Model-Assisted Uniformly Honest Inference for Optimal Treatment Regimes in High Dimension

Yunan Wu, Lan Wang and Haoda Fu

Abstract

This paper develops new tools to quantify uncertainty in optimal decision making and to gain insight into which variables one should collect information about given the potential cost of measuring a large number of variables. We investigate simultaneous inference to determine if a group of variables is relevant for estimating an optimal decision rule in a high-dimensional semiparametric framework. The unknown link function permits flexible modeling of the interactions between the treatment and the covariates, but leads to nonconvex estimation in high dimension and imposes significant challenges for inference. We first establish that a local restricted strong convexity condition holds with high probability and that any feasible local sparse solution of the estimation problem can achieve the near-oracle estimation error bound. We further rigorously verify that a wild bootstrap procedure based on a debiased

[1]Yunan Wu is Assistant Professor, Department of Mathematical Sciences, University of Texas at Dallas. Emails: yunan.wu@utdallas.edu. Lan Wang is Professor, Department of Management Science, University of Miami. Emails: lanwang@mbs.miami.edu. Wang and Wu's research was supported by NSF DMS-1712706. Dr. Haoda Fu is Research Fellow, Enterprise Lead for Machine Learning and AI, Eli Lilly and Company. Email: fu_haoda@lilly.com. Wang and Wu's research was partly supported by NSF DMS-1952373 and NSF OAC-1940160. The authors are grateful to the referees, the associate editor and the Co-editor for their valuable comments, which have significantly improved the paper.

version of the local solution can provide asymptotically honest uniform inference for the effect of a group of variables on optimal decision making. The advantage of honest inference is that it does not require the initial estimator to achieve perfect model selection and does not require the zero and nonzero effects to be well-separated. We also propose an efficient algorithm for estimation. Our simulations suggest satisfactory performance. An example from a diabetes study illustrates the real application.

Keywords: confidence interval; inference; kernel smoothing; multiplier bootstrap; high-dimensional data; optimal treatment regime; precision medicine.

1 Introduction

Precision medicine is an innovative practice for disease treatment that takes into account individual variability in genes, environment, and lifestyle for each patient. Substantial efforts have recently been devoted to studying how to estimate the optimal personalized treatment regime given the individual-level information, which aims to yield the best expected outcome if the treatment regime is followed by each individual in the population. Several successful approaches have been developed for this estimation problem, including Q-learning and A-learning based methods [Watkins and Dayan, 1992, Robins et al., 2000, Murphy, 2003, Moodie and Richardson, 2010, Qian and Murphy, 2011], and classification-based methods [Zhang et al., 2012, Zhao et al., 2012, 2015, Wang et al., 2018, Qi et al., 2018], among others. We refer to Chakraborty and Moodie [2013] and Kosorok and Moodie [2016] for a general introduction to this area and other relevant references.

Inference or uncertainty quantification is important in practice. This paper studies the following inference problem for optimal personalized decision making: suppose we have

a large number of covariates (e.g., hundreds of genes), how will we determine if a given subset of covariates (e.g., genes associated with a given biological pathway) is relevant for making the optimal treatment recommendation? Scientifically, this knowledge would enable the doctors and researchers to identify critical characteristics (e.g., gender, age, gene pathways) that are influential for the optimal decision. It also helps gain insight into what information is worth collecting to be more cost effective, given the possibility of measuring a large number of variables (genetic, clinic, etc).

In the last few years, important progress has been made in inference with optimal decision rules. Laber et al. [2014] developed a novel locally consistent adaptive confidence interval for the Q-learning approach. Chakraborty et al. [2013] proposed a practically convenient adaptive m-out-of-n bootstrap method for inference for Q-learning. Song et al. [2015] studied penalized Q-learning. Jeng et al. [2018] developed Lasso-based debiased procedure for A-learning. Different but related, Chakraborty et al. [2014] and Luedtke and van der Laan [2016], Zhu et al. [2019] developed confidence intervals for another quantity of interest: the value function. However, existing work mostly deals with the classical asymptotic setting of fixed p and large n, where p is the number of covariates and n is the sample size, and have not addressed the challenge of inference with high-dimensional variables. Moreover, the aforementioned work often assumes that the interaction between the covariates and the treatment has a known functional form.

Motivated by the overarching goal of precision medicine to incorporate genetic information (e.g, measurements on thousands of genes) in the decision making process, this paper investigates inference about the effect of a group of variables on the optimal decision rule in the high-dimensional setting. The existing frameworks are known to face challenges for the purpose of inference in high dimension. The Q-learning approach is prone to model-

misspecification. Robust model-free procedures that directly estimate the Bayes rule (e.g., Zhang et al. [2012]) have a nonstandard convergence rate, see for example, the recent analysis in Wang et al. [2018] on the cubic-root convergence rate. On the other hand, the theory of Hinge-loss based O-learning [Zhao et al., 2012] has been focused on the generalization error bound. Inference for the Bayes rule based on the nonsmooth surrogate loss is very challenging in high dimension. We alleviate the above difficulty by adopting a flexi-ble semiparametric model-assisted approach for optimal decision estimation and inference. The semiparametric structure permits nonparametric main effects and nonlinear interac-tion effect between the covariates and treatment via an unknown smooth link function. This semiparametric framework incorporates many existing models as special cases.

When the interaction effects are nonlinear, the parameter indexing the optimal decision rule does not necessarily correspond to the solution of a convex problem. For inference, we first propose and study a preliminary estimator based on a high-dimensional penalized profile estimation equation. This estimator is motivated by earlier work on classical single-index models (e.g., Powell et al. [1989], Duan and Li [1991], Ichimura [1993], Zhu and Xue [2006], Carroll et al. [1997], Xia et al. [1999], Yu and Ruppert [2002], Wang et al. [2010], Ma and Zhu [2013], Ma and He [2016], among others). Several paper recently studied estimation for high-dimensional single-index models (e.g., Radchenko [2015], Neykov et al. [2016], Yang et al. [2017], Lin et al. [2019], among others) but focused on statistical properties of the global solution which may not be numerically achieved due to the nonconvex nature of the problem. Adopting tools from modern empirical process and random matrix theory, we establish that a local restricted strong convexity condition holds with high probability in high dimension and that any local sparse solution of the penalized estimation equation can achieve desirable estimation accuracy. Moreover, we propose a new algorithm for efficient

computation in high dimension.

Our research also makes new contributions to statistical inference in high-dimensional semiparametric models. Recent work on inference has been mostly limited to linear regression or generalized linear regression, see Zhang and Zhang [2014], Van de Geer et al. [2014], Javanmard and Montanari [2014], Belloni et al. [2015], Cai et al. [2017], Ning et al. [2017], Zhang and Cheng [2017], Zhu and Bradic [2018], Shi et al. [2020], among others. High-dimensional inference in the semiparametric setting with estimated nonparametric components is a substantially harder problem and has been little studied. We have a particularly challenging setting where the parameter of interest and nonparametric component are bundled together, that is, the nuisance functions depend on the parameter of interest [Ding and Nan, 2011]. So far, statistical inference for single-index model has mostly been limited to the lower-dimensional setting (e.g., Liang et al. [2010]), Gueuning and Claeskens [2016]).

Our approach is inspired by the de-biasing (or de-sparsifying) idea proposed in Zhang and Zhang [2014] and Van de Geer et al. [2014], which intuitively can be thought of inverting the Karush-Kuhn-Tucker conditions [Van de Geer et al., 2014]. We generalize this idea to the semiparametric setting and prove that valid honest uniform inference can be obtained based on a debiased version of a local solution. Specifically, we derive simultaneous confidence intervals for inference on a group of variables while allowing the number of covariates to exceed the sample size. The confidence intervals enjoy the honest property in the following sense

$$\sup_{_0:\|_0\|_0\le s}\ \sup_{(0,1)}\ P\left(\sqrt{n}\max_{j\in G}|\hat{e}_{0j}| \le c_{1-}(1-)\right) = o(1);$$

where $_0 = (_{01}; : : : ; _{0p})^T$ is the population parameter indexing the optimal treatment regime, $_j$es denote debiased estimators that will be introduced later, G denotes the group of variables of interest, jj jj_0 denotes the l_0 norm of a vector, and s is a positive integer denoting the sparsity size. The significance of the honest property is that the coverage probability is asymptotically valid uniformly over a class of s-sparse models. An immediate implication is that it relaxes the assumption on signal strength and does not require the zero and nonzero effects to be well-separated (so-called $_{min}$ condition). In particular, this procedure does not require the initial estimator to achieve perfect model selection. It avoids the problems associated with the nonuniformity of the limiting theory for penalized estimators, see discussions in Li [1989], Pötscher [2009], Van de Geer et al. [2014], McKeague and Qian [2015], among others. It is also worth noting that the number of variables in G can be either small or large. For example, one may be interested in assessing how a group of genes corresponding to a particular biological pathway, the size of which can be comparable with or even larger than the sample size, affect optimal decision making. The critical value c_1 is obtained using a wild bootstrap procedure, which automatically accounts for the dependence of the coordinates for testing component-wise hypotheses and leads to more accurate finite-sample performance.

