Statistica Sinica **32** (2022), 2497-2519
doi:https://doi.org/10.5705/ss.202022.0047

DATA-GUIDED TREATMENT RECOMMENDATION WITH FEATURE SCORES

Zhongyuan Chen, Ziyi Wang, Qifan Song and Jun Xie

Purdue University

Abstract: Despite the availability of large amounts of genomics data, medical treatment recommendations have yet to use them successfully. In this study, we consider the utility of high-dimensional genomic-clinical data and nonparametric methods for making cancer treatment recommendations. Our work builds on the framework of the individualized treatment rule (Qian and Murphy (2011)) but we aim to overcome their method's limitations, specifically when the method encounters a large number of covariates and when the model is misspecified. We tackle this problem using a dimension reduction method, namely Sliced Inverse Regression (SIR, Li (1991)), with a rich class of models for the treatment response. Notably, the SIR defines a feature space for high-dimensional data, offering an advantage similar to those found in the popular neural network models. With the features obtained from the SIR, we use a simple visualization to compare different treatment options and recommend a treatment. Additionally, we derive the consistency and the convergence rate of the proposed recommendation approach using a value function. Lastly, we demonstrate the effectiveness of the proposed approach using simulation studies and a real-data example of the treatment of multiple myeloma.

Key words and phrases: Dimension reduction, individualized treatment rules, sliced inverse regression, visualization.

1. Introduction

The conventional approach to recommending treatments for a disease has been to use expert-driven guidance, based on knowledge built over decades. With the increasing availability of large amounts of data, there is growing interest in using such data to help choose different treatment options. For instance, cancer research has generated extensive genomics data, for example, on genetic mutations and mRNA expressions, along with clinical data including treatment options and clinical outcomes. These data add valuable information to support and complement expert knowledge for cancer treatments. In this study, we aim to develop a data-guided tool with simple visualizations to help doctors and patients evaluate different treatment options and to recommend a treatment.

Corresponding author: Jun Xie, Department of Statistics, Purdue University, West Lafayette, IN 47907, USA. E-mail: junxie@purdue.edu.

As a case study, we examine a data set of gene expressions and treatment responses from multi-center clinical trials of bortezomib for the treatment of multiple myeloma (Mulligan et al. (2007)). Multiple myeloma is a malignant bone marrow cancer. This disease is highly heterogeneous, meaning that different patients with diverse genomic information show different clinical outcomes (Mitra et al. (2017)). However, the current treatment strategy is limited to the experience of physicians and experts, mainly using the patient's clinical information, such as age and cancer stage. Using a specific genomic-clinical data set, we aim to make a treatment recommendation between two therapeutic choices, namely, a traditional chemotherapy drug called dexamethasone, and a targeted drug bortezomib.

Our goal is related to research on precision medicine, which has attracted a considerable amount of interests. A recent study on precision oncology for acute myeloid leukemia (Gerstung et al. (2017)) analyzed genomic-clinical data to support clinical decision-making. Zhu and Xie (2015) used a nonparametric method to identify patient subpopulations that would experience stronger treatment effects than the rest of the patient population. However, these studies were exploratory, with no formal framework to define an optimal treatment rule. A valuable formulation has been contributed by Qian and Murphy (2011). Formally, we consider a list of random variables $(\mathbf{X}, A, Y)$ from a genomic-clincial dataset, where Y denotes a treatment response variable (where larger values are better), $\mathbf{X} \in \mathcal{X} \subset \mathbb{R}^p$ denotes a set of clinical covariates plus genetic variables, for example, gene expressions, and $A \in \mathcal{A}$ denotes a treatment index taking values in a finite discrete space of treatment options, for example, $\mathcal{A} = \{-1, 1\}$ corresponding to a control and a treatment, or $\mathcal{A} = \{1, \ldots, M\}$ corresponding to M treatment options. A treatment recommendation rule is a function $d(\mathbf{X}) : \mathcal{X} \to \mathcal{A}$, and is called an individualized treatment rule in Qian and Murphy (2011). An individualized treatment rule that gives the highest mean response is the optimal one that we hope to find.

There are two types of approaches to constructing an optimal treatment rule, namely, direct methods and indirect methods. The indirect methods consist of two steps. The first step estimates a conditional mean of the treatment response, given clinical variables and the treatment index, $E(Y|\mathbf{X}, A)$. The optimal treatment rule is then defined as the one that maximizes the estimated conditional mean (Rosenwald et al. (2002); van't Veer and Bernards (2008); Qian and Murphy (2011); Cui, Zhu and Kosorok (2017); Hager, Tsiatis and Davidian (2018); Bai et al. (2017); Zhao et al. (2019); Guo, Zhou and Ma (2021)). The indirect methods rely heavily on a correct model specification for the conditional mean $E(Y|\mathbf{X}, A)$

(Qian and Murphy (2011)), which is often challenging to achieve. On the other hand, the direct methods are one-step procedures that circumvent the need for a conditional mean estimation. The basic idea is to obtain a treatment recommendation rule $d(\mathbf{X})$ directly by optimizing a criterion, called the Value function. The direct methods use a weighted classification framework and are based on the support vector machine (SVM) approaches. These methods include Outcome Weighted Learning (Zhao et al. (2012)), Residual Weighted Learning (Zhou et al. (2017)), and other variations (Dasgupta and Huang (2020); Mo, Qi and Liu (2021)). However, these methods are often confined by the limitations of the SVM procedure, such as difficulties with small separation margins, choices of kernels, and so on. Furthermore, there are also methods focusing on dynamic treatment regimes that consider treatment recommendations at multiple times and on Bayesian approaches of dynamic treatment regimes (Schulte et al. (2014); Zhang et al. (2013); Luckett et al. (2020); Liu et al. (2018); Yang, Tsiatis and Blazing (2018); Laber and Davidian (2017); Xu et al. (2016); Murray, Yuan and Thall (2018)). Nevertheless, the existing methods do not work well with high-dimensional data.

We focus on the indirect method, and improve it by developing an approach that contains a class of rich conditional mean models. Specifically, we apply Sliced Inverse Regression, or SIR (Li (1991)), to predict the treatment response. The SIR is designed to retrieve interesting features of high-dimensional data by using low-dimensional projections. The method can model the relationship between a treatment response and a set of genomic and clinical variables using an arbitrary unknown function. There is no linear model assumption about the conditional mean of the treatment response. In other words, the model space of $E(Y|\mathbf{X}, A)$ from the SIR method is often bigger than that of other indirect methods, such as Qian and Murphy (2011). Therefore, we gain robustness to the model specification.

An important strength of the SIR procedure is that it directly estimates the low-dimensional projection space and represents the high-dimensional data using only a few features. This resembles the feature definition component of the neural network models that are popularly used nowadays. We name the SIR projected data Feature Scores. Specifically, the SIR works well when the effects from individual clinical or genetic variables are weak, but the treatment response may depend on an unknown feature, which is common in cancer treatment. A simple scatter plot of the treatment response versus Feature Score allows users to visualize and compare different treatment options. We also obtain a nonparametric functional fitting of the treatment response versus Feature Score by LOESS

(Cleveland and Devlin (1988)). LOESS, namely, locally estimated scatterplot smoothing, was proposed by Cleveland and Devlin as a locally weighted regression method. It provides a simple estimation of $E(Y|\mathbf{X}, A)$. We further prove that the SIR plus LOESS procedure consistently estimates the optimal treatment rule under moderate assumptions. Thus, our method offers a tool for doctors, and even patients, to assess and confirm available treatment plans.

In summary, the primary contribution of this study is to define a small feature space in a framework of individualized treatment rules. The major advantages of the proposed method include 1) dimension reduction with feature detection, 2) rich conditional mean models for consistent estimation of the optimal treatment, 3) visualization of the optimal treatment recommendation, 4) a theoretical guarantee with a convergence rate.

The remainder of the article is organized as follows. In Section 2, we introduce the Value function, define the Feature Score, and show a visualization of the treatment recommendation. In Section 3, we prove the consistency and derive the convergence rate of the proposed recommendation approach. In Section 4, we present the results of our simulations and compare our proposed method with other methods. Section 5 demonstrates the results of applying the proposed method to the case study of treatments for multiple myeloma. Some discussions are given in Section 6. The Supplementary Materials include the information on the data, code, and technical proofs of the lemmas and theorems.

