JOURNAL OF THE AMERICAN STATISTICAL ASSOCIATION
2019, VOL. 114, NO. 526, 682–696: Theory and Methods
https://doi.org/10.1080/01621459.2018.1429275

Joint Mean and Covariance Estimation with Unreplicated Matrix-Variate Data

Michael Hornstein[a], Roger Fan[a], Kerby Shedden[a], and Shuheng Zhou[a,b]

[a]Department of Statistics, University of Michigan, Ann Arbor, MI; [b]Department of Statistics, University of California, Riverside, CA

ABSTRACT

It has been proposed that complex populations, such as those that arise in genomics studies, may exhibit dependencies among observations as well as among variables. This gives rise to the challenging problem of analyzing unreplicated high-dimensional data with unknown mean and dependence structures. Matrix-variate approaches that impose various forms of (inverse) covariance sparsity allow flexible dependence structures to be estimated, but cannot directly be applied when the mean and covariance matrices are estimated jointly. We present a practical method utilizing generalized least squares and penalized (inverse) covariance estimation to address this challenge. We establish consistency and obtain rates of convergence for estimating the mean parameters and covariance matrices. The advantages of our approaches are: (i) dependence graphs and covariance structures can be estimated in the presence of unknown mean structure, (ii) the mean structure becomes more efficiently estimated when accounting for the dependence structure among observations; and (iii) inferences about the mean parameters become correctly calibrated. We use simulation studies and analysis of genomic data from a twin study of ulcerative colitis to illustrate the statistical convergence and the performance of our methods in practical settings. Several lines of evidence show that the test statistics for differential gene expression produced by our methods are correctly calibrated and improve power over conventional methods. Supplementary materials for this article are available online.

ARTICLE HISTORY
Received November 2016
Revised December 2017

KEYWORDS
Generalized least squares; Genomics; Graphical modeling; Sparsity; Two-group comparison

1. Introduction

Understanding how changes in gene expression are related to changes in biological state is one of the fundamental tasks in genomics research, and is a prototypical example of "large-scale inference" (Efron 2010). While some genomics datasets have within-subject replicates or other known clustering factors that could lead to dependence among observations, most are viewed as population cross-sections or convenience samples, and are usually analyzed by taking observations (biological samples) to be statistically independent of each other. Countering this conventional view, Efron (2009) proposed that there may be unanticipated correlations between samples even when the study design would not suggest it. To identify and adjust for unanticipated sample-wise correlations, Efron (2009) proposed an empirical Bayes approach utilizing the sample moments of the data. In particular, sample-wise correlations may lead to inflated evidence for mean differences, and could be one explanation for the claimed lack of reproducibility in genomics research (Leek et al. 2010; Allen and Tibshirani 2012; Sugden et al. 2013).

A persistent problem in genomics research is that test statistics for mean parameters (e.g., t-statistics for two-group comparisons) often appear to be incorrectly calibrated (Efron 2005; Allen and Tibshirani 2012). When this happens, for example, when test statistics are uniformly overdispersed relative to their intended reference distribution, this is usually taken to be an indication of miscalibration, rather than reflecting a nearly global pattern of differential effects (Efron 2007). Adjustments such as genomic control (Devlin and Roeder 1999) can be used to account for this; a related approach is that of Allen and Tibshirani (2012). In this work, we address unanticipated sample-wise dependence, which can exhibit a strong effect on statistical inference. We propose a new method to jointly estimate the mean and covariance with a single instance of the data matrix, as is common in genetics. The basic idea of our approach is to alternate for a fixed number of steps between mean and covariance estimation. We exploit recent developments in two-way covariance estimation for matrix-variate data (Zhou 2014). We crucially combine the classical idea of generalized least squares (GLS) (Aitken 1936) with thresholding for model selection and estimation of the mean parameter vector. Finally, we use Wald-type statistics to conduct inference. We motivate this approach using differential expression analysis in a genomics context, but the method is broadly applicable to matrix-variate data having unknown mean and covariance structures, with or without replications. We illustrate, using theory and data examples, including a genomic study of ulcerative colitis, that estimating and accounting for the sample-wise dependence can systematically improve the calibration of test statistics, therefore reducing or eliminating the need for certain post-hoc adjustments.

CONTACT Michael Hornstein mdhorn@umich.edu Department of Statistics, University of Michigan, Ann Arbor, MI 48109.
Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/r/JASA.
Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JASA.

With regard to variable selection, another major challenge we face is that variables (e.g., genes or mRNA transcripts) have a complex dependency structure that exists together with any dependencies among observations. As pointed out by Efron (2009) and others, the presence of correlations among the samples makes it more difficult to estimate correlations among variables, and vice versa. A second major challenge is that due to dependence among both observations and variables, there is no independent replication in the data, that is, we have a single matrix to conduct covariance estimation along both axes. This challenge was addressed by Zhou (2014) when the mean structure is taken to be zero. A third major challenge that is unique to our framework is that covariance structures can only be estimated after removing the mean structure, a fact that is generally not considered in most work on high-dimensional covariance and graph estimation, where the population mean is taken to be zero. We elaborate on this challenge next.

1.1. Our Approach and Contributions

Two obvious approaches for removing the mean structure in our setting are to globally center each column of the data matrix (containing the data for one variable), or to center each column separately within each group of sample points to be compared (subsequently referred to as "group centering"). Globally centering each column, by ignoring the mean structure, may result in an estimated covariance matrix that is not consistent. Group centering all genes, by contrast, leads to consistent covariance estimation, as shown in Theorem 3 with regard to Algorithm 1. However, group centering all genes introduces extraneous noise when the true vector of mean differences is sparse. We find that there is a complex interplay between the mean and covariance estimation tasks, such that overly flexible modeling of the mean structure can introduce large systematic errors in the mean structure estimation. To mitigate this effect, we aim to center the data using a model selection strategy. More specifically, we adopt a model selection centering approach in which only mean parameters having a sufficiently large effect size (relative to the dimension of the data) are targeted for removal. This refined approach has theoretical guarantees and performs well in simulations. The estimated covariance matrix can be used in uncertainty assessment and formal testing of mean parameters, thereby improving calibration of the inference.

In Section 2, we define the two group mean model, which is commonly used in the genomics literature, and introduce the GLS algorithm in this context. We bound the statistical error for estimating each column of the mean matrix using the GLS procedure so long as each column of X shares the same covariance matrix B, for which we have a close approximation. It is commonly known that genes are correlated, so correlations exist across columns as well as rows of the data matrix. In particular, in Theorem 1 in Section 3.1, we establish consistency for the GLS estimator given a deterministic $\widehat{B}$ which is close to B in the operator norm, and present the rate of convergence for mean estimation for data generated according to a subgaussian model to be defined in Definition 2.1. Moreover, we do not impose a separable covariance model in the sense of (1).

What distinguishes our model from those commonly used in the genomics literature is that we do not require that individuals are independent. Our approach to covariance modeling builds on the Gemini method (Zhou 2014), which is designed to estimate a separable covariance matrix for data with two-way dependencies. For matrices $A \in \mathbb{R}^{m\times m}$ and $B \in \mathbb{R}^{n\times n}$, the Kronecker product $A \otimes B \in \mathbb{R}^{mn\times nm}$ is the block matrix for which the (i, j)th block is $a_{ij}B$, for $i, j \in \{1, \ldots, m\}$. We say that an $n \times m$ random matrix X follows a matrix variate distribution with mean $M \in \mathbb{R}^{n\times m}$ and a separable covariance matrix

$$X_{n\times m} \sim \mathcal{L}_{n,m}(M, A_{m\times m} \otimes B_{n\times n}), \tag{1}$$

if vec{ X } has mean vec{ M } and covariance $\Sigma = A \otimes B$. Here, vec{ X } is formed by stacking the columns of X into a vector in $\mathbb{R}^{mn}$. For the mean matrix M, we focus on the two-group setting to be defined in (4). Intuitively, A describes the covariance between columns while B describes the covariance between rows of X. Even with perfect knowledge of M, we can only estimate A and B up to a scaling factor, as $A\eta \otimes \frac{1}{\eta}B = A \otimes B$ for any $\eta > 0$, and hence this will be our goal and precisely what we mean when we say we are interested in estimating covariances A and B. However, this lack of identifiability does not affect the GLS estimate, because the GLS estimate is invariant to rescaling the estimate of B^{-1}.

1.2. Related Work

Efron (2009) introduced an approach for inference on mean differences in data with two-way dependence. His approach uses empirical Bayes ideas and tools from large-scale inference, and also explores how challenging the problem of conducting inference on mean parameters is when there are uncharacterized dependences among samples. We combine GLS and variable selection with matrix-variate techniques. Allen and Tibshirani (2012) also consider this question and develop a different approach that uses ordinary least squares (OLS) through the iterations, first decorrelating the residuals and then using OLS techniques again on this adjusted dataset. The confounder adjustment literature in genomics, including Sun, Zhang, and Owen (2012) and Wang et al. (2017), can also be used to perform large-scale mean comparisons in similar settings that include similarity structures among observations. These methods use the same general matrix decomposition framework, where the mean and noise are separated. They exploit low-rank structure in the mean matrix, as well as using sparse approximation of OLS estimates, for example where thresholding. Our model introduces row-wise dependence through matrix-variate noise, while the confounder adjustment literature instead assumes that a small number of latent factors also affect the mean expression, resulting in additional low-rank structure in the mean matrix. Web Supplement Section J contains detailed comparisons between our approach and these alternative methods.

Our inference procedures are based on Z-scores and associated FDR values for mean comparisons of individual variables. While we account for sample-wise correlations, gene–gene correlations remain, which we regard as a nuisance parameter. Our estimated correlation matrix among the genes can be used in future work in combination with the line of work that addresses FDR in the presence of gene correlations. This relies on earlier work for false discovery rate estimation using correlated data, including Owen (2005), Benjamini and Yekutieli (2001), Cai, Jessie Jeng, and Jin (2011), Li and Zhong (2014), Benjamini

and Hochberg (1995), and Storey (2003). Taking a different approach, Hall and Jin (2010) develop the innovated higher criticism test statistics to detect differences in means in the presence of correlations between genes. Our estimated gene–gene correlation matrix can be used in combination with this approach; we leave this as future work. Another line of relevant research has focused on hypothesis testing of high-dimensional means, exploiting assumed sparsity of effects, and developing theoretical results using techniques from high-dimensional estimation theory. Work of this type includes Cai and Xia (2014), Chen, Li, and Zhong (2014), Bai and Saranadasa (1996), and Chen and Qin (2010). Hoff (2011) adopts a Bayesian approach, using a model that is a generalization of the matrix-variate normal distribution.

Our method builds on the Gemini estimator introduced by Zhou (2014), which estimates covariance matrices when both rows and columns of the data matrix are dependent. In the setting where correlations exist along only one axis of the array, researchers have proposed various covariance estimators and studied their theoretical and numerical properties (Banerjee, Ghaoui, and d'Aspremont 2008; Fan, Feng, and Wu 2009; Friedman, Hastie, and Tibshirani 2008; Lam and Fan 2009; Meinshausen and Bühlmann 2006; Peng et al. 2009; Ravikumar et al. 2011; Rothman et al. 2008; Yuan and Lin 2007; Zhou, Lafferty, and Wasserman 2010). Although we focus on the setting of Kronecker products, or separable covariance structures, Cai et al. (2016) proposed a covariance estimator for a model with several populations, each of which may have a different variable-wise covariance matrix. Our methods can be generalized to this setting. Tan and Witten (2014) use a similar matrix-variate datasetting as in (1), but perform biclustering instead of considering a regression problem with a known design matrix.

1.3. Notation and Organization

Before we conclude this section, we introduce the notation needed for the technical sections. Let $e_1, \ldots, e_p$ be the canonical basis of $\mathbb{R}^p$. For a matrix $A = (a_{ij})_{1 \leqslant i,j \leqslant m}$, let $|A|$ denote the determinant and $\mathrm{tr}(A)$ be the trace of A. Let $\|A\|_{\max} = \max_{i,j} |a_{ij}|$ denote the entry-wise max norm. Let $\|A\|_1 = \max_j \sum_{i=1}^m |a_{ij}|$ denote the matrix ℓ_1 norm. The Frobenius norm is given by $\|A\|_F^2 = \sum_i \sum_j a_{ij}^2$. Let $\varphi_i(A)$ denote the ith largest eigenvalue of A, with $\varphi_{\max}(A)$ and $\varphi_{\min}(A)$ denoting the largest and smallest eigenvalues, respectively. Let $\kappa(A)$ be the condition number for matrix A. Let $|A|_{1,\mathrm{off}} = \sum_{i \neq j} |a_{ij}|$ denote the sum of the absolute values of the off-diagonal entries and let $|A|_{0,\mathrm{off}}$ denote the number of nonzero off-diagonal entries. Let $a_{\max} = \max_i a_{ii}$. Denote by $r(A)$ the stable rank $\|A\|_F^2 / \|A\|_2^2$. We write $\mathrm{diag}(A)$ for a diagonal matrix with the same diagonal as A. Let I be the identity matrix. We let $C, C_1, c, c_1, \ldots$ be positive constants which may change from line to line. For two numbers a, b, $a \wedge b := \min(a, b)$ and $a \vee b := \max(a, b)$. Let $(a)_+ := a \vee 0$. For sequences $\{a_n\}, \{b_n\}$, we write $a_n = O(b_n)$ if $|a_n| \leqslant C|b_n|$ for some positive absolute constant C which is independent of n and m or sparsity parameters, and write $a_n \asymp b_n$ if $c|a_n| \leqslant |b_n| \leqslant C|a_n|$. We write $a_n = \Omega(b_n)$ if $|a_n| \geqslant C|b_n|$ for some positive absolute constant C which is independent of n and m or sparsity parameters. We write $a_n = o(b_n)$ if $\lim_{n \to \infty} a_n / b_n = 0$. For random variables X and Y, let $X \sim Y$ denote that X and Y follow the same distribution.

