

Statistica Sinica Preprint No: SS-2018-0170

Title	Penalized linear regression with high-dimensional pairwise screening
Manuscript ID	SS-2018-0170
URL	http://www.stat.sinica.edu.tw/statistica/
DOI	10.5705/ss.202018.0170
Complete List of Authors	Siliang Gong Kai Zhang and Yufeng Liu
Corresponding Author	Siliang Gong
E-mail	siliang@penncmedicine.upenn.edu
Notice: Accepted version subject to English editing.	

Penalized linear regression with high-dimensional pairwise screening

Siliang Gong, Kai Zhang and Yufeng Liu

University of Pennsylvania

and The University of North Carolina at Chapel Hill

Abstract: In variable selection, most existing screening methods focus on marginal effects and ignore dependence between covariates. To improve the performance of selection, we incorporate pairwise effects in covariates for screening and penalization. We achieve this by studying the asymptotic distribution of the maximal absolute pairwise sample correlation among independent covariates. The novelty of the theory is in that the convergence is with respect to the dimensionality p , and is uniform with respect to the sample size n . Moreover, we obtain an upper bound for the maximal pairwise R squared when regressing the response onto two different covariates. Based on these extreme value results, we propose a screening procedure to detect covariates pairs that are potentially correlated and associated with the response. We further combine the pairwise screening with Sure Independence Screening [9] and develop a new regularized variable selection procedure. Numerical studies show that our method is very competitive in terms of both prediction accuracy and variable selection accuracy.

Key words and phrases: Pairwise Screening, Penalized Regression, Sure Independence Screening, Variable Selection

1. Introduction

In the era of big data, high dimensional problems are of interest in many scientific fields, where the number of variables may be comparable to or even much larger than the sample size. For example, in genetic studies, one often has tens of thousands of genes in the microarray datasets with only a few hundreds of patients; in neuroscience, fMRI images may contain millions of voxels.

In recent years, much research effort has been devoted to dealing with high dimensional data analysis. Among those methods developed, penalized least squares plays an important role. In particular, one of the most well-known method is the LASSO proposed by [21], which is the solution to the following penalized problem

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\mathbf{y} - X\boldsymbol{\beta}\|_2^2 + \lambda P(\boldsymbol{\beta}), \quad (1.1)$$

where $\lambda P(\boldsymbol{\beta}) = \lambda \sum_{j=1}^p |\beta_j|$ is the l_1 penalty. [21] showed that the LASSO leads to a sparse estimator that shrinks the OLS solution and sets some of the estimated coefficients to exact zero. Despite with good theoretical properties and practical performance, the LASSO has two major drawbacks: firstly, due to the shrinkage nature, LASSO may over-shrink the estimates and cause significant bias; secondly, if there is a group of variables that are highly correlated, LASSO tends to select only one of them. To address these issues, [27] introduced the elastic net method, using $\lambda_1 \|\boldsymbol{\beta}\|_1 + \lambda_2 \|\boldsymbol{\beta}\|_2^2$ as the regularization term in (1.1) and thus encouraging a grouping effect. Besides the elastic net, various penalized variable selection methods have been proposed as extensions to LASSO, including the Dantzig selector [7], the

smoothly clipped absolute deviation (SCAD) penalty [8], among many others. See [13] and [10] for a comprehensive overview.

For high dimensional variable selection, it is crucial to account for the dependency structure of covariates. Such structure information not only improves the accuracy of selection, but also have practical meanings. For instance, in gene expression data, genes usually function as biological pathways instead of working independently. Classical penalized variable selection methods, however, usually do not explicitly take into account the relationships between covariates. To address this problem, [23] proposed the group LASSO method, which takes advantage of the grouping information among the covariates. Extensions of group lasso include, but are not limited to [3]. Other methods use the structure information as predictor graph (see [14, 15, 26, 22] among others for reference).

A common assumption for the methods mentioned above is that the underlying predictor graph is given, which may not hold in practice. When the prior information is not available, the idea of clustering can be incorporated to improve regression performance. Specifically, [16] proposed to perform hierarchical clustering on the covariates first and take the cluster average as new predictors for regression. There are also methods using supervised clustering to encourage highly correlated pairs of covariates to be included or excluded simultaneously [2, 17]. Similarly, another type of methods aims to make correlated covariates have similar regression coefficients [18]. Nevertheless, a large sample correlation between two variables does not necessarily indicate that they are dependent in the population sense. When the dimensionality continues to increase, the maximal pairwise correlation among p independent covariates can be close to 1 [10]. Therefore, it is important to identify covariates that are

truly correlated and incorporate such information into variable selection procedures.

In this paper, we study the limiting behavior of the maximal absolute pairwise sample correlation among covariates when they are independent Gaussian random variables. Different from existing work, we investigate the limiting distribution as the dimensionality p diverges. Therefore, the proposed asymptotic results potentially can be applied to datasets with arbitrarily large dimensionality. We further discuss the extreme behavior of the maximal absolute Spearman's rho statistic for covariates with general distributions. On the other hand, we obtain the upper bound of maximal pairwise R squared when regressing the response onto pairs of covariates. With the extreme value results, we formulate a screening procedure to identify covariates pairs that are potentially dependent and associated with the response. We further combine the pairwise screening with the Sure Independence Screening (SIS) [9] and propose a novel penalized variable selection method. More specifically, we assign different penalties to each individual covariate according to the screening results. Numerical experiments show that the performance of our proposed method is competitive compared with existing approaches in terms of both variable selection and prediction accuracy.

The remainder of this paper is organized as follows: We first investigate the limiting distribution of the maximal pairwise sample correlation among covariates in Section 2.1. We also show that our asymptotic results cover that of [5] as a special case. Then we propose an upper bound for the maximal pairwise R squared in Section 2.2. In Section 3.1 we formulate our proposed variable selection approach as a penalized maximum likelihood problem, and discuss potential extensions of our method in Section 3.2. Theoretical properties are discussed in Section 4. We show with simulated experiments as well as two real datasets in

Section 5 that the proposed method has improved performance when important variables are highly correlated. Finally, we conclude this paper and discuss possible future work in Section 6. Proofs of the theoretical results are provided in the Appendix.

2. Pair Screening for covariates

Suppose we have the following linear model

$$\mathbf{y} = X\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad (2.2)$$

where $\mathbf{y} = (y_1, y_2, \dots, y_n)^T$ is the response vector, $X = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_p)$ is an $n \times p$ design matrix with \mathbf{x}_j being n independent and identical observations from the covariate X_j . We assume that the covariate vector $\mathbf{x} = (X_1, X_2, \dots, X_p)^T$ has a multivariate distribution with unknown covariance matrix Σ , and $\boldsymbol{\varepsilon} = (\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n)^T$ is a vector of i.i.d. random variables with mean 0 and standard deviation σ , and is independent of the covariate vector \mathbf{x} .

For the linear model (2.2), variable selection methods aim to identify the non-zero components of $\boldsymbol{\beta}$, in other words, the important variables among all candidate predictors. Particularly, if two covariates have a large pairwise correlation, we may want to include or exclude these two variables simultaneously when conducting variable selection. However, the sample correlation can be spurious, especially when the number of covariates p is relatively large. Therefore, it is important to identify covariates that are truly correlated. In other words, we need to find a threshold for the pairwise sample correlation among covariates to screen covariates pairs. In the following subsection, we will discuss in details the asymptotic results

that generate the screening rule.

2.1 Extreme laws of pairwise sample correlation among covariates

We propose to choose a bound based on the extreme laws of the pairwise sample correlation when the p covariates are independent. Our investigations are under two settings: (a) the covariates are normally distributed; (b) the covariates are non-Gaussian random variables.

2.1.1 Gaussian covariates

It has been recently studied that the maximal absolute Pearson sample correlation between p i.i.d. Gaussian covariates and an independent response has a Gumble-type limiting distribution as p goes to infinity [24]. Motivated by [24]'s work, we find that the maximal absolute pairwise sample correlation among p independent covariates also has a limiting distribution, as stated in the following theorem:

Theorem 1. *Suppose X_1, X_2, \dots, X_p are p independent Gaussian variables and we observe n independent samples from each of X_j 's. Let $W_{pn} = \max_{1 \leq i < j \leq p} |\rho_{i,j}|$, where $\rho_{i,j} = \widehat{\text{Corr}}(X_i, X_j)$ is the Pearson sample correlation between X_i and X_j . Then as $p \rightarrow \infty$,*

$$\lim_{p \rightarrow \infty} \left| P\left(\frac{W_{pn}^2 - a_{p,n}}{b_{p,n}} \leq x\right) - I\left(x \leq \frac{n-2}{2}\right) \exp\left\{-\frac{1}{2}\left(1 - \frac{2}{n-2}x\right)^{\frac{n-2}{2}}\right\} - I\left(x > \frac{n-2}{2}\right) \right| = 0, \quad (2.3)$$

which is uniformly for any $n \geq 3$. Here $a_{p,n} = 1 - p^{-4/(n-2)}c_{p,n}$, $b_{p,n} = \frac{2}{n-2}p^{-4/(n-2)}c_{p,n}$, and $c_{p,n} = \left(\frac{n-2}{2}B\left(\frac{1}{2}, \frac{n-2}{2}\right)\sqrt{1 - p^{-4/(n-2)}}\right)^{2/(n-2)}$ are the normalizing constants.

In random matrix theory, W_{pn} is also known as the coherence when the design matrix X

is random. Specifically, the coherence is defined as the largest magnitude of the off-diagonal entries of the sample correlation matrix associated with a random matrix. The limiting behavior of the coherence has been well studied when the sample size n goes to infinity. For example, [6] studied the asymptotic distribution under certain regularity conditions with application to the testing of covariance matrix. [5] further obtained the limiting laws of the coherence for different divergence rate of p with respect to n and summarized the results as phase transition phenomena. We can show that our result unifies the convergence in terms of the sample size, and covers [5]'s as special cases, described in the following corollary.