2. Treatment Recommendation through Feature Scores

Formally, we have a set of random variables $(\mathbf{X}, A, Y)$ in the data set, where $\mathbf{X} \in \mathcal{X} \subset \mathbb{R}^p$ denotes clinical covariates plus a big set of genetic variables, $A \in \mathcal{A}$ is the treatment index, taking values in a finite discrete space $\mathcal{A}$ of treatment options, and Y is the treatment response variable, with a larger values indicating a better treatment response. A treatment recommendation rule is a function $d(\mathbf{X})$ with values in $\mathcal{A}$. Denote the distribution of $(\mathbf{X}, A, Y)$ by P. Following the framework of individualized treatment rules (Qian and Murphy (2011)), we will first show that an optimal treatment recommendation rule must maximize $E(Y|\mathbf{X}, A = a)$ over $a \in \mathcal{A}$. This result justifies the use of indirect methods, which focus on the estimation of $E(Y|\mathbf{X}, A = a)$. Next, we will apply the SIR (Li (1991)) to estimate $E(Y|\mathbf{X}, A = a)$ and then obtain the optimal recommendation rule. The model space for the estimation of $E(Y|\mathbf{X}, A = a)$ in the SIR method is very large, which is the biggest advantage of our proposed method.

2.1. Value function and optimal recommendation

By convention, we use uppercase letters for random variables, and lowercase letters for the values of the random variables. The likelihood of $(\mathbf{X}, A, Y)$ under P is $f_0(\mathbf{x})p(a|\mathbf{x})f_1(y|\mathbf{x}, a)$, where f_0 is the unknown density of $\mathbf{X}$, $p(\cdot|\mathbf{x})$ is the randomization probability of A given $\mathbf{X} = \mathbf{x}$, and f_1 is the unknown distribution of Y conditional on $(\mathbf{X}, A)$. Let P^d denote the distribution of $(\mathbf{X}, A, Y)$ when a treatment recommendation rule $d(\mathbf{X})$ is used to assign treatments. Then, the likelihood becomes $f_0(\mathbf{x})\mathbb{1}(d(\mathbf{x}) = a)f_1(y|\mathbf{x}, a)$. We assume all treatment options, $a \in \mathcal{A}$, are available to the patient population, so the support spaces $\text{supp}(d(\mathcal{X})) = \text{supp}(\mathcal{A})$. Define the Value of d as $V(d) \triangleq E^d(Y)$. Assume $p(a|\mathbf{x}) > 0$, for any $a \in \mathcal{A}$ and $\mathbf{x} \in \mathcal{X}$. The Value of any treatment rule d can be expressed as

$$V(d) = \int Y dP^d = \int Y \frac{dP^d}{dP} dP = \int Y \frac{\mathbb{1}_{d(\mathbf{X})=A}}{p(A|\mathbf{X})} dP = E\left[Y \frac{\mathbb{1}_{d(\mathbf{X})=A}}{p(A|\mathbf{X})}\right].$$

An optimal treatment recommendation rule, denoted as d_0, is a rule that has the maximum Value over all possible treatment recommendation rules,

$$d_0 \in \underset{d}{\text{argmax}}\, V(d).$$

Moreover, denote $Q_0(\mathbf{X}, A) \triangleq E(Y|\mathbf{X}, A)$. We also have

$$\begin{aligned} V(d) &= E\left[\frac{\mathbb{1}(d(\mathbf{X}) = A)}{p(A|\mathbf{X})} E[Y|\mathbf{X}, A]\right] \\ &= E\left[\sum_{a\in\mathcal{A}} \mathbb{1}_{d(\mathbf{X})=a} Q_0(\mathbf{X}, a)\right] = E\left[Q_0(\mathbf{X}, d(\mathbf{X}))\right]. \end{aligned}$$

Note that the Value for the optimal treatment rule $V(d_0) = E[Q_0(\mathbf{X}, d_0(\mathbf{X}))] \leq E[\max_{a\in\mathcal{A}} Q_0(\mathbf{X}, a)]$. Meanwhile by the definition of d_0, $V(d_0) \geq V(d)|_{d(\mathbf{X})\in\text{argmax}_{a\in\mathcal{A}} Q_0(\mathbf{X},a)} = E[\max_{a\in\mathcal{A}} Q_0(\mathbf{X}, a)]$. Thus, the optimal treatment rule satisfies $d_0(\mathbf{X}) \in \text{argmax}_{a\in\mathcal{A}} Q_0(\mathbf{X}, a)$. Note that the optimal treatment recommendation rule d_0 is unique in many cases. For example, when the conditional mean function $Q_0(\mathbf{X}, a)$ has distinct values over $a \in \mathcal{A}$, the optimal rule is unique. Our ultimate goal is to estimate d_0, which will be achieved by first estimating the conditional mean $Q_0(\mathbf{X}, A)$.

More specifically, the estimated treatment recommendation rule is defined as

$$d(\mathbf{X}) \in \underset{a\in\mathcal{A}}{\text{argmax}}\, Q(\mathbf{X}, a), \tag{2.1}$$

where $Q(\mathbf{X}, A)$ is an estimator of the true conditional mean $Q_0(\mathbf{X}, A)$. The following result, modified from Qian and Murphy (2011), shows that the difference between the largest Value $V(d_0)$ and $V(d)$ is controlled by the mean squared error of the estimator $Q(\mathbf{X}, A)$.

We require an assumption similar to the margin condition in classification. Assume both the true conditional mean $Q_0(\mathbf{X}, A)$ and its estimator $Q(\mathbf{X}, A)$ are square integrable. Define $T(\mathbf{X}, A) = Q(\mathbf{X}, A) - E[Q(\mathbf{X}, A)|\mathbf{X}]$ and $T_0(\mathbf{X}, A) = Q_0(\mathbf{X}, A) - E[Q_0(\mathbf{X}, A)|\mathbf{X}]$. They are referred to as the treatment effect terms in Qian and Murphy (2011). The following assumption is about the margin of T_0, i.e., the difference in the mean responses between the optimal treatment and the suboptimal treatment.

(A.1) There exist some constants $C > 0$ and $\alpha > 0$ such that

$$\mathbf{P}\left(\max_{a\in\mathcal{A}} T_0(\mathbf{X}, a) - \max_{a\in\mathcal{A}\setminus \operatorname{argmax}_a T_0(\mathbf{X},a)} T_0(\mathbf{X}, a) \le \epsilon\right) \le C\epsilon^{\alpha},$$

for any $\epsilon > 0$.

Lemma 1. *Suppose $p(a|\mathbf{x}) \ge S^{-1}$ for a positive constant S for all $(\mathbf{x}, a)$ pairs and assume* (A.1)*. For any treatment rule $d : \mathcal{X} \mapsto \mathcal{A}$ and square integrable function $Q : \mathcal{X} \times \mathcal{A} \mapsto \mathbb{R}$, such that $d(\mathbf{X}) \in \operatorname{argmax}_{a\in\mathcal{A}} Q(\mathbf{X}, a)$, we have*

$$V(d_0) - V(d) \le C' \left[E(Q(\mathbf{X}, A) - Q_0(\mathbf{X}, A))^2\right]^{(1+\alpha)/(2+\alpha)},$$

where $C' = (2^{2+3\alpha} S^{1+\alpha} C)^{1/(2+\alpha)}$.

The proof is provided in the Supplementary Material.

