

# GENERALIZED MATRIX DECOMPOSITION REGRESSION: ESTIMATION AND INFERENCE FOR TWO-WAY STRUCTURED DATA

BY YUE WANG<sup>1</sup>, ALI SHOJAIE<sup>2,\*</sup>, TIMOTHY RANDOLPH<sup>3,†</sup>, PARKER  
KNIGHT<sup>4</sup> AND JING MA<sup>5,‡</sup>

<sup>1</sup>*Department of Biostatistics and Informatics, University of Colorado Anschutz Medical Campus, [yue.2.wang@cuanschutz.edu](mailto:yue.2.wang@cuanschutz.edu)*

<sup>2</sup>*Department of Biostatistics, University of Washington, [ashojaie@uw.edu](mailto:ashojaie@uw.edu)*

<sup>3</sup>*Clinical Research Division, Fred Hutchinson Cancer Center, [trandolp@fredhutch.org](mailto:trandolp@fredhutch.org)*

<sup>4</sup>*Department of Biostatistics, Harvard University, [pknight@g.harvard.edu](mailto:pknight@g.harvard.edu)*

<sup>5</sup>*Public Health Sciences Division, Fred Hutchinson Cancer Center, [jingma@fredhutch.org](mailto:jingma@fredhutch.org)*

Motivated by emerging applications in ecology, microbiology, and neuroscience, this paper studies high-dimensional regression with two-way structured data. To estimate the high-dimensional coefficient vector, we propose the generalized matrix decomposition regression (GMDR) to efficiently leverage auxiliary information on row and column structures. GMDR extends the principal component regression (PCR) to two-way structured data, but unlike PCR, GMDR selects the components that are most predictive of the outcome, leading to more accurate prediction. For inference on regression coefficients of individual variables, we propose the generalized matrix decomposition inference (GMDI), a general high-dimensional inferential framework for a large family of estimators that include the proposed GMDR estimator. GMDI provides more flexibility for incorporating relevant auxiliary row and column structures. As a result, GMDI does not require the true regression coefficients to be sparse, but constrains the coordinate system representing the regression coefficients according to the column structure. GMDI also allows dependent and heteroscedastic observations. We study the theoretical properties of GMDI in terms of both the type-I error rate and power and demonstrate the effectiveness of GMDR and GMDI in simulation studies and an application to human microbiome data.

**1. Introduction.** We consider the problem of regressing a scalar outcome from  $n$  observations on a vector of  $p$  predictors, formally,  $\mathbb{E}(y) = \mathbf{x}^\top \boldsymbol{\beta}$ , in settings where it may be implausible to assume that the  $p$  variables or the  $n$  samples are independent. To address this problem, we account for the sample- and variable-wise dependencies to provide a framework for estimation of the coefficient vector,  $\boldsymbol{\beta}$ , and inference on the individual coefficients,  $\beta_j$  ( $j = 1, \dots, p$ ). The proposed framework is motivated by the increasing occurrence of high-dimensional two-way structured data—that is, data with structures among the variables (columns) and samples (rows)—in ecology, microbiology, and neuroscience. Informative two-way structures can often be obtained from various auxiliary sources *a priori* (Allen, Grosebeck and Taylor, 2014; Li, Cai and Li, 2021). In many applications, the goal is to examine associations between such structured data and an outcome of interest. One application that motivated the current work comes from human microbiome data which record the composition and function of bacterial taxa. These data are used to investigate the role of human

---

\*Supported by National Institutes of Health R01GM133848.

†Supported by National Institutes of Health R01GM129512, R01HL1554178, and P50CA228944.

‡Supported by National Institutes of Health R01GM145772.

*Keywords and phrases:* dimensionality reduction, high-dimensional inference, microbiome data, prediction, two-way structured data.

microbiome in health and diseases. An interesting property of these data is that taxa are related to one another, both evolutionarily and functionally. Evolutionary relationships among taxa are typically characterized by a phylogenetic tree, or dendrogram, whose nodes represent taxonomic assignments based on genomic similarities (Washburne et al., 2018). Their functional relationships may be characterized by genomic content known to contribute to a biological process (Sharifi and Ye, 2017).

To motivate our regression framework, we consider data from a study investigating age-associated microbial signatures across geographic regions (Yatsunenko et al., 2012). In this example, stool samples from  $n = 100$  individuals from the Amazonas of Venezuela, rural Malawi, and US metropolitan areas were processed to identify  $p = 149$  genus-level bacterial abundances. Figure 1A shows a principal-component (PC) plot of the configuration of samples based on the first two PCs of the  $n \times p$  microbiome data matrix; samples are colored by the logarithm of each individual’s age, which range from a few months to over 50 years. This plot suggests a strong association between age and microbial composition. This is further supported by Fig. 1B, a volcano plot of the log 10-transformed  $p$ -values versus the estimated coefficients obtained from a univariate regression of each genus on age. Red dots represent bacteria that have statistically significant marginal associations with age after controlling the false discovery rate (FDR) at 0.1 using the Benjamini–Yekutieli procedure (Benjamini and Yekutieli, 2001); purple dots represent bacteria with  $p$ -values less than 0.05 that are no longer statistically significant after controlling the FDR; cyan dots represent bacteria for which  $p$ -values are greater than 0.05. Figure 1B shows that the majority of bacteria (105 out of 149) are marginally associated with age after controlling the FDR at 0.1. This type of analysis, however, does not account for the relationships between either the taxa or the individuals from which the samples were taken. As noted above, bacteria tend to be correlated via their phylogeny, and individuals also tend to be correlated in their microbial composition due to shared households, diets, and/or cultures (Zeevi et al., 2019; Hullar et al., 2021).

These structures are commonly acknowledged in the analysis of microbiome data. For example, phylogeny-aware distances between samples (e.g., UniFrac, Lozupone and Knight, 2005) are used in the principal coordinate analysis (PCoA) and in kernel-based association tests (Zhao et al., 2015). In an extension of PCoA, Wang et al. (2019) used the generalized matrix decomposition (GMD, Allen, Grosebeck and Taylor, 2014) to produce dimension-reduced plots like PCoA while leveraging similarities among the taxa and among the samples. This approach is illustrated in Fig. 1C, which shows a GMD-biplot of sample configurations (dots) and corresponding variable loadings (arrows) in these coordinates. Here, the coordinate system is derived by extending the singular value decomposition in a manner that accounts for both row and column structures. More specifically, the structure among taxa is characterized by a  $p \times p$  similarity kernel derived from the patristic distance between each pair of tips of a phylogenetic tree. The structure among samples is derived from extrinsic data based on bacterial genes: the functional protein content produced by the bacteria in each sample is estimated by classifying genes according to Enzyme Commission (EC) numbers (Cuesta et al., 2015); see Section 5 for more details. Then, an  $n \times n$  matrix of pairwise sample similarities based on EC numbers provides a biologically-informed auxiliary representation of sample-based structures. The two axes in Fig. 1C are the first two columns of the right GMD vectors. Each sample is represented by the coordinates of the projection of its microbial abundance vector onto the two axes and is colored by the logarithm of the subject’s age. An arrow is then plotted for each taxon, its coordinates coming from the first two columns of the right GMD vectors. Compared to Fig. 1A, the GMD-biplot provides an alternative two-dimensional configuration of samples; it shows a strong age-dependent variation and many tightly clustered arrows (genera) contributing to this configuration. Consistent with Fig. 1B, this biplot suggests that there are many correlated age-associated taxa. This analysis, however, is unsupervised and any inference made about the associations is circumstantial. It is

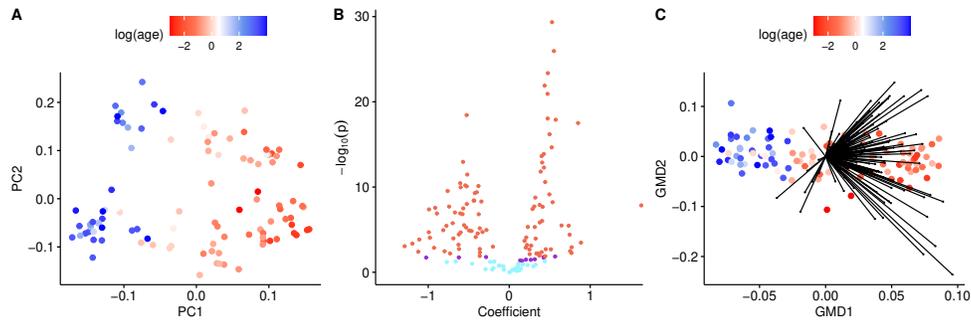


Fig 1: (A): The PC-plot of data from [Yatsunen et al. \(2012\)](#). (B): The volcano plot showing the log 10-transformed  $p$ -values for the associations of the bacteria with age versus the corresponding regression coefficients. Cyan dots represent bacteria for which  $p$ -values are greater than 0.05; purple dots represent bacteria for which  $p$ -values are less than 0.05 but no longer show statistically significant associations after controlling the FDR at 0.1; red dots represent bacteria that still show statistically significant associations after controlling the FDR at 0.1. (C): The GMD-plot of data from [Yatsunen et al. \(2012\)](#); metagenomic similarities among samples and phylogenetic similarities among taxa are considered.

desirable, therefore, to develop a supervised analytical framework of high-dimensional regression that leverages auxiliary row and column structures, and, importantly, provides valid inference for the associations between the taxa and a response variable.

1.1. *Our Contributions.* This paper introduces the GMD regression (GMDR), a dimension reduction-based estimation procedure that efficiently leverages pre-specified two-way structures. GMDR is built upon the generalized matrix decomposition (GMD, [Allen, Grose and Taylor, 2014](#); [Escoufier, 2006](#)), which extends the singular value decomposition (SVD) to incorporate auxiliary two-way structures and will be reviewed in Section 2. Thus, GMDR can be viewed as an extension of principal component regression (PCR) for analyzing two-way structured data. However, unlike PCR which uses top principal components as the predictors, our GMDR selects the GMD components that are most predictive of the outcome. This novel selection procedure ensures a more accurate prediction using GMDR.

We further define a broad class of estimators for high-dimensional regression on two-way structured data by leveraging the connection between dimension reduction-based regression (e.g., PCR) and penalized regression (e.g., ridge regression), which is discussed in detail in Section 2. This connection also allows us to develop the GMD inference (GMDI) framework, a high-dimensional inference (HDI) procedure that can assess the statistical significance of individual variables based on any arbitrary estimator in this class. As such, GMDI can be applied to not only the proposed GMDR but also many existing estimation procedures that lack inferential procedures for individual variables, such as PCR, generalized ridge regression ([Golub and Van Loan, 2013](#)), and the kernel penalized regression (KPR, [Randolph et al., 2018](#)). GMDI has three distinct features. First, unlike most existing HDI tools that assume *i.i.d* samples, which may not hold for two-way structured data, GMDI allows for dependent and heteroscedastic samples by efficiently leveraging auxiliary row structures. Ignoring sample correlations may lead to incorrect inference even in low-dimensional settings. Second, existing HDI tools, including [Bühlmann \(2013\)](#); [Zhang and Zhang \(2014\)](#); [Javanmard and Montanari \(2014a,b\)](#); [van de Geer et al. \(2014\)](#); [Belloni, Chernozhukov and Kato \(2015\)](#); [Zhao and Shojaie \(2016\)](#); [Mitra and Zhang \(2016\)](#); [Ning and Liu \(2017\)](#); [Zhu and Bradic \(2018\)](#), all require at least one of the following assumptions: (i) the regression coefficient

vector is sparse, (ii) the design matrix satisfies a restricted eigenvalue-type condition if a fixed design is considered, and (iii) the precision matrix of the variables in the design matrix has row sparsity if a random design is considered. However, these conditions may fail when strong correlations exist among variables, which is common for two-way structured data. Third, GMDI provides flexibility for users to specify relevant auxiliary row and column structures. In particular, we provide methods to avoid uninformative structures and to incorporate partially informative structures, leading to well-controlled type-I error rates and guaranteed power.

Regarding the second property, it may be that a majority of variables are *marginally* associated with the outcome, as appears to be the case in Fig. 1B. This has two possible explanations: (i) a large number of the variables are also *conditionally* associated with the outcome; (ii) these variables are highly correlated, but only a few of these are conditionally associated with the outcome. In the first situation, the vector of regression coefficients is not sparse; in the second situation, any restricted eigenvalue-type condition may fail (see [van de Geer et al., 2009](#)), and likely, the precision matrix of the variables is not sparse. As an alternative to these assumptions, GMDI assumes the pre-specified column structure informs the structure of the regression coefficients, which reduces to sparsity when no column structure is pre-specified.

GMDI follows the general idea of bias correction for ridge-type estimators ([Bühlmann, 2013](#)) but uses a novel initial estimator that efficiently leverages the pre-specified two-way structures. We derive the asymptotic distribution of the bias-corrected estimator. Based on this, we construct asymptotically valid two-sided  $p$ -values and provide sufficient conditions under which GMDI offers guaranteed power. We introduce a procedure that selects against uninformative sample structure. We also show that the GMDI results are robust to misspecification. Our numerical studies demonstrate the superior performance of GMDI for two-way structured regression compared to existing HDI methods, even when pre-specified structures are not fully informative.

**1.2. Organization and Notation.** The rest of the paper is organized as follows. In Section 2, we first introduce the GMDR estimation/prediction framework, accompanied by the novel procedure for the selection of GMD components. We then link the GMDR estimator to a broad class of estimators. In Section 3, we present the GMDI procedure for any arbitrary estimator in this class, explain the rationale behind the key assumptions for GMDI, and provide ways to assess the informativeness of the pre-specified structures and incorporate partially informative structures. Multiple simulation studies, including one based on real data, are presented in Section 4 to examine the finite-sample performance of GMDI. In Section 5, we demonstrate the effectiveness of GMDR and GMDI on an application to microbiome data. Section 6 summarizes our findings and outlines potential extensions. Technical proofs are provided in the supplement.

Throughout the paper, we use normal typeface to denote scalars, bold lowercase typeface to denote vectors, and bold uppercase typeface to denote matrices. For any vector  $\mathbf{v} \in \mathbb{R}^p$ , we use  $v_j$  to denote the  $j$ -th element of  $\mathbf{v}$  for  $j = 1, \dots, p$ . For any matrix  $\mathbf{M} \in \mathbb{R}^{n \times p}$ ,  $\mathbf{m}_j$  and  $m_{ij}$  denote, respectively, the  $j$ -th column and  $(i, j)$  entry of  $\mathbf{M}$  for  $i = 1, \dots, n$  and  $j = 1, \dots, p$ . For any index set  $\mathcal{I} \subset \{1, \dots, p\}$ ,  $\mathbf{v}_{\mathcal{I}}$  and  $\mathbf{M}_{\mathcal{I}}$  denote, respectively, the subvector of  $\mathbf{v}$  whose elements are indexed by  $\mathcal{I}$  and the submatrix of  $\mathbf{M}$  whose columns are indexed by  $\mathcal{I}$ . The indicator function  $\mathbb{1}(\mathcal{A})$  denotes the occurrence of the event  $\mathcal{A}$ ; i.e.,  $\mathbb{1}(\mathcal{A}) = 1$  if  $\mathcal{A}$  is true, and  $\mathbb{1}(\mathcal{A}) = 0$ , otherwise. We denote  $\|\mathbf{v}\|_0 = \sum_{j=1}^p \mathbb{1}(v_j \neq 0)$ ,  $\|\mathbf{v}\|_q = \left(\sum_{j=1}^p |v_j|^q\right)^{1/q}$  for any  $0 < q < \infty$ ,  $\|\mathbf{v}\|_\infty = \max_j |v_j|$ ,  $\|\mathbf{v}\|_{\mathbf{K}}^2 = \mathbf{v}^\top \mathbf{K} \mathbf{v}$  for any positive semi-definite matrix  $\mathbf{K}$ ,  $\|\mathbf{M}\|_q = \sup_{\|\mathbf{v}\|_q=1} \|\mathbf{M}\mathbf{v}\|_q$  for any  $q > 0$  and  $\|\mathbf{M}\|_F^2 = \sum_{i=1}^n \sum_{j=1}^p m_{ij}^2$ . Finally, for any square matrix  $\mathbf{S}$ , we denote the trace of  $\mathbf{S}$  as  $\text{tr}(\mathbf{S})$ .

