

# High-Dimensional Interactions Detection with Sparse Principal Hessian Matrix

**Cheng Yong Tang**

*Department of Statistical Science  
Fox School of Business  
Temple University  
Philadelphia, PA 19122-6083, USA*

YONGTANG@TEMPLE.EDU

**Ethan X. Fang**

*Department of Statistics  
Pennsylvania State University  
University Park, PA 16801, USA*

XXF13@PSU.EDU

**Yuexiao Dong**

*Department of Statistical Science  
Fox School of Business  
Temple University  
Philadelphia, PA 19122-6083, USA*

YDONG@TEMPLE.EDU

**Editor:** Jie Peng

## Abstract

In statistical learning framework with regressions, interactions are the contributions to the response variable from the products of the explanatory variables. In high-dimensional problems, detecting interactions is challenging due to combinatorial complexity and limited data information. We consider detecting interactions by exploring their connections with the principal Hessian matrix. Specifically, we propose a one-step synthetic approach for estimating the principal Hessian matrix by a penalized M-estimator. An alternating direction method of multipliers (ADMM) is proposed to efficiently solve the encountered regularized optimization problem. Based on the sparse estimator, we detect the interactions by identifying its nonzero components. Our method directly targets at the interactions, and it requires no structural assumption on the hierarchy of the interactions effects. We show that our estimator is theoretically valid, computationally efficient, and practically useful for detecting the interactions in a broad spectrum of scenarios.

**Keywords:** Interaction Detection; Principal Hessian Matrix; ADMM; Sparse M-Estimator

## 1. Introduction

The interaction effect is an important consideration in regression problems that are commonly encountered in practice. It refers to a general situation when it is not adequate to build a model with the original explanatory variables alone in a simple additive way – referred to as the main effects. At the same time, it could be more effective to build a model incorporating constructed variables from the products of the variables – referred to as the interactions between the variables; see, for example, the monograph of Jaccard and Turrisi (2003) for an overview. Since the interactions are constructed from the explanatory

variables, it is clear that they can be further expanded in a hierarchical manner. That is, the construction can continue searching for higher order interactions that can be defined based on all variables including those previously constructed ones. In current paradigm of scientific investigations such as genome-wide association studies (GWAS), attempting to find interactions has received substantial recent attentions; see, for example, Cordell (2009); Dato et al. (2018); Ritchie and Steen (2018). Indeed, as pointed out in Ritchie and Steen (2018), the development of modeling approaches and computational techniques for identifying interaction items will enable practitioners to effectively identify important genetic features for new advances in medical studies.

Recently, there has been a surge in the development of methods for solving high-dimensional regression problems. For an overview of the current state of knowledge, we refer to the monographs Bühlmann and van de Geer (2011), Hastie et al. (2015), the review by Fan and Lv (2010) and references therein. Conceptually speaking, existing penalized regression methods can be applied by incorporating the constructed interactions as new variables. However, because of the aforementioned interactive and hierarchical nature, the complexity of the problem scales up very quickly, both computationally for practical implementations and theoretically for analyzing the properties of the estimators. For example, in a typical situation with  $p$  explanatory variables, the number of the possible two-way interactions is  $p(p-1)/2$ , growing at a quadratic rate of  $p$  that itself is typically assumed growing at some exponential rate of the sample size  $n$ . Consequently, on one hand, with the total number of explanatory variables including interactions at the order of  $\mathcal{O}(p^2)$ , the computational complexity increases substantially for estimation. On the other hand, due to the increased data dimensionality, the relative signal strength from the explanatory variables becomes substantially weaker, and makes detection these signals much more challenging.