Corollary 1. *Let W_{pn} be defined as in Theorem 1, where we still assume X_j 's are independent normal random variables. Let $T_{pn} = \log(1 - W_{pn}^2)$.*

(a) (**Sub-Exponential Case**) *Suppose $p = p_n \rightarrow \infty$ as $n \rightarrow \infty$ and $(\log p)/n \rightarrow 0$, then as $n \rightarrow \infty$,*

$$\mathrm{P}(nT_{pn} + 4 \log p - \log \log p \leq x) \rightarrow 1 - e^{-\frac{1}{\sqrt{8\pi}}e^{x/2}}.$$

(b) (**Exponential Case**) *Suppose $p = p_n$ satisfies $(\log p)/n \rightarrow \beta \in (0, \infty)$ as $n \rightarrow \infty$. Then as $n \rightarrow \infty$,*

$$\mathrm{P}(nT_{pn} + 4 \log p - \log \log p \leq x) \rightarrow 1 - \exp\{K(\beta)e^{(x+8\beta)/2}\},$$

$$\text{where } K(\beta) = \left(\frac{\beta}{2\pi(1-4e^{-4\beta})}\right)^{1/2}.$$

(c) (**Super-Exponential Case**) *Suppose $p = p_n$ satisfies $(\log p)/n \rightarrow \infty$ as $n \rightarrow \infty$.*

Then as $n \rightarrow \infty$,

$$P \left(nT_{pn} + \frac{4n}{n-2} \log p - \log n \leq x \right) \rightarrow 1 - e^{-\frac{1}{\sqrt{2\pi}} e^{x/2}}.$$

Compared with previous work, our asymptotic distribution is novel in two aspects. First, the convergence in Theorem 1 is with respect to p instead of n , making it applicable to high dimensional data, or even ultrahigh dimensional problems. Moreover, the convergence result we have discovered is uniform for any $n \geq 3$, thus finite sample performance is guaranteed.

2.1.2 Non-Gaussian covariates

When the covariates are non-Gaussian random variables, it is more desirable to choose a distribution-free statistic for the screening rule. Therefore, instead of using the Pearson's sample correlation, we study the extreme behavior of the Spearman's rho statistic [20]. Recall that $\mathbf{x}_j = (X_{1j}, X_{2j}, \dots, X_{nj})^T$ are n i.i.d. observations from the covariate X_j . Let Q_{ni}^j and Q_{ni}^k be the ranks of X_{ij} and X_{ik} in $\{X_{1j}, \dots, X_{nj}\}$ and $\{X_{1k}, \dots, X_{nk}\}$ respectively. Then the Spearman's rho is defined as

$$\rho_{ij} = \frac{\sum_{i=1}^n (Q_{ni}^j - \bar{Q}_n^j)(Q_{ni}^k - \bar{Q}_n^k)}{\sqrt{\sum_{i=1}^n (Q_{ni}^j - \bar{Q}_n^j)^2 \sum_{i=1}^n (Q_{ni}^k - \bar{Q}_n^k)^2}}, \quad (2.4)$$

where $\bar{Q}_n^j = \bar{Q}_n^k = \frac{n+1}{2}$.

Similar to the normal setting, we are particularly interested in the limiting distribution of $S_{pn}^2 = \max_{1 \leq i < j \leq p} \rho_{ij}^2$ when the covariates are all independent, which has been studied in

[12]. The following proposition states that as n increases, S_{pn}^2 converges to a Gumble type distribution.

Proposition 1. *Suppose that X_1, \dots, X_p are independent and identically distributed random variables, and we have n independent samples for each of the covariates. Let $S_{pn}^2 = \max_{1 \leq i < j \leq p} \rho_{ij}^2$ be the squares of the maximal pairwise Spearman's rho statistics, then for $\log p = o(n^{1/3})$, we have*

$$\lim_{n \rightarrow \infty} |\mathbb{P}((n-1)S_{pn}^2 - 4 \log p + \log \log p \leq x) - \exp\{- (8\pi)^{-1/2} \exp(-x/2)\}| = 0. \quad (2.5)$$

Theorem 1 and Proposition 1 characterize the magnitude of the maximal pairwise correlation and Spearman's rho statistic respectively when the covariates are independent. Suppose a pair of covariates, say X_1 and X_2 , have a absolute sample correlation greater than the 95% quantile of the distribution given in Theorem 1 or Proposition 1, then they tend to be marginally dependent. Since we are only interested in pairs of truly important covariates, we further investigate the extreme behavior of the maximal pairwise R squared under the null model, i.e., β_j 's are all equal to zero.

2.2 R squared screening for pairs of covariates

With the asymptotic distributions introduced in the previous subsections, we can identify covariates pairs that are potentially dependent. However, such screening does not take into account the association between the covariates and the response. It is possible that an important variable has a large sample correlation with unimportant ones; or two highly

3. PENALIZED VARIABLE SELECTION USING PAIRWISE SCREENING

correlated covariates are both unrelated to the response. To address such an issue, we introduce another screening procedure based on the R squared from regressing the response Y onto the pairs of covariates.

Consider the linear regression where we regress Y onto a pair of covariates X_i and X_j with $i \neq j$, we can obtain the corresponding R squared R_{ij}^2 . Under the model setting (2.2), when all the coefficients are zeros, the maximal pairwise R squared $\max_{1 \leq i < j \leq p} R_{ij}^2$ cannot be too large. In fact, there exists an asymptotic bound for $\max_{1 \leq i < j \leq p} R_{ij}^2$, as described in the following theorem.

Theorem 2. *Let $R_{pn}^2 = \max_{1 \leq i < j \leq p} R_{ij}^2$, where R_{ij}^2 is the pairwise R squares from regressing Y onto X_i and X_j where $i \neq j$. Suppose that X_1, \dots, X_p and Y are from the model setting 2.2 and we further assume that Y is a normally distributed. Then when β_j 's are all zeros, we have for any fixed $n \geq 4$, $\delta > 0$, as $p \rightarrow \infty$, $P(R_{pn}^2 \geq 1 - p^{-(4+\delta)/(n-3)}) = O(p^{-\delta/2}) \rightarrow 0$.*

With the bound given by Theorem 2, we can design a screening rule to find pairs of covariates that are potentially associated with the response. In Section 3, we introduce how to make use of the theoretical results to benefit variable selection.

3. Penalized variable selection using pairwise screening

In this section, we propose a pairwise screening procedure that takes advantages of the asymptotic results in Section 2. We further establish a new penalization algorithm for variable selection.

3. PENALIZED VARIABLE SELECTION USING PAIRWISE SCREENING

3.1 Screening-based penalization

Given the limiting distribution of the maximal pairwise sample correlation described in Section 2, we propose the following screening rule to identify covariates pairs that are potentially correlated and related to the response:

$$\mathcal{G} = \{(i, j) : i < j, |\widehat{\text{Corr}}(X_i, X_j)| \geq a \text{ and } R_{ij}^2 \geq r_0\}, \quad (3.6)$$

where a is the $100(1 - \alpha)\%$ quantile of the distribution given in Theorem 1 (for Gaussian covariates) or Proposition 1 (for non-Gaussian covariates), and $r_0 = 1 - p^{-(4+\delta)/(n-3)}$. Note that the values of α and δ can affect the size of \mathcal{G} . The larger α and δ are, there are fewer pairs included in \mathcal{G} . In practice, we suggest to take $\alpha = 0.05$ and $\delta = 0.1$.

The group definition in (3.6) is a screening procedure with respect to covariates pairs. The idea of screening is prevalent for high dimensional data analysis. In particular, for penalized variable selection methods, increasing dimensionality makes it more difficult to capture the inherent sparsity structure. Therefore, dimension reduction is necessary when there are tens of thousands of candidate variables. To this end, [9] introduced the Sure Independence Screening (SIS) method, which ranks the covariates based on the magnitude of their sample correlation with the response. More specifically, let $\mathbf{w} = (w_1, w_2, \dots, w_p)^T$ be a vector such that $w_j = |\widehat{\text{Corr}}(X_j, Y)|$ and γ is a constant between $(0, 1)$, then a sub-model is defined as

$$\mathcal{M}_\gamma = \{j : w_j \text{ is amongst the largest } [\gamma n] \text{ of all}\}, \quad (3.7)$$

where $[\gamma n]$ denotes the integer part of γn . [9] further demonstrated that SIS is screening

3. PENALIZED VARIABLE SELECTION USING PAIRWISE SCREENING

consistent under some conditions. This guarantees that all those X_j 's with $\beta_j \neq 0$ are included in the subset of covariates.

To take advantage of the distribution information while implementing dimension reduction, we propose a new penalized variable selection approach that applies different penalties to each covariate based on the screening results. Let \mathcal{M} be the index set of covariates that have the largest $[n \setminus \log n]$ absolute sample correlation with the response among X_1, X_2, \dots, X_p . We also define the set of paired covariates as

$$\mathcal{C} = \{X_i : \exists j \text{ such that } (i, j) \in \mathcal{G}\}. \quad (3.8)$$

Our proposed method is established by solving the following optimization problem:

$$\min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|\mathbf{y} - X\boldsymbol{\beta}\|_2^2 + \lambda_1 \sum_{j: j \in \mathcal{C}^c \cap \mathcal{M}} |\beta_j| + \lambda_2 \sum_{j: j \in \mathcal{C} \cap \mathcal{M}} \beta_j^2 \quad (3.9)$$

subject to $\beta_j = 0$ for $j \notin \mathcal{M}$. In other words, we ignore the covariates that fail the marginal screening.

From the above penalty, it can be seen that we apply different penalties to covariates based upon the results from two types of screening. The intuitions behind the proposed penalty are

- For a covariate that is included in both \mathcal{C} and \mathcal{M} , we only apply the l_2 penalty because it tends to be an important variable that we need to include in the final model.
- For a covariate that is included in \mathcal{M} but not in \mathcal{C} , we only apply the l_1 penalty since

3. PENALIZED VARIABLE SELECTION USING PAIRWISE SCREENING

there is no significant multicollinearity between it and other covariates.

- For a covariate that is not included in \mathcal{M} , since it does not pass the marginal screening, we no longer consider it in the regression. This is because SIS enjoys screening consistency under certain assumptions, which implies that \mathcal{M} covers all important variables.

Our proposed method is connected with existing penalization approaches when the covariates have certain covariance structure. In particular, when the covariates are all independent, our method reduces to SIS-LASSO, which performs marginal screening first and then implements LASSO on the remaining covariates; when the predictors are all highly correlated such that \mathcal{G} includes all covariates pairs, our method is equivalent to SIS-Ridge.

