



HHS Public Access

Author manuscript

J Am Stat Assoc. Author manuscript; available in PMC 2022 January 01.

Published in final edited form as:

J Am Stat Assoc. 2021 ; 116(534): 984–998. doi:10.1080/01621459.2019.1699421.

Global and Simultaneous Hypothesis Testing for High-Dimensional Logistic Regression Models

Rong Ma, T. Tony Cai, Hongzhe Li

University of Pennsylvania Perelman School of Medicine, Philadelphia, 19104 United States

The Wharton School - Univ of Pennsylvania, Philadelphia, 19104 United States

University of Pennsylvania School of Medicine, 215 Blockley Hall, Philadelphia, 19104 United States

Abstract

High-dimensional logistic regression is widely used in analyzing data with binary outcomes. In this paper, global testing and large-scale multiple testing for the regression coefficients are considered in both single- and two-regression settings. A test statistic for testing the global null hypothesis is constructed using a generalized low-dimensional projection for bias correction and its asymptotic null distribution is derived. A lower bound for the global testing is established, which shows that the proposed test is asymptotically minimax optimal over some sparsity range. For testing the individual coefficients simultaneously, multiple testing procedures are proposed and shown to control the false discovery rate (FDR) and falsely discovered variables (FDV) asymptotically. Simulation studies are carried out to examine the numerical performance of the proposed tests and their superiority over existing methods. The testing procedures are also illustrated by analyzing a data set of a metabolomics study that investigates the association between fecal metabolites and pediatric Crohn's disease and the effects of treatment on such associations.

Keywords

False discovery rate; Global testing; Large-scale multiple testing; Minimax lower bound

1 INTRODUCTION

Logistic regression models have been applied widely in genetics, finance, and business analytics. In many modern applications, the number of covariates of interest usually

Corresponding author Hongzhe Li hongzhe@upenn.edu.

Rong Ma is PhD Candidate, Department of Biostatistics, Epidemiology and Informatics, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104.

T. Tony Cai is Daniel H. Silberberg Professor of Statistics, Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104.

Hongzhe Li is Professor of Biostatistics and Statistics, Department of Biostatistics, Epidemiology and Informatics, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104.

SUPPLEMENTARY MATERIALS

In the online Supplemental Materials, we prove Theorem 3, 5, Proposition 1, and the technical lemmas. The technical results and simulations concerning the two-sample tests discussed in Section 4 are also included.

grows with, and sometimes far exceeds, the number of observed samples. In such high-dimensional settings, statistical problems such as estimation, hypothesis testing, and construction of confidence intervals become much more challenging than those in the classical low-dimensional settings. The increasing technical difficulties usually emerge from the non-asymptotic analysis of both statistical models and the corresponding computational algorithms.

In this paper, we consider testing for high-dimensional logistic regression model:

$$\log\left(\frac{\pi_i}{1-\pi_i}\right) = X_i^\top \beta, \quad \text{for } i = 1, \dots, n. \quad (1)$$

where $\beta \in \mathbb{R}^p$ is the vector of regression coefficients. The observations are i.i.d. samples $Z_i = (y_i, X_i)$ for $i = 1, \dots, n$, and we assume $y_i | X_i \sim \text{Bernoulli}(\pi_i)$ independently for each $i = 1, \dots, n$.

1.1 Global and Simultaneous Hypothesis Testing

It is important in high-dimensional logistic regression to determine 1) whether there are any associations between the covariates and the outcome and, if yes, 2) which covariates are associated with the outcome. The first question can be formulated as testing the global null hypothesis $H_0: \beta = 0$; and the second question can be considered as simultaneously testing the null hypotheses $H_{0,i}: \beta_i = 0$ for $i = 1, \dots, p$. Besides such single logistic regression problems, hypothesis testing involving two logistic regression models with regression coefficients $\beta^{(1)}$ and $\beta^{(2)}$ in \mathbb{R}^p is also important. Specifically, one is interested in testing the global null hypothesis $H_0: \beta^{(1)} = \beta^{(2)}$, or identifying the differentially associated covariates through simultaneously testing the null hypotheses $H_{0,i}: \beta_i^{(1)} = \beta_i^{(2)}$ for each $i = 1, \dots, p$.

Estimation for high-dimensional logistic regression has been studied extensively. van de Geer (2008) considered high-dimensional generalized linear models (GLMs) with Lipschitz loss functions, and proved a non-asymptotic oracle inequality for the empirical risk minimizer with the Lasso penalty. Meier, van de Geer, and Bühlmann (2008) studied the group Lasso for logistic regression and proposed an efficient algorithm that leads to statistically consistent estimates. Negahban et al. (2010) obtained the rate of convergence for the ℓ_1 -regularized maximum likelihood estimator under GLMs using restricted strong convexity property. Bach (2010) extended tools from the convex optimization literature, namely self-concordant functions, to provide interesting extensions of theoretical results for the square loss to the logistic loss. Plan and Vershynin (2013) connected sparse logistic regression to one-bit compressed sensing and developed a unified theory for signal estimation with noisy observations.

In contrast, hypothesis testing and confidence intervals for high-dimensional logistic regression have only been recently addressed. van de Geer et al. (2014) considered constructing confidence intervals and statistical tests for single or low-dimensional components of the regression coefficients in high-dimensional GLMs. Mukherjee, Pillai, and Lin (2015) studied the detection boundary for minimax hypothesis testing in high-

dimensional sparse binary regression models when the design matrix is sparse. Belloni, Chernozhukov, and Wei (2016) considered estimating and constructing the confidence regions for a regression coefficient of primary interest in GLMs. More recently, Sur, Chen, and Candès (2017) and Sur and Candès (2019) considered the likelihood ratio test for high-dimensional logistic regression under the setting that $p/n \rightarrow \kappa$ for some constant $\kappa < 1/2$, and showed that the asymptotic null distribution of the log-likelihood ratio statistic is a rescaled χ^2 distribution. Cai et al. (2017) proposed a global test and a multiple testing procedure for differential networks against sparse alternatives under the Markov random field model. Nevertheless, the problems of global testing and large-scale simultaneous testing for high-dimensional logistic regression models with $p \gtrsim n$ remain unsolved.

In this paper, we first consider global and multiple testing for a single high-dimensional logistic regression model. The global test statistic is constructed as the maximum of squared standardized statistics for individual coefficients, which are based on a two-step standardization procedure. The first step is to correct the bias of the logistic Lasso estimator using a generalized low-dimensional projection (LDP) method, and the second step is to normalize the resulting nearly unbiased estimators by their estimated standard errors. We show that the asymptotic null distribution of the test statistic is a Gumbel distribution and that the resulting test is minimax optimal under the Gaussian design by establishing the minimax separation distance between the null space and alternative space. For large-scale multiple testing, data-driven testing procedures are proposed and shown to control the false discovery rate (FDR) and falsely discovered variables (FDV) asymptotically. The framework for testing for single logistic regression is then extended to the setting of testing two logistic regression models.

The main contributions of the present paper are threefold.

1. We propose novel procedures for both the global testing and large-scale simultaneous testing for high dimensional logistic regressions. The dimension p is allowed to be much larger than the sample size n . Specifically, we require $\log p = O(n^{c_1})$ for the global test and $p = O(n^{c_2})$ for the multiple testing procedure, with some constant $c_1, c_2 > 0$. For the global alternatives characterized by the ℓ_∞ norm of the regression coefficients, the global test is shown to be minimax rate optimal with the optimal separation distance of order $\sqrt{\log p/n}$.
2. Following similar ideas in Ren, Zhang, and Zhou (2016) and Cai et al. (2017), our construction of the test statistics depends on a generalized version of the LDP method for bias correction. The original LDP method (Zhang and Zhang 2014) relies on the linearity between the covariates and outcome variable. For logistic regression, the generalized approach first finds a linearization of the regression function, and the weighted LDP is then applied. Besides its usefulness in logistic regression, the generalized LDP method is flexible and can be applied to other nonlinear regression problems (see Section 7 for a detailed discussion).
3. The minimax lower bound is obtained for the global hypothesis testing under the Gaussian design. The lower bound depends on the calculation of the χ^2 -divergence between two logistic regression models. To the best of our

knowledge, this is the first lower bound result for high-dimensional logistic regression under the Gaussian design.

1.2 Other Related Work

We should note that a different but related problem, namely inference for high-dimensional linear regression, has been well studied in the literature. Zhang and Zhang (2014), van de Geer et al. (2014) and Javanmard and Montanari (2014a,b) considered confidence intervals and testing for low-dimensional parameters of the high-dimensional linear regression model and developed methods based on a two-stage debiased estimator that corrects the bias introduced at the first stage due to regularization. Cai and Guo (2017) studied minimaxity and adaptivity of confidence intervals for general linear functionals of the regression vector.

The problems of global testing and large-scale simultaneous testing for high-dimensional linear regression have been studied by Liu and Luo (2014), Ingster, Tsybakov, and Verzelen (2010) and more recently by Xia, Cai, and Cai (2018) and Javanmard and Javadi (2019). However, due to the nonlinearity and the binary outcome, the approaches used in these works cannot be directly applied to logistic regression problems. In the Markov random field setting, Ren, Zhang, and Zhou (2016) and Cai et al. (2017) constructed pivotal/test statistics based on the debiased LDP estimators for node-wise logistic regressions with binary covariates. However, the results for sparse high-dimensional logistic regression models with general continuous covariates remain unknown.

Other related problems include joint testing and false discovery rate control for high-dimensional multivariate regression (Xia, Cai, and Li 2018) and testing for high-dimensional precision matrices and Gaussian graphical models (Liu 2013; Xia, Cai, and Cai 2015), where the inverse regression approach and de-biasing were carried out in the construction of the test statistics. Such statistics were then used for testing the global null with extreme value type asymptotic null distributions or to perform multiple testing that controls the false discovery rate.

1.3 Organization of the Paper and Notations

The rest of the paper is organized as follows. In Section 2, we propose the global test and establish its optimality. Some comparisons with existing works are made in detail. In Section 3, we present the multiple testing procedures and show that they control the FDR/FDP or FDV/FWER asymptotically. The framework is extended to the two-sample setting in Section 4. In Section 5, the numerical performance of the proposed tests are evaluated through extensive simulations. In Section 6, the methods are illustrated by an analysis of a metabolomics study. Further extensions and related problems are discussed in Section 7. In Section 8, some of the main theorems are proved. The proofs of other theorems as well as technical lemmas, and some further discussions are collected in the online Supplementary Materials.

Throughout our paper, for a vector $\mathbf{a} = (a_1, \dots, a_n)^\top \in \mathbb{R}^n$, we define the ℓ_p norm $\|\mathbf{a}\|_p = (\sum_{i=1}^n |a_i|^p)^{1/p}$, and the ℓ_∞ norm $\|\mathbf{a}\|_\infty = \max_{1 \leq j \leq n} |a_j|$. $\mathbf{a}_{-j} \in \mathbb{R}^{n-1}$ stands for the subvector of \mathbf{a} without the j th component. We denote $\text{diag}(a_1, \dots, a_n)$ as the $n \times n$ diagonal

matrix whose diagonal entries are a_1, \dots, a_n . For a matrix $A \in \mathbb{R}^{p \times q}$, $\lambda_i(A)$ stands for the i -th largest singular value of A and $\lambda_{\max}(A) = \lambda_1(A)$, $\lambda_{\min}(A) = \lambda_{p \wedge q}(A)$. For a smooth function $f(x)$ defined on \mathbb{R} , we denote $\dot{f}(x) = df(x)/dx$ and $\ddot{f}(x) = d^2f(x)/dx^2$. Furthermore, for sequences $\{a_n\}$ and $\{b_n\}$, we write $a_n = o(b_n)$ if $\lim_{n \rightarrow \infty} a_n/b_n = 0$, and write $a_n = O(b_n)$, $a_n \lesssim b_n$ or $b_n \gtrsim a_n$ if there exists a constant C such that $a_n \leq Cb_n$ for all n . We write $a_n \asymp b_n$ if $a_n \lesssim b_n$ and $a_n \gtrsim b_n$. For a set A , we denote $|A|$ as its cardinality. Lastly, C, C_0, C_1, \dots are constants that may vary from place to place.

2 GLOBAL HYPOTHESIS TESTING

In this section, we consider testing the global null hypotheses

$$H_0: \beta = 0 \text{ vs. } H_1: \beta \neq 0,$$

under the logistic regression model with random designs. The global testing problem corresponds to the detection of any associations between the covariates and the outcome.

Our construction of the global testing procedure begins with a bias-corrected estimator built upon a regularized estimator such as the ℓ_1 -regularized M-estimator. For high-dimensional logistic regression, the ℓ_1 -regularized M-estimator is defined as

$$\hat{\beta} = \operatorname{argmin}_{\beta} \left\{ \frac{1}{n} \sum_{i=1}^n \left[-y_i \beta^\top X_i + \log(1 + e^{\beta^\top X_i}) \right] + \lambda \|\beta\|_1 \right\}, \tag{2}$$

which is the minimizer of a penalized log-likelihood function. Negahban et al. (2010) showed that, when X_j are i.i.d. sub-gaussian, under some mild regularity conditions, standard high-dimensional estimation error bounds for $\hat{\beta}$ under the ℓ_1 or ℓ_2 norm can be obtained by choosing $\lambda \asymp \sqrt{\log p/n}$. Once we obtain the initial estimator $\hat{\beta}$, our next step is to correct the bias of $\hat{\beta}$.