**2. The GMD regression.** Consider the following linear model

$$(1) \quad \mathbf{y} = \mathbf{X}\boldsymbol{\beta}^* + \boldsymbol{\epsilon},$$

where  $\mathbf{X} \in \mathbb{R}^{n \times p}$  denotes the structured design matrix,  $\mathbf{y} \in \mathbb{R}^n$  is the response variable, and  $\boldsymbol{\beta}^* \in \mathbb{R}^p$  is the underlying true regression coefficient. We allow  $p$  to be greater than  $n$ . In addition, we assume that  $\boldsymbol{\epsilon}$  is a vector of random noises with  $\mathbb{E}[\boldsymbol{\epsilon} | \mathbf{X}] = \mathbf{0}_n$  and  $\text{Cov}(\boldsymbol{\epsilon} | \mathbf{X}) = \boldsymbol{\Psi}$ , where  $\mathbf{0}_n$  is an  $n \times 1$  vector of zeros and  $\boldsymbol{\Psi}$  is an  $n \times n$  positive definite matrix. By considering the non-identity matrix  $\boldsymbol{\Psi}$ , we do not assume that entries of  $\boldsymbol{\epsilon}$  are *i.i.d.*, allowing for samples to be correlated and heteroscedastic. Let  $\mathbf{H} \in \mathbb{R}^{n \times n}$  and  $\mathbf{Q} \in \mathbb{R}^{p \times p}$  denote two auxiliary positive definite matrices, capturing similarities among rows and columns of  $\mathbf{X}$ , respectively. More specifically, we assume that entries of  $\mathbf{H}$  ( $\mathbf{Q}$ ) inform the *conditional similarity* between samples (variables); that is, the similarity between samples (variables) after the effects of other samples (variables) are removed. This implies that, for instance,  $\mathbf{H}$  provides information about  $\boldsymbol{\Psi}$ , and their connection will be made explicit in Assumption (A1). We assume that  $\mathbf{X}$ ,  $\mathbf{H}$  and  $\mathbf{Q}$  are deterministic quantities and refer to the triple  $(\mathbf{X}, \mathbf{H}, \mathbf{Q})$  as two-way structured data hereafter. Throughout the article, we assume that  $\mathbf{X}$  and  $\mathbf{y}$  are appropriately centered such that  $\mathbf{1}_n^\top \mathbf{H} \mathbf{y} = 0$  and  $\mathbf{1}_n^\top \mathbf{H} \mathbf{X} = \mathbf{0}_p^\top$ , where  $\mathbf{1}_n$  is an  $n \times 1$  vector of all ones. We will study the estimation and inference of the high-dimensional parameters  $\boldsymbol{\beta}^*$ , while leveraging the information from  $\mathbf{H}$  and  $\mathbf{Q}$ .

Our idea is built upon the generalized matrix decomposition (GMD), which we will review next. The GMD of  $\mathbf{X}$  with respect to  $\mathbf{H}$  and  $\mathbf{Q}$  is  $\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^\top$ , where the components are obtained by solving the optimization problem

$$(2) \quad \text{argmin}_{\mathbf{U}, \mathbf{S}, \mathbf{V}} \|\mathbf{X} - \mathbf{U}\mathbf{S}\mathbf{V}^\top\|_{\mathbf{H}, \mathbf{Q}},$$

subject to  $\mathbf{U}^\top \mathbf{H} \mathbf{U} = \mathbf{I}_K$ ,  $\mathbf{V}^\top \mathbf{Q} \mathbf{V} = \mathbf{I}_K$  and  $\mathbf{S} = \text{diag}(\sigma_1, \dots, \sigma_K)$ . Here,  $K \leq \min(n, p)$  is the rank of  $\mathbf{X}^\top \mathbf{H} \mathbf{X} \mathbf{Q}$  and  $\|\mathbf{M}\|_{\mathbf{H}, \mathbf{Q}}^2 = \text{tr}(\mathbf{M}^\top \mathbf{H} \mathbf{M} \mathbf{Q})$  for any matrix  $\mathbf{M} \in \mathbb{R}^{n \times p}$ . Note that unlike SVD, the GMD vectors  $\mathbf{U}$  and  $\mathbf{V}$  are not orthogonal in the Euclidean norm unless  $\mathbf{H} = \mathbf{I}_n$  and  $\mathbf{Q} = \mathbf{I}_p$ . GMD directly extends SVD by replacing the Frobenius norm with the  $(\mathbf{H}, \mathbf{Q})$ -norm  $\|\cdot\|_{\mathbf{H}, \mathbf{Q}}$ . As such, GMD preserves appealing properties of SVD such as ordering the component vectors according to a nonincreasing set of GMD values,  $\sigma_1, \dots, \sigma_K$ , indicating that the decomposition of the total variance of  $\mathbf{X}$  into each dimension is non-increasing. An efficient algorithm was proposed by [Allen, Groseknick and Taylor \(2014\)](#) to iteratively solve for each column of  $\mathbf{U}$ ,  $\mathbf{S}$  and  $\mathbf{V}$  in (2). Analogous to the SVD of  $\mathbf{X}$ , which is closely related to the eigen-decomposition of  $\mathbf{X}^\top \mathbf{X}$ , the GMD of  $\mathbf{X}$  with respect to  $\mathbf{H}$  and  $\mathbf{Q}$  is related to the eigen-decomposition of  $\mathbf{X}^\top \mathbf{H} \mathbf{X} \mathbf{Q}$ . In fact, [Escoufier \(1987\)](#) and [Allen, Groseknick and Taylor \(2014\)](#) show that the squared GMD values  $\sigma_1^2, \dots, \sigma_K^2$  are non-zero eigenvalues of  $\mathbf{X}^\top \mathbf{H} \mathbf{X} \mathbf{Q}$ , and columns of  $\mathbf{V}$  are the corresponding eigenvectors. Note that  $\mathbf{X}^\top \mathbf{H} \mathbf{X} \mathbf{Q}$  may not be symmetric, again implying that columns of  $\mathbf{V}$  may not be orthogonal in the Euclidean norm. Given  $\mathbf{V}$  and  $\mathbf{S}$ , the  $n \times K$  matrix  $\mathbf{U}$  can be uniquely defined by  $\mathbf{U}\mathbf{S} = \mathbf{X}\mathbf{Q}\mathbf{V}$ .

Similar to PCR, the GMDR estimate of  $\boldsymbol{\beta}^*$  in (1) is obtained by regressing  $\mathbf{y}$  on a reduced subset of GMD components. More specifically, let  $\boldsymbol{\nu}_j = \mathbf{u}_j \sigma_j$  be the  $j$ -th GMD component for  $j = 1, \dots, K$  and set  $\boldsymbol{\Upsilon} = [\boldsymbol{\nu}_1 \cdots \boldsymbol{\nu}_K] \in \mathbb{R}^{n \times K}$ . For any fixed index set  $\mathcal{I} \subset \{1, \dots, K\}$ , the GMDR estimator of  $\boldsymbol{\beta}^*$ ,  $\hat{\boldsymbol{\beta}}_{\text{GMDR}}(\mathcal{I})$ , can be obtained in two steps:

- (i) Regress  $\mathbf{y}$  on  $\boldsymbol{\Upsilon}_{\mathcal{I}}$  and obtain  $\hat{\boldsymbol{\gamma}}(\mathcal{I}) = \text{argmin}_{\boldsymbol{\gamma}} \|\mathbf{y} - \boldsymbol{\Upsilon}_{\mathcal{I}} \boldsymbol{\gamma}\|_{\mathbf{H}}^2$ .
- (ii) Calculate  $\hat{\boldsymbol{\beta}}_{\text{GMDR}}(\mathcal{I}) = (\mathbf{Q}\mathbf{V})_{\mathcal{I}} \hat{\boldsymbol{\gamma}}(\mathcal{I})$ .

Letting  $w_j = \mathbb{1}(j \in \mathcal{I})$  for  $j = 1, \dots, K$  and  $\mathbf{W}_{\mathcal{I}} = \text{diag}(w_1, \dots, w_K)$ ,  $\hat{\boldsymbol{\beta}}_{\text{GMDR}}(\mathcal{I})$  can be explicitly expressed as

$$(3) \quad \hat{\boldsymbol{\beta}}_{\text{GMDR}}(\mathcal{I}) = \mathbf{Q}\mathbf{V}\mathbf{W}_{\mathcal{I}}\mathbf{S}^{-1}\mathbf{U}^\top \mathbf{H}\mathbf{y},$$

where  $\mathbf{U}, \mathbf{S}, \mathbf{V}$  are the GMD components of  $\mathbf{X}$  with respect to  $\mathbf{H}$  and  $\mathbf{Q}$ .

**REMARK.** Similar to SVD, GMD is not invariant to a scale transformation of the variables unless the same scale transformation is applied to all variables. Thus, our GMDR estimator  $\hat{\beta}_{\text{GMDR}}(\mathcal{I})$  is not invariant to a scale transformation of the predictors. Therefore, we recommend standardizing each predictor before implementing GMDR, especially in high-throughput sequencing studies where different variables may have different scales. However,  $\hat{\beta}_{\text{GMDR}}(\mathcal{I})$  is invariant to a scale transformation of  $\mathbf{H}$  and  $\mathbf{Q}$ .

The prediction performance of  $\hat{\beta}_{\text{GMDR}}(\mathcal{I})$  depends on the choice of the index set  $\mathcal{I}$ , which can be seen as a tuning parameter. Note that, if  $\mathbf{Q} = \mathbf{I}_p$  and  $\mathbf{H} = \mathbf{I}_n$ , then GMDR reduces to PCR. Thus, analogous to PCR, a natural way to select  $\mathcal{I}$  is to consider GMD components that correspond to large GMD values, referred to as top GMD components hereafter. However, since PCs are constructed without using the outcome, top PCs are not necessarily more predictive of the outcome than tail PCs (Cook, 2007). Thus, we propose an alternative approach to find the most predictive  $\mathcal{I}$  among all subsets of  $\{1, \dots, K\}$ . Note that an exhaustive search over all  $2^K$  subsets of  $\{1, \dots, K\}$  is computationally infeasible even for moderate  $K$ . To address this problem, we propose a procedure that weighs the importance of each GMD component by its contribution to the prediction of the outcome. Our idea is to decompose the total  $R^2$  of the model into  $K$  terms, each corresponding to a GMD component. Specifically, we first regress  $\mathbf{y}$  on all GMD components  $\mathbf{Y}$  with respect to the  $\mathbf{H}$ -norm, and obtain

$$(4) \quad \hat{\gamma} = \operatorname{argmin}_{\gamma} \|\mathbf{y} - \mathbf{Y}\gamma\|_{\mathbf{H}}^2.$$

It can then be seen that the total  $R^2$  for the model is given by  $R^2 = \|\mathbf{Y}\hat{\gamma}\|_{\mathbf{H}}^2 / \|\mathbf{y}\|_{\mathbf{H}}^2$ . Letting  $\hat{\gamma} = (\hat{\gamma}_1, \dots, \hat{\gamma}_K)^\top$ , we can write  $R^2 = \sum_{j=1}^K r_j^2$ , with each  $r_j$  represented explicitly in terms of  $\nu_j$ ,  $\sigma_j^2$ , and  $\hat{\gamma}_j$  as

$$r_j^2 = \frac{\|\nu_j \hat{\gamma}_j\|_{\mathbf{H}}^2}{\|\mathbf{y}\|_{\mathbf{H}}^2} = \frac{\sigma_j^2 \hat{\gamma}_j^2}{\|\mathbf{y}\|_{\mathbf{H}}^2}, \text{ for } j = 1, \dots, K.$$

Here, we use the fact that  $\nu_i^\top \mathbf{H} \nu_j = 0$  for any  $i \neq j$ . Since  $r_1^2, \dots, r_K^2$  share the same denominator, we define the *variable importance* (VI) score of the  $j$ -th GMD component as  $\text{VI}_j = \sigma_j^2 \hat{\gamma}_j^2$  for  $j = 1, \dots, K$ , with a higher score being more predictive of the outcome.

Based on  $\text{VI}_1, \dots, \text{VI}_K$ , we select the most predictive  $\mathcal{I}$  in three steps:

- (i) Sort  $\{\text{VI}_j : j = 1, \dots, K\}$  in nonincreasing order:  $\text{VI}_{j_1} \geq \text{VI}_{j_2} \geq \dots \geq \text{VI}_{j_K}$ .
- (ii) For each  $k = 1, \dots, K$ , consider  $\mathcal{I}_k = \{j_1, \dots, j_k\}$  and calculate the generalized cross-validation (GCV) statistic:

$$(5) \quad \text{GCV}(k) = \frac{\|(\mathbf{I}_n - \mathbf{G}(k))\mathbf{y}\|_{\mathbf{H}}^2}{(\operatorname{tr}(\mathbf{I}_n - \mathbf{G}(k)))^2} = \frac{\|(\mathbf{I}_n - \mathbf{G}(k))\mathbf{y}\|_{\mathbf{H}}^2}{(n - k)^2},$$

where  $\mathbf{G}(k) = \mathbf{Y}_{\mathcal{I}_k} (\mathbf{Y}_{\mathcal{I}_k}^\top \mathbf{H} \mathbf{Y}_{\mathcal{I}_k})^{-1} \mathbf{Y}_{\mathcal{I}_k}^\top \mathbf{H}$ .

- (iii) Find  $k_{\text{opt}} = \operatorname{argmin}_k \text{GCV}(k)$ , and obtain  $\mathcal{I}_{k_{\text{opt}}} = \{j_1, \dots, j_{k_{\text{opt}}}\}$ .