So far we have established a new penalized variable selection. Now we discuss how to solve the optimization problem in (3.9). One can see that the penalty part of (3.9) is convex, so we can efficiently solve it by coordinate descent algorithm [11]. Specifically, the updating rule has the following form:

$$\hat{\beta}_j \leftarrow \begin{cases} S\left(\frac{1}{N} \sum_{i=1}^N x_{ij}(y_i - \tilde{y}_i^{(j)}), \lambda_1\right) & \text{for } j \in \mathcal{C}^c \cap \mathcal{M}, \\ \frac{\frac{1}{N} \sum_{i=1}^N x_{ij}(y_i - \tilde{y}_i^{(j)})}{1 + \lambda_2} & \text{for } j \in \mathcal{C} \cap \mathcal{M}, \end{cases} \quad (3.10)$$

where $\tilde{y}_i^{(j)} = \hat{\beta}_0 + \sum_{k \neq j} x_{ik} \hat{\beta}_k$ is the fitted value excluding the effect of x_{ij} , and $S(z) = \text{sign}(z)(|z| - \lambda)_+$ is the soft-thresholding function. In practice, we can first implement SIS to obtain \mathcal{M} when the dimension is high, then run the algorithm on the covariates X_j 's with $j \in \mathcal{M}$.

3. PENALIZED VARIABLE SELECTION USING PAIRWISE SCREENING

Remark 1. The computational cost of the pairwise screening procedure is $O(p^2)$, which can be very inefficient as p increases. In our proposed procedure, to reduce the computational complexity, we implement the marginal screening first to obtain \mathcal{M} . Since the cardinality of \mathcal{M} is $O(n/\log(n))$, the computational cost of applying pairwise screening to \mathcal{M} will reduce to $O((n/\log(n))^2)$.

3.2 Further extensions

As discussed in the previous subsection, we introduce a new penalized method that combines marginal screening with pairwise screening under the linear model setting. Note that the pairwise covariates screening does not involve the response. Therefore, our method can be further extended to generalized linear models (GLM), e.g., logistic regression for binary response, or cox model for survival data. Suppose the response Y is from the following one-parameter exponential family $f(y|\mathbf{x}, \theta) = h(y) \exp\{y\theta - b(\theta)\}$. Moreover, we assume $\theta = \mathbf{x}^T \boldsymbol{\beta}$ for generalized linear models.

Similar to 3.6, we define the pairwise screening as

$$\mathcal{G}_1 = \{(i, j) : i < j, |\widehat{\text{Corr}}(X_i, X_j)| \geq a\}. \quad (3.11)$$

The difference is that we do not consider the R squared screening for GLMs. This is because for GLMs, it is not reasonable to use the regression R squared to evaluate the associations between the covariates and the response. We further define the set of paired covariates as

follows

$$\mathcal{C}_1 = \{X_i : \exists j \text{ such that } (i, j) \in \mathcal{G}_1\}. \quad (3.12)$$

Let $P_{\lambda_1, \lambda_2}(\boldsymbol{\beta}) = \lambda_1 \sum_{j: j \in \mathcal{C}_1^c \cap \mathcal{M}} |\beta_j| + \lambda_2 \sum_{j: j \in \mathcal{C}_1 \cap \mathcal{M}} \beta_j^2$ be our proposed screening-based penalty. Then for logistic regression, we need to solve the following penalized maximum likelihood problem

$$\min_{\boldsymbol{\beta}} \sum_{i=1}^n (y_i(\mathbf{x}_i^T \boldsymbol{\beta}) - \log(1 + e^{\mathbf{x}_i^T \boldsymbol{\beta}})) + P_{\lambda_1, \lambda_2}(\boldsymbol{\beta}). \quad (3.13)$$

In the above optimization problem, the log likelihood part can be approximated by a quadratic function, which is a weighted least squares term [11]. Therefore, it can still be solved by coordinate descent algorithm. Similarly, we can use the algorithm proposed by [19] to solve the regularized Cox proportional hazard model using the screening based penalty $P_{\lambda_1, \lambda_2}(\boldsymbol{\beta})$.

4. Theoretical properties

In this section, we study the theoretical properties of the proposed pairwise correlation screening (PCS) method. More specifically, we investigate the conditions under which the PCS achieves the variable selection consistency.

Note that we implemented the marginal screening using SIS to the covariates set. Fan and Lv [9] demonstrated that under certain regularity conditions, SIS has the screening consistency, that is, the resulting subset of covariates includes all important variables. Due to space constraints, we only present the main result. The regularity conditions (A1) – (A4) are provided in the appendix.

Proposition 2 ([9]). *Under (A1) – (A4), if $2\kappa + \tau < 1$, then there is some $\theta < 1 - 2\kappa - \tau$ such that, when $\gamma \sim cn^{-\theta}$ with $c > 0$, we have, for some $C > 0$,*

$$P(\mathcal{M}^* \subset \mathcal{M}_\gamma) = 1 - O[\exp\{-C^{1-2\kappa}/\log(n)\}], \quad (4.14)$$

where \mathcal{M}_γ is the subset of covariates obtained from the sure independence screening.

The above proposition guarantees that all important variables survive the marginal screening with high probability. In order to achieve the selection consistency, we also need to ensure that only important variables can pass the pairwise screening. In the following theorem, we present the technical conditions that are required such that the event $\mathcal{C} \cap \mathcal{M} \subset \mathcal{M}^*$ occurs with high probability.

Theorem 3. *Suppose the following conditions holds*

(B1) $n/p^2 \rightarrow 0$.

(B2) *There exists $\eta > 0$ such that either one of the following two conditions holds:*

(a) $\lim_{n \rightarrow \infty} \log p/n \rightarrow \eta_0$, $\max_{i \in \mathcal{M}^*, j \in \mathcal{M} \setminus \mathcal{M}^*} |\text{Corr}(X_i, X_j)| < \min\{\eta, 1 - e^{-4\eta_0}\}$

(b) $\lim_{n \rightarrow \infty} \log p/n \rightarrow 0$, $\max_{i \in \mathcal{M}^*, j \in \mathcal{M} \setminus \mathcal{M}^*} |\text{Corr}(X_i, X_j)| < \eta$.

Here $\text{Corr}(X_i, X_j)$ denotes the population correlation between covariates X_i and X_j . Then under conditions (B1) and (B2)(a) or conditions (B1) and (B2)(b), we have that as $n \rightarrow \infty$,

$$P(\mathcal{C} \cap \mathcal{M} \subset \mathcal{M}^*) \rightarrow 1. \quad (4.15)$$

Given Proposition 2 and Theorem 3, to demonstrate the selection consistency of PCS, we only need to show that the l_1 penalty in (3.9) can identify the important variables in $\mathcal{C}^c \cap \mathcal{M}$ exactly. This relates to the selection consistency for the LASSO, which has been studied extensively. In particular, Zhao and Yu [25] have shown that the Irrepresentable Condition (to be clarified later) is almost necessary and sufficient for LASSO to select all important variables.

We first introduce some necessary notations. Let $C = \frac{1}{n}X^T X$. Without loss of generality, assume that $\beta = (\beta_1, \beta_2, \dots, \beta_p)^T$ where $\beta_j \neq 0$ for $j = 1, \dots, s$ and $\beta_j = 0$ otherwise. By Theorem 3, we further assume that $\mathcal{C} \cap \mathcal{M} = \{1, \dots, s_1\}$ where $1 \leq s_1 \leq s$. Then the design matrix X can be expressed as $X = (X_{(1)}^1, X_{(1)}^2, X_{(2)})$, where $X_{(1)}^1$ corresponds to the first s_1 columns, $X_{(1)}^2$ corresponds to the $(s_1 + 1)$ th to the s th columns and $X_{(2)}$ corresponds to the last $p - s$ columns of X respectively. Similarly, we can write $\beta_1^{(1)} = (\beta_1, \dots, \beta_{s_1})^T$, $\beta_2^{(1)} = (\beta_{s_1+1}, \dots, \beta_s)^T$, and $\beta^{(2)} = (\beta_{s+1}, \dots, \beta_p)^T$.

Set $C_{11}^{(11)} = \frac{1}{n}X_{(1)}^{1T} X_{(1)}^1$, $C_{11}^{(12)} = \frac{1}{n}X_{(1)}^{1T} X_{(1)}^2$, $C_{11}^{(21)} = \frac{1}{n}X_{(1)}^{2T} X_{(1)}^1$, $C_{11}^{(22)} = \frac{1}{n}X_{(1)}^{2T} X_{(1)}^2$, $C_{21}^{(1)} = \frac{1}{n}X_{(2)}^T X_{(1)}^1$, $C_{21}^{(2)} = \frac{1}{n}X_{(2)}^T X_{(1)}^2$, $C_{22} = \frac{1}{n}X_{(2)}^T X_{(2)}$, $C_{12}^{(1)} = \frac{1}{n}X_{(1)}^{1T} X_{(2)}$, $C_{12}^{(2)} = \frac{1}{n}X_{(1)}^{2T} X_{(2)}$.

Then C can be expressed in a block-wise form as follows:

$$\begin{pmatrix} C_{11}^{(11)} & C_{11}^{(12)} & C_{12}^{(1)} \\ C_{11}^{(21)} & C_{11}^{(22)} & C_{12}^{(2)} \\ C_{21}^{(1)} & C_{21}^{(2)} & C_{22} \end{pmatrix}$$

We impose the following assumption analogous to the Irrepresentable Condition introduced by Zhao and Yu [25]. Specifically, we assume that there exists a constant $\delta > 0$, such

that

$$\|C_{21}^{(2)}(C_{11}^{(22)})^{-1}\text{sign}(\boldsymbol{\beta}_2^{(1)})\|_{\max} \leq 1 - \delta, \quad (4.16)$$

where $\|\cdot\|_{\max}$ is the max norm.

In fact, we can show that the condition mentioned above is implied by the Irrepresentable Condition on the full covariates set \mathcal{M} under mild assumptions. We illustrate this result in the following theorem:

Theorem 4. *Assume that there exists $\lambda_0 > 0$ so that $\lambda_{\min}(C_{11}^{(11)}) \geq \lambda_0$, $\lambda_{\min}(C_{11}^{(22)}) \geq \lambda_0$, and conditions (B1) and (B2)(b) holds. Suppose the Irrepresentable Condition holds, i.e., $\exists \xi > 0$ s.t.*

$$\|C_{21}C_{11}^{-1}\text{sign}(\boldsymbol{\beta}_1)\|_{\max} \leq 1 - \xi, \quad (4.17)$$

where $C_{11} = \begin{pmatrix} C_{11}^{(11)} & C_{11}^{(12)} \\ C_{11}^{(21)} & C_{11}^{(22)} \end{pmatrix}$, $C_{21} = \begin{pmatrix} C_{21}^{(1)} & C_{21}^{(2)} \end{pmatrix}$, $\boldsymbol{\beta}_1 = (\beta_1, \dots, \beta_s)^T$ and ξ is a positive constant. $\lambda_{\min}(\cdot)$ Then with probability tending to 1, the condition (4.16) holds.