2.2. A rich conditional mean model

Lemma 1 justifies the use of indirect methods. When we have a consistent estimator of $Q_0(\mathbf{X}, A) = E(Y|\mathbf{X}, A)$, that is, an estimator $Q(\mathbf{X}, A)$ converges to $Q_0(\mathbf{X}, A)$, Lemma 1 shows that the Value of the estimated treatment recommendation rule, i.e., $V(d)$ of $d(\mathbf{X}) \in \operatorname{argmax}_{a\in\mathcal{A}} Q(\mathbf{X}, a)$, also converges to the optimal value $V(d_0)$. However, this does not happen if the conditional mean is modeled incorrectly. In fact, if the approximation space used to estimate Q_0 does not contain the truth, then the estimated treatment recommendation rule will not be consistent. Qian and Murphy (2011) pointed out this challenge but did not present methods to address it. We attempt to offer a solution via Sliced Inverse Regression (SIR) (Li (1991)). SIR is a novel method for reducing the dimension of $\mathbf{X}$, without going through any model-fitting process in the first place. It is

developed under a very general model, $Y = g(\beta_1\mathbf{X}, \beta_2\mathbf{X}, \ldots, \beta_k\mathbf{X}, \epsilon)$, where β's are unknown row vectors, k is a small number, ϵ is the error term independent of $\mathbf{X}$, and g is an unknown function. When applying the SIR, we make the following very general assumption:

(A.2) For each treatment group $a \in \mathcal{A}$, the conditional mean response depends on a low-dimensional projection of $\mathbf{X}$. That is, $E[Y|\mathbf{X}, A = a] = E[Y|\beta_{a,1}\mathbf{X}, \beta_{a,2}\mathbf{X}, \ldots, \beta_{a,k}\mathbf{X}, A = a]$, with k as a small number, for example, $k = 1$ or 2.

In other words, given a treatment $a \in \mathcal{A}$, the conditional mean response is assumed to be $E(Y|\mathbf{X}, A = a) = \eta_a(\beta_{a,1}\mathbf{X}, \beta_{a,2}\mathbf{X}, \ldots, \beta_{a,k}\mathbf{X})$, where η_a is an unknown function that can take a general form.

There is a slight difference between the model assumption of the SIR, i.e., $Y = g(\beta_1\mathbf{X}, \beta_2\mathbf{X}, \ldots, \beta_k\mathbf{X}, \epsilon)$, and (A.2). In this study, we aim to estimate the conditional mean function $E(Y|\mathbf{X}, A)$, and use (A.2) to achieve dimension reduction for the mean function. The space defined by β's in (A.2) is called the *central mean subspace* in the literature. On the other hand, SIR was originally developed to estimate the *central dimension-reduction subspace* (Cook (1994)), which, by definition, contains the *central mean subspace*. There is also a rich body of literature on dimension reduction and estimating of the *central dimension-reduction subspace*. Important approaches include the SIR, sliced average variance estimate (SAVE) (Cook (2000)), principal Hessian direction (Li (1992)), and their variations. We choose to use the SIR, owing to its simplicity.

With Assumption (A.2), the projection of a set of predictors $\mathbf{X}$ onto the k-dimensional subspace, $(\beta_{a,1}\mathbf{X}, \beta_{a,2}\mathbf{X}, \ldots, \beta_{a,k}\mathbf{X})$, captures all we need to know about Y for the given treatment $A = a$. The projection space and the arbitrary function η_a are allowed to be different for different treatment groups $a \in \mathcal{A}$. This assumption offers a rich class of models for the conditional mean $Q_0(\mathbf{X}, A) = E[Y|\mathbf{X}, A]$. Specifically, if we denote $\mathcal{Q}$ as the approximation space for Q_0, then $\mathcal{Q}$ contains the linear model, the commonly used generalized linear models, and many more. In fact, our assumption does not impose any structure on how the projected variable affects the mean response $E[Y|\mathbf{X}, A]$. Therefore, it gives a very general approximation space $\mathcal{Q}$, likely containing the true function Q_0. When $Q_0 \in \mathcal{Q}$, we obtain a consistent estimator Q, and hence the estimated recommendation rule in Formula (2.1) converges to the optimal rule, as per Lemma 1. In contrast, Qian and Murphy (2011) used linear models (with certain basis functions) to approximate Q_0. Our approximation space $\mathcal{Q}$ improves on the methods of Qian and Murphy (2011) when a low-dimensional

projection of $\mathbf{X}$ exists that contains most of the information required to predict Y. In conclusion, our rich class of approximation space $\mathcal{Q}$ eliminates the model misspecification problem, to a certain degree, because a big space $\mathcal{Q}$ is likely to contain the correct model.

The number k is supposed to be very small (e.g., 1 or 2), and $\beta_{a,1}\mathbf{X}$, or $(\beta_{a,1}\mathbf{X}, \beta_{a,2}\mathbf{X})$, provides summary information about a patient that we use to predict the treatment response. We name $\beta_{a,1}\mathbf{X}$, or $(\beta_{a,1}\mathbf{X}, \beta_{a,2}\mathbf{X})$ if $k = 2$, as Feature Score. Using the Feature Score enables us to represent the cancer treatment situation in which there is no strong effect from an individual genetic variable but the treatment response depends on unknown features. Li (1991) provided a direct etimator of β's through the SIR procedure. For each treatment group $A = a$, suppose we have patient samples of the treatment response and the covariate vector $\{(y_i, \mathbf{x}_i)\}$. We apply the SIR and obtain the first projection direction $\hat{\beta}_1$. The Feature Score is denoted as $u_i = \hat{\beta}_1\mathbf{x}_i$, which can be interpreted as a summary feature of a patient, and is supposed to capture the majority of data information for the prediction of the treatment response Y. More interestingly, this feature definition is analogous to that of the neural network model, and the SIR is able to directly estimate the features without knowing the link function η_a.

2.3. Simple visualization with the Feature Score

Suppose $\mathcal{A} = \{1, \ldots, M\}$. Therefore, there are M different treatment groups in a given data set. We conduct a SIR for each treatment group and obtain the first projection direction $\hat{\beta}_{a,1}$, for $a = 1, \ldots, M$. We can project all patients onto a one-dimensional space (line) and calculate their Feature Scores $u_i = \hat{\beta}_{a,1}\mathbf{x}_i$, where the Feature Scores vary across different treatment groups. We draw a simple scatter plot of y_i versus u_i for each treatment group, for $a = 1, \ldots, M$. Even though the Feature Scores of different treatment groups $u_i = \hat{\beta}_{a,1}\mathbf{x}_i$ are not comparable, we can still compare the treatment response via the vertical axis, which has the same scale over different scatter plots (see Figure 1). These plots provide a visualization of the treatment options, i.e., $a = 1, \ldots, M$, where a larger vertical value indicate a better treatment response.

We also obtain a nonparametric fitting of the function, $\hat{g}_a(u)$, for example, by local constant estimates, or LOESS (Cleveland and Devlin (1988)), for each of the treatment groups $a = 1, \ldots, M$. These nonparametric estimates provide the predicted treatment responses for each treatment option. Given a new patient with data vector $\mathbf{x}$, we first calculate its Feature Score, $u_a = \hat{\beta}_{a,1}\mathbf{x}$. Then, the best treatment option is the one that maximizes the predicted treatment responses.

More specifically, we recommend a treatment choice as

$$\underset{a=1,\ldots,M}{\operatorname{argmax}}\, \tilde{g}_a(\hat{\beta}_{a,1}\mathbf{x}), \tag{2.2}$$

where $\hat{\beta}_{a,1}$ is from the SIR procedure, and $\tilde{g}_a(\cdot)$ is the nonparametric function estimate based on the patient samples $\{(y_i, u_i)\}$, with $u_i = \hat{\beta}_{a,1}\mathbf{x}_i$. In general, we have the subspace dimension $k > 1$, and the SIR may project data from different treatment groups onto different subspaces. Nevertheless, we obtain a nonparametric estimate of the functional relationship, $\tilde{g}_a(\hat{\beta}_{a,1}\mathbf{x}, \ldots, \hat{\beta}_{a,k}\mathbf{x})$. The treatment recommendation is defined similarly to (2.2).

The visualization through the scatter plot of y_i versus Feature Score u_i is a very useful tool. We can locate a patient's Feature Score $u_a = \hat{\beta}_{a,1}\mathbf{x}$ on the horizontal axis (or on the projected space when the Feature Score is more than one dimensional), and then consider the treatment response values on the vertical axis in the scatter plots, as shown in Figure 2. We can also use the vertical axis to compare the predicted treatment responses between our proposed treatment plan and the plan based on current expert guidelines, which will indicate the improvements in the treatment response that may be achieved by using the proposed treatment recommendation.

2.4. Data preprocessing and the algorithm

Before implementing the SIR procedure, we need to preprocess the data. First, we need to confirm that a given genomics data set contains significant information for the prediction of the treatment response. We evaluate the overall dataset information through a global hypothesis testing method, namely, the Cauchy combination test developed by Liu and Xie (2019). The p-value from this test serves as evidence to support a data-guided treatment recommendation. If the test gives a large p-value, we should not consider using the genomic data to forecast a patient's prognosis and to recommend treatments.