Having selected the most predictive GMD components, we now return to the estimation of regression coefficients. It can be seen from (3) that our GMDR estimator  $\hat{\beta}_{\text{GMDR}}(\mathcal{I})$  belongs to the following class of estimators:

$$(6) \quad \mathcal{B}_{\text{GMD}} = \{\beta^w \in \mathbb{R}^p : \beta^w = \mathbf{Q}\mathbf{V}\mathbf{W}\mathbf{S}^{-1}\mathbf{U}^\top \mathbf{H}\mathbf{y}\}$$

for some weight matrix  $\mathbf{W} = \operatorname{diag}(w_1, \dots, w_K)$ , where  $w_j \geq 0$  for  $j = 1, \dots, K$ . In addition to letting  $\mathbf{W}$  depend on the tuning index set  $\mathcal{I}$ , as done for GMDR, one can instead

let  $\mathbf{W}$  depend on a tuning parameter  $\eta$ . For example, letting  $w_j = w_j(\eta) = (\sigma_j^2 + \eta)^{-2}\sigma_j^2$  and  $\mathbf{W}_\eta = \text{diag}(w_1(\eta), \dots, w_K(\eta))$ , one can obtain another estimator in  $\mathcal{B}_{\text{GMD}}$  as  $\beta^w(\eta) = \mathbf{Q}\mathbf{V}\mathbf{W}_\eta\mathbf{S}^{-1}\mathbf{U}^\top\mathbf{H}\mathbf{y}$ . It can be shown that (see Section 1 of the supplement)

$$(7) \quad \beta^w(\eta) = \operatorname{argmin}_{\beta} \left\{ \|\mathbf{y} - \mathbf{X}\beta\|_{\mathbf{H}}^2 + \eta \|\beta\|_{\mathbf{Q}^{-1}}^2 \right\} := \widehat{\beta}_{\text{KPR}}(\eta),$$

where  $\widehat{\beta}_{\text{KPR}}(\eta)$  is the estimator obtained from the kernel penalized regression (KPR, [Randolph et al., 2018](#)). Although the motivations behind KPR and GMDR are quite different, (7) implies that they share many features. First, both  $\widehat{\beta}_{\text{GMDR}}(\mathcal{I})$  and  $\widehat{\beta}_{\text{KPR}}(\eta)$  are in the column space of  $\mathbf{Q}$ , indicating that both estimators incorporate information from  $\mathbf{Q}$  in similar ways. Second, both estimators exert shrinkage effects on the GMD components through the weight matrix  $\mathbf{W}$ . The difference is that  $\widehat{\beta}_{\text{GMDR}}(\mathcal{I})$  exerts discrete shrinkage by truncation, nullifying the contribution of the GMD components that are not selected, while  $\widehat{\beta}_{\text{KPR}}(\eta)$  exerts a smooth shrinkage effect through the tuning parameter  $\eta$  inherently involved in its construction. This connection between GMDR and KPR is similar to that between PCR and the ridge regression (see Section 3.4 in [Friedman et al. \(2001\)](#) for more details).

**3. The GMD Inference.** In this section, we propose a high-dimensional inferential framework for testing  $H_0 : \beta_l^* = 0$  for  $l = 1, \dots, p$ , called the GMD inference (GMDI). The proposed framework is based on any arbitrary estimator in the class  $\mathcal{B}_{\text{GMD}}$ , given in (6). The GMDI procedure and its theoretical properties are presented in Section 3.1. In Section 3.2, we provide additional discussions on key assumptions made for GMDI. Section 3.3 introduces methods to assess the informativeness of the pre-specified  $\mathbf{H}$  and  $\mathbf{Q}$  to avoid violations of the assumptions that may impact type-I error and power. Section 3.4 proposes a robust GMDI procedure to incorporate partially informative structures for controlling type-I error rates and guaranteeing power.

Recall from (1) that  $\mathbf{y} = \mathbf{X}\beta^* + \epsilon$ , where  $\mathbb{E}[\epsilon | \mathbf{X}] = \mathbf{0}_n$  and  $\text{Cov}(\epsilon | \mathbf{X}) = \Psi$ . Letting  $\Psi = \mathbf{L}_\psi^\top \mathbf{L}_\psi$  and  $\epsilon = \mathbf{L}_\psi^\top \tilde{\epsilon}$  with  $\tilde{\epsilon} = (\tilde{\epsilon}_1, \dots, \tilde{\epsilon}_n)^\top$ , we assume that  $\tilde{\epsilon}_1, \dots, \tilde{\epsilon}_n$  are *i.i.d.* sub-Gaussian random variables with mean 0 and variance 1; that is, there exists a constant  $C > 0$  such that

$$(8) \quad \mathbb{E}[\exp(t\tilde{\epsilon}_i)] \leq \exp\left(\frac{Ct^2}{2}\right) \quad \text{for all } t \in \mathbb{R} \text{ and } i = 1, \dots, n.$$

This sub-Gaussianity assumption is only considered for ease of presentation; our results can be easily extended to other distributions with certain tail bounds, such as sub-exponential distributions (Chapter 2, [Wainwright, 2019](#)).

**3.1. The GMDI Procedure.** Let  $\beta^w = (\beta_1^w, \dots, \beta_p^w)^\top$  be an arbitrary estimator from  $\mathcal{B}_{\text{GMD}}$  in (6) with a fixed weight matrix  $\mathbf{W}$ . We first note that  $\beta_j^w$  can be a biased estimator of  $\beta_j^*$ . Letting  $B_j$  denote the bias of  $\beta_j^w$ , one can see that

$$B_j = (\mathbf{Q}\mathbf{V}\mathbf{W}\mathbf{V}^\top\beta^*)_j - \beta_j^* = \sum_{m \neq j} \xi_{jm}^w \beta_m^* + (\xi_{jj}^w - 1)\beta_j^*,$$

where  $\xi_{jm}^w = (\mathbf{Q}\mathbf{V}\mathbf{W}\mathbf{V}^\top)_{(j,m)}$ , for  $j, m = 1, \dots, p$ . Under  $H_{0,j}$ , it holds that for any  $h_j \in \mathbb{R}$ ,  $B_j = B_j(h_j) := \sum_{m \neq j} \xi_{jm}^w \beta_m^* + h_j(\xi_{jj}^w - 1)\beta_j^*$ . To construct a statistic for testing  $H_{0,j}$  based on  $\beta_j^w$ , we correct the bias  $B_j(h_j)$  using a consistent initial estimator of  $\beta^*$ . Denoting by  $\beta^{init} = (\beta_1^{init}, \dots, \beta_p^{init})^\top$  such an initial estimator (to be discussed in detail later in this section), we can estimate  $B_j(h_j)$  by

$$(9) \quad \widehat{B}_j(h_j) = \sum_{m \neq j} \xi_{jm}^w \beta_m^{init} + h_j(\xi_{jj}^w - 1)\beta_j^{init}.$$

Then, our bias-corrected estimator of  $\beta_j^*$  is given by

$$(10) \quad \widehat{\beta}_j^w(h_j) = \beta_j^w - \widehat{B}_j(h_j), \quad j = 1, \dots, p.$$

Our bias-correction procedure is motivated by the ridge test proposed in [Bühlmann \(2013\)](#) and the grace test proposed in [Zhao and Shojaie \(2016\)](#). Note that this is different from the widely used de-biased Lasso ([Zhang and Zhang, 2014](#); [van de Geer et al., 2014](#)), where the key step is to construct a projection direction that satisfies some “orthogonality property”. However, in the high-dimensional setting, such a projection direction may not exist for highly correlated variables, which is common for two-way structured data. Our bias-correction procedure overcomes this issue since it only requires a consistent initial estimator of  $\beta^*$ . This comes with the cost of not having an optimal test, which we discuss in detail in the remark below [Theorem 3.4](#).

**REMARK.** The two most intuitive choices of  $h_j$  are 0 and 1, which are, respectively, considered in [Bühlmann \(2013\)](#) and [Zhao and Shojaie \(2016\)](#). By considering  $h_j = 0$ , one only corrects the bias under the null hypothesis, while  $h_j = 1$  corrects the general bias regardless of  $\beta_j^*$ . While other choices of  $h_j$  are mathematically valid, they are practically less meaningful. Thus, we shall limit the following discussion to consider  $h_j = 0$  or 1.

Recall that for model (1),  $\text{Cov}(\epsilon \mid \mathbf{X}) = \Psi = \mathbf{L}_\psi^\top \mathbf{L}_\psi$ ,  $\epsilon = \mathbf{L}_\psi^\top \tilde{\epsilon}$ , and  $\tilde{\epsilon} = (\tilde{\epsilon}_1, \dots, \tilde{\epsilon}_n)^\top$ , where  $\tilde{\epsilon}_1, \dots, \tilde{\epsilon}_n$  are *i.i.d.* sub-Gaussian random variables with mean 0 and variance 1. The following result characterizes the asymptotic distribution of  $\widehat{\beta}_j^w(h_j)$  as  $n \rightarrow \infty$ .

**PROPOSITION 3.1.** *For  $j = 1, \dots, p$ , consider the bias-corrected estimator  $\widehat{\beta}_j^w(h_j)$  with any fixed weight matrix  $\mathbf{W}$ , given in (10). Letting  $\mathbf{A} = \mathbf{Q}\mathbf{V}\mathbf{W}\mathbf{S}^{-1}\mathbf{U}^\top\mathbf{H}\mathbf{L}_\psi^\top = (a_{ji})_{j=1, \dots, p \text{ and } i=1, \dots, n}$ , if*

$$(11) \quad \lim_{n \rightarrow \infty} \frac{\max_{i=1, \dots, n} |a_{ji}|}{\sqrt{\sum_{i=1}^n a_{ji}^2}} = 0,$$

then for  $h_j \in \{0, 1\}$ ,

$$(12) \quad \widehat{\beta}_j^w(h_j) = ((1 - h_j)\xi_{jj}^w + h_j)\beta_j^* + \sum_{m \neq j} \xi_{jm}^w(\beta_m^* - \beta_m^{init}) + h_j(\xi_{jj}^w - 1)(\beta_j^* - \beta_j^{init}) + z_j^w.$$

Here,  $z_j^w = \sum_{i=1}^n a_{ji}\tilde{\epsilon}_i$  and  $(\Omega_{jj}^w)^{-1/2}z_j^w \xrightarrow{d} N(0, 1)$  as  $n \rightarrow \infty$ , where  $\Omega_{jj}^w = (\mathbf{A}\mathbf{A}^\top)_{(j,j)}$ .

With a consistent initial estimator  $\beta^{init}$  that will be discussed later, [Proposition 3.1](#) suggests using  $\left| \widehat{\beta}_j^w(h_j) \right|$  as an asymptotically valid test statistic for testing  $H_{0,j}$ . However, its asymptotic variance  $\Omega_{jj}^w$  involves the unknown quantity  $\mathbf{L}_\psi$ , which is not estimable in high-dimensional settings. The GMDI overcomes this difficulty by leveraging the relationship between the auxiliary information  $\mathbf{H}$  and  $\mathbf{L}_\psi$ . More specifically, we assume

$$(A1) \quad \text{As } n \rightarrow \infty, \text{ there exists } \sigma^2 > 0 \text{ such that } \|\mathbf{L}_\psi\mathbf{H}\mathbf{L}_\psi^\top - \sigma^2\mathbf{I}_n\|_2 = o(1).$$

An alternative assumption is that  $\Psi = \mathbf{H}^{-1}$ , which, however, is stringent in practice because it requires  $\mathbf{H}$  to fully capture the unknown covariance  $\Psi$ . Our Assumption (A1) is thus more flexible because it only requires  $\mathbf{H}^{-1}$  to be close to  $\Psi$  in terms of the spectral norm up to a scale transformation. Here, we assume  $\mathbf{H}$  directly informs  $\Psi^{-1}$ , not  $\Psi$ ; that is,  $\mathbf{H}$  informs

the *conditional* similarities between samples. It is well-known that such conditional similarities can be characterized by partial correlations, which are closely related to the inverse covariance matrix. In the following discussions, we will first develop the GMDI procedure by assuming  $\sigma^2$  is known and then discuss procedures for estimating  $\sigma^2$ .

The next proposition states that if Assumption (A1) holds, then  $z_j^w$  (see Proposition 3.1) converges in distribution to  $N(0, R_{jj}^w)$  as  $n \rightarrow \infty$ , where  $R_{jj}^w = \sigma^2 \{\mathbf{Q}\mathbf{V}\mathbf{W}^2\mathbf{S}^{-2}\mathbf{V}^\top\mathbf{Q}\}_{(j,j)}$ .

**PROPOSITION 3.2.** *Consider the  $z_j^w$  defined in Proposition 3.1. Suppose that Assumption (A1) and condition (11) hold. Then, we have  $(R_{jj}^w)^{-1/2}z_j^w \xrightarrow{d} N(0, 1)$  as  $n \rightarrow \infty$ .*

The proofs of Propositions 3.1 and 3.2 are given in Section 2 of the supplement. Next, we elaborate on how to obtain a consistent estimator  $\beta^{init}$ . Existing HDI tools that also perform bias correction use the lasso estimator (Tibshirani, 1996) as the initial estimator (Bühlmann, 2013; Zhao and Shojaie, 2016). Consistency of the lasso estimator requires that (i) the true regression coefficient vector is sparse, and (ii) the design matrix satisfies a restricted eigenvalue-type condition (van de Geer et al., 2009). However, for two-way structured regression, due to potential strong correlations among variables, the true coefficients may not be sparse, and any restricted eigenvalue-type condition may fail; see van de Geer et al. (2009) for more discussions.

As an alternative to those assumptions, we assume that  $\beta^*$  is informed by the eigenvectors of  $\mathbf{Q}$ . Roughly speaking, we assume that the majority of the signals in  $\beta^*$  can be captured by a few eigenvectors of  $\mathbf{Q}$ . More specifically, denoting by  $\mathbf{Q} = \mathbf{D}\mathbf{\Delta}\mathbf{D}^\top$  the eigen-decomposition of  $\mathbf{Q}$  and  $\tilde{\beta}^* = \mathbf{D}^\top\beta^*$ , we assume

(A2) For some  $S_0 \subset \{1, \dots, p\}$  with  $s_0 = |S_0|$ ,  $\left\| \tilde{\beta}_{S_0^c}^* \right\|_1 \leq \eta_1$ , where  $S_0^c$  is the complement of  $S_0$ ,  $\eta_1 = O\left(\sqrt{n^{-1}s_0 \log p}\right)$  and  $s_0 = o\left\{(n/\log p)^r\right\}$  for some  $r \in (0, 1/2)$  as  $n \rightarrow \infty$ .

Under Assumption (A2),  $\|\beta^*\|_{\mathbf{Q}^{-1}}$ , the penalty term of KPR in (7) is likely to be small. Thus, Assumption (A2) is in fact aligned with the key idea of KPR. Indeed, in Section 3.2, we will show that any estimator from the class  $\mathcal{B}_{\text{GMD}}$  is less biased if  $\beta^*$  satisfies Assumption (A2).

Our third assumption characterizes how  $\mathbf{H}$  and  $\mathbf{Q}$ , respectively, inform the row and column structures of the design matrix  $\mathbf{X}$ . As mentioned earlier, any restricted eigenvalue-type condition may break down due to potentially strong correlations in  $\mathbf{X}$ . We assume that  $\mathbf{H}$  and  $\mathbf{Q}$  can help decorrelate the rows and columns of  $\mathbf{X}$ , respectively, so that the decorrelated design matrix satisfies a restricted eigenvalue-type condition. More specifically, we assume

(A3) For some constants  $0 < c_* < c^* < \infty$ ,

$$c_* \leq \frac{\|\tilde{\mathbf{X}}_A \mathbf{v}\|^2}{n\|\mathbf{v}\|^2} \leq c^* \quad \text{for any } A \subset \{1, \dots, p\} \text{ with } |A| = q^* \text{ and } \mathbf{v} \in \mathbb{R}^*,$$

where  $\tilde{\mathbf{X}} = \mathbf{H}^{1/2}\mathbf{X}\mathbf{D}\mathbf{\Delta}^{1/2}$ ,  $q^* \geq M_1^*s_0 + 1$  with  $s_0$  given in Assumption (A2) and  $M_1^*$  specified in Section 3 of the supplement.

Letting  $\tilde{\Sigma}_A = n^{-1}\tilde{\mathbf{X}}_A^\top\tilde{\mathbf{X}}_A$ , Assumption (A3) implies that all eigenvalues of  $\tilde{\Sigma}_A$  are inside the interval  $[c_*, c^*]$  when the size of  $A$  is no greater than  $q^*$ . This assumption is called the sparse Riesz condition (Zhang et al., 2008). According to Proposition 1 in Zhang et al. (2008), if there exists some  $q^*$  such that the maximum correlation between the variables in  $\tilde{\mathbf{X}}$  is bounded by  $\delta/(q^* - 1)$  for some  $\delta < 1$ , then this condition holds with rank  $q^*$ ,  $c_* = 1 - \delta$  and  $c^* = 1 + \delta$ .