The remainder of the paper is organized as follows. Section 2 introduces the new methodology. Section 3 studies the statistical properties. Section 4 provides the details on computation and reports numerical results from Monte Carlo studies. Section 5 illustrates the new methods on a real data example from a diabetes study. Section 6 discusses some extensions. The regularity conditions, several useful technical lemmas and all the proofs are given in the supplementary material.

2 Methodology

2.1 A Semiparametric Framework

For notational simplicity, we will focus on the binary decision setting. Let $A \in \mathcal{A} = \{0, 1\}$ denote a binary treatment and $x \in \mathcal{X}$ denote a p-dimensional vector of baseline covariates. Let Y denote the outcome of interest. Without loss of generality, we assume a larger value of the outcome is preferred. The observed data consist of $\{(x_i, A_i, Y_i) : i = 1, \ldots, n\}$. We are interested in the setting where $p \gg n$.

A treatment regime is an individualized decision rule that can be represented as a function $d(x) : \mathcal{X} \to \mathcal{A}$. The optimal treatment regime is defined as the decision rule which, if followed by the whole population, will achieve the largest average outcome. Formally, it is defined using the potential outcome framework in causal inference [Neyman, 1990, Rubin, 1978]. Let $Y(a)$ be the potential outcome had the subject been assigned to treatment $a \in \{0, 1\}$. Given a treatment regime $d(x)$, the corresponding potential outcome is $Y(d) = Y(1)d(x) + Y(0)(1 - d(x))$. The optimal treatment regime is defined as $d^{opt}(x) = \arg\max_d E\{Y(d)\}$. It is now well known that $d^{opt}(x) = \arg\max_{a \in \mathcal{A}} E(Y \mid x, A = a)$ [Qian and Murphy, 2011].

This paper considers a flexible semiparametric framework for optimal treatment regime estimation and inference in the high-dimensional setting. Specifically, we assume

$$Y_i = g(x_i) + (A_i - 1/2) f_0(x_i^T \beta_0) + \epsilon_i, \qquad i = 1, \ldots, n, \tag{1}$$

where $\beta_0 = (\beta_{01}, \beta_{02}, \ldots, \beta_{0p})^T$, $g(x_i)$ is the unknown main effect, and $f_0(\cdot)$ is an unknown function that describes the interaction between the treatment and covariates, and the ran-

dom error $_i$ satisfies $E(_i jx_i) = 0$, $i = 1;:::;n$. For identification purpose, we assume that there exists a relevant covariate which has a continuous density given the other covariates [Ichimura, 1993]. Such an identification condition is required even in the lower-dimensional setting when the true model is known. Without loss of generality, we assume that the first covariate x_1 satisfies this condition and normalize its coefficient $_{01}$ such that $_{01} = 1$, see Remark (c) in Section S2 of the online supplementary material for more discussions on the identifiability condition. We denote $B_0 = f = (_1; ;_p)^T : _1 = 1g$ as the candidate set for $_0$. Under model (1), the optimal treatment regime is $d^{opt}(x) = I f_0(x_i _0) > 0 ^T$, where $I()$ denotes the indicator function. Note that the class of index rules are popular in practice due to its interpretability.

Existing work on inference for optimal treatment regime is mostly based on a parametric generative model, which is prone to model misspecification. The semiparametric structure alleviates this difficulty. In particular, it allows for possible nonlinear interaction effects between the covariates and treatment. It also circumvents the curse of dimensionality associated with a fully nonparametric model.

Our goal is to estimate $_0$ and make inference on its components in the high-dimensional setting. In the special case $f_0(u) = u$, which is popularly used in practice, the problem can be formulated as a high-dimensional convex estimation problem. However, when f_0 is nonlinear, it generally leads to a high-dimensional nonconvex problem. Both estimation and inference need to overcome new challenges.

2.2 Profiled Semiparametric Estimation

We start with introducing a penalized profiled semiparametric estimation equation for estimating the parameter indexing the optimal treatment regime. We consider data from

a random experiment, that is, $P(A_i = 0) = P(A_i = 1) = 1/2$, $i = 1, \ldots, n$. Extension to data from observational studies is discussed in Section 6. Inspired by an observation made for the linear model (Tian et al. [2014]), we observe

$$2(2A_i - 1)Y_i = f_0(x_i^T \beta_0) + 2(2A_i - 1)\varepsilon_i + g(x_i). \tag{2}$$

Let $Y_i^e = 2(2A_i - 1)Y_i$ be the modified response, and let $e_i = 2(2A_i - 1)\varepsilon_i + g(x_i)$ be the modified error. We have

$$E\{Y_i^e | x_i\} = f_0(x_i^T \beta_0). \tag{3}$$

In the ideal situation where the link function f_0 is known, we have $\beta_0 = \arg\min E\{Y_i^e - f_0(x_i^T \beta)\}^2$. It is noteworthy that for a nonlinear function f_0, the objective function is usually nonconvex in β. Ichimura [1993] carefully studied the properties of the global minimizer for a semiparametric nonlinear least-squares approach in the classical finite-dimensional setting.

To estimate β_0 in the high-dimensional setting with an known f_0, we consider a penal-ized profiled semiparametric estimation equation. In the ideal situation where f_0 is known a prior, β_0 satisfies the following unbiased estimating equation

$$E\{Y_i^e - f_0(x_i^T \beta_0) f_0'(x_i^T \beta_0) x_i\} = 0, \tag{4}$$

where $f_0'(\cdot)$ denotes the derivative of $f_0(\cdot)$. We will replace the unknown f_0 and f_0' by their respective profiled nonparametric estimator, and consider an appropriately penalized version of the estimated score function to handle the high-dimensional covariates.

We summarize the main steps of estimation as follows. Define $G(t|\cdot) = E\{Y|\beta^T x = t\}$. Note that $G(t|\beta_0) = f_0(t)$. However, when $\beta = \beta_0$, $G(t|\cdot)$ usually has a functional form different from f_0. Ichimura [1993] showed that $\frac{\partial G(x|\cdot)}{\partial \beta} \approx f_0(x_i^T\beta)x_i$ for β close to β_0. Consider the Nadaraya-Watson kernel estimator for $G(t|\cdot)$:

$$\hat{G}(t|\cdot) = \sum_{i=1}^{n} W_{ni}(t;\beta)Y_i, \tag{5}$$

where $K_h(z) = h^{-1}K(z/h)$, and $W_{ni}(t;\beta) = \frac{K_h(t - x_i^T\beta)}{\sum_{j=1}^{n} K_h(t - x_j^T\beta)}$. Write $G^{(1)}(t|\cdot) = \frac{d}{dt}G(t|\cdot)$ and $W_{ni}^{(1)}(t;\beta) = \frac{d}{dt}W_{ni}(t;\beta)$. Then the kernel estimator for the derivative $G^{(1)}(t|\cdot)$ is

$$\hat{G}^{(1)}(t|\cdot) = \sum_{i=1}^{n} W_{ni}^{(1)}(t;\beta)Y_i. \tag{6}$$

Write $G(x^T\beta|\cdot) = E\{Y|x^T\beta\}$. To estimate $G(x^T\beta_j|\cdot)$ and $G^{(1)}(x^T\beta_j|\cdot)$, we employ the following leave-one-out estimators

$$\hat{G}_{-j}(x_j^T\beta) = \sum_{i=1;i\neq j}^{n} W_{nij}(x_j;\beta)Y_i, \quad \hat{G}_{-j}^{(1)}(x_j^T\beta) = \sum_{i=1;i\neq j}^{n} W_{nij}^{(1)}(x_j;\beta)Y_i, \tag{7}$$

where $W_{nij}(x^T;\beta) = \frac{K_h(x_j^T\beta - x_i^T\beta)}{\sum_{k\neq j} K_h(x_j^T\beta - x_k^T\beta)}$, and $W_{nij}^{(1)}(x^T;\beta) = \frac{d}{dt}W_{nij}(t;\beta)\big|_{t=x_j^T\beta}$. Similarly, we estimate $E(x|x^T\beta_0)$ by $E(x_j|x^T\beta) \approx \sum_{i=1;i\neq j}^{n} W_{nij}(x_j;\beta)x_i$. Denote $x_i = (x_{i,1}; x_{i,-1}^T)^T$. Motivated by the semiparaemtric efficient score derived in Liang et al. [2010], we consider the following profiled semiparametric estimating function

$$S_n(\beta; \hat{G}, \hat{E}) = n^{-1}\sum_{i=1}^{n} Y_i - \hat{G}_{-i}(x_i^T\beta)\hat{G}_{-i}^{(1)}(x_i^T\beta)[x_{i,-1} - \hat{E}_{-i}(x_{i,-1}|x_i^T\beta)]. \tag{8}$$

