Published in final edited form as:

Adv Drug Deliv Rev. 2013 June 30; 65(7): 987–1000. doi:10.1016/j.addr.2013.04.008.

Statistical Analysis of Big Data on Pharmacogenomics

Jianqing Fan and Han Liu

Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA

Abstract

This paper discusses statistical methods for estimating complex correlation structure from large pharmacogenomic datasets. We selectively review several prominent statistical methods for estimating large covariance matrix for understanding correlation structure, inverse covariance matrix for network modeling, large-scale simultaneous tests for selecting significantly differently expressed genes and proteins and genetic markers for complex diseases, and high dimensional variable selection for identifying important molecules for understanding molecule mechanisms in pharmacogenomics. Their applications to gene network estimation and biomarker selection are used to illustrate the methodological power. Several new challenges of Big data analysis, including complex data distribution, missing data, measurement error, spurious correlation, endogeneity, and the need for robust statistical methods, are also discussed.

Keywords

Big data; High dimensional statistics; Approximate factor model; Graphical model; Multiple testing; Variable selection; Marginal screening; Robust statistics

1. Introduction

The ultimate goal of pharmacogenomics is to improve health care based on individual genomic profiles [1, 2, 3, 4, 5, 6, 7]. Together with other factors that may affect drug response — such as diet, age, diseases, lifestyle, environment, and state of health — pharmacogenomics has the potential to facilitate the creation of individualized treatment plan for patients and leads to the overarching goal of personalized medicine. Similar to other areas of human genomics, pharmacogenomics is experiencing an explosion of data, specifically by overloads of omics information (genomes, transcriptomes and other data from cells, tissues and drug effects)[8]. The pharmacogenomics research is entering the era of “Big data” — a term that refers to the explosion of available information. The vast amount of data brings both opportunities and new challenges. Statistical analysis for Big data is becoming increasingly important [9, 10, 11, 12].

This paper selectively overviews several state of the art statistical methods for analyzing large pharmacogenomics data. In particular, we emphasize on the fundamental problems of estimating two types of correlation structures: *marginal* correlation and *conditional* correlation. The former represents the correlation between two variables by ignoring the

© 2013 Elsevier B.V. All rights reserved.

Corresponding: jqfan@princeton.edu, Telephone number: (609)2583255, Fax number: (609)2588551.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

remaining variables, while the latter represents the correlation between two variables by conditioning on the remaining ones. The marginal correlation can be estimated using the covariance matrix. The conditional correlation can be estimated using the inverse covariance matrix. We introduce several recently developed statistical methods for estimating high dimensional covariance and inverse covariance matrices. We also introduce cutting-edge methods to estimate false discovery proportions for large-scale simultaneous tests, and to select important molecules or SNPs in high dimensional regression models. The former corresponds to finding molecules or SNPs that have significant marginal correlations with biological outcomes, while the latter finds conditional correlations with biomedical responses in presence of many other molecules or SNPs.

The rationale behind this paper is that we believe Big data provides new opportunities for estimating complex correlation structures among a large number of variables. These methods are new to the pharmacogenomics community and have the potential to play important roles in analyzing the next-generation Big data within the pharmacogenomics area. In the following we list some applications that motivate the development of new statistical methods for estimating large covariance matrices.

In functional genomics, an important problem is to cluster genes into different groups based on the similarities of their microarray expression profiles. One popular measure of the similarity between a pair of genes is the correlation of their expression profiles. Thus, if d genes are being analyzed (with d ranges from the order of $\sim 1,000$ to $\sim 10,000$), a correlation matrix of size $d \times d$ needs to be estimated. Note that $1,000 \times 1,000$ covariance matrices involve already over half a million elements. Yet, the sample size n is of order ~ 100 , which is significantly smaller than the dimensionality d . Thus, the sample covariance matrix degenerates and needs regularization.

As another example, many multivariate statistical methods that are useful in pharmacogenomics, including principal component analysis (PCA), linear discriminant analysis (LDA), multivariate regression analysis, the Hotelling T^2 -statistic, and multivariate likelihood ratio tests as in finding quantitative trait loci based on longitudinal data [13] request the estimation of the covariance matrix as their first step. Thus a reliable estimate of the covariance matrix is of paramount importance.

Moreover, the correlation matrix itself can be used to construct co-expression gene network, which is an undirected graph whose edges indicate marginal correlations among the d genes. The network is built by drawing edges between those pairs of genes whose magnitude of pairwise correlation coefficients exceed a certain threshold. More applications of large covariance matrix estimation will be discussed in §7.

A notable feature of most methods introduced in this paper is the exploitation of sparsity assumption, which is an essential concept for modern statistical methods applied to high dimensional data. For covariance estimation, we briefly introduce the thresholding approach [14] and its extension called POET (Principal Orthogonal complement Thresholding) [15, 16] which provides a unified view of most previous methods. For inverse covariance estimation, we mainly focus on introducing two inverse covariance estimation methods named CLIME [17] and TIGER [18], which stand respectively for “Constrained L_1 -Minimization for Inverse Matrix Estimation” and “Tuning-Insensitive Graph Estimation and Regression”. Both methods estimate the inverse covariance matrix in a column-by-column fashion and achieve the minimax optimal rates of convergence under different norms. For applications, we discuss how the estimated covariance matrix and inverse covariance matrix can be exploited on gene network estimation and large-scale simultaneous tests. We introduce the principal factor approximation (PFA) in [19] for estimating and controlling

false discovery proportions in large-scale simultaneous tests that are dependent. In addition, we introduce penalized likelihood [20], variable screening [21], and their iterated version of screening and selection for high dimensional variable selection. All these problems are important and widely applicable in different subareas of pharmacogenomics. Besides surveying existing methods, we also discuss new challenges and possible solutions of Big data analysis. Topics include modeling nonGaussianity, handling missing data, dealing with measurement error, spurious correlation, endogeneity, and developing robust methods.

The rest of this paper is organized as follows. In §2, we summarize some notations. In §3, we introduce statistical methods for estimating large covariance matrices. In §4, we introduce statistical methods for estimating large inverse covariance matrices. In §5, we apply the results of §3 and §4 to large-scale multiple tests. In §6, we introduce high dimensional variable selection. In §7, we provide case studies of several described methods. In §8, we discuss more on the future directions.

2. Notation

In this section, we summarize some notations. Let $\mathbf{A} = (a_{jk}) \in \mathbb{R}^{d \times d}$, $\mathbf{B} = (b_{jk}) \in \mathbb{R}^{d \times d}$ and $\mathbf{v} = (v_1, \dots, v_d)^T \in \mathbb{R}^d$. Denote by $\lambda_{\min}(\mathbf{A})$ and $\lambda_{\max}(\mathbf{A})$ the smallest and largest eigenvalues of $\mathbf{A}$. The inner product of $\mathbf{A}$ and $\mathbf{B}$ is defined as $\langle \mathbf{A}, \mathbf{B} \rangle = \text{tr}(\mathbf{A}^T \mathbf{B})$. Define the vector

norms: $\|\mathbf{v}\|_1 = \sum_j |v_j|$, $\|\mathbf{v}\|_2^2 = \sum_j v_j^2$, $\|\mathbf{v}\|_\infty = \max_j |v_j|$. We also define matrix operator and

elementwise norms: $\|\mathbf{A}\|_2^2 = \lambda_{\max}(\mathbf{A}^T \mathbf{A})$, $\|\mathbf{A}\|_F^2 = \sum_{j,k} a_{jk}^2$. Notation $\mathbf{A} > \mathbf{0}$ means that $\mathbf{A}$ is positive definite and $a_n \asymp b_n$ implies there are positive constants c_1 and c_2 independent of n such that $c_1 b_n \leq a_n \leq c_2 b_n$.

3. Estimating Large Covariance Matrix

Estimating a large covariance or correlation matrix under a small sample size is a fundamental problem which has many applications in pharmacogenomics. This problem can be formulated as follows: Let $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^d$ be n independent observations of a d -dimensional random vector $\mathbf{X} = (X_1, \dots, X_d)^T$. Without loss of generality, we assume $\mathbb{E}\mathbf{X} = \mathbf{0}$. We want to find a reliable estimate of the population covariance matrix $\Sigma = \mathbb{E}\mathbf{X}\mathbf{X}^T$. At the first sight, this problem does not seem to be challenging. In the literature, Σ was traditionally estimated by the sample covariance matrix

$$\mathbf{S} = \frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i^T - \bar{\mathbf{x}}^T) \quad \text{with} \quad \bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i, \quad (3.1)$$

which has many good theoretical properties when the dimensionality d is small. However, in the more realistic settings where the dimensionality d is comparable or even larger than n (i.e., d/n goes to a nonzero constant or infinity), the sample covariance matrix in (3.1) is no longer a good estimate of the population covariance matrix Σ . More details will be explained as follows.

3.1. Inconsistency of Sample Covariance in High Dimensions

We use a simple simulation to illustrate the inconsistency of the sample covariance matrix in high dimensions. Specifically, we sample n data points from a d -dimensional spherical Gaussian distribution $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \Sigma)$ with $\Sigma = \mathbf{I}_d$. Here $\mathbf{I}_d$ is a d -dimensional identity matrix. We consider different setups including $d/n = 2$, $d/n = 1$, $d/n = 0.2$, and $d/n = 0.1$. The results are summarized in Figure 1.

Figure 1 shows the sorted eigenvalues of the sample covariance matrix $\mathbf{S}$ (black curve) with the true eigenvalues (dashed red curve) for $d = 1,000$. By examining these plots, we see that when the dimensionality d is large, the eigenvalues of $\mathbf{S}$ significantly deviate from their true values. In fact, even when n is reasonably large compared with d ($d/n = 0.1$), the result is still not accurate.

This phenomenon can be characterized by random matrix theory. Let $d/n \rightarrow \gamma$ with $\gamma \in (0, 1)$ and $\lambda_{\max}(\mathbf{S})$ be the largest eigenvalue of $\mathbf{S}$. It has been shown that $\lambda_{\max}(\mathbf{S})$ converges to $(1 + \sqrt{\gamma})^2$ almost surely [22, 23], i.e.,

$$\mathbb{P}\left(\lim_{n \rightarrow \infty} \lambda_{\max}(\mathbf{S}) = (1 + \sqrt{\gamma})^2\right) = 1.$$

This result illustrates that, for large d , the largest eigenvalue of the sample covariance matrix $\mathbf{S}$ does not converge to that of the population covariance matrix $\mathbf{\Sigma}$.