In the second step of the data preprocessing, we conduct initial variable selection before implementing the SIR when we analyze a large number of genomic variables. The SIR is a dimension reduction method involving principal component analysis (PCA). In general, some initial reduction in dimensionality is desirable before applying any PCA-type methods (Johnstone and Lu (2009)). In fact, as proven by Lin, Zhao and Liu (2018), when the dimension p is larger than the sample size n, the SIR estimate of the central space is inconsistent. Hence, we require certain structural assumptions, such as the sparsity of the projection vectors β's. Various high-dimensional sparse SIR regression methods have

been proposed in the literature (Jiang and Liu (2014); Lin, Zhao and Liu (2019); Tan, Shi and Yu (2020); Lin et al. (2021)). Here, we perform a screening plus low-dimensional SIR procedure.

In terms of computational cost, "screening + SIR" has a clear advantage, because high-dimensional sparse SIR methods always involve operations of the p-dimensional matrix $var(E(\mathbf{X}|Y))$. Our experience shows that "screening + SIR" reduces the computational time of a high-dimensional sparse SIR method by at least half.

In terms of assumptions for theoretical properties, a high-dimensional SIR regression requires that the whole p-dimensional distribution of $\mathbf{X}$ satisfies the linearity assumption, that is, for any ξ, $E(\xi^T\mathbf{X}|\beta_1^T\mathbf{X},\ldots,\beta_k^T\mathbf{X})$ is a linear combination of $(\beta_1^T\mathbf{X},\ldots,\beta_k^T\mathbf{X})$. It also requires coverage assumption, that is, the k nonzero eigenvalues λ_i's of $var(E(\mathbf{X}|Y))$ satisfy $\kappa\lambda \geq \lambda_1 \geq \cdots \geq \lambda_k \geq \lambda$, for some constants κ and λ. In contrast, "screening + SIR" requires a mild condition that the subset of variables in the true SIR regression model are all marginally correlated with Y, and hence will survive from the screening step. Then, only the low-dimensional joint distribution of the selected X_i's needs to satisfy the linearity and coverage assumptions.

We consider here two variable-screening methods, to select a subset of variables before running the SIR process. The first selects variables with the smallest p-values from a simple regression of Y over X_j and A, for $j = 1,\ldots,p$, at a false discovery rate (FDR) cutoff, for example, of 5%. The second method screens the important variables from a nonparametric local regression of Y over X_j and A using LOESS, with the smallest 5% residual errors. Either method can be used before implementing the SIR. These two methods are based on the idea of sure independence screening (Fan and Lv (2008)), which are computationally efficient and can attain the sure screening property. The sure screening property guarantees that important variables survive the screening approach with probability tending to one. With a proper screening threshold, the sure screening property holds as long as the correlation coefficients between the true active variables X_j's and Y are bounded away from zero (e.g., Assumption 3 in Fan and Lv (2008)).

To determine the number of Feature Scores k, which is the number of dimensions for the reduction in the proposed SIR model, we can use the χ^2 test suggested by Li (1991). On the other hand, as SIR is a PCA-type method, it is common practice to consider one or two Feature Scores, that is, one or two principal components, for visualization. The specific algorithm for our treatment recommendation is provided in Algorithm 1.

Algorithm 1 *Treatment recommendation procedure*

procedure $s = \mathsf{TreatRcmd}(Y, \mathbf{X}, A, \mathbf{x}_{new})$

Input: A training data set with observed $(\mathbf{X}, A, Y)$, where Y is the treatment response, $\mathbf{X} = (X_1, \ldots, X_p)$ is the set of genomic variables and clinical covariates, and A is the treatment index; A new observation with vector value $\mathbf{x}_{new}$ for treatment recommendation.

Output: Scatter plots of Y versus Feature Scores; The predicted response under each treatment option for $\mathbf{x}_{new}$ and the optimal treatment option.

▷ **Overall information summary**

Calculate p-value from the Cauchy combination test.

Alert if the overall p-value is large. Continue only if the p-value is small.

▷ **Subset selection (Optional)**

Select a subset of X_j's for the following SIR procedure, using either a linear regression or a nonparametric local regression of Y over X_j and A. A default cutoff is the false discovery rate 5%, or using LOESS with the smallest 5% residual errors.

▷ **Dimension reduction (SIR)**

For each treatment group $A = a$, conduct the SIR to obtain the low-dimensional projection directions $\hat{\beta}_a$.

Construct scatter plots of Y versus Feature Score $u_a = \hat{\beta}_a \mathbf{X}$ for each treatment group.

▷ **Prediction**

For the new data point $\mathbf{x}_{new}$, calculate its Feature Scores $u_a = \hat{\beta}_a \mathbf{x}_{new}$ under each treatment option $A = a$, and predict the response under the corresponding treatment.

Obtain the optimal treatment recommendation that gives the largest predicted response.

end procedure

3. Consistency and Convergence Rate

Our treatment recommendation rule is $d(\mathbf{X}) \in \operatorname{argmax}_{a \in \mathcal{A}} Q(\mathbf{X}, a)$, where $Q(\mathbf{X}, A)$ is an estimator of $Q_0(\mathbf{X}, A) = E(Y|\mathbf{X}, A)$ and is obtained by SIR and the nonparametric procedure LOESS. Recall the Value function defined in Section 2.1. The following theorem shows that we can have $V(d)$ converging to the optimal Value $V(d_0)$ at a certain rate. In addition to the margin condition (A.1), we require further assumptions for the SIR (Li (1991)) and the nonparametric LOESS estimator. We first rewrite the SIR assumption (A.2) by denoting the treatment index as $i \in \mathcal{A} = \{1, \ldots, M\}$ and the projection directions β's as $\mathbf{B}_i \in \mathbb{R}^{k \times p}$, for $k < p$.

(A.2) There exist some full-rank matrices $\mathbf{B}_i \in \mathbb{R}^{k \times p}$, for $k < p$, such that $E[Y|\mathbf{X}, A = i] = E[Y|\mathbf{B}_i\mathbf{X}, A = i] = \eta_i(\mathbf{B}_i\mathbf{X})$, where $\eta_i(\cdot)$'s are ρ-Lipschitz continuous and have continuous second derivatives. Furthermore, for any

row vector $\xi \in \mathbb{R}^p$, $E[\xi\mathbf{X}|\mathbf{B}_i\mathbf{X}]$ is a linear function of $\mathbf{B}_i\mathbf{X}$. In addition, the dimension of the central inverse curve $E[\mathbf{X}|y, A=i]$ is equal to the dimension of the space spanned by the columns of $\mathbf{B}_i$, $col(\mathbf{B}_i)$, and the variance $v_i(u) = \mathrm{Var}[Y|\mathbf{B}_i\mathbf{X} = u, A = i]$ is a continuous function.

(A.3) Denote the kernel function of LOESS by $K_H(u) = |H|^{-1/2}K(H^{-1/2}u)$, where $u \in \mathbb{R}^k$ and the bandwidth matrix $H \in \mathbb{R}^{k\times k}$. Assume the kernel function $K(\cdot)$ is ρ-Lipschitz, compactly supported, and satisfies $\int uu^\top K(u)\, du = \mu_2(K)\mathbf{I}$, where $\mathbf{I}$ is the identity matrix and $\mu_2(K)$ is a constant depending on K. Moreover, all odd-order moments of K are equal to zero, that is, $\int u_1^{l_1}\cdots u_d^{l_d}K(u)du = 0$ for all non-negative $l_1\cdots l_d$ when their sum is odd. Furthermore, the bandwidth matrix H is symmetric and positive definite with each entry, as well as $n^{-1}|H|$, tends to 0 as $n \to \infty$, and the ratio of the largest and the smallest eigenvalue of H is uniformly bounded for all n.

(A.4) For all $i \in \mathcal{A}$, let $f_i(\cdot)$ be the conditional density function of $\mathbf{B}_i\mathbf{X}$ given $A = i$. Assume that $f_i(\cdot)$ is uniformly bounded away from zero and has a continuous gradient function $D_{f_i}(\cdot)$.