In the high-dimensional setting, the estimating equation $S_n(\beta) = 0$ is ill-posed when $p > n$. Let $\hat{\beta} = (\hat{\beta}_1, \ldots, \hat{\beta}_p)^T = (1, \hat{\beta}_1^T)^T$ be a solution in B_0 that solves the following penalized semiparametric profiled estimating equation

$$S_n(\beta; \hat{G}, \hat{E}) + \lambda \hat{\kappa} = 0, \tag{9}$$

where $\lambda > 0$ is a tuning parameter, $\kappa = (\kappa_2, \ldots, \kappa_p)^T \in \partial||\beta_1||_1$ with $||\beta_1||_1$ denoting the l_1 norm of $\beta_1 = (\beta_2, \ldots, \beta_p)^T$ and $\partial||\beta_1||_1$ denoting the subdifferential of $||\beta_1||_1$, that is $\kappa_j = \text{sign}(\beta_j)$ if $\beta_j \neq 0$, and $\kappa_j \in [-1, 1]$ otherwise, $j = 2, \ldots, p$. In (9), $\hat{G}$ and $\hat{E}$ are evaluated at $\hat{\beta}^T$ the corresponding β in the estimating equations, hence here they stand for $\hat{G}(x_i^T \hat{\beta})$ and $\hat{E}(x_i^T \hat{\beta} | x_i)$, respectively. Note that (9) may have multiple solutions. The theory we develop in Section 3.1 provides a near-optimal error bound for any sparse local solution of the estimating equation. The satisfactory performance of the proposed profiled estimator is demonstrated in the numerical simulations in Section 4.2.

2.3 Inference on the Optimal Decision Rule

To quantify the importance of the covariates on optimal decision making, we will construct confidence intervals for the individual components of $\beta_0 = (1, \beta_{0,2}^T, \beta_1^T)^T$ via debiasing a local solution to the semiparametric estimating equation (9). This generalizes the work of debiased confidence intervals for high-dimensional linear regression in Zhang and Zhang [2014] and Van de Geer et al. [2014] to the semiparametric setting where the initial estimator is an estimating equation solution and an estimated infinite-dimensional functional is present. The theory for semiparametric inference in high dimension is highly nontrivial and is carefully studied in Section 3. We further investigate a wild bootstrap procedure for testing a

general group hypothesis, which aims to achieve accurate finite-sample performance.

Let $\hat{\beta} = (\hat{\beta}_1^T, \hat{\beta}_{-1}^T)^T$ denote a solution satisfying (9). In the high-dimensional linear regression setting, the main idea of debiased estimator is to invert the Karush–Kuhn–Tucker (KKT) condition of the lasso. Inspired by this idea, we consider the following debiased estimator of β_0, β_1:

$$\hat{\beta}_1^d = \hat{\beta}_1 - \hat{\Theta} S_n(\hat{\beta}; G, E), \quad (10)$$

where the $(p-1) \times (p-1)$ matrix $\hat{\Theta}$ is an approximation to the inverse of $\nabla S_n(\hat{\beta}; G, E)$, the derivative matrix of $S_n(\beta; G, E)$ with respect to β_1 evaluated at $\beta = \hat{\beta}$. To construct the approximate inverse $\hat{\Theta}$, we propose a nodewise Dantzig estimator. Specifically, given the initial estimator $\hat{\beta}$ and a positive number λ, for $j = 2, \ldots, p$, define

$$\hat{d}_j(\lambda, \hat{\beta}) = \arg\min_{v \in R^{p-2}} \|v\|_1 \text{ s.t. } \left\| \frac{1}{n} \sum_{i=1}^{n} G^{(1)}(x_i^T \hat{\beta}) (\tilde{x}_{i,j} - \tilde{x}_{i,-j}^T v) \tilde{x}_{i,-j} \right\|_\infty \le \lambda, \quad (11)$$

where $\|\cdot\|_\infty$ denotes the infinity norm of a vector, $x_i = x_i - E(x_i | x_i)$, $\tilde{x}_{i,j}$ denotes the j^{th} entry of the vector x_i, $x_{i,-1}$ denotes the $(p-1)$-subvector of x_i that excludes the 1^{st} entry, and the $x_{i,-j}$ denotes the $(p-2)$-subvector of x_i that excludes the 1^{st} and j^{th} entries. Furthermore, for $j = 2, \ldots, p$, we define

$$\hat{D}(\lambda) = \mathrm{diag}\{\hat{d}_2(\lambda), \ldots, \hat{d}_p(\lambda)\}^T; \quad (12)$$

$$\hat{\tau}_j^2(\lambda) = n^{-1} \sum_{i=1}^{n} G^{(1)}(x_i^T \hat{\beta})^2 (\tilde{x}_{i,j} - \tilde{x}_{i,-j}^T \hat{d}_j(\lambda))^2; \quad (13)$$

$$\hat{\Theta}_j(\lambda) = \hat{\tau}_j^{-2}(\lambda) \hat{d}_j(\lambda), \quad (14)$$

where for a vector $u = (u_1, \ldots, u_p)^T$, given $1 \le i \le j \le p$, $(u)_{i:j}$ returns the subvector $(u_i, \ldots, u_j)^T$, and for any $i > j$, $(u)_{i:j}$ returns the empty vector. For notational simplicity, denote $\hat{d}_j = \hat{d}_j(\lambda)$, $\hat{\beta}_j = \hat{\beta}_j(\lambda)$, and $\hat{d}_{j,j} = \hat{\sigma}_j^2(\lambda)$. The approximate inverse of $r S_n(\lambda, G, E)$ is then constructed as

$$\hat{\Theta} = (\hat{\beta}_2, \ldots, \hat{\beta}_p).$$

The validity of $\hat{\Theta}$ as an approximation to the inverse of $r S_n(\lambda, G, E)$ is given in Lemma 2 of Section 3.2. Section 3 will also present the statistical properties of the debiased estimator $\hat{\theta}_1 = (\hat{\beta}_2, \ldots, \hat{\beta}_p)^T$. This then leads to the following asymptotic $100(1-\alpha)\%$ confidence interval for θ_{0j},

$$\left[\hat{\theta}_j - \Phi^{-1}(1-\alpha/2)\, \hat{\tau}_{jj} n^{-1/2}, \; \hat{\theta}_j + \Phi^{-1}(1-\alpha/2)\, \hat{\tau}_{jj} n^{-1/2} \right], \tag{15}$$

where $j = 2, \ldots, p$, $\Phi^{-1}(\cdot)$ is the quantile function of the standard normal distribution, and $\hat{\tau}_{jj}$ denotes the $(j-1)^{th}$ diagonal entry of $\hat{\Omega}$, with

$$\hat{\Omega} = \hat{\Theta}^T \left(\frac{1}{n} \sum_{i=1}^{n} \frac{\hat{G}^{(1)}(x_i^T j)^2}{[\hat{G}(x_i^T j)(1 - \hat{G}(x_i^T j))]^2} x_{i,-1} x_{i,-1}^T \right) \hat{\Theta}. \tag{16}$$

Corollary 1 in Section 3 justifies the asymptotic uniform validity of this marginal confidence interval.