The remainder of the article is organized as follows. In Section 2, we present our matrix-variate modeling framework and methods on joint mean and covariance estimation. In particular, we propose two algorithms for testing mean differences based on two centering strategies. In Section 3, we present convergence rates for these methods. In Theorems 3 and 4, we provide joint rates of convergence for mean and covariance estimation using Algorithms 1 and 2, respectively. We also emphasize the importance of the design effect (see Equation (15)) in testing and present theoretical results for estimating this quantity in Corollary 2 and Corollary 5. In Section 4, we demonstrate through simulations that our algorithms can outperform OLS estimators in terms of accuracy and variable selection consistency. In Section 5, we analyze a gene expression dataset, and show that our method corrects test statistic overdispersion that is clearly present when using sample-mean-based methods (see Section 4.2). We conclude in Section 6. We place all technical proofs and additional simulation and data analysis results in the Web Supplement, which is organized as follows. Sections A and B contain additional simulation and data analysis results. Section C contains some preliminary results and notation. In Section D, we prove Theorem 1. In Sections E and F, we prove Theorem 3. In Section G, we derive entry-wise rates of convergence for the sample covariance matrices. In Sections H and I, we prove Theorem 4 and its auxiliary results. In Section J, we provide additional comparisons between our method and some related methods on both simulated and real data.

2. Models and Methods

In this section, we present our model and method for joint mean and covariance estimation. Our results apply to subgaussian data. Before we present the model, we define subgaussian random vectors and the ψ_2 norm. The ψ_2 condition on a scalar random variable V is equivalent to the subgaussian tail decay of V, which means $P(|V| > t) \leqslant 2\exp(-t^2/c^2)$, for all $t > 0$. For a vector $y = (y_1, \ldots, y_p) \in \mathbb{R}^p$, denote by $\|y\|_2 = \sqrt{\sum_{i=1}^p y_i^2}$.

Definition 2.1. Let Y be a random vector in $\mathbb{R}^p$. (a) Y is called isotropic if for every $y \in \mathbb{R}^p$, $E[|\langle Y, y\rangle|^2] = \|y\|_2^2$. (b) Y is ψ_2 with a constant α if for every $y \in \mathbb{R}^p$,

$$\|\langle Y, y\rangle\|_{\psi_2} := \inf\{t : E[\exp(\langle Y, y\rangle^2 / t^2)] \leqslant 2\} \leqslant \alpha \|y\|_2.$$

Our goal is to estimate the group mean vectors $\beta^{(1)}, \beta^{(2)}$, the vector of mean differences between two groups $\gamma = \beta^{(1)} - \beta^{(2)} \in \mathbb{R}^m$, the row-wise covariance matrix $B \in \mathbb{R}^{n \times n}$, and the column-wise covariance matrix $A \in \mathbb{R}^{m \times m}$. In our motivating genomics applications, the people by people covariance matrix B is often incorrectly anticipated to have a simple known structure, for example, B is taken to be diagonal if observations are assumed to be uncorrelated. However, we show by example in Section 5 that departures from the anticipated diagonal structure may occur, corroborating earlier claims of this type by Efron (2009) and others. Motivated by this example, we define the two-group mean model and the GLS algorithm, which takes advantage of the covariance matrix B.

The model. Our model for the matrix-variate data X can be expressed as a mean matrix plus a noise term,

$$X = M + \varepsilon, \tag{2}$$

where columns (and rows) of ε are sub-Gaussian. Let $u, v, \in \mathbb{R}^n$ be defined as

$$u = (\underbrace{1, \ldots, 1}_{n_1}, \underbrace{0, \ldots, 0}_{n_2}) \in \mathbb{R}^n \quad \text{and}$$
$$v = (\underbrace{0, \ldots, 0}_{n_1}, \underbrace{1, \ldots, 1}_{n_2}) \in \mathbb{R}^n. \tag{3}$$

Let $\mathbf{1}_n \in \mathbb{R}^n$ denote a vector of ones. For the two-group model, we take the mean matrix to have the form

$$M = D\beta = \begin{bmatrix} \mathbf{1}_{n_1}\beta^{(1)T} \\ \mathbf{1}_{n_2}\beta^{(2)T} \end{bmatrix} \in \mathbb{R}^{n\times m}, \quad \text{where} \quad D = [\, u\ v\,] \in \mathbb{R}^{n\times 2} \tag{4}$$

is the design matrix and $\beta = (\beta^{(1)}, \beta^{(2)})^T \in \mathbb{R}^{2\times m}$ is a matrix of group means. Let $\gamma = \beta^{(1)} - \beta^{(2)} \in \mathbb{R}^m$ denote the vector of mean differences. Let $d_0 = |\operatorname{supp}(\gamma)| = |\{j : \gamma_j \neq 0\}|$ denote the size of the support of γ. To estimate the group means, we use a GLS estimator,

$$\widehat{\beta}(\widehat{B}^{-1}) := (D^T\widehat{B}^{-1}D)^{-1}D^T\widehat{B}^{-1}X \in \mathbb{R}^{2\times m}, \tag{5}$$

where $\widehat{B}^{-1}$ is an estimate of the observation-wise inverse covariance matrix. Throughout the article, we denote by $\widehat{\beta}(B^{-1})$ the oracle GLS estimator, since it depends on the unknown true covariance B. Also, we denote the estimated vector of mean differences as $\widehat{\gamma}(\widehat{B}^{-1}) = \delta^T\widehat{\beta}(\widehat{B}^{-1}) \in \mathbb{R}^m$, where $\delta = (1, -1) \in \mathbb{R}^2$.

2.1. Matrix-Variate Covariance Modeling

In the previous section, we have not yet explicitly constructed an estimator of B^{-1}. To address this need, we model the data matrix X with a matrix-variate distribution having a separable covariance matrix, namely, the covariance of $\operatorname{vec}\{\,X\,\}$ follows a Kronecker product covariance model. When ε (2) follows a matrix-variate normal distribution $\mathcal{N}_{n,m}(0, A \otimes B)$, as considered in Zhou (2014), the support of B^{-1} encodes conditional independence relationships between samples, and likewise, the support of A^{-1} encodes conditional independence relationships among genes. The inverse covariance matrices A^{-1} and B^{-1} have the same supports as their respective correlation matrices, so edges of the dependence graphs are identifiable under the model $\operatorname{Cov}(\operatorname{vec}(\varepsilon)) = A \otimes B$. When the data is sub-Gaussian, the method is still valid for obtaining consistent estimators of A, B, and their inverses, but the interpretation in terms of conditional independence does not hold in general.

Our results do not assume normally distributed data; we analyze the subgaussian correspondent of the matrix variate normal model instead. In the Kronecker product covariance model we consider in the present work, the noise term has the form $\varepsilon = B^{1/2}ZA^{1/2}$ for a mean-zero random matrix Z with independent subgaussian entries satisfying $1 = \mathbb{E}Z_{ij}^2 \leqslant \|Z_{ij}\|_{\psi_2} \leqslant K$. Clearly, $\operatorname{vec}\{\,\varepsilon\,\} = \mathrm{A} \otimes \mathrm{B}$. Here, the matrix A represents the shared covariance among variables for each sample, while B represents the covariance among observations which in turn is shared by all genes.

For identifiability, and convenience, we define

$$A^* = \frac{m}{\operatorname{tr}(A)}A \quad \text{and} \quad B^* = \frac{\operatorname{tr}(A)}{m}B, \tag{6}$$

where the scaling factor is chosen so that A^* has trace m. For the rest of the article, A and B refer to A^* and B^*, as defined in (6). Let S_A and S_B denote sample covariance matrices to be specified. Let the corresponding sample correlation matrices be defined as

$$\widehat{\Gamma}_{ij}(A) = \frac{(S_A)_{ij}}{\sqrt{(S_A)_{ii}(S_A)_{jj}}} \quad \text{and} \quad \widehat{\Gamma}_{ij}(B) = \frac{(S_B)_{ij}}{\sqrt{(S_B)_{ii}(S_B)_{jj}}}. \tag{7}$$

The baseline Gemini estimators (Zhou 2014) are defined as follows, using a pair of penalized estimators for the correlation matrices $\rho(A) = (a_{ij}/\sqrt{a_{ii}a_{jj}})$ and $\rho(B) = (b_{ij}/\sqrt{b_{ii}b_{jj}})$,

$$\widehat{A}_\rho = \underset{A_\rho \succ 0}{\arg\min} \left\{ \operatorname{tr}\left(\widehat{\Gamma}(A)A_\rho^{-1}\right) + \log|A_\rho| + \lambda_B|A_\rho^{-1}|_{1,\text{off}} \right\}, \tag{8a}$$

and

$$\widehat{B}_\rho = \underset{B_\rho \succ 0}{\arg\min} \left\{ \operatorname{tr}\left(\widehat{\Gamma}(B)B_\rho^{-1}\right) + \log|B_\rho| + \lambda_A|B_\rho^{-1}|_{1,\text{off}} \right\}, \tag{8b}$$

where the input are a pair of sample correlation matrices as defined in (7).

Let $\widehat{M}$ denote the estimator of the mean matrix M in (1). Denote the centered data matrix and the sample covariance matrices as

$$X_{\text{cen}} = X - \widehat{M}, \quad \text{for } \widehat{M} \text{ to be specified in Algorithms 1 and 2,}$$
$$S_B = X_{\text{cen}}X_{\text{cen}}^T/m, \quad \text{and} \quad S_A = X_{\text{cen}}^TX_{\text{cen}}/n. \tag{9}$$

Define the diagonal matrices of sample standard deviations as

$$\widehat{W}_1 = \sqrt{n}\operatorname{diag}(S_A)^{1/2} \in \mathbb{R}^{m\times m}, \quad \widehat{W}_2 = \sqrt{m}\operatorname{diag}(S_B)^{1/2} \in \mathbb{R}^{n\times n}, \tag{10}$$

$$\text{and} \quad \widehat{A \otimes B} = \left(\widehat{W}_1\widehat{A}_\rho\widehat{W}_1\right) \otimes \left(\widehat{W}_2\widehat{B}_\rho\widehat{W}_2\right) / \|X_{\text{cen}}\|_F^2. \tag{11}$$

2.2. Group-Based Centering Method

We now discuss our first method for estimation and inference with respect to the vector of mean differences $\gamma = \beta^{(1)} - \beta^{(2)}$, for $\beta^{(1)}$ and $\beta^{(2)}$ as in (4). Our approach in Algorithm 1 is to remove all possible mean effects by centering each variable within every group.