(A.5) Denote $n_i = |\{j : A_j = i\}|$ as the number of observations in the treatment group $A = i$. Assume $\min_{i\in\mathcal{A}} P(A = i) > c$, for some positive constant c, and the support set of $\mathbf{X}$ is bounded.

As represented in (2.2) in Section 2.3, we write the treatment recommendation rule as $d(\mathbf{x}) \in \mathrm{argmax}_{i\in\mathcal{A}} Q(\mathbf{x}, i)$, where $Q(\mathbf{x}, i) = \tilde{g}_i(\widehat{\mathbf{B}_i}\mathbf{x})$, with $\widehat{\mathbf{B}_i}$ as the estimated projection directions from the SIR and $\tilde{g}_i(\cdot)$ as the LOESS function from the training data $\{\widehat{\mathbf{B}_i}\mathbf{x}_j, y_j\}_{\{j:A_j=i\}}$.

Theorem 1. *Assume* (A.1)-(A.5) *hold. Then, the difference between the optimal Value, $V(d_0)$, and $V(d)$ of our treatment recommendation rule converges to zero in probability as $n \to \infty$:*

$$V(d_0) - V(d) \le \left(|H|^{-1}\|H^{-1/2}\|_F^2\mathcal{O}_p(\frac{1}{n}) + \mathcal{O}_p\left(\frac{|H|^{-1/2}}{n} + \|H\|_1^2\right)\right)^{(1+\alpha)/(2+\alpha)}, \tag{3.1}$$

where $\|H\|_1$ denotes the maximum column absolute sum, and $\|\cdot\|_F^2$ denotes the Frobenius norm. When the bandwidth matrix $H = diag\{h, \ldots, h\}$, with $h = n^{-1/(k+3)}$, the upper bound on the right-hand side becomes $\mathcal{O}_p(n^{-(2(1+\alpha))/((k+3)(2+\alpha))})$.

The proof is provided in the Supplementary Material.

Remark 1. Theorem 1 is obtained by combining the estimation errors of the SIR procedure and the LOESS nonparametric regression. The second error term in (3.1) is the intrinsic estimation error of the LOESS regression, and the first error term is the additional estimation error induced by the uncertainty of the SIR procedure. With an appropriate choice of the kernel bandwidth, for example, $h = n^{-1/(k+3)}$, our treatment recommendation rule is consistent in the sense that $V(d)$ converges to the optimal Value $V(d_0)$.

Remark 2. The smoothness assumption of η_i and v_i, the requirements on the kernel choice in (A.3) and (A.4) ensure the consistency of the nonparametric estimation for each mean regression function η_i via the local linear regression approach. The bandwidth matrix H usually takes a simple form as $diag\{h, \ldots, h\}$, where $h > 0$. Given this simplification, the last statement in assumption (A.3) is satisfied automatically.

Remark 3. The compactness assumption on the support set of $\mathbf{X}$ in (A.5) greatly facilitates our theoretical analysis, for example, by trivially ensuring that $\|D_f(\cdot)\|$ is bounded. This assumption is reasonable for most medical treatment applications, because patient measurements, such as gene expression levels, are usually bounded or standardized. We conjecture that our theoretical results will still hold for an unbounded $\mathbf{X}$, such as the Gaussian design. A rigorous convergence analysis for such cases is left for future work. Assumption (A.5) also ensures that $n_i \asymp n$, in probability.

Remark 4. For simplicity of representation, our theorem considers only fixed p and k situation. If p and k increase with respect to n, then the corresponding convergence rates can be studied rigorously by utilizing the high-dimensional algorithm and the SIR theory developed by, e.g., Zhu, Miao and Peng (2006) and Lin, Zhao and Liu (2018), Lin et al. (2021), Lin, Zhao and Liu (2019). In Section 3 of the Supplmentary Material, we present a convergence result under $p \to \infty$ and $p/n \to 0$.

4. Simulation Studies

To assess the proposed method, we perform extensive simulations. We compare our method with three existing approaches, namely, Outcome Weighted Learning (OWL) (Zhao et al. (2012)), Residual Weighted Learning (RWL) (Zhou et al. (2017)), and a linear regression method with an ordinary least squares estimation of the conditional mean of the treatment response, denoted as OLS.

We generate p covariates $X_1, \ldots, X_p$ from uniform $[-1, 1]$, where a small

and a large covariate set are considered with $p = 8$ or 100. We consider two treatment options $\mathcal{A} = \{1, -1\}$ of a randomized controlled study. The response Y follows a normal distribution with mean $\mu(\mathbf{x}) + t_0(\mathbf{x})a$ and standard deviation one, where $\mu(\mathbf{x})$ represents the effect of the covariates $\mathbf{x} = (x_1, x_2, \ldots, x_p)$ and $t_0(\mathbf{x})a$ represents the treatment effect, which may depend on $\mathbf{x}$. We simulate two sample sizes, $n = 100$ and $n = 400$, with half of the samples in the treatment group and the other half in the control group. The terms $\mu(\mathbf{x})$ and $t_0(\mathbf{x})a$ are chosen from the following four scenarios:

1. $\mu(\mathbf{x}) = 2 + 4x_1 + 4x_2 + 4x_3$, when $a = 1$; $\mu(\mathbf{x}) = (2 + 4x_1 + 4x_2 + 4x_3)^2$, when $a = -1$; $t_0 = 0$.
2. $\mu(\mathbf{x}) = 2 + 2x_1 + 2x_2 + 4x_3 + 4x_4$; $t_0(\mathbf{x}) = 1.3(x_2 - 2x_1^2 + 0.3)$.
3. $\mu(\mathbf{x}) = 10x_1/(0.5 + (x_2 + 1.5)^2)$; $t_0(\mathbf{x}) = 1.3(x_2 - 2x_1^2 + 0.3)$.
4. $\mu(\mathbf{x}) = 10x_1/(0.5 + (x_2 + 1.5)^2)$; $t_0(\mathbf{x}) = 3.8(0.8 - x_1^2 - x_2^2)$.

Scenario 1 is modified from a simulation model of OWL (Zhao et al. (2012)). We define the mean function as a linear function for $a = 1$, and its quadratic function for $a = -1$. Scenario 2 is similar to the second scenario in RWL (Zhou et al. (2017)). Scenarios 3 and 4 have nonlinear functions, with $\mu(\mathbf{x})$ modified from a simulation model of SIR (Li (1991)).

We first conduct the Cauchy combination test on the entire set of covariates $X_1, \ldots, X_p$. The test gives significant results for all simulations. For data sets with a small number of covariates, i.e., $p = 8$, we directly implement the SIR method. On the other hand, for data sets in which the dimension is comparable to the sample size, i.e., $p = 100$, we conduct an initial variable selection before implementing the SIR (see Algorithm 1). A simulated data set typically has between two and seven selected variables, with the exact number varying for different model scenarios and different simulation replicates. We then conduct a SIR for each treatment group and obtain the first projection direction $\hat{\beta}_a$, for $a = 1$ or -1. Feature Scores are calculated for subjects in the corresponding treatment group, for $a = 1$ or -1, as $u_i = \hat{\beta}_a \mathbf{x}_i$.