The assumptions $\lambda_{\min}(C_{11}^{(11)}) \geq \lambda_0$, $\lambda_{\min}(C_{11}^{(22)}) \geq \lambda_0$ in Theorem 4 require that $C_{11}^{(11)}$ and $C_{11}^{(22)}$ have eigenvalues bounded below. Given the Irrepresentable Condition in (4.17), we need additional constraints on the random noise ε_i 's and the coefficients of important variables β_1, \dots, β_s .

(C1) ε_i 's are i.i.d. random variables with finite $2k$'s moment $E(\varepsilon_i)^{2k} < \infty$ for an integer

$k > 0$.

(C2) There exists $0 < \alpha \leq 1$ and $d_0 > 0$ such that $n^{\frac{1-\alpha}{2}} \min_{j=1, \dots, s} |\beta_j| \geq d_0$.

So far we have discussed all the theoretical assumptions required to ensure the selection consistency of the proposed PCS method. We conclude the consistency result in the following theorem:

Theorem 5. *Suppose conditions (A1)–(A4), (C1)–(C2) and inequality (4.17) hold, and the assumptions of Theorem 4 are satisfied, then for any λ_1 such that $\frac{\lambda_1}{\sqrt{n}} = o(n^\alpha/2)$ and $\frac{1}{p}(\frac{\lambda_1}{\sqrt{n}}) \rightarrow \infty$, we have,*

$$\mathbb{P} \left(\{j : \hat{\beta}_j \neq 0\} = \mathcal{M}^* \right) \rightarrow 1 \text{ as } n \rightarrow \infty, \quad (4.18)$$

where $\hat{\boldsymbol{\beta}} = (\hat{\beta}_1, \dots, \hat{\beta}_p)^T$ is the solution to (3.9).

The proof follows immediately from Proposition 2 and Theorems 3 and 4 as well as the selection consistency of the LASSO. It shows that under certain conditions, our proposed method is consistent in variable selection. In Section 5, we will show with numerical examples that our proposed method can perform well in practice.

5. Numerical Studies

In Section 3, we have established a new regularized variable selection approach for high-dimensional linear models. In this section, we demonstrate the performance of our proposed method using both simulations and real data examples.

5.1 Simulation study

In this section, we use several simulated examples to show that our method with pairwise correlation screening (PCS) or pairwise rank-based correlation screening (PRCS) outperforms some existing variable selection procedures. More specifically, PCS denotes our proposed method using the limiting distribution in Theorem 1, and PRCS uses the asymptotic result in Proposition 1.

For comparison, we consider LASSO, elastic net (Enet), SIS-LASSO, SIS-elastic net (SIS-Enet) and SIS-PACS. The SIS-PACS refers to applying the PACS method proposed by [17] after implementing the SIS procedure. In SIS-type methods, we first implement SIS and find those covariates with the largest $[n \setminus \log n]$ absolute sample correlations with the response, then perform LASSO, Enet or PACS on these variables. We evaluate the variable selection accuracy using False Negatives (FN) and False Positives (FP). FN is defined as $FN = \sum_{j=1}^p I(\hat{\beta}_j = 0) \times I(\beta_j \neq 0)$, where $I(\cdot)$ denotes the indicator function, and FP is defined as $FP = \sum_{j=1}^p I(\hat{\beta}_j \neq 0) \times I(\beta_j = 0)$. We use the following quantities to evaluate the prediction accuracy:

- $\|\hat{\beta} - \beta_0\|_2$: the l_2 distance between the estimated coefficient vector and the true coefficients β_0 ;
- Out of sample mean squared errors (MSE) on the independent test data;

We generate the simulated data from Model (2.2) and conduct 100 replications. Each simulated dataset includes a training set of size 100, an independent validation set of size 100 and an independent test set of size 400. Here we fix the sample size to be 100 throughout the

Table 1: Results for simulated Example 1. For each method, we report the average MSE, l_2 distance, FN and FP over 100 replications (with standard errors given in parentheses).

Method	MSE	$\ \hat{\beta} - \beta_0\ _2$	FN	FP
$p = 1000, \sigma = 2$				
Elnet	5.94 (0.07)	1.40 (0.03)	0.00 (0.00)	1.64 (0.24)
SIS-Elnet	5.47 (0.06)	1.30 (0.03)	0.00 (0.00)	1.15 (0.12)
LASSO	5.95 (0.07)	1.50 (0.03)	0.00 (0.00)	1.28 (0.18)
SIS-LASSO	5.47 (0.06)	1.42 (0.03)	0.00 (0.00)	0.85 (0.10)
SIS-Ridge	86.00 (0.76)	4.50 (0.01)	0.00 (0.00)	12.00 (0.00)
SIS-PACS	4.69 (0.07)	0.48 (0.02)	0.00 (0.00)	0.01 (0.01)
PCS	4.74 (0.05)	0.76 (0.02)	0.00 (0.00)	0.03 (0.02)
PRCS	4.91 (0.05)	0.93 (0.02)	0.00 (0.00)	2.55 (0.15)
$p = 5000, \sigma = 2$				
Elnet	6.42 (0.09)	1.57 (0.03)	0.00 (0.00)	2.45 (0.26)
SIS-Elnet	5.64 (0.06)	1.41 (0.03)	0.00 (0.00)	1.28 (0.12)
LASSO	6.41 (0.08)	1.64 (0.04)	0.00 (0.00)	2.06 (0.21)
SIS-LASSO	5.65 (0.06)	1.52 (0.03)	0.00 (0.00)	1.03 (0.10)
SIS-Ridge	88.74 (0.75)	4.59 (0.01)	0.00 (0.00)	12.00 (0.00)
SIS-PACS	4.97 (0.08)	0.72 (0.02)	0.00 (0.00)	1.78 (0.43)
PCS	4.77 (0.05)	0.81 (0.03)	0.00 (0.00)	0.02 (0.02)
PRCS	4.85 (0.06)	0.89 (0.03)	0.00 (0.00)	1.21 (0.11)

simulation study. In the next subsection, we also consider varying sample size for sensitivity study. We only fit models on the training data, and we use the validation data to select tuning parameters. Given the fitted model, we can calculate the FN, FP and the estimation error $\|\hat{\beta} - \beta_0\|_2$, and we make predictions and calculate the out of sample MSEs using the test data. We simulate the covariates from the multivariate Gaussian distribution $\mathcal{N}(0, \Sigma)$, with $\Sigma = (\sigma_{ij})_{p \times p}$ being the correlation matrix.

Details of the simulated examples are as follows:

Example 1: We consider $p = 1000$ or 5000 , $\sigma = 2$, and we take $\beta = (2, 2, \dots, 2, 0, \dots, 0)^T$ where the first 10 coefficients being non-zero and equal to 2. We set $\sigma_{ij} = 0.8$ for $1 \leq i \neq j \leq 5$, $6 \leq i \neq j \leq 10$ and 0 for all the other $i \neq j$. We also consider $\sigma = 6$ and present the

Table 2: Results for simulated Example 2. The format of this table is the same as Table 1.

Method	MSE	$\ \hat{\beta} - \beta_0\ _2$	FN	FP
$p = 1000, \sigma = 2$				
Enet	6.75 (0.08)	2.45 (0.02)	1.00 (0.01)	0.98 (0.25)
SIS-Enet	6.47 (0.10)	2.30 (0.03)	0.76 (0.05)	3.16 (0.41)
LASSO	6.75 (0.08)	2.45 (0.02)	1.00 (0.01)	0.98 (0.25)
SIS-LASSO	6.47 (0.10)	2.30 (0.03)	0.76 (0.05)	3.16 (0.41)
SIS-Ridge	14.14 (0.10)	3.85 (0.00)	0.27 (0.04)	19.27 (0.04)
SIS-PACS	6.53 (0.14)	2.43 (0.04)	1.06 (0.05)	3.39 (0.73)
PCS	5.24 (0.12)	1.41 (0.08)	0.34 (0.05)	1.63 (0.13)
PRCS	5.72 (0.13)	1.75 (0.08)	0.43 (0.05)	1.34 (0.24)
$p = 5000, \sigma = 2$				
Elnet	7.16 (0.08)	2.55 (0.02)	1.02 (0.01)	0.40 (0.09)
SIS-Elnet	7.02 (0.09)	2.49 (0.03)	0.94 (0.03)	1.31 (0.34)
LASSO	7.16 (0.08)	2.55 (0.02)	1.02 (0.01)	0.36 (0.08)
SIS-LASSO	7.03 (0.09)	2.49 (0.03)	0.94 (0.03)	1.31 (0.34)
SIS-Ridge	14.40 (0.11)	3.87 (0.00)	0.59 (0.05)	19.59 (0.05)
SIS-PACS	7.28 (0.16)	2.83 (0.04)	1.26 (0.07)	2.41 (0.95)
PCS	5.96 (0.14)	1.83 (0.09)	0.63 (0.06)	0.74 (0.08)
PRCS	6.48 (0.13)	2.14 (0.07)	0.68 (0.05)	0.73 (0.24)

Table 3: Results for simulated Example 3. The format of this table is the same as Table 1.

Method	MSE	$\ \hat{\beta} - \beta_0\ _2$	FN	FP
Enet	69.71 (0.88)	5.13 (0.03)	4.99 (0.13)	1.57 (0.37)
SIS-Enet	72.54 (0.88)	5.25 (0.03)	5.65 (0.10)	0.23 (0.12)
LASSO	72.78 (0.87)	5.41 (0.03)	6.06 (0.10)	0.09 (0.04)
SIS-LASSO	70.12 (0.86)	5.35 (0.04)	5.69 (0.12)	0.94 (0.19)
SIS-Ridge	109.66 (0.87)	5.74 (0.01)	4.46 (0.06)	16.46 (0.06)
SIS-PACS	71.27 (0.89)	5.58 (0.02)	5.06 (0.02)	3.45 (0.07)
PCS	58.87 (0.50)	4.80 (0.04)	4.95 (0.03)	0.06 (0.06)
PRCS	59.76 (0.56)	4.83 (0.04)	4.97 (0.02)	0.00 (0.00)

results in the supplementary. In other words, there are two groups in the covariates, where each group has 5 important variables.

Example 2: We consider $p = 1000$ or 5000 , $\sigma = 2$, $\beta_0 = (3, -1.5, 2, 0, \dots, 0, \dots, 0)^T$,

Table 4: Results for simulated Example 4. The format of this table is the same as Table 1.