There also has been active development on methods for detecting and estimating the interactions in high-dimensional regression problems. To address the challenges, a main strategy is to impose some hierarchical structure between the main effect and the interactions. That is, in the absence of the original variables, the interactions alone do not contribute to the model. This is referred to as the marginality principle (Nelder, 1977). In a class of the existing methods guided by this principle, the interactions are detected and estimated by some two-step or iterative two-step procedures that first select the main effects by some screening or penalized regression methods, and then perform a second run of selections by incorporating all potential interactions from the selected variables. For examples of the methods belonging to this class, we refer to Yuan et al. (2009), Zhao et al. (2009), Choi et al. (2010), Bien et al. (2013), Hao and Zhang (2014), Shah (2016), Haris et al. (2016), and Hao et al. (2018). More recently, She et al. (2018) investigate the problem by a group penalty approach incorporating the structural hierarchy. Requiring no hierarchical structure and in a setting with multiple response, Kong et al. (2017) propose to search for interactions by a sure screening procedure based on the square functions of the distance correlation measures. This approach is computationally efficient, and only requires ranking with pairwise quantities. From the inverse modeling perspective, Jiang and Liu (2014) consider testing procedures by the sliced inversion regression approach (Li, 1991) for variable selections of the first- and second-order effects. For classification problems and general index models, Li and Liu (2017) study some information criteria based approaches for variable and interaction selections. Recently, Fan et al. (2016) propose a method to detect the signals

from the interactions by examining the pairwise covariances between the squared response and each explanatory variable. Upon evaluating the covariances, Fan et al. (2016) propose a new two-step procedure that first conducts sure screening for the interactions, and then conducts a second step of penalized regression without imposing the hierarchical structure.

For those two-step procedures, detecting interactions in the second step upon selecting some variables such as the main effect has a clear limitation in the high-dimensional scenarios. Particularly for the ones with some hierarchical structures, if the signal strength relative to the noise from the main effects is weak, then the method in the first step may miss the main effect. Subsequently, it becomes impossible to detect the interactions in the second step. Under high-dimensional settings, the noise aggregates as the dimension increases, and decreases the relative signal strength. Such weaker signals make missing main effects in the first step more likely to happen.

In this paper, we propose a new method to detect the interaction effects in regression problems by a one-step penalized M-estimator. Our method does not assume a hierarchical structure, and it also requires no screening step. Our method is developed by utilizing the so-called principal Hessian matrix (Li, 1992), which is defined as the expectation of the second order derivative of the mean function. To solve the challenging high-dimensional problem, we assume that the principal Hessian matrix is sparse, reflecting the reality that given limited data information, only few meaningful interactions can be supported with desired estimation accuracy. A sparsity encouraging M-estimator is proposed by minimizing a dedicated crafted penalized squared loss function. As shown in our development, this procedure involves a matrix optimization problem, which is substantially more challenging than those vector optimization problems (Ding et al., 2014). To meet the challenge, we propose an alternating direction method of multipliers (ADMM) based algorithm to compute the estimator efficiently. We show by simulations that the proposed method outperforms existing methods, and our theory confirms that the estimator works satisfactorily, allowing the dimensionality of the explanatory variables growing exponentially with the sample size.

The rest of this paper is organized as follows. The proposed method using the principal Hessian matrix for detecting interactions is outlined in Section 2, followed by the alternating direction method of multipliers (ADMM) for solving the optimizations elaborated in Section 3. Numerical examples including simulation and a real data example are presented in Section 4 to demonstrate the promising performance of the method. Section 5 gives theoretical analysis, and the proofs are provided in the Appendix.

## 2. Methods

### 2.1. Principal Hessian Matrix

The key device in our method is the principal Hessian matrix. In Li (1992), the principal Hessian matrix is proposed as a powerful device for investigating dimension reduction and data visualization. We consider the regression problem with response variable  $Y$  and explanatory variable  $X = (X_1, \dots, X_p)^T \in \mathbb{R}^p$  that is taken as a random vector. Let  $m(X) = \mathbb{E}(Y|X)$  be the conditional mean function. The principal Hessian matrix of Li (1992) is then defined as  $\Psi = \mathbb{E}\{\partial^2 m(X)/\partial X \partial X^T\}$  where the expectation is taken with respect to the joint distribution of  $(Y, X^T)^T$ . By the Stein's Lemma (Stein, 1981), Li (1992)

shows that if  $X \sim N(0, \Sigma_X)$ , then the principal Hessian matrix  $\Psi$  satisfies

$$\Psi = \Sigma_X^{-1} \Sigma_{YXX} \Sigma_X^{-1}, \quad (1)$$

where  $\Sigma_{YXX} = \mathbb{E}[\{Y - \mathbb{E}(Y)\}XX^T]$ . The principal Hessian matrix is an effective device for the so-called sufficient dimension reduction; see Li (1992) and Cook (1998).

