Published in final edited form as:

Stat Sin. 2014 October ; 24(4): 1633–1654. doi:10.5705/ss.2013.019.

VARIABLE SELECTION FOR HIGH DIMENSIONAL MULTIVARIATE OUTCOMES

Tamar Sofer,

Department of Biostatistics, Harvard School of Public Health, Boston, MA, 02115, USA

Lee Dicker, and

Department of Statistics and Biostatistics, Rutgers University, Piscataway, NJ 08854, USA

Xihong Lin

Department of Biostatistics, Harvard School of Public Health, Boston, MA, 02115, USA

Abstract

We consider variable selection for high-dimensional multivariate regression using penalized likelihoods when the number of outcomes and the number of covariates might be large. To account for within-subject correlation, we consider variable selection when a working precision matrix is used and when the precision matrix is jointly estimated using a two-stage procedure. We show that under suitable regularity conditions, penalized regression coefficient estimators are consistent for model selection for an arbitrary working precision matrix, and have the oracle properties and are efficient when the true precision matrix is used or when it is consistently estimated using sparse regression. We develop an efficient computation procedure for estimating regression coefficients using the coordinate descent algorithm in conjunction with sparse precision matrix estimation using the graphical LASSO (GLASSO) algorithm. We develop the Bayesian Information Criterion (BIC) for estimating the tuning parameter and show that BIC is consistent for model selection. We evaluate finite sample performance for the proposed method using simulation studies and illustrate its application using the type II diabetes gene expression pathway data.

Keywords

BIC; Consistency; Correlation; Efficiency; Model selection; Multiple outcomes; Oracle Estimator

1. Introduction

Correlated multivariate responses are often observed in health science studies where the number of responses per subject may be large and correlated, and numerous covariates may be observed for each subject. For instance, in genetic pathway studies, one is often interested in associating gene expressions in a genetic pathway to such clinical covariates as exposure, treatment, and individual characteristics. A pathway may have tens and hundreds of genes and only a small number of gene expressions are likely to be associated with exposures. A question of particular interest is to identify a subset of genes that are associated with exposures while accounting for the fact that the gene expressions with the same pathway are likely to be correlated and the number of genes is large relative to the sample size. In this paper, we consider the problem of variables selection in the presence of multivariate

outcomes. We allow the numbers of outcomes and covariates to be large and the outcomes to be correlated.

There is a vast literature on variable selection for independent data with methods often based on penalized likelihoods. Examples include LASSO (Tibshirani (1996)), SCAD (Fan and Li (2001)), adaptive LASSO (Zou (2006)), SELO (Dicker and Lin (2012)) and MC+ (Zhang (2010)). A desirable property for variable selection is the oracle property (Fan and Li (1991)). It is known that LASSO is often not consistent for variable selection, while SCAD, Adaptive LASSO, SELO and MC+ all enjoy the oracle property. These methods have been extended to the situations where the number of independent variables increases with sample size (Fan and Peng (2004)). Tuning parameter selection is important for variable selection. BIC has been proposed for selecting the tuning parameter and has been shown to be consistent for variable selection (Wang (2009); Dicker and Lin (2012)).

Compared to independent data, a challenge in variable selection for multivariate outcomes is that multiple outcomes of each subject are generally correlated. If outcomes have the same unknown mean, sparse precision matrix (inverse covariance matrix) estimation has been a fertile research area. For instance, Schäfer and Strimmer (2005) proposed shrinkage estimators, while Meinshausen and Bühlmann (2006) and Yuan and Lin (2007) proposed Lasso penalized estimators for the covariance or precision matrix. Lam and Fan (2009) established asymptotic theory for inverse covariance estimation when the outcomes have mean zero. In general, this work dealt with mean zero outcome variables, and did not consider the effect of diverging number of mean parameters on the estimation of the precision matrix.

Limited research has been done on variable selection for regression coefficients for multiple outcomes. Several authors have proposed methods for variable selection under the assumption that a covariate either affects all or none of the outcomes predictors (Brown, Frean and Vannucci (1999); Turlach, Venables and Wright (2005); Peng et al. (2010)). These authors implicitly assumed working independence among outcomes by ignoring between-outcome correlation in their procedure, and did not present large sample properties of the resulting estimators. Rothman et al. (2010) proposed an iterative algorithm for variable selection and estimation in high-dimensional multivariate regression that alternately estimates the inverse covariance matrix and regression parameters under the ℓ_1 penalty. Asymptotic properties of these estimators were not provided.

We develop variable selection and estimation procedures for regression coefficients in a multivariate regression model when the number of outcomes and the number of regression coefficients are likely to be large. We first propose variable selection for a given working precision matrix, that allows the within-subject correlation to be misspecified. We then propose a “two-stage” estimation procedure to jointly estimate regression coefficients and the precision matrix using penalized likelihood methods, that results in efficiency gains in regression coefficient estimators.

We focus on variable selection for regression coefficients in multivariate regression while most existing literature has focused on covariance estimation. Second, unlike Rothman,

Levina and Zhu (2010) who primarily focused on the LASSO penalty for variable selection of regression coefficients, we can accommodate a wide range of concave penalty functions. We show that our regression parameter estimator is consistent for an arbitrary working precision matrix, has the oracle property, if an appropriate penalty is used, when the true precision matrix is used or is consistently estimated, and that the number of parameters might diverge to ∞ . We propose a BIC criterion for tuning parameter selection in the multivariate penalized regression problem, and show that it is consistent for selecting the true regression model even if the number of parameters diverges. Our “two-stage” algorithm is computationally more efficient than the MRCE algorithm introduced in Rothman et al. (2010).

The rest of the paper is organized as follows. We describe the model in Section 2. In Section 3.1, we discuss sparse multivariate regression estimation using an arbitrary working precision matrix. In Section 3.1.1, we propose a two-stage joint estimation procedure for regression parameters and the precision matrix, and a computationally efficient estimation algorithm. We study its properties and show model selection consistency can be achieved using BIC. In Section 6, we present simulation studies that evaluate the finite sample performance of the proposed methods, compare the use of different penalties in penalized multivariate regression, and demonstrate the performance of the BIC. In Section 7, we apply the proposed method to the analysis of type II diabetes gene expression data, and there is a discussion in Section 8.

2. The model

Consider a multivariate regression model for the i th subject

$$\mathbf{y}_i = \mathbf{X}_i \boldsymbol{\beta} + \boldsymbol{\varepsilon}_i, \quad i=1, \dots, n \quad (2.1)$$

where $\mathbf{y}_i$ and $\boldsymbol{\varepsilon}_i$ are m -dimensional vectors of outcomes and (unobserved) errors, respectively, $\mathbf{X}_i = \mathbf{x}_i^T \otimes \mathbf{I}_m$ is a $m \times p$ dimensional matrix of covariates, where $p = mp_0$, and $\mathbf{x}_i$ is a $p_0 \times 1$ covariate vector. Let $\boldsymbol{\beta}_{ok}$ be a vector of p_0 -dimensional regression parameters for the k th outcome, so $(\boldsymbol{\beta}_{01}^T, \dots, \boldsymbol{\beta}_{0m}^T)^T = \boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$ is a p -dimensional vector of unknown regression parameters. We assume that $\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_n$ are iid, with $E(\boldsymbol{\varepsilon}_i) = 0$ and $\text{Cov}(\boldsymbol{\varepsilon}_i) = \boldsymbol{\Sigma}_m$, for some unknown $m \times m$ covariance matrix $\boldsymbol{\Sigma}_m$. The total number of subjects is n and the total number of observations is $N = nm$. We will take $n \rightarrow \infty$. Our asymptotic framework will also allow $m \rightarrow \infty$ and $p_0 \rightarrow \infty$. We first require in Section 4 that p grows more slowly than n , and then extend the asymptotic results to $p > n$ in Section 5. Since p changes with n , it is implicit that $\boldsymbol{\beta}$ may vary with n as well.

Let $A = \{j; \beta_j \neq 0\}$ be the subset of $\{1, \dots, p\}$ corresponding to the nonzero entries of $\boldsymbol{\beta}$. We sometimes refer to A as the “true model.” Let $\mathbf{Y} = (\mathbf{y}_1^T, \dots, \mathbf{y}_n^T)^T \in \mathbb{R}^N$, $\mathbf{X} = (\mathbf{X}_1^T, \dots, \mathbf{X}_n^T)^T$, and $\boldsymbol{\varepsilon} = (\boldsymbol{\varepsilon}_1^T, \dots, \boldsymbol{\varepsilon}_n^T)^T \in \mathbb{R}^N$, then (2.1) may be rewritten as $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$, where $\mathbf{Y}$ and $\boldsymbol{\varepsilon}$ are $N \times 1$ vectors, $\boldsymbol{\beta}$ is a $p \times 1$ vector and $\mathbf{X}$ is a $N \times p$ matrix. Let $\boldsymbol{\Sigma}$ be the

$N \times N$ block diagonal matrix whose i -th diagonal component is $\Sigma_m(1 \leq i \leq n)$. We write the true precision matrix $\Omega_m = (\omega_{ij}) = \Sigma_m^{-1}$ and $\Omega = \Sigma^{-1}$.