Method	Classification Error	$\ \hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\ _2$	FN	FP
Enet	0.129 (0.003)	5.79 (0.01)	2.16 (0.17)	12.77 (1.54)
SIS-Enet	0.126 (0.003)	5.69 (0.03)	1.37 (0.15)	7.48 (0.39)
LASSO	0.136 (0.003)	5.83 (0.01)	4.19 (0.13)	4.25 (0.49)
SIS-LASSO	0.130 (0.003)	5.75 (0.02)	3.94 (0.12)	3.50 (0.32)
SIS-Ridge	0.311 (0.003)	6.28 (0.01)	0.11 (0.05)	12.11 (0.05)
PCS	0.098 (0.004)	5.39 (0.05)	1.73 (0.14)	2.92 (0.31)
PRCS	0.099 (0.004)	5.34 (0.06)	1.71 (0.13)	3.26 (0.32)

Table 5: Results for simulated Example 5. The format of this table is the same as Table 1.

Method	MSE	$\ \hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\ _2$	FN	FP
Enet	102.47 (1.84)	3.90 (0.08)	1.51 (0.12)	4.88 (0.86)
SIS-Enet	96.60 (2.74)	3.49 (0.09)	1.02 (0.12)	4.20 (0.37)
LASSO	103.11 (1.89)	4.42 (0.08)	2.30 (0.13)	3.74 (0.71)
SIS-LASSO	96.97 (2.78)	4.27 (0.08)	2.05 (0.14)	1.87 (0.20)
SIS-Ridge	226.52 (3.78)	4.95 (0.03)	0.26 (0.08)	12.26 (0.08)
SIS-PACS	89.82 (2.70)	3.54 (0.15)	0.26 (0.08)	7.32 (0.45)
PCS	79.79 (3.16)	2.42 (0.14)	0.42 (0.10)	1.29 (0.33)
PRCS	74.60 (1.24)	2.15 (0.12)	0.31 (0.08)	0.06 (0.03)

where the first 3 coefficients are non-zero ones. We also consider $\sigma = 6$ and present the results in the supplementary. We generated Gaussian covariates with $\sigma_{ij} = 0.5^{|i-j|}$ for $1 \leq i \neq j \leq 1000$.

Example 3: The coefficients have the same set up as in Example 1. But we set $\sigma_{ij} = 0.8$ for $1 \leq i \neq j \leq 5$ and 0 for all the other $i \neq j$. Therefore only part of the important variables are highly correlated. We consider $p = 5000$ and $\sigma = 6$ in this Example.

Example 4: In this example, we examine the performance of all methods under the logistic regression setting. We simulate the binary response Y from the binomial distribution $\text{Binom}(1, \frac{\exp\{X^T\boldsymbol{\beta} + \sigma\}}{1 + \exp\{X^T\boldsymbol{\beta} + \sigma\}})$, where X , and $\boldsymbol{\beta}$ follow the same set ups as in Example 1. We

consider $p = 5000$ and $\sigma = 6$ in this Example. Instead of comparing MSE, we calculate the classification errors on the test data. We did not compare with SIS-PACS in this example since the R program does not support GLM.

Example 5: In this example, we generate the covariates from a multivariate t distribution, where X_j 's are t distributed with degrees of freedom 5. The covariance structure of the covariates and the coefficients are set the same as in Example 1. We consider $p = 5000$ and $\sigma = 6$ in this Example.

The results for simulated Example 1 is shown in Table 1. We see that when there are groups in the covariates, the performance improvement of our approach is significant compared with other penalized methods. While elastic net-based procedures perform better than LASSO-type approaches in terms of FN, as illustrated by [27], they still miss approximately one important covariate on average. In contrast, the model selection results of our method are much closer to the correct model for this example. In addition, although SIS-PACS has competitive performance when σ is small, it tends to include more unimportant variables into the model when the noise level increases, and therefore may not work well.

Table 2 displays the performance comparisons for Example 2. Compared with Example 1, this setting is a more difficult one for our method, since correlation exists among all pairs of covariates. Nevertheless, PCS and PRCS perform better than, or as well as all the others in terms of estimation error and prediction accuracy. Moreover, besides SIS-Ridge, our proposed methods are able to identify more important variables than others in this example when the noise level is low.

Table 3 shows the results for Example 3, where correlation exists only within part of the

important variables. This example is more difficult compared with Example 1 due to the correlation structure of the covariates. One can see that the false negatives are significantly larger for all procedures. Nevertheless our method still outperforms all the others in terms of prediction and variable selection accuracy.

Example 4 considers the logistic regression setting, and the results are provided in Table 4. One can see that as the correlations among the covariates vary, the performance of our method is always competitive compared with the others.

Table 5 displays the results for all methods under the non-Gaussian covariates setting. Similar to Example 1, our proposed PCS and PRCS achieve much better performance compared with the competitors. Moreover, due to the non-Gaussian set ups, the nonparametric method PRCS outperforms PCS.

As a conclusion, our method can make use of the correlation structure among predictors. Compared with other penalized variable selection procedures, our method performs well, especially when the covariates are highly correlated.

5.2 Sensitivity Study

In this subsection, we investigate how the performance of our method depends on the sample size, dimensionality, and noise level. In particular, we consider $n = 100$ or 500 , $p = 500, 1000, 2000$ or 5000 and $\sigma = 2$ or 6 in the Simulated Example 1 as introduced in Section 5.1. We illustrate the MSE, $\|\hat{\beta} - \beta_0\|_2$, FN and FP against different values of p for each configuration of sample size and noise level in Figure 1.

One can see from the plots that the performance of PCS does not change much as the

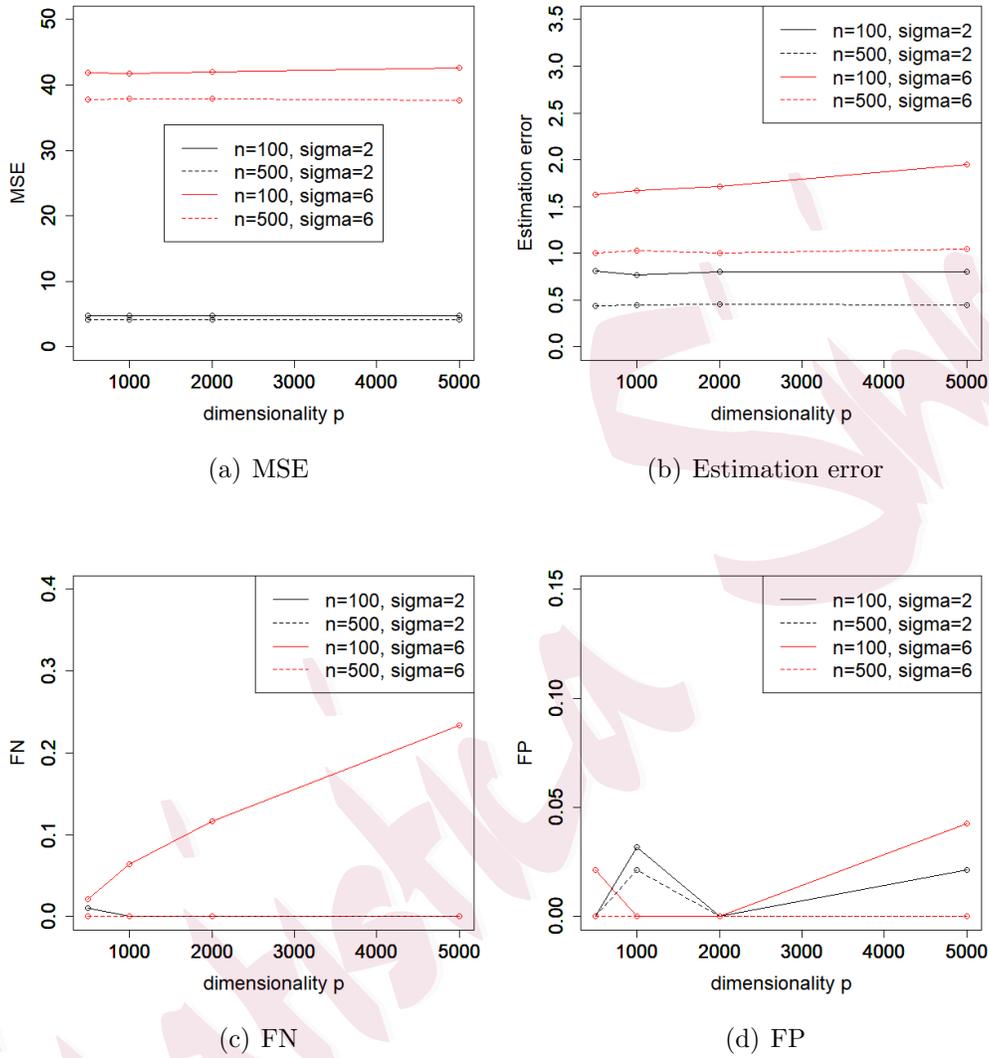


Figure 1: Performance of PCS against different dimensionality p .

dimensionality p increases from 500 to 5000, especially in terms of MSE and the estimation error of β_0 . Moreover, the performance is better when the sample size and signal to noise ratio (SNR) become larger, which is expected. In general, our proposed PCS method is robust to sample size, dimensionality and SNR.

5.3 Soil data

We first demonstrate the performance of our method in real applications using a small dataset. This dataset contains 15 covariates of soil characteristics for 20 plots with the same area in the Appalachian Mountain. The outcome variable is the forest diversity for each plot. More descriptions of the data can be found in [2]. To better demonstrate the correlation structure of covariates, we obtain the absolute pairwise correlation matrix and show the heatmap in Figure 2. One can see that some predictors are highly correlated. In particular, the magnitude of the pairwise correlations among Sum of Cations (SumCation), calcium, magnesium, Base Saturation (BaseSat), and cation exchange capacity (CEC) are as large as 0.9. The reason is SumCation, BaseSat, CEC are characteristics for cations; while calcium and magnesium are examples of cations [2].

We conduct a total of 100 replications. In each replication, 15 samples are randomly chosen as the training set and the remaining as the test set. As in the simulation experiments, we applied LASSO, Enet, Ridge and our proposed PCS, PRCS to the dataset. For each method, 5-fold cross-validation is used to choose the tuning parameters since the sample size is very small. We report the average prediction errors on the test data and the model

Method	MSE	Model Size
Enet	1.088 (0.047)	3.70 (0.38)
LASSO	1.068 (0.045)	2.08 (0.21)
Ridge	1.113 (0.044)	15.00 (0.00)
PCS	0.996 (0.062)	5.82 (0.37)
PRCS	1.028 (0.063)	5.96 (0.38)

Table 6: Average mean squared errors and model size (with standard errors in parenthesis) for Enet, LASSO, Ridge and our methods on the soil data.