Since  $m(X) = \mathbb{E}(Y|X)$ , a model for  $Y$  can be written as  $Y = m(X) + \varepsilon$  with  $\varepsilon$  being a zero mean random variable independent of  $X$ . For ease of presentation and without loss of generality, we assume that  $X$  is centered so that all of its components  $X_1, \dots, X_p$  are of mean zero. We now demonstrate the connection between the principal Hessian matrix  $\Psi$  in (1) and the interaction effects. Let us consider the following working model with both the main and interaction effects:

$$Y = m(X) = \sum_{i=1}^p \theta_i X_i + \sum_{i=1}^p \sum_{j=i}^p \beta_{ij} X_i X_j + \varepsilon. \quad (2)$$

Then, by (2), it holds that

$$\mathbb{E} \left\{ \frac{\partial^2 m(X)}{\partial X_i \partial X_j} \right\} = \begin{cases} \psi_{ij} = \psi_{ji} = \beta_{ij}, & \text{when } i \neq j, \\ 2\beta_{jj}, & \text{when } i = j, \end{cases} \quad (3)$$

where  $\psi_{ij}$  is the  $(i, j)$ -th component of  $\Psi$ . For the case  $i \neq j$ , the first equality in (3) is by definition (1), the second equality is due to that  $\Psi$  is symmetric, and the third equality is from the working model (2). Hence, it is intuitive that the principal Hessian matrix  $\Psi$  is a very informative device for detecting interactions. Immediately from (1) and (3), we have the following result for detecting interactions with  $\Psi$ .

**Proposition 1** *If  $X \sim N(0, \Sigma_X)$ , then under model (2),  $\psi_{ij} = \psi_{ji} \neq 0$  if and only if  $\beta_{ij} \neq 0$ .*

The normality assumption in Proposition 1 is inherited from the Stein's Lemma. For detecting interactions in model (2), the normality assumption can be relaxed. For example, in the ideal case that  $\Sigma_X = I$ , then  $\Psi = \Sigma_{YXX}$ , and  $\psi_{ij} = \psi_{ji} = \beta_{ij}$  remains true provided that  $\mathbb{E}(X_i X_j X_k) = 0$  ( $i, j, k = 1, \dots, p$ ). For a general  $\Sigma_X$ , the normality assumption in Proposition 1 can be replaced by the following moment conditions:

**Condition 2**  $X = (X_1, \dots, X_p)^T$  satisfies

C1.  $\mathbb{E}(X_i X_j X_k) = 0$  for any  $i, j, k = 1, \dots, p$ ;

C2. Let  $Z = (Z_1, \dots, Z_p)^T = \Sigma_X^{-1} X$ . It holds that  $\text{Cov}(X_k X_\ell, Z_i Z_j) \neq 0$  if and only if  $k = i, \ell = j$  or  $k = j, \ell = i$ .

**Proposition 3** *Under Condition 2 and model (2),  $\psi_{ij} = \psi_{ji} \neq 0$  if and only if  $\beta_{ij} \neq 0$ .*

The proof of Proposition 3 is given in the Appendix. Condition 2 is clearly valid for the normal distribution, and it holds more broadly. For example, it holds for the elliptical symmetric distributions whose characteristic function takes the form  $\phi(t^T \Sigma_X t)$ . Other

practically appealing situations for Condition 2 include when  $X$  contains independent components, and/or there are only few nonzero  $\beta_{ij}$  in (2).