If $\mathbf{e}_i \sim \mathcal{N}(0, \Sigma_m)$, the joint log-likelihood for estimating β and Ω_m is, apart from a constant,

$$\begin{aligned} \ell(\beta, \Omega_m) &= \ln |\Omega_m| - \frac{1}{n} \sum_{i=1}^N (\mathbf{y}_i - \mathbf{X}_i \beta)^T \Omega_m (\mathbf{y}_i - \mathbf{X}_i \beta) \\ &= \ln |\Omega_m| - \text{tr} \left\{ \hat{\Sigma}_m(\beta) \Omega_m \right\} \end{aligned} \quad (2.2)$$

where $\hat{\Sigma}_m(\beta) = n^{-1} \sum_{i=1}^n (\mathbf{y}_i - \mathbf{X}_i \beta)(\mathbf{y}_i - \mathbf{X}_i \beta)^T$. We do not need to assume that the $\mathbf{e}_i$ are normally distributed, but (2.2) is basic to our method.

3. Estimation Procedures for Variable Selection

In Section 3.1, we propose an estimation procedure for variable selection given a working precision matrix that might be misspecified; in Section 3.1.1 we introduce joint estimation of the regression parameters and the precision matrix; in Section 3.1.2, we propose a two-stage joint estimation procedure.

3.1. Estimation of regression parameters using a working precision matrix

Let

$$Q(\beta | \Lambda_m) = \frac{1}{n} (\mathbf{Y} - \mathbf{X}\beta)^T \Lambda (\mathbf{Y} - \mathbf{X}\beta) + 2 \sum_{j=1}^p P_\lambda(|\beta_j|), \quad (3.1)$$

where $P_\lambda : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is a concave and differentiable penalty function indexed by a nonnegative tuning parameter $\lambda \geq 0$ satisfying $P_\lambda^{(1)}(0) = 0$, with Λ a working precision matrix that might misspecify the true precision matrix Ω . This objective function might be referred to as a penalized general likelihood. By transforming outcomes and covariates using $\tilde{\mathbf{Y}} = \Lambda^{1/2} \mathbf{Y}$ and $\tilde{\mathbf{X}} = \Lambda^{1/2} \mathbf{X}$, we can minimize $Q(\beta | \Lambda_m) = n^{-1} (\tilde{\mathbf{Y}} - \tilde{\mathbf{X}}\beta)^T (\tilde{\mathbf{Y}} - \tilde{\mathbf{X}}\beta) + 2 \sum_{j=1}^p P_\lambda(|\beta_j|)$ by the penalized likelihood methods for independent data.

Several penalty functions have been proposed for variable selection and estimation for independent data. Examples are as follows:

- (P1) The L_1 penalty, $P_\lambda(|\theta|) = \lambda/|\theta|$, also known as the Lasso (Tibshirani (1996)).
- (P2) The Adaptive Lasso, $P_\lambda(|\theta|) = \lambda w^{-1} |\theta|$ for some data-dependent weight w (Zou (2006)).
- (P3) The SCAD penalty has $P_\lambda(0) = 0$ and the derivative

$$P'_{\lambda}(|\theta|) = \lambda \left\{ I(|\theta| \leq \lambda) + \frac{(a\lambda - |\theta|)^+}{(a-1)\lambda} I(|\theta| > \lambda) \right\},$$

where $a > 0$ is another tuning parameter often taken to be $a = 3.7$ (Fan and Li (2001)).

(P4) The seamless- L_0 (SELO) penalty

$$P_{\lambda}(\theta) = \frac{\lambda}{\log(2)} \log \left\{ \frac{|\theta|}{|\theta| + \tau} + 1 \right\},$$

where $\tau > 0$ is another tuning parameter often taken to be $\tau = 0.01$ when x 's are standardized (Dicker and Lin, 2012).

All these penalties satisfy $\lim_{|\theta| \rightarrow 0} P'_{\lambda}(|\theta|) > 0$ provided $\lambda > 0$, property that ensures sparsity of the estimated parameters.

To estimate standard errors for the regression parameters, we used a quadratic approximation of the penalty function (Fan and Li (2001)). Standard errors are estimated only for the parameters that are estimated as non-zeros, and are set to be zero if the estimates are zero. Let $\hat{A}$ be a set of indices with estimated non-zero parameters. Let $\beta_{\hat{A}}$ be the parameter sub-vector and $\mathbf{X}_{\hat{A}}$ be the covariate sub-matrix that corresponds to these indices.

Given λ , one can show that the estimated covariance matrix of $\hat{\beta}_{\hat{A}}$ is given by the sandwich estimator (Liang and Zeger (1986))

$$\widehat{cov}(\hat{\beta}_{\hat{A}}) = \left\{ \mathbf{X}_{\hat{A}}^T \mathbf{\Lambda} \mathbf{X}_{\hat{A}} + n \sum_{\lambda} (\hat{\beta}_{\hat{A}}) \right\}^{-1} \left\{ \sum_{i=1}^n \mathbf{X}_{\hat{A},i}^T \mathbf{\Lambda} \sum_i \mathbf{\Lambda} \mathbf{X}_{\hat{A},i} \right\} \left\{ \mathbf{X}_{\hat{A}}^T \mathbf{\Lambda} \mathbf{X}_{\hat{A}} + n \sum_{\lambda} (\hat{\beta}_{\hat{A}}) \right\}^{-1},$$

(3.2)

where $\sum_i = \sum_i (\hat{\beta}_{\hat{A}}) = (\mathbf{y}_i - \mathbf{X}_{\hat{A},i} \hat{\beta}_{\hat{A}})(\mathbf{y}_i - \mathbf{X}_{\hat{A},i} \hat{\beta}_{\hat{A}})^T$ and, for a set of indices $\mathcal{K}$ of size k , $\dots, \sum_{\lambda} (\beta_{\mathcal{K}}) = \text{diag}\{P'_{\lambda}(|\beta_{\mathcal{K},(1)}|)/|\beta_{\mathcal{K},(1)}|, \dots, P'_{\lambda}(|\beta_{\mathcal{K},(k)}|)/|\beta_{\mathcal{K},(k)}|\}$.

3.1.1. Joint estimation of regression coefficients and precision matrix—Results in Section 4.1 show that misspecification of the precision matrix leads to less efficient estimators of regression coefficients. To improve the efficiency, we propose to estimate β and $\mathbf{\Omega}$ simultaneously. It is well known that although the MLE of the error covariance matrix

$$\mathbf{S} = \sum (\hat{\beta}) = \frac{1}{n} \sum_{i=1}^n (\mathbf{y}_i - \mathbf{X}_i \hat{\beta})(\mathbf{y}_i - \mathbf{X}_i \hat{\beta})^T$$

is consistent, it is not positive definite when the number of outcomes m is larger than the number of observations, $m > n$, and even if $m < n$ and the number of parameters is large, the MLE is unstable (Schäfer and Strimmer (2005)). Likewise the estimator of the precision matrix $\mathbf{S}^{-1}$, is unstable, if it exists. This motivates regularized estimation of the precision matrix $\mathbf{\Omega}$ as well as of $\boldsymbol{\beta}$. Let $P_\gamma: \mathbb{R}^+ \rightarrow \mathbb{R}^+$ be a concave, differentiable penalty function, $\gamma > 0$, $P_\gamma(0) = 0$. We jointly estimate $\boldsymbol{\beta}$ and $\mathbf{\Omega}$ by minimizing the objective function

$$Q(\boldsymbol{\beta}, \mathbf{\Omega}_m) = -\ln|\mathbf{\Omega}_m| + \text{tr} \left[\sum_m (\boldsymbol{\beta}) \mathbf{\Omega}_m \right] + 2 \sum_{j=1}^p P_\lambda(|\beta_j|) + 2 \sum_{1 \leq i < j \leq m} P_\gamma(|\omega_{ij}|)$$

where $P_\gamma^{(2)}(\cdot)$ is a penalty for $\mathbf{\Omega}_m$.