Algorithm 1: GLS-Global group centering

Input: X; and $\mathcal{G}(1), \mathcal{G}(2)$: indices of group one and two, respectively.
Output: $\widehat{A}^{-1}, \widehat{B}^{-1}, \widehat{A \otimes B}, \widehat{\beta}(\widehat{B}^{-1}), \widehat{\gamma}, T_j$ for all j

1. *Group center the data.* Let Y_i denote the ith row of the data matrix. To estimate the group mean vectors $\beta^{(1)}, \beta^{(2)} \in$

$\mathbb{R}^m$: Compute sample mean vectors

$$\widetilde{\beta}^{(1)} = \frac{1}{n_1}\sum_{i\in\mathcal{G}(1)} Y_i \quad \text{and} \quad \widetilde{\beta}^{(2)} = \frac{1}{n_2}\sum_{i\in\mathcal{G}(2)} Y_i;$$
$$\text{set} \quad \widehat{\gamma}^{\text{OLS}} = \widetilde{\beta}^{(1)} - \widetilde{\beta}^{(2)}. \tag{12}$$

Center the data by

$$X_{\text{cen}} = X - \widehat{M}, \text{ with } \widehat{M} = \begin{bmatrix} 1_{n_1}\widetilde{\beta}^{(1)T} \\ 1_{n_2}\widetilde{\beta}^{(2)T} \end{bmatrix}.$$

2. *Obtain regularized correlation estimates.*
   (2a) The centered data matrix used to calculate S_A and S_B for Algorithm 1 is $X_{\text{cen}} = (I - P_2)X$, where P_2 is the projection matrix that performs within-group centering,
   $$P_2 = \begin{bmatrix} n_1^{-1}1_{n_1}1_{n_1}^T & 0 \\ 0 & n_2^{-1}1_{n_2}1_{n_2}^T \end{bmatrix} = uu^T/n_1 + vv^T/n_2, \tag{13}$$
   with u and v as defined in (3). Compute sample covariance matrices based on group-centered data: $S_A = \frac{1}{n}X_{\text{cen}}^T X_{\text{cen}} = \frac{1}{n}X^T(I-P_2)X$ and $S_B = \frac{1}{m}X_{\text{cen}}X_{\text{cen}}^T = \frac{1}{m}(I-P_2)XX^T(I-P_2)$.
   (2b) Compute (7) to obtain penalized correlation matrices $\widehat{A}_\rho$ and $\widehat{B}_\rho$ using the Gemini estimators as defined in (8a) and (8b) with tuning parameters to be defined in (23).
3. *Rescale the estimated correlation matrices to obtain penalized covariance*
   $$\widehat{B}^{-1} = m\widehat{W}_2^{-1}\widehat{B}_\rho\widehat{W}_2^{-1} \text{ and}$$
   $$\widehat{A}^{-1} = (\|X_{\text{cen}}\|_F^2/m)\widehat{W}_1^{-1}\widehat{A}_\rho\widehat{W}_1^{-1}. \tag{14}$$
4. *Estimate the group mean matrix* using the GLS estimator as defined in (5).
5. *Obtain test statistics.* The jth test statistic is defined as
   $$T_j = \frac{\widehat{\gamma}_j(\widehat{B}^{-1})}{\sqrt{\delta^T(D^T\widehat{B}^{-1}D)^{-1}\delta}}, \qquad \text{with } \delta = (1,-1) \in \mathbb{R}^2, \tag{15}$$
   and $\widehat{\gamma}_j(\widehat{B}^{-1}) = \delta^T\widehat{\beta}_j(\widehat{B}^{-1})$, for $j = 1,\ldots,m$. Note that T_j as defined in (15) is essentially a Wald test and the denominator is a plug-in standard error of $\widehat{\gamma}_j(B^{-1})$.

2.3. Model Selection Centering Method

In this section, we present Algorithm 2, which aims to remove mean effects that are strong enough to have an impact on covariance estimation. The strategy here is to use a model selection step to identify variables with strong mean effects.

Algorithm 2: GLS-Model selection centering

Input: X, and $\mathcal{G}(1), \mathcal{G}(2)$: indices of group one and two, respectively.
Output: $\widehat{A}^{-1}, \widehat{B}^{-1}, \widehat{A\otimes B}, \widehat{\beta}(\widehat{B}^{-1}), \widehat{\gamma}, T_j$ for all j

1. *Run Algorithm 1.* Use the group centering method to obtain initial estimates $\widehat{\gamma}_j^{\text{init}} = \widehat{\beta}_j^{(1)} - \widehat{\beta}_j^{(2)}$ for all $j = 1,\ldots,m$. Let $\widehat{B}_{\text{init}}^{-1}$ and $\widehat{B}_{\text{init}}$ be as obtained in (14).
2. *Select genes with large estimated differences in means.* Let $\widetilde{J}_0 = \{j : |\widehat{\gamma}_j^{\text{init}}| > 2\widehat{\tau}_{\text{init}}\}$ denote the set of genes which we consider as having strong mean effects, where
   $$\widehat{\tau}_{\text{init}} \asymp \left(\frac{\log^{1/2} m}{\sqrt{m}} + \frac{\|\widehat{B}_{\text{init}}\|_1}{n_{\min}}\right)\sqrt{\frac{n_{\text{ratio}}|\widehat{B}_{\text{init}}^{-1}|_{0,\text{off}}}{n_{\min}}} + \sqrt{\log m}\,\|\left(D^T\widehat{B}_{\text{init}}^{-1}D\right)^{-1}\|_2^{1/2}, \tag{16}$$
   with $n_{\min} = n_1 \wedge n_2$, $n_{\max} = n_1 \vee n_2$, and $n_{\text{ratio}} = n_{\max}/n_{\min}$.
3. *Calculate Gram matrices based on model selection centering.* Global centering can be expressed in terms of the projection matrix $P_1 = n^{-1}\mathbf{1}_n\mathbf{1}_n^T$. Compute the centered data matrix
   $$X_{\text{cen},j} = \begin{cases} X_j - P_2X_j & \text{if } j \in \widetilde{J}_0 \\ X_j - P_1X_j & \text{if } j \in \widetilde{J}_0^c, \end{cases}$$
   where $X_{\text{cen},j}$ denotes the jth column of the centered data matrix X_{cen}. Compute the sample covariance and correlation matrices with X_{cen} following (9) and (7).
4. *Estimate covariances and means.*
   (4a) Obtain the penalized correlation matrices $\widehat{B}_\rho$ and $\widehat{A}_\rho$ using Gemini estimators as defined in (8a) and (8b) with tuning parameters of the same order as those in (23).
   (4b) Obtain inverse covariance estimates $\widehat{B}^{-1}$, $\widehat{A}^{-1}$ using (14).
   (4c) Calculate the GLS estimator $\widehat{\beta}(\widehat{B}^{-1})$ as in (5), as well as the vector of mean differences $\widehat{\gamma}(\widehat{B}^{-1}) = \delta^T\widehat{\beta}(\widehat{B}^{-1})$, for $\delta = (1,-1) \in \mathbb{R}^2$.
5. *Obtain test statistics.* Calculate test statistics as in (15), now using $\widehat{B}^{-1}$ as estimated in Step 4.

Remarks. In the case that γ is sparse, we show that this approach can perform better than the approach in Section 2.2, in particular when the sample size is small. We now consider the expression $\widehat{\tau}_{\text{init}}$ in (16) as an upper bound on the threshold in the sense that it is chosen to tightly control false positives. In Section 4.2 we show in simulations that with this plug-in estimate $\widehat{\tau}_{\text{init}}$, Algorithm 2 can nearly reach the performance of GLS with the true B. Since this choice of $\widehat{\tau}_{\text{init}}$ acts as an order on the threshold we need, the plug-in method can also be applied with a multiplier between 0 and 1. When we set $\widehat{\tau}_{\text{init}}$ at its lower bound, namely, $\sqrt{\log m}\|(D^T\widehat{B}_{\text{init}}^{-1}D)^{-1}\|_2^{1/2}$, where $\widehat{B}_{\text{init}}^{-1}$ is obtained as in Step 3 from Algorithm 1, we anticipate many false positives. In Figure 3, we show that the performance of Algorithm 2 is stable in the setting of small n and sparse γ for different values of $\widehat{\tau}_{\text{init}}$, demonstrating robustness of our methods to the multiplier; there we observe that the performance can degrade if the threshold is set to be too small, eventually reaching the performance of Algorithm 1.

Second, if an upper bound on the number of differentially expressed genes is known a priori, one can select a set of genes

$\check{J}_0$ to group center such that the cardinality $|\check{J}_0|$ is understood to be chosen as an upper bound on $d_0 = |\operatorname{supp}(\gamma)|$ based on prior knowledge. We select the set $\check{J}_0$ by ranking the components of the estimated vector of mean differences $\widehat{\gamma}$. In the data analysis in Section 5 we adopt this strategy in an iterative manner by successively halving the number of selected genes, choosing at each step the genes with largest estimated mean differences from the previous step. We show in this data example and through simulation that the proposed method is robust to the choice of $|\check{J}_0|$.

Finally, it is worth noting that these algorithms readily generalize to settings with more than two groups, in which case we simply group center within each group. This is equivalent to applying the method with a different design matrix D. In fact, we can move beyond group-wise mean comparisons to a regression analysis with a fixed design matrix D, which includes the k-group mean analysis as a special case.

3. Theoretical Results

We first state Theorem 1, which provides the rate of convergence of the GLS estimator (5) when we use a fixed approximation of the covariance matrix B. We then provide in Theorems 3 and 4 the convergence rates for estimating the group mean matrix $\beta \in \mathbb{R}^{2\times m}$ for Algorithms 1 and 2, respectively. In Theorem 3 we state rates of convergence for the Gemini estimators of B^{-1} and A^{-1} when the input sample covariance matrices use the group centering approach as defined in Algorithm 1, while in Theorem 4, we state only the rate of convergence for estimating B^{-1}, anticipating that the rate for A^{-1} can be similarly obtained, using the model selection centering approach as defined in Algorithm 2.

3.1. GLS Under Fixed Covariance Approximation

We now state a theorem on the rate of convergence of the GLS estimator (5), where we use a fixed approximation $B_{n,m}^{-1}$ to B^{-1}, where the operator norm of $\Delta_{n,m} = B_{n,m}^{-1} - B^{-1}$ is small in the sense of (17). We will specialize Theorem 1 to the case where B^{-1} is estimated using the baseline method in Zhou (2014) when X follows subgaussian matrix-variate distribution as in (1). We prove Theorem 1 in Web Supplement Section D.

Theorem 1. Let Z be an $n \times m$ random matrix with independent entries Z_{ij} satisfying $\mathbb{E}Z_{ij} = 0$, $1 = \mathbb{E}Z_{ij}^2 \leqslant \|Z_{ij}\|_{\psi_2} \leqslant K$. Let $Z_1, \ldots, Z_m \in \mathbb{R}^n$ be the columns of Z. Suppose the jth column of the data matrix satisfies $X_j \sim B^{1/2}Z_j$. Suppose $B_{n,m} \in \mathbb{R}^{n\times n}$ is a positive definite symmetric matrix. Let $\Delta_{n,m} := B_{n,m}^{-1} - B^{-1}$. Suppose

$$\|\Delta_{n,m}\|_2 < \frac{1}{(n_{\max}/n_{\min})\,\|B\|_2},$$
$$\text{where} \qquad n_{\min} = n_1 \wedge n_2 \quad \text{and} \quad n_{\max} = n_1 \vee n_2. \tag{17}$$

Then, with probability at least $1 - 8/(m \vee n)^2$, for some absolute constants C, C',

$$\forall j, \quad \|\widehat{\beta}_j(B_{n,m}^{-1}) - \beta_j^*\|_2 \leqslant r_{n,m} := s_{n,m} + t_{n,m}, \quad \text{where} \tag{18}$$

$$s_{n,m} = C\sqrt{\log m \|B\|_2 / n_{\min}} \quad \text{and} \quad t_{n,m} = C'\|\Delta_{n,m}\|_2 / n_{\min}^{1/2}; \tag{19}$$

$$\text{and } \|\widehat{\gamma}(B_{n,m}) - \gamma\|_\infty \leqslant \sqrt{2}\left(C\sqrt{\frac{\log m\|B\|_2}{n_{\min}}} + C' n_{\min}^{-1/2}\|\Delta_{n,m}\|_2\right). \tag{20}$$

Remarks. If the operator norm of B is bounded, that is $\|B\|_2 < W$, then condition (17) is equivalent to $\|\Delta_{n,m}\|_2 < 1/(W n_{\text{ratio}})$. The term $t_{n,m}$ in (19) reflects the error due to approximating B^{-1} with $B_{n,m}^{-1}$, whereas $s_{n,m}$ reflects the error in estimating the mean matrix (5) using GLS with the true B^{-1} for the random design X. The term $s_{n,m}$ is $O(\sqrt{\log m/n})$, whereas $t_{n,m}$ is $O(1/\sqrt{n})$. The dominating term $s_{n,m}$ in (19) can be replaced by the tighter bound, namely, $s'_{n,m} = C' \log^{1/2}(m)\sqrt{\delta^T(D^TB^{-1}D)^{-1}\delta}$, with $\delta = (1, -1) \in \mathbb{R}^2$. This bound correctly drops the factor of $\|B\|_2$ present in (19) and (20), while revealing that variation aligned with the column space of D is especially important in mean estimation.

Note that the condition (17) is not stringent, and that the $\widehat{B}$ estimates used in Algorithms 1 and 2 have much lower errors than this. When $M = 0$ is known, S_A and S_B can be the usual Gram matrices, and the theory in Zhou (2014) guarantees that $t_{n,m}$ as defined in (19) has rate $C_A\sqrt{\log m/m}$, with $C_A = \sqrt{m}\|A\|_F/\operatorname{tr}(A)$. However, in our setting, M in general is nonzero. In Sections 2.2 and 2.3, we provide two constructions for S_A and S_B, which differ in how the data are centered. These constructions have a different bound $t_{n,m}$, as we will discuss in Theorems 3 and 4.

In Section 4, we present simulation results that demonstrate the advantage of the oracle GLS and GLS with estimated $\widehat{B}$ (5) over the sample-mean-based (OLS) method (c.f. (12) and (32)) for mean estimation as well as the related variable selection problem with respect to γ. There, we scrutinize this quantity and its estimation procedure in detail.