For technical reasons, we split the samples so that the initial estimation step and the bias correction step are conducted on separate and independent datasets. Without loss of generality, we assume there are $2n$ samples, divided into two subsets \mathcal{D}_1 and \mathcal{D}_2 , each with n independent samples. The initial estimator $\hat{\beta}$ is obtained from \mathcal{D}_1 . In the following, we construct a nearly unbiased estimator $\check{\beta}$ based on $\hat{\beta}$ and the samples from \mathcal{D}_2 , using the generalized LDP approach. Throughout the paper, the samples $Z_i = (X_i, Y_i)$, $i = 1, \dots, n$, are from \mathcal{D}_2 , which are independent of $\hat{\beta}$. We would like to emphasize that the sample splitting procedure is only used to simplify our theoretical analysis, which does not make it a restriction for practical applications. Numerically, as our simulations in Section 5 show, sample splitting is in fact not needed in order for our methods perform well (see further discussions in Section 7).

2.1 Construction of the Test Statistic via Generalized Low-Dimensional Projection

Let X be the design matrix whose i -th row is X_i . We rewrite the logistic regression model defined by (1) as

$$y_i = f(\beta^\top X_i) + \epsilon_i \quad (3)$$

where $f(u) = e^u / (1 + e^u)$ and ϵ_i is error term. To correct the bias of the initial estimator $\hat{\beta}$, we consider the Taylor expansion of $f(u_i)$ at \hat{u}_i for $u_i = \beta^\top X_i$ and $\hat{u}_i = \hat{\beta}^\top X_i$

$$f(u_i) = f(\hat{u}_i) + \dot{f}(\hat{u}_i)(u_i - \hat{u}_i) + Re_i$$

where Re_i is the reminder term. Plug this into the regression model (3), we have

$$y_i - f(\hat{u}_i) + \dot{f}(\hat{u}_i)X_i^\top \hat{\beta} = \dot{f}(\hat{u}_i)X_i^\top \beta + (Re_i + \epsilon_i). \quad (4)$$

By rewriting the logistic regression model as (4), we can treat $y_i - f(\hat{u}_i) + \dot{f}(\hat{u}_i)X_i^\top \hat{\beta}$ on the left hand side as the new response variable, whereas $\dot{f}(\hat{u}_i)X_i$ as the new covariates and $Re_i + \epsilon_i$ as the noise. Consequently, β can be considered as the regression coefficient of this approximate linear model.

The bias-corrected estimator, or, the generalized LDP estimator $\check{\beta}$ is defined as

$$\check{\beta}_j = \hat{\beta}_j + \frac{\sum_{i=1}^n v_{ij}(y_i - f(\hat{\beta}^\top X_i))}{\sum_{i=1}^n v_{ij}\dot{f}(\hat{\beta}^\top X_i)X_{ij}}, \quad j = 1, \dots, p, \quad (5)$$

where X_{ij} is the j -th component of X_i and $v_j = (v_{1j}, v_{2j}, \dots, v_{nj})$ is the score vector that will be determined carefully (Ren, Zhang, and Zhou 2016; Cai et al. 2017). More specifically, we define the weighted inner product $\langle \cdot, \cdot \rangle_n$ for any $a, b \in \mathbb{R}^n$ as $\langle a, b \rangle_n = \sum_{i=1}^n \dot{f}(\hat{u}_i) a_i b_i$, and denote $\langle \cdot, \cdot \rangle$ as the ordinary inner product defined in Euclidean space. Combining (4) and (5), we can write

$$\check{\beta}_j - \beta_j = \frac{\langle v_j, \epsilon \rangle}{\langle v_j, \mathbf{x}_j \rangle_n} + \frac{\langle v_j, Re \rangle}{\langle v_j, \mathbf{x}_j \rangle_n} - \frac{\langle v_j, \mathbf{h}_{-j} \rangle_n}{\langle v_j, \mathbf{x}_j \rangle_n}, \quad (6)$$

where $\mathbf{x}_j \in \mathbb{R}^n$ denote the j -th column of X , $\mathbf{h}_{-j} = X_{-j}(\hat{\beta}_{-j} - \beta_{-j})$ where $X_{-j} \in \mathbb{R}^n \times \mathbb{R}^{p-1}$ is the submatrix of X without the j -th column, and $Re = (Re_1, \dots, Re_n)$ with $Re_i = f(u_i) - f(\hat{u}_i) - \dot{f}(\hat{u}_i)(u_i - \hat{u}_i)$. We will construct score vector v_j so that the first term on the right hand side of (6) is asymptotically normal, while the second and third terms, which together contribute to the bias of the generalized LDP estimator $\hat{\beta}_j$, are negligible.

To determine the score vector v_j efficiently, we consider the following node-wise regression among the covariates

$$\mathbf{x}_j = X_{-j}\gamma_j + \eta_j, \quad j = 1, \dots, p, \tag{7}$$

where $\gamma_j = \operatorname{argmin}_{\gamma \in \mathbb{R}^{p-1}} \mathbb{E}[\|\mathbf{x}_j - X_{-j}\gamma\|_2^2]$ and η_j is the error term. Intuitively, if we set $v_j = \widehat{W}^{-1}\eta_j$ for $\widehat{W} = \operatorname{diag}(\dot{f}(\hat{u}_1), \dots, \dot{f}(\hat{u}_n))$, then it should follow that

$$\langle v_j, \mathbf{h}_{-j} \rangle_n \leq \max_{k \neq j} |\langle v_j, \mathbf{x}_k \rangle_n| \cdot \|\hat{\beta} - \beta\|_1 = \max_{k \neq j} |\langle \eta_j, \mathbf{x}_k \rangle| \cdot \|\hat{\beta} - \beta\|_1 \approx 0.$$

In practice, we use the node-wise Lasso to obtain an estimate of η_j . For X from \mathcal{D}_2 and $\hat{\beta}$ obtained from \mathcal{D}_1 , the score v_j is obtained by calibrating the Lasso-generated residue $\hat{\eta}_j$, i.e.

$$\begin{aligned} v_j(\lambda) &= \widehat{W}^{-1}\hat{\eta}_j(\lambda), \quad \hat{\eta}_j(\lambda) = \mathbf{x}_j - X_{-j}\hat{\gamma}_j(\lambda), \\ \hat{\gamma}_j(\lambda) &= \operatorname{argmin}_b \left\{ \frac{\|\mathbf{x}_j - X_{-j}b\|_2^2}{2n} + \lambda\|b\|_1 \right\}. \end{aligned} \tag{8}$$

Clearly, $v_j(\lambda)$ depends on the tuning parameter λ . Define the following quantities

$$\zeta_j(\lambda) = \max_{k \neq j} \frac{|\langle v_j(\lambda), \mathbf{x}_k \rangle_n|}{\|v_j(\lambda)\|_n}, \quad \tau_j(\lambda) = \frac{\|v_j(\lambda)\|_n}{|\langle v_j(\lambda), \mathbf{x}_j \rangle_n|}. \tag{9}$$

The tuning parameter λ can be determined through $\zeta_j(\lambda)$ and $\tau_j(\lambda)$ by the algorithm in Table 1, which is adapted from the algorithm in Zhang and Zhang (2014).

Once $\check{\beta}_j$ and τ_j are obtained, we define the standardized statistics

$$M_j = \check{\beta}_j / \tau_j,$$

for $j = 1, \dots, p$. The global test statistic is then defined as

$$M_n = \max_{1 \leq j \leq p} M_j^2. \tag{10}$$

2.2 Asymptotic Null Distribution

We now turn to the analysis of the properties of the global test statistic M_n defined in (10). For the random covariates, we consider both the Gaussian design and the bounded design. Under the Gaussian design, the covariates are generated from a multivariate Gaussian distribution with an unknown covariance matrix $\Sigma \in \mathbb{R}^{p \times p}$. In this case, we assume

(A1). $X_j \sim N(0, \Sigma)$ independently for each $i = 1, \dots, n$.

In the case of bounded design, we assume instead

(A2). X_j for $j = 1, \dots, n$ are i.i.d. random vectors satisfying $\mathbb{E}X_j = 0$ and $\max_{1 \leq j \leq n} \|X_j\|_\infty \leq T$ for some constant $T > 0$.

Define the ℓ_1 ball

$$\mathcal{B}_1(k) = \left\{ \mathbf{\Omega} = (\omega_{ij}) \in \mathbb{R}^{p \times p} : \max_{1 \leq i \leq p} \sum_{j=1}^p \min(|\omega_{ij}| \sqrt{\frac{n}{\log p}}, 1) \leq k \right\}.$$

In general, $\mathcal{B}_1(k)$ includes any matrix $\mathbf{\Omega}$ whose rows ω_i are ℓ_0 sparse with $\|\omega_i\|_0 \leq k$ or ℓ_1 sparse with $\|\omega_i\|_1 \leq k\sqrt{\log p/n}$ for all $i = 1, \dots, p$. The parameter space of the covariance matrix $\mathbf{\Sigma}$ and the regression vector β are defined as following.

(A3). The parameter space $\Theta(k)$ of $\theta = (\beta, \mathbf{\Sigma}) \in \mathbb{R}^p \times \mathbb{R}^{p \times p}$ satisfies

$$\Theta(k) = \{(\beta, \mathbf{\Sigma}) : \|\beta\|_0 \leq k, M^{-1} \leq \lambda_{\min}(\mathbf{\Sigma}) \leq \lambda_{\max}(\mathbf{\Sigma}) \leq M, \mathbf{\Sigma}^{-1} \in \mathcal{B}_1(k)\},$$

for some constant $M \geq 1$. For convenience, we denote $\Theta_1(k) = \{\beta \in \mathbb{R}^p : \|\beta\|_0 \leq k\}$ and $\Theta_2(k) = \{\mathbf{\Sigma} \in \mathbb{R}^{p \times p} : M^{-1} \leq \lambda_{\min}(\mathbf{\Sigma}) \leq \lambda_{\max}(\mathbf{\Sigma}) \leq M, \mathbf{\Sigma}^{-1} \in \mathcal{B}_1(k)\}$, so that $\Theta(k) = \Theta_1(k) \times \Theta_2(k)$.

The following theorem states that the asymptotic null distribution of M_n under either the Gaussian or bounded design is a Gumbel distribution.

Theorem 1. *Let M_n be the test statistic defined in (10), D be the diagonal of $\mathbf{\Sigma}^{-1}$ and $(\zeta_{ij}) = D^{-1/2} \mathbf{\Sigma}^{-1} D^{-1/2}$. Suppose $\max_{1 \leq i < j \leq p} |\zeta_{ij}| \leq c_0$ for some constant $0 < c_0 < 1$, $\log p = \mathcal{O}(n^r)$ for some $0 < r < 1/5$, and*

1. *under the Gaussian design, we assume (A1) (A3) and $k = o(\sqrt{n}/\log^3 p)$; or*
2. *under the bounded design, we assume (A2) (A3) and $k = o(\sqrt{n}/\log^{5/2} p)$.*

Then under H_0 , for any given $x \in \mathbb{R}$,

$$P_\theta(M_n - 2\log p + \log \log p \leq x) \rightarrow \exp\left(\frac{1}{\sqrt{\pi}} \exp(-x/2)\right), \text{ as } (n, p) \rightarrow \infty.$$

The condition that $\log p = \mathcal{O}(n^r)$ for some $0 < r < 1/5$ is consistent with those required for testing the global hypothesis in high-dimensional linear regression (Xia, Cai, and Cai 2018) and for testing two-sample covariance matrices (Cai, Liu, and Xia 2013). It allows the dimension p to be exponentially large comparing to the sample size n , which is much more flexible than the likelihood ratio test considered in Sur, Chen, and Candès (2017) and Sur and Candès (2019), where the dimension can only scale as $p < n$. Under the Gaussian design,

it is required that the sparsity k is $o(\sqrt{n}/\log^3 p)$ whereas for the bounded design, it suffices that the sparsity k to be $o(\sqrt{n}/\log^{5/2} p)$.

Remark 1. The analysis can be extended to testing $H_0: \beta_G = 0$ versus $H_1: \beta_G \neq 0$ for a given index set G . Specifically, we can construct the test statistic as $M_{G,n} = \max_{j \in G} M_j^2$ and obtain a similar Gumbel limiting distribution by replacing p by $|G|$, as $(n, |G|) \rightarrow \infty$. The sparsity condition thus should be forwarded to the set G .

Based on the limiting null distribution, the asymptotically α level test can be defined as

$$\Phi_\alpha(M_n) = I\{M_n \geq 2\log p - \log\log p + q_\alpha\},$$

where q_α is the $1 - \alpha$ quantile of the Gumbel distribution with the cumulative distribution function $\exp\left(-\frac{1}{\sqrt{x}}\exp(-x/2)\right)$, i.e.

$$q_\alpha = -\log(\alpha) - 2\log\log(1 - \alpha)^{-1}.$$

The null hypothesis H_0 is rejected if and only if $\Phi_\alpha(M_n) = 1$.