Here C1 of Condition 2 ensures that the linear contribution from  $X$  in  $m(X)$  does not interfere with that from the interactions. We further note that if there is no linear contribution from  $X$ , or the main effect has been consistently estimated and removed from  $m(X)$ , then C1 in Condition 2 is not required for detecting the interactions. That is, consistent estimation of the main effects helps alleviating the conditions on the distribution of  $X$ . Existing high-dimensional penalized regression methods, e.g., those in Fan and Lv (2010), Bühlmann and van de Geer (2011), and Hastie et al. (2015) can potentially be performed for such a purpose, but some cautions are needed. To ensure that the linear effects can be adequately recovered in a working model with no interactions, some conditions are required. A sufficient condition is that  $\mathbb{E}\{X_i(X_j X_k)\} = 0$  for all  $i$  such that  $\theta_i \neq 0$ , and all  $(j, k)$  such that  $\beta_{jk} \neq 0$ . Furthermore, C2 of Condition 2 is on the correlations between  $X_k X_\ell$  and  $Z_i Z_j$ . Since  $\text{Cov}(Z, X) = \text{Cov}(\Sigma_X^{-1} X, X) = I$ , zero correlations between  $X_k$  and  $Z_i$  for  $k \neq i$  are already ensured. Instead of imposing the normality assumption, C2 in Condition 2 is seen as on correlations between second order terms, which can be verified empirically; and transformation may be developed if concerns arise.

We note that by its definition,  $\Psi$  is suitable for studying interactions between numerical explanatory variables. For categorical variables, dummy variable coding is needed, resulting in sub-groups of the observations so that our methods and other penalized regression analysis can be equally applied within the resulting groups for detecting the interactions between numerical and categorical variables.

## 2.2. Sparse Estimation of $\Psi$

Our study intends to explore an effective estimator for  $\Psi$  under high-dimensional settings, and then to detect the interactions. In high-dimensional regression problems, sparsity is a useful notion for statistical inferences; see, among others, Hastie et al. (2015). In the context of interactions in  $m(X)$ , it means that the total number of contributing pairs of the explanatory variables is small, so that  $\Psi$  is sparse. Thus, exploring the interactions becomes a sparse matrix estimation problem.

By examining  $\Psi$  in (1), we observe that a sparse estimation of this  $p \times p$  matrix is challenging; and plugging-in individual components is not feasible for solving high-dimensional problems. First, estimating  $\Sigma_X^{-1}$  is difficult when  $p$  is large. Though existing methods can be applied, e.g., those reviewed in Fan et al. (2016), there are two major concerns. First is on whether or not it is suitable to impose assumptions such as  $\Sigma_X^{-1}$  being sparse. Second is that if one constructs an estimator of  $\Psi$  from sparse estimators of  $\Sigma_X^{-1}$  and  $\Sigma_{YXX}$ , the implied sparse components of  $\Psi$  may not be the desirable ones in  $\Psi$ . As an example, we consider  $Y = X_1 + X_1 X_2 + \epsilon$ ,  $X = (X_1, X_2, X_3)^T$ ,  $X \sim N(0, \Sigma_X)$  with  $\Sigma_X = (\sigma_{ij})_{i,j=1}^3$ , and  $\sigma_{ij} = \rho^{|i-j|}$  ( $i, j = 1, 2, 3$ ). Then we have

$$\Sigma_X^{-1} = \frac{1}{1-\rho^2} \begin{pmatrix} 1 & -\rho & 0 \\ -\rho & 1+\rho^2 & -\rho \\ 0 & -\rho & 1 \end{pmatrix} \text{ and } \Sigma_{YXX} = \begin{pmatrix} 2\rho & 1+\rho^2 & \rho+\rho^3 \\ 1+\rho^2 & 2\rho & 2\rho^2 \\ \rho+\rho^3 & 2\rho^2 & 2\rho^3 \end{pmatrix}.$$

Then, some algebra shows that in this case the principal Hessian matrix is

$$\Psi = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

That is,  $\Psi$  precisely reflects the signals from the interactions, despite that both  $\Sigma_X^{-1}$  and  $\Sigma_{YXX}$  are dense.