3.2. Sparse Covariance Matrix Estimation Methods

To handle the inconsistency issue of high dimensional sample covariance matrix, most existing methods assume the population covariance matrix is sparse, i.e., many off-diagonal entries of $\mathbf{\Sigma}$ are zero. Such a sparsity assumption is reasonable in many pharmacogenomics applications. For example, in a longitudinal study where variables have a natural order, it is natural to assume that variables are weakly correlated when they are far apart [24]. Under the sparsity assumption, many regularization based statistical methods have been proposed to estimate large covariance matrix [25, 26, 27, 28, 29, 30, 31, 32, 33, 34]. These methods usually only require elementwise thresholding procedures and are computationally efficient. The simplest thresholding estimator is the hard-thresholding estimator [35, 14]

$$\widehat{\mathbf{\Sigma}} = (s_{ij} I(|s_{ij}| \geq \tau_{ij})), \quad (3.2)$$

where $I(\cdot)$ is the indicator function and s_{ij} is the sample covariance between X_i and X_j . Here τ_{ij} is a thresholding parameter. Another example is the adaptive thresholding [29] which takes $\tau_{ij} = \text{SD}(s_{ij})\tau$. Here

$$\text{SD}(s_{ij}) = \sqrt{\frac{\sum_{k=1}^n (x_{ki}x_{kj} - s_{ij})^2}{n}}$$

is the estimated standard error of s_{ij} and τ is a user-specified parameter (e.g., $\sqrt{(2\log d)/n}$). [16] suggests a simplified version $\tau_{ij} = \sqrt{s_{ii}s_{jj}}\tau$ so that the correlation is thresholded at level τ (e.g., $\tau = 0.2$). Estimator (3.2) is not necessarily positive definite. However, when τ is sufficiently large, it is positive definite with high probability.

One additional example is the soft-thresholding estimator [16]:

$$\widehat{\mathbf{\Sigma}} = (\text{sgn}(s_{ij})(|s_{ij}| - \tau_{ij})_+), \quad (3.3)$$

where $(a)_+ = \max\{0, a\}$ and $\text{sgn}(s_{ij}) = \mathbb{I}(s_{ij} > 0) - \mathbb{I}(s_{ij} < 0)$. It makes the matrix (3.3) positive definite for a wider range of τ_{ij} than the hard-thresholding estimator (3.2) [16]. Although estimators (3.2) and (3.3) suffice for many applications, more general sparse covariance estimation can be found via the following penalized least-squares [36]:

$$\widehat{\Sigma} = \underset{\Sigma=(\sigma_{jk})}{\text{argmin}} \left\{ \frac{1}{2} \|\mathbf{S} - \Sigma\|_F^2 + \sum_{j \neq k} P_{\lambda, \gamma}(\sigma_{jk}) \right\}, \quad (3.4)$$

where the term $\|\mathbf{S} - \Sigma\|_F^2$ measures the goodness of fit and $\sum_{j \neq k} P_{\lambda, \gamma}(\sigma_{jk})$ is a sparsity-inducing penalty that encourages sparsity [20]. The tuning parameter λ controls the bias-variance tradeoff and γ is a possible fine-tune parameter which controls the concavity of the penalty. Popular choices of the penalty function $P_{\lambda, \gamma}(\cdot)$ include the hard-thresholding, soft-thresholding [37, 38], smoothly clipped absolute deviation (SCAD, [20]), and minimax concavity penalties (MCP, [39]). In the following we explain these penalties in more detail.

Penalized least-squares or more generally penalized likelihood is a generally applicable method for variable selection [20, 40]. We introduce it in the context of sparse covariance matrix estimation. First, it is easy to see that the optimization problem in (3.4) decomposes into many univariate subproblems:

$$\widehat{\sigma}_{jk} = \underset{u}{\text{argmin}} \left\{ \frac{1}{2} |s_{jk} - u|^2 + P_{\lambda, \gamma}(u) \right\}, \quad (3.5)$$

whose solution has been carefully studied in [41]. Each can be analytically solved for the following commonly used penalty functions:

1. Complexity penalty: $P_{\lambda, \gamma}(u) = \lambda \mathbb{I}(u \neq 0)$.
2. Hard-thresholding penalty: $p_{\lambda}(u) = \lambda^2 - (\lambda - |u|)_+^2$.
3. L_1 -penalty: $P_{\lambda, \gamma}(u) = \lambda |u|$.
4. Smoothly clipped absolute deviation penalty (SCAD): with $\gamma > 2$,

$$P_{\lambda, \gamma}(u) = \begin{cases} \lambda |u| & \text{if } |u| \leq \lambda \\ \frac{u^2 - 2\gamma\lambda|u| + \lambda^2}{2(1-\gamma)} & \text{if } \lambda < |u| \leq \gamma\lambda \\ \frac{(\gamma+1)\lambda^2}{2} & \text{if } |u| > \gamma\lambda \end{cases}.$$

5. Minimax concavity penalty (MCP): With $\gamma > 1$,

$$P_{\lambda, \gamma}(u) = \begin{cases} |u| - \frac{u^2}{2\lambda\gamma} & \text{if } |u| \leq \lambda\gamma \\ \frac{\lambda^2\gamma}{2} & \text{if } |u| > \lambda\gamma \end{cases}.$$

Note that for problem (3.5) both the complexity penalty and the hard-thresholding penalty produce the same solution, which is the hard-thresholding rule. Figure 2 visualizes these penalty functions for $\lambda = 1$. All penalty functions are folded concave, but L_1 -penalty is also convex. The parameter γ in SCAD and MCP controls the degree of concavity. From Figure 2(e) and Figure 2(g), we see that a smaller value of γ results in more concave penalties. When γ becomes larger, the SCAD and MCP converge to the L_1 -penalty. MCP is a generalization of the hard-thresholding penalty which corresponds to $\gamma = 1$.

With the aforementioned thresholding penalties, the closed-form solution to the penalized least-squares (3.5) can be found explicitly. The complexity and hard-thresholding penalties give hard-thresholding rule (3.2), and the L_1 -penalty yields the soft-thresholding rule (3.3). The SCAD thresholding rule is given by

$$\widehat{\sigma}_{jk} = \begin{cases} 0 & \text{if } |s_{jk}| \leq \lambda \\ \frac{\text{sgn}(s_{jk})(|s_{jk}| - \lambda)}{\frac{(\gamma-1)s_{jk} - \text{sgn}(s_{jk})\gamma\lambda}{\gamma-2}} & \text{if } \lambda < |s_{jk}| \leq 2\lambda \\ s_{jk} & \text{if } 2\lambda < |s_{jk}| \leq \gamma\lambda \\ s_{jk} & \text{if } |s_{jk}| > \gamma\lambda \end{cases},$$

and the MCP thresholding operator is

$$\widehat{\sigma}_{jk} = \begin{cases} 0 & \text{if } |s_{jk}| \leq \lambda \\ \text{sgn}(s_{jk}) \left(\frac{|s_{jk}| - \lambda}{1 - 1/\gamma} \right) & \text{if } \lambda < |s_{jk}| \leq \gamma\lambda \\ s_{jk} & \text{if } |s_{jk}| > \gamma\lambda \end{cases}.$$

When $\gamma \rightarrow \infty$, the last two operators become the soft-thresholding operator.

How shall we choose among these thresholding operators in applications? We provide rough insights and suggestions on this. From Figure 2(b), we see that the hard-thresholding operator does not introduce extra bias for the large nonzero entries. However, it is highly discontinuous. Unlike the hard-thresholding operator, the soft-thresholding operator in Figure 2(d) is continuous. However, it introduces biases for the nonzero entries. From Figure 2(f) and Figure 2(h), we see that the SCAD and MCP thresholding operators are both continuous and do not introduce extra bias for nonzero entries with large absolute values. In applications, we always recommend to use either SCAD or MCP thresholding since they combine the advantages of both hard- and soft-thresholding operators. All these thresholding based methods apply the corresponding thresholding operator on each entry of the sample covariance matrix. Their computing times scale linearly with the problem size. Thus they are suitable for very large-scale problems.

All the described methods in this section have tuning parameters (e.g., the λ in (3.5)). These tuning parameters can be selected either by cross-validation or a resampling scheme proposed by [27].

3.3. Positive Definite Sparse Covariance Matrix Estimation

The above thresholding methods, though simple, do not guarantee the positive definiteness of the estimated covariance matrix. This may cause trouble in downstream analysis such as evaluating the predictive likelihood or linear discriminant analysis. To handle this problem, we introduce a covariance estimation method named EC2 (Estimation of Covariance with Eigenvalue Constraints) which explicitly enforces the positive definiteness constraint [42].

EC2 solves the following constrained optimization

$$\widehat{\Sigma} = \underset{\lambda_{\min}(\widehat{\Sigma}) \geq \tau}{\text{argmin}} \left\{ \frac{1}{2} \|\mathbf{S} - \widehat{\Sigma}\|_F^2 + \sum_{j \neq k} P_{\lambda, \gamma}(\sigma_{jk}) \right\}, \quad (3.6)$$

where $\tau > 0$ is a pre-specified tuning parameter that controls the smallest eigenvalue of the estimated covariance matrix $\hat{\Sigma}$. Compared with (3.4), the optimization problem in (3.6) is not decomposable due to the constraint $\lambda_{\min}(\Sigma) \geq \tau$. [42] proves that the optimization problem in (3.6) is convex when the penalty function is convex and develops an efficient ISP (Iterative Soft-thresholding and Projection) algorithm to solve it. Thorough numerical benchmarks have been conducted to show that the ISP algorithm is more efficient than most competing methods especially when the problem scale is large, more details on computing time can be found in [42]. A similar algorithm is also proposed in [43].

Other sparse covariance matrix estimation methods with positive definiteness constraints include [28] and [44]. Both use a log-determinant function to enforce the positive definiteness of the estimated covariance matrix. Their main difference is that [44] adopts a convex least square formulation, while [28] adopts the penalized Gaussian log-likelihood approach.

3.4. Theory of Sparse Covariance Matrix Estimation

Under the sparsity assumption, the above sparse covariance estimators have good theoretical properties. In particular, for the positive definite sparse covariance matrix estimator $\hat{\Sigma}$ defined in (3.6), it has been proven by various authors (For more details, see [42]) that

$$\sup_{\Sigma \in \mathcal{M}_k} \mathbb{E} \|\hat{\Sigma} - \Sigma\|_2 \asymp k \sqrt{\frac{\log d}{n}},$$

where $\mathcal{M}_k$ represents the model class

$$\mathcal{M}_k = \left\{ \Sigma : \Sigma \succ \mathbf{0} \text{ and } \max_{1 \leq j \leq d} \sum_{\ell=1}^d I(\Sigma_{j\ell} \neq 0) \leq k \right\}. \quad (3.7)$$

This rate of convergence is minimax optimal. Similar results also hold for different types of thresholding estimators [36] defined in (3.4).

Using the triangle inequality, we get

$$\left| \mathbb{E} \lambda_{\max}(\hat{\Sigma}) - \lambda_{\max}(\Sigma) \right| \leq \mathbb{E} \|\hat{\Sigma} - \Sigma\|_2 \asymp k \sqrt{\frac{\log d}{n}}.$$

This result implies that $\lim_{n \rightarrow \infty} \mathbb{E} \lambda_{\max}(\hat{\Sigma}) = \lambda_{\max}(\Sigma)$. Thus the inconsistency issue of the sample covariance matrix discussed in §3.1 is avoided.