Design effect. The "design effect" is the variance of the "oracle" GLS estimator (5) of γ_j using the true B, that is,

$$\delta^T(D^TB^{-1}D)^{-1}\delta = \operatorname{Var}(\widehat{\gamma}_j(B^{-1})), \quad \forall j = 1, \ldots, m. \tag{21}$$

The design effect reflects the potential improvement of GLS over OLS. It appears as a factor above in $s'_{n,m}$, so it contributes to the rate of mean parameter estimation as characterized in Theorem 1. Lower variance in the GLS estimator of the mean difference contributes to greater power of the test statistics relative to OLS. The design effect also appears as a scale factor in the test statistics for $\widehat{\gamma}$ (15), and therefore it is particularly important that the design effect is accurately estimated in order for the test statistics to be properly calibrated. In a study focusing on mean differences, it may be desirable to assess the sample size needed to detect a given effect size using our methodology. Given the design effect, our tests for differential expression are essentially Z-tests based on the GLS fits, followed by some form of multiple comparisons adjustment.

Corollary 2. Let $\Omega = (D^TB^{-1}D)^{-1}$, $\widehat{\Omega} = (D^T\widehat{B}^{-1}D)^{-1}$, and $\Delta = \widehat{\Omega} - \Omega$. Under the conditions of Theorem 1, the relative error in estimating the design effect is bounded as

$$\frac{\left|\delta^T\widehat{\Omega}\delta - \delta^T\Omega\delta\right|}{\delta^T\Omega\delta} \leqslant 2C'\frac{\kappa(B)\,\|B\|_2\,\|\Delta\|_2}{n_{\text{ratio}}}, \tag{22}$$

with probability $1 - C/(m \vee n)^d$, for some absolute constants C, C'.

We prove Corollary 2 in Web Supplement Section D.2. Corollary 2 implies that given an accurate estimator of B^{-1}, the design effect is accurately estimated and therefore suggests that traditional techniques can be used to gain an approximate understanding of the power of our methods. We show that B^{-1} can be accurately estimated under conditions in Theorems 3 and 4. If pilot data are available that are believed to have similar between-sample correlations to the data planned for collection in a future study, Corollary 2 also justifies using this pilot data to estimate the design effect. If no pilot data are available, it is possible to conduct power analyses based on various plausible specifications for the B matrix.

3.2. Rates of Convergence for Algorithms 1 and 2

We state the following assumptions.

(A1) The number of nonzero off-diagonal entries of A^{-1} and B^{-1} satisfy

$$\left|A^{-1}\right|_{0,\text{off}} = o(n/\log(m \vee n)) \quad (n, m \to \infty) \quad \text{and}$$
$$\left|B^{-1}\right|_{0,\text{off}} = o\left(\left[m/\log(m \vee n)\right] \wedge \left[n_{\min}^2/\|B\|_1^2\right]\right) (n, m \to \infty).$$

(A2) The eigenvalues of A and B are bounded away from 0 and $+\infty$. We assume that the stable ranks satisfy $r(A), r(B) \geqslant 4\log(m \vee n)$, where $r(A) = \|A\|_F^2 / \|A\|_2^2$.

Theorem 3. Suppose that (A1) and (A2) hold. Consider the data as generated from model (2) with $\varepsilon = B^{1/2} Z A^{1/2}$, where $A \in \mathbb{R}^{m\times m}$ and $B \in \mathbb{R}^{n\times n}$ are positive definite matrices, and Z is an $n \times m$ random matrix as defined in Theorem 1. Let $C, C', C_1 C_2, C'', C'''$ be some absolute constants. Let $C_A = \sqrt{m}\|A\|_F / \operatorname{tr}(A)$ and $C_B = \sqrt{n}\|B\|_F / \operatorname{tr}(B)$. (I) Let λ_A and λ_B denote the penalty parameters for (8b) and (8a), respectively. Suppose

$$\lambda_A \geqslant C\left(C_A K \frac{\log^{1/2}(m \vee n)}{\sqrt{m}} + \frac{\|B\|_1}{n_{\min}}\right) \quad \text{and}$$
$$\lambda_B \geqslant C'\left(C_B K \frac{\log^{1/2}(m \vee n)}{\sqrt{n}} + \frac{\|B\|_1}{n_{\min}}\right). \tag{23}$$

Then, with probability at least $1 - C''/(m \vee n)^2$, for $\widehat{A \otimes B}$ as define in (11),

$$\|\widehat{A \otimes B} - A \otimes B\|_2 \leqslant \|A\|_2 \|B\|_2 \delta,$$
$$\|\widehat{A \otimes B}^{-1} - A^{-1} \otimes B^{-1}\|_2 \leqslant \|A^{-1}\|_2 \|B^{-1}\|_2 \delta',$$

where

$$\delta, \delta' = O\left(\lambda_A \sqrt{|B^{-1}|_{0,\text{off}} \vee 1} + \lambda_B \sqrt{|A^{-1}|_{0,\text{off}} \vee 1}\right).$$

Furthermore, with probability at least $1 - C'''/(m \vee n)^2$,

$$\|\widehat{A \otimes B} - A \otimes B\|_F \leqslant \|A\|_F \|B\|_F \eta, \tag{24}$$

where

$$\eta = O\left(\lambda_A \sqrt{|B^{-1}|_{0,\text{off}} \vee n}/\sqrt{n} + \lambda_B \sqrt{|A^{-1}|_{0,\text{off}} \vee m}/\sqrt{m}\right). \tag{25}$$

The same conclusions hold for the inverse estimate, with η being bounded in the same order as in (25). (II) Let $\widehat{\beta}$ be defined as in (5) with $\widehat{B}^{-1}$ being defined as in (14) and D as in (4). Then, with probability at least $1 - C/m^d$ the following holds for all j,

$$\|\widehat{\beta}_j(\widehat{B}^{-1}) - \beta_j^*\|_2 \leqslant C_1 \lambda_A \sqrt{\frac{n_{\text{ratio}}\left(|B^{-1}|_{0,\text{off}} \vee 1\right)}{n_{\min}}} + C_2 \sqrt{\log m}\|(D^T B^{-1} D)^{-1}\|_2^{1/2}. \tag{26}$$

We prove Theorem 3 part I in Web Supplement Section E; this relies on rates of convergence of $\widehat{B}^{-1}$ and $\widehat{A}^{-1}$ in the operator and the Frobenius norm, which are established in Lemma S7. We prove part II in Web Supplement Section E.2.

Remarks. We find that the additional complexity of estimating the mean matrix leads to an additional additive term of order $1/n$ appearing in the convergence rates for covariance estimation for B and A. In part I of Theorem 3, λ_A is decomposed into two terms, one term reflecting the variance of S_B, and one term reflecting the bias due to group centering. The variance term goes to zero as m increases, and the bias term goes to zero as n increases. To analyze the error in the GLS estimator based on $\widehat{B}^{-1}$, we decompose $\|\widehat{\beta}_j(\widehat{B}^{-1}) - \beta_j^*\|_2$ as

$$\|\widehat{\beta}_j(\widehat{B}^{-1}) - \beta_j^*\|_2 \leqslant \|\widehat{\beta}_j(\widehat{B}^{-1}) - \widehat{\beta}_j(B^{-1})\|_2 + \|\widehat{\beta}_j(B^{-1}) - \beta_j^*\|_2,$$

where the first term is the error due to not knowing B^{-1}, and the second term is the error due to not knowing β_j^*. The rate of convergence given in (26) reflects this decomposition. For Algorithm 2, we have analogous rates of convergence for both mean and covariance estimation. Simulations suggest that the constants in the rates for Algorithm 2 are smaller than those in (26).

We state the following assumptions for Theorem 4 to hold on Algorithm 2.

(A2') Suppose (A2) holds, and $n = \Omega((\log m)(\|A\|_2 \|B\|_2 b_{\max}/C_A^2))$.

(A3) Let $\operatorname{supp}(\gamma) = \{j : \gamma_j \neq 0\}$. Let $s = |\operatorname{supp}(\gamma)|$ denote the sparsity of γ. Assume that $s = O(\frac{C_A}{\|B\|_2} n \sqrt{\frac{m}{\log m}})$.

Remarks. Condition (A2') is mild, because the condition on the stable rank of B already implies that $n \geqslant \log m$.

Theorem 4. Suppose that (A1'), (A2'), and (A3) hold. Consider the data as generated from model (4) with $\varepsilon = B^{1/2} Z A^{1/2}$, where $A \in \mathbb{R}^{m\times m}$ and $B \in \mathbb{R}^{n\times n}$ are positive-definite matrices, and Z is an $n \times m$ random matrix as defined in Theorem 3. Let λ_A denote the penalty parameter for estimating B. Suppose λ_A is as defined in (23). Let

$$\tau_{\text{init}} \asymp \sqrt{\log m}\|(D^T B^{-1} D)^{-1}\|_2^{1/2}. \tag{27}$$

Then, with probability at least $1 - C''/(m \vee n)^2$, for output of Algorithm 2,

$$\left\| \mathrm{tr}\,(A)\left(\widehat{W}_2\widehat{B}_\rho\widehat{W}_2\right)^{-1} - B^{-1}\right\|_2 \leqslant \frac{C'\lambda_A\sqrt{|B^{-1}|_{0,\mathrm{off}} \vee 1}}{b_{\min}\varphi^2_{\min}(\rho(B))}, \quad \text{and} \tag{28}$$

$$\|\widehat{\beta}_j(\widehat{B}^{-1}) - \beta^*_j\|_2 \leqslant C_2\sqrt{\log m}\|(D^TB^{-1}D)^{-1}\|_2^{1/2}, \tag{29}$$

for all j, for absolute constants C, C_2, C', and C''.

We prove Theorem 4 in Web Supplement Section H.5. In Web Supplement Section H.4, we also show a standalone result, namely Theorem S21, for the case of fixed sets of group and globally centered genes. This result shows how the algorithm used in the preliminary step to choose which genes to group center can be decoupled from the rest of the estimation procedure in Algorithm 2, so long as certain conditions hold. The proof of Theorem 4 indeed validates that such conditions hold for the output of Algorithm 1. It is worth noting that a similar rate of convergence for estimating A could also be derived, but we focus on B in our methodology and applications, and therefore leave this as an exercise for interested readers.

We specialize Corollary 2 to the case where B^{-1} is estimated using Algorithm 2.

Corollary 5. Under the conditions of Theorem 4, we have with probability $1 - C/m^2$

$$\frac{\left|\delta^T\widehat{\Omega}\delta - \delta^T\Omega\delta\right|}{\delta^T\Omega\delta} \leqslant 2C'\frac{n_{\mathrm{ratio}}}{\lambda_{\min}(B)}\kappa(B)\lambda_A\sqrt{|B^{-1}|_{0,\mathrm{off}} \vee 1}, \tag{30}$$

for some absolute constants C and C'.

Remarks. The right-hand side of (30) goes to zero because of the assumptions (A1), (A2'), and (A3), which ensure that the factor $\lambda_A\sqrt{|B^{-1}|_{0,\mathrm{off}} \vee 1}$ goes to zero. We conduct simulations to assess the accuracy of estimating the design effect in Section 4.2.

4. Simulations

We present simulations to compare Algorithms 1 and 2 to both sample-mean-based analysis and oracle algorithms that use knowledge of the true correlation structures A and B. We show these results for a variety of population structures and sample sizes. We construct covariance matrices for A and B from one of:

- AR1(ρ) model. The covariance matrix is of the form $B = \{\rho^{|i-j|}\}_{i,j}$, and the graph corresponding to B^{-1} is a chain.
- Star-Block model. The covariance matrix is block-diagonal with equal-sized blocks whose inverses correspond to star structured graphs, where $B_{ii} = 1$, for all i. In each subgraph, a central hub node connects to all other nodes in the subgraph, with no additional edges. The covariance matrix for each block S in B is generated as in Ravikumar et al. (2011): $S_{ij} = \rho = 0.5$ if $(i, j) \in E$ and $S_{ij} = \rho^2$ otherwise.
- Erdős–Rényi model. We use the random concentration matrix model in Zhou, Lafferty, and Wasserman (2010). The graph is generated according to a type of Erdős–Rényi random graph. Initially we set $B^{-1} = 0.25I_{n\times n}$. Then, we randomly select d edges and update B^{-1} as follows: for each new edge (i, j), a weight $w > 0$ is chosen uniformly at random from $[w_{\min}, w_{\max}]$ where $w_{\min} = 0.6$ and $w_{\max} = 0.8$; we subtract w from B^{-1}_{ij} and B^{-1}_{ji}, and increase B^{-1}_{ii} and B^{-1}_{jj} by w. This keeps B^{-1} positive definite. We then rescale so that B^{-1} is an inverse correlation matrix.