Figure 1 shows two scatter plots of y_i versus u_i, one for each treatment group. These plots display the functional relationships between the response Y and Feature Score and are used to predict responses for a new observation $\mathbf{x}$. Figure 2 shows plots of the predicted response versus Feature Scores, where each sample $\mathbf{x}_i$, for $i = 1, \ldots, n$, is considered as a new observation (a test data). Each sample has two Feature Scores, $u_{a,i} = \hat{\beta}_a \mathbf{x}_i$, for $a = 1$ or -1, and two predicted

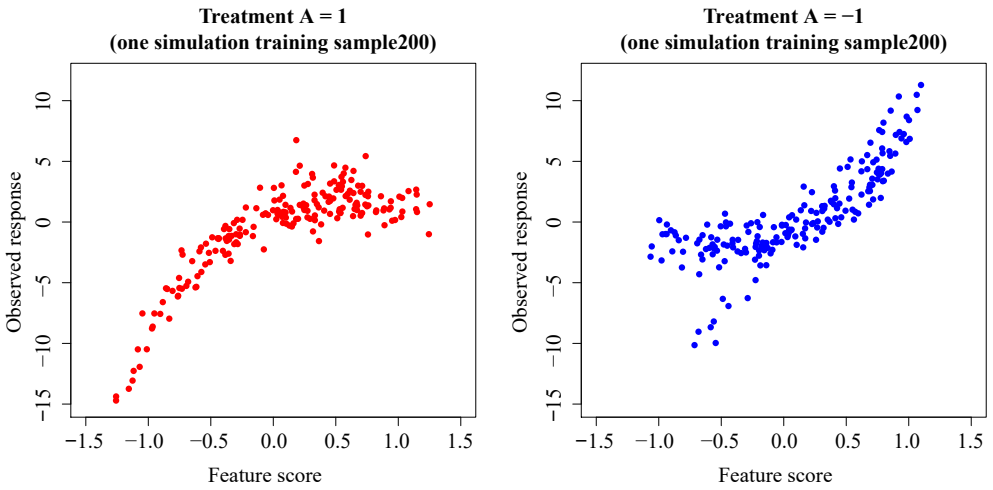

Figure 1. Scatter plot of Y versus Feature Score $u_a = \hat{\beta}_a \mathbf{x}$ under each treatment (Scenario 3).

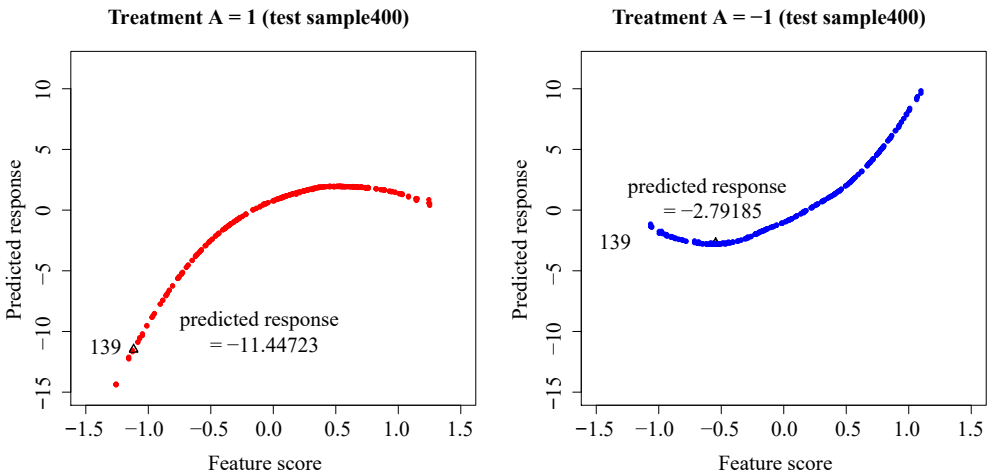

Figure 2. Predicted response value versus Feature Score $u_a = \hat{\beta}_a \mathbf{x}$ under each treatment (Scenario 3). A specific data point, ID 139, is marked for visualization of the optimal treatment.

treatment responses from the LOESS fits of Figure 1. We use the R package `loess()` with its default bandwidth parameter $h = 0.75$. The vertical axes of these plots use exactly the same scale for the treatment response, and thus are directly comparable. It clearly demonstrates the optimal treatment option, either $a = 1$ or -1, for each sample. Specifically, the subject with ID 139, marked by

a small triangle in the plots, has the predicted response value $\hat{Y} = -11.44723$ when assigned to the treatment $a = 1$, and $\hat{Y} = -2.79185$ when assigned to the treatment $a = -1$. This subject is then recommended to get treatment $a = -1$ due to the larger predicted response value.

For the simulation studies, we know the true optimal treatment recommendation, that is, the treatment option with the larger value of $\mu(\mathbf{x})+t_0(\mathbf{x})a$, for a given subject with covariate values $\mathbf{x}$. Thus, we can evaluate our method and compare it with existing methods by calculating a misclassification error. Specifically, if an approach recommends the same treatment option as the truth, then there is no misclassification error. Otherwise, the misclassification error is one. We apply four treatment recommendation methods: OWL (Zhao et al. (2012)); RWL (Zhou et al. (2017)); a linear regression to predict Y, and then to recommend the treatment with the larger predicted value (OLS); and our proposed method, denoted as SIR. We use an existing R package to perform OWL and RWL, available at `https://cran.r-project.org/web/packages/DynTxRegime/index.html`. We make a treatment recommendation for each sample, while considering all other samples as training data. Figure 3 displays the misclassification rates. The rate is the percentage of the number of misclassified treatments over the total number of patients (sample size n). We repeat the whole simulation procedure 1,000 times and plot the mean value and the standard deviation, with two error bars around the mean, in Figure 3.

In general, our approach (SIR) shows better performance with lower misclassification rates. In particular, our approach performs substantially better than RWL and OWL in Scenarios 2, 3, and 4. The results for SIR and RWL are comparable in Scenario 1. In addition, our approach shows lower misclassification rates than those using OLS in Scenario 1, 3, and 4. The results for SIR and OLS are comparable in Scenario 2. The favorable performance of SIR is because of the general assumption of the treatment response model, i.e., Assumption (A.2), which gives a large approximation space for the true conditional mean function Q_0. In other words, the model space of SIR is often bigger than those of existing methods with mostly linear models. We improve the treatment recommendation by obtaining a good estimator of Q_0.

5. A Case Study

In this section, we apply our proposed method to the study of bortezomib in the treatment of multiple myeloma (Mulligan et al. (2007)). Bortezomib is the first therapeutic proteasome inhibitor tested in humans. It is approved in the

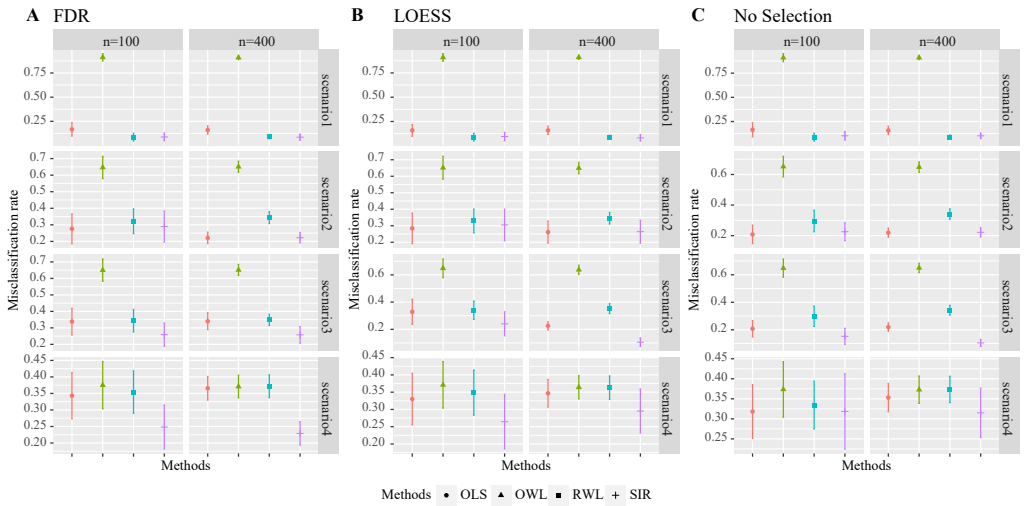

Figure 3. Comparison of different treatment recommendation methods in terms of the mean (center) and the standard deviation (error bars) of misclassification rates from 1,000 simulations: A. FDR is used to screen all $p = 100$ variables in the first step; B. LOESS is used to screen all $p = 100$ variables in the first step; C. No screening but $p = 8$.

United States for treating relapsed multiple myeloma. Because bortezomib is a therapeutic choice in addition to the standard chemotherapy, there is a need to be able to recommend a treatment for a given patient. Our goal is to provide a treatment recommendation, either dexamethasone (dex) or bortezomib, based on data information.