3.5. POET: New Insight on Large Covariance Matrix Estimation

All the aforementioned methods assume that Σ is sparse. Though this assumption is reasonable for many pharmacogenomics applications, it is not always appropriate. For example, all the genes from the same pathway may be co-regulated by a small amount of regulatory factors, which makes the gene expression data highly correlated; when genes are stimulated by cytokines, their expressions are also highly correlated. The sparsity assumption is obviously unrealistic in these situations. To solve this problem, we introduce

the POET method in recent paper with discussion [16], which provides an integrated framework for combining latent factor analysis and sparse covariance matrix estimation.

3.5.1. The POET Method—The POET estimator is formed by directly running the singular value decomposition on the sample covariance matrix $\mathbf{S}$. It keeps the covariance matrix formed by the first K principal components and applies the thresholding procedure to the residual covariance matrix. The final covariance matrix estimate is then obtained by combining these two components.

Let $\hat{\lambda}_1, \hat{\lambda}_2, \dots, \hat{\lambda}_d$ be the ordered eigenvalues of $\mathbf{S}$ and $\hat{\xi}_1, \dots, \hat{\xi}_d$ be their corresponding eigenvectors. Then $\mathbf{S}$ has the following spectral decomposition:

$$\mathbf{S} = \sum_{m=1}^K \hat{\lambda}_m \hat{\xi}_m \hat{\xi}_m^T + \hat{\mathbf{R}}_K \quad (3.8)$$

where the first K principal components $\{\hat{\xi}_m\}_{m=1}^K$ estimate the latent factors that drive the common dependence. We now apply sparse thresholding on $\hat{\mathbf{R}}_K$ by solving

$$\hat{\mathbf{R}}_K^l = \underset{\Sigma}{\operatorname{argmin}} \left\{ \frac{1}{2} \|\hat{\mathbf{R}}_K - \Sigma\|_F^2 + \sum_{j \neq k} P_{\lambda, \gamma}(\sigma_{jk}) \right\}. \quad (3.9)$$

For example, one can simply apply the soft-thresholding (3.3) to $\hat{\mathbf{R}}_K$ to obtain a $\hat{\mathbf{R}}_K^l$ in (3.9).

The POET estimator of Σ is defined as:

$$\hat{\Sigma}_K = \sum_{m=1}^K \hat{\lambda}_m \hat{\xi}_m \hat{\xi}_m^T + \hat{\mathbf{R}}_K^l. \quad (3.10)$$

It is a nonparametric estimator and can be positive definite when λ is large [16]. When $K > 0$ and we use the thresholding method in §3.2 with $\tau_{ij} = \sqrt{\hat{r}_{K,ii} \hat{r}_{K,jj}}$, namely, $\tau = 1$, where $\hat{r}_{K,ii}$ is a diagonal element of $\hat{\mathbf{R}}_K$. In this case, $\hat{\mathbf{R}}_K^l = \operatorname{diag}(\hat{\mathbf{R}}_K)$, and our nonparametric estimator $\hat{\Sigma}_K$ is positive definite even when $d \gg n$. It is called the estimator based on the strict factor model [26]. The POET method is available in the R-package named POET. When $K = 0$, the POET estimator reduces to the thresholding estimator in §3.2. From the computational perspective, the main cost of POET is to calculate the eigenvectors of the sample covariance matrix, which can be efficiently calculated by conducting singular value decomposition on the original data matrix and achieves a time complexity $O(n^2 d)$. Since $d \gg n$, the computing time of POET scales linearly with the dimensionality d . Thus the POET can be used to analyze very high dimensional data.

3.5.2. Approximate Factor Model—POET exploits an approximate factor structure [15, 16] and works the best under such a model. The approximate factor model assumes

$$\mathbf{X} = \mathbf{B}\mathbf{U} + \boldsymbol{\varepsilon} \quad (3.11)$$

where $\mathbf{B}$ is a $d \times K$ loading matrix, $\mathbf{U}$ is a $K \times 1$ latent factor vector, and $\boldsymbol{\varepsilon}$ is a d -dimensional random error term that is uncorrelated with $\mathbf{U}$. The model implied covariance structure is

$$\Sigma = \mathbf{B} \Sigma_U \mathbf{B}^T + \Sigma_e, \quad (3.12)$$

where $\Sigma_U = \text{Var}(U)$ and $\Sigma_e = \text{Var}(e)$. We assume Σ_e is sparse. This can be interpreted as the conditional sparse covariance model: Given the latent factor U , the conditional (after taking the linear projection out) covariance matrix of X is sparse.

Model (3.11) is not identifiable. The linear space spanned by the first K principal components of $\mathbf{B} \Sigma_U \mathbf{B}^T$ is the same as those spanned by the columns of $\mathbf{B}$. Without loss of generality, we impose the identifiability condition that the columns of $\mathbf{B}$ are orthogonal and $\Sigma_U = \mathbf{I}_K$.

3.5.3. Modeling Assumption—POET assumes that the factors U are pervasive and Σ_e is sparse. An example of pervasiveness is that the factor loadings $\{b_i\}_{i=1}^d$ (the transpose of the i^{th} row of $\mathbf{B}$) for $\{X_i\}_{i=1}^d$ are an independent realization from a certain population quantity b . In this case,

$$\frac{1}{d} \mathbf{B} \mathbf{B}^T = \frac{1}{d} \sum_{i=1}^d b_i b_i^T$$

converges to a non-degenerate matrix $\mathbb{E} b b^T$ (as $d \rightarrow \infty$). Therefore, the matrix $\mathbf{B} \mathbf{B}^T$ (recalling $\Sigma_U = \mathbf{I}_K$) in (3.12) has spiked eigenvalues. The sparseness is made so that $d^{-1} \|\Sigma_e\|_2 \rightarrow 0$. The rationale of this assumption is that, after taking out the common factors, many residual pairs become weakly correlated. Therefore, the first term in (3.12) dominates and the principal component analysis is approximately the same as the factor analysis when d is large.

The spiked eigenvalues assumption still needs to be verified on more applications. Since POET works very well in the presence of spiked principal eigenvalues where most of the aforementioned sparse covariance matrix estimation methods may fail, it has many potential applications in pharmacogenomics.

4. Large Inverse Covariance Matrix Estimation

Estimating a large inverse covariance matrix $\Theta = \Sigma^{-1}$ is another fundamental problem in Big data analysis. Unlike the covariance matrix Σ which only captures the marginal correlations among $X_1, \dots, X_d$, the inverse covariance matrix Θ captures the conditional correlations among these variables and is closely related to undirected graphs under a Gaussian model.

More specifically, we define an undirected graph $G = (V, E)$, where V contains nodes corresponding to the d variables in X and the edge $(j, k) \in E$ if and only if $\Theta_{jk} \neq 0$. Let $X_{-\{j,k\}} = \{X_\ell: \ell \neq j, k\}$. Under a Gaussian model $X \sim N(\mathbf{0}, \Sigma)$, X_j is independent of X_k given $X_{-\{j,k\}}$ for all $(j, k) \notin E$. Therefore, the graph estimation problem is reduced to estimating the inverse covariance matrix Θ [45].

Figure 3 illustrates the difference between the marginal and conditional uncorrelatedness. From Figure 3(b), we see the inverse covariance matrix Θ has many zero entries. Thus the undirected graph defined by Θ is sparse. The covariance matrix $\Sigma = \Theta^{-1}$ is

$$\Sigma = \begin{pmatrix} 1.05 & -0.23 & 0.05 & -0.02 & 0.05 \\ -0.23 & 1.45 & -0.25 & 0.10 & -0.25 \\ 0.05 & -0.25 & 1.10 & -0.24 & 0.10 \\ -0.02 & 0.10 & -0.24 & 1.10 & -0.24 \\ 0.05 & -0.25 & 0.10 & -0.24 & 1.10 \end{pmatrix}.$$

It is a dense matrix, which implies that every pair of variables are marginally correlated. Thus the covariance matrix and inverse covariance matrix encode different relationships. For example, in Figure 3(a), even though X_1 and X_5 are conditionally uncorrelated given the other variables, they are marginally correlated since both of them are correlated with X_2 .

In the following subsections, we first introduce the Gaussian graphical model which provides a general framework for inverse covariance matrix estimation. We then introduce several estimation methods which are computationally simple and suitable for Big data analysis. In particular, we highlight two class of methods: global methods such as penalized likelihood [28] and column-by-column methods such as CLIME [17] and TIGER [18], which have potential to be widely applied for different pharmacogenomics applications. The former provides positive definite inverse covariance matrix estimates and is more computation intensive. The latter is more scalable in computation and is suitable for graphical modeling. We also briefly explain how to combine the ideas of POET with TIGER.

We first introduce some additional notations. Let $\mathbf{A}$ be a symmetric matrix and I and J be index sets. We denote $\mathbf{A}_{I,J}$ to be the submatrix of $\mathbf{A}$ with rows and columns indexed by I and J . Let $\mathbf{A}_{*j}$ be the j^{th} column of $\mathbf{A}$ and $\mathbf{A}_{*-j}$ be the submatrix of $\mathbf{A}$ with the j^{th} column removed. We define the matrix norm:

$$\|\mathbf{A}\|_q = \max_{\|\mathbf{v}\|_q=1} \|\mathbf{A}\mathbf{v}\|_q.$$

It is easy to see that when $q = \infty$, $\|\mathbf{A}\|_\infty = \|\mathbf{A}\|_1$ for a symmetric matrix $\mathbf{A}$.

4.1. Gaussian Graphical Model

Let $\mathbf{X} \sim N(\mathbf{0}, \Sigma)$, $\alpha_j = (\Sigma_{-j,-j})^{-1} \Sigma_{-j,j}$ and $\sigma_j^2 = \Sigma_{jj} - \Sigma_{-j,j} (\Sigma_{-j,-j})^{-1} \Sigma_{-j,j}$. Then, we can decompose

$$X_j = \alpha_j^T \mathbf{X}_{-j} + \varepsilon_j \quad (4.1)$$

where $\varepsilon_j \sim N(0, \sigma_j^2)$ is independent of $\mathbf{X}_{-j} = \{X_\ell: \ell \neq j\}$. Using the block matrix inversion formula, we have

$$\Theta_{jj} = \sigma_j^{-2}, \quad (4.2)$$

$$\Theta_{-j,j} = -\sigma_j^{-2} \alpha_j. \quad (4.3)$$

Therefore, we can recover Θ by estimating the regression coefficient α_j and the residual variance σ_j^2 . Indeed, [46] estimates α_j by solving the Lasso problem

$$\hat{\alpha}_j = \underset{\alpha_j \in \mathbb{R}^{d-1}}{\operatorname{argmin}} \frac{1}{2n} \|\mathbf{X}_{*j} - \mathbf{X}_{*-j} \alpha_j\|_2^2 + \lambda_j \|\alpha_j\|_1, \quad (4.4)$$

where λ_j is a tuning parameter. Once $\hat{\alpha}_j$ is given, we get the neighborhood edges by reading out its nonzero coefficients. The final graph estimate $\hat{G}$ is obtained by combining the neighborhoods for all the d nodes. However, this method only estimates the graph G but not the inverse covariance matrix Θ .

