}^{-1}k_{\Gamma^*}(x_0, \mathbf{X})^T.$$

The square root of the posterior variance naturally controls the length of the confidence interval for $\eta_1(x_0)$. For example, setting $\delta_1(x_0) := 3\hat{\sigma}_1(x_0)$ gives the confidence interval for $\eta_1(x_0)$ of posterior probability $> 99\%$ (note that this may not correspond to the frequentist coverage).

Define $\hat{\eta}_{-1}(x_0)$ and the associated $\delta_{-1}(x_0)$ analogously. We thus obtain a confidence interval with the center $\hat{f}(x_0) := \eta_1(x_0) - \hat{\eta}_{-1}(x_0)$ and the radius $\delta(x_0) := \delta_1(x_0) + \delta_{-1}(x_0)$. The numerical implementation of this Bayesian inference procedure can easily be performed with the gpml Matlab toolbox (Rasmussen and Nickisch 2010).

3. Theoretical Analysis

In this section, we focus our theoretical analysis on the kernel smoothing approach. The frequentist property of the credible set for the nonparametric Bayesian method has not been well developed. Hence, a theoretical analysis of the Gaussian process regression method is beyond the scope of this article.

For simplicity, we suppose that the marginal distribution Π of the baseline variable vector X is known. Note that Algorithm 1 does not need to know or estimate Π explicitly, and this assumption is only for the theoretical analysis. In addition, we assume the following conditions on the kernel $K: \mathbb{R}^p \mapsto \mathbb{R}$ and the distribution Π :

- (A0) Both $\eta_1(x) := \mathbb{E}[R|A = 1, X = x]$ and $\eta_{-1}(x) := \mathbb{E}[R|A = -1, X = x]$ are Lipschitz continuous on $\mathbb{R}^p$ with Lipschitz constants bounded by L .
- (A1) The random variable $|R|$ is bounded by $0 < M < \infty$ a.s.
- (A2) $K(x)$ is a nonnegative, compactly supported, Lipschitz-continuous function with a Lipschitz constant L_K . Moreover, $K(x) \geq \ell_K I\{\|x\|_2 \leq 1\}$ for some $\ell_K > 0$.

- (A3) The “intrinsic dimension” of $\text{supp}(\Pi)$ is equal to d for some integer $d \leq p$, and Π is equivalent to the uniform distribution over its support; see Appendix 7 for a more precise statement.
- (A4) Margin condition: there exist $K_2 = K_2(\Pi)$, $\gamma = \gamma(\Pi) > 0$ such that for all $t > 0$

$$\Pi(x: |f^*(x)| \leq t) \leq K_2 t^\gamma.$$

Assumption (A3) says that over a “nice” set, Π is close to the uniform distribution; see condition (A.1) in the Appendix. The intrinsic dimension d is crucial here: in many applications, p is large but d is small (see Remark 1). Note that the rate in our main result, Theorem 3.1, depends on d but not on p . Assumption (A4) is an analogue of the well-known *margin condition* (Tsybakov 2004), which is commonly used to characterize the complexity of a binary classification problem. Larger values of γ mean that the two treatment effects are less likely to be similar, yet in nontrivial examples, $\gamma \in [0, d]$. In particular, as indicated in proposition 3.4 in Audibert and Tsybakov (2007), for a smooth contrast function $f^*(x)$, $\gamma \in [0, d]$ unless $f^*(x)$ does not cross 0 at any point in the interior of $\text{supp}(\Pi)$, that is, all the patients benefit from one treatment. Note that our analysis does not require γ to be known in advance.

We are now ready to present the (finite-sample) performance guarantee for our method.

Theorem 3.1

Let $h_k = \lceil \{\log(N/\alpha) + d \log(N_k)\} / N_k \rceil^{1/(d+2)}$, $k = 0, \dots, L$, where L is the total number of iterations in Algorithm 1. Set the associated $\delta_k = 4Ch_k$, where C is a constant specified in Lemma A.1. With probability greater than $1 - \alpha$, the estimator $\hat{D}$ of D^* returned by Algorithm 1 satisfies

$$|V(\hat{D}) - V(D^*)| \leq \tilde{C} N^{-\frac{1+\gamma}{2+d-\gamma}} (\log(N/\alpha))^\theta,$$

where N is the number of randomized subjects, $\theta = \frac{(4+2d-\gamma)(1+\gamma)}{(2+d)(2+d-\gamma)}$, and $\tilde{C}$ is a constant that depends on the kernel K and distribution Π .

It is worth noting that when γ is large (say, close to d), the rate of Theorem 3.1 is “almost” dimension-free. We also remark that the rate of Theorem 3.1 cannot be uniformly improved by any active learning technique, as shown in Minsker (2012). Qian and Murphy (2011) used a *parametric* modeling approach, that is, they fitted an L_1 -penalized regression model to estimate the optimal ITRs, and they obtained a rate of $(\log N/N)^{(1+\gamma)/(2+\gamma)}$ with an appropriate choice of tuning parameter. However, their model could be misspecified in practice. In contrast, our method is *nonparametric* with possibly minimal model assumptions.

4. Simulation Studies

In this section, we assess the empirical performance of the active clinical trial method. Let $X = (X_1, X_2, \dots, X_p)$, where $X_1, \dots, X_p$ are independent of each other. The distribution of X varies according to different scenarios detailed below. The treatment A is generated from $\{-1, 1\}$ with equal probability. The response R is generated from $\mathcal{N}(Q_0(X, A), 1)$, where

$$Q_0(X, A) = m_0(X) + T_0(X, A).$$

Here, $T_0(X, A)$ is the interaction between the treatment and the baseline variables. In what follows, $U[a, b]$ stands for the uniform distribution on the interval $[a, b] \subset \mathbb{R}$. Consider four scenarios for $T_0(X, A)$:

1. $X_l \sim U[-1, 1]$, $l = 1, 2$, $m_0(X) = 1 + 2X_1 + X_2$, $T_0(X, A) = 0.5(1 - X_1 - X_2)A$.
2. $X_l \sim U[-1, 1]$, $l = 1, 2$, $m_0(X) = 1 + 2X_1 + X_2$,
 $T_0(X, A) = 1/2 + (1 - X_1^2 - X_2^2)(X_1^2 + X_2^2 - 1)A$.
3. $p = 3$, $m_0(X) = 1 + 2X_1 + X_2 - X_3$, $T_0(X, A) = 1.5(X_1X_2(1 + X_3))A$,
 where X_1, X_2, X_3 are on the sphere generated as follows. Let $\tilde{X}_1, \dots, \tilde{X}_3 \sim U[-1, 1]$, and

$$X_l = \frac{\tilde{X}_l}{\sqrt{\sum_{l=1}^3 \tilde{X}_l^2}} I \left(\sum_{l=1}^3 \tilde{X}_l^2 \leq 1 \right), l = 1, 2, 3.$$

4. $p = 8$, $m_0(X) = 1 + 2X_1 + X_2 - X_3$, $T_0(X) = 0.2 \begin{cases} X_l & l \text{ is even} \\ -X_l & l \text{ is odd} \end{cases} A$, where $X_1, \dots, X_8$ have uniform distribution $U[-1, 1]$.

