Statistics and Its Interface Volume 16 (2023) 475–491

Confidence in the treatment decision for an individual patient: strategies for sequential assessment

Nina Orwitz*, Thaddeus Tarpey, and Eva Petkova

Evolving medical technologies have motivated the development of treatment decision rules (TDRs) that incorporate complex, costly data (e.g., imaging). In clinical practice, we aim for TDRs to be valuable by reducing unnecessary testing while still identifying the best possible treatment for a patient. Regardless of how well any TDR performs in the target population, there is an associated degree of uncertainty about its optimality for a specific patient. In this paper, we aim to quantify, via a confidence measure, the uncertainty in a TDR as patient data from sequential procedures accumulate in real-time. We first propose estimating confidence using the distance of a patient's vector of covariates to a treatment decision boundary, with further distances corresponding to higher certainty. We further propose measuring confidence through the conditional probabilities of ultimately (with all possible information available) being assigned a particular treatment, given that the same treatment is assigned with the patient's currently available data or given the treatment recommendation made using only the currently available patient data. As patient data accumulate, the treatment decision is updated and confidence reassessed until a sufficiently high confidence level is achieved. We present results from simulation studies and illustrate the methods using a motivating example from a depression clinical trial. Recommendations for practical use of the measures are proposed.

Keywords and phrases: Precision medicine, Treatment decision rules.

1. INTRODUCTION

Randomized controlled trials (RCTs) compare treatments on average in the target population. Treatment decisions based on average results from RCTs are a one-size-fits-all approach that do not address the goal of personalizing treatment assignments for patients. It has long been acknowledged that there is typically a wide heterogeneity in the response to a particular treatment. Even when one treatment is better than another on average, there are patients that would benefit more from the inferior (on average) treatment than from the superior one.

This has lead to heightened interest and a recent surge of developments in precision medicine, where a treatment decision for an individual is made based on the patient's baseline data, which might include characteristics such as clinical, demographic and biological features. The precision medicine approach continues to grow in popularity as technological advances allow for acquisition of complex biological data for deep patient phenotyping, which can potentially lead to the development of new therapeutics, more treatment options, and the need for optimal treatment selection. Patient-specific treatment decisions can be made by developing *treatment decision rules* (TDRs), first formalized by [15] and [17], using patient characteristics. As more detailed patient information becomes available with advances in technology, TDRs can become more accurate in selecting the best treatment for individual patients.

Performance of TDRs is commonly assessed by estimating their "value", i.e., the expected value of the outcome if every patient receives the treatment prescribed by the TDR. Specifically, a TDR $d(\boldsymbol{X})$ maps a p-dimensional vector of patient baseline covariates $\boldsymbol{X} = (X_1, \ldots, X_p)' \in \mathbb{R}^p$ to a treatment indicator A. In this paper, we consider the common case of a discrete treatment indicator A with only two possible treatment options $A = \{1, 2\}$. If Y denotes the outcome, the value of a decision rule d is defined as [16]:

$$V\big(d(\boldsymbol{X})\big) = E\big(E\big(Y|\boldsymbol{X}, A = d(\boldsymbol{X})\big)\big). \tag{1}$$

If (without loss of generality) a higher outcome of Y indicates an improvement, then TDRs with higher values are preferred as they lead to better outcomes across the entire target population. With more patient information available, TDRs are expected to improve and therefore have higher values than TDRs employing fewer patient characteristics. At the same time, it may be expected that optimal TDRs would require more assessments to be made on patients. This would elevate the cost and patient burden, both potential barriers for implementation of these TDRs in clinical practice.

Suppose an optimal TDR requires a large number of tests or procedures to be performed on the patients, for exam-

*Corresponding author.

ple, all or a subset of the following: demographics, medical history, clinical evaluation, lab tests, neuropsychological testing, structural and functional imaging. Each of these procedures or tests provides a set of patient characteristics from different data modalities: for example, m_1 demographic data; m_2 variables assessing medical history; m_3 variables from a clinical data modality; ...; m_K variables from the last procedure, say a functional brain imaging under a specific task, i.e., an imaging data modality. Since the K procedures above are not performed simultaneously, without loss of generality assume that they are performed in some specific order, $1:K$. The order may be determined by ranking or classifying the procedures by criteria such as ease of administration, cost, length of time, and relative contribution to the TDR. Let us also assume that TDRs based on data from an expanding number of procedures satisfy an assumption of increasing values. More precisely, if the k-th procedure generates a vector of m_k patient characteristics $\boldsymbol{X}_k$, and we let $d_k = d_k(\boldsymbol{X}_1, \ldots, \boldsymbol{X}_k)$ be the optimal TDR based on patient characteristics obtained from the first k procedures, then we assume $V(d_k) \leq V(d_{k+1})$ for all $k = 1, \ldots, K-1$.

While TDRs are constructed to perform well on average across the target patient population (see (1)), inherently decisions for individual patients will have different associated levels of uncertainty depending on two major factors: (i) the patients' specific baseline characteristics used for making a treatment recommendation and (ii) the uncertainty due to estimating the corresponding TDR based on available training data. Even if the TDRs are well estimated and validated (e.g., when a large training sample size is available), uncertainty in the patient's treatment recommendation may persist if the currently available baseline measures lack sufficient information to guide on an optimal treatment decision. This paper focuses primarily on the first component of uncertainty by deriving analytical expressions of distances to decision boundaries and probabilities of correct decisions (and the corresponding level of confidence) given a patient's currently available baseline predictors; the uncertainty related to estimating TDRs is handled in the simulation studies in Section 5.

Each TDR establishes a boundary in a potentially high-dimensional space of baseline covariates. If a vector of patient's baseline values, denoted $\boldsymbol{x}^{new}$, falls on one side of the boundary, one of the two treatments is recommended, while if $\boldsymbol{x}^{new}$ falls on the other side of the boundary, the other treatment is recommended. Intuitively, the further a patient's $\boldsymbol{x}^{new}$ lies from the decision boundary, the higher the confidence in a treatment decision.

For a new patient, it is of interest to quantify this certainty, or confidence, in the patient's individual treatment decision. Suppose d_k is the TDR, based on the first k baseline modalities, that has been sufficiently tested and validated for implementation in practice. In this setting, let

$$d_k^{new} = d_k(\boldsymbol{x}_1^{new}, \ldots, \boldsymbol{x}_k^{new}) \tag{2}$$

denote the treatment decision after k data modalities have been collected. When data from the next test or procedure are obtained, d_k^{new} is updated to d_{k+1}^{new} and, respectively, the measures of confidence in a patient's treatment decision are reassessed. When an acceptable level of confidence is achieved, the patient's treatment decision would be considered nearly optimal and we would recommend the patient begin this treatment, rather than undergo more testing. This would allow patients to avoid unnecessary or costly procedures while receiving the optimal treatment decision as early in the examination process as possible.

In this paper, we present novel strategies for estimating confidence in a patient's treatment decision based on a sequential collection of baseline data from different tests or procedures. In Section 2, we set up notations and describe the general framework for sequential optimal treatment decision making. We consider distance as a confidence measure such that if a patient's baseline measures $\boldsymbol{x}_1^{new}, \ldots, \boldsymbol{x}_k^{new}$ are close to a treatment decision boundary, we intuitively would not have high confidence in which treatment is optimal. In order to study the utility of distance as a confidence measure, in Section 3 we present theoretical results on distance to a decision boundary, a concept borrowed from the support vector machine literature [14]. Later, in Section 5.3, the distance of a patient's data to the decision boundary is used to determine confidence in a treatment decision based on a subset of predictors. In Section 4, a conditional probability approach to estimating confidence is presented. We derive two probabilities: 1) the probability that the ultimate treatment decision made with all available covariates agrees with the current decision made using a subset of covariates, i.e.,

$$P(d_K^{new} = a | d_k^{new} = a), \quad a = 1, 2, \text{ for } k < K$$

and 2) the probability that the ultimate treatment decision made with all covariates is a particular treatment given the patient's currently observed covariates, i.e.,

$$P(d_K^{new} = a | \boldsymbol{x}_1^{new}, \ldots, \boldsymbol{x}_k^{new}), \quad a = 1, 2, \text{ for } k < K.$$

In Section 5 we present results from simulation studies of the confidence measures, and evaluate cutoff thresholds for assessing if a patient has achieved a desired level of high confidence. An illustrative example using a depression clinical trial is presented in Section 6. The limitations and implications of these findings are discussed in Section 7.

2. A FRAMEWORK FOR SEQUENTIAL TREATMENT DECISIONS

We begin with a description of a general sequential procedure for estimating TDRs and subsequently estimating confidence in an individual patient's treatment decision based on a TDR. Assume it is possible to collect p total covariates from K separate data modalities, with each procedure generating m_k covariates such that $p = \sum_{k=1}^{K} m_k$. Suppose

data on these p covariates exist from a RCT from which K sequential TDRs $d_1, \ldots, d_K$ have been developed, satisfying $V(d_k) \leq V(d_{k+1})$ for $k = 1, \ldots, K-1$.

To illustrate the sequential procedure and notation, consider the case with $K = 2$ data modalities consisting of 2 sets of covariates $\boldsymbol{X}_1$ and $\boldsymbol{X}_2$, with p_1 and $p - p_1$ elements respectively. Write $\boldsymbol{X} = (\boldsymbol{X}_1', \boldsymbol{X}_2')'$ and correspondingly, partition the mean and covariance matrices as

$$\boldsymbol{\mu} = \begin{pmatrix} \boldsymbol{\mu}_1 \\ \boldsymbol{\mu}_2 \end{pmatrix}, \quad \boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} \\ \boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22} \end{pmatrix}$$

where $\boldsymbol{\Sigma}_{11} = \text{cov}(\boldsymbol{X}_1)$, $\boldsymbol{\Sigma}_{22} = \text{cov}(\boldsymbol{X}_2)$, and $\boldsymbol{\Sigma}_{12} = \boldsymbol{\Sigma}_{21}'$ is the matrix of covariances between $\boldsymbol{X}_1$ and $\boldsymbol{X}_2$. Let $Y_a \in \mathbb{R}$ denote the potential outcome for a patient receiving treatment a; assume without loss of generality that higher values of Y_a are better. Let $\boldsymbol{\Psi}_{1Y_1}$ and $\boldsymbol{\Psi}_{2Y_2}$ denote the vector of covariances between Y_a and $\boldsymbol{X}_1$ and $\boldsymbol{X}_2$ for $a = 1$ and 2 respectively. We will assume throughout that conditional expectations are linear (which often provides a good approximation) and also consider the case with multivariate normal data (where conditional expectations are exactly linear).

Starting with the first set of covariates $\boldsymbol{X}_1$, let us define the corresponding TDR d_1 using a linear relation

$$\boldsymbol{y}_a = \mu_{y_a} + \boldsymbol{\beta}_{a1}'(\boldsymbol{X}_1 - \boldsymbol{\mu}_1) + \epsilon_a. \tag{3}$$

Under this model, the decision rule for a new patient with $\boldsymbol{x}_1^{new}$ covariates is to assign treatment 1 if $\mu_{y_1} + \boldsymbol{\beta}_{11}'(\boldsymbol{x}_1^{new} - \boldsymbol{\mu}_1) > \boldsymbol{\mu}_{y_2} + \boldsymbol{\beta}_{21}'(\boldsymbol{x}_1^{new} - \boldsymbol{\mu}_1)$; otherwise, the patient is assigned to treatment 2. This is the optimal TDR when only the covariates $\boldsymbol{X}_1$ are available. To make this decision rule operational, we can replace the parameters by their estimates.