3.1.2. Two-stage joint estimation—Minimizing $Q(\boldsymbol{\beta}, \mathbf{\Omega})$ simultaneously over $\boldsymbol{\beta}$ and $\mathbf{\Omega}$ is computationally difficult as the number of parameters is very large. As an alternative, consider an objective function for estimation of $\mathbf{\Omega}$ given $\boldsymbol{\beta}$:

$$Q(\mathbf{\Omega}|\boldsymbol{\beta}) = -\ln|\mathbf{\Omega}_m| + \text{tr} \left[\sum_m (\boldsymbol{\beta}) \mathbf{\Omega}_m \right] + 2 \sum_{1 \leq i < j \leq m} P_\gamma(|\omega_{ij}|) \quad (3.3)$$

We propose a two-stage procedure for estimation of $\mathbf{\Omega}$ by iteratively minimizing $Q(\mathbf{\Omega}|\hat{\boldsymbol{\beta}})$ in (3.3) by setting $\boldsymbol{\beta} = \hat{\boldsymbol{\beta}}$, and estimating $\boldsymbol{\beta}$ by minimizing $Q(\boldsymbol{\beta}|\hat{\mathbf{\Omega}})$ in (3.1) by setting $\mathbf{\Lambda} = \hat{\mathbf{\Omega}}$.

Suppose the tuning parameters λ and γ are known. At the first stage, using the results in Section 3.1, we estimate a consistent estimator $\hat{\boldsymbol{\beta}}^{(1)}$ by assuming a working independent precision matrix; we then estimate $\mathbf{\Omega}$ using the obtained $\hat{\boldsymbol{\beta}}^{(1)}$ in minimization of $Q(\mathbf{\Omega}|\hat{\boldsymbol{\beta}})$ and call this estimator $\hat{\mathbf{\Omega}}^{(1)}$. At a second stage, we improve efficiency of estimation of $\boldsymbol{\beta}$ using the consistently-estimated $\hat{\mathbf{\Omega}}^{(1)}$ from stage 1 and call this estimator $\hat{\boldsymbol{\beta}}$. We further improve upon estimation of $\mathbf{\Omega}$ by using the efficient estimate of $\hat{\boldsymbol{\beta}}$. Hence, at the first stage consistent estimators of $\boldsymbol{\beta}$ and $\mathbf{\Omega}$ are calculated; at the second, their efficient counterparts are obtained.

The tuning parameters are not known and can be estimated e.g, using BIC or data validation procedure (see Section 3.2). We propose a two-stage joint estimation procedure that is computationally efficient and show that it has desirable properties.

Stage 1: Set $\hat{\mathbf{\Omega}}_m^{(0)} = \mathbf{I}_m$.

- Let $\hat{\boldsymbol{\beta}}^{(1)} = \hat{\boldsymbol{\beta}}^{(1)}(\hat{\lambda})$ be a model selection consistent working independence estimator of $\boldsymbol{\beta}$, where the penalty satisfies the regularity conditions in Section 3, and $\hat{\lambda}$ is estimated by some selection criterion, see Section 3.2.

- Set $\hat{\Omega}^{(1)} = \hat{\Omega}(\hat{\gamma})$ to be the sparse consistent estimator of Ω given $\hat{\beta}^{(1)}$ by minimizing $Q(\Omega|\hat{\beta}^{(1)})$, where the tuning parameter $\hat{\gamma} = \hat{\gamma}(\hat{\lambda})$ is estimated using a selection criterion, see Section 3.2.

Stage 2: For each tuning parameter λ :

- Set $\hat{\beta}(\lambda)$ to be the minimizer of $Q(\beta|\hat{\Omega}^{(1)})$, with $P_\lambda(\cdot)$ an oracle penalty function.
- Set $\hat{\Omega}(\lambda) = \hat{\Omega}\{\hat{\gamma}(\lambda)\}$ to be the minimizer of $Q\{\Omega|\hat{\beta}(\lambda)\}$, with $P_\gamma(\cdot)$, where $\hat{\gamma}(\lambda)$ for given λ is selected using a selection criterion.

Choose “the best” estimator $(\hat{\beta}, \hat{\Omega}) = (\hat{\beta}(\hat{\lambda}), \hat{\Omega}(\hat{\gamma}(\hat{\lambda})))$ by estimating λ using a selection criterion, see Section 5.

The first stage provides simple consistent estimators of β and Ω , the second stage calculates their efficient counterparts. The proposed “profile” tuning parameter estimation at the second stage mimics the profile likelihood idea and is computationally efficient. For instance, when performing a double grid search, one has to iteratively estimate both β and Ω for each combination of tuning parameters (λ, γ) . Here, not all combinations of tuning parameter values need to be searched, and less iterations are required for each estimated combination of the tuning parameter values.

We can use the GLASSO procedure (Friedman et al. (2008)) to minimize $Q(\Omega|\hat{\beta})$ if the L_1 penalty is used. Tuning parameter selection is discussed in Section 3.2. Standard errors of $\hat{\beta}$ can be estimated by a formula similar to (3.2). In Section 4.2, we show that the two-stage procedure produces the oracle estimators of (β, Ω) .

3.2. Tuning parameter selection

We propose BIC criteria for selecting the tuning parameter when a working (possibly misspecified) precision matrix is used, or when β and Ω are jointly estimated.

Denote by $\hat{\beta}(\lambda)$ the estimated regression coefficient given the tuning parameter λ . Let $\hat{s}$ be the number of estimated non-zero coefficients, $|\{j: \hat{\beta}_j \neq 0\}|$, and $k_n \rightarrow \infty$. When a known working precision matrix Λ is used for estimating β as in Section 3.1., the BIC for choosing λ is

$$BIC\{\hat{\beta}(\lambda)\} = \text{tr} \left[\sum_m (\hat{\beta}(\lambda)) \Lambda_m \right] + \hat{s} \frac{k_n}{n}. \quad (3.4)$$

This BIC is also appropriate when Λ is a consistent estimator of Ω , rather than a fixed matrix. This is relevant, for instance, when one first estimates Ω based on some consistent estimator of β and then uses the single initial estimator of Ω , denoted by Λ , to estimate β without iterations or a two dimensional grid search.

When β is jointly estimated with Ω for a given λ , $\hat{\Omega} = \hat{\Omega}\{\hat{\gamma}(\lambda)\}$ is a function of λ . This since, given $\hat{\beta}(\lambda)$, the tuning parameter γ is estimated as a function of $\hat{\beta}(\lambda)$ by minimizing with respect to γ (Gao et al, 2009)

$$BIC(\hat{\Omega}_m(\gamma)|\lambda) = -\log|\hat{\Omega}_m(\gamma)| + \text{tr} \left[\sum_m \{\hat{\beta}(\lambda)\} \hat{\Omega}_m(\gamma) \right] + \hat{t} \frac{\log(n)}{n}, \quad (3.5)$$

where $\sum_m$ is the sample covariance matrix and $\hat{t}$ is the estimated number of nonzero entries in the upper off-diagonal matrix of $\hat{\Omega}_m(\gamma)$. We denote the resulting estimator as $\hat{\gamma}(\lambda)$ and $\hat{\Omega}_m(\lambda) = \hat{\Omega}_m\{\hat{\gamma}(\lambda)\}$. Hence we define the BIC for λ when Ω is estimated together with β as

$$BIC(\hat{\beta}(\lambda)) = -\log(|\hat{\Omega}_m(\lambda)|) + \text{tr} \left[\hat{\Omega}_m(\lambda) \sum (\hat{\beta}(\lambda)) \right] + \hat{s} \frac{k_n}{n}, \quad (3.6)$$

The optimal tuning parameter λ is estimated by minimizing (3.6).

4. Asymptotic results when $p < n$

We first study the properties of the penalized multivariate regression estimator of β assuming an arbitrary working precision matrix, when the penalty function satisfies the oracle properties (Fan and Peng (2004)). We then extend the results to joint estimation of β and Ω . The proofs are given in the supplementary material.

4.1. Estimation of regression coefficients when the working precision matrix is given

We show that for an arbitrary working precision matrix, $\hat{\beta}$ is consistent and sparsistent, where “sparsistency” means that every true zero parameter is estimated as zero with probability tending to 1, uniformly over all parameters (Lam and Fan (2009)). We first provide results for the case where the number of parameters is smaller than the sample size, $p/n \rightarrow 0$; in this case we also show that $\hat{\beta}$ is asymptotically normally distributed and is most efficient when the working precision matrix Λ is correctly specified as Ω_m . We first outline the regularity conditions.

(C1) $n \rightarrow \infty$, $p = p_0 m$ may vary with n , and $p/n \rightarrow 0$.

(C2) There exists a positive constant R such that

$$0 < 1/R < \frac{\lambda_{\min}(\Lambda_m), \lambda_{\min}(\Omega_m), \lambda_{\min}(n^{-1} \mathbf{X}^T \mathbf{X})}{\lambda_{\max}(\Lambda_m), \lambda_{\max}(\Omega_m), \lambda_{\max}(n^{-1} \mathbf{X}^T \mathbf{X})} < R < \infty$$

where $\lambda_{\min}(\mathbf{B})$ and $\lambda_{\max}(\mathbf{B})$ are the smallest and largest eigenvalues of a matrix $\mathbf{B}$.