To achieve this, we use a genomic-clinical data set from the Gene Expression Omnibus (GEO) database (GSE9782). Data from two platforms of Affymetric microarrays (GPL96 and GPL97) are merged to obtain a large sample size, with a total of 477 patients, 338 of whom received bortezomib and 139 received dex. The merged data contain a smaller number of gene probesets (or simply genes) than each of the individual platform data. On the other hand, we verified that significant genes from each data set are included in the merged data. The variables considered in our analysis include:

- a set of clinical prognostic factors, i.e., gender, race, and age;
- a treatment index, either bortezomib or dex, denoted as A;
- gene expression measurements of 168 genes in the merged data, denoted as X_j, for $j = 1, \ldots, 168$; and
- clinical response, denoted as Y, with five levels coded as 1-5 corresponding to

progressive disease (PD), no change (NC), minimal response (MR), partial response (PR), and complete response (CR), respectively.

We first evaluate whether this data set provides significant information for the prediction of treatment response Y. The three clinical factors, i.e., gender, race, and age, have no significant effect on Y ($R^2 = 0.004179$), and hence are not considered in the following anlaysis. The Cauchy combination test (Liu and Xie (2019)) gives a p-value of 0.0004, suggesting that the genomic data set contributes to the treatment response and provides useful information for treatment recommendation.

Given that the sample size and the number of genes are comparable, we deem an initial variable selection to be necessary before running the SIR procedure. We select a subset of the eight most significant genes at a false discovery rate cutoff (0.002). These are the genes of ribosomal proteins and translation initiation factors. Interestingly, these genes match with the results reported in the literature that patients with perturbations of certain ribosomal proteins and translation initiation factors showed responses to the bortezomib treatment (Mulligan et al. (2007); Sulima and De Keersmaecker (2017); Hofman et al. (2016)). We then apply our SIR method of treatment recommendation using this set of eight genes and compare the performance with that of OLS and RWL. For the SIR method, Feature Score is calculated as a one-dimensional projection of the gene predictors for each treatment group.

More specifically, we randomly split the data into five equal-sized parts. Four parts (training data) are used to fit a model, either OLS, SIR, or RWL, and the remaining one part (test data) is used to evaluate the corresponding treatment recommendation methods. Different from the simulation examples, we do not know the true optimal treatment recommendation for this case study, and hence cannot calculate misclassification errors. Instead, we calculate an unbiased estimator of the Value function, as in Qian and Murphy (2011). We repeat the process 1,000 times, and report the mean and standard deviation of the estimated Value functions in Table 1. The observed treatment index A in the data also corresponds to a treatment recommendation rule. Its estimated Value function serves as a baseline for the performance comparison.

Table 1 shows that SIR improves the baseline Value function from 2.542 to 2.825, and is slightly better than the OLS and RWL methods, although the difference from OLS and RWL is minimal. Plots of Y versus Feature Score (plots not shown here) actually display a certain degree of a linear trend, and the predicted Y curves from OLS and SIR are not very different from each other.

Table 1. Comparison of the empirical Value function in a random testing dataset using different methods, OLS, SIR, and RWL. Mean (std) values of the empirical Value function through 1,000 resampling are reported.

Observed	OLS	SIR	RWL
2.542(0.127)	2.818(0.157)	2.825(0.158)	2.804(0.166)

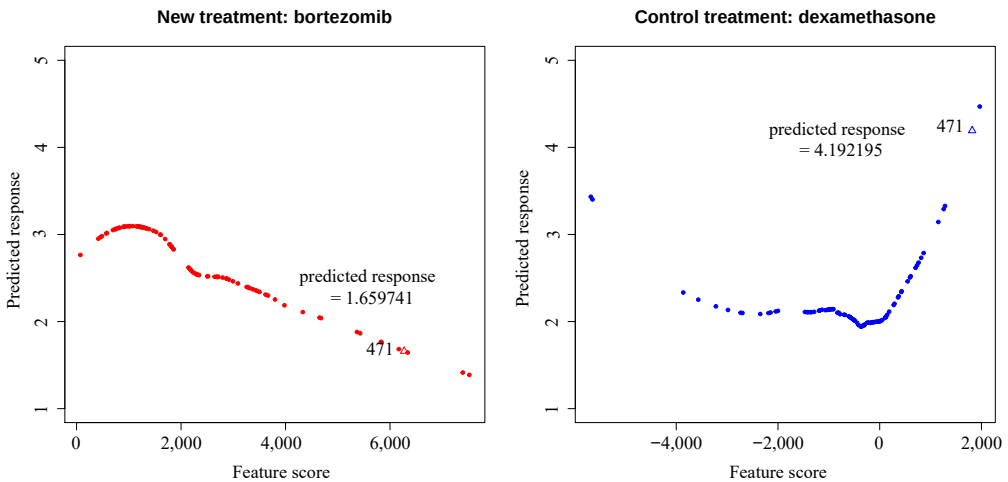

Figure 4. Scatter plot of treatment response versus Feature Score for each treatment group in a test data set of the real data example. A specific data point, patient ID 471, is marked for visualization of the optimal treatment.

This explains the similar result of different methods in Table 1. On the other hand, RWL is computationally expensive, costing about 300 times more than SIR and OLS.

Figure 4 is a plot of the predicted treatment response versus Feature Score in a random test data set. To demonstrate the projection directions $\hat{\beta}$ for Feature Scores, Table 2 shows the specific directions of the two treatment groups calculated from a training data set. There are eight values in $\hat{\beta}$, representing the weights of the corresponding eight genes for Feature Score. A larger absolute value indicates a greater contribution by the respective gene. Table 2 also shows that the two treatment groups have different directions for Feature Score. Despite this difference, we can still directly compare the predicted treatment responses on the vertical axes. A specific data point, patient ID 471, is marked to visualize the treatment recommendation. This patient has a lower predicted treatment response value under bortezomib than under dex. Therefore, the optimal treatment recommendation is the standard chemotherapy dex for this patient. This

Table 2. The estimated projection directions for Feature Score, denoted as $\hat{\beta}_{bort}$ and $\hat{\beta}_{dex}$ of the two treatment groups, in the case study example.

Gene	200010_at	200023_s_at	200082_s_at	200005_at	200017_at	200024_at	200036_s_at	200094_s_at
$\hat{\beta}_{bort}$	-0.038	-0.945	-0.079	0.312	0.013	-0.025	-0.010	0.041
$\hat{\beta}_{dex}$	-0.581	-0.581	0.455	-0.073	-0.203	0.234	0.079	0.103

recommendation is based on the gene expression data through the Feature Score generated by SIR. To conclude, our data-guided method is able to provide patients with multiple myeloma a treatment recommendation between bortezomib and dexamethasone. The data-guided method attempts to connect information from the gene expression data with that from treatment responses, and may reveal relationships between genes and a corresponding phenotype.

6. Discussion

A major advantage of the proposed method lies in its low-dimensional representation of data, i.e., the Feature Score definition, and the automatic detection of these features through the SIR approach. Compared with the lasso-type approaches, such as seen in Qian and Murphy (2011), SIR outperforms the variable selection methods when the effects from individual predictors are minimal. The features from the SIR approach resemble the feature definition of the popular neural network models, with a wide potential of applications. However, the SIR procedure is much simpler than learning a neural network model and does not require a very large sample size.

SIR is a novel method for reducing the dimension of $\mathbf{X}$ without going through any model-fitting process in the first place. It is developed under a very general model assumption that the treatment response Y depends on the covariates $\mathbf{X}$ through a low-dimensional projection space. Assumption (A.2) is the most critical assumption for the theoretical guarantee, whereas the other assumptions are standard. If the true conditional mean function Q_0 does not satisfy Assumption (A.2), we cannot obtain a consistent estimation of the optimal recommendation rule. However, Assumption (A.2) takes the weakest form when we believe a low-dimensional space exists in modeling the response variable Y. The existence of a low-dimensional projection space is also the basic idea behind all dimension reduction methods for high-dimensional data. This general assumption does offer a large approximation space for the true function Q_0, hence likely resulting in a consistent estimation of the optimal recommendation rule.

The proposed method does not consider dynamic treatment regimes that involve treatment recommendations being made at multiple times. However,

because there are far more datasets with only one-time treatment information than there are with multiple-time treatment information, the proposed method can be applied more widely than the methods of dynamic treatment regimes. In addition to recommending treatments, the proposed method can be applied to other precision medicine research, such as risk prediction, treatment effect estimation, and even causal inference. These will be the topics of our future work.

Supplementary Material

The reader is referred to the online Supplementary Materials for the information of data and code and technical proofs.