4.2. Penalized Likelihood Estimation

When the penalized likelihood approach is applied to estimate Θ , it becomes

$$\hat{\Theta} = \underset{\Theta = (\theta_{ij})}{\operatorname{argmin}} \left\{ \underbrace{\log|\Theta| + \operatorname{tr}(\mathbf{S}\Theta)}_{\text{negative Gaussian log-likelihood}} + \sum_{i \neq j} P_{\lambda, \gamma}(\theta_{ij}) \right\}, \quad (4.5)$$

where the first part is the negative Gaussian log-likelihood and $P_{\lambda, \gamma}(\cdot)$ is defined the same way as in §3.2. It is a penalty term that encourages sparse solutions. The optimization in (4.5) can be computed by the graphical lasso [47] and the first-order method [48] and have also been studied in [27, 28].

4.3. The CLIME Estimator

To estimate both the inverse covariance matrix Θ and graph G , [17] proposes the CLIME estimator, which directly estimates the j^{th} column of Θ by solving

$$\hat{\Theta}_{*j} = \underset{\Theta_{*j}}{\operatorname{argmin}} \|\Theta_{*j}\|_1 \text{ subject to } \|\mathbf{S}\Theta_{*j} - \mathbf{e}_j\|_\infty \leq \delta_j, \quad (4.6)$$

where $\mathbf{e}_j$ is the j^{th} canonical vector (i.e., the vector with the j^{th} element being 1, while the remaining elements being 0) and δ_j is a tuning parameter. The method borrows heavily the idea from the Dantzig selector [49]. [17] shows that this convex optimization can be formulated into a linear program and has the potential to scale to large problems. Once $\hat{\Theta}$ is obtained, we use another tuning parameter τ to threshold it to estimate the graph G . However, $\hat{\Theta}$ itself might not be positive definite.

4.4. The TIGER Estimator

There are several ways to implement model (4.1). Examples include the graphical Dantzig selector [50], the scaled-Lasso estimator [51], and the SICO estimator [52]. Each of these involves d tuning parameters such as $\{\lambda_j\}_{j=1}^d$ in (4.4) and $\{\delta_j\}_{j=1}^d$ in (4.6). Yet, model (4.1) is heteroscedastic, with noise variance depending on unknown σ_j^2 . Therefore, these tuning parameters should depend on j , which is difficult to implement in practice.

TIGER method [18] overcomes these scale problems. Let $\hat{\mathbf{D}} = \operatorname{diag}(\mathbf{S})$ be a d -dimensional diagonal matrix with the diagonal elements being the same as those in $\mathbf{S}$. Let us normalize the problem and the data. Set

$$\mathbf{Z}=(Z_1, \dots, Z_d)^T=\mathbf{X}\widehat{\mathbf{D}}^{-1/2},$$

$$\beta_j=\widehat{s}_{jj}^{-1/2}\widehat{\mathbf{D}}_{-j,-j}^{1/2}\alpha_j \text{ and } \tau_j^2=\sigma_j^2\widehat{s}_{jj}^{-1},$$

where $\widehat{s}_{jj}$ is the j^{th} diagonal element of $\widehat{\mathbf{S}}$. Therefore, model (4.1) can be standardized as

$$Z_j=\beta_j^T\mathbf{Z}_{-j}+\widehat{s}_{jj}^{-1/2}\varepsilon_j. \quad (4.7)$$

We define $\widehat{\mathbf{R}}=(\widehat{\mathbf{D}})^{-1/2}\mathbf{S}(\widehat{\mathbf{D}})^{-1/2}$ to be the sample correlation matrix and let $\mathbf{Z}\in\mathbb{R}^{n\times d}$ be the normalized data matrix, i.e., $\mathbf{Z}_{*j}=\widehat{s}_{jj}^{-1/2}\mathbf{X}_{*j}$ for $j=1, \dots, d$. Motivated by the model in (4.7), we propose the following inverse covariance matrix estimator. More specifically,

$$\widehat{\beta}_j=\operatorname{argmin}_{\beta_j\in\mathbb{R}^{d-1}}\left\{\frac{1}{\sqrt{n}}\|\mathbf{Z}_{*j}-\mathbf{Z}_{*-j}\beta_j\|_2+\lambda\|\beta_j\|_1\right\},$$

$$\widehat{\tau}_j=\frac{1}{\sqrt{n}}\|\mathbf{Z}_{*j}-\mathbf{Z}_{*-j}\widehat{\beta}_j\|_2, \quad (4.8)$$

$$\widehat{\Theta}_{jj}=\widehat{\tau}_j^{-2}\widehat{s}_{jj}^{-1} \text{ and } \widehat{\Theta}_{-j,j}=-\widehat{\tau}_j^{-2}\widehat{s}_{jj}^{-1/2}\widehat{\mathbf{D}}_{-j,-j}^{-1/2}\widehat{\beta}_j.$$

Once we have $\widehat{\Theta}$, the estimated graph $\widehat{G}=(V, \widehat{E})$ is $(j, k)\in\widehat{E}$ if and only if $\widehat{\Theta}_{jk}\neq 0$. The formulation in (4.8) is called a SQRT-Lasso problem. [53] has shown that the SQRT-Lasso is tuning-insensitive. This explains the tuning-insensitive property of the TIGER method. The TIGER method is available in the R package flare. Similar to the CLIME method, the TIGER method estimates the inverse covariance matrix by solving many regression subproblems. Since these subproblems can be solved parallelly, this method scales to very large datasets.

An equivalent form to estimating the j^{th} column of Θ is to solve

$$\widehat{\beta}_j=\operatorname{argmin}_{\beta_j\in\mathbb{R}^{d-1}}\left\{\sqrt{1-2\beta_j^T\widehat{\mathbf{R}}_{-j,j}+\beta_j^T\widehat{\mathbf{R}}_{-j,-j}\beta_j}+\lambda\|\beta_j\|_1\right\},$$

$$\widehat{\tau}_j=\sqrt{1-2\widehat{\beta}_j^T\widehat{\mathbf{R}}_{-j,j}+\widehat{\beta}_j^T\widehat{\mathbf{R}}_{-j,-j}\widehat{\beta}_j}, \quad (4.9)$$

In (4.9), λ is a tuning parameter. [18] shows that, by choosing $\lambda=\pi\sqrt{\frac{\log d}{2n}}$, the obtained estimator achieves the minimax optimal rates of convergence, thus this procedure is asymptotically tuning-free. For finite samples, [18] suggests to set

$$\lambda=\zeta\pi\sqrt{\frac{\log d}{2n}}, \quad (4.10)$$

and ζ can be chosen from a grid in $[\sqrt{2}/\pi, 1]$. Since the choice of ζ does not depend on any unknown quantities, we call the TIGER procedure *tuning-insensitive*. Empirically, we can simply set $\zeta=\sqrt{2}/\pi$ and the resulting estimator works well in many applications.

If a symmetric inverse covariance matrix estimate is preferred, we make the correction: $\widehat{\Theta}_{jk}\leftarrow\min\{\widehat{\Theta}_{jk}, \widehat{\Theta}_{kj}\}$ for all $k\neq j$. Another correction method is

$$\tilde{\Theta} \leftarrow \frac{\widehat{\Theta} + \widehat{\Theta}^T}{2}. \quad (4.11)$$

As has been shown by [17], $\tilde{\Theta}$ achieves the same rate of convergence as $\hat{\Theta}$.

4.5. Combining POET with TIGER: Conditional Sparse Graphical Model

The TIGER estimator can be integrated into the POET framework. Recall that the main idea of the POET method is to exploit conditional sparsity. The common factors $\mathbf{BU}$ are extracted out from (3.11) via a principal component analysis, resulting in the residual matrix $\hat{\mathbf{R}}_K$ in (3.8). The thresholding rules are applied directly to $\hat{\mathbf{R}}_K$ due to the assumption of the conditional sparsity of the covariance matrix. If we assume $(\Sigma_e)^{-1}$ is sparse, namely, the genomics network is sparse conditioned on unobservable factors, it is inappropriate to apply POET. In this case, it is more appropriate to apply the TIGER method on $\hat{\mathbf{R}}_K$ to obtain the final estimate in (3.10). The graph induced by the TIGER method is a conditional graph after taking out the common dependence on latent factors.

5. Large-Scale Simultaneous Tests and False Discovery Control

Selection of differently expressed genes and proteins as well as finding SNPs that are associated with phenotypes or gene expressions give rise to large-scale simultaneous tests in pharmacogenomics [54, 55, 56]. They examine the marginal effects of the treatments and gain substantial popularity in the last decades [57, 58, 59, 60]. These procedures are designed predominately based on the assumption that test statistics are independent or weakly dependent [61, 62, 63]. Yet, biological outcomes such as gene expressions are often correlated, so do the statistical tests in the genomewide association studies (GWAS). It is very important to incorporate the dependence information in controlling the false discovery proportion (FDP) and to improve the power of the tests [64, 65, 66, 67]. In the recent paper [19] with discussions, it is convincingly demonstrated that FDP varies substantially from data to data, but can still be well estimated and controlled for each given data when the dependence information is properly used.

5.1. Rises of Large-Scale Hypothesis Tests

Suppose that gene or protein expression profiles $\{\mathbf{X}_i\}_{i=1}^m$ and $\{\mathbf{Y}_i\}_{i=1}^n$ are collected for m and n individuals respectively from the treatment and control groups. The problem of selecting significantly expressed genes is the following large-scale two-sample testing:

$$H_{0j}: \mu_{X,j} = \mu_{Y,j} \quad \text{vs} \quad H_{1j}: \mu_{X,j} \neq \mu_{Y,j}, \quad j=1, \dots, d$$

based on the assumption that

$$\mathbf{X}_i \sim N(\mu_X, \Sigma) \quad \text{and} \quad \mathbf{Y}_i \sim N(\mu_Y, \Sigma), \quad (5.1)$$

where $\mu_{X,j}$ and $\mu_{Y,j}$ are the mean expressions of the j^{th} gene in the treatment and control groups, respectively. Letting $\bar{\mathbf{X}}$ and $\bar{\mathbf{Y}}$ be the sample means, the two-sample Z -statistic is

$$\mathbf{Z}_1 = \sqrt{mn/(m+n)}(\bar{\mathbf{X}} - \bar{\mathbf{Y}}) \sim N(\mu_1, \Sigma)$$

with $\mu_1 = \sqrt{mn/(m+n)}(\mu_x - \mu_y)$. Let $\mathbf{D}$ be the diagonal matrix with the pooled estimates of the marginal standard deviations on its diagonal. The vector of the two-sample t -test statistics has the approximate distribution as follows:

$$\mathbf{Z} = \mathbf{D}^{-1} \mathbf{Z}_1 \stackrel{a}{\sim} N(\boldsymbol{\mu}, \mathbf{R}) \quad (5.2)$$

where $\boldsymbol{\mu} = \mathbf{D}^{-1} \boldsymbol{\mu}_1$, $\mathbf{R}$ is the correlation matrix of $\boldsymbol{\Sigma}$, and $\stackrel{a}{\sim}$ means “distributed approximately”. Our testing problem then reduces to

$$H_{0j}: \mu_j = 0 \quad \text{vs} \quad H_{1j}: \mu_j \neq 0, \quad j=1, \dots, d \quad (5.3)$$

based on the vector of test statistics $\mathbf{Z}$. In pharmacogenomics applications, the correlation matrix $\mathbf{R}$ represents the co-expressions of molecules and is in general unknown.