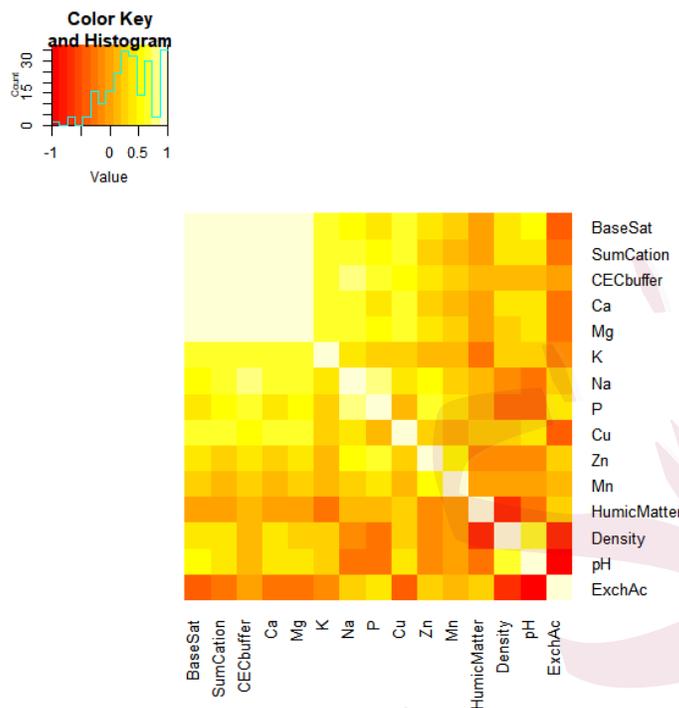


Figure 2: Heatmap for the absolute pairwise correlation matrix of the covariates for the soil data.

size in Table 6. One can see that PCS and PRCS outperform all the others in terms of prediction accuracy. Moreover, PCS and PRCS tend to include more covariates into the model compared with LASSO and Enet.

To further investigate the performance of variable selection, we summarize the frequency that each covariate is selected for LASSO, Enet and our method, which is displayed in Table 7. Note that among those variables that are most frequently selected by LASSO and Enet, for instance, CEC, Mn, HumicMatt, they also tend to be included for our method. Moreover, our method can identify covariates that are strongly correlated. For example, potassium, sodium and copper are variables related to cations, and all have a large sample correlation with CEC, which is a potentially important variable. These variables are frequently selected

Variables	PCS	Enet	LASSO
BaseSat	16	9	0
SumCaton	32	23	0
CECbuffer	86	62	48
Ca	37	32	11
Mg	6	10	0
K	49	27	12
Na	22	10	6
P	32	15	5
Cu	47	17	9
Zn	29	17	4
Mn	69	43	32
HumicMatt	89	70	69
Density	25	15	4
pH	27	11	4
ExchAc	16	9	4

Table 7: Frequency of each variable being selected for PCS, Enet and LASSO out of 100 replications.

by our method, but not by Enet and LASSO.

5.4 Riboflavin data

In this section, we consider a real data set about the riboflavin production in *Bacillus subtilis*. The data contain $n = 71$ samples, where the response variable is the logarithm of the riboflavin production rate, and the covariates are the logarithm of expression levels of $p = 4081$ genes. More descriptions about the dataset can be found in [4]. Before analysis, all covariates are standardized to have zero means and unit standard deviations.

For comparison purpose, we apply LASSO, Enet, SIS-LASSO, SIS-Enet, SIS-ridge and our method to the dataset. We conduct 100 replications, and we randomly split the dataset into a training set of size 50 with the remaining as the test data. For all methods, we

Method	MSE	Model Size
SIS-Enet	0.358 (0.015)	15.66 (0.46)
SIS-Lasso	0.356 (0.016)	9.12 (0.18)
SIS-Ridge	0.632 (0.024)	26.00 (0.00)
PCS	0.327 (0.014)	15.04 (0.39)
PRCS	0.361 (0.018)	12.77 (0.37)

Table 8: Average mean squared errors and model size (with standard errors in parenthesis) for SIS-Enet, SIS-LASSO, SIS-Ridge, PCS and PRCS applied to the riboflavin data.

implement 10-fold cross validation on the training data to select the penalty parameters.

The results are reported in Table 8. One can see that PCS has significant improvement in terms of out of sample mean squared errors compared with other competitors. On the other hand, PRCS does not perform well compared with PCS. A possible reason is that in this dataset all the variables have been taken log transformations and are approximated well by Gaussian distribution. Moreover, due to the assumption of Proposition 1 where $\log p = o(n^{1/3})$, PRCS is more sensitive to the dimensionality and the sample size of dataset. As a result, PRCS may not achieve good performance when the dimensionality is too high.

We also examine the gene selection results. There are 8 genes that are selected at least 50 times out of the 100 replications by our method, i.e., XTRA_at, YCKE_at, YDAR_at, YOAB_at, YWFO_at, YXLC_at, YXLD_at and YXLE_at. Besides YXLC_at, all the other genes also appear among the most frequently selected genes by SIS-Enet and SIS-LASSO with a frequency no less than 50. For YXLC_at, we find that the magnitude of the pairwise sample correlations between this gene and two other genes, YXLD_at and YXLE_at, are greater than 0.95. It indicates that our method is capable of identifying potentially important variables that are highly correlated with the others.

6. Discussion

In summary, we propose a novel variable selection method that regularizes covariates selectively based on the results from two screening procedures: pairwise screening and marginal screening. The screening process of covariates pairs takes advantage of the distribution information of the maximal absolute pairwise sample correlation among covariates, and is applicable to large scale problems. Simulation experiments and real data study demonstrate that the proposed method performs well when important variables are highly correlated compared with existing approaches. For future research, we can consider other extensions of our proposed method, for example, the Cox model for survival data.

A. Technical Proofs

We present some regularity conditions and key proofs in the appendix.

Regularity Conditions for Sure Independence Screening Define $\mathbf{z} = \Sigma^{-1/2}\mathbf{x}$, $Z = X\Sigma^{-1/2}$. Let \mathcal{M}^* be the index set of covariates with non-zero coefficient. The following assumptions are imposed:

(A1) $p > n$ and $\log(p) = O(n^\epsilon)$ for some $\epsilon \in (0, 1 - 2\kappa)$, where κ is given by condition (A3).

(A2) \mathbf{z} has a spherically symmetric distribution, and $\exists c_0, c_1 > 1, C_1 > 0$ such that

$$P\left(\lambda_{\max}(\tilde{p}\tilde{Z}\tilde{Z}^T) > c_1 \text{ or } \lambda_{\min}(\tilde{p}\tilde{Z}\tilde{Z}^T) < 1/c_1\right) \leq \exp(-C_1n)$$

holds for any $n \times \tilde{p}$ submatrix \tilde{Z} of Z with $c_0n < \tilde{p} \leq p$.

(A3) $Var(Y) = O(1)$, and for some $\kappa \geq 0$ and $c_2, c_3 > 0$,

$$\min_{j \in \mathcal{M}^*} |\beta_j| \geq \frac{c_2}{n^\kappa}, \quad \min_{j \in \mathcal{M}^*} Cov(\beta_j^{-1}Y, X_j) \geq c_3$$

(A4) There are some $\tau \geq 0$ and $c_4 > 0$ such that $\lambda_{max}(\Sigma) \leq c_4 n^\tau$.

Proof of Theorem 1. To prove Theorem 1, we need to use the following lemma, which is from [1].

Lemma 1. *Let I be an index set and $\{B_\alpha, \alpha \in I\}$ be a set of subsets of I , that is, $B_\alpha \subset I$ for each $\alpha \in I$. Let also $\{\eta_\alpha, \alpha \in I\}$ be random variables. For a given $t \in R$, set $\lambda = \sum_{\alpha \in I} P(\eta_\alpha > t)$. Then*

$$|P\left(\max_{\alpha \in I} \eta_\alpha < t\right) - e^{-\lambda}| \leq (1 \wedge \lambda^{-1})(b_1 + b_2 + b_3)$$

where $b_1 = \sum_{\alpha \in I} \sum_{\beta \in B_\alpha} P(\eta_\alpha > t)P(\eta_\beta > t)$, $b_2 = \sum_{\alpha \in I} \sum_{\alpha \neq \beta \in B_\alpha} P(\eta_\alpha > t, \eta_\beta > t)$ and $b_3 = \sum_{\alpha \in I} E|P(\eta_\alpha > t | \sigma(\eta_\beta, \beta \notin B_\alpha)) - P(\eta_\alpha > t)|$, and $\sigma(\eta_\beta, \beta \notin B_\alpha)$ is the σ -algebra generated by $\{\eta_\beta, \beta \notin B_\alpha\}$. In particular, if η_α is independent of $\{\eta_\beta, \beta \notin B_\alpha\}$ for each α , then $b_3=0$.

In our proof, we take $I = \{(i, j); 1 \leq i \leq j \leq p\}$. Let $\alpha = (i, j) \in I$, we define $B_\alpha = \{(k, l) \in I; \text{one of } k \text{ and } l = i \text{ or } j, \text{ but } (k, l) \neq \alpha\}$, and $A_\alpha = A_{ij} = \{|\rho_{i,j}|^2 \geq t\}$, where $\rho_{i,j} = |\widehat{Corr}(X_i, X_j)|$. Let $W_{pn} = \max_{1 \leq i < j \leq p} |\rho_{i,j}|$, by the Chen-Stein method (in particular, Lemma 6.2 in Cai and Jiang (2011)),

$$|P(W_{pn}^2 \leq t) - e^{-\lambda_{p,n}}| \leq b_1 + b_2, \quad (\text{A.19})$$

where $\lambda_{p,n} = \sum_{\alpha \in I} P(A_\alpha) = \frac{p(p-1)}{2}P(A_{12})$, and $b_1 = \sum_{\alpha \in I} \sum_{\beta \in B_\alpha} P(A_\alpha)P(A_\beta)$, $b_2 = \sum_{\alpha \in I} \sum_{\alpha \neq \beta \in B_\alpha} P(A_\alpha A_\beta)$.

Moreover, we have $b_1 \leq 2p^3P(A_{12})^2$ and $b_2 \leq 2p^3P(A_{12}A_{13})$.