It can be seen that the optimal ITR for Scenario 1 is linear, that is, $D^*(X) = \text{sign}(1 - X_1 - X_2)$. The optimal ITR for Scenario 2 is nonlinear, with $D^*(X) = I(0.3 \leq X_1^2 + X_2^2 \leq 1.7)$. Scenario 3 represents the case where the data are supported on the manifold, that is, a two-dimensional sphere in $\mathbb{R}^3$. The treatment effect in Scenario 3 is highly nonlinear with $D^*(X) = \text{sign}(X_1X_2(1 + X_3))$. Scenario 4 has a relatively high-dimensional covariate, that is, 8, with a linear treatment effect.

We apply the active clinical trial outlined in Algorithm 1. We implemented both the kernel method in Section 2.3.1, denoted AL-BV, and the Gaussian process regression method in Section 2.3.2, denoted AL-GP. We used the gpml Matlab toolbox (Rasmussen and Nickisch 2010) in the latter method. The active clinical trial proceeds by selectively recruiting subjects whose differential treatment effects are smaller than a threshold, that is, Steps 3 and 4 of Algorithm 1. This iterative procedure screens out a certain number of subjects, whose optimal treatments can be determined with high confidence, and retains the remaining N subjects for estimating the optimal ITR. The active learning algorithm proposed by Deng, Pineau, and Murphy (2011a) serves as a comparison. The designed trial therein focuses on simple ITRs that use a small number of subpopulation categories to personalize the treatment. In particular, they assume the subpopulations are known a priori. Throughout the

clinical trial, they sequentially select the (subpopulation, treatment) pair so as to minimize the maximal error of selecting a suboptimal treatment. For all the scenarios, we form four subgroups by dichotomizing X_1 and X_2 at 0. We use MINMAXPICS, the algorithm name in their paper, to denote this method.

In addition, we compare with two passive learning approaches: the outcome weighted learning (OWL) method (Zhao et al. 2012) and the ordinary least squares (OLS) method. Both methods recruit subjects upon arrival in the clinical trial and estimate the optimal ITR using the collected data after the trial ends. N subjects are randomly selected for both methods. In the OWL method, using the available data, we minimize the target function

$\mathbb{P}_n[R\phi(Af(X))/\pi(A;X)] + \lambda_n \|f\|_2^2$ where $\phi(t) = \max(1 - t, 0)$ is the hinge loss, $\mathbb{P}_n$ is the empirical measure, and λ_n is a tuning parameter controlling the amount of penalization. We consider a nonlinear functional space for the ITRs, and a Gaussian kernel is used in the implementation. The optimal ITR is estimated via $\hat{D}(x) = \text{sign}(\hat{f}(x))$, where $\hat{f}(x)$ are the minimizers of the above objective function. In the OLS method, we first regress R on (X, A, XA) , and then estimate the optimal ITR by finding the treatment that yields a larger predicted outcome for each individual.

The initial sample size N_0 is fixed at 50, while the additional sample size $N - N_0$ is 50, 100, 200, 300, 400, 500, or 800. To evaluate the empirical performance, we generated a testing dataset of sample size 10,000, mimicking a large pool of future subjects. The estimated ITRs $\hat{D}(X)$ using the different methods are validated on this large testing set. Since the main effect is invariant across different ITRs, we can calculate the average excess value $\text{AEV}(D^*, \hat{D})$ as

$$\text{AEV}(D^*, \hat{D}) = \frac{1}{n} \sum_{i=1}^n [T_0(X^{(i)}, D^*) - T_0(X^{(i)}, \hat{D})],$$

$$n=10,000$$

where the empirical average is taken over the validation set. This quantity directly reflects the expected clinical benefits for future subjects treated according to $\hat{D}(X)$, with a smaller value indicating a better treatment decision. We repeat the process 1000 times and average the resulting values over all the runs. In Figure 3, we plot $\log\{\text{AEV}(D^*, \hat{D})\}$ against $\log(N - N_0)$, where $\hat{D}$ was obtained using each method. The log-scale is used to give a better display of the polynomial convergence rates in the different methods.

In all the scenarios, our active clinical trials perform uniformly better than OWL. In Scenario 1, the treatment effect is linear, which indicates that OLS is the best possible method. However, AL-GP has comparable performance, especially when the sample size is large. The performance of AL-BV also improves with N . When the treatment effect is nonlinear, as in Scenarios 2 and 3, the strength of active learning is clearly demonstrated. Both methods initially perform better for small sample sizes and then converge much faster as the sample size grows. In contrast, the values of the estimated ITRs from the other two methods do not converge (Scenario 3) or do not converge to the optimal value (Scenario 2). In Scenario 4,

the number of covariates is increased to eight. This places severe difficulties on the kernel estimation because of the curse of dimensionality. However, with the linear treatment effect, the AL-GP results are satisfactory compared with those of OLS. The MINIMAXPICS algorithm performs the best in Scenario 3, where the subpopulations are correctly defined in advance. When the prespecification of the subpopulations is incorrect, each subpopulation contains both patients who benefit from treatment 1 and patients who benefit from treatment -1. In this case, the constructed ITRs are not ideal, even with large sample sizes. In general, we believe that active learning methods provide robust results for various treatment mechanisms, which are usually unknown in practice. We presented some additional simulation results in the online supplementary materials, where we investigated scenarios with biomarkers of normal distribution, and conducted several sensitivity analyses.

5. Real Data Analysis

5.1 Nefazodone–CBASP Clinical Trial

We apply the proposed active learning methods to analyze the data from the Nefazodone–CBASP clinical trial (Keller et al. 2000). The randomized trial was conducted to compare the efficacy of three treatments for nonpsychotic chronic major depressive disorder (MDD), namely Nefazodone, cognitive behavioral-analysis system of psychotherapy (CBASP), or a combination of Nefazodone and CBASP. CBASP requires twice-weekly onsite visits to the clinic, and thus imposes a significant burden on patients compared with Nefazodone alone. Hence, we compare Nefazodone with the combination treatment, and investigate whether CBASP is necessary for all patients. We perform a complete case analysis. The score on the 24-item Hamilton rating scale for depression (HRSD) was the primary outcome, with higher scores indicating more severe depression. Data from 441 patients were available, with 218 patients randomized to Nefazodone and 223 to the combined treatment group. We consider three covariates for tailoring the treatment: the baseline HRSD scores, the alcohol dependence, and the HAMA somatic anxiety scores. The latter two covariates were selected by Gunter, Zhu, and Murphy (2007) as important variables for optimal treatment decision making.