In our investigation, we consider the case that  $\Psi$  is sparse, with no sparsity requirements on  $\Sigma_X$ ,  $\Sigma_X^{-1}$ , and  $\Sigma_{YXX}$ . All our technical development is based on assumptions directly on  $\Psi$ . Specifically, we aim at a one-step sparse estimator of  $\Psi$ . The key development is the following. First, we observe that from (1),  $\Sigma_X \Psi \Sigma_X = \Sigma_{YXX}$ . This motivates to consider the weighted quadratic loss

$$\mathcal{L}(\Psi) = \text{tr}\{\Sigma_X^{-1}(\Sigma_X \Psi \Sigma_X - \Sigma_{YXX})^T \Sigma_X^{-1}(\Sigma_X \Psi \Sigma_X - \Sigma_{YXX})\}, \quad (4)$$

which is minimized at  $\Psi = \Sigma_X^{-1} \Sigma_{YXX} \Sigma_X^{-1}$ . With some elementary algebra, it can be shown that  $\mathcal{L}(\Psi)$  depends on  $\Psi$  only through

$$\text{tr}(\Psi^T \Sigma_X \Psi \Sigma_X) - 2\text{tr}(\Psi \Sigma_{YXX}).$$

We denote by  $\{y_i, x_i\}_{i=1}^n$  the observed data. Then, we propose to replace  $\Sigma_X$  and  $\Sigma_{YXX}$  by their sample counterparts:  $S = n^{-1} \sum_{i=1}^n x_i x_i^T$ ,  $Q = n^{-1} \sum_{i=1}^n (y_i - \bar{y}) x_i x_i^T$  with  $\bar{y} = n^{-1} \sum_{i=1}^n y_i$ , and consider sparse inducing  $\ell_1$ -penalized estimation. Concretely, we propose an M-estimator:

$$\hat{\Psi} = \arg \min_{\Psi} \left\{ \text{tr}(\Psi^T S \Psi S) / 2 - \text{tr}(\Psi Q) + \lambda \|\Psi\|_1 \right\}, \quad (5)$$

where  $\|A\|_1$  is the  $\ell_1$ -norm of the matrix  $A$ , and  $\lambda > 0$  is a tuning parameter. Given  $\hat{\Psi} = (\hat{\psi}_{ij})_{p \times p}$ , we propose to detect interactions as

$$\mathcal{M} = \{(i, j) : \hat{\psi}_{ij} \neq 0, i \leq j\}. \quad (6)$$

Here the function in (5) is quadratic in the matrix  $\Psi$ , which is different from those penalized estimators for handling vectors in the literature so that existing algorithms for solving penalized regressions do not apply. Because only linear operators are involved in (5) with no inverse of a large matrix, the calculation of the loss is computationally efficient and scalable. In the next section, we design an efficient algorithm to compute  $\hat{\Psi}$ . We also note that  $\hat{\Psi}$  may not symmetric. For practical applications, we suggest symmetrizing it by  $(\hat{\Psi} + \hat{\Psi}^T) / 2$ .

### 3. Algorithm

#### 3.1. Alternating Direction Method of Multipliers

We first observe that the Hessian of problem (5) is  $S \otimes S$  and positive-semidefinite, where  $\otimes$  denotes the Kronecker product. The loss function in (5) is thus convex. We propose to solve (5) by the alternating direction method of multipliers (ADMM). Our algorithm is inspired

by the algorithm developed in Jiang et al. (2018), in which a large matrix for classification with quadratic discriminant analysis is directly estimated in the same spirit of (5).