Under assumptions (A1)–(A3), we introduce the following three-step procedure to construct the bias-corrected estimator of  $\beta_j^*$  based on an arbitrary estimator  $\beta_j^w$  from  $\mathcal{B}_{\text{GMD}}$  for  $j = 1, \dots, p$ .

**GMDI bias-correction procedure:**

(B1) For a fixed tuning parameter  $\lambda$ , find

$$(13) \quad \tilde{\beta}(\lambda) = \operatorname{argmin}_{\beta} \left\{ \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{D}\beta\|_{\mathbf{H}}^2 + \lambda \|\Delta^{-1/2}\beta\|_1 \right\}.$$

(B2) Calculate  $\beta^{\text{init}} = \mathbf{D}\tilde{\beta}(\lambda)$ .

(B3) For a fixed  $h_j \in \{0, 1\}$ , let  $\hat{\beta}_j^w(h_j) = \beta_j^w - \hat{B}_j(h_j)$  with  $\hat{B}_j(h_j)$  defined in (9).

We use a weighted  $l_1$  penalty in (13) with the weights equal to the inverse of the square root of the eigenvalues of  $\mathbf{Q}$ . We will explain the rationale behind this weight choice in Section 3.2. Also, consistency of  $\beta^{\text{init}}$  requires certain conditions on  $\lambda$ , which will be specified in Theorem 3.3.

Letting  $\zeta_j(h_j) = \sum_{m \neq j} \xi_{jm}^w (\beta_m^* - \beta_m^{\text{init}}) + h_j (\xi_{jj}^w - 1) (\beta_j^* - \beta_j^{\text{init}})$  for  $j = 1, \dots, p$  and  $\Xi = \operatorname{diag}(\xi_{11}^w, \dots, \xi_{pp}^w)$ , the following result serves as the basis for constructing an asymptotically valid test for  $H_{0,j}$  using the bias-corrected estimator  $\hat{\beta}_j^w(h_j)$  given in (12). In the following theorems, without loss of generality, we assume that  $\mathbf{Q}$  is appropriately scaled such that  $\|\mathbf{Q}\|_2 = 1$ .

**THEOREM 3.3.** *Suppose the columns of  $\mathbf{X}$  are standardized such that  $\|\mathbf{X}\mathbf{d}_j\|_{\mathbf{H}}^2 = n$ , where  $\mathbf{d}_j$  is the  $j$ -th column of  $\mathbf{D}$ , for  $j = 1, \dots, p$ . For  $\tilde{\beta}(\lambda)$  in (13), consider  $\lambda = 2\sqrt{2c^*n \log p(1 + c_0)} \|\mathbf{L}_{\psi}\mathbf{H}\mathbf{L}_{\psi}^{\top}\|_2$  with any  $c_0 > 0$ , where  $c^*$  is given in Assumption (A3). For  $h_j \in \{0, 1\}$ , denote*

$$(14) \quad \Psi_j(h_j) = \left\| [(\mathbf{Q}\mathbf{V}\mathbf{W}\mathbf{V}^{\top} - (1 - h_j)\Xi - h_j\mathbf{I}_p)\mathbf{D}]_{(j,\cdot)} \right\|_{\infty} \left( \frac{\log p}{n} \right)^{1/2-r},$$

where for any matrix  $\mathbf{M}$ ,  $\mathbf{M}_{(j,\cdot)}$  denotes the  $j$ -th row of  $\mathbf{M}$ . Then, under condition (11) and Assumptions (A1)–(A3), we have  $\lim_{n \rightarrow \infty} \Pr(|\zeta_j(h_j)| \leq \Psi_j(h_j)) = 1$ . Furthermore, under  $H_{0,j}$ , for any  $\alpha > 0$ ,

$$(15) \quad \limsup_{n \rightarrow \infty} \Pr\left(|\hat{\beta}_j^w(h_j)| > \alpha\right) \leq \limsup_{n \rightarrow \infty} \Pr\left(|Z_j^w| + \Psi_j(h_j) > \alpha\right),$$

where  $Z_j^w$  is given in Proposition 3.1.

Combining Theorem 3.3 with Proposition 3.2, we can test  $H_{0,j}$  using the asymptotically valid two-sided  $p$ -value

$$(16) \quad P_j^w(h_j) = 2 \left\{ 1 - \Phi \left( (R_{jj}^w)^{-1/2} \left\{ |\hat{\beta}_j^w(h_j)| - \Psi_j(h_j) \right\}_+ \right) \right\},$$

where  $\Phi(\cdot)$  is the cumulative distribution function of the standard normal distribution and  $a_+ = \max(a, 0)$ . Calculating  $P_j^w(h_j)$  requires obtaining a consistent estimator of  $\sigma^2$ . In this paper, we use the organic lasso (Yu and Bien, 2019) to estimate  $\sigma^2$  by regressing  $\mathbf{H}^{1/2}\mathbf{y}$  against  $\tilde{\mathbf{X}}$  with  $\tilde{\mathbf{X}}$  defined in Assumption (A3), but other approaches, such as the scaled lasso (Sun and Zhang, 2012), may also be used.

Our next result guarantees the power of GMDI when the size of the true regression coefficient is sufficiently large.

**THEOREM 3.4.** *Assume the conditions in Theorem 3.3 hold. For  $h_j \in \{0, 1\}$ , if there exists some  $0 < \alpha < 1$  and  $0 < \psi < 1$  such that*

$$(17) \quad |\beta_j^*| \geq |(1 - h_j)\xi_{jj}^w + h_j|^{-1} \left( 2\Psi_j(h_j) + (q_{(1-\alpha/2)} + q_{(1-\psi/2)}) \sqrt{R_{jj}^w} \right),$$

where  $\Phi(q_t) = t$  for any  $t \in (0, 1)$  and  $\Psi_j(h_j)$  is defined in (14), then  $\lim_{n \rightarrow \infty} \Pr \left( P_j^w(h_j) \leq \alpha \right) \geq \psi$ .

It should be noted that condition (17) does not hold when  $h_j = 0$  and  $\xi_{jj}^w = 0$ ; however, this rarely happens and can be easily checked in advance. In cases where (17) is not true,  $h_j = 1$  can be used. Proofs of Theorems 3.3 and 3.4 are provided in Sections 3 and 4 of the supplement, respectively.

**REMARK.** Similar to the ridge test and the Grace test, (15) implies that GMDI may be conservative. Also, theoretical guarantees of GMDI require using a fixed weight matrix  $\mathbf{W}$ , but in practice, to achieve the optimal prediction performance,  $\mathbf{W}$  is chosen via cross-validation (e.g., the proposed VI-based approach in Section 2). When samples are *i.i.d.*, one could address this issue by splitting the data into two parts, and then use one part to select  $\mathbf{W}$  and the other part to perform inference. However, this data-splitting procedure becomes non-trivial, if not impossible, for two-way structured data. An alternative way is to select top GMD components for GMDR and a fixed tuning parameter for KPR. In these cases,  $\mathbf{W}$  becomes deterministic, but the prediction/estimation accuracy of GMDR/KPR may be compromised. Nonetheless, despite these two potential limitations, we show in Section 4, through extensive simulation studies, that the GMDI is more powerful than existing HDI methods with well-controlled type-I error rates.

**3.2. On GMDI Assumptions.** In this section, we discuss Assumptions (A2) and the weighted  $l_1$  penalty used in (B1) from the perspective of the bias of any arbitrary estimator in  $\mathcal{B}_{\text{GMD}}$ . Recall that the bias of  $\beta^w = \mathbf{Q}\mathbf{V}\mathbf{W}\mathbf{S}^{-1}\mathbf{U}^\top\mathbf{H}\mathbf{y}$  is given by  $\text{Bias}(\beta^w) = \mathbb{E}(\beta^w) - \beta^* = \mathbf{Q}\mathbf{V}\mathbf{W}\mathbf{V}^\top\beta^* - \beta^*$ , which can be rewritten as

$$(18) \quad \text{Bias}(\beta^w) = \mathbf{Q}(\mathbf{V}\mathbf{W}\mathbf{V}^\top\mathbf{Q} - \mathbf{I}_p)\mathbf{Q}^{-1}\beta^*.$$

Recalling  $K = \text{rank}(\mathbf{X}\mathbf{Q}\mathbf{X}^\top\mathbf{H})$ , we make the following observations from (18).

- (O1) Suppose  $K = p$ . Let  $\beta^w$  be the GMDR estimator with all GMD components selected. In this case,  $\mathbf{W} = \mathbf{I}_p$  and it can be seen that  $\mathbf{V}\mathbf{W}\mathbf{V}^\top\mathbf{Q} = \mathbf{I}_p$ . Thus,  $\text{Bias}(\beta^w) = 0$ . This demonstrates that in the low-dimensional case ( $K = p \leq n$ ), the GMDR estimator based on all GMD components is an unbiased estimator of  $\beta^*$  for any  $\beta^* \in \mathbb{R}^p$ .
- (O2) Suppose  $K < p$ , a common scenario in high-dimensional settings ( $n < p$ ). In this case, it can be seen that  $\mathbf{V}\mathbf{W}\mathbf{V}^\top\mathbf{Q} \neq \mathbf{I}_p$  for any weight matrix  $\mathbf{W}$ . Then, using (18), we have  $\|\text{Bias}(\beta^w)\|_2 \leq \|\mathbf{Q}\|_2 \|\mathbf{V}\mathbf{W}\mathbf{V}^\top\mathbf{Q} - \mathbf{I}_p\|_2 \|\mathbf{Q}^{-1}\beta^*\|_2$ , indicating that  $\beta^w$  is less biased if  $\|\mathbf{Q}^{-1}\beta^*\|_2$  is small. Since  $\mathbf{Q} = \mathbf{D}\mathbf{\Delta}\mathbf{D}^\top$ , it can be seen that

$$(19) \quad \|\mathbf{Q}^{-1}\beta^*\|_2^2 = \sum_{j=1}^p \delta_j^{-2} \left( \mathbf{d}_j^\top \beta^* \right)^2,$$

where  $\mathbf{d}_j$  is the  $j$ -th column of  $\mathbf{D}$ , i.e., the  $j$ -th eigenvector of  $\mathbf{Q}$ . Since  $\delta_1 \geq \dots \geq \delta_p > 0$ , (19) implies that  $\beta^w$  is less biased if (a) only a few  $|\mathbf{d}_j^\top \beta^*|$  are non-zero, or (b) for large  $j$  (small  $\delta_j$ ),  $\mathbf{d}_j^\top \beta^* = 0$ . Thus, Assumption (A2) aligns well with (a) because it indicates that the majority of the signals in  $\beta^*$  lie in the space spanned by a few eigenvectors of  $\mathbf{Q}$ . The weighted  $l_1$  penalty in (B1) encourages  $\mathbf{d}_j^\top \beta^*$  to be 0 for large  $j$  and thus aligns with (b). Note that (b) also aligns with the heuristic of KPR, where  $\beta^*$  is assumed to be informed by the top eigenvectors of  $\mathbf{Q}$ .

3.3. *Tests for informative  $\mathbf{H}$  and  $\mathbf{Q}$ .* Informative  $\mathbf{H}$  and  $\mathbf{Q}$  required by the proposed GMDR and GMDI can be obtained from auxiliary data sources, which are common in omics studies. For example, the row and column structures used to construct Fig. 1C are estimated from the phylogenetic tree and the metagenomics data, respectively. However, in practice, one may get uninformative  $\mathbf{H}$  and/or  $\mathbf{Q}$ , which may impact the type-I error and power.

To avoid uninformative external structures, we propose to use the kernel RV coefficient (KRV, Zhan et al., 2017) to examine the informativeness of  $\mathbf{Q}$  with respect to the column structure of  $\mathbf{X}$ . Specifically, we define  $\mathbf{Q}_x = \mathbf{X}^\top \mathbf{X}$  to measure the Euclidean similarities between variables. Since  $\mathbf{Q}$  is assumed to characterize conditional similarities, we test the association between  $\mathbf{Q}^{-1}$  and  $\mathbf{Q}_x$  using

$$\text{KRV}(\mathbf{Q}_x, \mathbf{Q}) = \frac{\text{tr}(\tilde{\mathbf{Q}}_x \tilde{\mathbf{Q}})}{\sqrt{\text{tr}(\tilde{\mathbf{Q}}_x^2) \text{tr}(\tilde{\mathbf{Q}}^2)}},$$

where  $\tilde{\mathbf{Q}}_x = \mathbf{C}_p \mathbf{Q}_x \mathbf{C}_p$  and  $\tilde{\mathbf{Q}} = \mathbf{C}_p \mathbf{Q}^{-1} \mathbf{C}_p$  with

$$(20) \quad \mathbf{C}_p = \mathbf{I}_p - p^{-1} \mathbf{1}_p \mathbf{1}_p^\top.$$

A permutation test with a fast approximation of the permutation null distribution is used to test whether the true KRV is 0 (Zhan et al., 2017). If the permutation  $p$ -value is less than a pre-selected significance level, say 0.05, then we consider  $\mathbf{Q}$  an informative column structure.

Similarly, defining  $\mathbf{H}_x = \mathbf{X} \mathbf{X}^\top$ , one can calculate  $\text{KRV}(\mathbf{H}_x, \mathbf{H})$  with a permutation-based  $p$ -value. As  $\mathbf{H}^{-1}$  captures sample-wise similarities, we also examine the association between  $\mathbf{H}^{-1}$  and the outcome  $\mathbf{y}$  using the microbiome regression-based association tests (MiRKAT, Zhao et al., 2015). MiRKAT is not performed for  $\mathbf{Q}$  because the dimension of  $\mathbf{Q}$  is incompatible with that of  $\mathbf{y}$ . MiRKAT is built upon a mixed-effect model, where the microbiome abundances are modeled as random effects with the covariance matrix  $\tau \mathbf{H}^{-1}$  for some  $\tau \geq 0$ . Thus, the statistical significance of the MiRKAT test (i.e.,  $\tau > 0$ ) rejects the hypothesis that the sample-wise covariance is substantially distinct from (a constant multiple of)  $\mathbf{H}^{-1}$ . Hence, this test is in the spirit of our Assumption (A1). If both the KRV and MiRKAT tests are statistically significant, we consider  $\mathbf{H}$  an informative row structure.

In Section 4, we will also demonstrate the effectiveness of the KRV and MiRKAT tests in terms of excluding uninformative row and column structures.