Next, if $\boldsymbol{X}_2$ is available, we can incorporate the second set of covariates into the estimation of an updated TDR containing all p covariates. The linear decision models, similar to above, are now derived from

$$\boldsymbol{y}_a = \mu_{y_a} + \boldsymbol{\alpha}_{a1}'(\boldsymbol{X}_1 - \boldsymbol{\mu}_1) + \boldsymbol{\alpha}_{a2}'(\boldsymbol{X}_2 - \boldsymbol{\mu}_2) + \nu_a. \tag{4}$$

The updated decision rule assigns treatment 1 to a new patient with $(\boldsymbol{x}_1^{new'}, \boldsymbol{x}_2^{new'})'$ covariates if $\mu_{y_1} + \boldsymbol{\alpha}_{11}'(\boldsymbol{x}_1^{new} - \boldsymbol{\mu}_1) + \boldsymbol{\alpha}_{12}'(\boldsymbol{x}_2^{new} - \boldsymbol{\mu}_2) > \mu_{y_2} + \boldsymbol{\alpha}_{21}'(\boldsymbol{x}_1^{new} - \boldsymbol{\mu}_1) + \boldsymbol{\alpha}_{22}'(\boldsymbol{x}_2^{new} - \boldsymbol{\mu}_2)$, and assigns treatment 2 otherwise.

This illustration with two consecutively assessed sets of covariates can be easily generalized to a setting with K total data modalities, in which case the modalities, or sets of covariates, are incorporated into TDR estimation sequentially. Table 1 provides an illustration of a rich data source with n total subjects, where predictors comprising $\boldsymbol{X}$ can be divided into K data modalities, or sets of covariates, such that $m_1 + m_2 + \cdots + m_K = p$. From the illustration in Table 1, K sequential TDRs can be developed. If we assume the sequential TDRs (indicated below the table) have been previously developed, tested, and validated, for a new patient with only $\boldsymbol{x}_1^{new} \in \mathbb{R}^{m_1}$ measured we can apply the TDR d_1 to obtain the patient's treatment decision based on the currently available patient information. We now aim to determine how confident we are in the treatment decision based on $\boldsymbol{x}_1^{new}$. To evaluate the decision made for this new patient, we propose measures of confidence, or certainty, in Sections 3 and 4. The measures of confidence are computed for a patient sequentially, incorporating each successive test or procedure as more patient data are collected, until a desired level of confidence is achieved.

3. DISTANCE TO A DECISION BOUNDARY

For a given TDR d, a patient's baseline characteristics $\boldsymbol{x}^{new}$ will lie some distance from the decision boundary separating recommended treatment classes. This is a concept related to the support vector machine literature [14], where the classification method uses the distance to the decision boundary to determine margins of the classifier. It is intuitively expected that there will be less certainty, or confidence, in decisions for patients whose data fall on or near the decision boundary. Greater certainty, or confidence, is expected in the decisions when a patient's vector of covariates lies farther away from the decision boundary. This is due to the notion that the estimated boundary could shift somewhat as a result of minor perturbations to the data, for example random sampling. If a patient's data lie close to the boundary, the side of the boundary on which the data lies could switch if the data are perturbed slightly. This corresponds to a switch in the assigned treatment, and therefore it is important to quantify distance in order to provide confidence in a patient's treatment decision. We consider Euclidean and (more generally) Mahalanobis distance metrics.

3.1 General equations for distance to a linear decision boundary

Let us consider the case in which we have a linear decision boundary, which can be viewed as a hyperplane. Denote a hyperplane as

$$H = \{\boldsymbol{x} \in \mathbb{R}^p : \boldsymbol{w}'\boldsymbol{x} + b = 0\},$$

with $b = 0$ indicating a hyperplane that runs through the origin. As derived in [14], the Euclidean distance from a new point $\boldsymbol{x}^{new} \in \mathbb{R}^p$ to a linear decision boundary H is expressed as

$$\delta_E(\boldsymbol{x}^{new}) = \frac{|\boldsymbol{w}'\boldsymbol{x}^{new} + b|}{\sqrt{\boldsymbol{w}'\boldsymbol{w}}}. \tag{5}$$

One criticism of using Euclidean distance is its inability to capture relationships between variables [3]. To address this shortcoming, Mahalanobis distance, a generalization of Euclidean distance, from a point $\boldsymbol{x}^{new}$ to a hyperplane can be used and is expressed as

$$\delta_M(\boldsymbol{x}^{new}) = \frac{|\boldsymbol{w}'\boldsymbol{x}^{new} + b|}{\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}^{-1}\boldsymbol{w}}},$$

Table 1. Sequential TDR Framework Example

i	$\mathbf{Y}$	$\mathbf{A}$	$\mathbf{X}$			
			$\mathbf{X}_1$	$\mathbf{X}_2$	$\cdots$	$\mathbf{X}_K$
1	Y_1	A_1	$X_{1,1}\cdots X_{1,m_1}$	$X_{1,m_1+1}\cdots X_{1,m_1+m_2}$	$\cdots$	$X_{1,p-m_K+1}\cdots X_{1,p}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
n	Y_n	A_n	$X_{n,1}\cdots X_{n,m_1}$	$X_{n,m_1+1}\cdots X_{n,m_1+m_2}$	$\cdots$	$X_{n,p-m_K+1}\cdots X_{n,p}$

Set 1 $\underbrace{\qquad}_{d_1}$

Sets 1 & 2 $\underbrace{\qquad\qquad}_{d_2}$

$\vdots$

Sets 1,2,...,K $\underbrace{\qquad\qquad\qquad}_{d_K}$

where $\mathbf{\Sigma}$ denotes the covariance matrix of the predictor variables and is assumed to be positive definite. (Note that when $\mathbf{\Sigma} = \sigma^2\mathbb{I}_p$, i.e., when the predictor variables are uncorrelated with one another, each with a common variance, the Mahalanobis distance corresponds to the usual Euclidean distance.)

3.2 Distribution of distances to a linear decision boundary

Our goal is to determine the utility of distance to a decision boundary as a measure of confidence in a patient's treatment decision. In our framework, the sequential TDRs d_k for $k = 1, \ldots, K$ have boundaries based on an increasing number of predictors, with the $((m_1 + \cdots + m_k) - 1)$-dimensional hyperplane containing all predictors in d_k. By studying the distribution of distances to a linear decision boundary, we can study the relationship between distances to boundaries based on a few covariates to distances to boundaries when incorporating additional covariates.

While previous literature describes the behavior of distance between two points in higher dimensions [e.g., 1], there is a paucity of investigations of distances from a point to a hyperplane decision boundary. In the sequential framework, it will be important to understand the behavior of Euclidean and Mahalanobis distances to the decision boundary with increasing number of covariates.

First, the distribution of the distances to a hyperplane with multivariate normal data is derived. Let $\boldsymbol{X} \sim \mathcal{MVN}(\boldsymbol{\mu}, \mathbf{\Sigma})$. Since linear transformations of multivariate normal random vectors are normal, the numerator of (5) follows a normal distribution, $\boldsymbol{w}'\boldsymbol{X} + b \sim \mathcal{N}(\boldsymbol{w}'\boldsymbol{\mu} + b, \boldsymbol{w}'\mathbf{\Sigma}\boldsymbol{w})$, where we denote $\mu = \boldsymbol{w}'\boldsymbol{\mu} + b$ and $\sigma^2 = \boldsymbol{w}'\mathbf{\Sigma}\boldsymbol{w}$. Denote $R = |\boldsymbol{w}'\boldsymbol{X} + b|$, where taking absolute value by definition means R is now distributed as folded normal. R/σ has a non-central χ distribution with one degree of freedom and non-centrality parameter μ/σ [e.g., see 20, Section 2.1]. Therefore,

$$\frac{|\boldsymbol{w}'\boldsymbol{X} + b|}{\sqrt{\boldsymbol{w}'\mathbf{\Sigma}\boldsymbol{w}}} \sim \chi_1'(\lambda),$$

where $\chi_1'(\lambda)$ denotes a non-central χ distribution with one degree of freedom and non-centrality parameter

$$\lambda = \left(\frac{\boldsymbol{w}'\boldsymbol{\mu} + b}{\sqrt{\boldsymbol{w}'\mathbf{\Sigma}\boldsymbol{w}}}\right).$$

Scaling a non-central χ distribution with one degree of freedom yields a folded normal distribution; by scaling and rearranging, we obtain

$$\begin{aligned}\delta_E(\boldsymbol{X}^{new}) :=& \frac{|\boldsymbol{w}'\boldsymbol{X}^{new} + b|}{\sqrt{\boldsymbol{w}'\boldsymbol{w}}} \\ =& \left(\frac{\sqrt{\boldsymbol{w}'\mathbf{\Sigma}\boldsymbol{w}}}{\sqrt{\boldsymbol{w}'\boldsymbol{w}}}\right)\left(\frac{|\boldsymbol{w}'\boldsymbol{X}^{new} + b|}{\sqrt{\boldsymbol{w}'\mathbf{\Sigma}\boldsymbol{w}}}\right) \\ \sim& \mathcal{FN}(\mu_E, \sigma_E^2),\end{aligned}$$

where $\mathcal{FN}$ denotes the folded normal distribution with parameters

$$\mu_E = \frac{\boldsymbol{w}'\boldsymbol{\mu} + b}{\sqrt{\boldsymbol{w}'\boldsymbol{w}}} \quad , \quad \sigma_E = \frac{\sqrt{\boldsymbol{w}'\mathbf{\Sigma}\boldsymbol{w}}}{\sqrt{\boldsymbol{w}'\boldsymbol{w}}}. \tag{6}$$

The expected value for the folded normal distribution is provided in [20, Equation 7]. For our application, the expected Euclidean distance from a point $\boldsymbol{X}^{new}$ to a $(p-1)$-dimensional hyperplane is

$$\begin{aligned}&\mathbb{E}[\delta_E(\boldsymbol{X}^{new})] \\ &\approx \sigma_E\sqrt{\frac{2}{\pi}}exp\left[-\left(\frac{\mu_E^2}{2\sigma_E^2}\right)\right] + \mu_E\left[1 - 2\Phi\left(-\frac{\mu_E}{\sigma_E}\right)\right]\end{aligned} \tag{7}$$

where $\Phi(\cdot)$ denotes the CDF of the standard normal distribution. For Mahalanobis distance, following a similar derivation, we can express

$$\delta_M(\boldsymbol{X}^{new}) := \frac{|\boldsymbol{w}'\boldsymbol{X}^{new} + b|}{\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}^{-1}\boldsymbol{w}}}$$
$$= \left(\frac{\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}}}{\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}^{-1}\boldsymbol{w}}}\right)\left(\frac{|\boldsymbol{w}'\boldsymbol{X}^{new} + b|}{\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}}}\right) \sim \mathcal{FN}(\mu_M, \sigma_M^2),$$

with

$$(8) \qquad \mu_M = \frac{\boldsymbol{w}'\boldsymbol{\mu} + b}{\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}^{-1}\boldsymbol{w}}} \quad , \quad \sigma_M = \frac{\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}}}{\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}^{-1}\boldsymbol{w}}}.$$

The expected value of the Mahalanobis distance from a point $\boldsymbol{X}^{new}$ to a $(p-1)$-dimensional hyperplane is

$$(9) \quad \mathbb{E}[\delta_M(\boldsymbol{X}^{new})] \approx \sigma_M\sqrt{\frac{2}{\pi}}exp\left[-\left(\frac{\mu_M^2}{2\sigma_M^2}\right)\right] + \mu_M\left[1 - 2\Phi\left(-\frac{\mu_M}{\sigma_M}\right)\right].$$

In the special case of multivariate standard normal with $\boldsymbol{\Sigma} = \mathbb{I}_p$ and $\boldsymbol{\mu} = 0$, $\delta_E(\boldsymbol{X}^{new}) = \delta_M(\boldsymbol{X}^{new})$. If additionally the hyperplane runs through the origin ($b = 0$), further simplifications of (7) and (9) are possible, since this leads to $\mu_E = \mu_M = 0$.

When $\boldsymbol{\Sigma} = \mathbb{I}_p$, $\sigma_E = \sigma_M = 1$. In this case, $\delta_E(\boldsymbol{X}^{new}) = \delta_M(\boldsymbol{X}^{new})$, with both expected values following a half-normal distribution, or equivalently a χ distribution with one degree of freedom. The expected values in particular reduce to $\mathbb{E}[\delta_E(\boldsymbol{X}^{new})] = \mathbb{E}[\delta_M(\boldsymbol{X}^{new})] \approx \sqrt{2/\pi}$.