(C3) If $\rho = \rho_n = \min\{|\beta_j|; j \in A\}$. $\rho / \sqrt{p/n} \rightarrow \infty$

$$(C4) \max_{1 \leq i \leq n} n^{-1} \|x_i x_i^T\|_2 \rightarrow 0, \text{ where } \|x_i x_i^T\|_2 = \|x_i\|_2^2$$

$$(C5) \text{ There exists a } \delta \text{ such that } E(\varepsilon_{ij}^{2+\delta}) < \infty, i = 1, \dots, n, j = 1, \dots, m.$$

$$(C6) \text{ The function } P_\lambda \text{ is concave on } [0, \infty) \text{ and differentiable on } (0, \infty), \text{ with } P_\lambda(0) =$$

$$0, P_\lambda(\theta) = P_\lambda(-v\theta) \text{ and } \lim_{\theta \rightarrow \infty} P_\lambda(\theta) \leq \frac{1}{n}.$$

$$(C7) \text{ If } r_n / \sqrt{p/n} \rightarrow \infty, \text{ then } P'_\lambda(r_n) = o(1/\sqrt{np}).$$

$$(C8) \text{ Let } k_n \text{ be such that } p/k_n^{\frac{2+\delta}{2}} \rightarrow 0. \text{ If } r_n = O(\sqrt{p/n}), \text{ then}$$

$$\lim_{n \rightarrow \infty} (\sqrt{n/\max(p, k_n)}) P'_\lambda(r_n) \rightarrow \infty$$

Conditions (C1) – (C5) are related to the likelihood function. Conditions (C6) – (C8) are related to the penalty function. The latter two conditions are specified so that the behavior of the derivative of the penalty function that is “close to zero” and “far from zero” in some rate sense behaves according to the rate of convergence of the estimators. These requirements are guaranteed to hold for the oracle penalty SELO, for instance, for some specifications of sequences of the tuning parameters.

Specifically, condition (C1) bounds the rate of the number of covariates p . Condition (C2) guarantees the stability of the estimator. Condition (C3) sets a required bound on the smallest parameter in the model and permits model selection consistency. Condition (C4) gives a very weak bound on the covariates. Notice that it could equivalently be written as $\max_{1 \leq i \leq n} n^{-1} x_i^T x_i \rightarrow 0$ which is not restrictive as covariates, for instance, are often naturally bounded. The last requirement on the likelihood function, (C5) is important for asymptotic normality of the regression coefficient estimators.

Under these conditions, the estimator $\hat{\beta}$ of the regression parameters obtained by minimizing $Q(\beta|\Lambda)$ in (3.1) has the oracle properties.

Theorem 1—*Suppose (C1) – (C8) hold. More specifically, if β^* is the true vector of regression parameters, then the following hold.*

a. (Consistency) *For every $c > 0$, there exists a positive constant M such that*

$$\liminf_{n \rightarrow \infty} P \left[\text{There exist a local minima } \hat{\beta} \text{ of } Q(\beta|\Lambda_m) \text{ such that} \right. \\ \left. \|\hat{\beta} - \beta^*\| < M \sqrt{p/n} \right] > 1 - c.$$

*Further, a global minimizer of $Q(\beta|\Lambda_m)$ is $\sqrt{p/n}$ consistent for β^**

b. (Sparsistency) *For every fixed $M > 0$,*

$$\lim_{n \rightarrow \infty} P \left[\left\{ \hat{\beta} \text{ is a local minimum of } L(\beta) \text{ and } \left\| \hat{\beta} - \beta^* \right\| \leq M \sqrt{p/n} \right\} \cap \left\{ \hat{\beta}_{A^c} \neq 0 \right\} \right] = 0.$$

- c. (Asymptotic normality) Let $s = |A| = |\{j; \beta_j^* \neq 0\}|$. Let $q \in \mathbb{N}$ be fixed and $B = B_n$ be a sequence of $q \times s$ matrices such that $BB^T \rightarrow G$, for some $q \times q$ symmetric matrix G . There exists a sequence of local minima, $\hat{\beta}$, of $Q(\beta; \Lambda_m)$ such that $\lim_{n \rightarrow \infty} P(\{j; \hat{\beta}_j \neq 0\} = A) = 1$, and

$$B \Upsilon^{-1/2} (\hat{\beta}_A - \beta_A^*) \xrightarrow{D} N(0, G),$$

where $\Upsilon = (X_A^T \Lambda X_A)^{-1} X_A^T \Lambda \sum \Lambda X_A (X_A^T \Lambda X_A)^{-1}$. If $\Lambda = \Omega$, then $\Upsilon = (X_A^T \Omega X_A)^{-1}$, and $\hat{\beta}$ is the most efficient estimator.

4.2. Joint estimation of regression coefficients and precision matrix

We show consistency and sparsistency of the two-stage estimator $(\hat{\beta}, \hat{\Omega})$. We first state additional regularity conditions required for estimation of the precision matrix. Here $P_\gamma(\cdot)$ is the penalty function used for the entries ω_{ij} , $i, j = 1, \dots, m$ of the precision matrix. Let $B = \{(i, j) | \omega_{ij} \neq 0, i < j\}$, the set of true non-zero entries in Ω , and $t = |B|$. Let $\tau = \tau_n = \min\{|\omega_{ij}| : (i, j) \in B\}$. Assume that $P_\gamma(\cdot)$ is an oracle penalty function that satisfies condition (C6), and the following.

$$(C9) \quad pm, m^2/n \rightarrow 0.$$

$$(C10) \quad \sup_{j,k} E(\varepsilon_{ij} \varepsilon_{ik})^2 < \infty.$$

$$(C11) \quad \text{If } r_n / \sqrt{(m+p)m/n} \rightarrow \infty, \text{ then } mP'_\gamma(r_n) = O(1).$$

$$(C12) \quad \text{If } r_n = O(\sqrt{(m+p)m/n}), \text{ then } \lim_{n \rightarrow \infty} \sqrt{(m+p)m/n} P'_\gamma(r_n) = \infty.$$

Conditions (C11) and (C12) are satisfied under a proper selection of the sequence $\gamma_n = \gamma$. Let $\|A\|_F$ be the Frobenius norm of the matrix A .

Lemma 1—Suppose that $\hat{\beta} = \hat{\beta}_n$ is a sequence of estimators such that

$\|\hat{\beta} - \beta^*\| = O_P(\sqrt{p/n})$, and that (C1)–(C8) hold.

a. If

$$\widehat{\sum}_m = \frac{1}{n} \sum_{i=1}^n (y_i - X_i \hat{\beta})(y_i - X_i \hat{\beta})^T, \text{ then } \|\widehat{\sum}_m - \sum_m\|_F = O_P\left\{\sqrt{(m+p)m/n}\right\}.$$

- b. If (C9)–(C12) also hold. There exists a $\sqrt{(m+p)m/n}$ consistent local minimizer, $\hat{\Omega}_m$ of $Q(\Omega_m | \sum_m)$.
- c. Any local minimizer of $Q(\Omega_m | \sum_m)$ satisfies $\|\hat{\Omega} - \Omega\|_F^2 = O_p((m+p)m/n)$, and $\hat{\omega}_{jj'} = 0$ with probability 1 if $\omega_{jj'} = 0$, for all $(i, j) \in B$.

Our results show that estimation of regression coefficients β increases the rate of convergence of the precision matrix estimator compared to that in non-regression settings (assuming mean 0 for y). We make a more robust assumption on the error distribution, and thus the estimator $\hat{\Omega}$ requires a smaller rate on the number of outcomes that could otherwise be achieved.

Theorem 2—Let $(\hat{\beta}, \hat{\Omega})$ be the estimators of (β, Ω) obtained by the two-stage procedure.

- i. If (C1)–(C8) hold, then $\hat{\beta}$ is $\sqrt{p/n}$ -consistent for β , $\hat{\beta}_j = 0$ with probability tending to 1 for all j such that $\beta_j = 0$, also, $\hat{\beta}$ is asymptotically normally distributed with parameters as in Theorem 1(c).
- ii. If also (C9)–(C12) hold, then $\hat{\Omega}$ is $\sqrt{(m+p)m/n}$ -consistent for Ω , and $\hat{\omega}_{ij} = 0$ with probability tending to 1 for all i, j such that $\omega_{ij} = 0$. If B is a sequence of matrices as in Theorem 1, then $B\Upsilon^{-1/2}(\hat{\beta}_A - \beta_A^*) \xrightarrow{D} N(0, G)$, where $\Upsilon = (X_A^T \Omega X_A)^{-1}$.

4.3. Model selection consistency of the BIC criteria for regression parameters

We show that, under regularity conditions, minimization of the BIC criterion leads to selection of the true model for β asymptotically, either when Λ_m is an arbitrary working precision matrix or when $\Lambda_m = \hat{\Omega}_m$ is jointly consistently estimated. Let $x_n = \Theta(y_n)$ denote sequences satisfying $x_n/y_n \rightarrow \infty$.