Next, we consider the following more general simultaneous testing problem

$$H_{0,G} : \theta_{0j} = 0 \text{ for all } j \in G \quad \text{versus} \quad H_{1,G} : \theta_{0j} = 0 \text{ for some } j \in G; \tag{17}$$

where G is a prespecified subset of $\{2, \ldots, p\}$. The size of G may depend on the sample size n. Such a hypothesis naturally arises in the high-dimensional setting. For example, researchers may want to test whether a gene pathway, consisting of multiple genes for the same biological functions, is important for optimal treatment regime recommendation. For this purpose, we propose an effective bootstrap procedure. Although the asymptotic normal distribution of the debiased estimator (see Theorem 2) allows for construction of confidence intervals for individual coefficients (or fixed-dimensional subvector of coefficients), applying it to make inference for groups of variables when the group size diverges (potentially larger than n) is not straightforward. Moreover, confidence intervals based on the asymptotic distribution have been observed to sometimes lead to undercoverage for nonzero coefficients in finite samples. The bootstrap procedure we study automatically accounts for the dependence structure of the variables in the group and provides more accurate critical value.

When deriving the asymptotic property of the debiased estimator (in the proof of Theorem 2), it is observed that the asymptotic property of $\sqrt{n}(\tilde{\beta}_{-1} - \beta_{0,-1})$ is determined by the leading term $\sqrt{n}\hat{\Theta}^T S_n(\beta_0; G; E)$. This suggests that we approximate the distribution of $\sqrt{n}(\tilde{\beta}_j - \beta_{0j})$, $j = 2, \ldots, p$, by the distribution of the following multiplier bootstrap statistic

$$\xi_j \triangleq \frac{1}{\sqrt{n}} \sum_{i=1}^{n} r_i \tilde{Y}_i \left[\hat{G}(x_i^T \hat{\beta}) - \hat{G}^{(1)}(x_i^T \hat{\beta}) x_i^T \hat{\beta}_{-1} \right] \hat{b}_{i,j} \tag{18}$$

where $r_1, \cdots, r_n$ are i.i.d. standard normal random variables, independent of the data. Let $c_{1-\alpha}$ be the upper α-quantile of the distribution of $\max_{j \in G} |\xi_j|$ conditional on the data, which can be easily simulated by generating multiple independent copies of the random weights. We reject the null hypothesis at level α if $\max_{j \in G} |\sqrt{n}\tilde{\beta}_j| > c_{1-\alpha}$. The asymptotic

validity of the bootstrap procedure is formally established in Section 3. Its performance is demonstrated in the numerical simulations in Section 4.2.

3 Statistical Properties

3.1 Theory for Estimation

Making inference about the optimal treatment regime requires an adequate initial estimator for $_0$. To obtain such an initial estimator in the high-dimensional semiparametric framework, a significant challenge is that the corresponding estimation problem is not necessarily convex. To tackle this, we first establish in Lemma 1 below that the estimated $(p \quad 1)$-dimensional gradient $S_n(\ ; \hat{G}; \hat{E})$ in (8) possesses an important local restricted strong convexity property with high probability. Theorem 1 then shows that all local sparse solutions within a small neighborhood of $_0$ enjoy a near-optimal error rate under mild conditions. In the sequel, we use $a \_ b$ to denote $\max(a; b)$, and $a \wedge b$ to denote $\min(a; b)$. Let $s = jj_0jj_0$ be the sparsity size of $_0$, the population parameter indexing the optimal treatment regime.

Lemma 1. (local restricted strong convexity property) Assume conditions (A1)–(A5) in Section S2 of the online supplementary material are satisfied. If $d_0 \frac{s \log(p \_ n)^{1=5}}{n}$ $h < 1$ for some constant $d_0 > 0$, then there exist universal positive constants c_0, c_1, c_2 and $r \quad 1$, which do not depend on n, p and $_0$, such that

$$P\big(S_n(\ ; G; E) \quad S_n(_0; G; E); \ _1 \quad _{0;\ 1} \quad c_0 jj \quad _0 jj^2{}_2 \quad c_1 h^2 jj \quad _0 jj_2; \ 8 \ 2 \ B\big) \quad 1 \quad \exp(\ c_2 \log p);$$

for all n sufficiently large, where $B = f \ 2 \ B_0 : jj \quad _0 jj_2 \quad r; jjjj_0 \quad ksg$ and $k > 1$ is a

positive constant.

Remark 1. Lemma 1 characterizes the local geometry of the profiled score function. For high-dimensional regression with convex loss function such as L_1 penalized linear regression, restricted strong convexity plays an important role on the theory of the regularized estimator [Negahban et al., 2012]. Local restricted strong convexity condition were investigated in Loh and Wainwright [2015] and Mei et al. [2018] for some specific nonconvex loss functions. Those results, however, do not apply to our setting due to the estimated infinite-dimensional parameter.

Theorem 1 below presents non-asymptotic high-probability error bounds for any local sparse solution $\widehat{\beta}$ that satisfies the penalized profiled estimation equation (9).

Theorem 1. Assume conditions (A1)–(A5) in Section S2 of the online supplementary material are satisfied. Suppose $\lambda = d_1 h^2$ for some constant $d_1 > 0$, and $d_0 \frac{s \log(p \vee n)^{1/5}}{n} \le h \le d_0 n^{-1/6}$ for some constant $d_0 > 0$. Then there exist universal positive constants c_0 and c_1 such that for any solution $\widehat{\beta}$ in B, we have

$$\|\widehat{\beta} - \beta_0\|_2 \le \frac{6\sqrt{s}\lambda}{c}, \quad \text{and} \quad \|\widehat{\beta} - \beta_0\|_1 \le \frac{24 s \lambda}{c}$$

with probability at least $1 - \exp(-c_1 \log p)$, for all n sufficiently large.

Remark 2. Theorem 1 shows that under some mild regularity conditions, local solutions of the profiled estimation equation (9) enjoy desirable estimation error rates, same as Lasso does for high-dimensional linear regression. For the purpose of inference, the initial estimator is not require to achieve perfect variable selection. The debiased estimator, however, can achieve the $n^{-1/2}$ rate for each individual coefficient, as we will show in Section 3.2. Carefully going through the proof of the theorem also reveals that the above error bounds

hold uniformly for all $_0$ such that jj_0jj_0 s.

Remark 3. Based on Theorem 1, Lemmas A5–A6 in the online supplement establish the uniform convergence rates for the nonparametric estimator $\hat{G}(x_i^T j)$ and $G^{(1}\hat{h}(x^T j)$, which are of independent interest. Under the assumptions of Theorem 1, we show that there exist universal positive constants c_0 and c_1 such that

$$P\ \max \sup_{2B} \hat{G}(x^T j) \quad G(x^T j)\ c_0h^2\ \exp[\ c_1 \log(p\_n)];$$
$$P\ \max_{1in} \sup \hat{G}^{(1)}(x_i^T j) \quad G^{(1)}(x_i^T j)\ c_0h\ \exp[\ c_1 \log(p\_n)]: \ _{2B}$$

3.2 Theory for Inference

We first introduce some additional notation. Let $x_i = x_i\ E(x_i j x_i^T{}_0)$, and let $x_{i,1}$ denote the $(p\ 1)$-subvector of x_i that excludes its 1st entry. Let $_i\ e\ e^T$ $= E\ [G^{(1)}(x^T{}_0 j_0)]^2 x_{i;\ 1} x_{i;\ 1}$. Assume the $(p\ 1)(p\ 1)$ matrix is positive definite and write its inverse ^{1}, $= (_2;:::;_p)$. For $j = 2;::::;p$, let $_{(j\ 1);\ (j\ 1)} 2 R^{(p\ 2)(p\ 2)}$ be the submatrix of with its $(j\ 1)^{th}$ row and $(j\ 1)^{th}$ column removed; similarly $^2_{(j\ 1);(j\ 1)} 2 R^{p\ 1}$ denotes the $(j\ 1)^{th}$ column of with its $(j\ 1)^{th}$ entry removed. Note that $_{(j\ 1);\ (j\ 1)}$ is positive definite. Define $d_{0j} = (\hat{}_{(j\ 1);\ (j\ 1)})^{\ 1}$ $_{(j\ 1);(j\ 1)}$, $s_j = jjd_{0j}jj_0$, $s = \max_{2jp} s_j$ and $_{0j} = {}_{(j\ 1);(j\ 1)}\ d_{0j}$ $_{(j\ 1);(j\ 1)} = (_{(j\ 1);(j\ 1)})^{\ 1}$, $j = 2;::::;p$.