Table 1. Assessment of the difficulty of estimating B^{-1} and the potential gain from GLS. The total correlation ρ_B is the average squared off-diagonal value of the correlation matrix $\rho(B)$. The fourth column is the design effect as defined in (21). The last column (sd ratio) presents the ratio of the standard deviation of the difference in sample means in (12) to the standard deviation of the GLS estimator of the difference in means. The first three columns of the table reflect the difficulty of estimating B, whereas the last two columns reflect the potential improvement of GLS over the sample-mean-based method (12). In the notation StarBlock(a, b), a refers to the number of blocks, and b refers to the block size.

	B	ρ_B^2	$\|B\|_F/\mathrm{tr}(B)$	$\|\rho(B)^{-1}\|_{1,\mathrm{off}}$	sd GLS	sd ratio
$n = 80$						
1	AR1(0.2)	0.00	0.12	32.92	0.27	1.00
2	AR1(0.4)	0.00	0.13	75.24	0.33	1.02
3	AR1(0.6)	0.01	0.16	148.12	0.40	1.07
4	AR1(0.8)	0.04	0.24	351.11	0.46	1.32
5	StarBlock(4, 20)	0.02	0.18	101.33	0.35	1.51
6	ER(0.6, 0.8)	0.01	0.14	92.75	0.17	1.21
$n = 40$						
1	AR1(0.2)	0.00	0.16	16.25	0.38	1.01
2	AR1(0.4)	0.01	0.19	37.14	0.45	1.03
3	AR1(0.6)	0.03	0.23	73.12	0.53	1.12
4	AR1(0.8)	0.08	0.33	173.33	0.53	1.47
5	StarBlock(2, 20)	0.04	0.25	50.67	0.50	1.51
6	ER(0.6, 0.8)	0.02	0.21	47.24	0.25	1.23

4.1. Accuracy of $\widehat{\gamma}$ and its Implication for Variable Ranking

Table 1 displays metrics that reflect how the choice of different population structures B can affect the difficulty of the mean and covariance estimation problems. Column 2 is a measure discussed by Efron (2007). Column 3 appears directly in the theoretical analysis, reflecting the entry-wise error in the sample correlation $\widehat{\Gamma}(B)$. Columns 4 analogously reflects the entry-wise error for the Flip-Flop procedure in Zhou (2014), and is included here for completeness. Column 5 displays the value of $\sqrt{\delta^T(D^TB^{-1}D)^{-1}\delta}$, where $\delta = (1, -1) \in \mathbb{R}^2$, which represents the standard deviation of the difference in means estimated using GLS with the true B^{-1}. Column 6 displays what we call the standard deviation ratio, namely

$$\sqrt{\frac{u^TBu}{\delta^T(D^TB^{-1}D)^{-1}\delta}}, \tag{31}$$

where $u = (\underbrace{1/n_1, \ldots, 1/n_1}_{n_1}, \underbrace{-1/n_2, \ldots, -1/n_2}_{n_2}) \in \mathbb{R}^n$ and $\delta = (1, -1) \in \mathbb{R}^2$, which reflects the potential efficiency gain for GLS over sample-mean-based method (12) for estimating γ. Note that the standard deviation ratio depends on the relationship between the covariance matrix B and the design matrix D. In Table 1, the first $n/2$ individuals are in group one, and the following $n/2$ are in group two. The values in Column 6 show that substantial improvement is possible in mean estimation. For an AR1 covariance matrix, the standard deviation ratio increases as the AR1 parameter increases; as the correlations get stronger, the potential improvement in mean estimation due to GLS grows. For the Star Block model with fixed block size, the standard deviation ratio is stable as n increases.

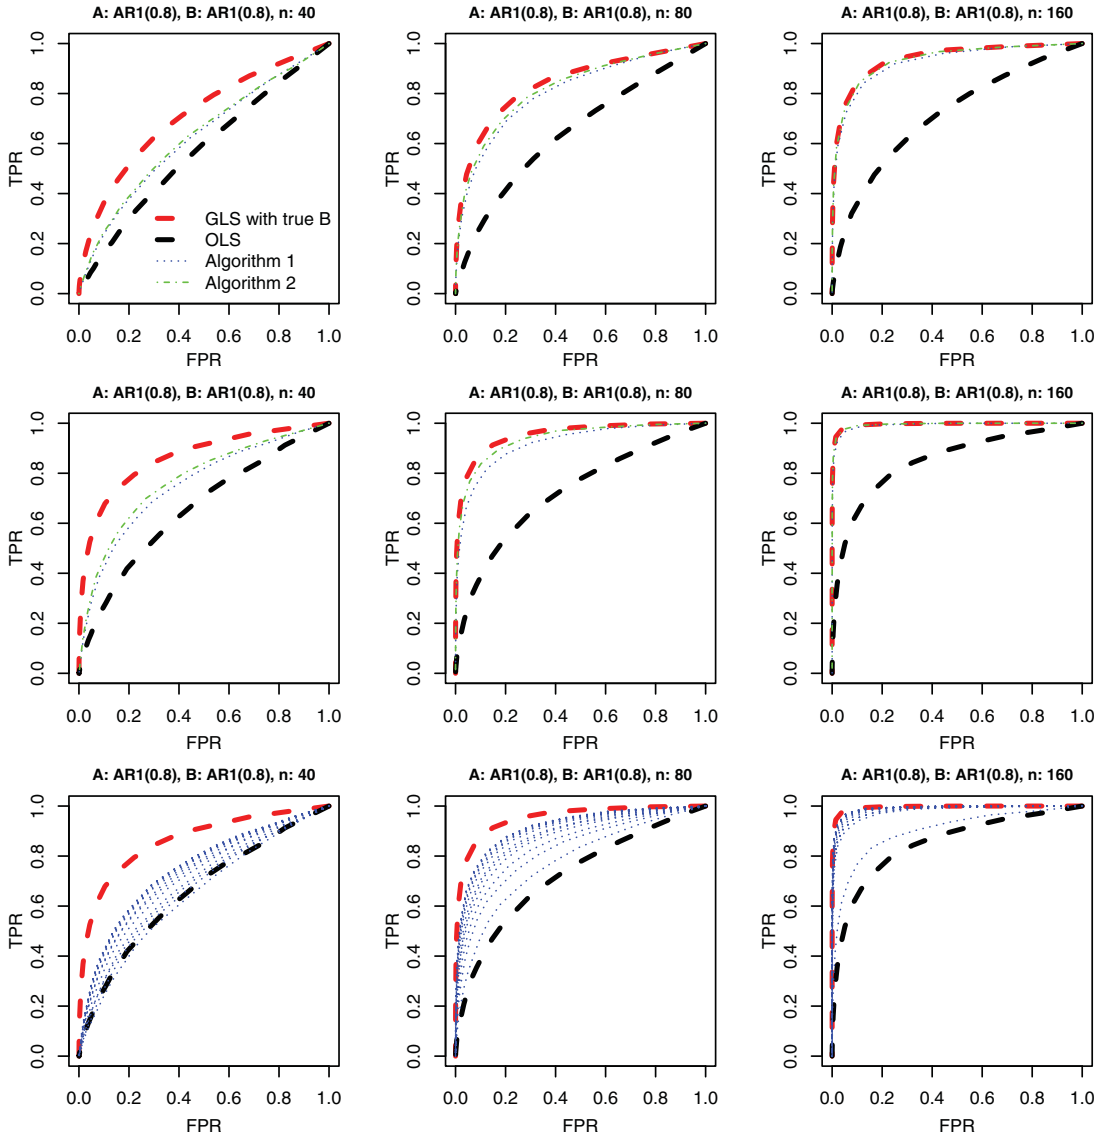

Figure 1. ROC curves. For each plot, the horizontal axis is false positive rate (FPR) and the vertical axis is true positive rate (TPR), as we vary a threshold for classifying variables as null or nonnull. The covariance matrices A and B are both AR1 with parameter 0.8, with $m = 2000$ and $n = 40$, 80, and 160 in column one, two, and three, respectively. Ten variables in γ have nonzero entries. On each trial, the group labels are randomly assigned, with equal sample sizes. The marginal variance of each entry of the data matrix is equal to 1. For the first row of plots, the magnitude of each nonzero entry of γ is 0.2, and for the second and third rows of plots, the magnitude of each nonzero entry of γ is 0.3. In the first two rows, we display ROC curves for Algorithms 1 and 2 with penalty parameters chosen to maximize area under the curve. The third row displays an ROC curve for Algorithm 1, sweeping out penalty parameters.

In Figure 1, we use ROC curves to illustrate the sensitivity and specificity for variable selection in the sense of how well we can identify the support for $\{i : \gamma_i \neq 0\}$ when we threshold $\widehat{\gamma}_i$ at various values. To evaluate and compare different methods, we let $\widehat{\gamma}$ be the output of Algorithm 1, Algorithm 2, the oracle GLS, and the sample-mean-based method (12). These correspond to the four curves on each plot of the top two rows of plots. We find that Algorithm 1 and Algorithm 2 perform better than the sample-mean-based method (12), and in some cases perform comparably to the oracle GLS. Plots in the third row of Figure 1 illustrate the sensitivity of Algorithm 1 to the choice of the graphical lasso (GLasso) penalty parameter (23); the simulations are run using the `glasso` R package (Friedman, Hastie, and Tibshirani 2008) to estimate B via (8b). The performance can degenerate to that of the sample-mean-based method (12), if the penalty is too high.

In the top row of Figure 2, we plot the root-mean-squared error (RMSE) when estimating the mean differences γ for Algorithm 1, Algorithm 2, OLS (i.e., sample means) and the oracle GLS estimate. The population structures for B are Erdős-Rényi and Star Block. Both Algorithms 1 and 2 consistently outperform the sample-mean-based method (12) for mean estimation, and Algorithm 2 even achieves comparable performance to the oracle GLS in some settings. The bottom row displays the relative Frobenius error for estimating B^{-1}. Algorithm 2 outperforms Algorithm 1 in terms of covariance estimation and is comparable to oracle model selection, which only centers the columns with a true mean difference.

In Figure 3, we illustrate that Algorithm 2 can perform well using a plug-in estimator $\widehat{\tau}_{\text{init}}$ as in (16). We compare the methods when the true mean structure is a decaying exponential; we display the correlation of the ranks of the entries of γ to the ranks of the estimates of γ. Algorithm 2 with a plugin estimator $\widehat{\tau}_{\text{init}}$ can nearly reach the performance of GLS with the true B. Furthermore, the plug-in version of Algorithm 2 also consistently outperforms Algorithm 1. We also assess sensitivity to the choice of threshold: the curve labeled "Algorithm 2" uses

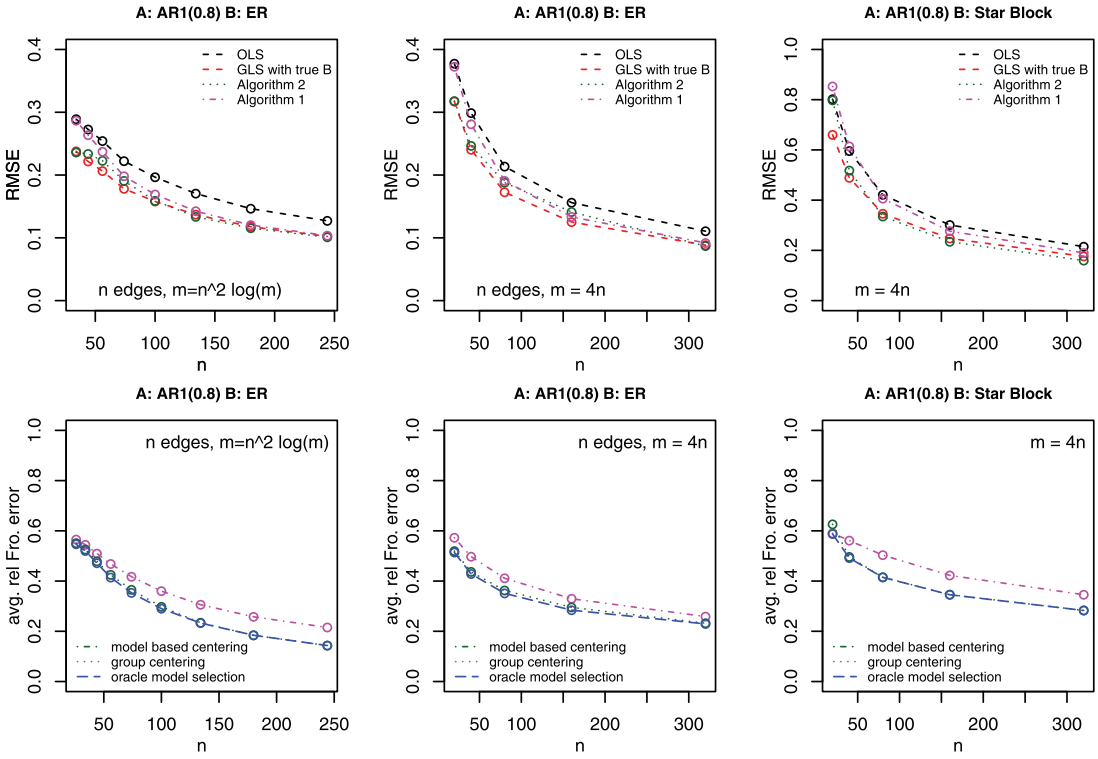

Figure 2. Performance of centering methods as n and m are varied, with n shown on the horizontal axis. In the first column of plots, the number of edges is proportional to $\sqrt{m/\log(m)}$. In the second and third columns of plots, the number of edges is proportional to m. In the first two columns of plots, B^{-1} is an Erdős-Rényi inverse covariance matrix. In the third column, B^{-1} is star block with blocks of size 10. The first row of plots shows RMSE for estimating γ, whereas the second row shows average relative Frobenius error in estimating B^{-1}. All panels are based on 250 simulation replications.