Theorem 3

- a. Suppose (C1)–(C8) hold, with $\beta_j \neq 0 \iff |\beta_j| = \Theta(\sqrt{\max\{k_n, p\}/n})$. Let $\hat{\beta}_A$ be a minimizer of $Q(\beta | \Lambda)$ for any tuning parameter value λ such that the non-zero entries in $\hat{\beta}_A$ correspond to a model $\hat{A} \subseteq A$. Then if Λ is a fixed positive definite matrix,

$$P \left\{ \sup_{\hat{A} \neq A, \hat{A} \in \mathcal{A}} \left(BIC(\hat{\beta}_A) - BIC(\hat{\beta}_A) \right) < 0 \right\} \rightarrow 1 \text{ as } n \rightarrow \infty.$$

- b. If Λ_m is a consistent estimator of Ω_m , obtained using any $\sqrt{p/n}$ -consistent estimator of the mean $X\hat{\beta}$ of Y .

$$P \left\{ \sup_{\hat{A} \neq A, \hat{A} \in \mathcal{A}} (BIC(\hat{\beta}_{\hat{A}}) - BIC(\hat{\beta}_A)) < 0 \right\} \rightarrow 1 \text{ as } n \rightarrow \infty.$$

This result shows that the BIC in (3.4) gives an estimator of β that is consistent for model selection when an arbitrary working precision matrix Λ is used, and in case Λ is a consistent estimator of the precision matrix.

5. Asymptotic results when $p > n$

The use of a working covariance matrix leads to an estimator of the regression coefficient vector that is consistent under suitable regularity conditions. We modify the conditions for the rate that $p/n \rightarrow 0$ to the following.

(C1') With $m < n < p$, $\log(p)/n \rightarrow 0$. The true model size $s = |A|$ satisfies $s < n$ and

$$\frac{s}{\log(p)} \rightarrow 0.$$

(C2') The eigenvalues of the positive definite matrices $\frac{1}{n} \mathbf{X} \mathbf{X}^T$, $\frac{1}{n} \mathbf{X}_A^T \mathbf{X}_A$ satisfy

$$0 < R^{-1} < \lambda_{\min} \left(\frac{1}{n} \mathbf{X}_A^T \mathbf{X}_A \right), \lambda_{\min} \left(\frac{1}{n} \mathbf{X} \mathbf{X}^T \right) < \lambda_{\max} \left(\frac{1}{n} \mathbf{X}_A^T \mathbf{X}_A \right),$$

$$\lambda_{\max} \left(\frac{1}{n} \mathbf{X} \mathbf{X}^T \right) < R < \infty.$$

$$(C3') \min_{j: \beta_j \neq 0} \frac{\beta_j}{\sqrt{\log(p)/n}} \rightarrow \infty.$$

$$(C7') \text{ If } \lim_n \frac{r_n}{\sqrt{\log(p)/n}} = \infty, \text{ then } n \sqrt{\log(p)/n} P'_\lambda(r_n) = o(1).$$

$$(C8') \text{ If } \lim_n \frac{r_n}{\sqrt{\log(p)/n}} \leq c, \text{ then } P'_\lambda(r_n)/m \rightarrow \infty.$$

$$(C10') \text{ The errors are normal, } \varepsilon \sim \mathcal{N}(\mathbf{0}, \Sigma).$$

Theorem 4

Suppose that Λ_m is given, and that conditions (C1') – (C8') hold. with the above modifications. Let β^* be the true vector of regression parameters.

a. (Consistency) For every $c > 0$, there exists a positive constant M such that

$$\liminf_{n \rightarrow \infty} P \left[\begin{array}{l} \text{There exist a local minima } \hat{\beta} \text{ of } Q(\beta | \Lambda_m) \\ \text{such that } \|\hat{\beta} - \beta^*\| < M \sqrt{\log(p)/n} \end{array} \right] > 1 - c.$$

b. (Sparsistency) For every fixed $M > 0$,

$$\lim_{n \rightarrow \infty} P \left[\left\{ \hat{\beta} \text{ is a local minimum of } L(\beta) \text{ and } \|\hat{\beta} - \beta^*\| \leq M \sqrt{\log(p)/n} \right\} \cap \left\{ \hat{\beta}_{Ac} \neq 0 \right\} \right] = 0.$$

If the number of outcomes satisfies $m^2/n \rightarrow 0$, then these asymptotic properties readily extend to joint estimation of β and Ω . Let (C9)–(C12) be modified, with p replaced by $\log(p)$, to get conditions (C9')–(C12').

Corollary

Let $(\hat{\beta}, \hat{\Omega})$ be the estimators of (β, Ω) obtained by the two-stage procedure.

- i. If (C1')–(C8') hold, then $\hat{\beta}$ is a $\sqrt{\log(p)/n}$ -consistent for β , and $\hat{\beta}_j = 0$ with probability tending to 1 for all j such that $\beta_j = 0$.
- ii. If (C1')–(C12') hold, then $\hat{\Omega}$ is $\sqrt{(m + \log(p))m/n}$ -consistent for Ω , and $\hat{\omega}_{ij} = 0$ with probability tending to 1 for all i, j such that $\omega_{ij} = 0$.

6. Simulation studies

We present simulation results as evaluation of the finite sample performance of the proposed methods. Additional simulation results are provided in the supplementary material.

6.1. Small- p simulations

We consider two types of correlation structures between outcomes: autoregressive (AR) and exchangeable (EX) outcome correlations. In both scenarios the correlation parameter was $\rho_y = 0.5$, and the variance was 9. The precision matrix is thus sparse in the AR case (tridiagonal) and dense in the EX case. In each of the simulations we generated five outcomes and five covariates, so $p = 25$. The parameter vector $\beta_{25 \times 1}$ was generated such that three entries were non-zeros, and their value were 3, 2, and 1.5 (as in Tibshirani (1996)). The predictor $\mathbf{X}$ was simulated by assuming its covariance to be autoregressive with $\rho_x = 0.7$.

For each scenario, $n = 50, 100$, and 200 simulations were run for each sample size. We estimated β using the Lasso, adaptive Lasso, SCAD (with $a = 3.7$) and SELO (with $\tau = 0.01$). The precision matrix Ω was estimated using GLASSO with the Lasso penalty. (The Lasso penalty does not satisfy (C11)–(C12)).

Tuning parameters were selected either by BIC or by a validation data set. For the latter, we minimized prediction error. Thus, for $\mathbf{X}_{valid}$ and $\mathbf{Y}_{valid}$ independent validation data sets, we estimated λ by minimizing

$$\|\mathbf{Y}_{valid} - \mathbf{X}_{valid}\hat{\beta}(\lambda)\|_2 \quad (6.1)$$

Given $\hat{\beta}(\lambda)$, to select the tuning parameter γ for estimating Ω , we chose the γ that maximizes the loglikelihood of the validation data as

$$\log |\hat{\mathbf{\Omega}}(\gamma)| - \text{tr}\{\hat{\mathbf{\Omega}}(\gamma) \sum_{\text{valid}}\}, \quad (6.2)$$

where $\hat{\mathbf{\Omega}}(\gamma) = \sum_{\text{train}}^{-1}(\gamma)$ is estimated from the training data, and $\sum_{\text{valid}}$ is the sample covariance matrix estimated from the validation data set. The same validation data set was used to calculate (6.1) and (6.2). For tuning parameters selection using a validation data set, we used twice the number of observations so that, for $n = 100$, we used 100 observations for parameters estimation and additional 100 observations for validation. For the BIC, we used only 100 observations.

To evaluate the performance of variable selection procedures, we summarize the results with averages of the following measures across the 200 simulations. Model size: $|\{\hat{\beta}_{ij} \neq 0\}|$; the true model size in all simulations was 3. True model chosen: $1(\{i, j: \hat{\beta}_{ij} \neq 0\} = \mathcal{A})$. False positives: $|\{\hat{\beta}_{ij} \neq 0, \beta_{ij} = 0\}|$. False negatives: $|\{\hat{\beta}_{ij} = 0, \beta_{ij} \neq 0\}|$. Model error (Yuan and Lin, (2007)): $\text{tr}[(\hat{\beta} - \beta)^T \sum_{xx} (\hat{\beta} - \beta)]$. Prediction error: for $\mathbf{X}_{\text{new}}, \mathbf{Y}_{\text{new}}$ validation data sets, $\|\mathbf{Y}_{\text{new}} - \mathbf{X}_{\text{new}}\hat{\beta}\|_2$.