Lemma 2 below establishes useful properties of the approximate inverse of $rS_n(;G;E)$, defined in Section 2.3.

Lemma 2. Assume the conditions of Theorem 1 are satisfied. Let $= d_2 h$ for some positive constant $d_2 > 0$. If $s\ d_0$ and $d_0 s \frac{\log(p\_n)^{1=5}}{n} h\ d_0 n^{\ 1=6}$ for some constant $d_0 > 0$,

then there exist some universal positive constants d_2, c_0 and c_1 such that results (1)-(3) below hold uniformly in $j = 2;\ldots;p$, with probability at least $1 - \exp(-c_1 \log p)$ for all n sufficiently large:

(1) $\|\hat{d}_j - d_{0j}\|_2 \leq \frac{8 s_j \sqrt{p}}{\phi_2}$ and $\|\hat{d}_j - d_{0j}\|_1 \leq \frac{16 s_j}{\phi_2}$;

(2) $|\hat{\sigma}_{0j}^2 - \hat{b}_j^2 \Sigma \hat{b}_j| \leq c_0 \sqrt{p}\, s_j$, and $|\hat{\sigma}_{0j}^2 - \hat{b}_j^2 \Sigma \hat{b}_j| \leq c_0 \sqrt{p}\, s_j$;

(3) $\|\hat{b}_j - \Sigma \hat{b}_j\|_2 \leq c_0 \sqrt{p}\, s_j$, and $\|\hat{b}_j - \Sigma \hat{b}_j\|_1 \leq c_0 s_j$;

where $\phi_2 > 0$ is the smallest eigenvalue of

.

Lemma 2 requires $s = \max_{2 \leq j \leq p} s_j$ to be of order $O(h^{-1})$. For high-dimensional generalized linear models (Theorem 3.1, Van de Geer et al. [2014]), the corresponding sparsity constraint is $s = o(\sqrt{n}/\log p)$. Our constrain is somewhat stricter due to the need to estimate the infinite-dimensional nuisance parameter. Building on Lemma 2, we prove the statistical property of the debiased estimator $\tilde{\theta}_1$ defined in (10).

Theorem 2. Assume the conditions of Lemma 2 are satisfied. Let $\varepsilon_{n,p} = sh^3 \sqrt{n} + s\sqrt{h}\log p$. Assume $\varepsilon_{n,p} = o(1)$ and $s \log(p \vee n) \leq d_0 n h^5$ for some constant $d_0 > 0$. Then for all n sufficiently large,

$$\sqrt{n}\,(\tilde{\theta}_{0j} - \theta_{0j}) = W_j + \Delta_j, \qquad j = 2;\ldots;p,$$

with

$$W_j = n^{-1/2} e_{j-1}^T \sum_{i=1}^{n} \varepsilon_i \hat{G}^{(1)}(x_{i0}^T \hat{\beta}_0) x_{i,-1},$$

$$P\Big(\max_{2 \leq j \leq p} |\Delta_j| \leq c_0 \varepsilon_{n,p}\Big) \geq 1 - \exp(-c_1 \log p),$$

where c_0, c_1 are universal positive constants, and e_{j-1} denotes the $(p-1)$-dimensional

vector with the $(j-1)^{th}$ entry being one and all the other entries equal to zero.

Remark 4. Theorem 2 suggests that if we consider a lower-dimensional linear combination of coefficients $\alpha^T\beta_{0,-1}$, where α is a $(p-1)$-dimensional nonzero vector of constants, then $\alpha^T(\hat{\beta}_{-1}-\beta_{0,-1})$ has the asymptotic distribution $N(0,\alpha^T\Sigma\alpha)$ with $\Sigma = E[\dot{G}^{(1)}(x_{i,1}^T\beta_0)]^2 x_{i,-1}x_{i,-1}^T$. The asymptotic covariance matrix resembles that in the literature on profiled estimation for index models in lower dimension, see Liang et al. [2010], Ma and He [2016], among other. The assumption $\delta_{n,p} = o(1)$ is a sufficient condition for the remaining term of the linear approximation of $\sqrt{n}(\hat{\beta}_j/\hat{\beta}_1 - \beta_{0j}/\beta_{01})$ to be uniformly negligible. It still allows p to grow at an exponential rate of n.

Remark 5. The proof Theorem 2 is given in the online supplement. To build the theory, we show that

$$\begin{aligned}\sqrt{n}(\hat{\beta}_{-1}-\beta_{0,-1}) &= \sqrt{n}\Sigma^{-1}S_n(\beta_0;G;E) + \sqrt{n}(I_{p-1}-\Sigma^{-1}J_1)(\hat{\beta}_{-1}-\beta_{0,-1}) \\ &\quad \sqrt{n}\Sigma^{-1}[S_n(\hat{\beta};\hat{G};\hat{E}) - S_n(\beta_0;G;E) - J_1(\hat{\beta}_{-1}-\beta_{0,-1})] \\ &\triangleq A_{n1}+A_{n2}+A_{n3},\end{aligned}$$

where $J_1 = n^{-1}\sum_{i=1}^n[\dot{G}^{(1)}(x_i^T\beta)]^2 x_{i,-1}x_{i,-1}^T$ is the leading term in the approximation to $\nabla S_n(\beta_0;G;E)$. In the proof, we carefully justify that: (1) The $(j-1)^{th}$ component of A_{n1} can be approximated by W_j in the theorem, for $2 \le j \le p$, (2) $P(\|A_{n2}\|_1 \ge c_0\delta_{n,p}) \le \exp(-c_1\log p)$, and (3) $P(\|A_{n3}\|_1 \ge c_0\delta_{n,p}) \le \exp(-c_1\log p)$, for some positive constants c_0 and c_1. Furthermore, to provide a deeper insight into the extension into the semiparametric setting, we consider the Gateaux functional derivative of the estimating function with respect to the infinite-dimensional nuisance parameters. Consider

the functional $M(z;\ ;G;E)_e = [Y \quad G(x^T j)G^{(1)}(x^T j)[x_{\ 1} \quad E(x_{\ 1} j x^T)]$; where $z = (A; X; Y_e)$ denotes a vector of random observations of the data. The Gateaux derivative of $M(z;\ ;G;E)$ at G in the direction $[G\_\_ \quad G]$ is defined as

$$\lim_{\ \to 0} \frac{EM(z;\ ;G + (G \quad \_\_ \quad G);E) \quad M(z;\ ;G;E)}{\ } :$$

It is easy to see that this Gateaux derivative at G is zero when evaluated at $= {}_0$. Similarly, the Gateaux derivative with respect to E vanishes at the true value ${}_0$. This orthogonality behavior suggests the insensitivity of the estimating function to the infinite-dimensional nuisance parameters.

The following corollary establishes uniform validity of the marginal confidence intervals (15) introduced in Section 2.3.

Corollary 1. Under the conditions of Theorem 2,

$$\sup_{{}_0 2 B_0 : jj_0 jj_0 s} \max_{2 j p} \sup_{2(0;1)} P\sqrt{n}(j \quad e \quad {}_{0j}) \ \hat{}_{jj}^{1=2} \quad {}^{1}(1 \qquad =2) \qquad (1 \quad) = o(1);$$

where $\hat{}_{jj}$ denotes the $(j \quad 1)^{th}$ diagonal entry of $\hat{}()$ defined in Section 2.3, and ${}^{1}()$ is the quantile function of $N(0; 1)$.