3.3 A simulation illustration on distances to a decision boundary

This section presents a simulation to illustrate the behavior of these distances to the decision boundary in high dimensions. In these simulations, data of sample size $n = 100$ with $p = 100$ covariates were generated from a multivariate normal distribution with zero means and a covariance matrix assumed to be in the form of a correlation matrix with all pairwise correlations equal to ρ, i.e.,

$$\boldsymbol{\Sigma} = \rho\mathbf{1}\mathbf{1}' + (1-\rho)\mathbf{I},$$

with $\mathbf{1}' = (1, 1, \ldots, 1)$. We will in particular study correlated data with $\rho = 0.5$ to compare differences between Euclidean and Mahalanobis distance metrics. From this simulated data, we randomly generated decision boundary hyperplanes by selecting a $(p-1)$-dimensional hyperplane $H(p)$ in $\mathbb{R}^p$ for each $p = 3, \ldots, 100$. This is done using the `OjaNP` package in `R` [5], which computes a $(p-1)$ dimensional hyperplane passing through p points in $\mathbb{R}^p$. The expected distances to the $(p-1)$-dimensional hyperplanes are computed using (7) and (9). The data and random hyperplane generation are repeated for $M = 50$ repetitions and results for the expected distances are averaged.

Figure 1 depicts the average distances to the hyperplanes as a function of the number of p dimensions. The Mahalanobis distances, which are scaled versions of the Euclidean distances, are smaller in magnitude than their Euclidean distance counterparts. The average distances level out to a constant limiting value; in the Euclidean distance case this would be

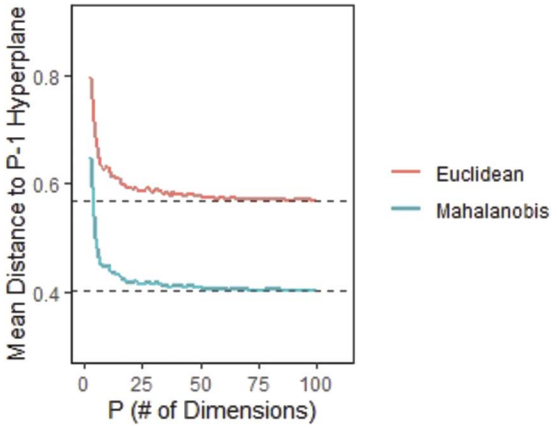

Figure 1. Behavior of distance from a point in $\mathbb{R}^p$ to a $(p-1)$-dimensional hyperplane in high dimensions. The colored curves represent the expected values for the distances averaged across fifty simulations. The horizontal dashed lines denote the theoretical limiting values.

$$\lim_{p\to\infty} \mathbb{E}[\delta_E] = \sigma_E\sqrt{\frac{2}{\pi}},$$

with the counterpart for Mahalanobis distance obtained by replacing σ_E with σ_M. The convergence value therefore depends on σ_E or σ_M, which are a function of the covariance matrix and hyperplane coefficient values $\boldsymbol{w}$. This limiting value is equivalent to the mean of a half-normal distribution, therefore indicating that the folded normal distributed distance converges to half-normal.

Specifically in the context of our sequential framework under the decision model in (3), the linear decision boundary based on $p_1 < p$ observed covariates is determined by $\mu_{y_1} - \mu_{y_2} + (\boldsymbol{\beta}_{11} - \boldsymbol{\beta}_{21})'(\boldsymbol{X}_1 - \boldsymbol{\mu}_1) = 0$. The Euclidean distance of a patient's $\boldsymbol{x}_1^{new}$ observed covariates to this decision boundary is

$$(10) \quad \delta_E(\boldsymbol{x}_1^{new}) = \frac{|\mu_{y_1} - \mu_{y_2} + (\boldsymbol{\beta}_{11} - \boldsymbol{\beta}_{21})'(\boldsymbol{x}_1^{new} - \boldsymbol{\mu}_1)|}{\sqrt{(\boldsymbol{\beta}_{11} - \boldsymbol{\beta}_{21})'(\boldsymbol{\beta}_{11} - \boldsymbol{\beta}_{21})}}.$$

The corresponding Mahalanobis distance to the decision boundary is

$$(11) \quad \delta_M(\boldsymbol{x}_1^{new}) = \frac{|\mu_{y_1} - \mu_{y_2} + (\boldsymbol{\beta}_{11} - \boldsymbol{\beta}_{21})'(\boldsymbol{x}_1^{new} - \boldsymbol{\mu}_1)|}{\sqrt{(\boldsymbol{\beta}_{11} - \boldsymbol{\beta}_{21})'\boldsymbol{\Sigma}_{11}^{-1}(\boldsymbol{\beta}_{11} - \boldsymbol{\beta}_{21})}}.$$

To use these distances to measure confidence, based on the results from the simulation studies, we will need to account for the number of observed covariates p_1 and the maximum number of covariates p that can be included. An additional factor impacting the distance are the correlations

between covariates, which may vary within and between data modalities. This poses a challenge in our sequential setting since we will need to determine what magnitude a patient's distance to the decision boundary should be in order to achieve high confidence in the treatment decision. The plateau seen in Figure 1 that is achieved with p_1 much smaller than p indicates that it may be possible to achieve high confidence for a small number of modalities when confidence is based on distance. The possible use of distance as a confidence measure will be further studied in the simulations in Section 5.3.

4. A CONDITIONAL PROBABILITY APPROACH

Let us define the ultimate treatment decision for a patient as the decision from d_K, the TDR obtained when using all possible data modalities for the patient. Using the framework described in Section 2, we derive conditional probabilities for the ultimate treatment decision for a patient given the treatment decision based on a patient's currently available data, or, alternatively, given the patient's currently available data from $k < K$ tests or procedures.

Assume sequential TDRs $d_1, \ldots, d_K$ have been previously estimated and validated, and assume a linear conditional expectation as done previously (which will hold in the case of elliptical distributions).

4.1 Reparameterization of a TDR based on observed patient characteristics

We start by deriving a reparameterized version of the TDR d_1, the decision rule derived from (3), based on the available patient covariates $p_1 < p$. If $\boldsymbol{X} \sim \mathcal{MVN}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, then we can reparameterize the TDR derived from (3) using the well-known property that a linear combination of normally distributed random variables is normal. Write

$$(\boldsymbol{\beta}_{11} - \boldsymbol{\beta}_{21})'\boldsymbol{X}_1 \sim \mathcal{N}(\tilde{\mu}_1, \tilde{\sigma}_1^2),$$

where

$$\tilde{\mu}_1 = (\boldsymbol{\beta}_{11} - \boldsymbol{\beta}_{21})'\boldsymbol{\mu}_1$$

$$\tilde{\sigma}_1^2 = (\boldsymbol{\beta}_{11} - \boldsymbol{\beta}_{21})'\boldsymbol{\Sigma}_{11}(\boldsymbol{\beta}_{11} - \boldsymbol{\beta}_{21}). \tag{12}$$

The linear combination is transformed to follow a standard normal distribution, such that

$$Z_{p_1} = \frac{(\boldsymbol{\beta}_{11} - \boldsymbol{\beta}_{21})'(\boldsymbol{X}_1 - \boldsymbol{\mu}_1)}{\tilde{\sigma}_1} \sim \mathcal{N}(0,1). \tag{13}$$

The slope coefficients $\boldsymbol{\beta}_{a1} (a = 1, 2)$ in (13) can be written in terms of population parameters, and (13) can be expressed as

$$Z_{p_1} = \frac{(\boldsymbol{\Psi}_{1,Y_1} - \boldsymbol{\Psi}_{1Y_2})'\boldsymbol{\Sigma}_{11}^{-1}(\boldsymbol{X}_1 - \boldsymbol{\mu}_1)}{\sqrt{(\boldsymbol{\Psi}_{1Y_1} - \boldsymbol{\Psi}_{1Y_2})'\boldsymbol{\Sigma}_{11}^{-1}(\boldsymbol{\Psi}_{1Y_1} - \boldsymbol{\Psi}_{1Y_2})}} \sim \mathcal{N}(0,1). \tag{14}$$

For a new patient with $\boldsymbol{X}_1^{new}$, we assign treatment 1 if $Z_{p_1} > c_{p_1}$ where

$$c_{p_1} = \frac{\mu_{y_2} - \mu_{y_1}}{\sqrt{(\boldsymbol{\Psi}_{1Y_1} - \boldsymbol{\Psi}_{1Y_2})'\boldsymbol{\Sigma}_{11}^{-1}(\boldsymbol{\Psi}_{1Y_1} - \boldsymbol{\Psi}_{1Y_2})}}.$$

To make this operational, parameter estimates can be substituted in for the parameters in this expression, but note that if the dimension is large, some type of regularization will be necessary in the estimation, e.g., if $\boldsymbol{\Sigma}_{11}$ is not of full rank.

4.2 Reparameterization of a TDR based on a full set of patient characteristics

Next, we can similarly reparameterize the TDR with all p covariates. Following the notation in Section 2, the slope coefficients are given by

$$\begin{pmatrix} \boldsymbol{\alpha}_{a1} \\ \boldsymbol{\alpha}_{a2} \end{pmatrix} = \begin{bmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} \\ \boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22} \end{bmatrix}^{-1} \begin{pmatrix} \boldsymbol{\Psi}_{1Y_a} \\ \boldsymbol{\Psi}_{2Y_a} \end{pmatrix}.$$

The full derivation is presented in A. After simplification, the final expression containing the regression coefficients is

$$\begin{pmatrix} \boldsymbol{\alpha}_{a1} \\ \boldsymbol{\alpha}_{a2} \end{pmatrix} = \begin{pmatrix} \boldsymbol{\Sigma}_{11}^{-1}(\boldsymbol{\Psi}_{1,Y_a} - \boldsymbol{\Sigma}_{12}\boldsymbol{\alpha}_{a2}) \\ (\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12})^{-1}(\boldsymbol{\Psi}_{2Y_a} - \boldsymbol{\Sigma}_{21}\boldsymbol{\beta}_{a1}) \end{pmatrix}. \tag{15}$$

Then similarly to the approach in Section 4.1, we can write

$$\begin{pmatrix} \boldsymbol{\alpha}_{11} - \boldsymbol{\alpha}_{21} \\ \boldsymbol{\alpha}_{12} - \boldsymbol{\alpha}_{22} \end{pmatrix}' \begin{pmatrix} \boldsymbol{X}_1 \\ \boldsymbol{X}_2 \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \boldsymbol{\alpha}_{11} - \boldsymbol{\alpha}_{21} \\ \boldsymbol{\alpha}_{12} - \boldsymbol{\alpha}_{22} \end{pmatrix}' \begin{pmatrix} \boldsymbol{\mu}_1 \\ \boldsymbol{\mu}_2 \end{pmatrix}, \begin{pmatrix} \boldsymbol{\alpha}_{11} - \boldsymbol{\alpha}_{21} \\ \boldsymbol{\alpha}_{12} - \boldsymbol{\alpha}_{22} \end{pmatrix}' \begin{pmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} \\ \boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22} \end{pmatrix} \begin{pmatrix} \boldsymbol{\alpha}_{11} - \boldsymbol{\alpha}_{21} \\ \boldsymbol{\alpha}_{12} - \boldsymbol{\alpha}_{22} \end{pmatrix} \right).$$

We express this as

$$(\boldsymbol{\alpha}_{11} - \boldsymbol{\alpha}_{21})'\boldsymbol{X}_1 + (\boldsymbol{\alpha}_{12} - \boldsymbol{\alpha}_{22})'\boldsymbol{X}_2 \sim \mathcal{N}(\tilde{\mu}_2, \tilde{\sigma}_2^2),$$

with

$$\tilde{\mu}_2 = (\boldsymbol{\alpha}_{11} - \boldsymbol{\alpha}_{21})'\boldsymbol{\mu}_1 + (\boldsymbol{\alpha}_{12} - \boldsymbol{\alpha}_{22})'\boldsymbol{\mu}_2$$

$$\begin{aligned} \tilde{\sigma}_2^2 = {} & (\boldsymbol{\alpha}_{11} - \boldsymbol{\alpha}_{21})'\boldsymbol{\Sigma}_{11}(\boldsymbol{\alpha}_{11} - \boldsymbol{\alpha}_{21}) \\ & + (\boldsymbol{\alpha}_{12} - \boldsymbol{\alpha}_{22})'\boldsymbol{\Sigma}_{21}(\boldsymbol{\alpha}_{11} - \boldsymbol{\alpha}_{21}) \\ & + (\boldsymbol{\alpha}_{11} - \boldsymbol{\alpha}_{21})'\boldsymbol{\Sigma}_{12}(\boldsymbol{\alpha}_{12} - \boldsymbol{\alpha}_{22}) \\ & + (\boldsymbol{\alpha}_{12} - \boldsymbol{\alpha}_{22})'\boldsymbol{\Sigma}_{22}(\boldsymbol{\alpha}_{12} - \boldsymbol{\alpha}_{22}). \end{aligned} \tag{16}$$