2.3 Minimax Separation Distance and Optimality

In this subsection, we answer the question: “What is the essential difficulty for testing the global hypothesis in logistic regression.” To fix ideas, we begin with defining the minimax separation distance that measures such an essential difficulty for testing the global null hypothesis at a given level and type II error. In particular, we consider the alternative

$$H_1: \beta \in \{\beta \in \mathbb{R}^p: \|\beta\|_\infty \geq \rho, \|\beta\|_0 \leq k\}$$

for some $\rho > 0$. This alternative concerns the detection of any discernible signals among the regression coefficients where the signals can be extremely sparse, which has interesting applications (see Xia, Cai, and Cai (2015)). Similar alternatives are also considered by Cai, Liu, and Xia (2013) and Cai, Liu, and Xia (2014).

By fixing a level $\alpha > 0$ and a type II error probability $\delta > 0$, we can define the δ -separation distance of a level α test procedure Φ_α for given design covariance Σ as

$$\begin{aligned} \rho(\Phi_\alpha, \delta, \Sigma) &= \inf\left\{\rho > 0: \inf_{\beta \in \Theta_1(k), \|\beta\|_\infty \geq \rho} P_\beta(\Phi_\alpha = 1) \geq 1 - \delta\right\} \\ &= \inf\left\{\rho > 0: \sup_{\beta \in \Theta_1(k), \|\beta\|_\infty \geq \rho} P_\beta(\Phi_\alpha = 0) \leq \delta\right\}. \end{aligned} \tag{11}$$

The δ -separation distance $\rho(\Phi_\alpha, \delta, \Theta(k))$ over $\Theta(k)$ can thus be defined by taking the supremum over all the covariance matrices $\Sigma \in \Theta_2(k)$, so that

$$\rho(\Phi_\alpha, \delta, \Theta(k)) = \sup_{\Sigma \in \Theta_2(k)} \rho(\Phi_\alpha, \delta, \Sigma),$$

which corresponds to the minimal ℓ_∞ distance such that the null hypothesis H_0 is well separated from the alternative H_1 by the test Φ_α . In general, δ -separation distance is an analogue of the statistical risk in estimation problems. It characterizes the performance of a specific α -level test with a guaranteed type II error δ . Consequently, we can define the (α, δ) -minimax separation distance over $\Theta(k)$ and all the α -level tests as

$$\rho^*(\alpha, \delta, \Theta(k)) = \inf_{\Phi_\alpha} \rho(\Phi_\alpha, \delta, \Theta(k)).$$

The definition of (α, δ) -minimax separation distance generalizes the ideas of Ingster (1993), Baraud (2002) and Verzelen (2012). The following theorem establishes the minimax lower bound of the (α, δ) -separation distance under the Gaussian design for testing the global null hypothesis over the parameter space $\Theta'(k) \subset \Theta(k)$ defined as

$$\Theta'(k) = \left(\Theta_1(k) \cap \left\{ \beta \in \mathbb{R}^p : \|\beta\|_2 \lesssim (n^{1/4} \log p)^{-1} \right\} \right) \times \Theta_2(k).$$

Theorem 2. *Assume that $\alpha + \delta \leq 1$. Under the Gaussian design, if (A1) and (A3) hold, $(\beta, \Sigma) \in \Theta'(k)$ and $k \lesssim \min\{p^\gamma, \sqrt{n}/\log^3 p\}$ for some $0 < \gamma < 1/2$, then the (α, δ) -minimax separation distance over $\Theta'(k)$ has the lower bound*

$$\rho^*(\alpha, \delta, \Theta'(k)) \geq c \sqrt{\frac{\log p}{n}} \tag{12}$$

for some constant $c > 0$.

In order to show the above lower bound is asymptotically sharp, we prove that it is actually attainable under certain circumstances, by our proposed global test Φ_α . In particular, for the bounded design, we make the following additional assumption.

(A4). It holds that $P_\theta(\max_{1 \leq j \leq n} |\beta^\top X_j| \leq C) = O(p^{-c})$ for some constant $C, c > 0$.

Theorem 3. *Suppose that $\log p = O(n^r)$ for some $0 < r < 1$. Under the alternative $H_1: \|\beta\|_\infty \geq c_2 \sqrt{\log p/n}$ for some $c_2 > 0$, and*

(i) *under the Gaussian design, assume that (A1) and (A3) hold, $\|\beta\|_2 \leq C(\log \log p)/\sqrt{\log n}$ for $C \leq \min\left\{\sqrt{2/\lambda_{\max}(\Sigma)}, (2r\sqrt{2\lambda_{\max}(\Sigma)})^{-1}\right\}$, $\log p \gtrsim \log^{1+\delta} n$ for some $\delta > 0$ and $k = o(\sqrt{n}/\log^3 p)$;*
 or

(ii) *under the bounded design, assume that (A2), (A3), and (A4) hold, and $k = o(\sqrt{n}/\log^{5/2} p)$.*

Then we have $P_\theta(\Phi_\alpha(M_n) = 1) \rightarrow 1$ as $(n, p) \rightarrow \infty$.

In Theorem 3, (A4) is assumed for the bounded case and $\|\beta\|_2 = O(\log \log p / \sqrt{\log n})$ is required for the Gaussian case. In particular, since $\log p = O(n^r)$ for some $0 < r < 1$, the upper bound $\log \log p / \sqrt{\log n}$ for $\|\beta\|_2$ can be as large as $\sqrt{\log n}$. In Theorem 2, the minimax lower bound is established over $(\beta, \Sigma) \in \Theta'(k)$, so that the same lower bound holds over a larger set

$$(\beta, \Sigma) \in (\Theta_1(k) \cap \{\beta \in \mathbb{R}^p: \|\beta\|_2 \leq \log \log p / \sqrt{\log n}\}) \times \Theta_2(k), \tag{13}$$

since $\log \log p / \sqrt{\log n} \gtrsim (n^{1/4} \log p)^{-1}$. On the other hand, Theorem 3 (i) indicates an upper bound $\rho^* \lesssim \sqrt{\log p / n}$ attained by our proposed test under the Gaussian design over the set (13). These two results imply the minimax rate $\rho^* \asymp \sqrt{\log p / n}$ and the minimax optimality of our proposed test over the set (13).

2.4 Comparison with Existing Works

In this section, we make detailed comparisons and connections with some existing works concerning global hypothesis testing in the high-dimensional regression literature.

Ingster, Tsybakov, and Verzelen (2010) addressed the detection boundary for high-dimensional sparse linear regression models, and more recently Mukherjee, Pillai, and Lin (2015) studied the detection boundary for hypothesis testing in high-dimensional sparse binary regression models. However, although both works obtained the sharp detection boundary for the global testing problem $H_0: \beta = 0$, their alternative hypotheses are different from ours. Specifically, Mukherjee, Pillai, and Lin (2015) considered the alternative hypothesis $H_1: \beta \in \{\beta \in \mathbb{R}^p: \|\beta\|_0 \geq k, \min\{|\beta_j|: \beta_k \neq 0\} \geq A\}$, which implies that β has at least k nonzero coefficients exceeding A in absolute values. Ingster, Tsybakov, and Verzelen (2010) considered the alternative hypothesis $H_1: \beta \in \{\beta \in \mathbb{R}^p: \|\beta\|_0 \leq k, \|\beta\|_2 \geq \rho\}$, which concerns k sparse β with ℓ_2 norm at least ρ . In fact, the proof of our Theorem 2 can be directly extended to such an alternative concerning the ℓ_2 norm, which amounts to obtaining a lower bound of order $\sqrt{\frac{k \log p}{n}}$ for high dimensional logistic regression. However, developing a minimax optimal test for such alternative is beyond the scope of the current paper.

Additionally, in contrast to the minimax separation distance considered in this paper, the papers by Ingster, Tsybakov, and Verzelen (2010) and Mukherjee, Pillai, and Lin (2015) considered the minimax risk (or the minimax total error probability) given by

$$\inf_{\Phi} \sup_{\Sigma \in \Theta_2(k)} \text{Risk}(\Phi, \Sigma) = \inf_{\Phi} \sup_{\Sigma \in \Theta_2(k)} \left\{ \max_{\beta \in H_0} P_{\theta}(\Phi = 1) + \max_{\beta \in \Theta_1(k)} P_{\theta}(\Phi = 0) \right\}, \tag{14}$$

where the infimum is taken over all tests Φ . This minimax risk can be also written as

$$\inf_{\Phi} \sup_{\Sigma \in \Theta_2(k)} \text{Risk}(\Phi, \Sigma) = \inf_{\alpha \in (0, 1)} \left\{ \alpha + \inf_{\Phi_{\alpha}} \sup_{\Sigma \in \Theta_2(k)} \sup_{\beta \in \Theta_1(k)} P_{\theta}(\Phi_{\alpha} = 0) \right\}. \tag{15}$$

A comparison of (11) and (15) yields the slight difference between the two criteria, as one depends on a given Type I error α and the other doesn't.

Moreover, these two papers considered different design scenarios from ours. In Ingster, Tsybakov, and Verzelen (2010), only the isotropic Gaussian design was considered. As a result, the optimal tests proposed therein rely highly on the independence assumption. In Mukherjee, Pillai, and Lin (2015), the general binary regression was studied under fixed sparse design matrices. In particular, the minimax lower and upper bounds were only derived in the special case of design matrices with binary entries and certain sparsity structures.

In comparison with the recent works of Sur, Chen, and Candès (2017), Candès and Sur (2018) and Sur and Candès (2019), besides the aforementioned difference in the asymptotics of (p, n) , these two papers only considered the random Gaussian design, whereas our work also considered random bounded design as in van de Geer et al. (2014). In addition, Sur, Chen, and Candès (2017) and Sur and Candès (2019) developed the Likelihood Ratio (LLR) Test for testing the hypothesis $H_0: \beta_{j_1} = \beta_{j_2} = \dots = \beta_{j_k} = 0$ for any finite k . Intuitively, a valid test for the global null and $p/n \rightarrow \kappa \in (0, 1/2)$ can be adapted from the individual LLR tests using the Bonferroni procedure. However, as our simulations show (Section 5), such a test is less powerful compared to our proposed test.

Lastly, our minimax results focus on the highly sparse regime $k \lesssim p^\gamma$ where $\gamma \in (0, 1/2)$. As shown by Ingster, Tsybakov, and Verzelen (2010) and Mukherjee, Pillai, and Lin (2015), the problem under the dense regime where $\gamma \in (1/2, 1)$ can be very different from the sparse regime. Mostly likely, the fundamental difficulty of the testing problem changes in this situation so that different methods need to be carefully developed. We leave these interesting questions for future investigations.

3 LARGE-SCALE MULTIPLE TESTING

Denote by β the true coefficient vector in the model and denote $\mathcal{H}_0 = \{j: \beta_j = 0, j = 1, \dots, p\}$, $\mathcal{H}_1 = \{j: \beta_j \neq 0, j = 1, \dots, p\}$. In order to identify the indices in \mathcal{H}_1 , we consider simultaneous testing of the following null hypotheses

$$H_{0,j}: \beta_j = 0 \quad \text{vs.} \quad H_{1,j}: \beta_j \neq 0, \quad 1 \leq j \leq p.$$

Apart from identifying as many nonzero β_j as possible, to obtain results of practical interest, we would like to control the false discovery rate (FDR) as well as the false discovery proportion (FDP), or the number of falsely discovered variables (FDV).

3.1 Construction of Multiple Testing Procedures

Recall that in Section 2, we define the standardized statistics $M_j = \check{\beta}_j / \tau_j$, for $j = 1, \dots, p$. For a given threshold level $t > 0$, each individual hypothesis $H_{0,j}: \beta_j = 0$ is rejected if $|M_j| \geq t$. Therefore for each t , we can define

$$\text{FDP}_{\theta}(t) = \frac{\sum_{j \in \mathcal{H}_0} I\{|M_j| \geq t\}}{\max\{\sum_{j=1}^p I\{|M_j| \geq t\}, 1\}}, \quad \text{FDR}_{\theta}(t) = \mathbb{E}_{\theta}[\text{FDP}(t)],$$

and the expected number of falsely discovered variables $\text{FDV}_{\theta}(t) = \mathbb{E}_{\theta}[\sum_{j \in \mathcal{H}_0} I\{|M_j| \geq t\}]$

Procedure Controlling FDR/FDP.—In order to control the FDR/FDP at a pre-specified level $0 < \alpha < 1$, we can set the threshold level as

$$\tilde{t}_1 = \inf \left\{ 0 \leq t \leq b_p : \frac{\sum_{j \in \mathcal{H}_0} I\{|M_j| \geq t\}}{\max\{\sum_{j=1}^p I\{|M_j| \geq t\}, 1\}} \leq \alpha \right\}, \quad (16)$$

for some b_p to be determined later.