All the patients are used for the OLS and OWL analysis. To mimic an active clinical trial, we first randomly select 50 patients as the initial dataset. We then sequentially examine the remaining patients, and drop those who do not satisfy the selection criteria. The recruitment stops once an additional 300 patients are enrolled, or all the patients have been examined. The eligible patients are included in the dataset, assuming that they have been randomized in an active clinical trial. The treatments they are randomized to in this hypothetical active clinical trial are the actual treatments they received in the completed trial. In particular, 289 and 350 patients were used to construct the optimal ITRs using AL-BV and AL-GP, respectively. The estimated ITRs from the different methods were applied to the whole dataset to calculate the average HRSD scores due to the ITRs, with smaller values being preferable. Here, the “average HRSD scores due to the ITRs” are defined to be $\mathbb{P}_n[\{\tilde{R}(A = D(X))/\pi(A)\}]/\mathbb{P}_n[\{I(A = D(X))/\pi(A)\}]$, where $\tilde{R}$ denotes the HRSD scores and $\pi(A)$ is the probability of being assigned treatment A . AL-BV recommends combination therapy for 299 patients, giving an average HRSD score of 8.86. AL-GP recommends this therapy for

386 patients, giving an average HRSD score of 9.72. OLS and OWL recommend this therapy for all the patients, yielding a higher score of 10.89. Hence, the treatment rules produced by the active clinical trials not only lead to a higher overall benefit, but also reduce the time and monetary commitments for the patients.

We then used a five-fold cross-validation analysis to avoid potential overfitting. The dataset was partitioned into five subsets. Four of the five subsets were used as training data to construct the optimal ITR, and the remaining subset was used as the validation set to evaluate the estimated rule. In the training subset, we applied both active learning and passive learning methods, that is, OLS and OWL, to construct the optimal ITRs. The initial sample size was set to 50 for the active learning methods. The number of additional recruited patients was 20, 30, ..., 200; these were adaptively selected from the rest of the training samples. For the passive learning methods, we randomly chose 70, 80, ..., 250 patients from the training data and conducted the estimation. The process was repeated 200 times, and we recorded the average cross-validated values for each sample size. The results are presented in Figure 4. The active learning methods initially lead to higher HRSD scores, but they catch up and continue to improve as the sample size grows. In particular, we can see that the HRSD scores from our methods decrease faster than those of the other methods do. After the additional sample size reaches 120, with the total sample size at 170, the ITR identified by AL-BV yields the lowest score, and this value can be further improved with larger sample sizes.

5.2 Twelve-Step Intervention on Stimulant Drug Use

These data come from an RCT that aims to evaluate the effectiveness of 8-week group intervention plus individual 12-step facilitative intervention for reducing stimulant drug use (Donovan et al. 2013). Individuals with stimulant-use disorders were randomly assigned to treatment as usual (TAU) or to TAU integrated with Stimulant Abuser Groups to Engage in 12-step (STAGE-12) intervention.

The primary outcome of interest is the number of days of self-reported stimulant drug use over the 3- to 6-month post-randomization period, where a smaller value is preferable. We use seven baseline variables to evaluate the patients and to construct the optimal ITR: age, average number of days per month of self-reported stimulant drug use in the 3 months prior to randomization, baseline alcohol use, drug use, employment status, medical status, and psychiatric status composite scores on the addiction severity index (ASI), where the ASI composite score, ranging from 0 (no endorsement of any problems) to 1 (maximal endorsement of all problems), is usually taken as an indication of problem severity; it is perceived to guide the treatment decision (McGahan et al. 1986).

After removing the missing data, we have 305 participants in total. We evaluated the different methods on the whole dataset, with the initial sample size set to 50 and the additional sample size set to 200. AL-BV and AL-GP assigned 123 and 90 patients to the STAGE-12 group, with expected outcomes of 11.3 and 12, respectively. OLS and OWL gave overall averages of 11.8 and 12.7 days. We also calculated the cross-validated number of days of drug use over the 3- to 6-month post-randomization period. Since the outcome was count data, with considerable zero-inflation, overdispersion, and a nonlinear trend, we

anticipated that the active learning method using a bias-variance tradeoff for the estimation would lead to the best results. Indeed, as shown in Figure 5, AL-BV outperforms the other methods with a fast overall decreasing trend. While the Gaussian assumption is severely violated in this example, AL-GP still improves with sample size. We note that OLS and OWL do not improve noticeably as the sample size grows.

6. Sample Size Consideration

To design a future study, it is important to plan ahead a sample size, which is sufficient to guarantee that we will obtain an ITR very close to the optimal one using the proposed active clinical trial. Here, we provide a data-driven method that gives a preliminary assessment of the required sample size. Further work is warranted to derive the sample size formula for such an exploratory trial. Let V_0 be an average patient outcome based on the standard care. Our goal is to determine a minimal sample size N^* such that: (1) the power to reject the hypothesis $V(D^*) = V_0$ is at least $1 - \beta$ when $V(D^*) \leq (1 + \rho)V_0$; and (2) $\Pr(|V(\hat{D}_{N^*}) - V(D^*)| \leq \epsilon) \geq 1 - \alpha$ (Laber et al. 2015), where ϵ , α , and ρ are prespecified constants and $\hat{D}_N$ is the resulting estimator based on N samples from the active clinical trial. Since we have constructed a nonasymptotic error bound for $|V(\hat{D}_N) - V(D^*)|$ in terms of the sample size N in Theorem 3.1, we suggest to obtain N^* by inverting the above error bound such that $\Pr(|V(\hat{D}_{N^*}) - V(D^*)| \leq \min(\rho V_0, \epsilon)) \geq 1 - \alpha$. By doing so, we are able to provide finite sample evidence that the optimal ITR yields a larger average benefit than the standard care, and, moreover, the average benefit of the estimated rule is approximately optimal. The constant $\gamma \in [0, d]$ will be determined based on prior knowledge, or can be set to a range of values to examine the resulting sample sizes. The unknown constant $\tilde{C}$ can be determined via simulations. Given d , γ , θ , and $\tilde{C}$, the required sample size N^* , such that with probability greater than $1 - \alpha$, $|V(\hat{D}_{N^*}) - V(D^*)| \leq \min(\rho V_0, \epsilon)$, can be found by setting

$\tilde{C} N^{*-\frac{1+\gamma}{2+d-\gamma}} (\log(N^*/\alpha))^\theta = \min(\rho V_0, \epsilon)$. The required sample sizes vary by the allocated initial sample size N_0 . Therefore, our sample size calculation is conditional on a specified N_0 . We will illustrate the proposed strategy by designing a future study to explore the optimal ITRs for patients with MDD. The Nefazodone-CBASP clinical trial is used as the basis for planning.