Acknowledgments

The authors thank the associate editor and the two anonymous referees for their comments and suggestions that have helped improve the paper. The corresponding author also thanks Rena Liu for her editorial assistance.

References

Bai, X., Tsiatis, A. A., Lu, W. and Song, R. (2017). Optimal treatment regimes for survival endpoints using a locally-efficient doubly-robust estimator from a classification perspective. *Lifetime Data Analysis* **23**, 585–604.

Cleveland, W. S. and Devlin, S. J. (1988). Locally weighted regression: An approach to regression analysis by local fitting. *Journal of the American Statistical Association* **83**, 596–610.

Cook, R. D. (1994). On the interpretation of regression plots. *Journal of the American Statistical Association* **89**, 177–189.

Cook, R. D. (2000). SAVE: A method for dimension reduction and graphics in regression. *Communications in Statistics, Part A - Theory and Methods* **29**, 2109–2121.

Cui, Y., Zhu, R. and Kosorok, M. R. (2017). Tree based weighted learning for estimating individualized treatment rules with censored data. *Electronic Journal of Statistics* **11**, 3927–3953.

Dasgupta, S. and Huang, Y. (2020). Selecting biomarkers for building optimal treatment selection rules by using kernel machines. *Journal of the Royal Statistical Society, Series C (Applied Statistics)* **69**, 69–88.

Fan, J. and Lv, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society, Series B (Statistical Methodology)* **70**, 849–911.

Gerstung, M., Papaemmanuil, E., Martincorena, I., Bullinger, L., Gaidzik, V.I., Paschka, P. et al. (2017). Precision oncology for acute myeloid leukemia using a knowledge bank approach. *Nature Genetics* **49**, 332–327.

Guo, W., Zhou, X. and Ma, S. (2021). Estimation of optimal individualized treatment rules using a covariate-specific treatment effect curve with high-dimensional covariates. *Journal of the American Statistical Association* **116**, 309–321.

Hager, R., Tsiatis, A. A. and Davidian, M. (2018). Optimal two-stage dynamic treatment regimes from a classification perspective with censored survival data. *Biometrics* **74**, 1180–1192.

Hofman, I. J., Duin, M., De Bruyne, E., Fancello, L., Mulligan, G., Geerdens, E. et al. (2016). RPL5 on 1p22.1 is recurrently deleted in multiple myeloma and its expression is linked to bortezomib response. *Leukemia* **31**, 1706–1714.

Jiang, B. and Liu, J. S. (2014). Variable selection for general index models via sliced inverse regression. *The Annals of Statistics* **42**, 1751–1786.

Johnstone, I. M. and Lu, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association* **104**, 682–693.

Laber, E. B. and Davidian, M. (2017). Dynamic treatment regimes, past, present and future: A conversation with experts. *Statistical Methods in Medical Research* **26**, 1605–1610.

Li, K. C. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association* **86**, 316–327.

Li, K. C. (1992). On principal Hessian directions for data visualization and dimension reduction: another application of Stein's lemma. *Journal of the American Statistical Association* **87**, 1025–1039.

Lin, Q., Li, X., Huang, D. and Liu, J. S. (2021). On the optimality of sliced inverse regression in high dimensions. *The Annals of Statistics* **49**, 1–20.

Lin, Q., Zhao, Z. and Liu, J. S. (2018). On consistency and sparsity for sliced inverse regression in high dimensions. *The Annals of Statistics* **46**, 580–610.

Lin, Q., Zhao, Z. and Liu, J. S. (2019). Sparse sliced inverse regression via Lasso. *Journal of the American Statistical Association* **114**, 1726–1739.

Liu, Y. and Xie, J. (2019). Cauchy combination test: A powerful test with analytic p-value calculation under arbitrary dependency structures. *Journal of the American Statistical Association* **114**, 384–392.

Liu, Y., Wang, Y., Kosorok, M. R., Zhao, Y. and Zeng, D. (2018). Augmented outcome-weighted learning for estimating optimal dynamic treatment regimens. *Statistics in Medicine* **37**, 3776–3788.

Luckett, D., Laber, E. B., Kahkoska, A. R., Maahs, D. M., Mayer-Davis, E. and Kosorok, M. R. (2020). Estimating Dynamic Treatment Regimes in Mobile Health Using V-Learning. *Journal of the American Statistical Association* **115**, 692–706.

Mitra, A. K., Harding, T., Mukherjee, U .K., Jang, J. S., Li, Y., HongZheng, R. et al. (2017). A gene expression signature distinguishes innate response and resistance to proteasome inhibitors in multiple myeloma. *Blood Cancer Journal* **7**, e581

Mo, W., Qi, Z. and Liu, Y. (2021). Learning optimal distributionally robust individualized treatment rules. *Journal of the American Statistical Association* **116**, 659–674.

Mulligan, G., Mitsiades, C., Bryant, B., Zhan, F., Chng, W., Roels, S. et al. (2007). Gene expression profiling and correlation with outcome in clinical trials of the proteasome inhibitor bortezomib. *Blood* **109**, 3177–3188.

Murray, T. A., Yuan, Y. and Thall, P.F. (2018). A Bayesian machine learning approach for optimizing Dynamic Treatment Regimes. *Journal of the American Statistical Association* **113**, 1255–1267.

Qian, M. and Murphy, S. A. (2011). Performance guarantees for individualized treatment rules. *The Annals of Statistics* **39**, 1180–1210.

Rosenwald, A., Wright, G., Chan, W. C., Connors, J. M., Campo, E., Fisher, R. I. et al. (2002). The use of molecular profiling to predict survival after chemotherapy for diffuse large-B-cell lymphoma. *The New England Journal of Medicine* **346**, 1937–47.

Schulte, P. J., Tsiatis, A. A., Laber, E.,B. and Davidian, M. (2014). Q- and A-learning methods for estimating optimal dynamic treatment regimes. *Statistical Science* **29**, 640–661.

Sulima, S. O. and De Keersmaecker, K. (2017). Ribosomal proteins: A novel class of oncogenic drivers. *Oncotarget* **8**, 89427–89428.

Tan, K., Shi, L. and Yu, Z. (2020). Sparse SIR: Optimal rates and adaptive estimation. *The Annals of Statistics* **48**, 64–85.

van't Veer, L. and Bernards, R. (2008). Enabling personalized cancer medicine through analysis of gene-expression patterns. *Nature* **452**, 564–570.

Xu, Y., Müller, P., Wahed, A. S. and Thall, P. F. (2016). Bayesian nonparametric estimation for dynamic treatment regimes with sequential transition times. *Journal of the American Statistical Association* **111**, 921–950.

Yang, S., Tsiatis, A. A. and Blazing, M. (2018) Modeling survival distribution as a function of time to treatment discontinuation: A dynamic treatment regime approach. *Biometrics* **74**, 900–909.

Zhang, B., Tsiatis, A. A., Laber, E. B. and Davidian, M. (2013). Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika* **3**, 681–694.

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.

Zhao, Y., Zhu, R., Chen, G. and Zheng, Y. (2019). Constructing stabilized dynamic treatment regimes for censored data. *arXiv:1808.01332 [stat.ME]*

Zhou, X., Mayer-Hamblett, N., Khan, U. and Kosorok, M. R. (2017). Residual weighted learning for estimating individualized treatment rules. *Journal of the American Statistical Association* **112**, 169–187.

Zhu, L., Miao, B. and Peng, H. (2006). On sliced inverse regression with high-dimensional covariates. *Journal of the American Statistical Association* **101**, 630–643.

Zhu, J. and Xie, J. (2015). Nonparametric variable selection for predictive models and subpopulations in clinical trials. *Journal of Biopharmaceutical Statistics* **25**, 781–794.

Zhongyuan Chen

Departments of Statistics, Purdue University, West Lafayette, IN 47907, USA.

E-mail: chen3490@purdue.edu

Ziyi Wang

Departments of Statistics, Purdue University, West Lafayette, IN 47907, USA.

E-mail: wang4538@purdue.edu

Qifan Song

Departments of Statistics, Purdue University, West Lafayette, IN 47907, USA.

E-mail: qfsong@purdue.edu

Jun Xie

Departments of Statistics, Purdue University, West Lafayette, IN 47907, USA.

E-mail: junxie@purdue.edu

(Received January 2022; accepted June 2022)