In general, the ideal choice \tilde{t}_1 is unknown and needs to be estimated because it depends on the knowledge of the true null \mathcal{H}_0 . Let $G_0(t)$ be the proportion of the nulls falsely rejected by the procedure among all the true nulls at the threshold level t , namely, $G_0(t) = \frac{1}{p_0} \sum_{j \in \mathcal{H}_0} I\{|M_j| \geq t\}$, where $p_0 = |\mathcal{H}_0|$. In practice, it is reasonable to assume that the true alternatives are sparse. If the sample size is large, we can use the tails of normal distribution $G(t) = 2 - 2\Phi(t)$ to approximate $G_0(t)$. In fact, it will be shown that, for $b_p = \sqrt{2\log p - 2\log\log p}$ in probability as $(n, p) \rightarrow \infty$. To summarize, we have the following logistic multiple testing (LMT) procedure controlling the FDR and the FDP.

Procedure 1 (LMT). Let $0 < \alpha < 1$, $b_p = \sqrt{2\log p - 2\log\log p}$ and define

$$\hat{t} = \inf \left\{ 0 \leq t \leq b_p : \frac{pG(t)}{\max\{\sum_{j=1}^p I\{|M_j| \geq t\}, 1\}} \leq \alpha \right\}. \quad (17)$$

If \hat{t} in (17) does not exist, then let $\hat{t} = \sqrt{2\log p}$. We reject $H_{0,j}$ Whenever $|M_j| \geq \hat{t}$.

Procedure Controlling FDV.—For large-scale inference, it is sometimes of interest to directly control the number of falsely discovered variables (FDV) instead of the less stringent FDR/FDP, especially when the sample size is small (Liu and Luo 2014). By definition, the FDV control, or equivalently, the per-family error rate control, provides an intuitive description of the Type I error (false positives) in variable selection. Moreover, controlling $\text{FDV} = r$ for some $0 < r < 1$ is related to the family-wise error rate (FWER) control, which is the probability of at least one false positive. In fact, FDV control can be achieved by a suitable modification of the FDP controlling procedure introduced above. Specifically, we propose the following FDV (or FWER) controlling logistic multiple testing (LMT_v) procedure.

Procedure 2 (LMT_v). For a given tolerable number of falsely discovered variables $r < p$ (or a desired level of FWER $0 < r < 1$), let $\hat{t}_{FDV} = G^{-1}(r/p)$. H_{0j} is rejected whenever $|M_j| \geq \hat{t}_{FDV}$.

3.2 Theoretical Properties for Multiple Testing Procedures

In this section we show that our proposed multiple testing procedures control the theoretical FDR/FDP or FDV asymptotically. For simplicity, our theoretical results are obtained under the bounded design scenario. For FDR/FDP control, we need an additional assumption on the interplay between the dimension p and the parameter space $\Theta(k)$.

Recall that $\eta_j = (\eta_{j1}, \dots, \eta_{jn})$ for $j = 1, \dots, p$ defined in (7). We define $F_{jk} = \mathbb{E}_\theta[\eta_{ij}\eta_{ik}f'(u_i)]$ for $1 \leq j, k \leq p$, and $\rho_{jk} = F_{jk}/\sqrt{F_{ij}F_{kk}}$. Denote $\mathcal{B}(\delta) = \{(j, k) : |\rho_{jk}| \geq \delta, i \neq j\}$ and $\mathcal{A}(\epsilon) = \mathcal{B}((\log p)^{-2} - \epsilon)$.

(A5). Suppose that for some $\epsilon > 0$ and $q > 0$,

$$\sum_{(j, k) \in \mathcal{A}(\epsilon): j, k \in \mathcal{K}_0} p \frac{2^{|\rho_{jk}|}}{1 + |\rho_{jk}|} + q = O(p^2/(\log p)^2).$$

The following proposition shows that M_j is asymptotically normal distributed and $G_0(t)$ is well approximated by $G(t)$.

Proposition 1. Under (A2) (A3) and (A4), suppose $p = O(n^c)$ for some constant $c > 0$, $k = o(\sqrt{n}/\log^{5/2} p)$, then as $(n, p) \rightarrow \infty$,

$$\sup_{j \in \mathcal{K}_0} \sup_{0 \leq t \leq \sqrt{2 \log p}} \left| \frac{P_\theta(|M_j| \geq t)}{2 - 2\Phi(t)} - 1 \right| \rightarrow 0. \tag{18}$$

If in addition we assume (A5), then

$$\sup_{0 \leq t \leq b_p} \left| \frac{G_0(t)}{G(t)} - 1 \right| \rightarrow 0 \tag{19}$$

in probability, where Φ is the cumulative distribution function of the standard normal distribution and $b_p = \sqrt{2 \log p - 2 \log \log p}$.

The following theorem provides the asymptotic FDR and FDP control of our procedure.

Theorem 4. Under the conditions of Proposition 1, for \hat{t} defined in our LMT procedure, we have

$$\lim_{(n, p) \rightarrow \infty} \frac{\text{FDR}_\theta(\hat{t})}{\alpha p_0/p} \leq 1, \quad \lim_{(n, p) \rightarrow \infty} P_\theta \left(\frac{\text{FDP}_\theta(\hat{t})}{\alpha p_0/p} \leq 1 + \epsilon \right) = 1 \tag{20}$$

for any $\epsilon > 0$.

For the FDV/FWER controlling procedure, we have the following theorem.

Theorem 5. *Under (A2) (A3) and (A4), assume $p = O(n^c)$ for some $c > 0$ and $k = o(\sqrt{n}/\log^{5/2} p)$.*

Let $r < p$ be the desired level of FDV. For $\hat{\pi}_{FDV}$ defined in our LMT_V procedure, we have

$$\lim_{(n, p) \rightarrow \infty} \frac{FDV_{\theta}(\hat{\pi}_{FDV})}{r p_0/p} \leq 1. \text{ In addition, if } 0 < r < 1, \text{ we have } \lim_{(n, p) \rightarrow \infty} \frac{FWER_{\theta}(\hat{\pi}_{FDV})}{r p_0/p} \leq 1.$$

The above theoretical results are obtained under the dimensionality condition $p = O(n^c)$, which is stronger than that of the global test. Essentially, the condition is needed to obtain the uniform convergence (18), whose form (as ratio) is stronger than the convergence in distribution in the ordinary sense (as direct difference).

4 TESTING FOR TWO LOGISTIC REGRESSION MODELS

In some applications, it is also interesting to consider hypothesis testing that involves two separate logistic regression models of the same dimension. Specifically, for $\ell = 1, 2$ and $i = 1, \dots, n_{\ell}$, where $n_1 \asymp n_2$, $y_i^{(\ell)} = f(\beta^{(\ell)\top} X_i^{(\ell)}) + \epsilon_i^{(\ell)}$, where $f(u) = e^u / (1 + e^u)$, and $\epsilon_i^{(\ell)}$ is a binary random variable such that $y_i^{(\ell)} | X_i^{(\ell)} \sim \text{Bernoulli}(f(\beta^{(\ell)\top} X_i^{(\ell)}))$. The global null hypothesis $H_0: \beta^{(1)} = \beta^{(2)}$ implies that there is overall no difference in association between covariates and the response. If this null hypothesis is rejected, we are interested in simultaneously testing the hypotheses $H_{0,j}: \beta_j^{(1)} = \beta_j^{(2)}$ for each $j = 1, \dots, p$.

To test the global null $H_0: \beta^{(1)} = \beta^{(2)}$ against $H_1: \beta^{(1)} \neq \beta^{(2)}$, we can first obtain $\check{\beta}_j^{(\ell)}$ and $\check{\tau}_j^{(\ell)}$ for each model, and then calculate the coordinate-wise standardized statistics

$$T_j = \frac{\check{\beta}_j^{(1)}}{\sqrt{2}\check{\tau}_j^{(1)}} - \frac{\check{\beta}_j^{(2)}}{\sqrt{2}\check{\tau}_j^{(2)}}, \text{ for } j = 1, \dots, p. \text{ Define the global test statistic as } T_n = \max_{1 \leq j \leq p} T_j^2, \text{ it}$$

can be shown that the limiting null distribution is also a Gumbel distribution. The α level global test is thus defined as $\Phi_{\alpha}(T_n) = I\{T_n \leq 2 \log p - \log \log p + q_{\alpha}\}$, where $q_{\alpha} = -\log(\pi) - 2 \log \log(1 - \alpha)^{-1}$. For multiple hypotheses testing of two regression vectors $H_{0,j}: \beta_j^{(1)} = \beta_j^{(2)}$, we consider the test statistics T_j defined above. The two-sample multiple testing procedure controlling FDR/FDP is given as follows.

Procedure 3. *Let $0 < \alpha < 1$ and define $\hat{t} = \inf \left\{ 0 \leq t \leq b_p: \frac{pG(t)}{\max\{\sum_{j=1}^p I\{|T_j| \geq t\}, 1\}} \leq \alpha \right\}$. If the above \hat{t} does not exist, let $\hat{t} = \sqrt{2 \log p}$. We reject $H_{0,j}$ whenever $|T_j| \geq \hat{t}$.*

5 SIMULATION STUDIES

In this section we examine the numerical performance of the proposed tests. Due to the space limit, for both global and multiple testing problems, we only focus on the single

regression setting, and report the results on two logistic regressions in the Supplementary Materials. Throughout our numerical studies, sample splitting was not used.

5.1 Global Hypothesis Testing

In the following simulations, we consider a variety of dimensions, sample sizes, and sparsity of the models. For alternative hypotheses, the dimension of the covariates p ranges from 100, 200, 300 to 400, and the sparsity k is set as 2 or 4. The sample sizes n are determined by the ratio $r = p/n$ that takes values of 0.2, 0.4 and 1.2. To generate the design matrix X , we consider the Gaussian design with the blockwise-correlated covariates so that $\Sigma = \Sigma_B$, where Σ_B is a $p \times p$ blockwise diagonal matrix including 10 equal-sized blocks, whose diagonal elements are 1's and off-diagonal elements are set as 0.7. Under the alternative, suppose \mathcal{S} is the support of the regression coefficients β and $|\mathcal{S}| = k$, we set $|\beta_j| = \rho 1\{j \in \mathcal{S}\}$ for $j = 1, \dots, p$ and $\rho = 0.75$ with equal proportions of ρ and $-\rho$. We set $\kappa_0 = 0$ and $\kappa_1 = 0.5$.

To assess the empirical performance of our proposed test (“Proposed”), we compare our test with (i) a Bonferroni procedure applied to the p-values from univariate screening using MLE statistic (“U-S”), and (ii) to the method of Sur, Chen, and Candès (2017); Sur and Candès (2019) (“LLR”) in the setting where $r = 0.2$ and 0.4.

Table 2 shows the empirical type I errors of these tests at level $\alpha = 0.05$ based on 1000 simulations. Figure 1 shows the corresponding empirical powers under various settings. As we expected, our proposed method outperforms the other two alternatives in all the cases (including the moderate dimensional cases where $r = 0.2$ and 0.4), and the power increases as n or p grows. In the rather lower dimensional setting where $r = 0.2$, the LLR performs almost as well as our proposed method.

5.2 Multiple Hypotheses Testing

FDR Control.—In this case, we set $p = 800$ and let n vary from 600, 800, 1000, 1200 to 1400, so that all the cases are high-dimensional in the sense that $p > n/2$. The sparsity level k varies from 40, 50 to 60. For the true positives, given the support \mathcal{S} such that $|\mathcal{S}| = k$, we set $|\beta_j| = \rho 1\{j \in \mathcal{S}\}$ for $j = 1, \dots, p$ with equal proportions of ρ and $-\rho$. The design covariates X_j 's are generated from a $(|X_i^\top \beta| < 3)$ -truncated multivariate Gaussian distribution with covariance matrix $\Sigma = 0.01 \Sigma_M$, where Σ_M is a $p \times p$ blockwise diagonal matrix of 10 identical unit diagonal Toeplitz matrices whose off-diagonal entries descend from 0.1 to 0 (see Supplementary Material for the explicit form). The choice of κ_0 and κ_1 are the same as the global testing. Throughout, we set the desired FDR level as $\alpha = 0.2$.

We compare our proposed procedure (denoted as “LMT”) with following methods: (i) the basic LMT procedure with b_p in (17) replaced by ∞ (“LMT0”), which is equivalent to applying the BH procedure (Benjamini and Hochberg 1995) to our debiased statistics M_j , (ii) the BY procedure (Benjamini and Yekutieli 2001) using our debiased statistics M_j (“BY”), implemented using the R function `p.adjust(...,method=“BY”)`, (iii) a BH procedure applied to the p-values from univariate screening using the MLE statistics (“U-S”), and (iv) the knockoff method of Candès et al. (2018) (“Knockoff”). Figure 2 shows boxplots of the pooled empirical FDRs (see Supplementary Material for the case-by-case FDRs) and Figure

3 shows the empirical powers of these methods based on 1000 replications. Here the power is defined as the number of correctly discovered variables divided by the number of truly associated variables. As a result, we find that LMT and LMT0 correctly control FDRs and have the greatest power among all the cases. In particular, the power of LMT and LMT0 are almost the same, which increases as the sparsity decreases, the signal magnitude ρ increases, or the sample size n increases, while LMT0 has slightly inflated FDRs. The U-S method, although correctly controls the FDRs, has poor power, which is largely due to the dependence among the covariates.