Finally, Theorem 3 below establishes the validity of the bootstrap procedure introduced in Section 2.3 for testing the group hypothesis (17). Given a group of variables $G \ f2; : : : ; pg$, the wild bootstrap test statistic is defined as $\sqrt{n} \max_{j 2 G} j_j j$, where $_j \ , \ n^{\ 1} \sum_{i=1}^{n} r_i f Y_i \quad \hat{G}(x_i^T j)\hat{}\hat{G}^{(1)}(x_i^T j) x_{ij}$, and $r_1; \ ; r_n$ are standard normal random variables that are independent of the data. Denote $r = fr_1; \ ; r_n g$, and let $w = fw_1; \ ; w_n g$ denote the random sample $w_i = (A_i; x_i; Y_i)$. Given $0 < \ < 1$, recall that the

bootstrap critical value for a level- test is defined as

$$c_1 = \inf\left\{ t \in R : P\left(\sqrt{n} \max_{j \in G} |j| \le t \,|\, \{w_i\}_{i=1}^n \right) \ge 1 \right\} : \quad (19)$$

Theorem 3. Assume the conditions of Theorem 2 are satisfied. If $_{n;p}\sqrt{\log p} = o(1)$, h $d_0 \frac{s\log(p\_n)}{n}^{1=5}$ for some constant $d_0 > 0$, and $\sqrt{s}h\log^2 p = o(1)$, then

$$\sup_{_0 \in B_0 : |j_0| j_0 s \in (0;1)} \sup P\left(\sqrt{n} \max_{j \in G} |e_{0j}| \le c_{1} (G) \right) (1) = o(1):$$

Theorem 3 ensures that the multiplier bootstrap procedure is valid for the simultaneous testing problem (17). It is also honest in the sense of being valid uniformly over s-sparse models of the form (1). It does not require the nonzero components of $_0$ to be well-separated from zero. In particular, the multiple bootstrap procedure does not require the local solution of the profiled estimation to achieve perfect variable selection, which is usually unrealistic in practice.

4 Monte Carlo Studies

4.1 Algorithm for Estimation

To solve the penalized high-dimensional profiled estimating equation for the initial estimator ,b we extend the composite gradient algorithm [Nesterov, 2007, Agarwal et al., 2012] for high-dimensional M-estimator without nuisance parameters. A summary of the proposed algorithm is given in Algorithm 1 in Section S10.1 of the supplementary material.

Specifically, given a current estimator $^t = (1; (^t{}_1)^T)^T$ at step t, we update the esti-

mate by

$$\beta_1^{t+1} = \arg\min_{\beta_1 \in \mathbb{R}^p, \|\beta_1\|_1 \le \gamma} \left\{ \frac{\mu}{2} \|\beta_1 - \beta_1^t\|_2^2 + \frac{1}{2} [S_n(\beta^t; G, E)]^T (\beta_1 - \beta_1^t) + \lambda \|\beta_1\|_1 \right\}, \tag{20}$$

where μ is the step size, γ is a positive constant such that $\|\beta_{0,1}\|_1 \le \gamma$. An appealing practical property of the algorithm is that the update in step (20) can be done efficiently by the following formula:

$$\beta_1^{t+1} = T_s\left(\beta_1^t - \frac{1}{\mu} S_n(\beta^t; G, E), \lambda\right), \tag{21}$$

where the function $T_s(\beta_1, \lambda) = \mathrm{sgn}(\beta_j) \max(|\beta_j| - \lambda, 0)_{j=2,\ldots,p}$ is the soft-threshold operator. Then to ensure the constraint $\|\beta_1\|_1 \le \gamma$, we employ the projection method introduced in Duchi et al. [2008], which is described in Algorithm 2 in Section S10.1 of the online supplement.

In implementation, we choose the kernel function $K(\cdot)$ as the distribution function of the standard normal distribution. The bandwidth is set to be $h = 0.9 n^{-1/6} \min\{\mathrm{std}(x_i^T), \mathrm{IQR}(x_i^T)/1.34\}$, as motivated by Silverman [1986], where "std" denotes the standard deviation, and "IQR" denotes the interquartile range. For the step-size parameter, inspired by Agarwal et al. [2012], we employ an expanding series for μ, which ensures that the stepsize diminishes during the update process. Given a set of candidate tuning parameters $\{\lambda_k\}$ and the corresponding estimators $\hat{\beta}_k$, we employ 5-fold cross-validation to select the optimal tuning parameter λ by minimizing $\mathrm{MSE}(\lambda) = n^{-1} \sum_{i=1}^n \{Y_i - \hat{G}(x_i^T; \hat{\beta}_\lambda)\}^2$.

To obtain the debiased estimator $\hat{\beta}$, the nodewise Dantzig estimator $d_j(\cdot)$ in (11) is computed via linear programming, see details in Section S10.2 of the supplementary

material.

4.2 Monte Carlo Results

We generate random data from the model $Y = (x^T)^2 + \frac{1}{}(A \quad _2)f_0(x^T{}_0) + \ ,$ where $N(0; 1)$, A Bernoulli(0:5), and x follows a p-dimensional multivariate normal distribution with mean zero and identity covariance matrix, $= (0{:}5; 0{:}5; \ 0{:}5; \ 0{:}5; 0; \ ; 0)^T$, $_0 = (1; \ 1; 0{:}5; 0{:}4; \ 0{:}3; 0; \ ; 0)^T$, and $f_0(u) = 20 \ f[1 + \exp(\ u)] \ ^1 \ 0{:}5g$. We consider $n = 300; 500$ and $p = 200; 800; 2000$ in the Monte Carlo experiment.

We first investigate the finite-sample performance of the penalized profiled semiparametric estimator in Section 2.2. Table 1 reports the average l_1- and l_2-estimation errors, the average number of false negatives (nonzero components incorrectly identified as zero) and false positives (zero components incorrectly identified as nonzero), with their standard errors in the parentheses, based on 500 simulation runs. Results in Table 1 demonstrate satisfactory performance of the profiled estimator for both the scenarios $p < n$ and $p > n$.

Table 1: Performance of the penalized profile least-squares estimator

n	p	l_1 error	l_2 error	False Negative	False Positive
300	200	0.85 (0.02)	0.31 (0.00)	0.01 (0.01)	10.95 (0.32)
	800	1.10 (0.03)	0.37 (0.00)	0.07 (0.01)	19.05 (1.13)
	2000	1.32 (0.03)	0.40 (0.00)	0.09 (0.01)	31.25 (1.57)
500	200	0.58 (0.01)	0.22 (0.00)	0.00 (0.00)	9.30 (0.30)
	800	0.79 (0.02)	0.27 (0.00)	0.00 (0.00)	17.39 (0.66)
	2000	0.94 (0.02)	0.31(0.00)	0.01 (0.00)	25.60 (1.18)

Next we investigate the wild bootstrap procedure introduced in Section 2.3 for testing the group hypothesis (17). We consider the following six different choices for the groups: $G_1 = f6; 7; 8; 9g$, $G_2 = f5; 6; 7; 8; 9g$, $G_3 = f4; 6; 7; 8; 9g$, $G_4 = f4; 5; 6; 7; 8; 9g$,

$G_5 = \{3; 6; 7; 8; 9\}$ and $G_6 = \{2; 6; 7; 8; 9\}$. Note that G_1 consists of only zero entries in $_0$, while all the other groups include at least one non-zero elements. Table 2 summarizes the average Type I errors and powers for each scenario, based on 1000 Bootstrap samples and 500 simulation runs.

Table 2: Performance of the bootstrap procedure in Section 2.3 for simultaneous testing.

n	p	Type I error	Power				
		G_1	G_2	G_3	G_4	G_5	G_6
300	200	5.6%	96.4%	96.2%	97.8%	98.6%	100%
	800	5.4%	94.6%	97.6%	99.0%	99.6%	100%
	2000	3.2%	92.4%	96.8%	98.4%	99.0%	100%
500	200	4.4%	100%	100%	100%	100%	100%
	800	5.0%	99.6%	99.6%	100%	99.2%	100%
	2000	4.6%	98.8%	98.6%	99.0%	99.2%	100%

Table 2 indicates that type I errors are reasonable controlled for all scenarios. Power performance generally depends on the number and magnitudes of the nonzero components. The hypothesis corresponding to G_2 represents a more challenging situation where the only non-zero element is -0:3, close to 0. The average powers for this case for different values of p are still over 90%.

Note that for inference, we need to estimate the approximate inverse of $r S_n(\hat{p})$ which involves an additional tuning parameter . We observe that the inference procedure is not overly sensitive to its choice and fix it at the value = 25h to save computational time. Alternatively, it can also be selected via cross-validation similarly as what has been done for selection. We provide additional simulation results in Section S10.3 of the online supplement, including investigation on the choice of and comparing with alternative procedures for estimating the optimal value function.