Consider two treatment options, Nefazodone only versus the combined treatment. The average HRSD score in the combined treatment group is 10.96. It is desirable to develop an ITR that can at least reduce the expected HRSD score that would be achieved by 15% with probability greater than 80%, that is, $\rho V_0 \approx 1.7$ and $\alpha = 0.2$. We will use three tailoring variables with $d = 3$. To determine $\tilde{C}$, we adopt a bootstrap method. We will bootstrap the data, a total of 441 patients, 10,000 times. An active clinical trial, with 50 initial and 100 additional patients, will be implemented with each bootstrapped sample to estimate the optimal ITR. The average HRSD score under each estimated ITR will be recorded. We

denote these scores as $\hat{V}_1^B, \dots, \hat{V}_{10,000}^B$. Let $k^* = \arg\min_{k=1, \dots, 10,000} \hat{V}_k^B$. $\hat{V}_{k^*}^B$ will be used as an estimate of $V(D^*)$. Denote the 80 percentile of $|\hat{V}_k^B - \hat{V}_{k^*}^B|$, $k = 1, \dots, 10,000$ as

$\widehat{\text{Diff}}$. $\tilde{C}$ can be obtained using $\tilde{C} = N^{\frac{1+\gamma}{2+d-\gamma}} (\log(N/\alpha))^{-\theta} \times \widehat{\text{Diff}}$, where $N = 150$ and $\alpha = 0.2$. We will assess different γ 's with $\gamma = 0.5, 1, 1.5, 2, 2.5$, and 3 , which lead to different θ 's

and different estimates of $\tilde{C}$. Using the outlined strategy, we suggest to plan the sample size for the future study according to Table 1.

7. Discussion

In this article, we propose an active clinical trial with the goal of constructing favorable ITRs at a minimal cost. This new paradigm is distinct from the standard clinical trial framework that is designed for treatment evaluation. Along with the new designs, we also present new analysis and inference tools that are often practically useful and theoretically efficient. Two methods are presented to construct the confidence interval in the algorithm. The kernel smoothing method is better in the situation where a few important tailoring variables are known a priori. However, due to the curse of dimensionality, it cannot handle a covariate space of moderate to high dimension. The Gaussian process regression approach performs better in this situation. However, the kernel smoothing method does not assume a specific underlying distribution of the outcomes, whereas the Gaussian process regression approach requires the outcome to follow a Gaussian distribution, which may not be appropriate for discrete outcomes. Since the goal of the proposed trial is to explore the optimal treatment rules, it is a learning stage. The results should be further validated in a confirmatory trial.

In practice, we may also want to include individuals who would be known to benefit from the treatment. Instead of dropping certain patients whose confidence interval for the contrast function does not contain 0, we can enroll all patients with different priority (quantified by probability). Specifically, we can prioritize those patients, whose confidence interval of the contrast function contains 0, by assigning larger probabilities, while enrolling patients, whose confidence interval for the contrast function does not contain 0, with lower probabilities. Under this more flexible framework, individuals who benefit from the treatment are also included in the trial such that the marginal treatment effects can be investigated. In addition, when there is a small difference between two medications, treatment choice will likely depend on other considerations such as cost, side effects, etc. Our framework can take into account those factors that can be quantified. By allowing a ratio of the cost per treatment relative to the cost of a worse disease outcome, say δ , which reflects the patients' tolerance of treatment burden relative to tolerance of the disease burden (Vickers, Kattan, and Sargent 2007; Huang and Laber 2014), we maximize $V(D) - \delta P(A = 1)$. The optimal rule is instead given by $D^*(x) = \text{sign}\{f^*(x) - \delta\}$. In this case, the active clinical trial will enroll patients from whom the value of $f^*(x)$ is close to δ , since these patients are close to the decision boundary and thus most informative.

We next discuss two directions for future research. First, our method can easily be extended to incorporate multi-category treatments, where patients are recruited based on the minimum differential treatment effect of all the pairwise comparisons. For continuous treatments, one possibility is to discretize the continuous treatment into different percentiles and then apply the extended method for multi-category treatments. Second, in this article, we assume that the outcomes of the previous patients have been observed before we enroll the next patient. In fact, the proposed active clinical trials can be conducted with delayed responses, provided the required estimates (for the contrast function) can be updated as data become available, that is, some responses can be collected during the study period. We can also update the

estimates after groups of responses instead of individual responses. There have been some recent developments on handling the problem of delayed outcomes in phase I–II trials, for example, by treating them as missing values and applying imputation strategies (Jin et al. 2014). We would like to explore such options in the future. There are two other interesting extensions that we are pursuing: (i) a more general contrast function that can accommodate a high-dimensional or discrete covariate; and (ii) a dynamic treatment regime for a sequence of treatment rules (Murphy 2003; Robins 2004).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Guang Cheng was visiting SAMSI and on sabbatical at Princeton while this work was carried out and revised; he would like to thank SAMSI and Princeton ORFE department for their hospitality and support. We thank Dr. A. John Rush and the investigators for use of their data from the Nefazodone CBASP trial. The stimulant data used in this article were obtained from the datasets distributed by the NIDA.

Funding

Research Sponsored by NSF (DMS-0906497, CAREER Award DMS-1151692, DMS-1418042), Simons Fellowship in Mathematics, Office of Naval Research (ONR 11845813), and a grant from Indiana Clinical and Translational Sciences Institute.

Appendix A: Technical Proofs

A.1. Intrinsic Dimension of $\text{supp}(\Pi)$

We explain the meaning of “intrinsic dimension” introduced in Assumption (A3) here. We say that $\text{supp}(\Pi)$ possesses a *tree decomposition* $\mathcal{T} = \{T_{i,j} : i = 1, j = 1 \dots J(i)\}$ if

1. $T_{1,1} = \text{supp}(\Pi)$, and $\{T_{i,j}\}_{j=1}^{J(i)}$ forms a disjoint partition of $\text{supp}(\Pi)$ for all $i \geq 1$.
2. Nested partition: $\forall i \geq 2, j = 1, \dots, J(i)$, there exists a unique $1 \leq k \leq J(i-1)$ such that $T_{i,j} \subset T_{i-1,k}$.
3. Bounded diameter: for all $i \geq 1, 1 \leq j \leq J(i)$,

$$\text{diam}(T_{i,j}) := \sup_{x,y \in T_{i,j}} \|x - y\|_2 \leq K_1 2^{-i}$$

for some $K_1 = K_1(\Pi)$.