FDV Control.—For our proposed test that controls FDV (denoted as LMT $_{\nu}$), by setting desired FDV level $r=10$, we apply our method to various settings. Specifically, we set $\rho=3$, $p \in \{800, 1000, 1200\}$, set $k \in \{40, 50, 60\}$, and let n vary from 400, 600, 800 to 1000. The design covariates are generated similarly as the previous part. The resulting empirical FDV and powers are summarized in Table 3. Our proposed LMT $_{\nu}$ has the correct control of FDV in all the settings and the power increases as n grows, k decreases, or p decreases.

6 REAL DATA ANALYSIS

We illustrate our proposed methods by analyzing a dataset from the Pediatric Longitudinal Study of Elemental Diet and Stool Microbiome Composition (PLEASE) study, a prospective cohort study to investigate the effects of inflammation, antibiotics, and diet as environmental stressors on the gut microbiome in pediatric Crohn's disease (Lewis et al. 2015; Lee et al. 2015; Ni et al. 2017). The study considered the association between pediatric Crohn's disease and fecal metabolomics by collecting fecal samples of 90 pediatric patients with Crohn's disease at baseline, 1 week, and 8 weeks after initiation of either anti-tumor necrosis factor (TNF) or enteral diet therapy, as well as those from 25 healthy control children (Lewis et al. 2015). In details, an untargeted fecal metabolomic analysis was performed on these samples using liquid chromatography-mass spectrometry (LC-MS). Metabolites with more than 80% missing values across all samples were removed from the analysis. For each metabolite, samples with the missing values were imputed with its minimum abundance across samples. To avoid potential large outliers, for each sample, the metabolite abundances were further normalized by dividing 90% cumulative sum of the abundances of all metabolites. The normalized abundances were then log transformed and used in all analyses. The metabolomics annotation was obtained from Human Metabolome Database (Lee et al. 2015). In total, for each sample, abundances of 335 known metabolites were obtained and used in our analysis.

6.1 Association Between Metabolites and Crohn's Disease Before and After Treatment

We first test the overall association between 335 characterized metabolites and Crohn's disease by fitting a logistic regression using the data of 25 healthy controls and 90 Crohn's disease patients at the baseline. We obtain a global test statistic of 433.88 with a p-value < 0.001 , indicating a strong association between Crohn's disease and fecal metabolites. At the FDR $< 5\%$, our multiple testing procedure selects four metabolites, including C14:0.sphingomyelin, C24:1.Ceramide.(d18:1) and 3-methyladipate/pimelate (see Table 4). Recent studies have demonstrated that sphingolipid metabolites, particularly ceramide

and sphingosine-1-phosphate, are signaling molecules that regulate a diverse range of cellular processes that are important in immunity, inflammation and inflammatory disorders (Maceyka and Spiegel 2014). In fact, ceramide acts to reduce tumor necrosis factor (TNF) release (Rozenova et al. 2010) and has important roles in the control of autophagy, a process strongly implicated in the pathogenesis of Crohn's disease (Barrett et al. 2008; Sewell et al. 2012).

We next investigate whether treatment of Crohn's disease alters the association between metabolites and Crohn's disease by fitting two separate logistic regressions using the metabolites measured one week or 8 weeks after the treatment. At each time point, a significant association is detected based on our global test (p -value < 0.001). One week after the treatment, we observe six metabolites associated with Crohn's disease, including all four identified at the baseline and two additional metabolites, beta-alanine and adipate (see Table 4). The beta-alanine and adipate associations are likely due to that beta-alanine and adipate are important ingredients of the enteral nutrition treatment of Crohn's disease. However, it is interesting that at 8 weeks after the treatment, valine, C16.carnitine and C18.carnitine are identified to be associated with Crohn's disease together with 3-methyladipate/pimelate and beta-alanine. It is known that carnitine plays an important role in Crohn's disease, which might be a consequence of the underlying functional association between Crohn's disease and mutations in the carnitine transporter genes (Peltekova et al. 2004; Fortin 2011). Deficiency of carnitine can lead to severe gut atrophy, ulceration and inflammation in animal models of carnitine deficiency (Shekhawat et al. 2013). Our results may suggest that the treatment increases carnitine, leading to reduction of inflammation.

6.2 Comparison of Metabolite Associations Between Responders and Non-Responders

To compare the metabolic association with Crohn's disease for responders ($n = 47$) and non-responders ($n = 34$) eight weeks after treatment, we fit two logistic regression models, responder versus normal control and non-responder versus normal control. Our global test shows that there is an overall difference in regression coefficients for responders and for non-responders when compared to the normal controls (p -value < 0.001). We next apply our proposed multiple testing procedure to identify the metabolites that have different regression coefficients in these two different logistic regression models. At the $FDR < 0.05$, our procedure identifies 9 metabolites with different regression coefficients (see Table 5). It is interesting that all these 9 metabolites have the same signs of the refitted coefficients, while the actual magnitudes of the associations between responders and non-responders when compared to the normal controls are different. Besides C24:4.cholesteryl ester, beta-alanine, valine, C18.carnitine and 3-methyladipate/pimelate that we observe in previous analyses, metabolites 5-hydroxytryptophan, nicotinate, and succinate also have differential associations between responders and non-responders when compared to the controls.

7 DISCUSSION

In this paper, for both global and multiple testing, the precision matrix $\mathbf{\Omega} = \mathbf{\Sigma}^{-1}$ of the covariates is assumed to be sparse and unknown. Node-wise regression among the covariates is used to learn the covariance structure in constructing the debiased estimator. However,

if the prior knowledge of $\mathbf{\Omega} = \mathbf{I}$ is available, the algorithm can be simplified greatly. Specifically, instead of incorporating the Lasso estimators as in (8), we let $v_j = \widehat{\mathbf{W}}^{-1} \mathbf{x}_j$ and $\tau_j = \|v_j\|_1 / \langle v_j, \mathbf{x}_j \rangle$ for each $j = 1, \dots, p$. The theoretical properties of the resulting global testing and multiple testing procedures still hold, while the computational efficiency is improved dramatically. However, from our theoretical analysis, even with the knowledge of $\mathbf{\Omega} = \mathbf{I}$, the theoretical requirement for the model sparsity ($k = o(\sqrt{n}/\log^3 p)$ in the Gaussian case and $k = o(\sqrt{n}/\log^{5/2} p)$ in the bounded case) cannot be relaxed due to the nonlinearity of the problem.

Sample splitting was used in this paper for theoretical purpose. This is different from other works on inference in high-dimensional linear/logistic regression models, including Ingster, Tsybakov, and Verzelen (2010), van de Geer et al. (2014), Mukherjee, Pillai, and Lin (2015) and Javanmard and Javadi (2019), where sample splitting is not needed. However, as we discussed throughout the paper, the assumptions and the alternatives that we considered are different from those previous papers. In the case of high-dimensional logistic regression model, a sample splitting procedure seems unavoidable under the current framework of our technical analysis without making additional strong structural assumptions such as the sparse inverse Hessian matrices used in van de Geer et al. (2014) or the weakly correlated design matrices used in Mukherjee, Pillai, and Lin (2015). Our simulations showed that the sample splitting is actually not needed in order for our proposed methods to perform well. It is of interest to develop technical tools that can eliminate sample splitting in inference for high dimensional logistic regression models.

As mentioned in the introduction, the logistic regression model can be viewed as a special case of the single index model $y = f(\beta^\top x) + \epsilon$ where f is a known transformation function (Yang et al. 2015). Based on our analysis, it is clear that the theoretical results are not limited to the sigmoid transfer function. In fact, the proposed methods can be applied to a wide range of transformation functions satisfying the following conditions: (C1) f is continuous and for any $u \in \mathbb{R}$, $0 < f(u) < 1$; (C2) for any $u_1, u_2 \in \mathbb{R}$, there exists a constant $L > 0$ such that $|f(u_1) - f(u_2)| \leq L|u_1 - u_2|$; and (C3) for any constant $C > 0$, there exists $\delta > 0$ such that for any $|u| \leq C$, $f(u) \geq \delta$. Examples include but are not limited to the following function classes

- *Cumulative density functions:* $f(x) = P(X \leq x)$ for some continuous random variable X supported on \mathbb{R} . In particular, when $X \sim N(0, 1)$, the resulting model becomes the probit regression.
- *Affine hyperbolic tangent functions:* $f(x) = \frac{1}{2} \tanh(ax + b) + 1$ for some parameter $a, b \in \mathbb{R}$. In particular, $(a, b) = (1, 0)$ corresponds to $f(x) = e^x / (1 + e^x)$.
- *Generalized logistic functions:* $f(x) = (1 + e^{-x})^{-\alpha}$ for some $\alpha > 0$.

Besides the problems we considered in this paper, it is also of interest to construct confidence intervals for functionals of the regression coefficients, such as $\|\beta\|_1, \|\beta\|_2$, or $\theta^\top \beta$ for some given loading vector θ . In modern statistical machine learning, logistic regression is considered as an efficient classification method (Abramovich and Grinshtein 2018). In

practice, a predicted label with an uncertainty assessment is usually preferred. Therefore, another important problem is the construction of predictive intervals of the conditional probability π^* associated with a given predictor X^* . These problems are related to the current work and are left for future investigations.

8 PROOFS OF THE MAIN THEOREMS

In this section, we prove Theorems 1, Theorem 2 and Theorem 4 in the paper. The proofs of other results, including Theorems 3 and 5, Proposition 1 and the technical lemmas, are given in our Supplementary Materials.

Proof of Theorem 1

Define $F_{ij} = \mathbb{E}[\eta_{ij}^2 / \dot{f}(u_i)]$. Under H_0 , $F_{ij} = 4\mathbb{E}[\eta_{ij}^2] = 4/\omega_{jj}$, and by (A3), $c < F_{jj} < C$ for $j = 1, \dots, p$ and some constant $C > 0$. Define statistics

$$\tilde{M}_j = \frac{\langle v_j, \epsilon \rangle}{\|v_j\|_n}, \quad \text{and} \quad \tilde{M}_j = \frac{\sum_{i=1}^n \eta_{ij} \epsilon_i / \dot{f}(u_i)}{\sqrt{n F_{jj}}}, \quad j = 1, \dots, p.$$

and $\tilde{M}_n = \max_j \tilde{M}_j^2$, $\tilde{M}_n = \max_j \tilde{M}_j^2$. The following lemma shows that \tilde{M}_n and therefore \tilde{M}_n are good approximations of M_n .

Lemma 1. *Under the condition of Theorem 1, the following events*

$$B_1 = \{|\tilde{M}_n - \tilde{M}_n| = o(1)\}, \quad B_2 = \{|\tilde{M}_n - M_n| = o\left(\frac{1}{\log p}\right)\},$$

hold with probability at least $1 - O(p^{-c})$ for some constant $c > 0$.

It follows that under the event $B_1 \cap B_2$, let $y_p = 2\log p - \log \log p + x$ and $\epsilon_n = o(1)$, we have

$$P_\theta(\tilde{M}_n \leq y_p - \epsilon_n) \leq P_\theta(M_n \leq y_p) \leq P_\theta(\tilde{M}_n \leq y_p + \epsilon_n)$$

Therefore it suffices to prove that for any $t \in \mathbb{R}$, as $(n, p) \rightarrow \infty$,

$$P_\theta(\tilde{M}_n \leq y_p) \rightarrow \exp\left(-\frac{1}{\sqrt{\tau}} \exp(-x/2)\right). \tag{21}$$

Now define $\hat{M}_j = \frac{\sum_{i=1}^n \hat{Z}_{ij}}{\sqrt{n F_{ij}}}$, $j = 1, \dots, p$. where $\hat{Z}_{ij} = v_{ij}^0 \epsilon_i 1\{|v_{ij}^0 \epsilon_i| \leq \tau_n\} - \mathbb{E}[v_{ij}^0 \epsilon_i 1\{|v_{ij}^0 \epsilon_i| \leq \tau_n\}]$

for $\tau_n = \log(p + n)$, $v_{ij}^0 = \eta_{ij} / \dot{f}(u_i)$ and $\hat{M}_n = \max_j \hat{M}_j^2$. The following lemma states that \hat{M}_n is close to \tilde{M}_n .

Lemma 2. *Under the condition of Theorem 1, $|\tilde{M}_n - \hat{M}_n| = o(1)$ with probability at least $1 - O(p^{-c})$ for some constant $c > 0$.*

By Lemma 2, it suffices to prove that for any $t \in \mathbb{R}$, as $(n, p) \rightarrow \infty$,

$$P_\theta(\widehat{M}_n \leq y_p) \rightarrow \exp\left(-\frac{1}{\sqrt{\pi}}\exp(-x/2)\right). \tag{22}$$

To prove this, we need the classical Bonferroni inequality.