5 A Real Data Example

We illustrate the application on a clinical data set introduced by Charbonnel et al. [2005]. This is a randomized, double-blind, parallel treatment arm, phase III clinical trial to compare the efficacy and safety of pioglitazone versus gliclazide on metabolic control in naive patients with Type 2 diabetes mellitus. This data set we consider contains information on clinical characteristics for 813 individuals with Type 2 diabetes. The patients were randomized into two treatment arms: pioglitazone (treatment 0) and gliclazide (treatment 1). Their glycosylated haemoglobin A_{1c} (HbA_{1c}) and fasting plasma glucose (FPG) levels were recorded every four weeks, up to week 52.

The primary efficacy endpoint is the change of HbA_{1c} from baseline to the last available post-treatment value. We consider the main effects of 22 baseline covariates and their two-way interactions in the model. The dimension of the model is over 250. In the analysis, we standardize the covariates to have mean zero and sample variance one.

We consider testing the significance of six different groups of variables. Table 3 summarizes these six different groups and their respective p-values, based on the bootstrap procedure in Section 2.3. The estimated coefficients are reported in Section S10.3 of the supplementary.

Table 3: Real data analysis: evaluation of the significance of different groups of variables

Group	Variables	p value
1	HbA_{1c}, creatinine, BMI, waist circumference, HomaS	0:003
2	all variables in Group 1, all their two-way interactions, and their interactions with fasting insulin	0:011
3	HbA_{1c}, HomaS	$<$ 0:001
4	BMI, creatinine, waist circumference,	0:242
5	LDL-C, total cholesterol, age, weight	0:494

Based on the scientific literature and suggestions from our clinical collaborators, fasting insulin is important for estimating the optimal treatment regime. We normalize its coefficient as 1 in our model. The first group includes the main effects of five characteristics, which are the baseline average levels for HbA_{1c}, creatinine, BMI, waist circumference and homeostatic model assessment insulin sensativity (HomaS). The variables in this group are those identified by diabetes experts to be potentially important for optimal treatment regime estimation. The bootstrap procedures suggests a significant p-value (0.003) for this group, which indicates that at one variable in this group is influential for making an optimal personalized decision in the choice of the two treatments. Group 2 augments Group 1 by including all the two-way interaction of these six characteristics (including fasting insulin), hence includes 20 variables in total. The estimated p value is 0.011. Group 3 and Group 4 are subgroups of Group 1. The third group only includes two main effects: baseline HbA_{1c} and HomaS, while the fourth group includes the remaining three main effects. The estimated p values suggest that the significant characteristics are among those in Group 3 rather than Group 4. Group 5 consists of four variables: the baseline average levels for the low-density lipoprotein cholesterol (LDL-C), total cholesterol, age and weight. This group of variables is of interest because Glucose and lipid metabolism are linked to each other in many ways [Parhofer, 2015]. Age and weight are also always taken into account for optimal treatment regime estimation. Our test suggests that Group 5 does not appear to be influential in optimal treatment recommendation.

6 Discussions

We propose a flexible semi-parametric approach for making honest simultaneous inference about the importance of a group of variables on optimal treatment regime estimation. We develop new statistical theory to overcome the challenges of nonconvexity, high dimensionality and infinite-dimensional nonparametric components.

In this paper, we focus on a randomized trial. For observation studies, let $(x) = P(A = 1|x)$ be the propensity score. Observing that $Ef[A \quad (x)]g(x)g = 0$, we have

$$4[A_i \quad (x_i)]Y_i = 4[A_i \quad (x_i)]g(x_i) + 4[A_i \quad (x_i)](A_i \quad 1=2)f_0(x_i^T{}_0) + 4[A_i \quad (x_i)]_i:$$

Let $Y_i^e = 4[A_i \quad (x_i)]Y_i$, $e_i = 4[A_i \quad (x_i)][_i + g(x_i)]$, then we have

$$EY_i^e = 4[A_i \quad (x_i)](A_i \quad 1=2)f_0(x_i^T{}_0):$$

Denote $G(t|j) = E(Y^e|x^T = t) = 2Ef[A \quad (x)](2A \quad 1)f_0(x^T{}_0)|x^T = tg$, $G^{(1)}(t|j) = \frac{d}{dt}G(t|j)$, and define $G(t|\hat{}_j)$, $G^{(1)}(t|\hat{}_j)$ similarly as in Section 2.2. Assume $(x) = P(A = 1|x)$ can be modeled as $(x;)$, where is a finite-dimensional parameter. Let be an estimate of , such as the one based on the regularized logistic regression. Define the profiled semiparametric estimating function $S_n(; G; E;) = n^{1} \sum_{i=1} f4[A_i \quad (x_i;)]Y_i \quad \hat{G}(x_i^T j)gG^{(1)}(x^T j)[x_{i;1} \quad E(x_{i;1}|x_i^T)]$. We then estimate $_0$ through the following penalized semiparametric profiled estimating equation $S_n(; \hat{G}; \hat{E};) + = 0$. Promising numerical performance of this estimator is reported in Section S10.3 of the supplementary. Our approach can still be applied to investigate the theory but is it more complex due to the additional nuisance parameter. We will explore the complete theory for the above

estimator in the future work. Alternative approaches that can potentially be extended to our setting include Nie and Wager [2020], Künzel et al. [2018], among ohers.

Our approach for high-dimensional inference generalizes the “inverting K K T condition" technique in Van de Geer et al. [2014]. An alternative approach, which is more suitable if one is interested in some targeted lower-dimensional parameter is based on the idea of orthogonalization, see for example Belloni et al. [2015], Ning et al. [2017], Chernozhukov et al. [2018]. In contrast, our approach is able to achieve debiasing for the p-dimensional coefficient vector simultaneously. The main idea of the orthogonalization approach is to construct a lower-dimensional estimating equation which is locally insensitive to the nuisance parameters. The construction of such a lower-dimensional moment condition is non-trivial for high-dimensional semiparametric setting, particularly for index model, where the challenge of bundled parameter arises.

References

A. Agarwal, S. Negahban, and M. J. Wainwright. Fast global convergence of gradient methods for high-dimensional statistical recovery. The Annals of Statistics, 40(5):2452–2482, 2012.

A. Belloni, V. Chernozhukov, and K. Kato. Uniform post-selection inference for least absolute deviation regression and other z-estimation problems. Biometrika, 102(1):77–94, 2015.

T. T. Cai, Z. Guo, et al. Confidence intervals for high-dimensional linear regression: Minimax rates and adaptivity. The Annals of Statistics, 45(2):615–646, 2017.

R. J. Carroll, J. Fan, I. Gijbels, and M. P. Wand. Generalized partially linear single-index models. Journal of the American Statistical Association, 92(438):477–489, 1997.

B. Chakraborty and E. E. Moodie. Statistical Methods for Dynamic Treatment Regimes: Reinforcement Learning, Causal Inference, and Personalized Medicine. Springer Science & Business Media, 2013.

B. Chakraborty, E. B. Laber, and Y. Zhao. Inference for optimal dynamic treatment regimes using an adaptive m-out-of-n bootstrap scheme. Biometrics, 69(3):714–723, 2013.

B. Chakraborty, E. B. Laber, and Y.-Q. Zhao. Inference about the expected performance of a data-driven dynamic treatment regime. Clinical Trials, 11(4):408–417, 2014.

B. H. Charbonnel, D. R. Matthews, G. Schernthaner, M. Hanefeld, P. Brunetti, and on behalf of the QUARTET Study Group. A long-term comparison of pioglitazone and gliclazide in patients with type 2 diabetes mellitus: a randomized, double-blind, parallel-group comparison trial. Diabetic Medicine, 22(4):399–405, 2005.

V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins. Double/debiased machine learning for treatment and structural parameters. The Econometrics Journal, 21(1):C1–C68, 2018.

Y. Ding and B. Nan. A sieve M-theorem for bundled parameters in semiparametric models, with application to the efficient estimation in a linear model for censored data. The Annals of Statistics, 39(6):3032–3061, 2011.

N. Duan and K.-C. Li. Slicing regression: A link-free regression method. The Annals of Statistics, 19(2):505–530, 1991.

J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra. Efficient projections onto the l 1-ball for learning in high dimensions. In Proceedings of the 25th international conference on Machine Learning, pages 272–279, 2008.

T. Gueuning and G. Claeskens. Confidence intervals for high-dimensional partially linear single-index models. Journal of Multivariate Analysis, 149:13–29, 2016.