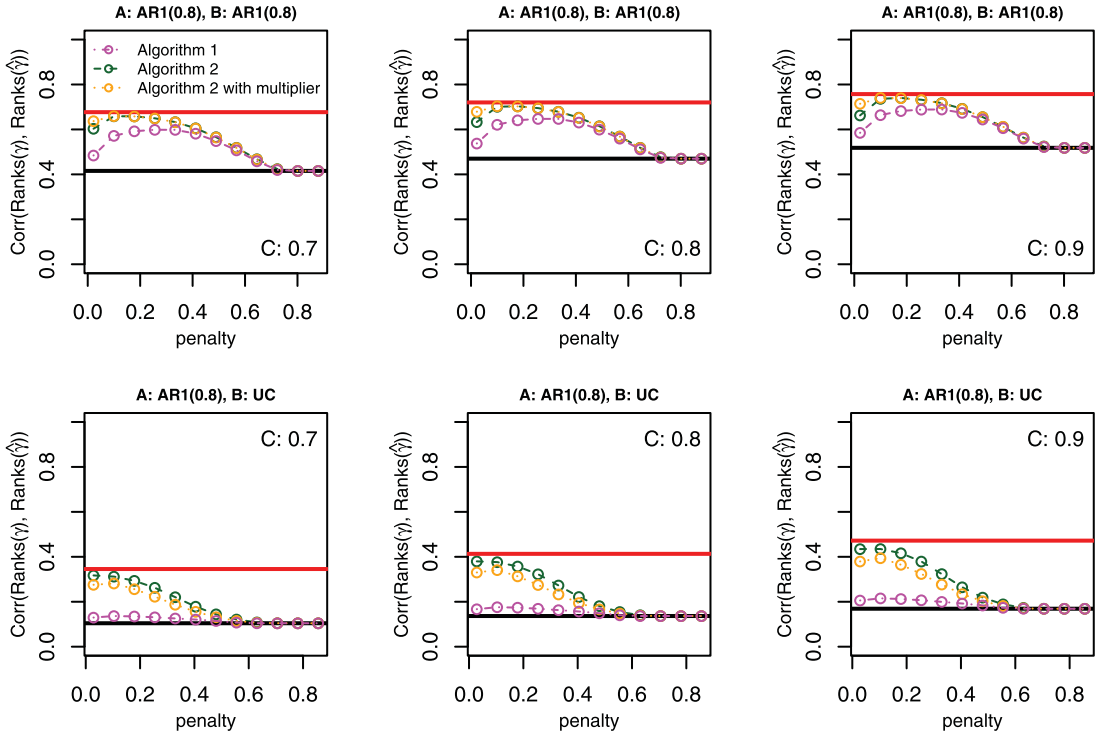

Figure 3. This figure displays the correlation between the rankings of the components of γ and $\hat{\gamma}$, sorted by magnitude, denoted Corr(Ranks(γ), Ranks($\hat{\gamma}$)) in the axis label. The vector of mean differences is chosen as $\gamma_j = C\exp(-(3/2000)j)$, for $j = 1, \ldots, 2000$. We also present the Algorithm 2 results with a multiplier on the threshold as described in Section 2.3. In the top row, the true B is AR1(0.8), with $n = 40$ and $m = 2000$. In the bottom row, the true B is chosen as an estimate from the UC data, with $n = 20$ and $m = 2000$. For the top row, the group labels are randomly assigned; for the bottom row, the first ten rows of the data are in group one, and the other ten are in group two. The figure is averaged over 200 replications. The top and bottom horizontal lines represent GLS with true B and OLS, respectively. The vertical axis displays the correlation of ranks between $\hat{\gamma}$ and γ, and the horizontal axis displays the GLasso penalty parameter.

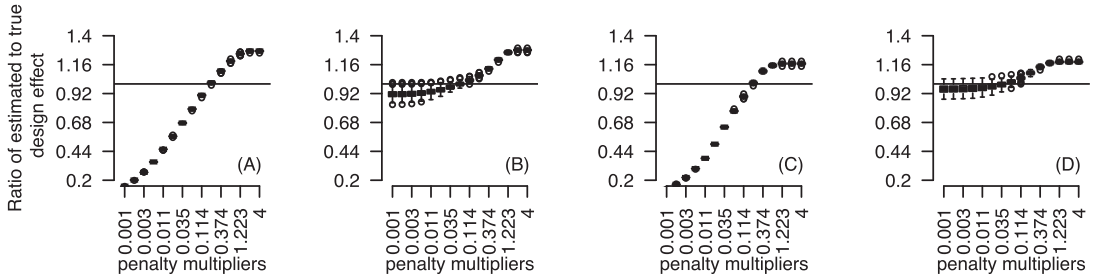

Figure 4. Ratio of estimated design effect to true design effect when B^{-1} is Erdős-Rényi, and A is AR1(0.8). Figures (A) and (B) correspond to sample size $n = 80$; (C) and (D) correspond to $n = 40$. Figures (A) and (C) correspond to Algorithm 1; Figures (B) and (D) correspond to Algorithm 2, with 10 columns group centered. These results are based on dimension parameter $m = 2000$ and 250 simulation replications.

the plug-in estimate $\widehat{\tau}_{\text{init}}$, whereas "Algorithm 2 with threshold multiplier" uses a plug-in estimate of the lower bound given in (27) in Theorem 4. These two-plug in estimators exhibit similar performance, showing robustness of Algorithm 2 to the choice of the threshold parameter. In real data analysis, we validate this further. For the top row (AR1), the ratio of thresholds (27) to (16) is 0.75, and for the bottom row (UC), the ratio is 0.17.

In Web Supplement Section J, we perform additional simulations to compare Algorithm 2 to two similar methods using ROC curves, namely, the sphering method of Allen and Tibshirani (2012), which uses a matrix-variate model similar to ours, and the confounder adjustment method of Wang et al. (2017), which uses a latent factor model. Our simulations show that Algorithm 2 consistently outperforms these competing methods in a variety of simulation settings using matrix-variate data.

4.2. Inference for the Mean Difference $\widehat{\gamma}$

Two basic approaches to conducting inference for mean differences are paired and unpaired t statistics. The unpaired t statistic is defined as follows. Let $X = (X_{ij})$. Then the jth unpaired t statistic is

$$T_j = \left(\widetilde{\beta}_j^{(1)} - \widetilde{\beta}_j^{(2)}\right) \widehat{\sigma}_j^{-1}(n_1^{-1} + n_2^{-1})^{-1/2}, \text{ where}$$
$$\widehat{\sigma}_j^2 = (n_1 + n_2 - 2)^{-1} \sum_{k=1}^{2} \sum_{i \in \mathcal{G}_k} \left(X_{ij} - \widetilde{\beta}_j^{(k)}\right)^2, \tag{32}$$

where $\widetilde{\beta}_j^{(k)}$, $k = 1, 2$, and $j = 1, \ldots, m$, denotes the sample mean of group k and variable j as defined in (12), and $\mathcal{G}_k$ is the set of indices corresponding to group k. When there is a natural basis for pairing the observations, and paired units are anticipated to be positively correlated, we can calculate paired t statistics. For the paired t statistic, suppose observations i and $i' = i + n/2$ are paired, for $i \in \{1, \ldots, n/2\}$. Note that samples can always be permuted so as to be paired in this way. Define the paired differences $d_{ij} = X_{ij} - X_{i'j}$, for $i \in \{1, \ldots, n/2\}$. Then the paired t statistic is $\overline{d}_j(n/2 - 1)^{1/2}/(\sum_{i=1}^{n/2}(d_{ij} - \overline{d}_j)^2)^{1/2}$, where $\overline{d}_j = (n/2)^{-1}\sum_{i=1}^{n/2} d_{ij}$.

Figure 4 considers estimation of the "design effect" $\delta^T(D^T B^{-1} D)^{-1}\delta$, as previously defined in (21), with $\delta = (1, -1)^T$. The importance of this object is discussed in Sections 3.1 and 3.2. The design effect is estimated via $\delta^T(D^T \widehat{B}^{-1} D)^{-1}\delta$, with $\widehat{B}^{-1}$ from Algorithm 1 or 2. The GLasso penalty parameters are chosen as

$$\lambda_A = f_A \left(C_A K \frac{\log^{1/2}(m \vee n)}{\sqrt{m}} + \frac{\|B\|_1}{n_{\min}} \right) \tag{33}$$

where we sweep over the factor f_A, referred to as the penalty multiplier. Figure 4 displays boxplots of the ratio $\delta^T(D^T \widehat{B}^{-1} D)^{-1}\delta / \delta^T(D^T B^{-1} D)^{-1}\delta$ over 250 replications for each setting of the penalty multiplier f_A. In Figure 4, B^{-1} follows the Erdős–Rényi model, and A is AR1(0.8), with $m = 2000$, and $n = 40$ and 80. Figure 4 shows that Algorithm 2 (plots B and D) estimates the design effect to high accuracy and is quite insensitive to the penalty multiplier as long as it is less than 1, as predicted by the theoretical analysis. Algorithm 1 also estimates the design effect with high accuracy, but with somewhat greater sensitivity to the tuning parameter. The best penalty parameter for Algorithm 1 is around 0.1, whereas reasonable penalty parameters for Algorithm 2 are in the range 0.01 to 0.1. This is consistent

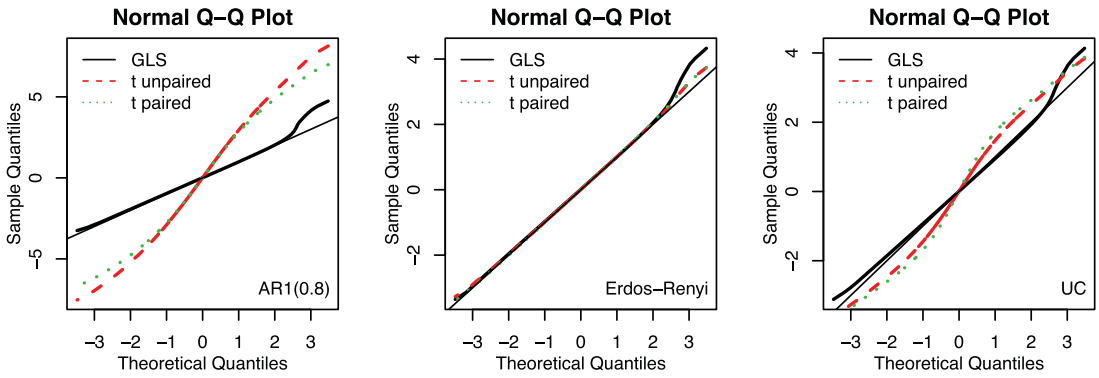

Figure 5. Quantile plots of test statistics. Ten genes have nonzero mean differences equal to 2, 0.8, and 1 in the three plots, respectively. In each plot A is AR1(0.8). Covariance structures for B are as indicated. In the third plot, the true B is set to $\widehat{B}$ for the ulcerative colitis data, described in Section 5. For the first two plots there are $n = 40$ samples and $m = 2000$ variables. For the third plot there are $n = 20$ samples and $m = 2000$ variables. Each plot has 250 simulation replications.