Lemma 3. **(Bonferroni inequality)** Let $B = \bigcup_{t=1}^p B_t$. For any integer $k < p/2$, we have

$$\sum_{t=1}^{2k} (-1)^{t-1} A_t \leq P(B) \leq \sum_{t=1}^{2k-1} (-1)^{t-1} A_t, \tag{23}$$

where $A_t = \sum_{1 \leq i_1 < \dots < i_t \leq p} P(B_{i_1} \cap \dots \cap B_{i_t})$.

By Lemma 3, for any integer $0 < q < p/2$,

$$\begin{aligned} & \sum_{d=1}^{2q} (-1)^{d-1} \sum_{1 \leq j_1 \leq \dots \leq j_d \leq p} P_\theta\left(\bigcap_{k=1}^d A_{j_k}\right) \leq P_\theta\left(\max_{1 \leq j \leq p} \widehat{M}_j^2 \geq y_p\right) \\ & \leq \sum_{d=1}^{2p-1} (-1)^{d-1} \sum_{1 \leq j_1 < \dots < j_d \leq p} P_\theta\left(\bigcap_{k=1}^d A_{j_k}\right), \end{aligned} \tag{24}$$

where $A_{j_k} = \{\widehat{M}_{j_k}^2 \geq y_p\}$. Now let $w_{ij} = \widehat{Z}_{ij}/\sqrt{F_{jj}}$ for $j = 1, \dots, p$, and $\mathbf{W}_i = (w_{i,j_1}, \dots, w_{i,j_d})^\top$ for $1 \leq i \leq n$. Define $\|\mathbf{a}\|_{\min} = \min_{1 \leq i \leq d} |a_i|$ for any vector $\mathbf{a} \in \mathbb{R}^d$. Then we have

$$P_\theta\left(\bigcap_{k=1}^d A_{j_k}\right) = P\left(\sum_{i=1}^n \mathbf{W}_i(\min \geq y_p^{1/2})\right).$$

Then it follows from Theorem 1.1 in Zaitsev (1987) that

$$\begin{aligned} & P_\theta\left(n^{-1/2} \sum_{i=1}^n \mathbf{W}_i(\min \geq y_p^{1/2})\right) \leq P_\theta\left(\|\mathbf{N}_d\|_{\min} \geq y_p^{1/2} - \epsilon_n(\log p)^{-1/2}\right) \\ & + c_1 d^{5/2} \exp\left\{-\frac{n^{1/2} \epsilon_n}{c_2 d^3 \tau_n(\log p)^{1/2}}\right\}, \end{aligned} \tag{25}$$

where $c_1 > 0$ and $c_2 > 0$ are constants, $\epsilon_n \rightarrow 0$ which will be specified later, and $\mathbf{N}_d = (N_{m_1}, \dots, N_{m_d})$ is a normal random vector with $E(\mathbf{N}_d) = 0$ and $\text{cov}(\mathbf{N}_d) = \text{cov}(\mathbf{W}_1)$. Here d is a fixed integer that does not depend on n, p . Because $\log p = o(n^{1/5})$, we can let $\epsilon_n \rightarrow 0$ sufficiently slow, say, $\epsilon_n = \sqrt{\log^5 p/n}$, so that for any large $c > 0$,

$$c_1 d^{5/2} \exp\left\{-\frac{n^{1/2} \epsilon_n}{c_2 d^3 \tau_n(\log p)^{1/2}}\right\} = O(p^{-c}). \tag{26}$$

Combining (24), (25) and (26), we have

$$P_\theta \left(\max_{1 \leq j \leq p} \widehat{M}_j^2 \geq y_p \right) \leq \sum_{d=1}^{2p-1} (-1)^{d-1} \sum_{1 \leq j_1 < \dots < j_d \leq p} P_\theta \left(\|\mathbf{N}_d\|_{\min} \geq y_p^{1/2} - \epsilon_n(\log p)^{-1/2} \right) + o(1). \tag{27}$$

Similarly, one can derive

$$P_\theta \left(\max_{1 \leq j \leq p} \widehat{M}_j^2 \geq y_p \right) \geq \sum_{d=1}^{2p} (-1)^{d-1} \sum_{1 \leq j_1 < \dots < j_d \leq p} P_\theta \left(\|\mathbf{N}_d\|_{\min} \geq y_p^{1/2} + \epsilon_n(\log p)^{-1/2} \right) + o(1). \tag{28}$$

Now we use the following lemma from Xia, Cai, and Cai (2018).

Lemma 4. For any fixed integer $d \geq 1$ and real number $t \in \mathbb{R}$,

$$\sum_{1 \leq j_1 < \dots < j_d \leq p} P_\theta \left(\|\mathbf{N}_d\|_{\min} \geq y_p^{1/2} \pm \epsilon_n(\log p)^{-1/2} \right) = \frac{1}{d!} \left(\frac{1}{\sqrt{\pi}} \exp(-t/2) \right)^d (1 + o(1)).$$

It then follows from the above lemma, (27) and (28) that

$$\limsup_{n, p \rightarrow \infty} P_\theta \left(\max_{1 \leq j \leq p} \widehat{M}_j^2 \geq y_p \right) \leq \sum_{d=1}^{2p} (-1)^{d-1} \frac{1}{d!} \left(\frac{1}{\sqrt{\pi}} \exp(-t/2) \right)^d,$$

$$\liminf_{n, p \rightarrow \infty} P_\theta \left(\max_{1 \leq j \leq p} \widehat{M}_j^2 \geq y_p \right) \geq \sum_{d=1}^{2p-1} (-1)^{d-1} \frac{1}{d!} \left(\frac{1}{\sqrt{\pi}} \exp(-t/2) \right)^d,$$

for any positive integer p . By letting $p \rightarrow \infty$, we obtain (22) and the proof is complete. \square

Proof of Theorem 2.

The proof essentially follows from the general Le Cam’s method described in Section 7.1 of Baraud (2002). The key elements can be summarized as the following lemma that reduces the lower bound problem to calculation of the total variation distance between two posterior distributions.

Lemma 5. Let \mathcal{H}_1 be some subset in an ℓ_2 bounded Hilbert space and ρ some positive number. Let μ_ρ be some probability measure on $\mathcal{H}_1 = \{\theta \in \Theta, \|\theta\| = \rho\}$. Set $P_{\mu_\rho} = \int P_\theta d\mu_\rho(\theta)$, P_0 as the (posterior) distribution at the null, and denote by Φ_α the level- α tests, we have

$$\inf_{\Phi_\alpha} \sup_{\theta \in \mathcal{H}_1} P_\theta(\Phi_\alpha = 0) \geq \inf_{\Phi_\alpha} P_{\mu_\rho}(\Phi_\alpha = 0) \geq 1 - \alpha - TV(P_{\mu_\rho}, P_0).$$

where $TV(P_{\mu_\rho}, P_0)$ denotes the total variation distance between P_{μ_ρ} and P_0 .

Now since by definition $\rho^*(\Phi_\alpha, \delta, \Theta(k)) = \rho^*(\Phi_\alpha, \delta, \Sigma)$ for any $\Sigma \in \Theta_2(k)$, by Lemma 5, it suffices to construct the corresponding \mathcal{H}_1 for $\beta \in \Theta_\beta(k)$ and find a lower bound $\rho_1 = \rho(\eta)$ such that

$$\forall \rho \leq \rho_1 \quad \inf_{\Phi_\alpha} P_{\mu_\rho}(\Phi_\alpha = 0) \geq 1 - \alpha - \eta = \delta. \tag{29}$$

for fixed covariance $\Sigma = \mathbf{I}$. In this case, an upper bound for the χ^2 -divergence between P_{μ_ρ} and P_0 , defined as $\chi^2(P_{\mu_\rho}, P_0) = \int \frac{(dP_{\mu_\rho})^2}{dP_0} - 1$, can be obtained by carefully constructing the alternative space \mathcal{H}_1 . Since $TV(f, g) \leq \sqrt{\chi^2(f, g)}$ (see p.90 of Tsybakov (2009)), it follows that $\inf_{\Phi_\alpha} P_{\mu_\rho}(\Phi_\alpha = 0) \geq 1 - \alpha - \sqrt{\chi^2(P_{\mu_\rho}, P_0)}$. By choosing $\rho_1 = \rho(\eta)$ such that for any $\rho \leq \rho_1$, $\chi^2(P_{\mu_\rho}, P_0) \leq \eta^2 = (1 - \alpha - \delta)^2$, we have (29) holds. In the following, we will construct the alternative space \mathcal{H}_1 and derive an upper bound of $\chi^2(P_{\mu_\rho}, P_0)$ where P_0 corresponds to the null space \mathcal{H}_0 defined at a single point $\beta = 0$. We divide the proofs into two parts. Throughout, the design covariance matrix is chosen as $\Sigma = \mathbf{I}$.

Step 1: Construction of \mathcal{H}_1 .—Firstly, for a set M , we define $\ell(M, n)$ as the set of all the n -element subsets of M . Let $[1: p] \equiv \{1, \dots, p\}$, so $\ell([1: p], k)$ contains all the k -element subsets of $[1: p]$. We define the alternative parameter space $\mathcal{H}_1 = \{\beta \in \mathbb{R}^p: \beta_j = \rho 1_{\{j \in I\}} \text{ for } I \in \ell([1: p], k)\}$. In other words, \mathcal{H}_1 contains all the k -sparse vectors $\beta(I)$ whose nonzero components ρ are indexed by I . Apparently, for any $\beta \in \mathcal{H}_1$, it follows $\|\beta\|_\infty = \rho$ and $\mathcal{H}_1 \subseteq \Theta_1(k)$.

Step 2: Control of $\chi^2(P_{\pi_{\mathcal{H}_1}}, P_0)$.—Let π denote the uniform prior of the random index set I over $\ell([1: p], k)$. This prior induces a prior distribution $\pi_{\mathcal{H}_1}$ over the parameter space \mathcal{H}_1 . For $\{\mathbf{0}_p\} = \mathcal{H}_0$, the corresponding joint distribution of the data $\{(X_i, y_i)\}_{i=1}^n$ is

$$f = \prod_{i=1}^n p(X_i, y_i) = \frac{1}{(2\pi)^{np/2}} \prod_{i=1}^n \frac{1}{2} e^{-\|X_i\|_2^2/2}.$$

Similarly, the posterior distribution of the samples over the prior $\pi_{\mathcal{H}_1}$ is denoted as

$$g = \prod_{i=1}^n \int_{\mathcal{H}_i} p(X_i, y_i; \beta) \pi_{\mathcal{H}_i} = \frac{1}{\binom{p}{k}} \sum_{\beta \in \mathcal{H}_1} \prod_{i=1}^n p(X_i, y_i; \beta).$$

As a result, we have the following lemma controlling $\chi^2(P_{\pi_{\mathcal{H}_1}}, P_0) = \chi^2(g, f)$.

Lemma 6. Let $\rho^2 = \frac{1}{n} \log \left(1 + \frac{p}{h(\eta)k^2} \right)$ where $h(\eta) = [\log(\eta^2 + 1)]^{-1}$ and $\eta = 1 - \alpha - \delta$, then we have $\chi^2(g, f) \leq (1 - \alpha - \delta)^2$.

Combining Lemma 5 and Lemma 6, we know that for $\alpha, \delta > 0$ and $\alpha + \delta < 1$, if

$\rho = \sqrt{\frac{1}{n} \log \left(1 + \frac{p}{h(\eta)k^2} \right)}$, then $\forall \rho' \leq \rho$, $\inf_{\Phi_{\alpha}\beta \in \Theta(k): \|\beta\|_{\infty} \geq \rho'} P_{\theta}(\Phi_{\alpha} = 0) \geq \delta$. Therefore, it follows that

$$\rho^*(\alpha, \delta, \Theta(k)) \geq \rho^*(\alpha, \delta, \mathbf{I}) \gtrsim \sqrt{\frac{1}{n} \log \left(1 + \frac{p}{k^2} \right)}. \tag{30}$$

Lastly, note that for the above chosen ρ , $\mathcal{H}_1 \subset \Theta_1(k) \cap \left\{ \beta \in \mathbb{R}^p: \|\beta\|_2 \lesssim (n^{1/4} \log p)^{-1} \right\}$ when $k \lesssim \min \left\{ p^{\gamma}, \sqrt{n} / \log^3 p \right\}$ for some $0 < \gamma < 1/2$. This completes the proof. \square

Proof of Theorem 4.