Since X_1, \dots, X_p are independent, A_{12} and A_{13} are also independent with equal probability. Therefore we have $b_1 \vee b_2 \leq 2p^3P(A_{12})^2$.

On the other hand, $|\rho_{i,j}|^2 \sim B(\frac{1}{2}, \frac{n-2}{2})$. Take $t^* = a_{p,n} + b_{p,n}x$ ($x \leq \frac{n-2}{2}$), where $a_{p,n} = 1 - p^{-4/(n-2)}c_{p,n}$, $b_{p,n} = \frac{2}{n-2}p^{-4/(n-2)}c_{p,n}$, and $c_{p,n} = (\frac{n-2}{2}B(\frac{1}{2}, \frac{n-2}{2})\sqrt{1 - p^{-4/(n-2)}})^{2/(n-2)}$.

Then

$$\begin{aligned} P(A_{12}^*) &= \frac{2(1-t^*)^{(n-2)/2}}{B(\frac{1}{2}, \frac{n-2}{2})(n-2)\sqrt{t^*}}(1 + O(\frac{1}{\log(p)})). \\ &= p^{-2}(1 - \frac{2}{n-2}x)^{\frac{n-2}{2}} \sqrt{\frac{1 - p^{-4/(n-2)}}{a_{p,n}}} (1 + (\frac{b_{p,n}}{a_{p,n}}x)^{-1/2}(1 + O(\log^{-1}(p))))). \\ &= p^{-2}(1 - \frac{2}{n-2}x)^{\frac{n-2}{2}} (1 + O(\frac{\log \log(p)}{\log(p)}))(1 + O(\log^{-1}(p)))^2 \\ &= p^{-2}(1 - \frac{2}{n-2}x)^{\frac{n-2}{2}} (1 + O(\frac{\log \log(p)}{\log(p)})) \end{aligned} \quad (\text{A.20})$$

Therefore, uniformly for any $n \geq 3$, $b_1 \vee b_2 = O(1/p)$, and $\lim_{p \rightarrow \infty} \lambda_{p,n} = \frac{1}{2}(1 - \frac{2}{n-2}x)^{\frac{n-2}{2}}$

Then it follows from (A.19) that uniformly for any $n \geq 3$ and $x \leq \frac{n-2}{2}$,

$$\lim_{p \rightarrow \infty} |P(W_{pn}^2 \leq t^*) - \exp\{-\frac{1}{2}(1 - \frac{2}{n-2}x)^{\frac{n-2}{2}}\}| = 0. \quad (\text{A.21})$$

When $x \geq \frac{n-2}{2}$, $t^* = 1 + (\frac{2}{n-2}x - 1)p^{-4/(n-2)}c_{p,n} \geq 1$. Therefore, uniformly for any $n \geq 3$,

$$\lim_{p \rightarrow \infty} P(W_{pn} \leq t^*) = 1 \quad (\text{A.22})$$

Combining (A.21) and (A.22) we have uniformly for any $n \geq 3$,

$$\lim_{p \rightarrow \infty} |P(W_{pn} \leq t^*) - I(x \leq \frac{n-2}{2}) \exp\left\{-\frac{1}{2}\left(1 - \frac{2}{n-2}x\right)^{\frac{n-2}{2}}\right\} - I(x > \frac{n-2}{2})| = 0. \quad (\text{A.23})$$

Or equivalently,

$$\lim_{p \rightarrow \infty} \left| P\left(\frac{W_{pn}^2 - a_{p,n}}{b_{p,n}} \leq x\right) - I\left(x \leq \frac{n-2}{2}\right) \exp\left\{-\frac{1}{2}\left(1 - \frac{2}{n-2}x\right)^{\frac{n-2}{2}}\right\} - I\left(x > \frac{n-2}{2}\right) \right| = 0. \quad (\text{A.24})$$

□

Proof of Theorem 3. Let event $A = \{R_{ij}^2 \leq 1 - p^{-(4+\delta)/(n-3)} \text{ for all } i, j \in \mathcal{M} \setminus \mathcal{M}^*\}$, event $B = \{\hat{\rho}_{ij} \leq f(n, p, \alpha) \text{ for } i \in \mathcal{M}^*, j \in \mathcal{M} \setminus \mathcal{M}^*\}$ where $\hat{\rho}_{ij} = |\widehat{\text{Corr}}(X_i, X_j)|$, $f(n, p, \alpha)$ is the screening threshold for pairwise correlation screening. Then A implies that no pairs of unimportant variables passed the R squares screening. B implies that important and unimportant variables can not be too highly correlated.

By the definition of \mathcal{C} , we have

$$P(\mathcal{C} \cap \mathcal{M} \subset \mathcal{M}^*) \geq P(A \cap B) \geq P(A) + P(B) - 1. \quad (\text{A.25})$$

For the event A , we have

$$\begin{aligned}
 P(A) &= 1 - P\left(\bigcup_{i \neq j \in \mathcal{M} \setminus \mathcal{M}^*} R_{ij}^2 \geq 1 - p^{-(4+\delta)/(n-3)}\right) \geq 1 - \sum_{i \neq j \in \mathcal{M} \setminus \mathcal{M}^*} P(R_{ij}^2 \geq 1 - p^{-(4+\delta)/(n-3)}) \\
 &= 1 - (n/\log(n))^2 P(\text{Beta}(1, \frac{n-3}{2}) \geq 1 - p^{-(4+\delta)/(n-3)}) \\
 &= 1 - (n/\log(n))^2 p^{-(4+\delta)/2}
 \end{aligned}$$

Under the assumption (B1), $(n/\log(n))^2 p^{-(4+\delta)/2} \rightarrow 0$ as $n \rightarrow \infty$. Therefore we have $P(A) \rightarrow 1$.

Next we show that $P(B) \rightarrow 1$ as $n \rightarrow \infty$. We have

$$\begin{aligned}
 P(B) &= 1 - p\left(\bigcup_{i \in \mathcal{M}^*, j \in \mathcal{M} \setminus \mathcal{M}^*} \hat{\rho}_{ij} \geq f(n, p, \alpha)\right) \geq 1 - \sum_{i \in \mathcal{M}^*, j \in \mathcal{M} \setminus \mathcal{M}^*} P(\hat{\rho}_{ij} \geq f(n, p, \alpha)) \\
 &= 1 - (n/\log(n))^2 P(\hat{\rho}_{ij} \geq \max\{a_{p,n} + b_{p,n}F_n(\alpha), \eta\}) \\
 &= 1 - (n/\log(n))^2 P(\hat{\rho}_{ij} \geq \delta_{p,n}),
 \end{aligned}$$

where $F_n(\alpha)$ is the $100(1 - \alpha)$ quantile of the limiting cumulative distribution function of the maximal pairwise correlation statistic, and we denote $\max\{a_{p,n} + b_{p,n}F_n(\alpha), \eta\}$ by $\delta_{p,n}$.

Note that

$$\begin{aligned}
 a_{p,n} + b_{p,n}F_n(\alpha) &= 1 - p^{-4/(n-2)} c_{p,n} \left(1 - \frac{2}{n-2} F_n(\alpha)\right) \\
 &= 1 - p^{-4/(n-2)} c_{p,n} \{-2 \log(1 - \alpha)\}^{2/(n-2)} \\
 &= 1 - \left(C_\alpha p^{-2} \frac{n-2}{2} B\left(\frac{1}{2}, \frac{n-2}{2}\right) \sqrt{1 - p^{-4/(n-2)}}\right)^{\frac{2}{n-2}} \\
 &= 1 - O\left(\frac{C_\alpha^2 (n-2) (1 - p^{-4/(n-2)})}{p^4}\right)^{\frac{1}{n-2}} \quad \text{for large enough } n \\
 &= 1 - O\left(e^{-\frac{\log p}{n}}\right) \quad \text{for large enough } n
 \end{aligned}$$

Let ρ_{ij} be the population correlation coefficient between X_i and X_j . Write $z(n) =$

$\frac{1}{2} \log \frac{1+\hat{\rho}_{ij}}{1-\hat{\rho}_{ij}}$, $\xi = \frac{1}{2} \log \frac{1+\rho_{ij}}{1-\rho_{ij}}$. It has been shown that as $n \rightarrow \infty$, $n^{1/2}(z(n) - \xi) \rightarrow \mathcal{N}(0, 1)$.

We have

$$\begin{aligned} P(\hat{\rho}_{ij} \geq \delta_{p,n}) &= P\left(n^{1/2}(z(n) - \xi) \geq n^{1/2}\left(\frac{1}{2} \log \frac{1+\delta_{p,n}}{1-\delta_{p,n}} - \xi\right)\right) \\ &= P\left(Z \geq n^{1/2}\left(\frac{1}{2} \log \frac{1+\delta_{p,n}}{1-\delta_{p,n}} - \xi\right) + o_n(1)\right) \leq \frac{e^{-C_{p,n}n}}{\sqrt{2\pi n C_{p,n}}}, \end{aligned} \quad (\text{A.26})$$

where $C_{p,n} = \frac{1}{2} \log \frac{1+\delta_{p,n}}{1-\delta_{p,n}} - \xi$.

If $\log(p)/n \rightarrow \infty$ as $n \rightarrow \infty$, then $a_{p,n} + b_{p,n}F_n(\alpha) \rightarrow 1$. Therefore $\delta_{p,n} \rightarrow 1$, which yields $C_{p,n} \rightarrow \infty$. Then the tail probability in (A.26) goes to zero as $n \rightarrow \infty$. It follows that $P(B) \rightarrow 1$ as $n \rightarrow \infty$.

If $\log(p)/n \rightarrow \eta_0$ as $n \rightarrow \infty$, then $\delta_{p,n} \rightarrow \max\{1 - e^{-4\eta_0}, \eta\}$. Under assumption (B2) that $\rho_{ij} < \max\{1 - e^{-4\eta_0}, \eta\}$, $\lim_{n \rightarrow \infty} C_{p,n} = \lim_{n \rightarrow \infty} \frac{1}{2} \log \frac{1+\max\{1-e^{-4\eta_0}, \eta\}}{1-\max\{1-e^{-4\eta_0}, \eta\}} - \xi > 0$. Again the tail probability in (A.26) goes to zero as $n \rightarrow \infty$. It follows that $P(B) \rightarrow 1$ as $n \rightarrow \infty$.