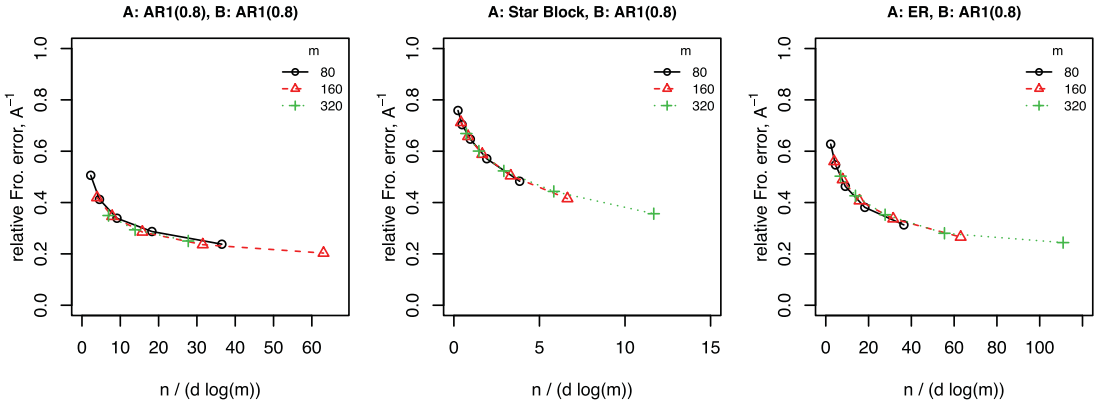

Figure 6. Relative Frobenius error in estimating A^{-1}, as n varies. In each plot the matrix B is AR1(0.8) and A is as indicated. The vertical axis is relative Frobenius error, and the horizontal axis $n/(d\log(m))$, where d is the maximum node degree. The GLasso penalty is chosen to minimize the relative Frobenius error. Each point is based on 250 Monte Carlo replications.

with smaller entrywise error in the sample covariance for model selection centering than for group centering.

We next compare the results from Algorithm 2 to results obtained using paired and unpaired t statistics. Figure 5 illustrates the calibration and power of plug-in Z-scores, $\widehat{\gamma}_j/\widehat{\mathrm{SE}}(\widehat{\gamma}_j)$ derived from Algorithm 2 for three population settings. The standard error is calculated as $\sqrt{\delta^T(D^T\widehat{B}^{-1}D)^{-1}\delta}$, with $\delta = (1, -1)$. In the first and second plots, the data was simulated from AR1(0.8) and Erdős-Rényi, respectively. In the third plot, the data was simulated from $\widehat{B}$ for ulcerative colitis data described in Section 5. To obtain $\widehat{B}$, we apply Algorithm 2 to the ulcerative colitis data, using a Glasso penalty of $\lambda \approx 0.5[(\log(m)/m) + 3/n]$ in Step 1, followed by group centering the top ten genes in Step 2, and using a Glasso penalty of $\lambda \approx 0.1[(\log(m)/m) + 3/n]$ in Step 4. In all cases A is AR1(0.8). In each case, we introduce 10 variables with different population means in the two groups, by setting $\gamma = 0.8$ for those variables, with the remaining γ values equal to zero. The ideal Q-Q plot would follow the diagonal except at the upper end of the range, as do our plug-in GLS test statistics. The t statistics (ignoring dependence) are seen to be overly dispersed throughout the range, and are less sensitive to the real effects.

4.3. Covariance Estimation for A

Figure 6 shows the relative Frobenius error in estimating A^{-1} as n grows, for fixed m. The horizontal axis is $n/(d\log(m))$, scaled so that the curves align, where d is the maximum node degree. Because $\|A^{-1}\|_F$ is of order $\sqrt{m}$, the vertical axis essentially displays $\|\widehat{A}^{-1} - A^{-1}\|_F/\sqrt{m}$. For estimating A^{-1}, the rate of convergence is of order $\sqrt{\log(m)/n}$. For each of the three population structures, accuracy increases with respect to n.

5. Genomic Study of Ulcerative Colitis

Ulcerative colitis (UC) is a chronic form of inflammatory bowel disease (IBD), resulting from inappropriate immune cell infiltration of the colon. As part of an effort to better understand the molecular pathology of UC, Lepage et al. (2011) reported on a study of mRNA expression in biopsy samples of the colon mucosal epithelium, with the aim of being able to identify gene transcripts that are differentially expressed between people with UC and healthy controls. The study subjects were discordant identical twins, that is, monozygotic twins such that one twin has UC and the other does not. This allows us to simultaneously explore dependences among samples (both within and between twins), dependences among genes, and mean differences between the UC and non-UC subjects. The dataset is available on the Gene Expression Omnibus, GEO accession GDS4519 (Edgar, Domrachev, and Lash 2002).

The data consist of 10 discordant twin pairs, for a total of 20 subjects. Each subject's biopsy sample was assayed for mRNA expression, using the Affymetrix UG 133 Plus 2.0 array, which has 54,675 distinct transcripts. Previous analyses of this data did not consider twin correlations or unanticipated non-twin correlations, and used very different methodology (e.g., Wilcoxon testing). Roughly 70 genes were found to be differentially expressed (Lepage et al. 2011).

We applied our Algorithm 2 to the UC genomics data as follows. First we selected the 2000 most variable genes based on marginal variance and then rescaled each gene to have unit marginal variance. We then applied Step 1 of Algorithm 2, setting $\lambda = 0.1 \approx 0.5(\sqrt{\frac{\log(m)}{m}} + \frac{3}{n})$, with $m = 2000$ and $n = 20$. For Step 2 of the algorithm, we ranked the estimated mean differences, group centered the top ten, and globally centered the remaining genes. We then re-calculated the Gram matrix S_B using the centered data. In Step 3, following the Gemini approach, we applied the GLasso to S_B using a regularization parameter $\lambda \approx 0.25(\sqrt{\log(m)/m} + 3/n)$. We obtain estimated differences in means and test statistics via Steps 4 through 6.

A natural analysis of these data using more standard methods would be a paired t-test for each mRNA transcript (paired by twin pair). Such an approach is optimized for the situation where there is a constant level of correlation within all of the twin pairs, with no non twin correlations. However as in Efron (2009), we wish to accommodate unexpected correlations, which in this case would be correlations between non-twin subjects or a lack of correlation between twin subjects. Our approach, developed in Section 2, does not require pre-specification or parameterization of the dependence structure, thus we were able to consider twin and non-twin correlations simultaneously. Lepage et al. noted that UC has lower heritability than other forms of IBD. If UC has a relatively stronger environmental component, this

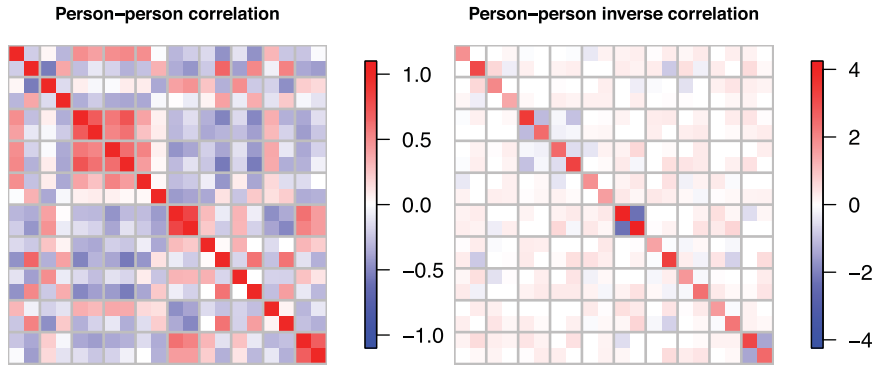

Figure 7. Estimated person-person correlation matrix and its inverse, estimated using the 2000 genes with largest marginal variance.

could explain the pattern of correlations that we uncovered, as shown in Figure 7. The samples are ordered so that twins are adjacent, corresponding to 2-by-2 diagonal blocks. The penalized inverse sample correlation matrix contains nonzero entries both within twin pairs and between twin pairs.

To also handle these unexpected nontwin correlations, we performed testing using Algorithm 2. We found only a small amount of evidence for differential gene expression between the UC and non-UC subjects. Four of the adjusted p-values fell below a threshold of 0.1, using the Benjamini-Hochberg adjustment; that is, four genes satisfied $2000\widehat{p}_{(i)}/i < 0.1$, where $\widehat{p}_{(i)}$ is the ith order statistic of the p-values calculated using Algorithm 2, for $i = 1, \ldots, 2000$. Based on our theoretical and simulation work showing that our procedure can successfully recover and accommodate dependence among samples, we argue that this is a more meaningful representation of the evidence in the data for differential expression compared to methods that do not adapt to dependence among samples. Specifically, in Section 5.1 we demonstrate that our test statistics are properly calibrated and as a result have weaker (but more accurate) evidence for differential expression results. Below we argue that the sample-wise correlations detected by our approach would be expected to artificially inflate the evidence for differential expression.

5.1. Calibration of Test Statistics

As noted above, based on the test statistics produced by Algorithm 2, we find evidence for only a small number of genes being differentially expressed. This conclusion, however, depends on the test statistics conforming to the claimed null distribution whenever the group-wise means are equal. In this section, we consider this issue in more detail.

The first plot of Figure 8 compares the empirical quantiles of $\Phi^{-1}(T_j)$ to the corresponding quantiles of a standard normal distribution, where Φ is the standard normal cdf and the T_j's are as defined in (32). Plots 2 and 3 show the same information for successive non overlapping blocks of two thousand genes sorted by marginal variance. Since this is a discordant twins study, we also show results for the standard paired t statistics, pairing by twin. In all cases, the paired and unpaired statistics are more dispersed relative to the reference distribution. By contrast, the central portion of the GLS test statistics coincide with the reference line. Overdispersion of test statistics throughout their range is often taken to be evidence of miscalibration (Devlin and Roeder 1999). In this setting, the GLS statistics are calibrated correctly under the null hypothesis, but the paired and unpaired t statistics are not.

5.2. Stability of Gene Sets

The motivation of our Algorithm 2 is that in many practical settings a relatively small fraction of variables may have differential means, and therefore it is advantageous to avoid centering variables presenting no evidence of a strong mean difference. Here, we assess the stability of the estimated mean differences as we vary the number of group centered genes in Algorithm 2. To do so, we successively group center fewer genes, globally centering the remaining genes.

The iterative process is as follows. Let $\widehat{B}^{-1}_{(i)} \in \mathbb{R}^{n\times n}$ denote the estimate of B^{-1} at iteration i, let $\widehat{\beta}_{(i)} \in \mathbb{R}^{2\times m}$ denote the estimates of the group means β on the ith iteration, let $\widehat{\gamma}_{(i)} \in \mathbb{R}^m$

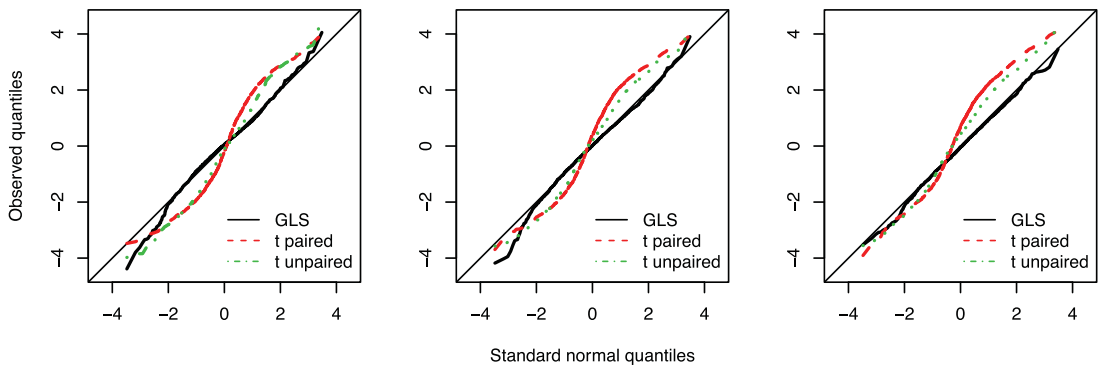

Figure 8. Quantile plots of test statistics for three disjoint gene sets, each consisting of 2000 genes. The genes are partitioned based on marginal variance. GLS statistics are taken from Step 5 of Algorithm 2; in Step 2, the 10 genes with greatest mean differences are selected for group centering.

Table 2. Each iteration k of the algorithm produces a ranking of all 2000 genes. For the top ten genes on each iteration, entry (i, j) of the table shows the number of genes in common in iterations i and j of the algorithm. Note that the maximum possible value for any entry of the table is 10; if entry (i, j) is 10, then iterations i and j selected the same top ten genes.

	1	2	3	4	5	6	7	8	9
1	10	10	7	5	5	3	3	3	3
2	10	10	7	5	5	3	3	3	3
3	7	7	10	6	5	3	3	3	3
4	5	5	6	10	8	5	5	5	5
5	5	5	5	8	10	7	7	7	7
6	3	3	3	5	7	10	10	10	10
7	3	3	3	5	7	10	10	10	10
8	3	3	3	5	7	10	10	10	10
9	3	3	3	5	7	10	10	10	10

denote the vector of differences in group means between the two groups, and let $\widehat{\mu}_{(i)} \in \mathbb{R}^m$ denote vector of global mean estimates. Let $\widehat{\mu}(B^{-1}) \in \mathbb{R}^m$ denote the result of applying GLS with design matrix $D = 1_n$ to estimate the global means.