ADMM was first proposed in Glowinski and Marroco (1975), which is essentially a splitting version of the augmented Lagrangian method to solve optimization problems with a separable objective under a linear constraint:

$$\min_{\beta, \eta \in \mathbb{R}^p} g_1(\beta) + g_2(\eta), \text{ subject to } A\beta + B\eta = b,$$

where  $A, B \in \mathbb{R}^{k \times p}$ ,  $b \in \mathbb{R}^p$ , and  $g_1, g_2 : \mathbb{R}^p \rightarrow \mathbb{R}$  are continuous functions. Recently, ADMM finds wide applications in different fields such as statistics, image processing and machine learning. This is due to the algorithm's easy implementation and practical efficiency; see Boyd et al. (2011) and Eckstein and Yao (2012) for some reviews.

To apply the ADMM algorithm to compute our estimator, we rewrite problem (5) into the following equivalent form to facilitate the algorithm design that

$$\min_{\Psi, \Phi} \text{tr}(\Psi^T S \Psi S) / 2 - \text{tr}(\Psi Q) + \lambda \|\Phi\|_1, \text{ subject to } \Psi - \Phi = 0. \quad (7)$$

The augmented Lagrangian dual problem associated with the above problem is

$$\mathcal{L}_\rho(\Psi, \Phi, \Lambda) = \text{tr}(\Psi^T S \Psi S) / 2 - \text{tr}(\Psi Q) + \lambda \|\Phi\|_1 + \langle \Lambda, \Psi - \Phi \rangle + \frac{\rho}{2} \|\Psi - \Phi\|_F^2,$$

where  $\Lambda$  is the dual variable associated with the equality constraint, and  $\rho > 0$  is a penalty parameter. The ADMM algorithm runs iteratively, at the  $(t + 1)$ -th iteration, we update the solutions by

$$\begin{aligned} \Psi^{t+1} &= \arg \min_{\Psi} \mathcal{L}_\rho(\Psi, \Phi^t, \Lambda^t), \\ \Phi^{t+1} &= \arg \min_{\Phi} \mathcal{L}_\rho(\Psi^{t+1}, \Phi, \Lambda^t), \\ \Lambda^{t+1} &= \Lambda^t + \rho(\Psi^{t+1} - \Phi^{t+1}). \end{aligned} \quad (8)$$

Note that since our problem is convex, and the objective value is lower bounded, the convergence result of the algorithm has been well established in existing literature; see Fang et al. (2015) and Sun et al. (2015) for examples. We further point out that in our simulation studies in Section 4, the algorithm performs satisfactorily empirically.

### 3.2. Solving the Subproblems

The key to implement the ADMM algorithm developed above is to solve the  $\Psi$ - and  $\Phi$ -subproblems in (8) efficiently. In this subsection, we derive efficient solutions for the two subproblems. For the  $\Psi$ -subproblem, we have that

$$\begin{aligned} \Psi^{t+1} &= \arg \min_{\Psi} \text{tr}(\Psi^T S \Psi S) / 2 - \text{tr}(\Psi Q) + \text{tr}(\Lambda^t, \Psi) + \frac{\rho}{2} \|\Psi - \Phi^t\|_F^2 \\ &= \arg \min_{\Psi} \text{vec}(S \Psi)^T \text{vec}(S \Psi) / 2 - \text{vec}(Q)^T \text{vec}(\Psi) + \text{vec}(\Lambda^t)^T \text{vec}(\Psi) + \frac{\rho}{2} \|\text{vec}(\Psi) - \text{vec}(\Phi^t)\|_2^2 \\ &= \arg \min_{\Psi} \text{vec}(\Psi)^T \{(S \otimes I_p)^T (S \otimes I_p)\} \text{vec}(\Psi) / 2 - \text{vec}(Q)^T \text{vec}(\Psi) + \text{vec}(\Lambda^t)^T \text{vec}(\Psi) \\ &\quad + \frac{\rho}{2} \|\text{vec}(\Psi) - \text{vec}(\Phi^t)\|_2^2, \end{aligned}$$

where  $\otimes$  denotes Kroneker product. Solving this problem is computationally expensive as it does not admit a closed-form solution. We also point out that the problem becomes not tractable for modestly large  $d$  as the dimension for the problem above  $d^2$ . Inspired by Fang et al. (2015), we propose to add a proximal term to improve the computational efficiency. Let