H. Ichimura. Semiparametric least squares (sls) and weighted sls estimation of single-index models. Journal of Econometrics, 58(1-2):71–120, 1993.

A. Javanmard and A. Montanari. Confidence intervals and hypothesis testing for high-dimensional regression. Journal of Machine Learning Research, 15(1):2869–2909, 2014.

X. J. Jeng, W. Lu, H. Peng, et al. High-dimensional inference for personalized treatment decision. Electronic Journal of Statistics, 12(1):2074–2089, 2018.

M. R. Kosorok and E. E. Moodie. Adaptive Treatment Strategies in Practice: Planning Trials and Analyzing Data for Personalized Medicine. ASA-SIAM Series on Statistics and Applied Probability, SIAM, Philadelphia, ASA, Alexandria, VA., 2016.

S. R. Künzel, S. J. S. Walter, and J. S. Sekhon. Causaltoolbox—estimator stability for heterogeneous treatment effects. arXiv: Methodology, 2018.

E. B. Laber, D. J. Lizotte, M. Qian, W. E. Pelham, and S. A. Murphy. Dynamic treatment regimes: Technical challenges and applications. Electronic Journal of Statistics, 8(1): 1225 – 1272, 2014.

K.-C. Li. Honest Confidence Regions for Nonparametric Regression. The Annals of Statistics, 17(3):1001–1008, 1989.

H. Liang, X. Liu, R. Li, and C.-L. Tsai. Estimation and testing for partially linear single-index models. The Annals of Statistics, 38(6):3811–3836, 2010.

Q. Lin, Z. Zhao, and J. S. Liu. Sparse sliced inverse regression via lasso. Journal of the American Statistical Association, 114(528):1726–1739, 2019.

P.-L. Loh and M. J. Wainwright. Regularized m-estimators with nonconvexity: Statistical and algorithmic theory for local optima. Journal of Machine Learning Research, 16: 559–616, 2015.

A. R. Luedtke and M. J. van der Laan. Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy. The Annals of Statistics, 44(2):713–742, 2016.

S. Ma and X. He. Inference for single-index quantile regression models with profile optimization. The Annals of Statistics, 44(3):1234–1268, 2016.

Y. Ma and L. Zhu. Doubly robust and efficient estimators for heteroscedastic partially linear single-index models allowing high dimensional covariates. Journal of the Royal Statistical Society: Series B, 75(2):305–322, 2013.

I. W. McKeague and M. Qian. An adaptive resampling test for detecting the presence of significant predictors. Journal of the American Statistical Association, 110(512):1422–1433, 2015.

S. Mei, Y. Bai, and A. Montanari. The landscape of empirical risk for nonconvex losses. The Annals of Statistics, 46(6A):2747–2774, 2018.

E. E. Moodie and T. S. Richardson. Estimating optimal dynamic regimes: Correcting bias under the null. Scandinavian Journal of Statistics, 37(1):126–146, 2010.

S. A. Murphy. Optimal dynamic treatment regimes. Journal of the Royal Statistical Society: Series B, 65(2):331–366, 2003.

S. N. Negahban, P. Ravikumar, M. J. Wainwright, B. Yu, et al. A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. Statistical Science, 27(4):538–557, 2012.

Y. Nesterov. Gradient methods for minimizing composite objective function. Core discussion papers, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE), 2007.

M. Neykov, J. S. Liu, and T. Cai. L1-regularized least squares for support recovery of high dimensional single index models with gaussian designs. The Journal of Machine Learning Research, 17(1):2976–3012, 2016.

J. Neyman. On the application of probability theory to agricultural experiments. Essay on principles. Section 9. Statistical Science, 5(4):465–472, 1990.

X. Nie and S. Wager. Quasi-oracle estimation of heterogeneous treatment effects. Biometrika, 09 2020.

Y. Ning, H. Liu, et al. A general theory of hypothesis tests and confidence regions for sparse high dimensional models. The Annals of Statistics, 45(1):158–195, 2017.

K. G. Parhofer. Interaction between glucose and lipid metabolism: more than diabetic dyslipidemia. Diabetes & Metabolism Journal, 39(5):353–362, 2015.

B. M. Pötscher. Confidence sets based on sparse estimators are necessarily large. Sankhyā: The Indian Journal of Statistics, Series A, pages 1–18, 2009.

J. L. Powell, J. H. Stock, and T. M. Stoker. Semiparametric estimation of index coefficients. Econometrica, pages 1403–1430, 1989.

Z. Qi, Y. Liu, et al. D-learning to estimate optimal individual treatment rules. Electronic Journal of Statistics, 12(2):3601–3638, 2018.

M. Qian and S. A. Murphy. Performance guarantees for individualized treatment rules. The Annals of Statistics, 39(2):1180–1210, 2011.

P. Radchenko. High dimensional single index models. Journal of Multivariate Analysis, 139:266–282, 2015.

J. Robins, M. Hernan, and B. Brumback. Marginal structural models and causal inference in epidemiology. Epidemiology, 11:550–560, 2000.

D. B. Rubin. Bayesian inference for causal effects: the role of randomization. The Annals of Statistics, 6:34–58, 1978.

C. Shi, R. Song, W. Lu, and R. Li. Statistical inference for high-dimensional models via recursive online-score estimation. Journal of the American Statistical Association, 2020.

B. Silverman. Density Estimation for Statistics and Data Analysis. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis, 1986.

R. Song, W. Wang, D. Zeng, and M. Kosorok. Penalized q-learning for dynamic treatment regimens. Statistica Sinica, 25:901–920, 2015.

L. Tian, A. A. Alizadeh, A. J. Gentles, and R. Tibshirani. A simple method for estimat-ing interactions between a treatment and a large number of covariates. Journal of the American Statistical Association, 109(508):1517–1532, 2014.

S. Van de Geer, P. Bühlmann, Y. Ritov, R. Dezeure, et al. On asymptotically optimal confidence regions and tests for high-dimensional models. The Annals of Statistics, 42 (3):1166–1202, 2014.

J.-L. Wang, L. Xue, L. Zhu, Y. S. Chong, et al. Estimation for a partial-linear single-index model. The Annals of Statistics, 38(1):246–274, 2010.

L. Wang, Y. Zhou, R. Song, and B. Sherwood. Quantile-optimal treatment regimes. Journal of the American Statistical Association, 113(523):1243–1254, 2018.

C. J. Watkins and P. Dayan. Q-learning. Machine Learning, 8(3-4):279–292, 1992.

Y. Xia, H. Tong, and W. K. Li. On extended partially linear single-index models. Biometrika, 86(4):831–842, 1999.

Z. Yang, K. Balasubramanian, and H. Liu. High-dimensional non-gaussian single index models via thresholded score function estimation. In International Conference on Machine Learning, pages 3851–3860, 2017.

Y. Yu and D. Ruppert. Penalized spline estimation for partially linear single-index models. Journal of the American Statistical Association, 97(460):1042–1054, 2002.

B. Zhang, A. A. Tsiatis, E. B. Laber, and M. Davidian. A robust method for estimating optimal treatment regimes. Biometrics, 68(4):1010–1018, 2012.

C.-H. Zhang and S. S. Zhang. Confidence intervals for low dimensional parameters in high dimensional linear models. Journal of the Royal Statistical Society: Series B, 76(1): 217–242, 2014.

X. Zhang and G. Cheng. Simultaneous inference for high-dimensional linear models. Journal of the American Statistical Association, 112(518):757–768, 2017.

Y. Zhao, D. Zeng, A. J. Rush, and M. R. Kosorok. Estimating individualized treatment rules using outcome weighted learning. Journal of the American Statistical Association, 107(499):1106–1118, 2012.

Y. Zhao, D. Zeng, E. B. Laber, and M. R. Kosorok. New statistical learning methods for estimating optimal dynamic treatment regimes. Journal of the American Statistical Association, 110:583–598, 2015.

L. Zhu and L. Xue. Empirical likelihood confidence regions in a partially linear single-index model. Journal of the Royal Statistical Society: Series B, 68(3):549–570, 2006.

W. Zhu, D. Zeng, and R. Song. Proper inference for value function in high-dimensional q-learning for dynamic treatment regimes. Journal of the American Statistical Association, 114(527):1404–1417, 2019.

Y. Zhu and J. Bradic. Linear hypothesis testing in dense high-dimensional linear models. Journal of the American Statistical Association, 113(524):1583–1600, 2018.