For comparison, we also used the MRCE algorithm (Rothman et al. (2008)) to estimate the regression parameters. This algorithm performs a double grid search over the tuning parameter values for the regression parameters and the precision matrix using the Lasso penalty. To select the best combination of tuning parameters, cross validation was used. As such, the results of the MRCE are most appropriately compared to the results of the two-stage algorithm implemented with the Lasso penalty. Since cross validation was used for tuning parameter selection with MRCE, the total sample size used in each simulation was identical to that used with the two-stage algorithm with BIC, while the two-stage algorithm with validation used twice the sample size.

The simulation results in Table 6.1 show that selecting the tuning parameter by minimizing prediction error using independent validation data yields larger models than that selected using BIC. Similarly, model error was usually better when using validation data for model selection. However, BIC selected a higher proportion of true models. Lasso tended to select larger models, and SCAD and SELO selected the true model more often. In terms of variable selection, MRCE performed better than the two-stage approach using the Lasso penalty when the tuning parameter was selected using validation, but worse when the tuning parameter was selected using BIC. The prediction error with MRCE was worse than when using Lasso with validation.

In addition to estimating $\mathbf{\Omega}$ sparsely using GLASSO, we also examined the effect of other estimators of the precision matrix on the estimator of β . We compared the estimator of $\mathbf{\Omega}$ using ridge, and the shrinkage estimator of Schafer and Strimmer (2005). The GLASSO yielded the most stable results compared to other regularized methods of $\mathbf{\Omega}$. In addition, we also compared the GLASSO to a parametric estimator of the precision matrix, when the true

covariance structure was assumed, and when independent correlation structure was assumed. Results for these scenarios are provided in the supplementary material. When using a parametric estimator of the precision matrix, model selection results were comparable to those using GLASSO in both EX (dense precision matrix) and AR (sparse precision matrix) settings. When assuming working independence, the estimated models were slightly larger, with the true model selected fewer times on average.

Table 6.2 compares the empirical SEs with the estimated SEs using the sandwich formula for the non-zero estimated regression parameters. The results show that the standard error estimates had mixed performance. When $n = 100$, usually they were slightly larger than the empirical SEs and, in general, closer to the empirical ones when the β coefficients were larger. When $n = 50$, the pattern was less uniform, with especially larger differences between the estimated and empirical SEs when $\beta = 1.5$, suggesting when β is relatively small, the SE estimates tend to underestimate the true variability.

6.2. Large- p simulations

We adapt the small- p simulations to large- p scenarios. In all simulations we had $n = 50$ observations. The number of covariates and the number of outcomes in $\{5, 20\} \times \{5, 20\}$. With $(p_0, m = 5)$ this is a “small- p ” (25 coefficients) scenario; two scenarios were with 100 coefficients, and one had 400 coefficients to estimate. We simulated AR and EX outcome correlations, each with $\rho_Y = \text{cor}(Y_k, Y_{k-1}) = 0.5$ and $\text{var}(Y_k) = 3$. The covariates $\mathbf{X}$ had an autoregressive covariance structure, with $\text{cor}(X_k, X_l) = 0.3$, $k, l = 1, \dots, p$. We conducted 200 simulations for each scenario, and applied the penalty functions Lasso, adaptive Lasso (with the initial estimator of β estimated as the weight using ridge regression with the General Cross Validation criterion (GCV)), SCAD (with $a = 3.7$), and SELO (with $\tau = 0.01$).

Tuning parameters were selected using data validation, with minimization of prediction error. Thus, 50 additional observations were sampled from the same distribution and used for tuning parameter selection. Here BIC was not used for tuning parameter selection. The true model consisted of three non-zero parameter values, $(3, 1.5, 2)$. Table 6.3 provides the proportion of the true model chosen, the average numbers of false positives and false negatives, and the average squared model error defined as $\text{tr}[(\hat{\beta} - \beta)^T \sum_{xx} (\hat{\beta} - \beta)]$.

The size of the precision matrix was larger when $m = 20$ than when $m = 5$. Therefore, the simulation results were better when $m = 5$, $p_0 = 20$ than when $m = 20$, $p_0 = 5$, with a higher proportion of true models chosen and lower numbers of false positives.

The outcome correlations lead to roughly similar results. Comparing the performance of different penalties, SELO had the best results, SCAD performed the second best with more false positive detections, and the Lasso and adaptive Lasso had considerably higher numbers of false positives. Adaptive Lasso had the worst results, with the highest false positive rate and model error.

Results from additional simulations, in which parametric and working independent correlation structures were used, are provided in the supplementary material. The results show similar patterns to those in Table 6.3.

7. Analysis of the type 2 diabetes data

The type 2 diabetes data set (Mootha et al. (2003)) consists of gene expression profiles of 43 males of similar age with three levels of Glucose Tolerance: Normal (NGT, 17 subjects), Impaired (IGT, 9 subjects) and type 2 Diabetes Mellitus (DM2, 17 subjects). The authors provided a list of gene sets, and performed gene set analysis by testing for the global effect of glucose tolerance levels on each gene set. As many genes in a gene set are likely not affected by glucose levels, we are interested in selecting a subset of genes in a gene set that are associated with glucose tolerance levels. We focused on the porphyrin and chlorophyll metabolism pathway, which has gene expressions of 35 probes of genes. After removing unexpressed probes, 18 of them were used for analysis. We treated the expresses probes as outcomes whereas the covariates were two dummy variables for glucose tolerance levels, with NGT as the baseline. We considered the model

$$y_{ij} = \beta_{0j} + \beta_{1j} \mathbf{1}_{IGT_i} + \beta_{2j} \mathbf{1}_{DM_i} + e_{ij},$$

where y_{ij} is the j^{th} outcome (probe expression) of subject i , $\mathbf{1}_{IGT_i} = 1$ if a subject has impaired glucose tolerance and 0 otherwise, and $\mathbf{1}_{DM_i}$ takes value 1 if a subject has DM2 and 0 otherwise. Hence $m = 18$, $p = 18 \times 3 = 54$ and $n = 43$. We applied the proposed two-stage procedure with the penalties Lasso, adaptive Lasso, SCAD, and SELO for estimation of β . The 18×18 precision matrix was estimated using the Lasso penalty via the GLASSO algorithm. The tuning parameters were selected using BIC for both regression coefficients and precision matrix estimation.

Table 7.4 presents the results of the four penalized regression models. Estimates of the β_{1j} and β_{2j} , the effects of glucose tolerance groups IGT and DM2 on probe expression relative to the baseline NGT group, are presented. The P -values of non-zero estimated coefficients are given in parentheses. The baseline effect (intercept estimates) are omitted. There are more non-zero estimates for the coefficients of DM than those for the coefficients of IGT. This makes sense biologically, as IGT is more similar to NGT than DM.

With all penalties, DM was found to be a significant predictor of the expressions of EPRS(1), HMBS, and ADH6, IGT is a significant predictor for HMBS and CP. IGT was also a significant predictor of BLVRB in the models selected by SELO, Lasso, SCAD, but not by adaptive Lasso. DM was found to be significantly associated with HMOX2 using SELO. The gene HMBS was found to be associated with with IGT and DM.

Lasso produces a large number of non-zero parameter estimates compared to the oracle penalties. The Lasso parameter estimates are shrunk more than those estimated by the oracle penalties and the P -values are relatively high. The SELO penalty produced the smallest number of non-zero parameter estimates.

The data set and code used in this analysis are provided in the journal website.

8. Discussion

We proposed two BIC criteria for selecting the tuning parameter. The first resembles the one used in univariate variable selection, and is used in estimating β with a fixed, or consistently estimated precision matrix. It is consistent for model selection in both cases when an oracle penalty function is used. Our proof allows both p and n go to infinity. The second BIC criterion is used when β and Ω are jointly estimated, as in the two-stage procedure. Simulation results show that the second BIC performs well, and outperforms the tuning parameter selection procedure using an external data set. Future research is needed for theoretical results for this BIC, and for $p > n$.

Our primary focus in this paper is on variable selection for regression coefficients. The optimal regression coefficient estimator that is consistent for model selection and has the oracle properties only requires a consistent estimator of the precision matrix, e.g. using GLASSO. Future research is needed to develop a more efficient estimator of the precision matrix. In the non-regression setting assuming y has mean 0, $\log(m)/n \rightarrow 0$, and sub-gaussian errors, Lam and Fan (2009) provided the rate of convergence for the Ω estimator for Lasso and oracle penalized estimators. Their results could be extended to regression settings in the presence of covariates X for $p_0/n \rightarrow 0$. When $p_0 \gg n$, say $\log(p_0)/n \rightarrow 0$, additional assumptions are needed on the design matrix and the sparsity of β . This is a topic for future research.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work is supported by the NCI grants R37 CA076404 and P01 CA134294.