Transforming this to a standard normal gives

$$Z_p = \frac{(\boldsymbol{\alpha}_{11} - \boldsymbol{\alpha}_{21})'(\boldsymbol{X}_1 - \boldsymbol{\mu}_1) + (\boldsymbol{\alpha}_{12} - \boldsymbol{\alpha}_{22})'(\boldsymbol{X}_2 - \boldsymbol{\mu}_2)}{\tilde{\sigma}_2}, \tag{17}$$

with $Z_p \sim \mathcal{N}(0,1)$. Thus, the TDR would assign treatment 1 to a new patient if $Z_p > c_p$, where

$$c_p = \frac{\mu_{y_2} - \mu_{y_1}}{\tilde{\sigma}_2}.$$

4.3 Joint probability

In this section, we use the joint distribution of Z_{p_1} and Z_p to derive the conditional probability that the treatment decision for a particular patient with all p covariates (ultimate decision) agrees with the current decision based only on a subset of $p_1 < p$ observed covariates. We then can also derive the conditional probability that the treatment decision with all p covariates is $A = a$ $(a = 1, 2)$, given the patient's currently observed $p_1 < p$ covariates.

The expressions for these conditional probabilities follow from the joint probability density function (pdf) of Z_{p_1} from (14) and Z_p from (17). Define $\rho_z = Cor(Z_{p_1}, Z_p) = Cov(Z_{p_1}, Z_p)$. We can write

$$Cov(\boldsymbol{a}^T(\boldsymbol{X}_1 - \boldsymbol{\mu}_1), \boldsymbol{b}^T(\boldsymbol{X}_1 - \boldsymbol{\mu}_1) + \boldsymbol{c}^T(\boldsymbol{X}_2 - \boldsymbol{\mu}_2)) = \boldsymbol{a}^T Var(\boldsymbol{X}_1)\boldsymbol{b} + \boldsymbol{a}^T Cov(\boldsymbol{X}_1, \boldsymbol{X}_2)\boldsymbol{c},$$

such that ρ_z can be expressed as

$$\rho_z = Cov(Z_{p_1}, Z_p) = \boldsymbol{a}^T \boldsymbol{\Sigma}_{11} \boldsymbol{b} + \boldsymbol{a}^T \boldsymbol{\Sigma}_{12} \boldsymbol{c},$$

where, using results for $\tilde{\sigma}_1$ from (12) and $\tilde{\sigma}_2$ from (16), we have

$$\boldsymbol{a} = \frac{(\boldsymbol{\beta}_{11} - \boldsymbol{\beta}_{21})}{\tilde{\sigma}_1} \quad , \quad \boldsymbol{b} = \frac{(\boldsymbol{\alpha}_{11} - \boldsymbol{\alpha}_{21})}{\tilde{\sigma}_2} \quad \text{and} \quad \boldsymbol{c} = \frac{(\boldsymbol{\alpha}_{12} - \boldsymbol{\alpha}_{22})}{\tilde{\sigma}_2}.$$

By properties of the multivariate normal distribution, $(Z_{p_1}, Z_p)'$ is bivariate normal:

$$\begin{pmatrix} Z_{p_1} \\ Z_p \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho_z \\ \rho_z & 1 \end{pmatrix} \right).$$

4.4 Conditional probability of the ultimate treatment decision agreeing with current decision

In the following, the conditional probabilities will be derived in terms of an ultimate treatment decision of $A = 1$, without loss of generality. The conditional probability that a particular patient's treatment decision with all p covariates agrees with the current treatment decision made with a subset of $p_1 < p$ covariates, specifically in the case where the treatment is $A = 1$, is expressed as

(18)
$$P(Z_p > c_p | Z_{p_1} > c_{p_1}) = \frac{\int_{c_p}^{\infty}\int_{c_{p_1}}^{\infty} f_{Z_{p_1},Z_p}(z_{p_1}, z_p) dz_{p_1} dz_p}{\Phi(-c_{p_1})} = \frac{\frac{1}{2\pi\sqrt{1-\rho_z^2}}\int_{c_p}^{\infty}\int_{c_{p_1}}^{\infty} exp\left(-\frac{z_{p_1}^2 - 2\rho_z z_{p_1} z_p + z_p^2}{2(1-\rho_z^2)}\right) dz_{p_1} dz_p}{\Phi(-c_{p_1})}$$

where c_p and c_{p_1} are cutoff values for the corresponding decision boundaries.

In [12] and [4], the authors describe and compare several ways to evaluate the integral in (18). Denote the numerator of (18) as

$$L(c_p, c_{p_1}, \rho_z) = \frac{1}{2\pi\sqrt{1-\rho_z^2}} \int_{c_p}^{\infty}\int_{c_{p_1}}^{\infty} exp\left(-\frac{z_{p_1}^2 - 2\rho_z z_{p_1} z_p + z_p^2}{2(1-\rho_z^2)}\right) dz_{p_1} dz_p.$$

As noted in [4], this reduces to a one-dimensional integral

(19)
$$L(c_p, c_{p_1}, \rho_z) = \frac{1}{2\pi}\int_{cos^{-1}\rho_z}^{\pi} exp\left(-\frac{c_{p_1}^2 - 2\rho_z c_{p_1} c_p + c_p^2}{2sin^2 r}\right) dr,$$

which can be expressed as

(20)
$$L(c_p, c_{p_1}, \rho_z) = \Phi(-c_p)\Phi(-c_{p_1}) + \frac{1}{2\pi}\int_0^{\rho_z} \frac{1}{\sqrt{1-r^2}} exp\left(-\frac{c_{p_1}^2 - 2\rho_z c_{p_1} c_p + c_p^2}{2(1-\rho_z^2)}\right) dr,$$

which can be evaluated by numerical integration. Alternatively, the *pmvnorm* function in the `mvtnorm` package in `R` provides a straightforward computation of the probability from the joint bivariate normal distribution [8]. This probability gives essentially the same result as the numerical integration of (20).

4.5 Conditional probability for a specific ultimate treatment decision, given current patient covariates

An alternative measure of confidence can be obtained using estimates of the probability of a particular ultimate treatment decision with all p covariates given a patient's observed set of p_1 covariates, or $P(d_K^{new} = a | \boldsymbol{x}_1^{new} \ldots \boldsymbol{x}_k^{new})$. To compute this, note that the conditional distribution of Z_p given $Z_{p_1} = z_{p_1}$ is $Z_p | Z_{p_1} = z_{p_1} \sim \mathcal{N}(\rho_z z_{p_1}, 1 - \rho_z^2)$. In the case with $a = 1$, this probability can be expressed as

(21)
$$P(Z_p > c_p | Z_{p_1} = z_{p_1}) = 1 - \int_{-\infty}^{c_p} f_{Z_p|Z_{p_1}}(z_p | z_{p_1}) dz_p = 1 - \Phi\left(\frac{c_p - \rho_z z_{p_1}}{\sqrt{1-\rho_z^2}}\right).$$

Similar results can be obtained for $a = 2$.

5. SIMULATION STUDIES

In the previous sections, potential measures of confidence in sequential treatment decisions based on distances and conditional probabilities for TDRs using baseline patient

characteristics were derived in terms of population parameters. This section investigates the performance of the confidence measures when the TDRs are estimated based on sample data with various sample sizes and number of predictors. TDRs are often developed using data from RCTs with sample sizes that can vary substantially depending on the type of study. For example, many RCTs in psychiatric research have sample sizes in the range of 100-150 patients per treatment arm, while other fields (e.g., cardiovascular disease) more regularly enroll 1000-2000 patients per arm. Fortunately, some recent international efforts have resulted in much larger RCTs with upwards of 5000 patients per arm.

In order to obtain realistic results, data was simulated to mimic a large depression RCT designed to discover biomarkers for treatment response: the EMBARC (Establishing Moderators and Biosignatures of Antidepressant Response in Clinical Care) study [19]. This is a multi-site placebo-controlled RCT designed to study sertraline, an antidepressant given to patients with major depressive disorder (MDD). MDD is a leading cause of disability globally [7]. Treatment options for MDD are often ineffectual, or just marginally effective, with 30-40% of patients having an inadequate response after multiple treatments [11]. The challenges posed in the treatment of depression mean that patients could benefit from the proposed personalized approaches. We use the EMBARC study to demonstrate the utility of the methods developed in the previous sections.

5.1 Overview of the simulation strategy and goals

For illustration, we let $K = 10$ procedures (i.e., modalities) to be performed on the patients, each of which produces 10 covariates $\boldsymbol{X}_k = (X_{10(k-1)+1}, \ldots, X_{10k})'$, for a total of $p = 100$ covariates. We assume that the outcome Y is observed in a RCT with 1:1 randomization to two treatments $a = 1$ and $a = 2$. This simulated data will be used to estimate the sequential TDRs, $d_1, d_2, \ldots, d_{10}$, based on $10, 20, \ldots, 100$ covariates respectively. Test ("new") data will be generated to investigate the following quantities:

1. Across all subjects, the probability that the ultimate treatment decision of $A = 1$ based on the decision d_{10} agrees with the current treatment decision of $A = 1$ based on d_k, for all $k = 1, \ldots, 9$. In other words, we compute $P(d_{10}^{new} = 1 | d_k^{new} = 1)$, using the notation defined in (2).
2. For an individual subject, the probability of their ultimate treatment decision based on d_K being, say $a = 1$, given the patient's covariates from the first $k < K$ tests, i.e., $P(d_{10}^{new} = 1 | \boldsymbol{x}_1^{new}, \ldots, \boldsymbol{x}_k^{new})$. We will refer to this as a "personalized probability". Different cutoffs for this probability will be investigated to judge whether confidence in the treatment decision is sufficiently high and how various cutoffs affect the operating characteristics of the sequential procedure.
3. The relationship between the personalized probabilities and the distances to the treatment decision boundaries as k increases.

The above quantities will be studied as functions of the sample size used to estimate the sequential TDRs $d_1, \ldots, d_K$, and of the error variance of the outcome as a function of all p covariates.

5.2 Data generation

The training sets used to generate the linear decision rules contain covariates $\boldsymbol{X} = (X_1, \ldots, X_{100})' \sim \mathcal{MVN}(\boldsymbol{0}, \boldsymbol{\Sigma})$ where $\boldsymbol{\Sigma}$ is obtained from an estimate using the EMBARC data. Six simulation configurations are studied, combining two outcome generation models with three different sample sizes. The regression coefficient parameters, given by $\boldsymbol{\beta}_a = \boldsymbol{\Sigma}^{-1} \boldsymbol{\Psi}_{\boldsymbol{X} Y_a}$, where $\boldsymbol{\Psi}_{\boldsymbol{X} Y_a} = \text{cov}(\boldsymbol{X}, Y_a)$, are also based on the EMBARC data. The first response generation model does not contain an error term and is given by $Y_a = \boldsymbol{X}\boldsymbol{\beta}_a$; the second model scenario incorporates an error term: $Y_a = \boldsymbol{X}\boldsymbol{\beta}_a + \epsilon_a$; the errors are modeled as $\epsilon_1 \sim \mathcal{N}(0, 0.16^2)$, $\epsilon_2 \sim \mathcal{N}(0, 0.18^2)$, similar to those in the EMBARC data.