4. Regularity: for any $i \geq 1, 1 \leq j \leq J(i)$ and $0 < r \leq 2^{-i}$ the following holds: there exists a $1 \leq d \leq p$ (d is the intrinsic dimension) such that for all $x \in T_{i,j}$

$$c_1 r^d \leq \Pi(B(x, r) \cap T_{i,j}) \leq c_2 r^d \quad (\text{A.1})$$

for some $0 < c_1(\Pi) \leq c_2(\Pi) < \infty$, which are independent of i, j . Here, $B(x, r)$ is the Euclidean ball of radius r centered at x .

A simple example that gives a good intuition to the tree decomposition is the uniform distribution over the unit cube in $\mathbb{R}^p$. In this case, the tree decomposition is given by partitioning the unit cube into dyadic cubes and $d = p$. If $\text{supp}(\Pi)$ is contained in a proper subspace W of $\mathbb{R}^p$, then $d = \dim(W)$.

Let $\mathcal{B}_j$ be the sigma-algebra generated by the collection of sets $\{T_{i,j}, j = 1, \dots, J(i)\}$ (a partition of $\text{supp}(\Pi)$ on the level j). A regular approximation of AS_k in Algorithm 1 is given by $\text{act}_k := \cap \{A : A \in \mathcal{B}_{m_{k-1}}, A \supset \text{AS}_k\}$.

A.2. Properties of Kernel Estimate

A.2.1. Preliminaries

For a measurable set $S \subset \text{supp}(\Pi)$, define $\Pi_S(dx) := \Pi(dx|x \in S)$ as the conditional distribution on S , and set

$$Q_h(x|S) := \int_{\mathbb{R}^p} K_h(x - y) d\Pi_S(y).$$

Since Π is assumed to be known, we can directly compute $Q_h(x|S)$ now. Accordingly, we modify the original kernel estimate for η_j , that is, (2.3), as follows: let $\{(X^{(i)}, A^{(i)}, R^{(i)}), i = 1, \dots, N\}$ be an iid sample from the conditional joint distribution of (X, A, R) given that $X \in S$, and set

$$\hat{\eta}_j(x; h, S) = \frac{1}{N} \sum_{i=1}^N \frac{R^{(i)} I\{A^{(i)} = j\} K_h(x - X^{(i)})}{Q_h(x|S) P(A^{(i)} = j)}, j = \pm 1,$$

$$\hat{f}(x; h, S) = \hat{\eta}_1(x; h, S) - \hat{\eta}_{-1}(x; h, S). \quad (\text{A.2})$$

We will discuss properties of these estimators in below.

Let $h > 0$, $S \in \mathcal{B}_j$, and $h = 2^{-j}$, and define

$$Q_{h,m}(x|S) := \int_{\mathbb{R}^p} \|x - y\|_2^m K_h(x - y) d\Pi_S(y). \quad (\text{A.3})$$

We next study the upper and lower bounds of $Q_{h,m}(x|S)$ based on Assumptions (A1)–(A4). Since K is bounded and compactly supported, there exists $R = R_K > 0$ such that $K(x) \leq \|K\|_\infty I\{x \in B(0, R_K)\}$. Let $F > 0$ be a large enough constant, namely, $F^d \geq 2c_2/c_1$. Recall that c_1, c_2 are defined in (A.1). Note that Assumption (A2) implies the following:

$$\begin{aligned}
\mathbb{Q}_{h,m}(x|S) &\geq \ell_K \int_{B(x,h) \cap S} \|x - y\|_2^m d\Pi_S(y) \\
&\geq \ell_K (h/F)^m \left(\int_{B(x,h) \cap S} d\Pi_S(y) - \int_{B(x,h/F) \cap S} d\Pi_S(y) \right) \\
&\geq \ell_K (h/F)^m \frac{(c_1 h^d - c_2 (h/F)^d)}{\Pi(S)} \geq \frac{1}{2F^m} \ell_K c_1 \frac{h^{d+m}}{\Pi(S)} \\
&:= c_3 \frac{h^{d+m}}{\Pi(S)},
\end{aligned} \tag{A.4}$$

and

$$\mathbb{Q}_{h,m}(x|S) \leq \|K\|_\infty \int_{B(x, R_K h) \cap A} \|x - y\|_2^m d\Pi_S(y) \leq \|K\|_\infty R_K^{m+d} c_2 \times \frac{h^{d+m}}{\Pi(S)} := c_4 \frac{h^{d+m}}{\Pi(S)}.$$

(A.5)

In what follows, we will set $Q_{h,x|S} := Q_{h,0}(x)$ for brevity.

A.2.2. Some Bounds for the Kernel Estimators

In this subsection, we derive basic concentration inequalities for the kernel estimators of $\eta_j(x) = \mathbb{E}[R|A=j, X=x]$, $j = \pm 1$ restricted to S , that is, $\hat{\eta}_j(x; h, S)$ defined in (A.2). The proof of the results can be found in the online supplementary materials.

Lemma A.1—For all $t > 0$ satisfying $t + d^2 \log(1/h) \leq nh^d$, with probability $1 - 2e^{-t}$,

$$\sup_{x \in \text{supp}(\Pi) \cap S} |\hat{\eta}_j(x; h) - \eta_j(x)| \leq C \left(h + \sqrt{\frac{\Pi(S)(t + d^2 \log(1/h))}{nh^d}} \right),$$

where $C = C(M, c_1, c_2, L, L_K, \|K\|_\infty, \ell_K, R_K)$ is a constant.

The following corollary is immediate:

Corollary A.1—Set $h_n := \{\Pi(S)(t + d \log(n/\Pi(S)))/n\}^{1/(d+2)}$. Then, under assumptions of Lemma A.1, with probability $1 - 4e^{-t}$,

$$\sup_{x \in \text{supp}(\Pi) \cap S} |\hat{f}(x; h_n) - f^*(x)| \leq 4Ch_n,$$

where constant C is the same as in Lemma A.1.

A.3. Proof of Theorem 3.1

A.3.1. Comparison Inequality

Our Lemma A.2 below illustrates the connection between the risk $V(\hat{D}) - V(D^*)$ of a treatment rule $\hat{D}(x) = \text{sign}(\hat{f}(x))$ and the sup-norm $\|\hat{f} - f^*\|_{\infty, \text{supp}(\Pi)}$.

Lemma A.2—Under the margin assumption (A4),

$$V(\hat{D}) - V(D^*) \leq C(\gamma) \|\hat{f} - f^*\|_{\infty, \text{supp}(\Pi)}^{\frac{1+\gamma}{1+\gamma}} \left\{ \text{sign}(\hat{f}) \neq \text{sign}(f^*) \right\}.$$

Proof: It is easy to see that $V(\hat{D}) - V(D^*) = 2\mathbb{E}(|f^*(X)|I\{\hat{D}(X) \neq D^*(X)\})$. The rest of the argument repeats lemma 5.1 in Audibert and Tsybakov (2007).