Another concrete example is that in GWAS, we wish to associate the SNPs with phenotypes or gene expressions. For individual i , let X_{ij} be the genotype of the j^{th} SNP and Y_i be its associated outcome. The association between the j^{th} SNP and the response is measured through the marginal regression [19]:

$$(\alpha_j, \beta_j) = \underset{a_j, b_j}{\operatorname{argmin}} n^{-1} \sum_{i=1}^n \mathbb{E}(Y_i - a_j - b_j X_{ij})^2.$$

Our simultaneous test in GWAS becomes testing

$$H_{0j}: \beta_j = 0 \quad \text{vs} \quad H_{1j}: \beta_j \neq 0, \quad j=1, \dots, d. \quad (5.4)$$

Let $\hat{\boldsymbol{\beta}}$ be the vector of the least-squares estimates of the regression coefficients $\boldsymbol{\beta} = (\beta_1, \dots, \beta_d)^T$, which are obtained by the marginal regressions based on the data $\{(x_{ij}, y_i)\}_{i=1}^n$. Denoting by $\mathbf{Z} = (Z_1, \dots, Z_d)^T$ the t -test statistics for the problem (5.4), it is easy to show that $\mathbf{Z} \stackrel{a}{\sim} N(\boldsymbol{\beta}, \mathbf{R})$, where $\mathbf{R}$ is the correlation matrix of $\{X_{ij}\}_{i=1}^n$. In this case, the correlation matrix is known.

The above problem can be summarized as follows. Let $\mathbf{Y} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with unit standard deviations ($\sigma_{11}^2 = \dots = \sigma_{dd}^2 = 1$) be a vector of dependent test statistics. We wish to simultaneously test

$$H_{0j}: \mu_j = 0 \quad \text{vs} \quad H_{1j}: \mu_j \neq 0, \quad j=1, \dots, d. \quad (5.5)$$

In pharmacogenomics applications, the number of hypotheses d is large and the number of interested genes or proteins or SNPs are small so that most of $\{\mu_j\}_{j=1}^d$ are zero.

5.2. False Discovery Rate and Proportion

Let $S_0 = \{j: \mu_j = 0\}$ be the set of true nulls and $\Theta(\cdot)$ be the cumulative distribution function of the standard Gaussian random variable. We denote $P_j = 1 - 2\Theta(|Z_j|)$ to be the P-value for the j^{th} test of problem (5.5) based on the dependent test statistics (5.2). Using a threshold t , we define

$$\text{FDP}(t) = V(t)/R(t) \quad \text{and} \quad \text{FDR}(t) = \mathbb{E}\{\text{FDP}(t)\}$$

to be the false discovery proportion and false discovery rate, where $V(t) = \#\{j \in S_0 : P_j \leq t\}$ and $R(t) = \#\{j : P_j \leq t\}$ are the number of false rejections and total rejections, respectively. Our aim is to accurately estimate $\text{FDP}(t)$ for a large class of Σ . Clearly, FDP is more relevant, since it is related to the data at hand whereas FDR only controls the average FDP. To control FDP at a prescribed level α , we choose the smallest t_0 such that estimated $\text{FDP}(t_0) \leq \alpha$ [56, 19].

Note that $R(t)$ is the number of rejected hypotheses or the number of discoveries. It is an observable random variable. Yet, $V(t)$ is unknown to us and needs to be estimated. When test statistics are independent, $V(t)$ is the number of success in a sequence of independent Bernoulli trials, with the number of trials $d_0 = |S_0|$ (the number of true nulls) and probability of success t . Therefore, by the law of large numbers, FDP and FDR are close, both approximately $d_0 t / R(t) \approx d t / R(t)$ due to sparsity. For dependent test statistics, however, FDP varies substantially from one data to another, even when FDR is the same. We use a simple example in [19] to illustrate this point.

For equally correlated matrix Σ with a common correlation ρ , we can decompose the test statistic $\{Y_j\}_{j=1}^d$ as

$$Y_j = \mu_j + \rho U + \sqrt{1 - \rho^2} \varepsilon_j, \quad j = 1, \dots, d.$$

where $U \sim N(0, 1)$ is independent of $\varepsilon_j \sim N(0, 1)$. Without loss of generality, we assume that S_0 is the first d_0 genes. Let $z_{t/2}$ be the upper $t/2$ -quantile of the standard normal distribution and $|Y_j| \geq z_{t/2}$ be a critical region. Then, by the law of large numbers,

$$V(t) = \sum_{j=1}^{d_0} I(|Y_j| \geq z_{t/2}) \approx d_0 P(t, \eta, \rho) \quad (5.6)$$

where $P(t, \eta, \rho) = \mathbb{P}(|Y_j| \geq z_{t/2} \mid U, \mu_j = 0)$ is given by

$$P(t, \eta, \rho) = \Phi\left(\frac{-z_{t/2} - \eta}{\sqrt{1 - \rho^2}}\right) + 1 - \Phi\left(\frac{z_{t/2} - \eta}{\sqrt{1 - \rho^2}}\right), \quad (5.7)$$

with $\eta = \rho U$. It is clear that the result depends critically on the realization of U . For example, when $\rho = 0.8$, $d_0 = 1,000$ and $z_{t/2} = 2.5$,

$$V(t) \approx 1000 \times \left\{ \Phi\left(\frac{-2.5 - 0.8U}{0.6}\right) + 1 - \Phi\left(\frac{2.5 - 0.8U}{0.6}\right) \right\}$$

which are 0, 2.3, 66.8, 433.6 respectively for $U = 0, 1, 2$, and 3. Clearly, it depends on the realization of U . Yet, the realized factor U is estimable. For example, using sparsity, $\bar{Y} \approx \rho U$

and thus U can be estimated by $\bar{Y}/\rho$. On the other hand, $\mathbb{E} V(t) = d_0 t = 12.4$. Clearly, $V(t)$ or $FDP(t)$ is far more relevant than $\mathbb{E} V(t)$ or $FDR(t)$.

In summary, FDP is a quantity of the primary interest for tests that are dependent.

5.3. Principal Factor Approximation

When the correlation matrix Σ is known as in GWAS, [19] gives a principal factor approach to approximate $FDP(t)$ for arbitrary Σ . It generalizes formula (5.6). Let $\{\lambda_j\}_{j=1}^d$ be ordered eigenvalues of Σ and $\{\xi_j\}_{j=1}^d$ be their corresponding eigenvectors. We can decompose Σ as

$$\Sigma = \sum_{j=1}^d \lambda_j \xi_j \xi_j^T = \mathbf{B} \mathbf{B}^T + \mathbf{A}$$

where $\mathbf{B} = (\sqrt{\lambda_1} \xi_1, \dots, \sqrt{\lambda_K} \xi_K)$ consists of the first K unnormalized principal components and $\mathbf{A} = \sum_{j=K+1}^d \lambda_j \xi_j \xi_j^T$. Then, as in (3.11), we can write

$$\mathbf{Y} = \boldsymbol{\mu} + \mathbf{B} \mathbf{U} + \boldsymbol{\varepsilon}, \quad \mathbf{U} \sim N(0, \mathbf{I}_K), \quad \boldsymbol{\varepsilon} \sim N(0, \mathbf{A}). \quad (5.8)$$

Since principal factors $\mathbf{U}$ that drive the dependence are taken out, $\boldsymbol{\varepsilon}$ can be assumed to be weakly correlated and $FDP(t)$ can be approximated by

$$FDP_A(t) = \sum_{j=1}^d P(t, \eta_j, \|\mathbf{b}_j\|_2) / R(t), \quad \eta_j = \mathbf{b}_j^T \mathbf{U} \quad (5.9)$$

under some mild conditions [19], where $\mathbf{b}_j^T$ is the j^{th} row of $\mathbf{B}$ and the function $P(\cdot)$ is given in (5.7). Note that the covariance matrix is used to calculate η_j and $\|\mathbf{b}_j\|_2$ and that formula (5.9) is a generalization of (5.6). Thus, we need only to estimate the realized but unobserved factors $\mathbf{U}$ in order to use $FDP_A(t)$. [19] provides a penalized estimator as described below.

Since $\boldsymbol{\mu}$ is sparse, by (5.8), a natural estimator is the minimizer of

$$\min_{\boldsymbol{\mu}, \mathbf{U}} \|\mathbf{Y} - \boldsymbol{\mu} - \mathbf{B} \mathbf{U}\|_2^2 + \lambda \|\boldsymbol{\mu}\|_1 \quad (5.10)$$

which is equivalent [68] to minimizing the Huber's ψ -loss with respect to $\mathbf{U}$: $\min_{\mathbf{U}} \psi(\mathbf{Y} - \mathbf{B} \mathbf{U})$. An alternative is to find the minimizer of the L_1 -loss:

$$\min_{\mathbf{U}} \|\mathbf{Y} - \mathbf{B} \mathbf{U}\|_1. \quad (5.11)$$

The sampling property of the estimator based on (5.10) has been established in [68].

With estimated factors $\hat{\mathbf{U}}$, the estimated FDP can be computed by using (5.9) for each given critical value t . This provides useful information about the FDP for a given set of selected molecules. To control the FDP at a prescribed level α , we choose the smallest t such that the estimated $FDP(t)$ does not exceed α .

5.4. Factor Adjustment: an Alternative Ranking of Molecules

With $\hat{U}$ estimated from (5.10), the common factors can be subtracted out from model (5.8), resulting in

$$Z = Y - B\hat{U} \approx \mu + \epsilon. \quad (5.12)$$

The test statistics Z are more powerful than those based on Y since ϵ has marginal variances no larger than 1. In fact, $Z_i \sim N(\mu_i, 1 - \|b_i\|_2^2)$, after ignoring the approximation error. Hence, Z_i has a larger signal-to-noise ratio than $Y_i \sim N(\mu_i, 1)$. In addition, ranking of important genes/proteins based on the standardized test statistics $\{|Z_i|/(1 - \|b_i\|_2^2)^{1/2}\}_{i=1}^d$ can be very different from that based on $\{|Y_i|\}_{i=1}^d$. This method is called dependent-adjusted procedure in [19]. It provides pharmaco-scientists a chance to discover important genes, proteins, and SNPs that are not highly ranked by the conventional methods, based on ranking of $\{|Y_i|\}_{i=1}^d$.