Three training set sample sizes were used: $N = 105, 1000$ and 10^5 per arm (the largest sample size is studied with the goal of evaluating the benefit of ideal, albeit unrealistic, data sets for the development of treatment decision rules). The mean difference between the outcome under treatment 2 and treatment 1 is $\mu_{y_2} - \mu_{y_1} = 0.05$. The effect size (Cohen's d) is computed as the difference in the mean outcomes divided by the pooled standard deviations of the groups [2]. In our simulated data, this is estimated as $0.05/0.17 = 0.29$, a small effect size similar to what is often seen in depression trials. An independent test data with 100 subjects is generated with $\boldsymbol{X} = (X_1, \ldots, X_{100})' \sim \mathcal{MVN}(\boldsymbol{0}, \boldsymbol{\Sigma})$. Starting with the first modality (i.e., the first set of covariates), covariates are incorporated sequentially 10-at-a-time into the TDR estimation: $p_1 = 10, 20, 30, 40, 50, 60, 70, 80, 90$. For each p_1 and the remaining $p - p_1$ covariates, $\boldsymbol{\beta}_{a1}$ from (3) and $\boldsymbol{\alpha}_{a1}$ and $\boldsymbol{\alpha}_{a2}$ from (4) are obtained by least squares estimation using the training data. The conditional probabilities defined in (18) and (21), and the distance to the decision boundary based on p_1 observed covariates as defined in (10), are computed. For each combination of data generation (with and without error) and sample size, $M = 100$ training datasets are generated and results are averaged over those M iterations.

5.3 Results

5.3.1 Probability of agreement between the current and the ultimate treatment decisions

In the following, we show simulation results for $P(d_{10}^{new} = 1 | d_k^{new} = 1)$, where $k < K = 10$, which is one example of agreement, the other being $P(d_{10}^{new} = 2 | d_k^{new} = 2)$. Here we only focus on the particular treatment $A = 1$, to keep

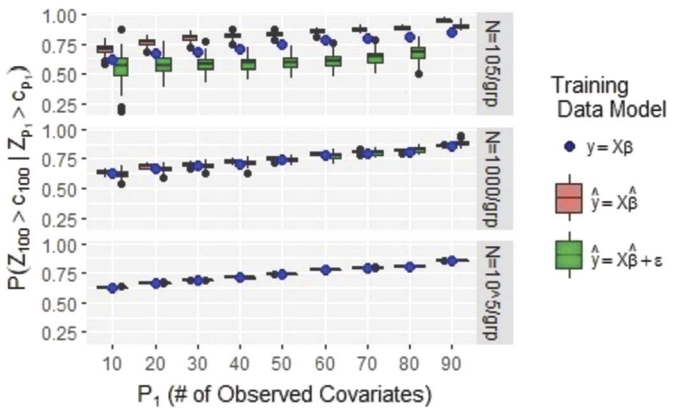

Figure 2. Probability of agreement of ultimate and current treatment assignment of $A = 1$. The blue dots represent the true conditional probabilities, found by using the true coefficients β_{a1} for the sequential TDRs. Z_{100} is the reparameterized TDR for all $p = 100$ covariates and c_{100} is the cutoff value for this decision boundary. The boxplots represent the distribution of the probabilities across the 100 repeated Monte Carlo simulations, where the response is generated without error (pink) and with error (green).

consistent with the results on personalized probabilities in Section 5.3.2.

Figure 2 depicts the probability of an ultimate treatment assignment of $A = 1$ given a current assignment of $A = 1$ based on p_1 covariates. The figure indicates that as more covariates are incorporated into the treatment decision, the probability of the ultimate decision (of $A = 1$) agreeing with the current decision based on the observed covariates increases. There is an approximate linear relationship between the number of predictors measured for a patient and the probability of agreement. With a very small training data set, the variability in the probabilities is greater and less aligned with the true probabilities, indicating that it may be difficult to estimate useful measures of confidence when TDRs are estimated with a small sample size and/or noisy data.

5.3.2 Probability of the ultimate treatment decision being treatment 1, given the currently observed covariates

The personalized conditional probability (21), along with Monte Carlo 95% confidence bands, is depicted for three distinct individuals in Figure 3, and shows the utility of this measure of confidence. The widths of the 95% confidence bands decrease with increasing sample sizes, indicating less variability in the estimated probabilities. In Figure 3, panel A shows the optimal treatment decision for the patient is $A = 1$, whereas in panel B, the optimal treatment decision for the patient is $A = 2$, indicating the ability of our approach to distinguish between two patients who may have very different sequential trajectories. Panel C alternatively

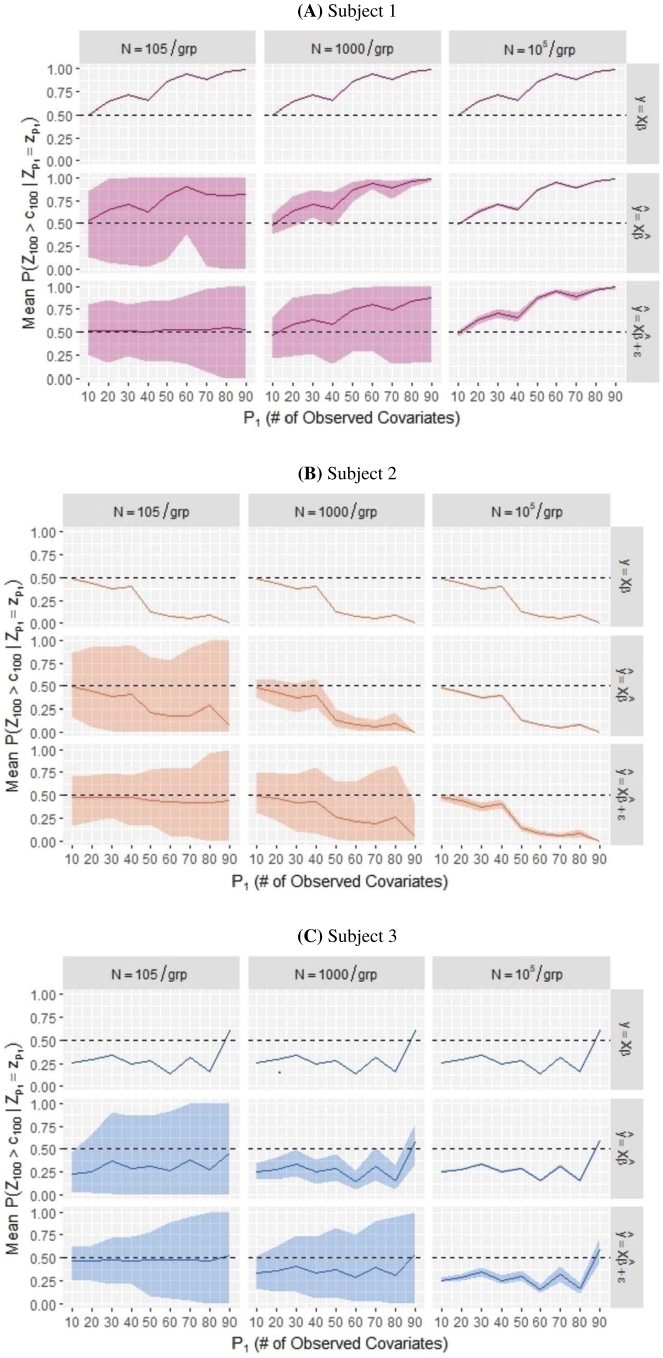

Figure 3. Personalized probabilities for three subjects. Subject 1 has an optimal treatment decision of $A = 1$, subject 2 has an optimal treatment decision of $A = 2$, and subject 3 has no clear optimal treatment. Z_{100} is the reparameterized TDR for all $p = 100$ covariates and c_{100} is the cutoff value for this decision boundary. The mean conditional probability is indicated by the bold line. The bands represent the 95% Monte Carlo confidence intervals (based on the 100 training datasets generated).

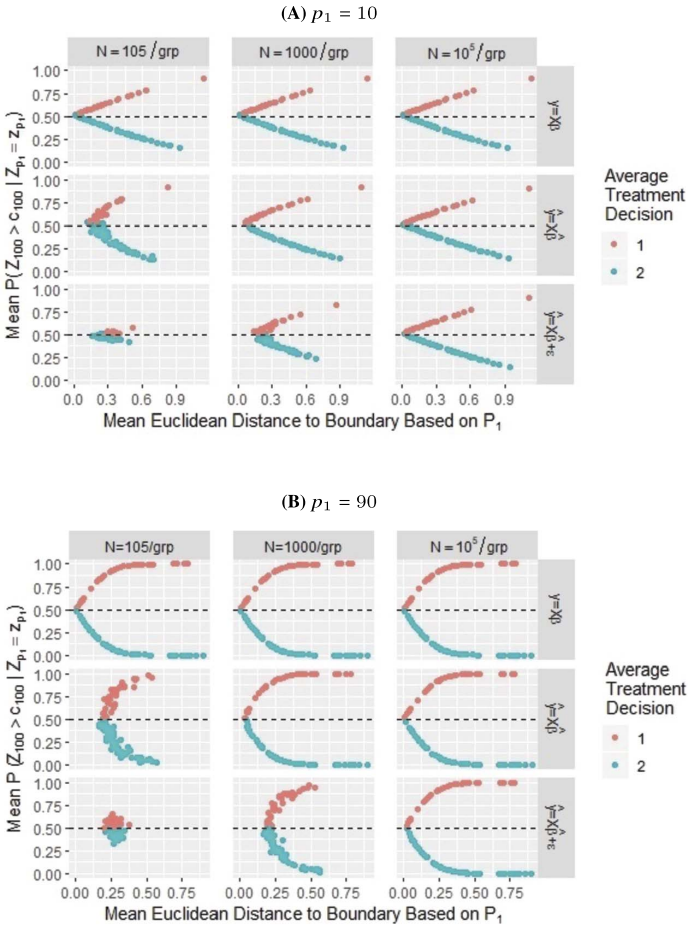

Figure 4. Personalized probability compared to Euclidean distance to the decision boundary based on p_1 covariates. Each point in the plot represents a patient in the test data ($n = 100$). The pink and blue indicators in the plot represent, respectively, the average treatment decision of $A = 1$ and $A = 2$ that were assigned to each patient across the 100 repeated simulations. The horizontal dashed line represents an equal probability of a patient being assigned to either treatment.

shows a patient for whom neither treatment decision appears to be optimal based on the conditional probabilities. For this patient, obtaining additional covariates does not result in greater confidence in either treatment option. In this case, we would be inclined not to recommend either treatment with high confidence based on our method. In practice, this may represent an example where both treatments are similarly effective (or ineffective) for the patient. Accordingly, the cheaper treatment, or one that is beneficial in some other way to the patient, could be considered.

Subjects 1 and 2, depicted in Figure 3, have a tendency on average to be assigned the same treatment as additional covariates are collected and incorporated into the sequential TDRs. However, the figure illustrates variability across patients and, as such, there may be patients with unstable treatment assignments that switch often as more data is obtained. In these cases, the utility of our measure of confidence can be to identify such patients and determine what patient characteristics may be driving the instability of the treatment assignments.

5.3.3 Relationship between conditional probabilities and distances to the decision boundaries

The relationship between the Euclidean distance of the patient's covariates vector to the decision boundary based on $p_1 < p$ covariates and the patient's personalized probability $P(d_K^{new} = 1 | \boldsymbol{x}_1^{new}, \ldots, \boldsymbol{x}_k^{new})$ is shown in Figure 4; $p_1 = 10$ and $p_1 = 90$ are depicted. When this probability is greater than 0.5, the ultimate treatment decision is taken to be $a = 1$; these points on Figure 4 are colored in red, while the points for which the conditional probability is less than 0.5 are colored in blue.

The figure shows that the farther away a patient's data lie from the decision boundary, the greater the conditional probability of being ultimately assigned to treatment 1 or treatment 2, indicated by the probabilities spreading out toward 1 and 0, respectively. As expected, when the probabilities are larger than 0.5, the average ultimate decision

Table 2. Percent correct decisions (PCD) among those that reached the cutoff threshold for high confidence using "personalized probability" of ultimate treatment assignment of $A = 1$ with fewer than p covariates; number of patients (out of 100) achieving cutoff for sufficient confidence in their treatment decision with $p_1 < p$ covariates (N); number of covariates in the TDR when the patient reached the cutoff for high confidence (p_1). The results are averaged across 100 Monte Carlo runs for each of the sample size and outcome generation model combinations.