Initialize $\widehat{\beta}_{(1)}$, $\widehat{\mu}_{(1)}$ and $\widehat{\gamma}_{(1)}$ using the sample means. On the ith iteration,

1. Rank the genes according to $|\widehat{\gamma}_{(i-1)}|$. Center the highest ranked n_i' genes around $\widehat{\beta}_{(i-1)}$. Center the remaining genes around $\widehat{\mu}_{(i-1)}$.
2. Obtain $\widehat{B}^{-1}_{(i)}$ by applying GLasso to the centered data matrix from Step 1.
3. Set $\widehat{\beta}_{(i)} = \widehat{\beta}(\widehat{B}^{-1}_{(i)})$, $\widehat{\mu}_{(i)} = \widehat{\mu}(\widehat{B}^{-1}_{(i)})$, and $\widehat{\gamma}_{(i)} = (1, -1)\widehat{\beta}_{(i)}$.

We assess the stability of the mean estimates by comparing the rankings of the genes across iterations of the algorithm. Table 2 displays the number of genes in common out of the top ten genes on each pair of iterations of the algorithm. For example, three genes ranked in the top ten on the first iteration of the algorithm are also ranked in the top ten on the last iteration. Iterations six through nine produce the same ranking of the top ten genes. Three genes are ranked among the top ten on every iteration of the algorithm: DPP10-AS1, OLFM4, and PTN. Web Supplement Table S1 shows simulations confirming these results.

5.3. Stability Analysis

Table 3 shows the number of genes that fall below an FDR threshold of 0.1 on each iteration, for several values of the GLasso penalty λ. The number of genes below the threshold is more sensitive to the number of group-centered genes than to the GLasso penalty parameter. This is consistent with the first plot of Web Supplement Figure S2a where the design effect (in the denominator of the test statistics) is likewise more sensitive to the number of group centered genes than to the GLasso penalty. When fewer than 128 genes are group centered, the number of genes below an FDR threshold of 0.1 is stable across the penalty parameters from $\lambda = 0.1$ to $\lambda = 0.8$.

Table 3. For the algorithm, this table shows the number of genes that are significant at an FDR level of 0.1 on each iteration of the algorithm, for different values of the GLasso penalty λ. The top row shows the number of genes group centered on each iteration.

n.group	2000	1024	512	256	128	64	32	16	8
$\lambda = 0.1$	1006	913	327	14	3	1	1	1	1
$\lambda = 0.2$	865	806	262	2	1	1	1	1	0
$\lambda = 0.3$	778	789	303	3	1	1	0	0	0
$\lambda = 0.4$	706	774	452	3	1	0	0	0	0
$\lambda = 0.6$	657	751	587	19	1	1	0	0	0
$\lambda = 0.8$	628	699	493	30	1	1	1	1	1

6. Conclusion

It has long been known that heteroscedasticity and dependence between observations impacts the precision and degree of uncertainty for estimates of mean values and regression coefficients. Further, data that are modeled for convenience as being independent observations may in fact show unanticipated dependence (Kruskal 1988). This has motivated the development of numerous statistical methods, including generalized/weighted least squares (GLS/WLS), mixed effect models, and generalized estimating equations (GEE). Our approach utilizes recent advances in high dimensional statistics to permit estimation of an inter-observation dependence structure (reflected in the matrix B in our model). Like GLS/GEE, we use an approach that alternates between mean and covariance estimation, but limit it in Algorithm 1 to a mean estimation step, followed by a covariance update, followed by a mean update, with an additional covariance and mean update if Algorithm 2 is used. We provide convergence guarantees and rates for both algorithms.

Estimation of dependence or covariance structures usually requires some form of replication, and/or strong models. We require a relatively weak form of replication and a relatively weak model. In our framework, the dependence among observations must be common (up to proportionality) across a set of "quasi-replicates" (the columns of X, or the genes in our UC example). These quasi-replicates may be statistically dependent, and may have different means. We also require the precision matrices for the dependence structures to be sparse, which is a commonly used condition in recent high-dimensional analyses.

In addition to providing theoretical guarantees, we also show through simulations and a genomic data analysis that the approach improves estimation accuracy for the mean structure, and appears to mitigate test statistic overdispersion, leading to test statistics that do not require post-hoc correction. The latter observation suggests that undetected dependence among observations may be one reason that genomic analyses are sometimes less reproducible than traditional statistical methods would suggest, an observation made previously by Efron (2009) and others.

Although our theoretical analysis guarantees the convergence of our procedure even with a single observation of the random matrix X, there are reasons to expect this estimation problem to be fundamentally challenging. One reason for this as pointed out by Efron (2009) and subsequently explored by Zhou (2014), is that the row-wise and column-wise dependence structures are somewhat non-orthogonal, in that row-dependence can "leak" into the estimates of column-wise dependence, and vice-versa. Our results suggest that while row-wise correlations make it more difficult to estimate column-wise correlations (and vice-versa), when the emphasis is on mean structure estimation, even a somewhat rough estimate of the dependence structure (B) can substantially improve estimation and inference.

Acknowledgments

The research is supported in part by NSF under Grant DMS-1316731 and the Elizabeth Caroline Crosby Research Award from the Advance Program at the University of Michigan. The authors thank the Editor, the Associate editor and three referees for their constructive comments that led to improvements in the article. Roger Fan, Kerby Shedden, and Shuheng Zhou are listed alphabetically. This work was done while all authors were affiliated with the Department of Statistics at the University of Michigan.

Funding

Division of Mathematical Sciences [NSF13-16731]; University of Michigan [Elizabeth C. Crosby Research to S. Zhou].

References

Aitken, A. C. (1936), "IV.–On Least Squares and Linear Combination of Observations)," *Proceedings of the Royal Society of Edinburgh*, 55, 42–48. [682]

Allen, G. I., and Tibshirani, R. (2012), "Inference with Transposable Data: Modelling the Effects of Row and Column Correlations," *Journal of the Royal Statistical Society*, Series B, 74, 721–743. [682,683,692]

Bai, Z., and Saranadasa, H. (1996), "Effect of High Dimension: By an Example of a Two Sample Problem," *Statistica Sinica*, 6, 311–329. [684]

Banerjee, O., Ghaoui, L. E., and d'Aspremont, A. (2008), "Model Selection Through Sparse Maximum Likelihood Estimation for Multivariate Gaussian or Binary Data," *Journal of Machine Learning Research*, 9, 485–516. [684]

Benjamini, Y., and Hochberg, Y. (1995), "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing," *Journal of the Royal Statistical Society*, Series B, 57, 289–300. [684]

Benjamini, Y., and Yekutieli, D. (2001), "The Control of the False Discovery Rate in Multiple Testing Under Dependency," *Annals of Statistics*, 29, 1165–1188. [683]

Cai, T., Jessie Jeng, X., and Jin, J. (2011), "Optimal Detection of Heterogeneous and Heteroscedastic Mixtures," *Journal of the Royal Statistical Society*, Series B, 73, 629–662. [683]

Cai, T. T., Li, H., Liu, W., and Xie, J. (2016), "Joint Estimation of Multiple High-Dimensional Precision Matrices," *Statistica Sinica*, 26, 445–464. [684]

Cai, T. T., and Xia, Y. (2014), "High-Dimensional Sparse Manova," *Journal of Multivariate Analysis*, 131, 174–196. [684]

Chen, S. X., Li, J., and Zhong, P. S. (2014), "Two-Sample Tests for High Dimensional Means with Thresholding and Data Transformation," arXiv:1410.2848. [684]

Chen, S. X., and Qin, Y.-L. (2010), "A Two-Sample Test for High-Dimensional Data with Applications to Gene-Set Testing," *The Annals of Statistics*, 38, 808–835. [684]

Devlin, B., and Roeder, K. (1999), "Genomic Control for Association Studies," *Biometrics*, 55, 997–1004. [682,694]

Edgar, R., Domrachev, M., and Lash, A. E. (2002), "Gene Expression Omnibus: Ncbi Gene Expression and Hybridization Array Data Repository," *Nucleic Acids Research*, 30, 207–210. [693]

Efron, B. (2005), "Large-Scale Simultaneous Hypothesis Testing: The Choice of a Null Hypothesis," *Journal of the American Statistical Association*, 99, 96–104. [682]

——— (2007), "Correlation and Large-Scale Simultaneous Significance Testing," *Journal of the American Statistical Association*, 93–103. [682,689]

——— (2009), "Are a Set of Microarrays Independent of Each Other?" *The Annals of Applied Statistics*, 3, 922–942. [682,683,684,693,695]

——— (2010), *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*, vol. 1., New York: Cambridge University Press. [682]

Fan, J., Feng, Y., and Wu, Y. (2009), "Network Exploration Via the Adaptive LASSO and SCAD Penalties," *The Annals of Applied Statistics*, 3, 521–541. [684]

Friedman, J., Hastie, T., and Tibshirani, R. (2008), "Sparse Inverse Covariance Estimation with the Graphical Lasso," *Biostatistics*, 9, 432–441. [684,693]

Hall, P., and Jin, J. (2010), "Innovated Higher Criticism for Detecting Sparse Signals in Correlated Noise," *The Annals of Statistics*, 38, 1686–1732. [684]

Hoff, P. (2011), "Separable Covariance Arrays Via the Tucker Product, with Applications to Multivariate Relational Data," *Bayesian Analysis*, 6, 179–196. [684]

Kruskal, W. (1988), "Miracles and Statistics: The Casual Assumption of Independence," *Journal of the American Statistical Association*, 83, 929–940. [695]

Lam, C., and Fan, J. (2009), "Sparsistency and Rates of Convergence in Large Covariance Matrices Estimation," *Annals of Statistics*, 37, 4254–4278. [684]

Leek, J. T., Scharpf, R. B., Bravo, H. C., Simcha, D., Langmead, B., Johnson, W. E., Geman, D., Baggerly, K., and Irizarry, R. A. (2010), "Tackling the Widespread and Critical Impact of Batch Effects in High-Throughput Data," *Nature Reviews Genetics*, 11, 733–739. [682]

Lepage, P., Häsler, R., Spehlmann, M. E., Rehman, A., Zvirbliene, A., Begun, A., Ott, S., Kupcinskas, L., Dore, J., Raedler, A., and Schreiber, S. (2011), "Twin Study Indicates Loss of Interaction Between Microbiota and Mucosa of Patients with Ulcerative Colitis," *Gastroenterology*, 141, 227–236. [693]

Li, J., and Zhong, P.-S. (2014), "A Rate Optimal Procedure for Sparse Signal Recovery Under Dependence," *arXiv preprint arXiv:1410.2839*. [683]

Meinshausen, N., and Bühlmann, P. (2006), "High Dimensional Graphs and Variable Selection with the Lasso," *Annals of Statistics*, 34, 1436–1462. [684]

Owen, A. B. (2005), "Variance of the Number of False Discoveries," *Journal of the Royal Statistical Society*, Series B, 67, 411–426. [683]

Peng, J., Wang, P., Zhou, N., and Zhu, J. (2009), "Partial Correlation Estimation by Joint Sparse Regression Models," *Journal of the American Statistical Association*, 104, 735–746. [684]

Ravikumar, P., Wainwright, M., Raskutti, G., and Yu, B. (2011), "High-Dimensional Covariance Estimation By Minimizing ℓ_1-Penalized Log-Determinant Divergence," *Electronic Journal of Statistics*, 4, 935–980. [684]

Rothman, A., Bickel, P., Levina, E., and Zhu, J. (2008), "Sparse Permutation Invariant Covariance Estimation," *Electronic Journal of Statistics*, 2, 494–515. [684]

Storey, J. D. (2003), "The Positive False Discovery Rate: A Bayesian Interpretation and the q-Value," *Annals of Statistics* 2013–2035. [684]

Sugden, L. A., Tackett, M. R., Savva, Y. A., Thompson, W. A., and Lawrence, C. E. (2013), "Assessing the Validity and Reproducibility of Genome-Scale Predictions," *Bioinformatics*, 29, 2844–2851. [682]

Sun, Y., Zhang, N. R., and Owen, A. B. (2012), "Multiple Hypothesis Testing Adjusted for Latent Variables, with an Application to the Agemap Gene Expression Data," *The Annals of Applied Statistics*, 6, 1664–1688. [683]

Tan, K. M., and Witten, D. M. (2014), "Sparse Biclustering of Transposable Data," *Journal of Computational and Graphical Statistics*, 23, 985–1008. [684]

Wang, J., Zhao, Q., Hastie, T., and Owen, A. B. (2017), "Confounder Adjustment in Multiple Hypothesis Testing," *The Annals of Statistics*, 45, 1863–1894. [683,692]

Yuan, M., and Lin, Y. (2007), "Model Selection and Estimation in the Gaussian Graphical Model," *Biometrika*, 94, 19–35. [684]

Zhou, S. (2014), "Gemini: Graph Estimation with Matrix Variate Normal Instances," *Annals of Statistics*, 42, 532–562. [682,683,684,685,687,689,695]

Zhou, S., Lafferty, J., and Wasserman, L. (2010), "Time Varying Undirected Graphs," *Machine Learning*, 80, 295–319. [684]