After common factors being taken out, the elements of ϵ are weakly correlated. There is no need to apply PFA to the test statistics Z in (5.12). In other words, FDP should be

approximately the same as FDR. But application of PFA to $\{|Z_i|/(1 - \|b_i\|_2^2)^{1/2}\}_{i=1}^d$ will not result in very different results from the case with $K = 0$.

5.5. Estimating FDP with Unknown Dependence

In many genomic applications such as selecting significant genes from microarray data, the covariance matrix Σ is typically unknown. A natural approach is to estimate it from the data by POET as described in §3.5.1. Then proceed as if the correlation matrix is given. The resulting method is named POET-PFA. More details on the development of this method can be found in [69]. A R package, called PFA, implements this procedure. The main computational cost of the POET-PFA method is the POET step, thus it computationally efficient and scales to large problems as has been discussed in §3.5.1.

6. High Dimensional Variable Selection

With proper coding of the treatment and control groups, it is easy to see that the two-sample test statistics in (5.2) are equivalent to the marginal correlations between the expressions of molecules and response [21]. To examine the conditional contribution of a molecule to the response Y given the rest of molecules, one often employs a sparse generalized linear model under canonical link: Up to a normalization constant, the conditional density of Y given X is

$$\log p(y|x, \beta_0) = (x^T \beta_0)y - b(x^T \beta_0) \quad (6.1)$$

for a given function $b(\cdot)$ that depends on the model. Here, the sparsity means that a majority of regression coefficients in $\beta_0 = (\beta_{0,1}, \dots, \beta_{0,d})^T$ are zero. The regression coefficient $\beta_{0,j}$ is frequently regarded as a measure of conditional contribution of the molecule X_j to the response Y given the rest of molecules, similar to the interpretation of the sparse inverse covariance matrix in §4.

When the response Y is a continuous measurement such as in eQTL studies, one takes $b(\theta) = \theta^2/2$, corresponding to a Gaussian model. For binary response, one often employs a sparse logistic regression in which $b(\theta) = -\log(1 + \exp(\theta))$. The latter is closely related to classification problems as will be described in §6.4.

Over the last decade, there is a surged interest in sparse regression. For an overview, see [40, 70, 71]. The basic principles are screening and penalization. We give only a brief summary due to the space limitation.

6.1. Penalized Likelihood Method

As has been shown in [20], sparse vector β_0 in (6.1) can be estimated by penalized likelihood which minimizes

$$n^{-1} \sum_{i=1}^n b(x_i^T \beta) - (x_i^T \beta) y_i + \sum_{j=1}^d P_{\lambda, \gamma}(|\beta_j|) \quad (6.2)$$

for a folded concave penalty function $P_{\lambda, \gamma}(\cdot)$. It has been systematically studied by [72, 73].

There are many algorithms to compute the minimizer of (6.2). Examples include local quadratic approximation [20], local linear approximation [74], coordinate decent algorithm [75], iterative shrinkage-thresholding algorithm [76, 77]. For example, given estimate

$\widehat{\beta}^{(k)} = (\beta_1^{(k)}, \dots, \beta_d^{(k)})^T$ at the k^{th} iteration, by Taylor's expansion,

$$P_{\lambda, \gamma}(|\beta_j|) \approx P_{\lambda, \gamma}(|\beta_j^{(k)}|) + P'_{\lambda, \gamma}(|\beta_j^{(k)}|)(|\beta_j| - |\beta_j^{(k)}|). \quad (6.3)$$

Thus, at the $(k+1)^{th}$ iteration, we minimize

$$n^{-1} \sum_{i=1}^n b(x_i^T \beta) - (x_i^T \beta) y_i + \sum_{j=1}^d w_{k,j} |\beta_j|, \quad (6.4)$$

where $w_{k,j} = P'_{\lambda, \gamma}(|\beta_j^{(k)}|)$. Note that problem (6.4) is convex so that a convex solver can be used. [26] suggests using initial values $\beta^{(0)} = \mathbf{0}$, which corresponds to the unweighted L_1 penalty. If one further approximates the likelihood part in (6.4) by a quadratic function via Taylor expansion, then the LARS algorithm [78] and pathwise coordinate descent algorithm [75] can be used. Both algorithms are efficient and scale to very high dimensional problems.

6.2. Screening Method

The sure independent screening [21, 79] is another effective method to reduce the dimensionality in sparse regression problems. It utilizes the marginal contribution of a covariate to probe its importance in the joint regression model. For example, assuming each covariate has been standardized, the marginal contribution can be measured by the magnitude of $\widehat{\beta}_j^M$, which, along with $\widehat{\alpha}_j^M$, minimizes the negative marginal likelihood:

$$n^{-1} \sum_{i=1}^n b(\alpha_j + \beta_j x_{ij}) - (\alpha_j + \beta_j x_{ij}) y_i. \quad (6.5)$$

The set of covariates that survive the marginal screening is

$$\widehat{\mathcal{S}} = \{j: |\widehat{\beta}_j^M| \geq \delta\}, \quad (6.6)$$

for a given threshold δ . One can also measure the importance of a covariate X_j by using its deviance reduction. For the least-squares problem, both methods reduce to ranking importance of predictors by using the magnitudes of their marginal correlations with the response Y . [21, 79] give conditions under which sure screening property can be established and false selection rate are controlled.

Since the computational complexity of sure screening scales linearly with the dimensionality d , the idea of sure screening is very effective in dramatic reduction of the computational burden of Big data analysis. It has been extended in various directions. For example, generalized correlation screening was used in [80], nonparametric screening was proposed by [81], and principled SIS was introduced in [82]. In addition, [83] utilizes the distance correlation to conduct screening and [84] employs rank correlation.

6.3. Iterative Sure Independence Screening

An iterative sure independence screening method was developed in [85]. The basic idea is to iteratively use the large-scale screening (6.6) and the moderate-scale selection, which applies the penalized likelihood method (6.2) to the survived variables (6.6). The software, ISIS, is available in R-package, which implements this idea. It also allows us to compute penalized likelihood (6.2) or (6.6) alone.

6.4. Sparse Classification

Classification, also known as supervised learning, is to predict the class label for a given covariate. In pharmacogenetic studies, one often wishes to not only predict well the class labels, but also understand the molecule mechanisms for the effectiveness of drugs. Therefore, it is desirable to select a small group of molecules that have a high classification power.

Sparse logistic regression (6.2) provides an effective approach to high dimensional classification. Other approaches include the support vector machine [86] and the AdaBoost algorithm [87, 88]. The former replaces the log-likelihood by the hinge loss $L(\theta, y) = (1 - \theta y)_+$ and the latter replaces the log-likelihood by the exponential loss $L(\theta, y) = \exp(-y)$. In both cases, the response Y is recoded as ± 1 .

When data from both classes follow normal distributions (5.1), the optimal classifier is the Fisher's discriminant rule, which classifies a new data point x to class "X" if

$$(\mu_x - \mu_y)^T \Sigma^{-1} (x - (\mu_x + \mu_y)/2) > 0, \quad (6.7)$$

where Σ is the common covariance matrix in (5.1). When the dimensionality d is large, we could estimate Σ or its inverse Σ^{-1} by the regularized estimators introduced in §3 or §4, depending on the sparsity assumption. More details about these classification methods and their computational aspects can be found in [89, 90, 91, 92].

7. Applications

In this section, we apply the aforementioned methods to estimate large gene networks and select significantly differently expressed genes when test statistics are dependent.

7.1. Network Estimation

As explained in §4, an important application of large inverse covariance matrix estimation is to reconstruct the undirected graph of a high dimensional distribution based on observational

data. In this section, we apply the TIGER method on a microarray data to illustrate the main idea.

This dataset includes 118 gene expression arrays from *Arabidopsis thaliana* originally appeared in [93]. Our analysis focuses on expression profiles from 39 genes involved in two isoprenoid metabolic pathways: 16 from the mevalonate (MVA) pathway are located in the cytoplasm, 18 from the plastidial (MEP) pathway are located in the chloroplast, and 5 are located in the mitochondria. While the MVA and MEP pathways generally operate independently, crosstalk is known to happen [93]. Our goal is to reconstruct the gene network by estimating the undirected graph, with special interest in crosstalk.

We first examine whether the data actually satisfies the Gaussian distribution assumption. In Figure 4(a) and Figure 4(b), we plot the histogram and normal QQ plot of the expression levels of a gene named MECPS. From the histogram, we see that the distribution is left-skewed compared to the Gaussian distribution. From the normal QQ plot, we see that the empirical distribution has a heavier tail compared to Gaussian. To apply the TIGER method to analyze this data, we need to first transform the data so that its distribution is closer to Gaussian. For this, we Gaussianize the expression values of each gene by converting them to the corresponding normal-scores. This is automatically done by the `huge.npn` function in the R package `huge`[94].

We apply the TIGER method on the transformed data using the default tuning parameter $\zeta = \sqrt{2}/\pi$. The estimated network is shown in Figure 4(c). We see the estimated network is very sparse with only 44 edges. We draw the within-pathway connections using solid lines and the between-pathway connections using dashed lines. Our result is consistent with previous investigations, which suggest that the connections from genes AACT1 and HMGR2 to gene MECPS indicate a primary sources of the crosstalk between the MEP and MVA pathways and these edges are presented in the estimated network. MECPS is clearly a hub gene for this pathway.

For the MEP pathway, the genes DXPS2, DXR, MCT, CMK, HDR, and MECPS are connected as in the true metabolic pathway. Similarly, for the MVA pathway, the genes AACT2, HMGR2, MK, MPDC1, MPDC2, FPPS1 and FPP2 are closely connected. Our analysis suggests 11 cross-pathway links, which is consistent to previous investigation in [93]. This result suggests that there might exist rich inter-pathway crosstalks.

7.2. Select Significantly Expressed Genes under Dependency

We now apply the POET-PFA method in §5.5 and the dependence-adjusted method in §5.4 to analyze a microarray data. The data contains expression profiles from a well-known breast cancer study [95, 96]. This study involves 15 woman subjects with two genetic mutations: BRCA1 (7 subjects) and BRCA2 (8 subjects). These two genetic mutations are known to increase the lifetime risk of hereditary breast cancer. We want to find a set of genes that are associated with these genetic mutations. This allows us to identify cases of hereditary breast cancer on the basis of gene-expression profiles.

We make two assumptions: (i) A large proportion of the genes are not differently expressed; (ii) The gene expression follows an approximate k -factor model. Both assumptions have gained increasing popularity among biologists in the past decade, since it has been widely acknowledged that gene activities are usually driven by a small number of latent variables and the genetic mutations are only caused by a small amount of genes. See [65, 97] for more details.