3.4. *Robust GMDI with partially informative structures.* While KRV and MiRKAT can help avoid uninformative row structures, they may identify partially informative structures that do not guarantee valid inference results. To address this issue, we propose a robust procedure to determine how much information from the external structures should be incorporated. The main idea is to find a linear combination of a partially informative structure and the identity matrix through an optimal weighting scheme. More specifically, consider model (1) with a partially informative structure  $\mathbf{H}$  and a fully informative structure  $\mathbf{Q}$ . Without loss of generality, we assume  $\|\mathbf{H}\|_2 = 1$ . In this case, we define a weighted structure  $\mathbf{H}(\tau) = \tau \mathbf{H} + (1 - \tau) \mathbf{I}_n$  with  $\tau \in (0, 1)$ . Motivated by the connection among  $l_2$ -penalized regression, dimension reduction-based regression, and linear mixed models (LMM) (Liu, Lin and Ghosh, 2007; Zhang and Pan, 2015; Randolph et al., 2018), we find the optimal value of  $\tau$  by considering the following LMM:

$$(21) \quad \mathbf{y} = \mathbf{X} \boldsymbol{\beta}^* + \boldsymbol{\epsilon}, \text{ with } \boldsymbol{\beta}^* \sim N_p(\mathbf{0}, c_Q \mathbf{Q}), \boldsymbol{\epsilon} \sim N_n(\mathbf{0}, c_H \mathbf{H}(\tau)^{-1})$$

for some  $c_H, c_Q > 0$ . Letting  $\Omega(\tau) = c_Q \mathbf{X} \mathbf{Q} \mathbf{X}^\top + c_H \mathbf{H}(\tau)^{-1}$ , one can see that  $\mathbf{y} \sim N_n(\mathbf{0}, \Omega(\tau))$ , leading to the following likelihood function:

$$l_n(c_H, c_Q, \tau) = \frac{\exp(-1/2 \mathbf{y}^\top \Omega(\tau)^{-1} \mathbf{y})}{\sqrt{(2\pi)^n |\Omega(\tau)|}}.$$

Since  $c_H$  and  $c_Q$  are identifiable only up to a scale transformation, we reparametrize the likelihood by defining  $\lambda_{HQ} = c_H/c_Q$  and  $l_n(\lambda_{HQ}, \tau) = l_n(c_H, c_Q, \tau)$ . Then, the maximum likelihood estimate (MLE) of  $\lambda_{HQ}$  and  $\tau$  is

$$\{\hat{\lambda}_{HQ}, \hat{\tau}\} = \operatorname{argmin} \{ \mathbf{y}^\top \Omega(\tau)^{-1} \mathbf{y} + \log |\Omega(\tau)| \} \text{ subject to } \lambda_{HQ} > 0, 1 > \tau > 0.$$

We use an augmented Lagrangian method to solve the optimization problem. Having found  $\hat{\tau}$ , one can implement GMDI with  $\mathbf{H}(\hat{\tau})$  and  $\mathbf{Q}$ , referred to as the robust GMDI procedure (r-GMDI) hereafter. We will demonstrate the effectiveness of r-GMDI using simulations and real data applications.

**4. Simulation Studies.** We conducted two simulation studies, each containing multiple settings, to compare the proposed GMDI with five existing high-dimensional inferential procedures: (i) the low-dimensional projection estimator (LDPE, [Zhang and Zhang, 2014](#)); (ii) the ridge-based high-dimensional inference (Ridge, [Bühlmann, 2013](#)); (iii) the de-correlated score test (dscore, [Ning and Liu, 2017](#)); (iv) inference for the graph-constrained estimator (Grace, [Zhao and Shojaie, 2016](#)) and (v) the non-sparse high-dimensional inference (ns-hdi, [Zhu and Bradic, 2018](#)). In the first study, we performed data-driven simulations based on a real microbiome data set. In the second study, we simulated two-way structured data using a matrix variate normal distribution ([Gupta and Nagar, 2018](#)) with pre-specified row and column covariance matrices. We used a two-sided significance level  $\alpha = 0.05$  for all tests.

As GMDI works for the entire family of estimators  $\mathcal{B}_{\text{GMD}}$ , we considered two specific estimators from  $\mathcal{B}_{\text{GMD}}$ : (i) the proposed GMDR estimator in (3) and (ii) the KPR estimator in (7). We denote the resulting tests for the GMDR and KPR estimators by GMDI-d and GMDI-k, respectively, because GMDR exerts *discrete* shrinkage effects on GMD components, whereas KPR exerts continuous shrinkage effects through a *kernel* function. For the selection of the index set  $\mathcal{I}$  of the GMDR estimator  $\hat{\beta}_{\text{GMDR}}(\mathcal{I})$ , GMD components that explain less than 0.1% of the total variance are excluded because the estimated coefficients corresponding to those components with low variances may be unstable. To see this, recall from (4) that  $\hat{\gamma} = \operatorname{argmin}_{\gamma} \|\mathbf{y} - \mathbf{Y}\gamma\|_{\mathbf{H}}^2$ . Then,  $\hat{\gamma}_l = \sigma_l^{-1} \mathbf{u}_l^\top \mathbf{H} \mathbf{y}$  and  $\operatorname{Var}(\hat{\gamma}_l) = \sigma_\epsilon^2 \sigma_l^{-2}$ , for  $l = 1, \dots, K$ . This indicates that when the total  $R^2$  is low ( $\sigma_\epsilon^2$  is relatively large), for large  $l$  (small  $\sigma_l$ ),  $\hat{\gamma}_l$  may be unstable due to its large variance. The index set  $\mathcal{I}$  is then selected by the proposed GCV procedure based on the remaining GMD components. For the KPR estimator  $\hat{\beta}_{\text{KPR}}(\eta)$ , the tuning parameter  $\eta$  is selected by 10-fold cross validation. For GMDI, the bias-correction parameter  $h_j$  (see Proposition 3.1) is set to be 1 for all  $j$ , as done for Grace; the tuning parameter  $\lambda$  in (13) is set to be  $2\sqrt{3n \log p}$ , and the sparsity parameter  $r$  is set to be 0.05. For LDPE and Ridge, we used the implementation in the R package `hdi`, and for the Grace test, we used the implementation in the R package `Grace`. For LDPE, Ridge, and Grace, the tuning parameters are selected using 10-fold cross-validation.

4.1. *Simulation 1.* In this study, we performed *data-driven simulations* using data collected as part of the ‘‘Carbohydrates and Related Biomarkers’’ (CARB) study, conducted between June 2006 and July 2009 at the Fred Hutchinson Cancer Center. CARB was a randomized, controlled, crossover feeding study aimed at evaluating the effects of glycemic load on a variety of biomarkers, such as systemic inflammation, insulin resistance, and adipokines ([Neuhouser et al., 2012](#)). Participants were randomized based on body mass index and sex, and fed two controlled diets (randomly assigned order) for 28 days, with a 28-day washout period between diets. The 16S rRNA genus abundance data used here are from 58 participants sampled at each of the three time points, resulting in 174 observations. To classify bacterial taxonomy, sequences were processed using QIIME ([Caporaso et al., 2010](#)). This processing produced a complete phylogenetic tree with 1054 leaves corresponding to level-7

taxa (species) defined by 97% similarity and 151 genera (level 6 of the tree). Our simulation used 114 genera after filtering out those that did not appear in at least 30% of the 174 samples. We correspondingly trimmed the tree back to the genus level with 114 leaves.

Let  $\mathbf{X} \in \mathbb{R}^{174 \times 114}$  be the sample-by-taxon matrix with entries being taxon counts. Let  $g(\mathbf{z}) = \left(\prod_{k=1}^p z_k\right)^{1/p}$  denote the geometric mean of  $\mathbf{z} = (z_1, \dots, z_p)^\top$ . The centered log-ratio (CLR) transformation of  $\mathbf{z}$  is defined as

$$(22) \quad \text{clr}(\mathbf{z}) = \left[ \log \frac{z_1}{g(\mathbf{z})}, \dots, \log \frac{z_p}{g(\mathbf{z})} \right].$$

Since the CLR transformation is not well defined when  $\mathbf{z}$  contains zero entries, we added a pseudo count of 1 to all entries in  $\mathbf{X}$  and then constructed the CLR transformed data matrix  $\tilde{\mathbf{X}}$  by applying the CLR transformation (22) to each row of  $\mathbf{X}$ . The auxiliary row structure was derived from the weighted UniFrac distance between observations (Lozupone and Knight, 2005). Specifically, letting  $\Delta_{\mathbf{U}} \in \mathbb{R}^{n \times n}$  be the squared weighted UniFrac distance matrix, we obtained  $\mathbf{H} = \left(-\frac{1}{2}\mathbb{C}_n \Delta_{\mathbf{U}} \mathbb{C}_n\right)^{-1}$ , where the centering matrix  $\mathbb{C}_n$  is defined in (20) in Section 3.3. The column structure  $\mathbf{Q} = \left(-\frac{1}{2}\mathbb{C}_p \Delta_{\mathbf{P}} \mathbb{C}_p\right)^{-1}$ , where  $\Delta_{\mathbf{P}}$  is the squared patristic distance between taxa obtained from the phylogenetic tree. The KRV test yields a zero  $p$ -value for  $\text{KRV}(\mathbf{H}_x, \mathbf{H})$  and a  $p$ -value of 0.025 for  $\text{KRV}(\mathbf{Q}_x, \mathbf{Q})$ , indicating that  $\mathbf{H}$  and  $\mathbf{Q}$  are informative for the row and column structures of  $\mathbf{X}$ , respectively.

Letting  $\mathbf{d}_j$  denote the  $j$ -th eigenvector of  $\mathbf{Q}$ , we set  $\beta_0 = 5 \sum_{j=1}^{10} \{2 + 3(j-1)\}^{-1/2} \mathbf{d}_j$ . We then defined the true signal  $\beta^*$  as a thresholded version of  $\beta_0$ :  $\beta^* = s(\beta_0, 0.1)$ , where  $s(x, \tau)$  is the hard-thresholding operator; i.e.,  $s(x, \tau) = x \mathbb{1}(|x| > \tau)$ , and the threshold  $\tau = 0.1$  was selected so that 81 entries of  $\beta^*$  are non-zeros. The reason why we considered this thresholded parameter as our true parameter is two-fold. First,  $\beta^*$  has both zero and non-zero entries, allowing us to evaluate the type-I error rate from testing the zero coefficients and the power from testing the non-zero coefficients. In comparison, all entries of  $\beta_0$  are non-zero due to the structure of  $\mathbf{Q}$ . Second, the thresholded parameter  $\beta^*$  is no longer fully informed by the top eigenvectors of  $\mathbf{Q}$ , which is more realistic in practice.

Let  $\mathbf{H} = \sum_{j=1}^n \lambda_{j,H} \mathbf{d}_{j,H} \mathbf{d}_{j,H}^\top$  denote the eigen-decomposition of  $\mathbf{H}$ , where  $\lambda_{1,H} \geq \lambda_{2,H} \geq \dots \geq \lambda_{n,H} > 0$  are the eigenvalues, and  $\mathbf{d}_{1,H}, \dots, \mathbf{d}_{n,H}$  are the corresponding eigenvectors. Defining  $\Psi = \sum_{j=1}^n \left(\lambda_{j,H}^{-1} + \delta \lambda_{1,H}^{-1}\right) \mathbf{d}_{j,H} \mathbf{d}_{j,H}^\top$ , we generated  $\epsilon$  from a multivariate normal distribution with mean  $\mathbf{0}$  and covariance  $\Psi$ , and simulated the response  $\mathbf{y} = \tilde{\mathbf{X}}\beta^* + \epsilon$ . In this case, we can calculate  $\|\mathbf{L}_\psi \mathbf{H} \mathbf{L}_\psi^\top - \mathbf{I}_n\| = \delta$ , where  $\Psi = \mathbf{L}_\psi^\top \mathbf{L}_\psi$ . Thus, according to Assumption (A1), a smaller  $\delta$  indicates that  $\mathbf{H}$  better informs  $\Psi$ ; in particular,  $\delta = 0$  means that  $\mathbf{H}$  fully informs  $\Psi$ .

We consider four values of  $\delta$ : 0.2, 0.5, 1, and 2. The results are summarized in Fig. 2. All existing HDI methods fail to differentiate between zero and non-zero entries. More specifically, LDPE, Ridge, and Grace have almost no power, while dscore and ns-hdi have highly inflated type-I error rates. This is because none of these methods can handle correlated samples. The proposed GMDI-k and GMDI-d show better performances. Both the GMDI-k and GMDI-d show decent power with roughly controlled type-I error rates.

*4.2. Simulation 2.* We considered four settings in this study. In Settings I and II, we considered data with column structures and examined how different choices of  $\mathbf{Q}$  affect the performance of GMDI and the Grace test. In Setting III, we demonstrated the effectiveness of the KRV and MiRKAT in terms of detecting informative structures. In Setting IV, we demonstrated the effectiveness of the proposed robust GMDI in terms of handling partially informative structures.

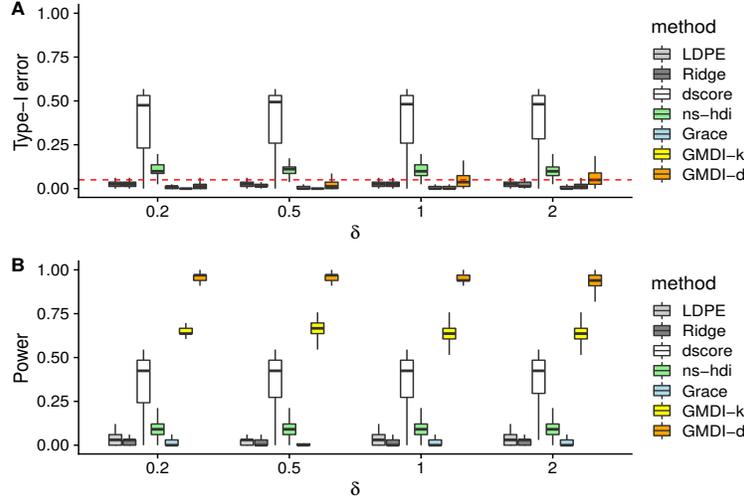


Fig 2: Boxplots of the type-I error (A) and power (B) over 500 replications for Simulation II with  $\delta = 0.2, 0.5, 1,$  and  $2$ : Both GMDI-d and GMDI-k can roughly control the type-I error and have considerably higher power than the existing HDI methods.

**Setting I:** We first simulated  $\mathbf{X} \in \mathbb{R}^{200 \times 300}$  from a matrix variate normal distribution with mean  $\mathbf{0}$ , row covariance  $\mathbf{I}_{200}$  and column covariance  $\Sigma$ , where

$$(\Sigma^{-1})_{(i,j)} = \begin{cases} 1, & i = j \\ 0.9^{|i-j|}, & i \neq j, i \leq 150, j \leq 150 \\ 0.5^{|i-j|}, & i \neq j, i > 150, j > 150 \\ 0, & \text{otherwise.} \end{cases}$$

Letting  $\mathbf{Q} = \Sigma^{-1}$  and denoting by  $\mathbf{f}_j$  the  $j$ -th eigenvector of  $\mathbf{Q}$ , for  $j = 1, \dots, 300$ , we defined  $\beta^* = \sum_{j=1}^{10} j^{-1/2} \mathbf{f}_j$ , which aligns with the top 10 eigenvectors of  $\mathbf{Q}$ . The response  $\mathbf{y}$  was generated according to  $\mathbf{y} = \mathbf{X}\beta^* + \epsilon$ , where  $\epsilon$  was simulated from a multivariate normal distribution with mean  $\mathbf{0}$  and covariance  $\Psi = \sigma_\epsilon^2 \mathbf{I}_{200}$  with  $\sigma_\epsilon^2$  selected to achieve an  $R^2$  of 0.4, 0.6 or 0.8. Our GMDI was implemented using  $\mathbf{H} = \mathbf{I}_{200}$  and  $\mathbf{Q} = \Sigma$ , and  $\sigma_\epsilon^2$  was estimated using the organic lasso (Yu and Bien, 2019). One can easily check that the pre-specified  $\mathbf{H}$  and  $\mathbf{Q}$  satisfy Assumptions (A1)-(A3). By the block diagonal design of  $\mathbf{Q}$ , we know that the first 150 coefficients of  $\beta^*$  are non-zero, while the rest are zero. This enables us to evaluate the power from testing the non-zero coefficients and the type-I error rate from testing the zero coefficients.