References

- Arias-Castro E, Candès EJ, Plan Y. Global testing under sparse alternatives: ANOVA, multiple comparisons and the higher criticism. *The Annals of Statistics*. 2011; 39:2533–2556.
- Bhatia, R. *Matrix Analysis*. Springer-Verlag; New York: 1997.
- Brown PJ, Freen R, Vannucci M. The choice of variables in multivariate regression: A non-conjugate Bayesian decision theory approach. *Biometrika*. 1999; 86:635–648.
- Dicker L, Huang B, Lin X. Variable selection and estimation with the seamless- L_0 penalty. *Statistica Sinica*. 2013; 23:929–962.
- Fan J, Feng Y, Wu Y. Network exploration via the adaptive lasso and scad penalties. *The Annals of Applied Statistics*. 2009; 3:521–541. [PubMed: 21643444]
- Fan J, Li R. Variable selection via nonconvex penalized likelihood and its oracle properties. *Journal of the American Statistical Association*. 2001; 96:1348–1360.
- Fan J, Peng H. Nonconcave penalized likelihood with a diverging number of parameters. *Annals of Statistics*. 2004; 32:928–961.
- Friedman J, Hastie T, Tibshirani R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*. 2008; 9:432–441. [PubMed: 18079126]
- Gao X, Pu Q, Wu Y, Xu H. Tuning parameter selection for penalized likelihood estimation of inverse covariance matrix. *Arxiv preprint*. 2009 arXiv:0909.0934.

- Lam C, Fan J. Sparsistency and rates of convergence in large covariance matrix estimation. *Annals of Statistics*. 2009; 37:4254–4278. [PubMed: 21132082]
- Mootha VK, et al. PGC-1 α -responsive Genes involved in oxidative phosphorylation are coordinately downregulated in human diabetes. *Nature Genetics*. 2003; 34:267–273. [PubMed: 12808457]
- Lian H. Shrinkage tuning parameter selection in precision matrices estimation. *Journal of Statistical Planning and Inference*. 2011; 141:2839–2848.
- Liang KY, Zeger SL. Longitudinal data analysis using generalized linear models. *Biometrika*. 1986; 73:13–22.
- Meinshausen N, Bühlmann P. High-dimensional graphs and variable selection with the lasso. *Annals of Statistics*. 2006; 34:1436–1462.
- Peng J, Wang P, Zhou N, Zhu J. Partial correlation estimation by joint sparse regression models. *Journal of the American Statistical Association*. 2009; 104:735–746. [PubMed: 19881892]
- Peng J, et al. Regularized Multivariate Regression for Identifying Master Predictors with Application to Integrative Genomics Study of Breast Cancer. *The Annals of Applied Statistics*. 2010; 4:53–77. [PubMed: 24489618]
- Rothman AJ, Bickel PJ, Levina E, Zhu J. Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics*. 2008; 2:595–515.
- Rothman AJ, Levina E, Zhu J. Sparse multivariate regression with covariance estimation. *Journal of Computational and Graphical Statistics*. 2010
- Schäfer J, Strimmer K. A shrinkage approach to large-scale covariance matrix estimation and implications for functional Genomics. *Statistical Applications in Genetics and Molecular Biology*. 2005; 4 Article 32.
- Stewart, GW., Sun, J. Matrix perturbation theory. Academic Press; New York: 1990.
- Tibshirani R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Association*. 1996; 58:267–288.
- Turlach BA, Venables WN, Wright SJ. Simultaneous variable selection. *Technometrics*. 2005; 47:349–363.
- Wang H, Li B, Leng C. Shrinkage tuning parameters selection with a diverging number of parameters. *Journal of the Royal Statistical Society, Series B*. 2009; 71:617–683.
- White H. Maximum likelihood estimation of misspecified models. *Econometrica*. 1982; 50:1–25.
- Yuan M, Lin Y. Model selection and estimation in the Gaussian graphical model. *Biometrika*. 2007; 95:19–35.
- Zhang CH. Nearly unbiased variable selection under minimax concave penalty. *Journal of the American Statistical Association*. 2010; 101:1418–1429.
- Zou H. The adaptive lasso and its oracle properties. *Annals of Statistics*. 2006; 38:894–942.

Table 6.1

Simulation results over 200 simulations using the two-stage joint estimation procedure with the penalties LASSO, adaptive LASSO, SCAD, and SELO. The precision matrix was estimated using GLASSO with the LASSO penalty. Sample sizes were $n = 50, 100$. The number of outcomes was five and the number of covariates was five. The correlation structures considered were: AR(1) and exchangeable (EX). The tuning parameters was selected either using a validation data set with the same sample size, or by BIC. The average selected model size, percentage of times the true model being selected, average number of false positives, average number of false negatives, average model error and prediction error are reported. The results are compared to the MRCE algorithm.

Simulation	Method	size	True model	false pos	false neg	ME	pred err
AR cov, $n = 50$	LASSO - BIC	5.31	0.14	2.39	0.08	1.70	9.53
	Adaptive LASSO - BIC	3.88	0.32	1.26	0.38	1.59	9.60
	SCAD - BIC	4.22	0.19	1.71	0.48	2.06	9.71
	SELO - BIC	3.82	0.38	1.14	0.32	1.68	9.66
	LASSO - validation	9.13	0.04	6.17	0.04	1.41	9.51
	Adaptive Lasso - validation	6.46	0.08	3.69	0.24	1.39	9.56
	SCAD - validation	5.18	0.10	2.54	0.36	1.52	9.62
	SELO - validation	3.75	0.46	1.10	0.35	1.39	9.60
	MRCE	8.35	0.01	5.41	0.06	1.71	11.59
AR cov, $n = 100$	Lasso - BIC	4.88	0.19	1.88	0.00	0.88	9.42
	Adaptive LASSO - BIC	3.66	0.48	0.80	0.14	0.63	9.45
	SCAD - BIC	3.54	0.60	0.73	0.18	0.59	9.46
	SELO - BIC	3.65	0.57	0.76	0.10	0.58	9.46
	LASSO - validation	9.80	0.02	6.80	0.00	0.73	9.42
	Adaptive LASSO - validation	6.45	0.14	3.50	0.06	0.60	9.44
	SCAD - validation	5.22	0.30	2.31	0.09	0.51	9.45
	SELO - validation	4.00	0.65	1.09	0.10	0.50	9.44
	MRCE	8.74	0.02	5.74	0.00	0.82	11.41
EX cov, $n = 50$	LASSO - BIC	5.25	0.17	2.29	0.05	1.54	9.36
	Adaptive LASSO - BIC	4.13	0.30	1.43	0.30	1.37	9.28
	SCAD - BIC	3.81	0.30	1.21	0.40	1.51	9.36
	SELO - BIC	3.79	0.44	1.04	0.25	1.36	9.27

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Simulation	Method	size	True model	false pos	false neg	ME	pred err
EX cov, $n = 100$	LASSO - validation	9.53	0.02	6.54	0.02	1.33	9.29
	Adaptive LASSO - validation	6.54	0.10	3.73	0.19	1.30	9.25
	SCAD - validation	5.34	0.16	2.63	0.28	1.26	9.25
	SELO - validation	3.75	0.50	1.04	0.30	1.15	9.23
	MRCE	8.26	0.01	5.29	0.03	1.39	10.98
	LASSO - BIC	4.88	0.20	1.88	0.00	0.82	9.11
	Adaptive LASSO - BIC	3.92	0.50	1.01	0.10	0.54	9.13
	SCAD - BIC	3.42	0.66	0.54	0.12	0.44	9.11
	SELO - BIC	3.54	0.63	0.58	0.04	0.42	9.10
	LASSO - validation	10.29	0.02	7.29	0.00	0.68	9.11
	Adaptive LASSO - validation	6.80	0.18	3.83	0.03	0.54	9.12
	SCAD - validation	5.16	0.36	2.23	0.07	0.45	9.12
	SELO - validation	4.12	0.71	1.18	0.06	0.40	9.11
	MRCE	8.91	0.01	5.92	0.00	0.78	10.90

Table 6.2

Median estimated SEs, denoted by SE_EST for the four simulation configurations described in Table 6.1 over 200 simulations. The SE estimators were calculated using the sandwich formula. Empirical SEs calculated using the median absolute deviation of the parameter estimates (SE_EMP) are provided in parenthesis.