		Sample Size of Data for Developing the TDRs					
		105 per arm		1000 per arm		10^5 per arm	
Cutoff		$y = X\beta$	$y = X\beta + \epsilon$	$y = X\beta$	$y = X\beta + \epsilon$	$y = X\beta$	$y = X\beta + \epsilon$
0.05/0.95	PCD	0.72	0.68	0.95	0.97	0.97	0.97
	N	97.9	86.0	60.0	62.2	57.0	56.8
	p_1	50.3	85.1	65.8	70.1	67.6	67.9
0.1/0.9	PCD	0.69	0.66	0.90	0.93	0.92	0.92
	N	99.3	90.2	74.8	74.9	70.6	72.1
	p_1	41.0	81.5	59.2	63.3	60.9	61.1
0.2/0.8	PCD	0.66	0.63	0.82	0.85	0.85	0.85
	N	100	95.5	92.6	91.1	91.8	91.1
	p_1	28.6	71.5	45.7	49.5	49.7	49.4

assigned to the patient is $A = 1$, and when less than 0.5, the average ultimate decision is $A = 2$. When the decision rule is poorly estimated (i.e., sample size of 105 per arm), the majority of the personalized probabilities lie around 0.5, as shown in the bottom left panel. When 90% of the covariates are included in the treatment decision, the average probabilities are more extreme towards 0 and 1 than when fewer ($p_1 = 10$) of the covariates are included, indicating greater confidence or certainty when more measures are obtained on patients. We note that, as seen in the simulation results in Section 3, the distance values tend to shrink in magnitude as the number of covariates increases – the distances range from 0 to 1.2 when $p_1 = 10$ (top panel), while the range is from 0 to 1 when $p_1 = 90$ (bottom panel). The results with Mahalanobis distances are depicted in B and show the same probability-distance relationship, with Mahalanobis distances simply a scaled version of a smaller magnitude than the Euclidean distances.

The simulations contain variables with moderate correlations ranging from -0.28 to 0.49, but all variables came from a single data modality, whereas in practice we expect data from multiple modalities. Variables within one data modality (e.g., obtained from the same procedure) are likely to be more highly correlated than variables from different modalities. In C and D, respectively, simulation results are shown with data generated from equal correlation structures within data modalities and varying correlations between data modalities. The results are similar to those presented here, and indicate that variability in the estimated probabilities may be more likely due to differences between data modalities rather than differences among variables within a data modality.

5.4 Personalized sequential confidence assessment scheme

In this section, we operationalize our methods by studying cutoff threshold values for sufficient confidence in a treatment decision such that further acquisition of patient's characteristics is deemed unnecessary. The goal is to provide a cutoff that can hopefully prevent some patients from undergoing unnecessary or costly tests to obtain more data, while still ultimately providing the best possible treatment decision.

Recall the personalized probability $P(d_K^{new} = 1|\boldsymbol{x}_1^{new} \cdots, \boldsymbol{x}_k^{new})$, derived in (21). We evaluate three cutoff schemes for this probability (given by upper and lower thresholds): $\{> 0.95$ or $< 0.05\}$; $\{> 0.9$ or $< 0.1\}$; and $\{> 0.8$ or $< 0.2\}$ under the six data generation settings described previously. Each test subject has an ultimate treatment decision found by applying d_{10}, the TDR with all covariates. For each patient, we determine if the confidence in the patient's treatment decision (i.e., the probability given in (21)) reaches the specific cutoff value with fewer than $p = 100$ covariates, and if so, we obtain the number of observed covariates p_1 when the cutoff is reached. The treatment decision made when reaching the cutoff is compared to the ultimate treatment decision. The *percent correct decision* (PCD) is defined as the proportion of individuals with the same ultimate treatment decision as the treatment assigned with the cutoff threshold value out of the number of patients who reached the cutoff threshold. By using PCD, we aim to evaluate the proportion of patients who will be assigned to their optimal treatment if we control the $P(d_K^{new} = 1|\boldsymbol{x}_1^{new}, \cdots, \boldsymbol{x}_k^{new})$ using various cutoff thresholds. The PCD is then averaged across 100 simulations. We additionally extract the average number of subjects that ever reach the confidence cutoff and the average p_1 covariates in the decision when patients reach the cutoff.

The results are depicted in Table 2 and indicate the overall ability of the personalized probability to align with the true (i.e., optimal) decision. There are fairly high average PCD values even in cases when we relax the cutoff values to be less stringent. There are few differences in PCD when responses are generated with and without error, indicating that perhaps moderate noise in the data will not greatly impact the performance of the method. As an example, when we choose the cutoff for high confidence to be $\{> 0.9 \text{ or } < 0.1\}$ with 1000 patients per arm in the training data set and error in the outcome generation, on average 85% of subjects achieving a cutoff threshold have the same decision with their observed data as they would have had if all possible covariates were obtained and used to make the treatment decision. On average approximately 91 subjects (out of 100) achieved a cutoff threshold, and did so with approximately half of the covariates observed. This shows the utility of the personalized probability as a measure of confidence by indicating that we may not need all data modalities measured on subjects in order to recommend a treatment decision with confidence.

6. APPLICATION TO DATA FROM A DEPRESSION CLINICAL TRIAL

Here we illustrate a potential use of the sequential treatment decision process in practice, using the EMBARC depression study. We use as a criteria for confidence the personalized probability, or $P(d_K^{new} = 1 | \boldsymbol{x}_1^{new}, \ldots, \boldsymbol{x}_k^{new})$, with cutoff threshold $\{> 0.9 \text{ or } < 0.1\}$. With this example, we demonstrate the utility of measuring confidence in personalized treatment decisions.

Of the 242 subjects in this dataset, 125 were randomized to placebo ($A = 1$) and 117 were randomized to the antidepressant drug sertraline ($A = 2$). Baseline covariates that were considered as potential treatment effect modifiers are from four data modalities and are assessed in the following order: (1) eighteen clinical variables; (2) eight behavioral phenotypes (BP) from neuropsychological testing; (3) three variables obtained from electroencephalography (EEG); and (4) ten variables characterizing brain functioning, obtained from functional magnetic resonance imaging (fMRI). The treatment outcome Y was the improvement in symptom severity, measured by the Hamilton Rating Scale for Depression, from week 0 (baseline) to week 8, where positive values of Y indicate improvement from baseline and larger values are preferred. The clinical variables are a mixture of binary (0/1), integer, and continuous variables, while all other data modalities contain continuous variables only.

Using these data, we construct the following TDRs with multiple linear regression models: a) d_1: TDR based on clinical covariates, with (estimated) value $\hat{V}(d_1) = 8.57$; b) d_2: TDR based on clinical + BP covariates, with value $\hat{V}(d_2) = 8.92$, c) d_3: TDR based on clinical + BP + EEG covariates, with $\hat{V}(d_3) = 9.17$; and d) d_4: TDR based on all clinical + BP + EEG + fMRI covariates, $\hat{V}(d_4) = 9.44$. Estimates of value, see (1), of the TDRs are obtained by computing the average of the outcomes of patients whose assigned treatment in the RCT coincides with the treatment recommended by the TDR. We consider d_4 to be the ultimate optimal decision, while d_1, d_2 and d_3 are the linear TDRs when only the respective data modalities are available. The values agree with our expectation that value increases with increasing number of procedures (i.e., patient characteristics), here $p = 39$ total variables. A new patient would undergo the following procedure:

1. Obtain the clinical variables. Estimate the probability that the ultimate the treatment decision will be $A = 2$ (antidepressant) given the patient's set of *clinical* covariates, $P(d_4 = \text{drug} | \boldsymbol{x}_{clinical})$.
   - If $P(d_4 = \text{drug} | \boldsymbol{x}_{clinical}) > 0.9$, recommend the antidepressant treatment; if $P(d_4 = \text{drug} | \boldsymbol{x}_{clinical}) < 0.1$, recommend the placebo treatment.
   - If the probability $P(d_4 = \text{drug} | \boldsymbol{x}_{clinical})$ does not achieve one of the thresholds $\{> 0.9 \text{ or } < 0.1\}$, continue to step 2.
2. Obtain the "BP" variables from neuropsychological testing. Estimate the probability of an ultimate treatment assignment of the antidepressant, given the patient's set of *clinical* and *BP* covariates, or $P(d_4 = \text{drug} | \boldsymbol{x}_{clinical}, \boldsymbol{x}_{bp})$.
   - If the probability achieves one of the cutoff thresholds $\{> 0.9 \text{ or } < 0.1\}$, recommend treatment 1 or treatment 2, respectively, as above.
   - If the probability does not achieve the threshold, continue to step 3.
3. Obtain variables from EEG testing. Estimate the probability of an ultimate treatment assignment of the antidepressant, given the patient's set of *clinical*, *BP* and *EEG* covariates, or $P(d_4 = \text{drug} | \boldsymbol{x}_{clinical}, \boldsymbol{x}_{bp}, \boldsymbol{x}_{eeg})$.
   - If the probability achieves one of the cutoff thresholds $\{> 0.9 \text{ or } < 0.1\}$, recommend treatment 1 or treatment 2, respectively.
   - If the probability does not achieve the threshold, continue to step 4.
4. Obtain variables from fMRI testing and recommend treatment based on d_4 using all four modalities: *clinical*, *BP*, *EEG* and *fMRI*.

Table 3 depicts the sequential confidence procedure for three subjects. Subject A achieved a threshold of high confidence with just the first clinical modality observed and Subject B achieved a high confidence threshold with 3 of the 4 modalities observed (i.e., with fewer than the total number of modalities). Subjects A and B were assigned to the same treatment with our sequential confidence approach

Table 3. Personalized probabilities of an ultimate treatment assignment of $A = 2$ (antidepressant) given individual subject data using the sequential procedure for assessing confidence. A bold value indicates the point at which high confidence was achieved. The treatment assignment using sequential confidence corresponds to the treatment assigned if the subject achieved the threshold. The treatment assigned using the optimal TDR is the assignment including patient's data from all four data modalities.

	Data Modalities in the Sequential TDRs			Treatment Assignment	
Subject	Clinical	Clinical BP	Clinical BP EEG	Using Sequential Confidence	Using Ultimate TDR (d_4)
A	**1**	-	-	Drug	Drug
B	0.28	0.10	**0.09**	Placebo	Placebo
C	0.76	0.61	0.32	-	Placebo

as they would have been if we had obtained all of their data, indicating the utility in minimizing unnecessary acquisition of more complex data modalities. Subject C did not achieve a threshold for high confidence. In cases where high confidence in a decision is not met, the optimal treatment either remains unknown or perhaps the patient would respond similarly to either treatment.

We applied the procedure described above to all subjects in the EMBARC data. Out of the 242 subjects, the percentages of individuals who achieved high confidence, using the $\{> 0.9 \text{ or } < 0.1\}$ cutoff thresholds, with fewer than the maximum number of covariates ($p = 39$) are: 78.5% with clinical variables only, 6.2% with clinical and BP variables, and 7.0% with clinical, BP, and EEG. These percentages sum to 91.7% and thus nearly all subjects achieve high confidence of their treatment decision with information from fewer than 4 procedures. The percentages of subjects whose decision with all four modalities agrees with the decision made if and when the subject achieves high confidence are 75.8% with clinical data, 86.7% with clinical and BP data, and 88.2% with clinical, BP, and EEG data. This reveals that the optimal treatment decision for a subject in this data can, to a large extent, be determined with our approach using a limited set of modalities, and that this treatment decision agrees with the decision if we were to measure all possible data modalities.

7. DISCUSSION

In this paper, we presented novel strategies for assessing confidence, or certainty, in an individual treatment decision. While studied in the context of a depression clinical trial, the methods are general and can be applied in other medical fields.