We first apply the POET-PFA method to obtain a consistent FDP estimator $\widehat{\text{FDP}}(t)$ for a given threshold value t and a fixed number of factors k .

Let $\hat{V}(t)$ be the estimated number of false rejections, which is the numerator of (5.9). The results of our analysis are summarized in Figure 5(a), which has two subfigures that plot $(R(t), \widehat{\text{FDP}}(t))$ and $(R(t), \hat{V}(t))$ for $k = 2, 3, 4, 5$. We see that $\widehat{\text{FDP}}(t)$ is close to zero when $R(t)$ is below 200, suggesting that the rejected hypotheses in this range have high accuracy to be the true discoveries. In addition, when 1,000 hypotheses, almost 1/3 of the total number, have been rejected, the estimated FDPs are still as low as 25%. Finally, it is worth noting that although our procedure seems robust under different choices of number of factors, the estimated FDP tends to be relatively small with larger number of factors.

We also apply the dependence-adjusted procedure to the data. The results are shown in Figure 5(b). Comparing with Figure 5(a), both the estimated FDP and the estimated number of false rejections become smaller.

The selected gene sets based on the conventional two-sample t tests and the factor-adjusted tests are different (not presented here, but see [69]). This provides pharmaco-scientists a chance to discover important genes, proteins, and SNPs that are not highly ranked by the conventional methods. In fact, the factor-adjusted method is more powerful than the conventional two-sample t -tests when test statistics are dependent.

8. Discussion and Future Directions

The paper is written in the context that massive amounts of pharmacogenomic data combined with quantitative statistical approaches are beginning to shed lights towards personalized medicine. While the so-called Big data movement has received a lot of attention, and the pharmacogenomics applications are prime examples, we highlight three characteristics of modern pharmacogenomic data that worth equal attention. These include: (1) Complex Data - the data distribution can not be characterized by simple parametric models; (2) Noisy Data - the data are usually aggregated from numerous sources and we must deal with large biological and measurement heterogeneity with possible outliers, measurement errors, and missing data; (3) Dependent Data - the data exhibit complicated correlations with relatively weak signals. These problems are now so ubiquitous, with new variations accruing at such an alarming rate, one might refer to them as simply the **Modern data** problem. The large amount of data are usually obtained by using high-throughput technologies and they contain inevitably measurement errors. Measurement errors distort statistical conclusion, reducing correlation with clinical outcomes. Spurious correlations arise inevitably in high dimensional data. They induce the so-called endogenous covariates. When some collected variables are correlated with the residual noise of a response, namely, when the part of the response that can not be explained by relevant molecules are correlated with irrelevant molecules, endogeneity arises. Like measurement errors, endogeneity causes model selection inconsistency, leading to erroneous scientific conclusions.

To handle the challenges of modern pharmacogenomic data, we need statistical methods that are simultaneously robust to the main issues of scalability, complexity, noise, and dependence. However, general methodological development for pharmacogenomic data analysis, mirroring numerous other scientific settings, has lagged behind the rapid development of new technologies and new datasets. For example, most existing methods assume the data are independently sampled from a parametric distribution (e.g., Gaussian), and most of them use Pearson's sample covariance matrix as the sufficient statistic and thus are not robust to outliers and possible data contamination. To bridge these gaps, new statistical methodology is urgently needed. Some rank based statistical methods that are

adaptive within the elliptical [98], nonparanormal [99, 100, 101], and transelliptical [102] families have shed lights on this.

In summary, statistical analysis of Big pharmacogenomics data is both promising and challenging. On one hand, significant progress has been made towards developing new statistical method and theory to explore and predict massive amounts of high dimensional data. On the other hand, there are many new challenges and open problems remain to be solved.

Acknowledgments

We thank Rongling Wu for his helpful comments and discussions. Jianqing Fan is supported by NSF Grant DMS-1206464 and NIH Grants R01GM100474 and R01-GM072611. Han Liu is supported by NSF Grant III-1116730 and a NIH sub-award from Johns Hopkins University.

References

1. Evans WE, Relling MV. Pharmacogenomics: translating functional genomics into rational therapeutics. *Science*. 1999; 286(5439):487–491. [PubMed: 10521338]
2. Wood AJ, Evans WE, McLeod HL. Pharmacogenomics-drug disposition, drug targets, and side effects. *New England Journal of Medicine*. 2003; 348(6):538–549. [PubMed: 12571262]
3. Jain K. Applications of biochip and microarray systems in pharmacogenomics. *Pharmacogenomics*. 2000; 1(3):289–307. [PubMed: 11256580]
4. Mishra PJ, Bertino JR. MicroRNA polymorphisms: the future of pharmacogenomics, molecular epidemiology and individualized medicine. *Pharmacogenomics*. 2009; 10(3):399–416. [PubMed: 19290790]
5. Winkelmann BR, März W, Boehm BO, Zotz R, Hager J, Hellstern P, Senges J. Rationale and design of the luric study—a resource for functional genomics, pharmacogenomics and long-term prognosis of cardiovascular disease. *Pharmacogenomics*. 2001; 2(1):1–73.
6. Wheeler HE, Maitland ML, Dolan ME, Cox NJ, Ratain MJ. Cancer pharmacogenomics: strategies and challenges. *Nature Reviews Genetics*. 2013:23–34.
7. Ross S, Anand SS, Joseph P, Paré G. Promises and challenges of pharmacogenetics: an overview of study design, methodological and statistical issues. *JRSM Cardiovascular Disease*. 2012; 1(1):1–7.
8. Wu, R.; Lin, M. Statistical and computational pharmacogenomics. Chapman & Hall/CRC; 2008.
9. Grady BJ, Ritchie MD. Statistical optimization of pharmacogenomics association studies: key considerations from study design to analysis. *Current pharmacogenomics and personalized medicine*. 2011; 9(1):41–66. [PubMed: 21887206]
10. Wang SJ, O'Neill RT, Hung HJ. Statistical considerations in evaluating pharmacogenomics-based clinical effect for confirmatory trials. *Clinical Trials*. 2010; 7(5):525–536. [PubMed: 20595242]
11. Turner SD, Crawford DC, Ritchie MD. Methods for optimizing statistical analyses in pharmacogenomics research. *Expert review of clinical pharmacology*. 2009; 2(5):559–570. [PubMed: 20221410]
12. Topi, E. The 7th EFCC Continuous Postgraduate Course in Clinical Chemistry: New trends in diagnosis, monitoring and management using molecular diagnosis methods. *Pharmacogenomics and personalized medicine*.
13. Yap JS, Fan J, Wu R. Nonparametric modeling of longitudinal covariance structure in functional mapping of quantitative trait loci. *Biometrics*. 2009; 65(4):1068–1077. [PubMed: 19302406]
14. Bickel P, Levina E. Covariance regularization by thresholding. *Annals of Statistics*. 2008; 36(6): 2577–2604.
15. Fan J, Liao Y, Mincheva M. High dimensional covariance matrix estimation in approximate factor models. *Annals of statistics*. 2011; 39(6):3320–3356. [PubMed: 22661790]
16. Fan J, Liao Y, Mincheva M. Large covariance estimation by thresholding principal orthogonal complements (with discussion). *Journal of the Royal Statistical Society Series B*. to appear.

17. Cai T, Liu W, Luo X. A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association*. 2011; 106(494):594–607.
18. Liu, H.; Wang, L. Tech rep. Massachusetts Institute of Technology; 2012. Tiger: A tuning-insensitive approach for optimally estimating gaussian graphical models.
19. Fan J, Han X, Gu W. Estimating false discovery proportion under arbitrary covariance dependence (with discussion). *Journal of American Statistical Association*. 2012; 107(3):1019–1048.
20. Fan J, Li R. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*. 2001; 96(456):1348–1360.
21. Fan J, Lv J. Sure independence screening for ultrahigh dimensional feature space (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2008; 70(5):849–911.
22. Geman S. A limit theorem for the norm of random matrices. *Annals of Probability*. 1980; 8(2): 252–261.
23. Yin YQ, Bai ZD, Krishnaiah PR. On the limit of the largest eigenvalue of the large-dimensional sample covariance matrix. *Probability Theory Related Fields*. 1988; 78:509–521.
24. Wu W, Pourahmadi M. Nonparametric estimation of large covariance matrices of longitudinal data. *Biometrika*. 2003; 90(1):831–844.
25. Bickel P, Levina E. Some theory for fisher's linear discriminant function, "naive bayes", and some alternatives when there are many more variables than observations. *Bernoulli*. 2004; 10(6):989–1010.
26. Fan J, Fan Y, Lv J. High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics*. 2008; 147(1):186–197.
27. Bickel P, Levina E. Regularized estimation of large covariance matrices. *Annals of Statistics*. 2008; 36(1):199–227.
28. Lam C, Fan J. Sparsistency and rates of convergence in large covariance matrix estimation. *Annals of Statistics*. 2009; 37:42–54.
29. Cai T, Liu W. Adaptive thresholding for sparse covariance matrix estimation. *Journal of the American Statistical Association*. 2011; 106(494):672–684.
30. Furrer R, Bengtsson T. Estimation of high-dimensional prior and posterior covariance matrices in kalman filter variants. *Journal of Multivariate Analysis*. 2007; 98(2):227–255.
31. Huang J, Liu N, Pourahmadi M, Liu L. Covariance matrix selection and estimation via penalised normal likelihood. *Biometrika*. 2006; 93(1):85–98.
32. Levina E, Rothman A, Zhu J. Sparse estimation of large covariance matrices via a nested lasso penalty. *The Annals of Applied Statistics*. 2008:245–263.
33. Rothman A, Levina E, Zhu J. A new approach to cholesky-based covariance regularization in high dimensions. *Biometrika*. 2010; 97(3):539–550.
34. Cai T, Zhang C, Zhou H. Optimal rates of convergence for covariance matrix estimation. *Annals of Statistics*. 2010; 38(4):2118–2144.
35. Donoho DL, Johnstone JM. Ideal spatial adaptation by wavelet shrinkage. *Biometrika*. 1994; 81(3): 425–455.
36. Rothman A, Levina E, Zhu J. Generalized thresholding of large covariance matrices. *Journal of the American Statistical Association*. 2009; 104(485):177–186.
37. Tibshirani R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*. 1996; 58(1):267–288.
38. Chen S, Donoho D, Saunders M. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*. 1998; 20(1):33–61.
39. Zhang C. Nearly unbiased variable selection under minimax concave penalty. *Annals of Statistics*. 2010; 38(2):894–942.
40. Hastie, T.; Tibshirani, R.; Friedman, J. *The elements of statistical learning: data mining, inference, and prediction*. 2. Springer; New York: 2009.
41. Antoniadis A, Fan J. Regularization of wavelet approximations. *Journal of the American Statistical Association*. 2001; 96(455):939–967.
42. Liu H, Wang L, Zhao T. Sparse covariance estimation with eigenvalue constraints. *Journal of Computational and Graphical Statistics*. to appear.