Simulation	Method	$\beta_1 = 3$ SE_EST(SE_EMP)	$\beta_2 = 1.5,$ SE_EST(SE_EMP)	$\beta_3 = 2,$ SE_EST(SE_EMP)
AR cov, $n = 50$	LASSO - BIC	0.39 (0.43)	0.33 (0.37)	0.35 (0.34)
	Adaptive LASSO - BIC	0.41 (0.55)	0.33 (0.71)	0.35 (0.41)
	SCAD - BIC	0.48 (0.62)	0.44(1.05)	0.38 (0.37)
	SELO - BIC	0.5 (0.49)	0.49 (0.59)	0.38 (0.36)
	LASSO - validation	0.44 (0.38)	0.43 (0.36)	0.37 (0.35)
	Adaptive LASSO - validation	0.44 (0.49)	0.43 (0.6)	0.37 (0.37)
	SCAD - validation	0.46 (0.59)	0.46 (0.93)	0.38 (0.38)
	SELO - validation	0.46 (0.52)	0.45 (0.69)	0.35 (0.3)
AR cov, $n = 100$	LASSO - BIC	0.32 (0.25)	0.29 (0.25)	0.27 (0.24)
	Adaptive LASSO - BIC	0.33 (0.34)	0.3 (0.33)	0.27 (0.26)
	SCAD - BIC	0.37 (0.32)	0.38 (0.36)	0.27 (0.19)
	SELO - BIC	0.38 (0.27)	0.38 (0.31)	0.28 (0.2)
	LASSO - validation	0.34 (0.26)	0.34 (0.26)	0.29 (0.24)
	Adaptive LASSO - validation	0.34 (0.3)	0.34 (0.34)	0.28 (0.25)
	SCAD - validation	0.35 (0.31)	0.36 (0.34)	0.26 (0.19)
	SELO - validation	0.35 (0.29)	0.35 (0.32)	0.26 (0.19)
EX cov, $n = 50$	LASSO - BIC	0.35 (0.37)	0.3 (0.36)	0.32 (0.33)
	Adaptive LASSO - BIC	0.36 (0.48)	0.28 (0.61)	0.32 (0.36)
	SCAD - BIC	0.42 (0.52)	0.42 (0.8)	0.33 (0.36)
	SELO - BIC	0.43 (0.43)	0.43 (0.48)	0.34 (0.32)
	LASSO - validation	0.4 (0.36)	0.39 (0.33)	0.34 (0.31)
	Adaptive LASSO - validation	0.4 (0.46)	0.39 (0.53)	0.33 (0.36)
	SCAD - validation	0.42 (0.53)	0.42 (0.75)	0.33 (0.33)
	SELO - validation	0.41 (0.47)	0.4 (0.6)	0.31 (0.29)
EX cov, $n = 100$	LASSO - BIC	0.27 (0.23)	0.26 (0.21)	0.25 (0.21)
	Adaptive LASSO - BIC	0.29 (0.29)	0.27 (0.32)	0.24 (0.22)
	SCAD - BIC	0.33 (0.26)	0.33 (0.27)	0.24 (0.19)
	SELO - BIC	0.33 (0.23)	0.33 (0.24)	0.25 (0.2)
	LASSO - validation	0.3 (0.21)	0.31 (0.23)	0.27 (0.21)
	Adaptive LASSO - validation	0.31 (0.27)	0.31 (0.3)	0.25 (0.21)
	SCAD - validation	0.31 (0.27)	0.32 (0.28)	0.23 (0.18)
	SELO - validation	0.31 (0.24)	0.32 (0.24)	0.23 (0.18)

Table 6.3

Large- p simulation results averaged over 200 simulations using the two-stage joint estimation procedure with LASSO, adaptive LASSO, SCAD, and SELO, with the precision matrix estimated using GLASSO with the Lasso penalty. The sample size was $n = 50$. The number of covariates and the number of outcomes were, independently, 5 and 20. The true model size were three. EX and AR correlations were considered. The tuning parameters were selected using an independent validation data set with the same sample size by minimizing the prediction error. The percentage of times the true model being chosen (mean T), the number of variables that were selected as false positives (mean FP), the average number of variables that were selected as false negatives (mean FN), and the average model error (mean ME) are reported.

penalty	mean T	mean FP	mean FN	mean ME
AR cov				
$p_0 = 5, m = 5$				
LASSO	0.01	8.65	0.00	0.57
Adaptive LASSO	0.00	16.05	0.05	1.39
SCAD	0.31	3.41	0.00	0.29
SELO	0.67	0.76	0.00	0.21
$p_0 = 20, m = 5$				
LASSO	0.00	13.83	0.00	0.80
Adaptive LASSO	0.00	26.68	0.16	2.21
SCAD	0.14	6.19	0.00	0.30
SELO	0.72	0.51	0.00	0.24
$p_0 = 5, m = 20$				
LASSO	0.00	18.36	0.00	1.03
Adaptive LASSO	0.00	45.55	0.11	3.03
SCAD	0.07	7.21	0.00	0.39
SELO	0.56	1.38	0.00	0.37
$p_0 = 20, m = 20$				
LASSO	0.00	23.81	0.00	1.37
Adaptive LASSO	0.00	41.94	0.48	4.12
SCAD	0.03	14.63	0.00	0.58
SELO	0.65	0.87	0.01	0.36
EX cov				
$p_0 = 5, m = 5$				
LASSO	0.01	9.01	0.00	0.52
Adaptive LASSO	0.00	15.63	0.06	1.24
SCAD	0.29	3.42	0.00	0.25

penalty	mean T	mean FP	mean FN	mean ME
SELO	0.66	0.86	0.00	0.19
$p_0 = 20, m = 5$				
LASSO	0.00	14.12	0.00	0.75
Adaptive LASSO	0.00	26.56	0.19	2.19
SCAD	0.16	6.24	0.00	0.28
SELO	0.70	0.61	0.00	0.23
$p_0 = 5, m = 20$				
LASSO	0.00	17.41	0.00	0.75
Adaptive LASSO	0.00	44.29	0.06	1.97
SCAD	0.12	6.43	0.00	0.28
SELO	0.62	1.52	0.00	0.27
$p_0 = 20, m = 20$				
LASSO	0.00	22.50	0.00	0.98
Adaptive LASSO	0.00	52.02	0.21	2.83
SCAD	0.09	12.53	0.00	0.36
SELO	0.61	1.31	0.00	0.29

Table 7.4

Analysis results of the diabetes data set ($n = 43$) with the porphyrin and chlorophyll metabolism pathway, which has 18 genes ($m = 18$). Variable selection was performed using the penalized multivariate regression with the two-stage joint estimation procedure with the penalties LASSO, adaptive LASSO, SCAD and SELO. The baseline expression values (intercepts), representing the mean expression levels of the normal glucose tolerance group are not presented. IGT and DM correspond to the mean expression differences between subjects with Impaired Glucose Tolerance (IGT) and subjects with type II diabetes (DM) compared the normal group. The number of regression coefficients was $p = 18 \times 3 = 54$. The tuning parameter was estimated using the BIC. The P -values are in parentheses.

Gene (probe)	Lasso			Adaptive Lasso			Scad			SELO		
	IGT	DM		IGT	DM		NGT	DM		IGT	DM	
EPRS (1)	0.01 (0.88)	0.62 (0.01)		0	0.69 (0.005)		0	0.73 (0.004)		0	0.79 (0.002)	
EPRS (2)	0.27 (0.19)	0.15 (0.41)		0.29 (0.18)	0.12 (0.49)		0.11 (0.49)	0 (0.83)		0	0	
EPRS (3)	0	0.06 (0.53)		0	<0.01 (0.88)		0	0.02 (0.75)		0	0	
BLVRB	0.35 (0.007)	-0.07 (0.61)		0.36 (0.478)	-0.03 (0.98)		0.44 (0.003)	0		0.5 (0.001)	0	
GUSB	-0.02 (0.86)	-0.13 (0.5)		0	-0.11 (0.82)		0	-0.07 (0.67)		0	0	
UROS	-0.08 (0.56)	0		-0.06 (0.84)	0		-0.02 (0.76)	0		0	0	
HMBS	-0.16 (0.15)	-0.23 (0.05)		-0.18 (0.17)	-0.26 (0.05)		-0.13 (0.27)	-0.26 (0.04)		-0.25 (0.03)	-0.32 (0.006)	
FECH	-0.03 (0.63)	0		0	0		0	0		0	0	
HMOX1	0.03 (0.76)	0		0	0		0	0		0	0	
HCCS (1)	0	-0.03 (0.69)		0	0		0	0		0	0	
HCCS (2)	0	-0.14 (0.39)		0	-0.08 (0.78)		0	-0.07 (0.62)		0	0	
BLVRA (1)	0	-0.15 (0.19)		0	-0.16 (0.27)		0	-0.13 (0.29)		0	-0.18 (0.14)	
CP	0.52 (0.03)	0		0.57 (0.01)	0		0.57 (0.01)	0		0.69 (0.003)	0	
UROD (1)	0	0		0	0		0	0		0	0	
UROD (2)	0	0		0	0		0	0		0	0	
BLVRA (2)	-0.03 (0.68)	0		0	0		0 (0.82)	0		0	0	
ADH6	-0.24 (0.41)	0.53 (0.034)		-0.22 (0.63)	0.58 (0.04)		0	0.78 (0.001)		0	0.76 (0.001)	
HMOX2	0	0.25 (0.22)		0	0.29 (0.34)		0	0.24 (0.26)		0	0.43 (0.04)	