The sequential procedure introduced assumes that the "ultimate" TDR involves patient characteristics from several (K) tests or procedures each giving rise to variables from different data modalities. We assume throughout this paper that the order of the K modalities is based on the order in which the data would be collected on a typical patient, and is selected with the guidance of a clinician. However, an interesting avenue for future research would be to determine an optimal sequential ordering of the modalities to add to the TDR to maximize value. One approach is to perform a forward stepwise procedure that begins with estimating TDRs for each of the K modalities individually, and selecting the modality, denoted as m_{1^*}, associated with the largest value. Next, TDRs containing m_{1^*} and each remaining modality (i.e. $m_{1^*} + m_k, k \neq 1^*$) will be estimated and the modality that leads to the greatest increase in value will be selected for inclusion into the TDR. Then this procedure is repeated until all modalities are entered into the TDR. If we let $V_1^*, V_2^*, \ldots, V_K^*$ denote the corresponding values for the TDRs in this sequence, then a threshold p^* could then be chosen to determine the number of modalities to add to the ultimate TDR, based on, for example $V_k^*/V_K^* > p^*$. An alternative approach to consider is to base the order of adding modalities on maximizing the probability of agreement (18) of an ultimate treatment assignment given data using all possible subsets of modalities (which is usually computationally feasible since K tends to be small in practice).

If only a TDR based on one modality is available, for example with just a set of EEG measurements, the probability approach will no longer be relevant, since all necessary variables can be obtained from one test. Instead, using distance to the decision boundary will allow for an absolute measure of certainty. In our simulation studies, on average, distances are smaller with more covariates in the TDR or with increasing correlations between variables. Determining high confidence in a patient's treatment decision using distance will require cutoff values that depend on p_1, the number of observed covariates. Future studies should focus on exploring how far away a patient's data lie from a decision boundary in order for given level of confidence to be reached. Additionally, the current use of the Mahalanobis distance results in only a scaled version of the Euclidean distance as noted in Section 3.1. For our working multivariate normal model, the conditional distribution given a subset of baseline covariates impacts the conditional mean only and not the conditional covariance matrix (which is constant). To add greater flexibility, future research could investigate regression and conditional probability approaches where the

baseline covariates may impact both the conditional mean and covariance matrix [e.g., 10].

In the context of decision rules based on linear boundaries, the proposed conditional probabilities are straightforward to interpret, which is especially important for implementation in clinical practice. They have the added benefit of computationally simple closed-form expressions under normality assumptions. In particular, the personalized probability for a new subject, $P(d_K^{new} = a|\boldsymbol{x}_1^{new} \dots \boldsymbol{x}_k^{new})$, is recommended as the most practical approach from a precision medicine perspective, given the ability to sequentially measure confidence in real-time for an individual patient.

The simulation studies indicated that the distance and probability approaches both rely heavily on how well the sequential decision rules are estimated. In cases with large amounts of noise and smaller samples, more rigorous cutoffs for high confidence are needed to offset the weak TDR estimation. The choice of cutoff thresholds also should depend on a clinician's evaluation of the value of gathering additional information, i.e., how do we maximize assigning patients to the optimal treatments while minimizing the number of time-consuming and costly procedures a patient will need. Based on the simulation results and the illustrative EMBARC example, we propose the following scheme for assessing confidence:

1. Apply d_k, the estimated decision based on k modalities, to a new patient with covariates observed from k procedures $\boldsymbol{x}_1^{new} \dots \boldsymbol{x}_k^{new}$, and obtain the treatment decision.
2. Estimate $P(d_K^{new} = a|\boldsymbol{x}_1^{new} \dots \boldsymbol{x}_k^{new})$, using (21) for this new patient.
3. Compare the personalized probability to a cutoff threshold for high confidence. For TDRs estimated with very small samples, use the strictest cutoff threshold ($\{> 0.95, < 0.05\}$); otherwise, choose a reasonable threshold based on the medical implications, consulting with a clinician.

If this probability lies beyond the confidence threshold, we recommend the appropriate treatment based on our confidence measure. Otherwise, we suggest that the patient undergo additional testing and repeat steps 1-3 updating d_k with additional data modalities until high confidence is achieved.

The methods in this paper were derived in the multivariate normal data setting. It would be useful to examine the robustness of our results in cases where multivariate normality does not hold, e.g., for skewed data or a mixture of both continuous and discrete variables. There are many other possible extensions such as deriving corresponding confidence measures when there are more than two treatment options and exploring ways to assess confidence using TDRs derived with other approaches that produce nonlinear decision rule boundaries, such as including nonlinear terms like squares and cross-products in the TDRs producing quadratic decision boundaries in $\boldsymbol{X}$ [9], or an outcome weighted learning (OWL) approach with a Gaussian kernel [21]. We note that determining TDRs is a classification problem and that simple classification rules producing linear decision boundaries often perform competitively compared to more complex approaches in terms of reducing prediction error [6] due to their simplicity (a trade-off of variance for bias). It would also be useful to explore regularization methods for the regression-based linear TDRs such as ridge regression or the lasso [18] when the dimension of the predictors becomes large relative to the sample size.

APPENDIX A. DERIVATION OF EQUATION 15

The complete derivation for (15) follows. Assume $\boldsymbol{\Sigma}_{11}$ is non-singular and the Schur complement of $\boldsymbol{\Sigma}_{11}$ is invertible, where the Schur complement is expressed as $\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12}$. By well-known results on inverting partitioned matrices [e.g., 13, Theorem 2], the matrix $\boldsymbol{\Sigma}$ is invertible and it's inverse can be expressed by

$$\begin{bmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} \\ \boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22} \end{bmatrix}^{-1} = \begin{bmatrix} \boldsymbol{A}_{11} & \boldsymbol{A}_{12} \\ \boldsymbol{A}_{21} & \boldsymbol{A}_{22} \end{bmatrix},$$

where

$$\boldsymbol{A}_{11} = \boldsymbol{\Sigma}_{11}^{-1} + \boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12}(\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12})^{-1}\boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}$$

$$\boldsymbol{A}_{12} = -\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12}(\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12})^{-1}$$

$$\boldsymbol{A}_{21} = -(\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12})^{-1}\boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}$$

$$\boldsymbol{A}_{22} = (\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12})^{-1}.$$

We can write

$$\begin{pmatrix} \boldsymbol{\alpha}_{a1} \\ \boldsymbol{\alpha}_{a2} \end{pmatrix} = \begin{bmatrix} \boldsymbol{A}_{11} & \boldsymbol{A}_{12} \\ \boldsymbol{A}_{21} & \boldsymbol{A}_{22} \end{bmatrix} \begin{pmatrix} \boldsymbol{\Psi}_{1Y_a} \\ \boldsymbol{\Psi}_{2Y_a} \end{pmatrix} = \begin{pmatrix} \boldsymbol{A}_{11}\boldsymbol{\Psi}_{1Y_a} + \boldsymbol{A}_{12}\boldsymbol{\Psi}_{2Y_a} \\ \boldsymbol{A}_{21}\boldsymbol{\Psi}_{1Y_a} + \boldsymbol{A}_{22}\boldsymbol{\Psi}_{2Y_a} \end{pmatrix}.$$

Plugging in the population parameters for $\boldsymbol{A}_{11}$, $\boldsymbol{A}_{12}$, $\boldsymbol{A}_{21}$, and $\boldsymbol{A}_{22}$ and simplifying gives

$$\begin{pmatrix} \boldsymbol{\alpha}_{a1} \\ \boldsymbol{\alpha}_{a2} \end{pmatrix} = \begin{pmatrix} (\boldsymbol{\Sigma}_{11}^{-1} + \boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12}(\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12})^{-1}\boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1})\boldsymbol{\Psi}_{1Y_a} \\ -\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12}(\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12})^{-1}\boldsymbol{\Psi}_{2Y_a} \\ -(\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12})^{-1}\boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Psi}_{1Y_a} \\ +(\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12})^{-1}\boldsymbol{\Psi}_{2Y_a} \end{pmatrix}$$

which can be written as

$$\begin{pmatrix} \boldsymbol{\alpha}_{a1} \\ \boldsymbol{\alpha}_{a2} \end{pmatrix} = \begin{pmatrix} \boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Psi}_{1Y_a} - \\ \boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12}(\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12})^{-1}(\boldsymbol{\Psi}_{2Y_a} - \boldsymbol{\Sigma}_{21}\boldsymbol{\beta}_{a1}) \\ (\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12})^{-1}(\boldsymbol{\Psi}_{2Y_a} - \boldsymbol{\Sigma}_{21}\boldsymbol{\beta}_{a1}) \end{pmatrix}.$$

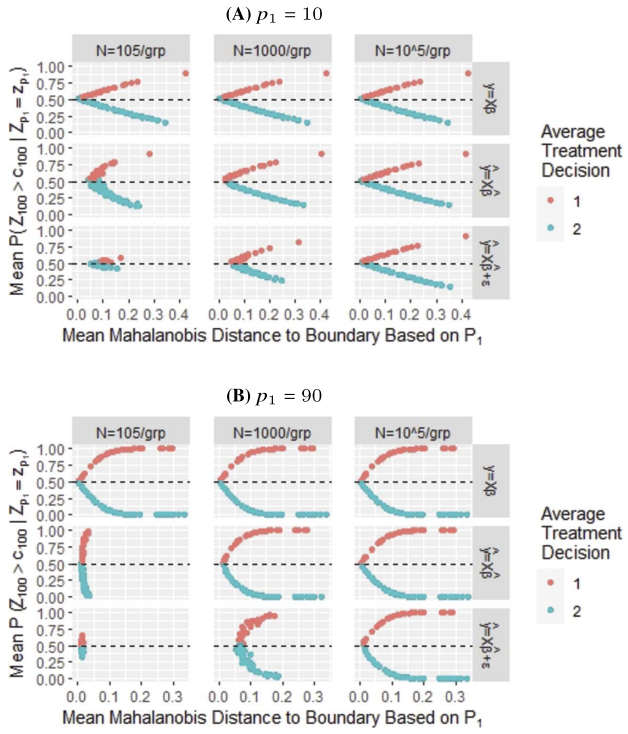

Figure 5. Personalized probability compared to Mahalanobis distance to the decision boundary based on p_1 covariates. Each point in the plot represents a patient in the test data ($n = 100$). The pink and blue indicators in the plot represent, respectively, the average treatment decision of $A = 1$ and $A = 2$ that were assigned to each patient across the 100 repeated simulations. The horizontal dashed line represents an equal probability of a patient being assigned to either treatment.

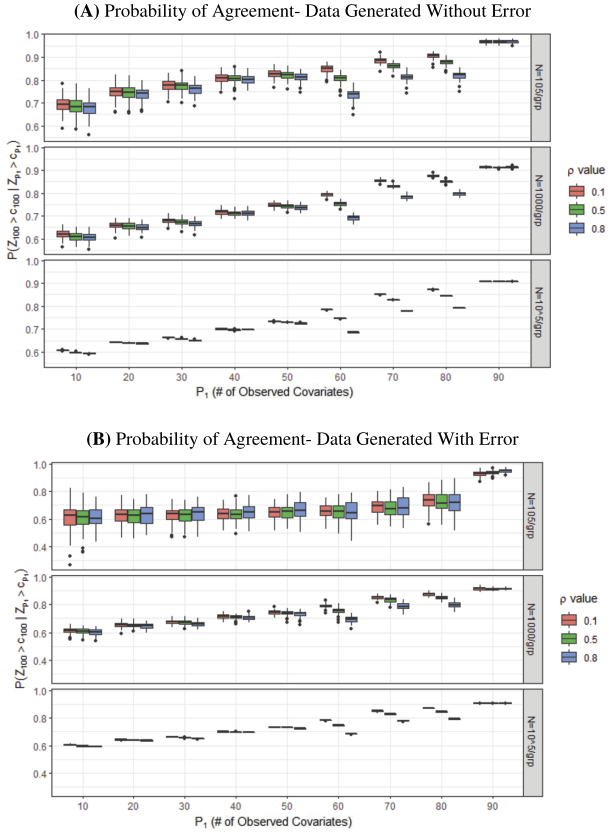

Figure 6. Probability of agreement of ultimate and current treatment assignment of $\mathcal{A} = 1$ with equally correlated predictors, with no error in the outcome generation in (a) and error in the outcome generation in (b). Z_{100} is the reparameterized TDR for all $p = 100$ covariates and c_{100} is the cutoff value for this decision boundary. The boxplots represent the probabilities across the 100 repeated Monte Carlo simulations, where the colors represent the three correlations.