A.3.2. Main Proof

Our main goal is to control the size of the set act_k defined by Algorithm 1. In turn, these bounds depend on the size of the confidence bands for $f^*(x)$ (denoted by δ_k). Suppose $L = N$ is the number of iterations performed by the algorithm before termination.

Let $N_k^{\text{act}} := \lfloor N_k \cdot \Pi(\text{act}_k) \rfloor$ be the number of labels requested on the k th iteration of the algorithm. We first claim that the following bounds hold uniformly for all $1 \leq k \leq L$ with probability at least $1 - \alpha$:

$$\|f^* - \hat{f}_k\|_{\infty, \text{act}_k} \leq C_1 \left(\frac{\log(N/\alpha) + d \log(N_k)}{N_k} \right)^{1/(d+2)},$$

$$\Pi(\text{act}_k) \leq C_2 \left(\frac{\log(N/\alpha) + d \log(N_{k-1})}{N_{k-1}} \right)^{\gamma/(d+2)} \quad (\text{A.6})$$

where $C_j = C_j(M, c_1, c_2, L, L_K, \|K\|_{\infty}, \ell_K, R_K, \gamma)$, $j = 1, 2$. This claim will be proved later.

Let $\mathcal{E}$ be the event of probability $1 - \alpha$ on which both inequalities of (A.6) hold, and assume that it occurs. Second inequality of (A.6) implies, together with the fact that $N_k = 2N_{k-1}$ by definition, that the number of randomized subjects on each step $1 \leq k \leq L$ satisfies

$$N_k^{\text{act}} = \lfloor N_k \Pi(\text{act}_k) \rfloor \leq 2N_{k-1}^{\frac{2+d-\gamma}{2+d}} \left(\log(N/\alpha) + d \log(N_{k-1}) \right)^{\gamma/(d+2)}$$

with probability $1 - \alpha$. If N is the maximum number of randomized subjects the algorithm is allowed to request, then

$$N \leq \sum_{k=0}^L N_k^{\text{act}} \leq 2(\log(N/\alpha) + d \log(N_L))^{\gamma/(d+2)} \sum_{k=0}^L N_k^{\frac{2+d-\gamma}{2+d}} \leq C_3(\gamma, d)(\log(N/\alpha) + d \log(N_L))^{\gamma/(d+2)} N_L^{\frac{2+d-\gamma}{2+d}},$$

and one easily deduces that on the last iteration L we have

$$N_L \geq c(\gamma, \Pi, d) \left(\frac{N}{\log(N/\alpha)} \right)^{\frac{2+d}{2+d-\gamma}}. \quad (\text{A.7})$$

Recall that N_L is defined in Algorithm 1.

To obtain the risk bound of the theorem from (A.7), we apply Lemma A.2:

$$V(\hat{D}) - V(D^*) \leq C(\gamma) \left\| (\hat{f}_L - f^*) \cdot I \left\{ \text{sign}(\hat{f}_L) \neq D^* \right\} \right\|_{\infty, \text{supp}(\Pi)}^{1+\gamma}. \quad (\text{A.8})$$

Since $\left\{ \text{sign}(\hat{f}_L) \neq D^* \right\} \subseteq \text{act}_L$ whenever bounds (A.6) hold, it remains to estimate

$\|\hat{f}_L - f^*\|_{\infty, \text{act}_L}$. Recalling the first inequality of (A.6) once again (for $k = L$), we get

$$\|\hat{f}_L - f^*\|_{\infty, \text{act}_L} \leq C_1 \left(\frac{\log(N/\alpha) + d \log(N_L)}{N_L} \right)^{1/(d+2)} \leq \tilde{C} N^{-\frac{1}{2+d-\gamma}} (\log(N/\alpha))^q,$$

where $q = \frac{4+2d-\gamma}{(2+d)(2+d-\gamma)}$, which together with (A.8) yields the final result.

It remains to show both inequalities of (A.6). We start with the bound on $\|\hat{f}_k - f^*\|_{\infty, \text{act}_k}$. First, note that by construction, for every $k \geq 1$ the samples $(X^{(i,k)}, A^{(i,k)}, R^{(i,k)})$, $i = 1 \dots \lfloor N_k \rfloor$

$\Pi(\text{act}_k)$ are conditionally independent given the data $\cup_{i=1}^{k-1} S_i$ collected on steps 1, ..., $k-1$, with conditional distribution of $X^{(i,k)}$ being Π_{act_k} . Thus, we can apply Corollary A.1

conditionally on $\cup_{i=1}^{k-1} S_i$ with $t = \log \frac{4N}{\alpha}$ to get that with probability $1 - \alpha/N$,

$$\|\hat{f}_k - f^*\|_{\infty, \text{act}_k} \leq 4C \left(\frac{\log \frac{\alpha}{4N} + d \log(\lfloor N_k \Pi(\text{act}_k) \rfloor / \Pi(\text{act}_k))}{\lfloor N_k \Pi(\text{act}_k) \rfloor / \Pi(\text{act}_k)} \right)^{1/(d+2)} \leq 8Ch_k.$$

It remains to integrate the bound with respect to the distribution of $\cup_{i=1}^{k-1} S_i$:

$$P \left(\|\hat{f}_k - f^*\|_{\infty, \text{act}_k} \geq 8Ch_k \right) = \mathbb{E} P \left(\|\hat{f}_k - f^*\|_{\infty, \text{act}_k} \geq 8Ch_k \mid \cup_{i=1}^{k-1} S_i \right) \leq \frac{\alpha}{N}.$$

The union bound over all $1 \leq k \leq L$ gives the result.

Finally, we will prove the second inequality of (A.6), the bound for the size of the active sets act_k . This is the place where assumption (A3) on the tree decomposition and margin assumption (A4) play the key role. To obtain the bound, we will compare two estimators of f^* : the first is the kernel estimator $\widehat{f}_k$ constructed by the Algorithm 1 on step k , and the second is the piecewise-constant estimator $\bar{f}_k$ with similar approximation properties to $\widehat{f}_k$. Namely, $\bar{f}_k$ is the $L_2(\Pi)$ -projection of f^* on the linear space of piecewise-constant functions

of the form $g(x) = \sum_{j=1}^{J(m_k)} \alpha_j I\{T_{m_k,j}\}(x)$, $\alpha_j \in \mathbb{R}$. Recall that $T_{i,j}$ is defined in the tree decomposition of Section 7. As a result, we will be able to relate the “active sets” associated to these estimators, taking advantage of the fact that the active set associated to $\bar{f}_k$ is always a union of the sets from a collection $\{T_{m_k,j} : j = 1 \dots J(m_k)\}$.