43. Xue L, Ma S, Zou H. Positive definite ℓ_1 penalized estimation of large covariance matrices. *Journal of the American Statistical Association*. 2012; 107(500):1480–1491.
44. Rothman A. Positive definite estimators of large covariance matrices. *Biometrika*. to appear.
45. Dempster A. Covariance selection. *Biometrics*. 1972; 28:157–175.
46. Meinshausen N, Bühlmann P. High dimensional graphs and variable selection with the lasso. *Annals of Statistics*. 2006; 34(3):1436–1462.
47. Friedman J, Hastie T, Tibshirani R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*. 2008; 9(3):432–441. [PubMed: 18079126]
48. d'Aspremont A, Banerjee O, El Ghaoui L. First-order methods for sparse covariance selection. *SIAM Journal on Matrix Analysis and Applications*. 2008; 30(1):56–66.
49. Candès E, Tao T. The dantzig selector: Statistical estimation when p is much larger than n . *The Annals of Statistics*. 2007; 35(6):2313–2351.
50. Yuan M. High dimensional inverse covariance matrix estimation via linear programming. *Journal of Machine Learning Research*. 2010; 11(8):2261–2286.
51. Sun, T.; Zhang, C-H. Tech rep. Department of Statistics, Rutgers University; 2012. Sparse matrix inversion with scaled lasso.
52. Liu, W.; Luo, X. High-dimensional sparse precision matrix estimation via sparse column inverse operator. arXiv/1203.3896
53. Belloni A, Chernozhukov V, Wang L. Square-root lasso: Pivotal recovery of sparse signals via conic programming. *Biometrika*. 2012; 98:791–806.
54. Efron B, Tibshirani R, Storey JD, Tusher V. Empirical bayes analysis of a microarray experiment. *Journal of the American Statistical Association*. 2001; 96(456):1151–1160.
55. Dudoit S, Fridlyand J, Speed TP. Comparison of discrimination methods for the classification of tumors using gene expression data. *Journal of the American statistical association*. 2002; 97(457): 77–87.
56. Storey JD, Tibshirani R. Statistical significance for genomewide studies. *Proceedings of the National Academy of Sciences*. 2003; 100(16):9440–9445.
57. Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*. 1995:289–300.
58. Genovese C, Wasserman L. A stochastic process approach to false discovery control. *The Annals of Statistics*. 2004; 32(3):1035–1061.
59. Lehmann E, Romano JP, Shaffer JP. On optimality of stepdown and stepup multiple test procedures. *Annals of statistics*. 2005:1084–1108.
60. Efron, B. Large-scale inference: empirical Bayes methods for estimation, testing, and prediction. Vol. 1. Cambridge University Press; 2010.
61. Benjamini Y, Yekutieli D. The control of the false discovery rate in multiple testing under dependency. *Annals of statistics*. 2001:1165–1188.
62. Storey JD, Taylor JE, Siegmund D. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: a unified approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2003; 66(1):187–205.
63. Clarke S, Hall P. Robustness of multiple testing procedures against dependence. *The Annals of Statistics*. 2009; 37(1):332–358.
64. Efron B. Correlation and large-scale simultaneous significance testing. *Journal of the American Statistical Association*. 2007; 102(477):93–103.
65. Leek JT, Storey JD. A general framework for multiple testing dependence. *Proceedings of the National Academy of Sciences*. 2008; 105(48):18718–18723.
66. Efron B. Correlated z-values and the accuracy of large-scale statistical estimates. *Journal of the American Statistical Association*. 2010; 105(491):1042–1055. [PubMed: 21052523]
67. Schwartzman A, Lin X. The effect of correlation in false discovery rate estimation. *Biometrika*. 2011; 98(1):199–214. [PubMed: 23049127]
68. Fan, J.; Tang, R.; Shi, X. Partial consistency with sparse incidental parameters.
69. Fan, J.; Han, X. Estimation of false discovery proportion with unknown dependence. *Manuscript*

70. Fan J, Lv J. A selective overview of variable selection in high dimensional feature space. *Statistica Sinica*. 2010; 20(1):101–148. [PubMed: 21572976]
71. Bühlmann, P.; Van De Geer, S. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer; 2011.
72. Van de Geer SA. High-dimensional generalized linear models and the lasso. *The Annals of Statistics*. 2008; 36(2):614–645.
73. Fan J, Lv J. Nonconcave penalized likelihood with NP-dimensionality. *Information Theory. IEEE Transactions on*. 2011; 57(8):5467–5484.
74. Zou H, Li R. One-step sparse estimates in nonconcave penalized likelihood models. *Annals of statistics*. 2008; 36(4):1509–1533. [PubMed: 19823597]
75. Friedman J, Hastie T, Höfling H, Tibshirani R. Pathwise coordinate optimization. *The Annals of Applied Statistics*. 2007; 1(2):302–332.
76. Daubechies I, Defrise M, De Mol C. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on pure and applied mathematics*. 2004; 57(11):1413–1457.
77. Beck A, Teboulle M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*. 2009; 2(1):183–202.
78. Efron B, Hastie T, Johnstone I, Tibshirani R. Least angle regression. *The Annals of statistics*. 2004; 32(2):407–499.
79. Fan J, Song R. Sure independence screening in generalized linear models with np-dimensionality. *The Annals of Statistics*. 2010; 38(6):3567–3604.
80. Hall P, Miller H. Using generalized correlation to effect variable selection in very high dimensional problems. *Journal of Computational and Graphical Statistics*. 2009; 18(3):533–550.
81. Fan J, Feng Y, Song R. Nonparametric independence screening in sparse ultra-high-dimensional additive models. *Journal of the American Statistical Association*. 2011; 106(494):544–557. [PubMed: 22279246]
82. Zhao SD, Li Y. Principled sure independence screening for cox models with ultra-high-dimensional covariates. *Journal of multivariate analysis*. 2012; 105(1):397–411. [PubMed: 22408278]
83. Li R, Zhong W, Zhu L. Feature screening via distance correlation learning. *Journal of the American Statistical Association*. 2012; 107(499):1129–1139.
84. Li G, Peng H, Zhang J, Zhu L. Robust rank correlation based screening. *The Annals of Statistics*. 2012; 40(3):1846–1877.
85. Fan J, Samworth R, Wu Y. Ultrahigh dimensional feature selection: beyond the linear model. *The Journal of Machine Learning Research*. 2009; 10:2013–2038.
86. Vapnik V, Golowich SE, Smola A. Support vector method for function approximation, regression estimation, and signal processing. *Advances in neural information processing systems*. 1997:281–287.
87. Freund, Y.; Schapire, R. *Computational learning theory*. Springer; 1995. A decision-theoretic generalization of on-line learning and an application to boosting; p. 23–37.
88. Breiman L. Arcing classifier (with discussion and a rejoinder by the author). *The annals of statistics*. 1998; 26(3):801–849.
89. Witten DM, Tibshirani R. Covariance-regularized regression and classification for high dimensional problems. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2009; 71(3):615–636.
90. Shao J, Wang Y, Deng X, Wang S. Sparse linear discriminant analysis by thresholding for high dimensional data. *The Annals of Statistics*. 2011; 39(2):1241–1265.
91. Cai T, Liu W. A direct estimation approach to sparse linear discriminant analysis. *Journal of the American Statistical Association*. 2011; 106(496):1566–1577.
92. Fan J, Feng Y, Tong X. A road to classification in high dimensional space: the regularized optimal affine discriminant. *Journal of the Royal Statistical Society: Series B*. 2012:745–771.
93. Wille A, Zimmermann P, Vranova E, Frholz A, Laule O, Bleuler S, Hennig L, Prelic A, von Rohr P, Thiele L, Zitzler E, Gruissem W, Bühlmann P. Sparse graphical gaussian modeling of the

- isoprenoid gene network in arabidopsis thaliana. *Genome Biology*. 2004; 5(11):R92. [PubMed: 15535868]
94. Zhao T, Liu H, Roeder K, Lafferty J, Wasserman L. The huge package for high-dimensional undirected graph estimation in r. *Journal of Machine Learning Research*. 2012; 13:1059–1062.
95. Hedenfalk I, Duggan D, Chen Y, Radmacher M, Bittner M, Simon R, Meltzer P, Gusterson B, Esteller M, Raffeld M, et al. Gene-expression profiles in hereditary breast cancer. *New England Journal of Medicine*. 2001; 344(8):539–548. [PubMed: 11207349]
96. Efron B. Correlation and large-scale simultaneous significance testing. *Journal of the American Statistical Association*. 2007; 102(477):93–103.
97. Desai KH, Storey JD. Cross-dimensional inference of dependent high-dimensional data. *Journal of the American Statistical Association*. 2012; 107(497):135–151.
98. Fang, K.; Kotz, S.; Ng, K. Symmetric multivariate and related distributions. Chapman & Hall; 1990.
99. Liu H, Lafferty J, Wasserman L. The nonparanormal: Semiparametric estimation of high dimensional undirected graphs. *Journal of Machine Learning Research*. 2009; 10:2295–2328.
100. Liu H, Han F, Yuan M, Lafferty J, Wasserman L. High dimensional semiparametric gaussian copula graphical models. *Annals of Statistics*. 2012; 40(40):2293–2326.
101. Xue L, Zou H. Regularized rank-based estimation of high-dimensional nonparanormal graphical models. *Annals of Statistics*. 2012; 40(5):2541–2571.
102. Liu H, Han F, Zhang C. Transelliptical graphical models. *Advances in Neural Information Processing Systems*. 2012; 25:809–817.

Figure 1. Sorted eigenvalues of the sample covariance matrix $\mathbf{S}$ (black curve) and that of the population covariance matrix $\mathbf{\Sigma}$ (dashed red line). In the simulation, we always use $d = 1,000$ but different ratios of d/n . Subfigures correspond to $d/n = 2$, $d/n = 1$, $d/n = 0.2$ and $d/n = 0.1$.

Figure 2.

Visualization of the penalty functions (above) and the corresponding thresholding operators (below) of Hard-thresholding, Soft-thresholding, SCAD and MCP penalties. In all cases, $\lambda = 1$. For SCAD and MCP, different values of γ are chosen as shown in graphs (Best viewed in color).

Figure 3.

A sparse inverse covariance matrix Θ and the corresponding undirected graph defined by Θ

Figure 4.

(a) and (b): The histogram and normal QQ plots of the marginal expression levels of the gene MECPS. We see that the data are not exactly Gaussian distributed. (c) The estimated gene networks of the Arabidopsis data. The within-pathway links are denoted by solid lines and between-pathway links are denoted by dashed lines. Adapted from [18].

(a) Results from the POET-PFA method

(b) Results from the dependence-adjusted method

Figure 5.

The estimated false discovery proportion and the estimated number of false discoveries as functions of the number of total discoveries for $d = 3, 226$ genes, where the number of factors are chosen to be $k = 2, 3, 4, 5$. (a) Results from the POET-PFA method; (b) Results from the dependence-adjusted method. Adapted from [69].