The results are summarized in Fig. 3. Figure 3A shows that all methods except ns-hdi can control the type-I error rate. This is likely because in this setting, the precision matrix of the variables,  $\Sigma^{-1}$ , does not satisfy the row sparsity condition required by ns-hdi. The power comparison in Fig. 3B shows that both GMDI-k and GMDI-d have considerably higher power than the existing methods. More specifically, LDPE, Ridge, and dscore have very low power since they completely ignore the column structure of  $\mathbf{X}$  and  $\beta^*$  is not sparse. Because the Grace estimator can incorporate the column structure (Grace is implemented using  $\mathbf{L} = \Sigma$ ; see Zhao and Shojaie, 2016 for details), the Grace test gains more power than LDPE and Ridge. However, since the Grace test still requires the sparsity of  $\beta^*$ , which is not satisfied in this setting, it is not as powerful as GMDI-d or GMDI-k. These results clearly demonstrate the importance of incorporating informative column structures for gaining more power. As  $R^2$  increases, GMDI-k and GMDI-d both yield more stringent control of the type-I error and

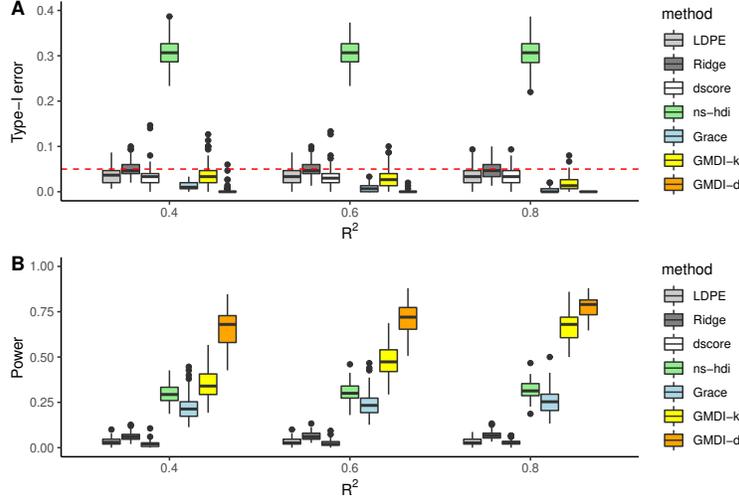


Fig 3: Boxplots of the type-I error (A) and power (B) over 500 replications for Setting I with  $R^2 = 0.4, 0.6$  and  $0.8$ : Both GMDI-d and GMDI-k can control the type-I error, and have considerably higher power than other methods.

TABLE 1  
The mean (sd) of the RMSEs for the methods of VI and TOP over 500 replications.

Method \ $R^2$	0.4	0.6	0.8
VI	0.946 (0.082)	0.895 (0.088)	0.832 (0.091)
TOP	0.967 (0.065)	0.934 (0.093)	0.859 (0.102)

more power at the same time. GMDI-d has higher power than GMDI-k, especially for low  $R^2$  values; this is accompanied by the observation that GMDI-k yields more conservative control of the type-I error rate than GMDI-d. This difference between GMDI-d and GMDI-k may be attributed to the fact that GMDI-k shrinks all components, whereas GMDI-d only selects a subset of components without adding any shrinkage effect.

We also evaluated the prediction performance of GMDR by considering two methods for selecting the GMD components: the proposed VI-based procedure and the classical procedure that selects top GMD components, referred to as VI and TOP, respectively. Specifically, for each  $i = 1, \dots, 200$ , we obtained a prediction of  $y_i$  based on the leave-one-out cross-validation (LOOCV), denoted by  $\hat{y}_i$ . Letting  $\hat{\mathbf{y}} = (\hat{y}_1, \dots, \hat{y}_{200})^\top$ , we calculated the relative mean squared error (RMSE) according to  $\text{RMSE} = \|\mathbf{y} - \hat{\mathbf{y}}\|^2 / \|\mathbf{y}\|^2$ . Table 1 shows the mean and standard deviation (sd) of the RMSEs over 500 replications. As  $R^2$  increases, both methods show better prediction performance. For all values of  $R^2$ , the VI method shows lower average prediction errors than the TOP method with similar standard deviations, demonstrating the effectiveness of the proposed VI method.

**Setting II:** In the previous setting, our GMDI was implemented using correctly specified  $\mathbf{H}$  and  $\mathbf{Q}$ . In practice, the auxiliary structures may be mis-specified. In this simulation, we examined how different choices of  $\mathbf{Q}$  affect the performance of GMDI and the Grace test. The simulation setting is mostly the same as in Setting I, except that instead of using  $\Sigma^{-1}$  as  $\mathbf{Q}$ , we considered two perturbed matrices:  $\mathbf{Q}^{(1)}$  and  $\mathbf{Q}^{(2)}$ . Here,  $\mathbf{Q}^{(1)}$  is defined similar to  $\Sigma^{-1}$ , except that  $\mathbf{Q}_{(i,j)}^{(1)} = 0.1^{|i-j|}$  for all  $(i, j) \in \{(a, b) : (a - 150)(b - 150) < 0\}$ , and  $\mathbf{Q}^{(2)} =$

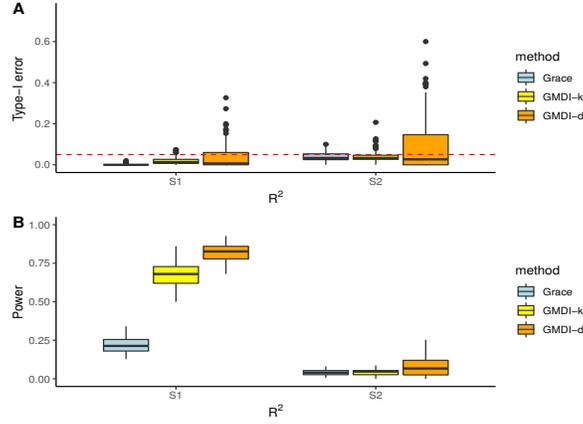


Fig 4: Boxplots of the type-I error (A) and the power (B) over 500 replications for Setting II with  $R^2 = 0.8$ . The S1 and S2 on the x-axis represent  $\mathbf{Q}^{(1)}$  and  $\mathbf{Q}^{(2)}$  respectively: Both GMDI-d and GMDI-k work well under small perturbations of  $\mathbf{Q}$ . With a completely mis-specified  $\mathbf{Q}$ , GMDI-d and GMDI-k have limited power. This mis-specified choice of  $\mathbf{Q}$  can be avoided by the KRV test.

$0.9 \times \mathbf{I}_{300} + 0.1 \times \mathbf{1}_{300}\mathbf{1}_{300}^T$ . Under the significance level 0.05, 492 out of 500 independent realizations of  $\mathbf{X}$  lead to statistically significant results for testing  $\text{KRV}(\mathbf{Q}_x, \mathbf{Q}^{(1)})$ , whereas only five are statistically significant for testing  $\text{KRV}(\mathbf{Q}_x, \mathbf{Q}^{(2)})$ . This indicates that  $\mathbf{Q}^{(1)}$  is still informative in spite of small perturbations, but  $\mathbf{Q}^{(2)}$  is completely mis-specified.

The results of Grace, GMDI-d and GMDI-k for  $R^2 = 0.8$  are summarized in Fig. 4. It can be seen that with small perturbations, i.e.,  $\mathbf{Q}^{(1)}$ , all three methods can still control the type-I error, and GMDI has higher power than Grace. When  $\mathbf{Q}$  is uninformative, i.e.,  $\mathbf{Q}^{(2)}$ , none of the three methods can differentiate between zero and non-zero entries. This simulation also indicates the importance and effectiveness of using the KRV test to examine the informativeness of the column structures before implementing the GMDI.

**Setting III:** Next, we assessed the effectiveness of KRV and MiRKAT in terms of identifying informative sample (row) structures. We simulated  $\mathbf{X}$  from the matrix variate normal distribution with mean  $\mathbf{0}$ , row covariance  $\mathbf{R}$  and column covariance  $\mathbf{\Sigma}$ , where  $\mathbf{\Sigma}$  is defined in Setting I, and

$$(\mathbf{R}^{-1})_{(i,j)} = \begin{cases} 1, & i = j \\ 0.9^{|i-j|}, & i \neq j, i \leq 100, j \leq 100 \\ 0.5^{|i-j|}, & i \neq j, i > 100, j > 100 \\ 0, & \text{otherwise.} \end{cases}$$

Finally, we simulated  $\mathbf{y} = 5\mathbf{X}\beta^* + \epsilon$ , where  $\beta^*$  is the same as defined in Setting I, and  $\epsilon$  follows a multivariate normal distribution with mean  $\mathbf{0}$  and covariance  $\mathbf{R}$ . Here, we multiplied  $\beta^*$  by 5 such that the model  $R^2$  is approximately 0.5. We considered six choices of  $\mathbf{H}$ :  $\mathbf{H}^{(1)} = \mathbf{R}^{-1}$ , the true row structure;  $\mathbf{H}^{(2)}$  has slightly mis-specified off-diagonal entries, defined similar to  $\mathbf{H}^{(1)}$  except that  $\mathbf{H}_{(i,j)}^{(2)} = 0.1^{|i-j|}$  for all  $(i, j) \in \{(a, b) : (a-100)(b-100) < 0\}$ ;  $\mathbf{H}^{(3)}$  captures the block diagonal structure of the true row correlation but has mis-

TABLE 2

The proportion of statistically significant KRV and MiRKAT tests based on 500 independent data sets under the significance level 0.01 (%).

	$\mathbf{H}^{(1)}$	$\mathbf{H}^{(2)}$	$\mathbf{H}^{(3)}$	$\mathbf{H}^{(4)}$	$\mathbf{H}^{(5)}$	$\mathbf{H}^{(6)}$
KRV	100	100	0	100	100	0
MiRKAT	100	100	0	100	21	0

specified entries:

$$\left(\mathbf{H}^{(3)}\right)_{(i,j)} = \begin{cases} 1, & i = j \\ (-0.4)^{|i-j|}, & i \neq j, i \leq 100, j \leq 100 \\ (-0.8)^{|i-j|}, & i \neq j, i > 100, j > 100 \\ 0, & \text{otherwise;} \end{cases}$$

$\mathbf{H}^{(4)}$  correctly specifies the correlation structure among the first 100 individuals but has a mis-specified structure for the other individuals:

$$\left(\mathbf{H}^{(4)}\right)_{(i,j)} = \begin{cases} 1, & i = j \\ 0.9^{|i-j|}, & i \neq j, i \leq 100, j \leq 100 \\ (-1)^{|i-j|} \times 0.002, & \text{otherwise;} \end{cases}$$

$\mathbf{H}^{(5)}$  correctly specifies the correlation structure among the first 20 individuals but has a mis-specified structure for the other individuals:

$$\left(\mathbf{H}^{(5)}\right)_{(i,j)} = \begin{cases} 1, & i = j \\ 0.9^{|i-j|}, & i \neq j, i \leq 20, j \leq 20 \\ (-1)^{|i-j|} \times 0.005, & \text{otherwise;} \end{cases}$$

$\mathbf{H}^{(6)}$  has completely mis-specified structures with  $\left(\mathbf{H}^{(6)}\right)_{ij} = (-0.5)^{|i-j|}$  for  $i, j = 1, \dots, 200$ . Here, the coefficients 0.002 and 0.005 were selected such that the smallest eigenvalues of  $\mathbf{H}^{(4)}$  and  $\mathbf{H}^{(5)}$  are both around 0.05. To test whether the six choices of  $\mathbf{H}$  are informative, we applied the KRV and MiRKAT tests using the R functions `KRV()` and `MiRKAT()`, respectively (Zhao et al., 2015). Table 2 summarizes the proportion of the statistically significant tests based on 500 simulated data sets under the significance level 0.01. As expected, both  $\mathbf{H}^{(1)}$  and  $\mathbf{H}^{(2)}$  are informative, because they are the same as or very close to the true row structure. Notably,  $\mathbf{H}^{(4)}$  is also deemed informative in spite of only capturing the true correlations among half of the total individuals. Since  $\mathbf{H}^{(6)}$  is completely mis-specified, its lack of informativeness can be foreseen. However,  $\mathbf{H}^{(3)}$  is also deemed uninformative in spite of correctly capturing the block-diagonal structure of the true correlation matrix. As we will see in Fig. 5,  $\mathbf{H}^{(1)}$ ,  $\mathbf{H}^{(2)}$ , and  $\mathbf{H}^{(4)}$  can lead to well-controlled type-I error rates and decent powers for GMDI-k, whereas  $\mathbf{H}^{(3)}$  and  $\mathbf{H}^{(6)}$  can yield highly inflated type-I error rates. Among the six choices,  $\mathbf{H}^{(5)}$  is the most special because all of the KRV tests are statistically significant but only 21% of the MiRKAT tests are statistically significant. This indicates that  $\mathbf{H}^{(5)}$  is informative of the row structure of  $\mathbf{X}$  but not predictive of the outcome  $\mathbf{y}$ . As discussed in Section 3.3, such a structure is not regarded as informative and should not be used in practice. Indeed, as we will see in Fig. 5,  $\mathbf{H}^{(5)}$  can lead to inflated type-I error rates.

We implemented the proposed GMDI-k and GMDI-d with respect to  $\mathbf{Q} = \mathbf{\Sigma}^{-1}$  and all six choices of  $\mathbf{H}$ . We only reported the performance of existing HDI methods under  $\mathbf{H}^{(1)}$ , because these methods are not affected by the selection of  $\mathbf{H}$ . All the existing methods fail to differentiate non-zero coefficients from zero ones because they assume *i.i.d* samples, which is violated in this setting. In particular, the dscore test can control the type-I error

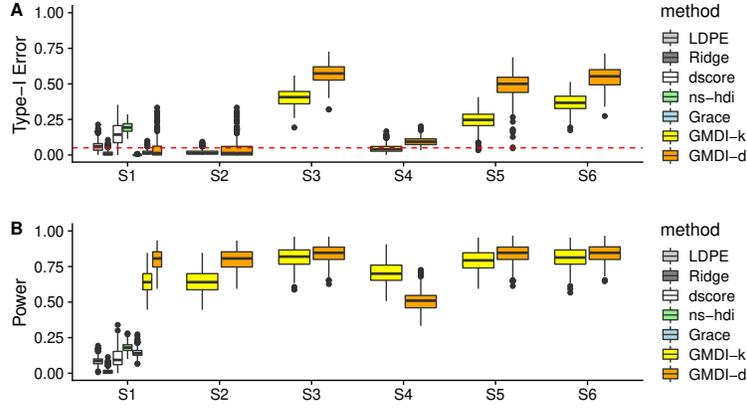


Fig 5: Boxplots of the type-I error (A) and power (B) over 500 replications for Setting III (S1):  $\mathbf{H}^{(1)}$ ; (S2):  $\mathbf{H}^{(2)}$ ; (S3):  $\mathbf{H}^{(3)}$ ; (S4):  $\mathbf{H}^{(4)}$ ; (S5):  $\mathbf{H}^{(5)}$ ; (S6):  $\mathbf{H}^{(6)}$ . None of the existing HDI methods can differentiate between zero and non-zero entries. GMDI-k and GMDI-d have highly inflated type-I error rates for  $\mathbf{H}^{(3)}$ ,  $\mathbf{H}^{(5)}$ , and  $\mathbf{H}^{(6)}$ , which are uninformative structures according to the KRV and MiRKAT tests in Table 2.

in Setting I, but it fails in this setting where samples are correlated. When the selected  $\mathbf{H}$  is correctly specified (e.g.,  $\mathbf{H}^{(1)}$ ) or has small perturbations (e.g.,  $\mathbf{H}^{(2)}$ ), both GMDI-k and GMDI-d show well-controlled type-I error rates, and GMDI-d shows the highest power; this is consistent with Fig. 3. When the selected  $\mathbf{H}$  is partially informative (e.g.,  $\mathbf{H}^{(4)}$ ), GMDI-k shows better controlled type-I error rates and higher power, compared to GMDI-d. This may indicate GMDI-k is more robust regarding partially informative structures. When the selected  $\mathbf{H}$  is uninformative (e.g.,  $\mathbf{H}^{(3)}$ ,  $\mathbf{H}^{(5)}$ , and  $\mathbf{H}^{(6)}$ ), both GMDI-d and GMDI-k suffer from a large inflation of the type-I error rate. This simulation demonstrates the effectiveness of using the KRV and MiRKAT tests to avoid uninformative row structures before implementing the GMDI.