The proof follows similar arguments of the proof of Theorem 3.1 in Javanmard and Javadi (2019). We first consider the case when \hat{t} , given by (17), does not exist. In this case, $\hat{t} = \sqrt{2 \log p}$ and we consider the event $\Omega_0 = \left\{ \sum_{j \in \mathcal{H}_0} I(|M_j| \geq \sqrt{2 \log p}) \geq 1 \right\}$ that there are at least one false positive. In order to show the FDR/FDP can be controlled in this case, we show that

$$P_{\theta}(\Omega_0) \rightarrow 0, \quad \text{as } (n, p) \rightarrow \infty. \tag{31}$$

Note that for $j \in \mathcal{H}_0$, we have $M_j = \frac{\check{\beta}_j}{\tau_j} = \frac{\langle v_j, \epsilon \rangle}{\|v_j\|_n} + \frac{\langle v_j, Re \rangle}{\|v_j\|_n} - \frac{\langle v_j, \mathbf{h}_{-j} \rangle_n}{\|v_j\|_n}$. Then

$$\begin{aligned} P_{\theta}(\Omega_0) &\leq P_{\theta} \left(\sum_{j \in \mathcal{H}_0} I \left(\frac{\langle v_j, \epsilon \rangle}{\|v_j\|_n} + \frac{\langle v_j, Re \rangle}{\|v_j\|_n} - \frac{\langle v_j, \mathbf{h}_{-j} \rangle_n}{\|v_j\|_n} \geq \sqrt{2 \log p} \right) \geq 1 \right) \\ &+ P_{\theta} \left(\sum_{j \in \mathcal{H}_0} I \left(\frac{\langle v_j, \epsilon \rangle}{\|v_j\|_n} + \frac{\langle v_j, Re \rangle}{\|v_j\|_n} - \frac{\langle v_j, \mathbf{h}_{-j} \rangle_n}{\|v_j\|_n} \leq -\sqrt{2 \log p} \right) \geq 1 \right). \end{aligned} \tag{32}$$

For any $\epsilon > 0$, we can bound the first term by

$$\begin{aligned}
 & P_\theta \left(\sum_{j \in \mathcal{H}_0} I \left(\frac{\langle v_j, \epsilon \rangle}{\|v_j\|_k} + \frac{\langle v_j, Re \rangle}{\|v_j\|_v} - \frac{\langle v_j, \mathbf{h}_{-j} \rangle_n}{\|v_j\|_n} \geq \sqrt{2 \log p} \right) \geq 1 \right) \\
 &= P_\theta \left(\sum_{j \in \mathcal{H}_0} I \left(\tilde{M}_j \geq \sqrt{2 \log p} + \frac{\langle v_j, \mathbf{h}_{-j} \rangle_n}{\|v_j\|_n} - \frac{\langle v_j, Re \rangle}{\|v_j\|_n} \geq 1 \right) \right) \\
 &\leq P_\theta \left(\sum_{j \in \mathcal{H}_0} I(\tilde{M}_j \geq \sqrt{2 \log p} - \epsilon) \geq 1 \right) + P_\theta \left(\max_{j \in \mathcal{H}_0} \left| \frac{\langle v_j, \mathbf{h}_{-j} \rangle_n}{\|v_j\|_n} - \frac{\langle v_j, Re \rangle}{\|v_j\|_n} \right| \geq \epsilon \right) \\
 &\leq p \max_{j \in \mathcal{H}_0} P_\theta(\tilde{M}_j \geq \sqrt{2 \log p} - \epsilon) + P_\theta \left(\max_{j \in \mathcal{H}_0} \left| \frac{\langle v_j, \mathbf{h}_{-j} \rangle_n}{\|v_j\|_n} - \frac{\langle v_j, Re \rangle}{\|v_j\|_n} \right| \geq \epsilon \right)
 \end{aligned}$$

By the proof of Lemma 1, we know that $P_\theta \left(\max_{j \in \mathcal{H}_0} \left| \frac{\langle v_j, \mathbf{h}_{-j} \rangle_n}{\|v_j\|_n} - \frac{\langle v_j, Re \rangle}{\|v_j\|_n} \right| \geq \epsilon \right) \rightarrow 0$. In addition, for $j \in \mathcal{H}_0$, $P_\theta(\tilde{M}_j \geq \sqrt{2 \log p} - \epsilon) \leq P_\theta(\tilde{M}_j \geq \sqrt{2 \log p} - 2\epsilon) + P_\theta(|\tilde{M}_j - \tilde{M}_j| \geq \epsilon)$, where $\max_{j \in \mathcal{H}_0} P_\theta(|\tilde{M}_j - \tilde{M}_j| \geq \epsilon) = O(p^{-c})$ for some sufficiently large $c > 0$. Now since

$$\begin{aligned}
 \tilde{M}_j &= \frac{\sum_{i=1}^n \eta_{ij} \epsilon_i / \hat{f}(u_i)}{\sqrt{n F_{jj}}} \text{ where } \mathbb{E} \frac{\eta_{ij} \epsilon_i / \hat{f}(u_i)}{\sqrt{F_{jj}}} = 0 \text{ and } \text{Var} \left(\frac{\eta_{ij} \epsilon_i / \hat{f}(u_i)}{\sqrt{F_{jj}}} \right) = 1, \text{ by Lemma 6.1 of Liu} \\
 (2013), \text{ we have } \sup_{0 \leq t \leq 4\sqrt{\log p}} \left| \frac{P_\theta(\tilde{M}_j \geq t)}{G(t)} - 1 \right| &\leq C(\log p)^{-1}. \text{ Now let } t = \sqrt{2 \log p} - 2\epsilon, \text{ we have} \\
 P_\theta(\tilde{M}_j \geq \sqrt{2 \log p} - 2\epsilon) &\leq G(\sqrt{2 \log p} - 2\epsilon) + C \frac{G(\sqrt{2 \log p} - 2\epsilon)}{\log p}.
 \end{aligned}$$

Hence $p \max_{j \in \mathcal{H}_0} P_\theta(\tilde{M}_j \geq \sqrt{2 \log p} - \epsilon) \leq CpG(\sqrt{2 \log p} - 2\epsilon) + O(p^{-c})$, which goes to zero as $(n, p) \rightarrow \infty$. By symmetry, we know that the second term in (32) also goes to 0. Therefore we have proved (31).

Now consider the case when $0 \leq \hat{t} \leq b_p$ holds. We have

$$\text{FDP}_{\theta(\hat{t})} = \frac{\sum_{j \in \mathcal{H}_0} I\{|M_j| \geq \hat{t}\}}{\max\{\sum_{j=1}^p I\{|M_j| \geq \hat{t}\}, 1\}} \leq \frac{p_0 G(\hat{t})}{\max\{\sum_{j=1}^p I\{|M_j| \geq \hat{t}\}, 1\}} (1 + A_p),$$

where $A_p = \sup_{0 \leq t \leq b_p} \left| \frac{\sum_{j \in \mathcal{H}_0} I\{|M_j| \geq t\}}{p_0 G(t)} - 1 \right|$. Note that by definition

$$\frac{p_0 G(\hat{t})}{\max\{\sum_{j=1}^p I\{|M_j| \geq \hat{t}\}, 1\}} \leq \frac{p_0 \alpha}{p}.$$

The proof is complete if $A_p \rightarrow 0$ in probability, which has been shown by Proposition 1. \square

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

- Abramovich F, and Grinshtein V (2018), “High-dimensional classification by sparse logistic regression,” *IEEE Transactions on Information Theory*, 65, 3068–3079.
- Bach F (2010), “Self-concordant analysis for logistic regression,” *Electronic Journal of Statistics*, 4, 384–414.
- Baraud Y (2002), “Non-asymptotic minimax rates of testing in signal detection,” *Bernoulli*, 8, 577–606.
- Barrett JC, Hansoul S, Nicolae DL, Cho JH, Duerr RH, Rioux JD, Brant SR, Silverberg MS, Taylor KD, Barmada MM, et al. (2008), “Genome-wide association defines more than 30 distinct susceptibility loci for Crohn’s disease,” *Nature Genetics*, 40, 955. [PubMed: 18587394]
- Belloni A, Chernozhukov V, and Wei Y (2016), “Post-selection inference for generalized linear models with many controls,” *Journal of Business & Economic Statistics*, 34, 606–619.
- Benjamini Y, and Hochberg Y (1995), “Controlling the false discovery rate: a practical and powerful approach to multiple testing,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 289–300.
- Benjamini Y, and Yekutieli D (2001), “The control of the false discovery rate in multiple testing under dependency,” *The Annals of Statistics*, 29, 1165–1188.
- Cai TT, and Guo Z (2017), “Confidence intervals for high-dimensional linear regression: Minimax rates and adaptivity,” *The Annals of Statistics*, 45, 615–646.
- Cai TT, Li H, Ma J, and Xia Y (2017), “Differential Markov random field analysis with applications to detecting differential microbial community structures,” *Unpublished Manuscript*.
- Cai TT, Liu W, and Xia Y (2013), “Two-sample covariance matrix testing and support recovery in high-dimensional and sparse settings,” *Journal of the American Statistical Association*, 108, 265–277.
- (2014), “Two-sample test of high dimensional means under dependence,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76, 349–372.
- Candès E, Fan Y, Janson L, and Lv J (2018), “Panning for gold: model-X knockoffs for high dimensional controlled variable selection,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80, 551–577.
- Candès EJ, and Sur P (2018), “The phase transition for the existence of the maximum likelihood estimate in high-dimensional logistic regression,” *arXiv preprint arXiv:1804.09753*.
- Fortin G (2011), “L-Carnitine and intestinal inflammation,” in *Vitamins & Hormones*, Elsevier, vol. 86, pp. 353–366. [PubMed: 21419279]
- Ingster YI (1993), “Asymptotically minimax hypothesis testing for nonparametric alternatives. I, II, III,” *Mathematical Methods of Statistics*, 2, 85–114.
- Ingster YI, Tsybakov AB, and Verzelen N (2010), “Detection boundary in sparse regression,” *Electronic Journal of Statistics*, 4, 1476–1526.
- Javanmard A, and Javadi H (2019), “False discovery rate control via debiased lasso,” *Electronic Journal of Statistics*, 13, 1212–1253.
- Javanmard A, and Montanari A (2014a), “Confidence intervals and hypothesis testing for high-dimensional regression,” *Journal of Machine Learning Research*, 15, 2869–2909.
- (2014b), “Hypothesis testing in high-dimensional regression under the gaussian random design model: Asymptotic theory,” *IEEE Transactions on Information Theory*, 60, 6522–6554.
- Lee D, Baldassano RN, Otley AR, Albenberg L, Griffiths AM, Compher C, Chen EZ, Li H, Gilroy E, Nessel L, et al. (2015), “Comparative effectiveness of nutritional and biological therapy in North American children with active Crohn’s disease,” *Inflammatory Bowel Diseases*, 21, 1786–1793. [PubMed: 25970545]
- Lewis JD, Chen EZ, Baldassano RN, Otley AR, Griffiths AM, Lee D, Bittinger K, Bailey A, Friedman ES, Hoffmann C, et al. (2015), “Inflammation, antibiotics, and diet as environmental stressors of the gut microbiome in pediatric Crohn’s disease,” *Cell Host & Microbe*, 18, 489–500. [PubMed: 26468751]

- Liu W (2013), “Gaussian graphical model estimation with false discovery rate control,” *The Annals of Statistics*, 41, 2948–2978.
- Liu W, and Luo S (2014), “Hypothesis testing for high-dimensional regression models,” Technical report.
- Maceyka M, and Spiegel S (2014), “Sphingolipid metabolites in inflammatory disease,” *Nature*, 510, 58. [PubMed: 24899305]
- Meier L, van de Geer S, and Bühlmann P (2008), “The group lasso for logistic regression,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70, 53–71.
- Mukherjee R, Pillai NS, and Lin X (2015), “Hypothesis testing for high-dimensional sparse binary regression,” *Annals of Statistics*, 43, 352–381. [PubMed: 26246645]
- Negahban S, Ravikumar P, Wainwright MJ, and Yu B (2010), “A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers,” Technical Report Number 979.
- Ni J, Shen T-CD, Chen EZ, Bittinger K, Bailey A, Roggiani M, Sirota-Madi A, Friedman ES, Chau L, Lin A, et al. (2017), “A role for bacterial urease in gut dysbiosis and Crohn’s disease,” *Science Translational Medicine*, 9, eaah6888. [PubMed: 29141885]
- Peltekova VD, Wintle RF, Rubin LA, Amos CI, Huang Q, Gu X, Newman B, Van Oene M, Cescon D, Greenberg G, et al. (2004), “Functional variants of OCTN cation transporter genes are associated with Crohn’s disease,” *Nature Genetics*, 36, 471. [PubMed: 15107849]
- Plan Y, and Vershynin R (2013), “Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach,” *IEEE Transactions on Information Theory*, 59, 482–494.
- Ren Z, Zhang C-H, and Zhou HH (2016), “Asymptotic normality in estimation of large Ising graphical model,” Unpublished Manuscript.
- Rozenova KA, Deevska GM, Karakashian AA, and Nikolova-Karakashian MN (2010), “Studies on the role of acid sphingomyelinase and ceramide in the regulation of tumor necrosis factor α (TNF α)-converting enzyme activity and TNF α secretion in macrophages,” *Journal of Biological Chemistry*, 285, 21103–21113.
- Sewell GW, Hannun YA, Han X, Koster G, Bielawski J, Goss V, Smith PJ, Rahman FZ, Vega R, Bloom SL, et al. (2012), “Lipidomic profiling in Crohn’s disease: abnormalities in phosphatidylinositols, with preservation of ceramide, phosphatidylcholine and phosphatidylserine composition,” *The International Journal of Biochemistry & Cell Biology*, 44, 1839–1846. [PubMed: 22728312]
- Shekhawat PS, Sonne S, Carter AL, Matern D, and Ganapathy V (2013), “Enzymes involved in L-carnitine biosynthesis are expressed by small intestinal enterocytes in mice: Implications for gut health,” *Journal of Crohn’s & Colitis*, 7, e197–e205.
- Sur P, and Candès EJ (2019), “A modern maximum-likelihood theory for high-dimensional logistic regression,” *Proceedings of the National Academy of Sciences*, 116, 14516–14525.
- Sur P, Chen Y, and Candès EJ (2017), “The likelihood ratio test in high-dimensional logistic regression is asymptotically a rescaled chi-square,” *Probability Theory and Related Fields*, 1–72.
- Tsybakov AB (2009), *Introduction to Nonparametric Estimation*, Springer Series in Statistics. Springer, New York.
- van de Geer S (2008), “High-dimensional generalized linear models and the lasso,” *The Annals of Statistics*, 36, 614–645.
- van de Geer S, Bühlmann P, Ritov Y, and Dezeure R (2014), “On asymptotically optimal confidence regions and tests for high-dimensional models,” *The Annals of Statistics*, 42, 1166–1202.
- Verzelen N (2012), “Minimax risks for sparse regressions: Ultra-high dimensional phenomena,” *Electronic Journal of Statistics*, 6, 38–90.
- Xia Y, Cai T, and Cai TT (2015), “Testing differential networks with applications to the detection of gene-gene interactions,” *Biometrika*, 102, 247–266. [PubMed: 28502988]
- (2018), “Two-sample tests for high-dimensional linear regression with an application to detecting interactions,” *Statistica Sinica*, 28, 63–92. [PubMed: 29386856]
- Xia Y, Cai TT, and Li H (2018), “Joint testing and false discovery rate control in high-dimensional multivariate regression,” *Biometrika*, 105, 249–269. [PubMed: 30799872]