The final expression is

$$\begin{pmatrix} \boldsymbol{\alpha}_{a1} \\ \boldsymbol{\alpha}_{a2} \end{pmatrix} = \begin{pmatrix} \boldsymbol{\Sigma}_{11}^{-1}(\boldsymbol{\Psi}_{1,Y_a} - \boldsymbol{\Sigma}_{12}\boldsymbol{\alpha}_{a2}) \\ (\boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12})^{-1}(\boldsymbol{\Psi}_{2Y_a} - \boldsymbol{\Sigma}_{21}\boldsymbol{\beta}_{a1}) \end{pmatrix}.$$

APPENDIX B. MAHALANOBIS DISTANCE TO THE DECISION BOUNDARY IN SEQUENTIAL FRAMEWORK

Using the same simulation study described in Section 5, we additionally studied the relationship between the personalized probability $P(d_K^{new} = 1|\boldsymbol{x}_1^{new}, \ldots, \boldsymbol{x}_k^{new})$ and the Mahalanobis distance to the decision boundary based on p_1, the number of observed covariates, computed with (11). The results in Figure 5 show that Mahalanobis distances are a scaled version of the Euclidean distance with smaller magnitudes, so the relationships in the plots are similar to that in Figure 4.

APPENDIX C. PROBABILITY OF AGREEMENT WITH EQUALLY CORRELATED COVARIATES

In our sequential framework, we expect that variables may be more highly correlated with others in the same subset compared with variables in other subsets. For example, variables measured from the same test such as an imaging scan may be more correlated than an imaging measure with a demographic variable. The following simulation studies aim to capture the impact of different correlations on the conditional probabilities. Here we focus for simplicity on presenting results for the probability of agreement as described in (18).

Simulations that varied the amount of equal correlation within subsets were conducted. The equal correlation setting

is similar to the simulation settings described in Section 5.2. We assume the same correlation matrix as (3.3), in this case with $\rho = \{0.1, 0.5, 0.8\}$.

We observe in Figure 6A that increasing ρ values correspond to a decrease in the overall average probabilities of agreement. However, the amount of variability across the $M = 100$ simulations remains relatively stable. Thus, the variability in the probability estimates may not be attributable specifically to the amount of correlation among variables. In Figure 6B, we observe that when the training data is small and the outcome is generated with error the probabilities are fairly similar across different correlations, especially with fewer p_1 covariates included in the treatment decision. With a larger sample size, we see the similarity to Figure 6A.

APPENDIX D. PROBABILITY OF AGREEMENT WITH BLOCK CORRELATED COVARIATES

Further simulations on the estimated probabilities and distances to the decision boundary were conducted in which the amount of correlation within subsets was fixed to be larger than the amount between subsets, which is the most realistic to real world clinical data.

The simulation settings described in Section 5.2 are very similar this case. The only difference in the data generation is in the covariance structure, as here we designate three block correlation settings, where the blocks are defined in subgroups of 10 covariates. The within correlation (within each block) is $\rho_w = \{0.1, 0.5, 0.8\}$, while the correlation between covariate subgroups is $\rho_b = 0.1$. The correlation matrix $\mathbf{R}$ can be written as:

$$\mathbf{R} = \begin{pmatrix} \mathbf{B}_1 & \rho_b \mathbf{1}\mathbf{1}' & \rho_b \mathbf{1}\mathbf{1}' & \cdots & \rho_b \mathbf{1}\mathbf{1}' \\ \rho_b \mathbf{1}\mathbf{1}' & \mathbf{B}_2 & \rho_b \mathbf{1}\mathbf{1}' & \cdots & \rho_b \mathbf{1}\mathbf{1}' \\ \rho_b \mathbf{1}\mathbf{1}' & \rho_b \mathbf{1}\mathbf{1}' & \mathbf{B}_3 & \cdots & \rho_b \mathbf{1}\mathbf{1}' \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho_b \mathbf{1}\mathbf{1}' & \rho_b \mathbf{1}\mathbf{1}' & \cdots & \rho_b \mathbf{1}\mathbf{1}' & \mathbf{B}_{10} \end{pmatrix}$$

and block matrix $\mathbf{B}_k$, where k is the subset, is written as

$$\mathbf{B}_k = \begin{pmatrix} 1 & \rho_w & \rho_w & \cdots & \rho_w \\ \rho_w & 1 & \rho_w & \cdots & \rho_w \\ \rho_w & \rho_w & 1 & \cdots & \rho_w \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho_w & \rho_w & \cdots & \rho_w & 1 \end{pmatrix}.$$

We observe in panel A of Figure 7 that increasing ρ_w values correspond to a decrease in the overall average probabilities of agreement, but that the amount of variability across the 100 simulations remains stable. The variability in the probability estimates may be more attributed to differences between subsets rather than within subsets of covariates. On the other hand, in panel B of Figure 7 we observe that when

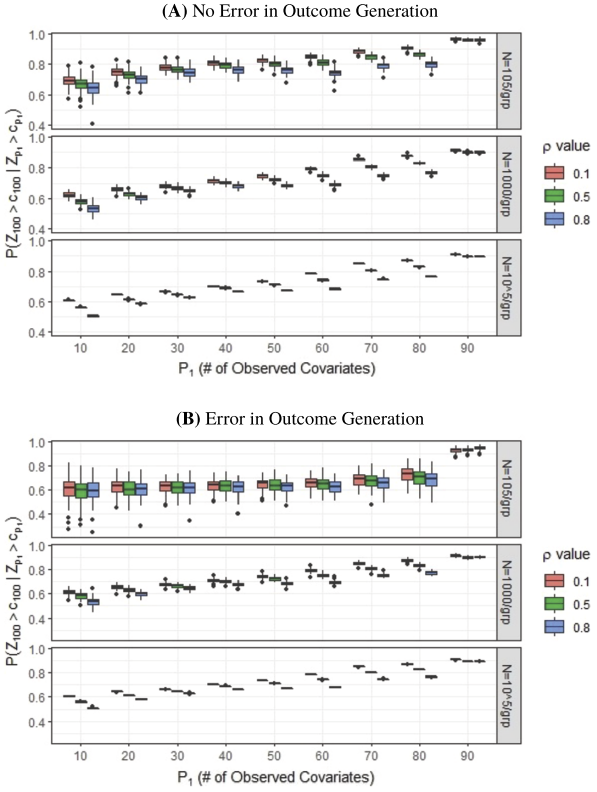

Figure 7. Probability of agreement of ultimate and current treatment assignment of $\mathcal{A} = 1$ with block correlation structures, with no error in the outcome generation in (a) and error in the outcome generation in (b). Z_{100} is the reparameterized TDR for all $p = 100$ covariates and c_{100} is the cutoff value for this decision boundary. The boxplots represent the probabilities across the 100 repeated Monte Carlo simulations, where the colors represent the three within subgroup correlations.

the training data is small and generated with error the probabilities are fairly similar across different correlations. This shifts once p_1 is around 50 covariates to be more similar to case with no error in the outcome generation.

ACKNOWLEDGMENTS

This work was supported by the National Institutes of Health (NIH) training grant 5TL1TR001447-05 and the National Institute of Mental Health of the NIH grant R01MH099003.

Received 23 September 2021

REFERENCES

[1] Aggarwal, C. C., Hinneburg, A. and Keim, D. A. (2001). On the Surprising Behavior of Distance Metrics in High-Dimensional Space. In *Database Theory — ICDT 2001* (J. Van den Bussche and V. Vianu, eds.). *Lecture Notes in Computer Science* **1973** 420–434. Springer Berlin Heidelberg, Berlin, Heidelberg.

[2] Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. L. Erlbaum Associates, Hillsdale, NJ.
[3] De Maesschalck, R., Jouan-Rimbaud, D. and Massart, D. L. (2000). The Mahalanobis distance. *Chemometrics and Intelligent Laboratory Systems* **50** 1–18.
[4] Drezner, Z. and Wesolowsky, G. O. (1990). On the computation of the bivariate normal integral. *Journal of Stat Comp and Sim* **35** 101–107. MR1041725
[5] Fischer, D., Mosler, K., Mottonen, J., Nordhausen, K., Pokotylo, O. and Vogel, D. (2020). Computing the Oja Median in R: The Package OjaNP. *Journal of Statistical Software, Articles* **92** 1–36.
[6] Friedman, J. H. (1997). On bias, variance, 0/1-loss, and the curse-of-dimensionality. *Data Mining and Knowledge Discovery* **1** 55–77.
[7] Friedrich, M. J. (2017). Depression Is the Leading Cause of Disability Around the World. *JAMA* **317** 1517–1517.
[8] Genz, A., Bretz, F., Miwa, T., Mi, X., Leisch, F., Scheipl, F. and Hothorn, T. (2020). mvtnorm: Multivariate Normal and t Distributions R package version 1.1-1.
[9] Hastie, T., Tibshirani, R. and Friedman, J. H. (2009). *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer, New York. MR2722294
[10] Hoff, P.D. and Niu, X. (2012). A covariance regression model. *Statistica Sinica* **22** 729-753. MR2954359
[11] Iniesta, R., Malki, K., Maier, W., Rietschel, M., Mors, O., Hauser, J., Henigsberg, N., Dernovsek, M. Z., Souery, D., Stahl, D., Dobson, R., Aitchison, K. J., Farmer, A., Lewis, C. M., McGuffin, P. and Uher, R. (2016). Combining clinical variables to optimize prediction of antidepressant treatment outcomes. *Journal of Psychiatric Research* **78** 94–102.
[12] Kotz, S., Johnson, N. L. and Balakrishnan, N. (2000). *Continuous Multivariate Distributions - Models and Applications*. John Wiley and Sons, Incorporated, Hoboken, NJ.
[13] Lu, T. T. and Shiou, S. H. (2002). Inverses of 2 x 2 Block Matrices. *Comp and Math with Apps* **43** 119–129.
[14] Manning, C. D., Raghavan, P. and Schutze, H. (2008). *Introduction to Information Retrieval*. Cambridge University Press, New York.
[15] Murphy, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **65** 331–355. MR1983752
[16] Qian, M. and Murphy, S. A. (2011). Performance guarantees for individualized treatment rules. *The Annals of Statistics* **39** 1180–1210. MR2816351
[17] Robins, J. M. (2004). *Optimal Structural Nested Models for Optimal Sequential Decisions* In *Proceedings of the Second Seattle Symposium in Biostatistics: Analysis of Correlated Data* 11, 189–326. Springer New York, New York, NY. MR2129402
[18] Tibshirani, R. (1996). Regression Shrinkage and Selection Via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* **58** 267–288. MR1379242
[19] Trivedi, M. H., McGrath, P. J., Fava, M., Parsey, R. V., T., K. B., Phillips, M. L., Oquendo, M. A., Bruder, G., Pizzagalli, D., Toups, M., Cooper, C., Adams, P., Weyandt, S., Morris, D. W., Grannemann, B. D., Ogden, R. T., Buckner, R., McInnis, M., Kraemer, H. C., Petkova, E., Carmody, T. J. and Weissman, M. M. (2016). Establishing moderators and biosignatures of antidepressant response in clinical care (EMBARC): Rationale and design. *J Psychiatr Res* **78** 11–23.
[20] Tsagris, M., Beneki, C. and Hassani, H. (2014). On the folded normal distribution. *Mathematics* **2** 12–28.
[21] Zhao, Y., Zeng, D., Rush, A. J. and Kosorok, M. R. (2012). Estimating Individualized Treatment Rules Using Outcome Weighted Learning. *Journal of the American Statistical Association* **107** 1106–1118. MR3010898

Nina Orwitz
180 Madison Avenue
New York, NY 10016
United States of America
E-mail address: no683@nyu.edu

Thaddeus Tarpey
180 Madison Avenue
New York, NY 10016
United States of America
E-mail address: Thaddeus.Tarpey@nyulangone.org

Eva Petkova
180 Madison Avenue
New York, NY 10016
United States of America
E-mail address: Eva.Petkova@nyulangone.org