**Setting IV:** We examine the robust GMDI procedure in Section 3.4 using a simulation study with partially informative row structures. Similar to Setting III, we simulated  $\mathbf{X}$  from the matrix variate normal distribution with mean  $\mathbf{0}$ , row covariance  $\mathbf{R}$ , and column covariance  $\mathbf{\Sigma}$ , where  $\mathbf{\Sigma}$  and  $\mathbf{R}$  are, respectively, defined in Setting I and III. We then generated the response  $\mathbf{y} = 10\mathbf{X}\beta^* + \epsilon$ , where  $\beta^*$  is defined in Setting I, and  $\epsilon \sim N_n(\mathbf{0}, \mathbf{R})$ . By design, the model  $R^2$  is approximately 0.85. According to Assumptions (A1)-(A3),  $\mathbf{R}^{-1}$  and  $\mathbf{\Sigma}^{-1}$  are fully informative row and column structures, respectively. We next constructed partially informative row structures by thresholding the tail eigenvalues of  $\mathbf{R}$ . Specifically, letting  $\mathbf{R} = \sum_{i=1}^n d_{r,i} \mathbf{v}_{r,i} \mathbf{v}_{r,i}^T$  denote the eigen-decomposition of  $\mathbf{R}$ , we defined  $\mathbf{H}(\theta) = \sum_{i=1}^{k(\theta)} d_{r,i}^{-1} \mathbf{v}_{r,i} \mathbf{v}_{r,i}^T$ , where  $k(\theta)$  is the smallest integer such that  $\sum_{i=1}^{k(\theta)} d_{r,i} / \sum_{i=1}^n d_{r,i} \geq \theta$  for any given threshold  $\theta \in (0, 1]$ . Note that  $\mathbf{H}(1) = \mathbf{R}^{-1}$ , which is a fully informative row structure. When  $\theta < 1$ ,  $\mathbf{H}(\theta)$  is partially informative with larger values of  $\theta$  leading to a more informative structure.

We implemented the GMDI-k and GMDI-d with respect to  $\mathbf{Q} = \mathbf{\Sigma}^{-1}$  and  $\mathbf{H} = \mathbf{H}(\theta)$  for  $\theta = 0.5, 0.8$ , and 1. For  $\theta = 0.5$  and 0.8, we also implemented the proposed robust GMDI procedure with  $\mathbf{Q} = \mathbf{\Sigma}^{-1}$  and  $\mathbf{H} = \tau \mathbf{H}(\theta) + (1 - \tau) \mathbf{I}_n$ , as described in Section 3.4. We denote the robust procedures for GMDI-k and GMDI-d by r-GMDI-k and r-GMDI-d, respectively. Figure 6 shows the type-I error rates and powers for all the scenarios over 500 independent replications. When  $\mathbf{H}$  is partially informative, i.e.,  $\theta = 0.5$  and 0.8, both GMDI-k and GMDI-d have inflated type-I error rates, and GMDI-d has compromised power. GMDI-k shows more

robustness to partially informative row structures than GMDI-d, which is consistent with S4 in Setting III. The robust GMDI procedures have significantly better performance in terms of better-controlled type-I error rates and enhanced powers. In particular, the robust GMDI-k procedure even has higher power than the GMDI with a fully informative row structure. This may be due to the fact that GMDI yields conservative  $p$ -values (the type-I error rates are mostly 0 when  $\theta = 1$ ), which could be alleviated by the robust GMDI procedure.

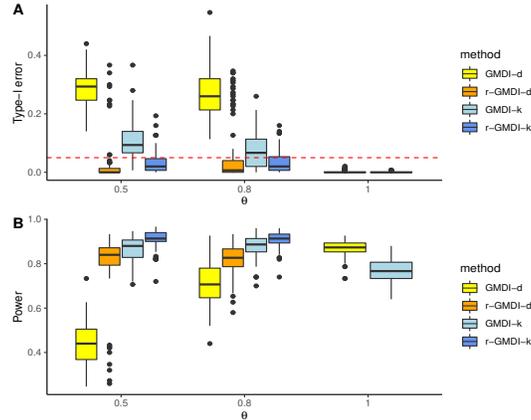


Fig 6: Boxplots of the type-I error (A) and power (B) over 500 replications for Setting IV with  $\theta = 0.5, 0.8, 1$ . The proposed robust GMDI procedure can control the type-I error rate and enhance power when the auxiliary row structure is partially informative.

**5. Analysis of Gut Microbiome Data.** In this section, we illustrate the proposed GMDR and GMDI by analyzing a gut microbiome data set from [Yatsunenko et al. \(2012\)](#), which was described briefly in the Introduction. We kept  $p = 149$  bacterial genera that were present in at least 25% of the  $n = 100$  samples. To make the measurements comparable between subjects, we applied the CLR transformation to obtain a  $100 \times 149$  data matrix  $\mathbf{X}$ , as done in Section 4.1. For the column structure, we used the inverse of the  $p \times p$  matrix of patristic similarities between the tips of the phylogenetic tree, as in Section 4.1. The row structure is derived from sample similarities based on Enzyme Commission (EC) numbers which provide insights into the microbial function: counts of EC numbers specify enzyme-catalyzed reactions based on bacterial genomic content. This gives a reasonable auxiliary view of microbial community similarity since evolutionary diversity in bacteria is correlated with metabolic diversity. Specifically, these EC data represent counts of 432 classes of enzymes observed in the bacteria from the same  $n = 100$  individuals. We applied the CLR transformation to rows of the EC data and centered its columns to have a mean of zero. The resulting  $100 \times 432$  matrix is denoted by  $\mathbf{Z}$ . The row similarity structure is then estimated by the inverse Euclidean kernel  $\mathbf{H} = n(\mathbf{Z}\mathbf{Z}^T)^{-1}$ . For clarity, in this example we denote the row and column structure respectively by  $\mathbf{H}^M$  and  $\mathbf{Q}^M$ . The KRV test yields zero  $p$ -values for both  $\mathbf{H}^M$  and  $\mathbf{Q}^M$ , indicating the informativeness of  $\mathbf{H}^M$  and  $\mathbf{Q}^M$ .

We aim to identify bacterial taxa associated with age. The human microbiome is a complex ecosystem and plays a crucial role in the host's development, nutrition, and immunity ([Belkaid and Hand, 2014](#); [Bana and Cabreiro, 2019](#)). The human microbiome has been found to be associated with many age-related diseases, including cancer and neurodegenerative disorders ([Sepich-Poore et al., 2021](#); [Fang et al., 2020](#)). Therefore, identifying age-associated taxa is important for uncovering the mechanistic link between the microbiome and aging. In

this dataset, the individuals' ages range from 6 months to 53 years. As the distribution of age is highly skewed (around 70% of the samples are below 3 years of age), we use the logarithm of age as our response variable, denoted by  $y$ . MiRKAT yields a zero  $p$ -value when testing the association between  $\mathbf{H}^M$  and  $y$ , indicating the row structure  $\mathbf{H}^M$  also informs the outcome  $y$ .

Besides the marginal analysis result shown in Fig. 1B, it is more interesting to examine the conditional association between each bacterial genus and age, as bacteria do not live independently. We implemented r-GMDI-k and r-GMDI-d to detect conditional associations between bacterial genera and age; the estimated robust row structure was  $\mathbf{H}^R = 0.996\mathbf{H}^M + 0.004\mathbf{I}_n$ . This again indicates the strong informativeness of  $\mathbf{H}^M$ . The GMDI bias-correction procedure yielded a sparse estimator  $\tilde{\beta}(\lambda) \in \mathbb{R}^{149}$  with 13 non-zero entries scattered over the index space  $\{1, \dots, 149\}$  (see (13) for the definition of  $\tilde{\beta}(\lambda)$ ). This indicates that the initial estimate  $\beta^{\text{init}}$  aligns with the space spanned by 13 eigenvectors of  $\mathbf{Q}^M$ . For r-GMDI-d, only 2 out of the 100 GMD components were excluded for having less than 0.1% of the total variance, and 31 GMD components were selected by the proposed VI-based procedure. We found that the organic lasso procedure for estimating  $\sigma^2$  (see the definition of  $\sigma^2$  in Assumption (A1)) is numerically unstable, which may yield slightly different GMDI results for different runs. Thus, we fitted the organic lasso 100 times and obtained the average estimate of  $\sigma^2$ , based on which we implemented the robust GMDI-d and GMDI-k with  $\mathbf{H}^R$  and  $\mathbf{Q}^M$ . As a reference, we also implemented the Grace test (Zhao and Shojaie, 2016), Ridge test (Bühlmann, 2013), and LDPE (Zhang and Zhang, 2014). The dscore and ns-hdi tests were not implemented because they failed to control the type-I error rates in Fig. 2. The Grace test was implemented using  $\mathbf{L} = (\mathbf{Q}^M)^{-1}$ . We considered a two-sided significance level  $\alpha = 0.05$  for all the tests.

Genera found statistically significantly associated with age after controlling for FDR at 0.1 are reported in Table 3. While the Ridge test results in no statistically significant genera, the Grace test and LDPE are able to detect 10 and 3 statistically significant microbes, respectively. By incorporating the auxiliary information, r-GMDI-d can detect more genera, whereas r-GMDI-k appears conservative. This is consistent with the results in Fig. 3. In addition, all the microbes detected by LDPE and r-GMDI-k are also detected by r-GMDI-d; five out of the ten microbes detected by Grace are also detected by r-GMDI-d. However, compared to the vast majority of taxa that are marginally associated with age shown in Fig. 1B, the number of statistically significant conditional associations is relatively small. This may indicate that only a limited number of microbes are near the end of the causal pathways linking the microbiome and age. However, without adjusting for potential confounders, we have to be cautious about making any causal interpretations, such as, which microbes are drivers or followers of the detected age-microbiome associations.

The bacterial genus *Staphylococcus*, detected by LDPE, Grace, and r-GMDI-d, is known as a dominant microbe in newborns delivered by Cesarean section (Dominguez-Bello et al., 2010). *Bifidobacterium*, identified by Grace and r-GMDI-d, was highlighted in Yatsunenکو et al. (2012) as one of the four dominant baby gut microbes. This may indicate the informativeness of  $\mathbf{Q}^M$  for identifying age-associated bacterial genera. *Dialister*, detected by Grace and GMDI, has been shown to play a role in age-related diseases, such as obesity and diabetes (Xu et al., 2020; Gurung et al., 2020). *Veillonella*, identified only by the two GMDI methods, is a signature of infant (4-month-old) microbiome and breastfeeding (Bäckhed et al., 2015). One particular genus only detected by r-GMDI-d, *Catenibacterium*, has been shown to be associated with decreased lifetime cardiovascular disease risk (Kelly et al., 2016).

**6. Discussion.** This paper proposes estimation and inference procedures for high-dimensional linear regression with two-way structured data. For estimation, we develop

TABLE 3

Genera found to be associated with age after controlling for FDR at 0.1 using the Ridge test, LDPE, the Grace test,  $r$ -GMDI-d, and  $r$ -GMDI-k. The microbes are arranged alphabetically according to their names.

	Genus	Total
Ridge	(none)	0
LDPE	<i>Desulfovibrio, Methanobrevibacter, Staphylococcus</i>	3
Grace	<i>Abiotrophia, Bifidobacterium, Desulfovibrio, Dialister, Holdemania, Lachnobacterium, Methanobrevibacter, Roseburia, Rothia, Staphylococcus</i>	10
$r$ -GMDI-d	<i>Adlercreutzia, Anaerococcus, Anaerotruncus, Atopobium, Bifidobacterium, Catenibacterium, Desulfovibrio, Dialister, Diaphorobacter, Erwinia, Kocuria, Limnohabitans, Methanobrevibacter, Mitsuokella, Plesiomonas, Proteus, Pseudobutyrvibrio, Staphylococcus, Streptococcus, Veillonella</i>	20
$r$ -GMDI-k	<i>Atopobium, Dialister, Erwinia, Veillonella</i>	4

GMDR which accounts for arbitrary pre-specified two-way structures. For inference of individual regression coefficients, we propose GMDI, a general high-dimensional inferential framework for a large family of estimators that include the GMDR estimator. Compared to existing high-dimensional inferential tools, GMDI does not require the true regression coefficients to be sparse, it allows dependent and heteroscedastic samples, and it provides flexibility for users to specify relevant auxiliary row and column structures.

We have also proposed a robust GMDI procedure for incorporating a partially informative row structure. In practice, one may have multiple partially informative row structures obtained from different data sources. We can extend the weighting scheme in Section 3.4 to this scenario. Suppose we observe  $N - 1$  informative structures  $\mathbf{H}_1, \dots, \mathbf{H}_{N-1}$ , for some  $N \geq 2$ . Let  $\boldsymbol{\pi} = (\pi_1, \dots, \pi_{N-1})^\top$  with  $\pi_l \geq 0$  for  $l = 1, \dots, N - 1$  and  $\sum_{l=1}^{N-1} \pi_l \leq 1$ , and one can consider  $\mathbf{H}(\boldsymbol{\pi}) = \sum_{l=1}^{N-1} \pi_l \mathbf{H}_l + \left(1 - \sum_{l=1}^{N-1} \pi_l\right) \mathbf{I}_n$ . One can find the  $\boldsymbol{\pi}$  that yields the best prediction accuracy using a constrained optimization method. The proposed robust GMDI procedure may be extended to handle a partially informative column structure  $\mathbf{Q}$ . However, simply taking a linear combination  $\mathbf{Q}(\tau) = \tau \mathbf{Q} + (1 - \tau) \mathbf{I}_p$  may not be effective because  $\mathbf{Q}(\tau)$  has the same set of eigenvectors as  $\mathbf{Q}$  for any  $\tau \in (0, 1]$ . As a result,  $\mathbf{Q}(\tau)$  would not satisfy Assumption (A2) better than  $\mathbf{Q}$ . We leave these extensions as future investigations.

The proposed GMDR and GMDI also provide a framework for supervised integrative analysis of multi-view data, i.e., data collected from multiple sources on the same subjects, which are becoming increasingly common in biology, neuroscience, and engineering (Li, Yang and Zhang, 2018; Zhang et al., 2019; Mars, Jbabdi and Rushworth, 2021). As demonstrated in Section 5, an informative row structure can be obtained from another data view that collects different features on the same set of samples. Analogously, when there are additional studies addressing the same scientific question, in other words, measuring the same set of variables, one can obtain the column structure from these studies in a similar way.