- Yang Z, Wang Z, Liu H, Eldar YC, and Zhang T (2015), “Sparse nonlinear regression: Parameter estimation and asymptotic inference,” arXiv preprint arXiv:1511.04514.
- Zaitsev AY (1987), “On the Gaussian approximation of convolutions under multidimensional analogues of S.N. Bernstein’s inequality conditions,” *Probability Theory and Related Fields*, 74, 535–566.
- Zhang C-H, and Zhang SS (2014), “Confidence intervals for low dimensional parameters in high dimensional linear models,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76, 217–242.

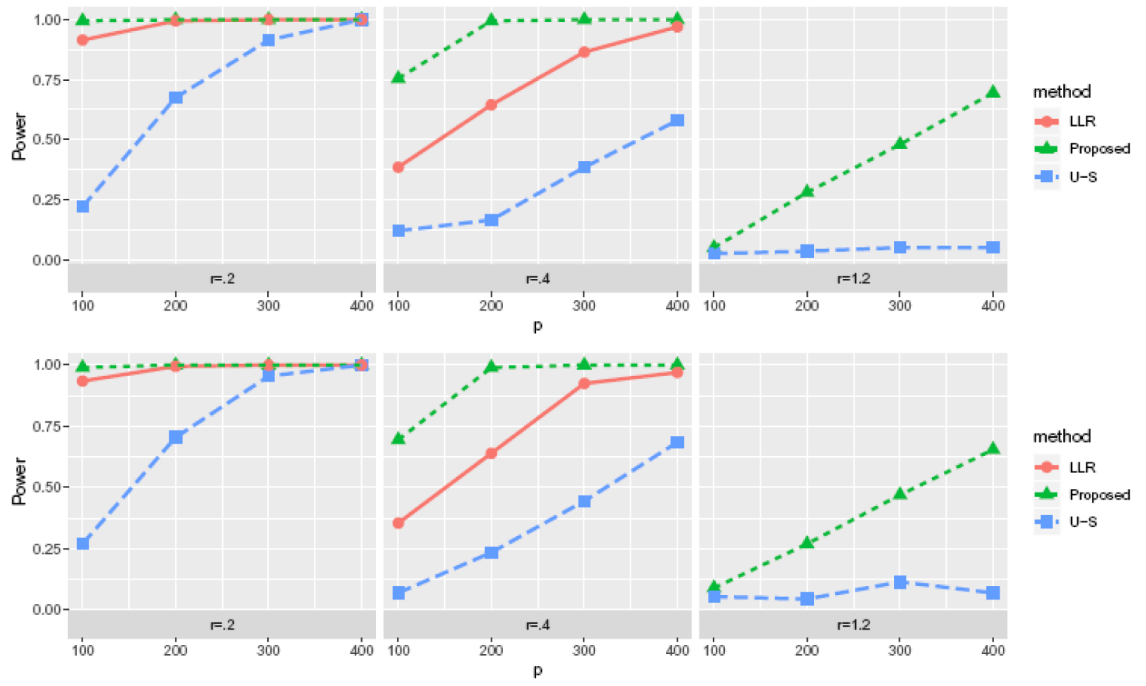


Fig. 1. Empirical power with $\alpha = 0.05$ for the proposed method (Proposed), the Bonferroni corrected univariate screening method (U-S) and the Bonferroni corrected likelihood ratio based method of Sur and Candès (2019) (LLR). Top panel: $k = 2$; bottom panel: $k = 4$.




Fig. 2. Boxplots of the empirical FDRs across all the settings for $\alpha = 0.2$.

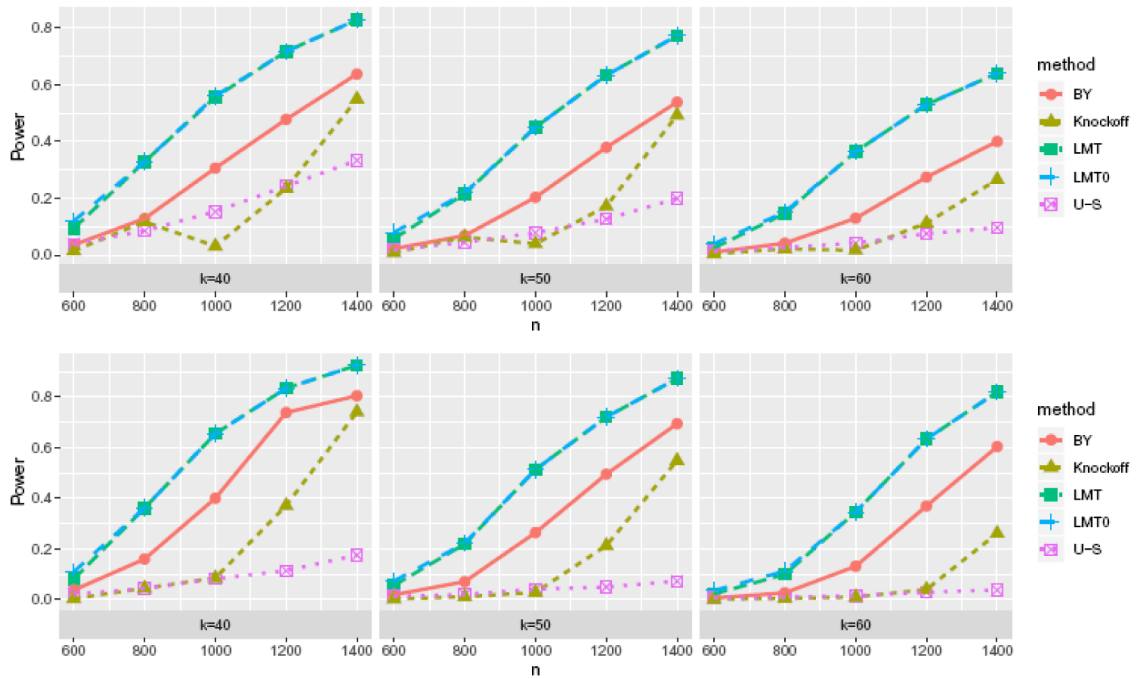


Fig. 3. Empirical power under FDR $\alpha = 0.2$ for $\rho = 3$ (top) and $\rho = 4$ (bottom).

Table 1

Computation of v_j from the Lasso (8)

Input:	An upper bound ζ_j^* for ζ_j with default value $\zeta^* = \sqrt{2\log p}$.
	tuning parameters $\kappa_0 \in [0,1]$ and $\kappa_1 \in (0,1]$;
Step 1:	If $\zeta_j(\lambda) > \zeta_j^*$ for all $\lambda = 0$, set $\zeta_j^* = (1 + \kappa_1) \inf_{\lambda > 0} \zeta_j(\lambda)$;
	$\lambda \leftarrow \max\{\lambda: \zeta_j(\lambda) \leq \zeta_j^*\}, \zeta_j^* \leftarrow \zeta_j(\lambda), \tau_j^* \leftarrow \tau_j(\lambda)$;
Step 2:	$\lambda_j \leftarrow \min\{\lambda: \tau_j(\lambda) \leq (1 + \kappa_0)\tau_j^*\}$;
	$v_j \leftarrow v_j(\lambda_j), \tau_j \leftarrow \tau_j(\lambda_j), \zeta_j \leftarrow \zeta_j(\lambda_j)$
Input:	An upper bound ζ_j^* for ζ_j with default value $\zeta^* = \sqrt{2\log p}$.
Output:	$\lambda_j, v_j, \tau_j, \zeta_j$

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2

Type I error with $\alpha = 0.05$ for the proposed method (Proposed), the Bonferroni corrected univariate screening method (U-S) and the Bonferroni corrected likelihood ratio based method of Sur and Candès (2019) (LLR), for different n , p and k .

p/n	$k = 2$				$k = 4$			
	$p = 100$	200	300	400	$p = 400$	600	800	1000
Proposed								
0.2	0.052	0.066	0.042	0.054	0.058	0.050	0.046	0.070
0.4	0.038	0.054	0.062	0.054	0.046	0.050	0.060	0.074
1.2	0.026	0.044	0.042	0.045	0.014	0.044	0.054	0.054
U-S								
0.2	0.040	0.032	0.024	0.018	0.018	0.022	0.028	0.034
0.4	0.050	0.032	0.024	0.020	0.028	0.028	0.032	0.046
1.2	0.028	0.038	0.024	0.020	0.032	0.018	0.034	0.014
LLR								
0.2	0.050	0.050	0.068	0.040	0.058	0.044	0.046	0.034
0.4	0.084	0.070	0.048	0.056	0.062	0.042	0.058	0.064

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 3

Empirical performance of LMT_V with FDV level $r = 10$.

ρ	p	k	Empirical FDV				Empirical Power			
			$n = 400$	600	800	1000	400	600	800	1000
		40	4.07	5.45	6.44	7.11	0.08	0.23	0.40	0.59
	800	50	4.30	6.29	7.27	8.26	0.06	0.16	0.32	0.49
		60	4.33	6.63	7.48	8.42	0.05	0.12	0.25	0.42
		40	3.30	4.59	5.79	6.82	0.06	0.18	0.35	0.52
3	1000	50	3.49	5.42	6.43	7.03	0.05	0.13	0.26	0.43
		60	3.68	5.47	7.29	7.97	0.03	0.09	0.20	0.34
		40	2.69	4.36	5.00	5.68	0.05	0.15	0.31	0.46
	1200	50	2.97	4.22	5.73	6.43	0.03	0.11	0.21	0.36
		60	2.78	4.91	5.91	7.25	0.02	0.07	0.16	0.27

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 4

Significant metabolites associated with Crohn's disease (coded as 1 in logistic regression) at the baseline, one week and 8 weeks after treatment with FDR < 5%. The refitted regression coefficients show the direction of the association.

Disease Stage	HMDB ID	Synonyms	Refitted Coefficient
Baseline	00885	C16:0.cholesteryl ester	4.45
	12097	C14:0.sphingomyelin	1.74
	04953	C24:1.Ceramide.(d18:1)	4.25
Week 1	00555	3-methyladipate/pimelate	-12.82
	06726	C20:4.cholesteryl ester	2.17
	12097	C14:0.sphingomyelin	2.06
	04949	C16:0.Ceramide.(d18:1)	0.87
	00555	3-methyladipate/pimelate	-6.10
Week 8	00056	beta-alanine	2.95
	00448	adipate	-4.50
	00883	valine	1.40
	00222	C16.carnitine	0.58
	00848	C18.carnitine	0.39
	00555	3-methyladipate/pimelate	-5.95
	00056	beta-alanine	0.63

Table 5

Significant metabolites identified via logistic regression of responder vs normal control and non-responder vs normal control for FDR = 5%.

HMDB ID	Synonyms	Refitted Coefficients	
		Responder vs. Normal	Non-Responder vs. Normal
06726	C20:4.cholesteryl ester	0.139	1.854
01043	Linoleic.acid	-0.686	-0.388
00472	5-hydroxytryptophan	1.000	1.034
00056	beta-alanine	0.503	2.298
00883	valine	0.628	0.530
00848	C18.carnitine	1.100	0.457
01488	nicotinate	-1.936	-4.312
00254	succinate	0.750	1.508
00555	3-methyladipate/pimelate	-1.989	-4.209