If $\log(p)/n \rightarrow 0$ as $n \rightarrow \infty$, then $a_{p,n} + b_{p,n}F_n(\alpha) \rightarrow 0$. Hence $\delta_{p,n} \rightarrow \eta$. Under the assumption (B2), we have $\lim_{n \rightarrow \infty} C_{p,n} = \log \frac{1+\eta}{1-\eta} - \xi > 0$. Therefore $P(B) \rightarrow 1$ as $n \rightarrow \infty$.

Given $P(A) \rightarrow 1$ and $P(B) \rightarrow 1$, we have $P(\mathcal{C} \cap \mathcal{M} \subset \mathcal{M}^*) \rightarrow 1$ as $n \rightarrow \infty$. \square

Proof of Theorem 4. It follows from (4.17) directly that

$$\|(C_{21}^{(2)} - C_{21}^{(1)}(C_{11}^{(11)})^{-1}C_{11}^{(12)})(C_{11}^{(22)} - C_{11}^{(21)}(C_{11}^{(11)})^{-1}C_{11}^{(12)})^{-1} \text{sign}(\beta_1^{(2)})\|_{\max} \leq 1 - \xi, \quad (\text{A.27})$$

where $\|\cdot\|_{\max}$ denotes the max norm of a matrix. Based the definition of \mathcal{C} , we have the following element wise inequalities $\|C_{11}^{(12)}\|_{\max} \leq c_{n,p,\alpha}$, $\|C_{11}^{(21)}\|_{\max} \leq c_{n,p,\alpha}$. Here $c_{n,p,\alpha}$ is the pairwise correlation screening bound. Since $C_{11}^{(11)}$ is positive definite, there exists an orthogonal

matrix Q s.t. $C_{11}^{(11)} = Q\Lambda Q^T$, where Λ is a diagonal matrix consists of the eigenvalues of $C_{11}^{(11)}$.

By assumption, we have $\lambda_{\min}(C_{11}^{(11)}) \geq \lambda_0$. Therefore $\|C_{11}^{(21)}(C_{11}^{(11)})^{-1}C_{11}^{(12)}\|_{\max} \leq \lambda_0^{-1}c_{n,p,\alpha}^2s_1^2$.

Under the assumption that $\log(p)/n \rightarrow 0$, $c_{n,p,\alpha} = o_n(1)$. It follows that $\lambda_0^{-1}c_{n,p,\alpha}^2s_1^2 = o_n(1)$.

By assumption (B2), $\|C_{21}^{(1)}\|_{\max} \leq \eta$. Thus $\|C_{21}^{(1)}(C_{11}^{(11)})^{-1}C_{11}^{(12)}\|_{\max} \leq \lambda_0^{-1}\eta c_{n,p,\alpha}^2s_1^2$, then

$\|C_{21}^{(1)}(C_{11}^{(11)})^{-1}C_{11}^{(12)}\|_{\max} = o_p(1)$ as $n \rightarrow \infty$. Therefore

$$\begin{aligned} & \| (C_{21}^{(2)} - C_{21}^{(1)}(C_{11}^{(11)})^{-1}C_{11}^{(12)})(C_{11}^{(22)} - C_{11}^{(21)}(C_{11}^{(11)})^{-1}C_{11}^{(12)})^{-1}\text{sign}(\beta_1^{(2)}) - C_{21}^{(2)}(C_{11}^{(22)})^{-1}\text{sign}(\beta_1^{(2)}) \|_{\max} \\ &= \| (C_{21}^{(2)}(C_{11}^{(22)})^{-1}C_{11}^{(21)}(C_{11}^{(11)})^{-1}C_{11}^{(12)} - C_{21}^{(1)}(C_{11}^{(11)})^{-1}C_{11}^{(12)})(C_{11}^{(22)} - C_{11}^{(21)}(C_{11}^{(11)})^{-1}C_{11}^{(12)})^{-1}\text{sign}(\beta_1^{(2)}) \|_{\max} \end{aligned}$$

Write $A = C_{21}^{(2)}(C_{11}^{(22)})^{-1}C_{11}^{(21)}(C_{11}^{(11)})^{-1}C_{11}^{(12)}$, $B = C_{21}^{(1)}(C_{11}^{(11)})^{-1}C_{11}^{(12)}$, $D = C_{11}^{(21)}(C_{11}^{(11)})^{-1}C_{11}^{(12)}$,

and $Y = \text{sign}(\beta_1^{(2)})$. Then the above term becomes $\|(A - B)(C_{11}^{(22)} - D)^{-1}Y\|_{\max}$. Moreover,

we have

$$\|(A - B)(C_{11}^{(22)} - D)^{-1}Y\|_{\max} \leq (s - s_1)\|A - B\|_{\max}\|(C_{11}^{(22)} - D)^{-1}Y\|_{\max}.$$

Since $\|A\|_{\max} \leq \lambda_0^{-1}(s - s_0)^2\|C_{21}^{(2)}\|_{\max}\|C_{11}^{(21)}(C_{11}^{(11)})^{-1}C_{11}^{(12)}\|_{\max} \leq \lambda_0^{-2}\eta c_{n,p,\alpha}^2s_1^2(s - s_1)^2$, $\|B\|_{\max} \leq$

$\lambda_0^{-1}\eta c_{n,p,\alpha}^2s_1^2$, and

$$\|(C_{11}^{(22)} - D)^{-1}Y\|_{\max} \leq (s - s_1)\|(C_{11}^{(22)} - D)^{-1}\|_{\max} \leq (s - s_1)(\lambda_0 - \lambda_0^{-1}c_{n,p,\alpha}^2s_1^2)^{-1}.$$

Therefore we have

$$\begin{aligned} \|(A - B)(C_{11}^{(22)} - D)^{-1}Y\|_{\max} &\leq (s - s_0)^2(\lambda_0^{-2}\eta c_{n,p,\alpha}^2s_1^2(s - s_1)^2 + \lambda_0^{-1}\eta c_{n,p,\alpha}^2s_1^2)(\lambda_0 - \lambda_0^{-1}c_{n,p,\alpha}^2s_1^2)^{-1} \\ &= o_p(1), \end{aligned}$$

as $n \rightarrow \infty$. It follows that $C_{21}^{(2)}(C_{11}^{(22)})^{-1}\text{sign}(\boldsymbol{\beta}_1^{(2)}) < 1 - \xi/2$ with probability tending to 1 as $n \rightarrow \infty$ which concludes the proof if we take $\delta = \xi/2$. \square

References

- [1] Richard Arratia, Larry Goldstein, and Louis Gordon. Two moments suffice for poisson approximations: the chen-stein method. *The Annals of Probability*, 17(1):9–25, 1989.
- [2] Howard D Bondell and Brian J Reich. Simultaneous regression shrinkage, variable selection, and supervised clustering of predictors with oscar. *Biometrics*, 64(1):115–123, 2008.
- [3] Patrick Breheny and Jian Huang. Group descent algorithms for nonconvex penalized linear and logistic regression models with grouped predictors. *Statistics and Computing*, 25(2):173–187, 2015.
- [4] Peter Bühlmann, Markus Kalisch, and Lukas Meier. High-dimensional statistics with a view toward applications in biology. *Annual Review of Statistics and Its Application*, 1:255–278, 2014.
- [5] T Tony Cai and Tiefeng Jiang. Phase transition in limiting distributions of coherence of high-dimensional random matrices. *Journal of Multivariate Analysis*, 107:24–39, 2012.
- [6] T Tony Cai, Tiefeng Jiang, et al. Limiting laws of coherence of random matrices with applications to testing covariance structure and construction of compressed sensing matrices. *The Annals of Statistics*, 39(3):1496–1525, 2011.
- [7] Emmanuel Candès and Terence Tao. The dantzig selector: Statistical estimation when p is much larger than n . *The Annals of Statistics*, 35(6):2313–2351, 2007.

- [8] Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, 2001.
- [9] Jianqing Fan and Jinchi Lv. Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(5):849–911, 2008.
- [10] Jianqing Fan and Jinchi Lv. A selective overview of variable selection in high dimensional feature space. *Statistica Sinica*, 20(1):101, 2010.
- [11] Jerome Friedman, Trevor Hastie, and Rob Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1, 2010.
- [12] Fang Han and Han Liu. Distribution-free tests of independence with applications to testing more structures. *arXiv preprint arXiv:1410.4179*, 2014.
- [13] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The elements of statistical learning*, volume 1. Springer series in statistics Springer, Berlin, 2003.
- [14] Caiyan Li and Hongzhe Li. Network-constrained regularization and variable selection for analysis of genomic data. *Bioinformatics*, 24(9):1175–1182, 2008.
- [15] Wei Pan, Benhuai Xie, and Xiaotong Shen. Incorporating predictor network in penalized regression with application to microarray data. *Biometrics*, 66(2):474–484, 2010.
- [16] Mee Young Park, Trevor Hastie, and Robert Tibshirani. Averaged gene expressions for regression. *Biostatistics*, 8(2):212–227, 2007.

- [17] Dhruv B Sharma, Howard D Bondell, and Hao Helen Zhang. Consistent group identification and variable selection in regression with correlated predictors. *Journal of Computational and Graphical Statistics*, 22(2):319–340, 2013.
- [18] Yiyuan She. Sparse regression with exact clustering. *Electronic Journal of Statistics*, 4:1055–1096, 2010.
- [19] Noah Simon, Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Regularization paths for coxs proportional hazards model via coordinate descent. *Journal of Statistical Software*, 39(5):1, 2011.
- [20] Charles Spearman. The proof and measurement of association between two things. *The American Journal of Psychology*, 15(1):72–101, 1904.
- [21] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.
- [22] Guan Yu and Yufeng Liu. Sparse regression incorporating graphical structure among predictors. *Journal of the American Statistical Association*, 111(514):707–720, 2016.
- [23] Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.
- [24] Kai Zhang. Spherical cap packing asymptotics and rank-extreme detection. *IEEE Transactions on Information Theory*, 63(7):4572–4584, 2017.
- [25] Peng Zhao and Bin Yu. On model selection consistency of lasso. *Journal of Machine learning research*, 7(Nov):2541–2563, 2006.

- [26] Yunzhang Zhu, Xiaotong Shen, and Wei Pan. Simultaneous grouping pursuit and feature selection over an undirected graph. *Journal of the American Statistical Association*, 108(502):713–725, 2013.
- [27] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.

Siliang Gong

University of Pennsylvania

E-mail: (siliang@pennmedicine.upenn.edu)

Kai Zhang

The University of North Carolina at Chapel Hill

E-mail: (zhangk@email.unc.edu)

Yufeng Liu

The University of North Carolina at Chapel Hill

E-mail: (yfliu@email.unc.edu)