While the proposed method is motivated and illustrated using microbiome data, our method is generally applicable to arbitrary two-way structured data, such as gene expression data and neuroimaging data. It is often possible to obtain informative auxiliary row and/or column structures for these data. For the analysis of gene expression data, one can obtain the gene pathway information from, for example, Kyoto Encyclopedia of Genes and Genomes (KEGG, Kanehisa, 2000) or NCI Pathway Interaction Database (Schaefer et al., 2009) and define  $\mathbf{Q}$  as the graph Laplacian of the gene pathway. For the analysis of neuroimaging data, these structures are often defined as smoothing matrices relevant to the spatial/temporal structure of the images. Specifically, for functional MRI (fMRI) studies that measure images of the brain over time, one can take  $\mathbf{Q}$  to be the graph Laplacian of the graph connecting voxels in the brain (Karas et al., 2019), and  $\mathbf{H} = (h_{ij})$  to be an exponentially smoothing matrix with

$h_{ij} = \exp\{-(t_i - t_j)^2/\kappa_i\}$ , where  $t_i$  and  $t_j$  are the  $i$ -th and  $j$ -th time points, respectively, and  $\kappa > 0$  is a tuning parameter (Allen, Grosenick and Taylor, 2014).

It would be useful to extend GMDR and GMDI to account for confounders. Letting  $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_n)^\top$  denote the low-dimensional matrix of confounders, we consider the following semi-parametric model

$$(23) \quad \mathbf{y} = g(\mathbf{Z}) + \mathbf{X}\beta^* + \epsilon,$$

where  $g(\mathbf{Z}) = (g(\mathbf{z}_1), \dots, g(\mathbf{z}_n))^\top$  with  $g(\cdot)$  being an unknown smooth function,  $\mathbb{E}[\epsilon | \mathbf{Z}, \mathbf{X}] = \mathbf{0}$ , and  $\text{Cov}[\epsilon | \mathbf{Z}, \mathbf{X}] = \Psi$ . To extend GMDR and GMDI to model (23), we leverage the connection between model (23) and the following linear mixed model (Liu, Lin and Ghosh, 2007):

$$(24) \quad \mathbf{y} = \mathbf{g} + \mathbf{X}\beta^* + \epsilon;$$

here,  $g$  is an  $n \times 1$  vector of random effects with mean  $\mathbf{0}$  and covariance  $\sigma_z^2 \mathbf{K}_Z$ , where  $\mathbf{K}_Z = (K(\mathbf{z}_i, \mathbf{z}_j))_{i,j=1,\dots,n}$  for some pre-specified kernel  $K(\cdot, \cdot)$ . Popular choices of  $K(\cdot, \cdot)$  include the Gaussian kernel  $K(\mathbf{z}_i, \mathbf{z}_j) = \exp\{-\|\mathbf{z}_i - \mathbf{z}_j\|^2/\rho\}$  and the  $d$ -th polynomial kernel  $K(\mathbf{z}_i, \mathbf{z}_j) = (\mathbf{z}_i^\top \mathbf{z}_j + \rho)^d$ , where  $\rho$  and  $d$  are tuning parameters. Letting  $\delta = \mathbf{g} + \epsilon$ , we obtain the marginal representation of model (24):  $\mathbf{y} = \mathbf{X}\beta^* + \delta$ , where  $\mathbb{E}[\delta | \mathbf{X}, \mathbf{Z}] = \mathbf{0}$  and  $\text{Cov}[\delta | \mathbf{X}, \mathbf{Z}] = \sigma_z^2 \mathbf{K}_Z + \Psi$ . Since  $(\sigma_z^2 \mathbf{K}_Z + \Psi)^{-1} = \Psi^{-1} (\sigma_z^2 \mathbf{K}_Z \Psi^{-1} + \mathbf{I}_n)^{-1}$ , one can then implement GMDR and GMDI with the row structure  $\sigma^{-2} \mathbf{H} (\sigma_z^2 \mathbf{K}_Z \sigma^{-2} \mathbf{H} + \mathbf{I}_n)^{-1}$  and the column structure  $\mathbf{Q}$  for some  $\mathbf{H}$  and  $\mathbf{Q}$  satisfying Assumptions (A1)-(A3), where  $\sigma^2$  is introduced in Assumption (A1). Assuming the normality of  $\mathbf{g}$  and  $\epsilon$ , the variance components  $\sigma^2$  and  $\sigma_z^2$  may be obtained by using penalized maximum likelihood estimation, which we leave for future investigation.

Finally, it would be interesting to extend GMDR and GMDI to analyze two-way structured categorical predictors. However, since the GMD incorporates  $\mathbf{H}$  and  $\mathbf{Q}$  through the  $\mathbf{H}, \mathbf{Q}$ -norm in (2), which is not suitable for categorical data, the current GMDR and GMDI framework are not directly applicable to categorical data. To address this issue, an extension of GMD that replaces the  $\mathbf{H}, \mathbf{Q}$ -norm with some appropriate norm for categorical variables is essential, which could be a fruitful future research direction.

**Acknowledgement.** We thank the editor, A.E., and reviewers for many helpful comments.

## REFERENCES

- ALLEN, G. I., GROSEINICK, L. and TAYLOR, J. (2014). A Generalized Least-Square Matrix Decomposition. *Journal of the American Statistical Association* **109** 145-159.
- BÄCKHED, F., ROSWALL, J., PENG, Y., FENG, Q., JIA, H., KOVATCHEVA-DATCHARY, P., LI, Y., XIA, Y., XIE, H., ZHONG, H. et al. (2015). Dynamics and stabilization of the human gut microbiome during the first year of life. *Cell host & microbe* **17** 690-703.
- BANA, B. and CABREIRO, F. (2019). The microbiome and aging. *Annual Review of Genetics* **53** 239-261.
- BELKAID, Y. and HAND, T. W. (2014). Role of the microbiota in immunity and inflammation. *Cell* **157** 121-141.
- BELLONI, A., CHERNOZHUKOV, V. and KATO, K. (2015). Uniform post-selection inference for least absolute deviation regression and other Z-estimation problems. *Biometrika* **102** 77-94.
- 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.
- BÜHLMANN, P. (2013). Statistical significance in high-dimensional linear models. *Bernoulli* **19** 1212-1242.
- CAPORASO, J. G., KUCZYNSKI, J., STOMBAUGH, J., BITTINGER, K., BUSHMAN, F. D., COSTELLO, E. K., FIERER, N., PENA, A. G., GOODRICH, J. K., GORDON, J. I. et al. (2010). QIIME allows analysis of high-throughput community sequencing data. *Nature methods* **7** 335-336.
- COOK, R. D. (2007). Fisher Lecture: Dimension Reduction in Regression. *Statistical Science* **22** 1-26.

- CUESTA, S. M., RAHMAN, S. A., FURNHAM, N. and THORNTON, J. M. (2015). The classification and evolution of enzyme function. *Biophysical journal* **109** 1082–1086.
- DOMINGUEZ-BELLO, M. G., COSTELLO, E. K., CONTRERAS, M., MAGRIS, M., HIDALGO, G., FIERER, N. and KNIGHT, R. (2010). Delivery mode shapes the acquisition and structure of the initial microbiota across multiple body habitats in newborns. *Proceedings of the National Academy of Sciences* **107** 11971–11975.
- ESCOUFIER, Y. (1987). The duality diagram: a means for better practical applications. In *Developments in Numerical Ecology* 139–156. Springer.
- ESCOUFIER, Y. (2006). Operator related to a data matrix: a survey. In *Compstat 2006 - Proceedings in Computational Statistics* (A. RIZZI and M. VICHI, eds.) 285–297. Physica-Verlag HD, Heidelberg.
- FANG, P., KAZMI, S., JAMESON, K. and HSIAO, E. (2020). The microbiome as a modifier of neurodegenerative disease risk. *Cell Host & Microbe* **28** 201–222.
- FRIEDMAN, J., HASTIE, T., TIBSHIRANI, R. et al. (2001). *The elements of statistical learning* **1**. Springer series in statistics New York.
- GOLUB, G. H. and VAN LOAN, C. F. (2013). *Matrix computations*. JHU press.
- GUPTA, A. K. and NAGAR, D. K. (2018). *Matrix variate distributions*. Chapman and Hall/CRC.
- GURUNG, M., LI, Z., YOU, H., RODRIGUES, R., JUMP, D. B., MORGUN, A. and SHULZHENKO, N. (2020). Role of gut microbiota in type 2 diabetes pathophysiology. *EBioMedicine* **51** 102590.
- HULLAR, M. A., JENKINS, I. C., RANDOLPH, T. W., CURTIS, K. R., MONROE, K. R., ERNST, T., SHEPHERD, J. A., STRAM, D. O., CHENG, I., KRISTAL, B. S. et al. (2021). Associations of the gut microbiome with hepatic adiposity in the Multiethnic Cohort Adiposity Phenotype Study. *Gut microbes* **13** 1965463.
- JAVANMARD, A. and MONTANARI, A. (2014a). Confidence Intervals and Hypothesis Testing for High-Dimensional Regression. *Journal of Machine Learning Research* **15** 2869–2909.
- JAVANMARD, A. and MONTANARI, A. (2014b). Hypothesis testing in high-dimensional regression under the gaussian random design model: Asymptotic theory. *IEEE Transactions on Information Theory* **60** 6522–6554.
- KANEHISA, M. (2000). *Post-genome informatics*. OUP Oxford.
- KARAS, M., BRZYSKI, D., DZEMIDZIC, M., GOŃI, J., KAREKEN, D. A., RANDOLPH, T. W. and HAREZLAK, J. (2019). Brain connectivity-informed regularization methods for regression. *Statistics in Biosciences* **11** 47–90.
- KELLY, T. N., BAZZANO, L. A., AJAMI, N. J., HE, H., ZHAO, J., PETROSINO, J. F., CORREA, A. and HE, J. (2016). Gut microbiome associates with lifetime cardiovascular disease risk profile among bogalusa heart study participants. *Circulation research* **119** 956–964.
- LI, S., CAI, T. T. and LI, H. (2021). Inference for high-dimensional linear mixed-effects models: A quasi-likelihood approach. *Journal of the American Statistical Association* 1–12.
- LI, Y., YANG, M. and ZHANG, Z. (2018). A survey of multi-view representation learning. *IEEE transactions on knowledge and data engineering* **31** 1863–1883.
- LIU, D., LIN, X. and GHOSH, D. (2007). Semiparametric regression of multidimensional genetic pathway data: Least-squares kernel machines and linear mixed models. *Biometrics* **63** 1079–1088.
- LOZUPONE, C. and KNIGHT, R. (2005). UniFrac: a new phylogenetic method for comparing microbial communities. *Applied and environmental microbiology* **71** 8228–8235.
- MARS, R. B., JBABDI, S. and RUSHWORTH, M. F. (2021). A common space approach to comparative neuroscience. *Annual Review of Neuroscience* **44**.
- MITRA, R. and ZHANG, C.-H. (2016). The benefit of group sparsity in group inference with de-biased scaled group Lasso. *Electronic Journal of Statistics* **10** 1829–1873.
- NEUHOUSER, M. L., SCHWARZ, Y., WANG, C., BREYMEYER, K., CORONADO, G., WANG, C.-Y., NOAR, K., SONG, X. and LAMPE, J. W. (2012). A low-glycemic load diet reduces serum C-reactive protein and modestly increases adiponectin in overweight and obese adults. *The Journal of nutrition* **142** 369–374.
- NING, Y. and LIU, H. (2017). A general theory of hypothesis tests and confidence regions for sparse high dimensional models. *The Annals of Statistics* **45** 158–195.
- RANDOLPH, T. W., ZHAO, S., COPELAND, W., HULLAR, M. and SHOJAIE, A. (2018). Kernel-penalized regression for analysis of microbiome data. *The Annals of Applied Statistics* **12** 540–566.
- SCHAEFER, C. F., ANTHONY, K., KRUPA, S., BUCHOFF, J., DAY, M., HANNAY, T. and BUETOW, K. H. (2009). PID: the pathway interaction database. *Nucleic acids research* **37** D674–D679.
- SEPICH-POORE, G. D., ZITVOGEL, L., STRAUSSMAN, R., HASTY, J., WARGO, J. A. and KNIGHT, R. (2021). The microbiome and human cancer. *Science* **371** eabc4552.
- SHARIFI, F. and YE, Y. (2017). From gene annotation to function prediction for metagenomics. In *Protein Function Prediction* 27–34. Springer.
- SUN, T. and ZHANG, C.-H. (2012). Scaled sparse linear regression. *Biometrika* **99** 879–898.
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* **58** 267–288.

- VAN DE GEER, S., BÜHLMANN, P. et al. (2009). On the conditions used to prove oracle results for the Lasso. *Electronic Journal of Statistics* **3** 1360–1392.
- 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.
- WAINWRIGHT, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint* **48**. Cambridge University Press.
- WANG, Y., RANDOLPH, T. W., SHOJAIE, A. and MA, J. (2019). The Generalized Matrix Decomposition Biplot and Its Application to Microbiome Data. *mSystems* **4**.
- WASHBURNE, A. D., MORTON, J. T., SANDERS, J., McDONALD, D., ZHU, Q., OLIVERIO, A. M. and KNIGHT, R. (2018). Methods for phylogenetic analysis of microbiome data. *Nature microbiology* **3** 652–661.
- XU, Y., WANG, N., TAN, H.-Y., LI, S., ZHANG, C. and FENG, Y. (2020). Function of *Akkermansia muciniphila* in obesity: interactions with lipid metabolism, immune response and gut systems. *Frontiers in microbiology* **11** 219.
- YATSUNENKO, T., REY, F. E., MANARY, M. J., TREHAN, I., DOMINGUEZ-BELLO, M. G., CONTRERAS, M., MAGRIS, M., HIDALGO, G., BALDASSANO, R. N., ANOKHIN, A. P., HEATH, A. C., WARNER, B., REEDER, J., KUCZYNSKI, J., CAPORASO, J. G., LOZUPONE, C. A., LAUBER, C., CLEMENTE, J. C., KNIGHTS, D., KNIGHT, R. and GORDON, J. I. (2012). Human gut microbiome viewed across age and geography. *Nature* **486** 222–227.
- YU, G. and BIEN, J. (2019). Estimating the error variance in a high-dimensional linear model. *Biometrika* **106** 533–546.
- ZEEVI, D., KOREM, T., GODNEVA, A., BAR, N., KURILSHIKOV, A., LOTAN-POMPAN, M., WEINBERGER, A., FU, J., WIJMENGA, C., ZHERNAKOVA, A. et al. (2019). Structural variation in the gut microbiome associates with host health. *Nature* **568** 43–48.
- ZHAN, X., PLANTINGA, A., ZHAO, N. and WU, M. C. (2017). A fast small-sample kernel independence test for microbiome community-level association analysis. *Biometrics* **73** 1453–1463.
- ZHANG, C.-H., HUANG, J. et al. (2008). The sparsity and bias of the lasso selection in high-dimensional linear regression. *The Annals of Statistics* **36** 1567–1594.
- ZHANG, Y. and PAN, W. (2015). Principal component regression and linear mixed model in association analysis of structured samples: competitors or complements? *Genetic epidemiology* **39** 149–155.
- ZHANG, C. H. and ZHANG, S. S. (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.
- ZHANG, X., LI, L., BUTCHER, J., STINTZI, A. and FIGEYS, D. (2019). Advancing functional and translational microbiome research using meta-omics approaches. *Microbiome* **7** 1–12.
- ZHAO, S. and SHOJAIE, A. (2016). A significance test for graph-constrained estimation. *Biometrics* **72** 484–493.
- ZHAO, N., CHEN, J., CARROLL, I. M., RINGEL-KULKA, T., EPSTEIN, M. P., ZHOU, H., ZHOU, J. J., RINGEL, Y., LI, H. and WU, M. C. (2015). Testing in microbiome-profiling studies with MiRKAT, the microbiome regression-based kernel association test. *The American Journal of Human Genetics* **96** 797–807.
- ZHU, Y. and BRADIC, J. (2018). Linear Hypothesis Testing in Dense High-Dimensional Linear Models. *Journal of the American Statistical Association* **113** 1583–1600.