Let $\mathcal{E}_1$ be the event of probability $1 - \alpha$ on which $\|\widehat{f}_k - f^*\|_{\infty, \text{act}_k} \leq \delta_k$ for any $k \geq 0$, where $\delta_k = 4Ch_k$. Assume that $\mathcal{E}_1$ occurs.

The following inclusions hold (for the definition of AS_{k+1} , see Algorithm 1):

$$\{x : |f^*(x)| < \delta_k/2\} \subseteq \text{AS}_{k+1} \subseteq \{x : |f^*(x)| < 5\delta_k/2\}. \quad (\text{A.9})$$

Indeed,

$$|f^*(x)| < \delta_k/2 \Rightarrow |\widehat{f}_k(x)| < \delta_k/2 + |f^*(x) - \widehat{f}_k(x)| < \frac{3}{2}\delta_k \Rightarrow x \in \text{AS}_{k+1}$$

and

$$x \in \text{AS}_{k+1} \Rightarrow |\widehat{f}_k(x)| < \frac{3}{2}\delta_k \Rightarrow |f^*(x)| < \frac{5}{2}\delta_k.$$

For all $x \in T_{m_k,j}$ set $\bar{f}_k(x) := \frac{1}{\Pi(T_{m_k,j})} \int_{T_{m_k,j}} f^*(y) d\Pi(y)$, and note that

$$\begin{aligned} |f^*(x) - \bar{f}_k(x)| &\leq \frac{1}{\Pi(T_{m_k,j})} \int_{T_{m_k,j}} |f^*(y) - f^*(x)| d\Pi(y) \\ &\leq \frac{2L}{\Pi(T_{m_k,j})} \int_{T_{m_k,j}} |x - y| d\Pi(y) \\ &\leq 2L \text{diam}(T_{m_k,j}) \leq 2LK_1 2^{-m_k} \leq 4LK_1 h_k, \end{aligned}$$

where the last two inequalities follow from part 3 of assumption (A3) given in Appendix 7, and from the definition of m_k . Define $\tau_k := \max(5\delta_k, 4LK_1 h_k)$. $\mathcal{E}_{k+1} := \{f : |f(x) - \bar{f}_k(x)| \leq (3/2)\tau_k, \forall x \in \text{act}_k\}$ to be the band of size $(3/2)\tau_k$ around $\bar{f}_k$, and

$$\bar{A}_{k+1} := \left\{ x \in \text{act}_k : \exists f_1, f_2 \in \bar{\mathcal{F}}_{k+1}, \text{sign}(f_1(x)) \neq \text{sign}(f_2(x)) \right\}.$$

By a reasoning similar to above, we have the inclusions

$$\{x: |f^*(x)| < \tau_k/2\} \subseteq \bar{A}_{k+1} \subseteq \{x: |f^*(x)| < 5\tau_k/2\}. \quad (\text{A.10})$$

Moreover, by the definition of τ_k , we have the inequality $5\delta_k/2 \leq \tau_k/2$. Hence (A.9) and (A.10) imply that $\text{AS}_{k+1} \subseteq \bar{A}_{k+1}$. It remains to note that

$\bar{A}_{k+1}$ is the union of the sets from a collection $\{T_{m_k, j}, j = 1 \dots J(m_k)\}$, hence $\bar{A}_{k+1} \supseteq \text{act}_{k+1}$;

By (A.10) and assumption (A4),

$$\Pi(\text{act}_{k+1}) \leq \Pi(\bar{A}_{k+1}) \leq \Pi(\{x: |f^*(x)| < 5\tau_k/2\}) \leq K_2(5\tau_k/2)^\gamma \leq C_6\delta_k^\gamma,$$

hence proving the claim.

References

- Allard WK, Chen G, Maggioni M. Multi-Scale Geometric Methods for Data Sets II: Geometric Multi-Resolution Analysis. *Applied and Computational Harmonic Analysis*. 2012; 32:435–462.
- Audibert J-Y, Tsybakov AB. Fast Learning Rates for Plug-In Classifiers. *The Annals of Statistics*. 2007; 35:608–633.
- Balcan, M-F.; Hanneke, S.; Wortman, J. The True Sample Complexity of Active Learning; *Proceedings of the Conference on Learning Theory*; 2008. p. 45-56.
- Cai T, Tian L, Wong PH, Wei L. Analysis of Randomized Comparative Clinical Trial Data for Personalized Treatment Selections. *Biostatistics*. 2011; 12:270–282. [PubMed: 20876663]
- Castro RM, Nowak RD. Minimax Bounds for Active Learning. *IEEE Transactions on Information Theory*. 2008; 54:2339–2353.
- Cui L, Hung H, Wang SJ, Tsong Y. Issues Related to Subgroup Analysis in Clinical Trials. *Journal of Biopharmaceutical Statistics*. 2002; 12:347–358. [PubMed: 12448576]
- Dasgupta S, Monteleoni C, Hsu DJ. A General Agnostic Active Learning Algorithm. *Advances in Neural Information Processing Systems*. 2007:353–360.
- Deng, K.; Bourke, C.; Scott, S.; Sunderman, J.; Zheng, Y. Bandit-Based Algorithms for Budgeted Learning; *Proceedings of the 2007 Seventh IEEE International Conference on Data Mining*; 2007. p. 463-468.
- Deng, K.; Pineau, J.; Murphy, S. Active Learning for Developing Personalized Treatment; *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*; 2011a. p. 161-168.
- Deng, K.; Pineau, J.; Murphy, S. Active Learning for Personalizing Treatment; *2011 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*; 2011b. p. 32-39.
- DiMasi JA, Feldman L, Seckler A, Wilson A. Trends in Risks Associated With New Drug Development: Success Rates for Investigational Drugs. *Clinical Pharmacology & Therapeutics*. 2010; 87:272–277. [PubMed: 20130567]
- Donovan DM, Daley DC, Brigham GS, Hodgkins CC, Perl HI, Garrett SB, Doyle SR, Floyd AS, Knox PC, Botero C, Kelly TM, Killeen TK, Hayes C, Kau'i Baumhofer N, Seamans C, Zammarelli L. Stimulant Abuser Groups to Engage in 12-Step: A Multisite Trial in the National Institute on Drug Abuse Clinical Trials Network. *Journal of Substance Abuse Treatment*. 2013; 44:103–114. [PubMed: 22657748]