$$\begin{aligned} \Psi^{t+1} = \arg \min_{\Psi} \text{vec}(\Psi)^T \{ (S \otimes I_p)^T (S \otimes I_p) \} \text{vec}(\Psi) / 2 - \text{vec}(Q)^T \text{vec}(\Psi) + \text{vec}(\Lambda^t)^T \text{vec}(\Psi) \\ + \frac{\rho}{2} \|\text{vec}(\Psi) - \text{vec}(\Phi^t)\|_2^2 + \frac{1}{2} \|\text{vec}(\Psi) - \text{vec}(\Psi^t)\|_G^2, \end{aligned}$$

where  $\|\cdot\|_G$  denotes the matrix induced norm  $\|x\|_G = \sqrt{x^T G x}$  for positive definite matrix  $G$ . Letting  $G = \tau I_{p^2} - \{(S \otimes I_p)^T (S \otimes I_p)\}$ , where  $\tau$  is greater than the largest eigenvalue of  $\{(S \otimes I_p)^T (S \otimes I_p)\}$ , we have

$$\Psi^{t+1} = \arg \min_{\Psi} \frac{\rho + \tau}{2} \|\text{vec}(\Psi)\|_2^2 - \left\{ \text{vec}(Q) - \text{vec}(\Lambda^t) + (\rho + \tau) \text{vec}(\Phi^t) - (S \otimes I_p) \text{vec}(\Phi^t) \right\}.$$

Then, we have

$$\text{vec}(\Psi^{t+1}) = \left\{ \text{vec}(Q) - \text{vec}(\Lambda^t) + (\rho + \tau) \text{vec}(\Phi^t) - (S \otimes I_p) \text{vec}(\Phi^t) \right\} / (\rho + \tau).$$

Thus, we have that by adding the proximal term, we have a closed-form solution for solving the  $\Psi$ -subproblem.

Next, considering the  $\Phi$ -subproblem, we have

$$\Phi^{t+1} = \arg \min_{\Phi} \mathcal{L}_{\rho}(\Psi^{t+1}, \Phi, \Lambda^t) = \arg \min_{\Phi} \rho \text{tr}(\Phi^T \Phi) / 2 - \rho \text{tr}(\Phi^T \Psi^{t+1}) - \text{tr}(\Phi^T \Lambda^t) + \lambda \|\Phi\|_1.$$

It is not difficult to see that this step admits a closed-form solution by soft thresholding:

$$\Phi^{t+1} = \text{shrinkage}(\Psi^{t+1} + \rho^{-1} \Lambda^t, \rho^{-1} \lambda),$$

where  $\text{shrinkage}(\cdot, \cdot) : \mathbb{R}^{p \times p} \times \mathbb{R} \rightarrow \mathbb{R}^{p \times p}$  is the elementwise shrinkage operator that for a matrix  $A \in \mathbb{R}^{p \times p}$  and  $\tau > 0$ , the  $(j, k)$ -th entry  $[\text{shrinkage}(A, \tau)]_{jk} = \text{sign}(A_{jk}) \cdot \max\{0, |A_{jk}| - \tau\}$ .

For the stopping criterion, we look at the primal and dual residuals. The primal residual is a measure of the feasibility for problem (7), which is defined as

$$\eta_P^t = \|\Psi^t - \Phi^t\|_F.$$

Meanwhile, the dual residual measures the convergence of the algorithm, where we take

$$\eta_D^t = \max\{\|\Psi^{t+1} - \Psi^t\|_F, \|\Phi^{t+1} - \Phi^t\|_F, \|\Lambda^{t+1} - \Lambda^t\|_F\}.$$

In our implementation, we stop the algorithm when both primal and dual residuals are small that

$$\max\{\eta_P^t, \eta_D^t\} \leq 10^{-3}.$$

We summarize the pseudo-code of the algorithm in Algorithm 1. As we have derived the closed-form solutions for the  $\Psi$ - and  $\Phi$ -subproblems in (8), which can be easily computed, Algorithm 1 can be implemented efficiently.