- Foster JC, Taylor JM, Ruberg SJ. Subgroup Identification From Randomized Clinical Trial Data. *Statistics in Medicine*. 2011; 30:2867–2880. [PubMed: 21815180]
- Gunter, L.; Zhu, J.; Murphy, S. Variable Selection for Optimal Decision Making. In: Bellazzi, R.; Abu-Hanna, A.; Hunter, J., editors. *Artificial Intelligence in Medicine*. Berlin: Springer; 2007. p. 149-154.
- Hamburg M, Collins F. The Path to Personalized Medicine. *The New England Journal of Medicine*. 2010; 363:301–304. [PubMed: 20551152]
- Hanneke S. Rates of Convergence in Active Learning. *The Annals of Statistics*. 2011; 39:333–361.
- Hu, F.; Rosenberger, WF. *The Theory of Response-Adaptive Randomization in Clinical Trials*. New York: Wiley; 2006.
- Huang Y, Laber E. Personalized Evaluation of Biomarker Value: A Cost-Benefit Perspective. *Statistics in Biosciences*. 2014;1–23. [PubMed: 27594925]
- Jin IH, Liu S, Thall PF, Yuan Y. Using Data Augmentation to Facilitate Conduct of Phase I-II Clinical Trials with Delayed Outcomes. *Journal of the American Statistical Association*. 2014; 109:525–536. [PubMed: 25382884]
- Keller MB, McCullough JP, Klein DN, Arnow B, Dunner DL, Gelenberg AJ, Markowitz JC, Nemeroff CB, Russell JM, Thase ME, Trivedi MH, Zajecka J. A Comparison of Nefazodone, The Cognitive Behavioral-Analysis System of Psychotherapy, and Their Combination for the Treatment of Chronic Depression. *The New England Journal of Medicine*. 2000; 342:1462–1470. [PubMed: 10816183]
- Koltchinskii V. Rademacher Complexities and Bounding the Excess Risk in Active Learning. *Journal of Machine Learning Research*. 2010; 11:2457–2485.
- Laber E, Zhao Y-Q, Regh T, Davidian M, Tsiatis A, Stanford JB, Zeng D, Song R, Kosorok MR. Sizing a Phase II Trial to Find a Nearly Optimal Personalized Treatment Strategy. 2015 submitted.
- Lagakos SW. The Challenge of Subgroup Analyses-Reporting Without Distorting. *New England Journal of Medicine*. 2006; 354:1667–1669. [PubMed: 16625007]
- Little, AV.; Jung, Y-M.; Maggioni, M. Multiscale Estimation of Intrinsic Dimensionality of Data Sets; AAAI Fall Symposium:Manifold Learning and Its Applications; 2009. p. 26-33.
- Madani, O.; Lizotte, DJ.; Greiner, R. The Budgeted Multi-Armed Bandit Problem. In: Shawe-Taylor, J.; Singer, Y., editors. *Learning Theory*. Berlin: Springer; 2004. p. 643-645.
- McGahan, P.; Griffith, J.; Parente, R.; McLellan, A. *Addiction Severity Index Composite Scores Manual*. Philadelphia, PA: Treatment Research Institute; 1986.
- Minsker S. Plug-In Approach to Active Learning. *Journal of Machine Learning Research*. 2012; 13:67–90.
- Murphy SA. Optimal Dynamic Treatment Regimes. *Journal of the Royal Statistical Society*. 2003; 65:331–366.
- Qian M, Murphy SA. Performance Guarantees for Individualized Treatment Rules. *The Annals of Statistics*. 2011; 39:1180–1210. [PubMed: 21666835]
- Raghavan H, Madani O, Jones R. Active Learning With Feedback on Features and Instances. *Journal of Machine Learning Research*. 2006; 7:1655–1686.
- Rasmussen CE, Nickisch H. Gaussian Processes for Machine Learning (GPML) Toolbox. *Journal of Machine Learning Research*. 2010; 11:3011–3015.
- Rasmussen, CE.; Williams, CKI. *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press; 2006.
- Robbins H. Some Aspects of the Sequential Design of Experiments. *Bulletin of the American Mathematical Society*. 1952; 58:527–535.
- Robins, JM. Optimal Structural Nested Models for Optimal Sequential Decisions; Proceedings of the Second Seattle Symposium on Biostatistics; 2004. p. 189-326.
- Rosenberger WF, Lachin JM. The Use of Response-Adaptive Designs in Clinical Trials. *Controlled Clinical Trials*. 1993; 14:471–484. [PubMed: 8119063]
- Rosenberger WF, Sverdlov O, Hu F. Adaptive Randomization for Clinical Trials. *Journal of Biopharmaceutical Statistics*. 2012; 22:719–736. [PubMed: 22651111]

- Simon N, Simon R. Adaptive Enrichment Designs for Clinical Trials. *Biostatistics*. 2013; 14:613–625. [PubMed: 23525452]
- Singer E. Personalized Medicine Prompts Push to Redesign Clinical Trials. *Nature Medicine*. 2005; 11:462.
- Tsybakov A. Optimal Aggregation of Classifiers in Statistical Learning. *The Annals of Statistics*. 2004; 32:135–166.
- Vickers AJ, Kattan MW, Sargent DJ. Method for Evaluating Prediction Models That Apply the Results of Randomized Trials to Individual Patients. *Trials*. 2007; 8:14. [PubMed: 17550609]
- Wang S-J, Hung H, O'Neill RT. Adaptive Patient Enrichment Designs in Therapeutic Trials. *Biometrical Journal*. 2009; 51:358–374. [PubMed: 19358222]
- Zhang B, Tsiatis AA, Laber EB, Davidian M. A Robust Method for Estimating Optimal Treatment Regimes. *Biometrics*. 2012; 68:1010–1018. [PubMed: 22550953]
- Zhao YQ, Zeng D, Rush AJ, Kosorok MR. Estimating Individualized Treatment Rules Using Outcome Weighted Learning. *Journal of American Statistical Association*. 2012; 107:1106–1118.

Figure 1.
Confidence band for the contrast function.

Figure 2.
 $X^{(1)}$ belongs to the active set; $X^{(2)}$ does not.

Figure 3.
Excess values (log scale). Initial size was set to 50.

Figure 4.
Mean cross-validated HRSD scores (the lower the better) against additional sample sizes.
The initial sample size was set to 50.

Figure 5. Mean cross-validated outcomes (number of days of self-reported stimulant drug use over the 3- to 6-month post-randomization period, the lower the better) against additional sample sizes. The initial sample size was set to 50.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 1

Planned sample size for different combinations of (γ, ϵ) .

ϵ	$\gamma = 0.5$	$\gamma = 1$	$\gamma = 1.5$	$\gamma = 2$	$\gamma = 2.5$
1.7	165	160	155	155	150
1.6	215	190	175	165	160
1.5	280	225	200	180	170