---

**Algorithm 1** ADMM Algorithm to Estimate  $\Psi$

---

- 1: **Input:**  $S, Q \in \mathbb{R}^{p \times p}$ ,  $\lambda > 0$ ,  $\Psi^0, \Phi^0, \Lambda^0 \in \mathbb{R}^{p \times p}$ ,  $\rho, \tau > 0$ ,  $t = 0$
  - 2: **Output:**  $\hat{\Psi}$
  - 3: **while** stopping criterion not satisfied **do**
  - 4:    $\Psi^{t+1} \leftarrow \left\{ \text{vec}(Q) - \text{vec}(\Lambda^t) + (\rho + \tau)\text{vec}(\Phi^t) - (S \otimes I_p)\text{vec}(\Phi^t) \right\} / (\rho + \tau)$
  - 5:    $\Phi^{t+1} \leftarrow \text{shrinkage}(\Psi^{t+1} + \Lambda^t / \rho, \lambda / \rho)$
  - 6:    $\Lambda^{t+1} \leftarrow \Lambda^t + \rho(\Psi^{t+1} - \Phi^{t+1})$
  - 7:    $t \leftarrow t + 1$
  - 8: **end while**
  - 9:  $\hat{\Psi} \leftarrow \Psi^t$
- 

To conclude our method and algorithm, we make some remarks. Our algorithm is designed for solve the matrix optimization problem (5). Our framework can be readily extended to a more general form of the weight loss function. That is, by letting  $Z_A = A^{1/2}(\Sigma_X \Psi \Sigma_X - \Sigma_{YXX})A^{1/2}$  with some symmetric weighting matrix  $A$ , we define the loss

$$\mathcal{L}_A(\Psi) = \text{tr}(Z_A^T Z_A) = \text{tr}\{A(\Sigma_X \Psi \Sigma_X - \Sigma_{YXX})^T A(\Sigma_X \Psi \Sigma_X - \Sigma_{YXX})\},$$

whose minimizer is also the principal Hessian matrix. Then  $\mathcal{L}_A(\Psi)$  depends on  $\Psi$  via

$$\text{tr}\{(A \Sigma_X \Psi^T \Sigma_X)(A \Sigma_X \Psi \Sigma_X)\} - 2\text{tr}(A \Sigma_X \Psi \Sigma_X A \Sigma_{YXX}).$$

By choosing  $A = \Sigma_X^{-1}$ , the loss becomes (5). Appropriate choice of  $A$  may improve the speed and stability of the algorithm. For example, when matrix  $S$  is ill-conditioned, the choice of  $A$  could be a viable device to solve the problem.

## 4. Numerical Studies

### 4.1. Synthetic Data

In this section, we conduct extensive numerical studies to demonstrate and validate the performance of our proposed method. We first conduct investigations using synthetic data. In our simulation setup, we fix the sample size as  $n = 100$ , and we consider different dimensions for  $p = 100, 200$  and  $300$ . Meanwhile, we generate the design matrix  $\mathbb{X} = (x_1, x_2, \dots, x_n)^T \in \mathbb{R}^{n \times p}$  by generating each sample  $x_i \in \mathbb{R}^p$  independently from a  $p$ -dimensional Gaussian distribution  $X \sim N(0, \Sigma)$ , where the covariance matrix  $\Sigma$  is either the identity matrix, or a Toeplitz matrix, i.e.  $\Sigma_{jk} = \rho^{|j-k|}$  for some  $\rho \in (0, 1)$ . We then generate the noises  $\epsilon_i$ 's independently from a normal random variable  $N(0, \sigma^2)$ , and we consider different  $\sigma$ 's. To thoroughly compare the proposed method with other methods, we consider the following nine models:

Model 1:  $Y = X_1 + X_5 + \epsilon$ ,

Model 2:  $Y = 0.6X_1X_2 + 0.8X_4X_5 + \epsilon$ ,

Model 3:  $Y = 0.6X_1X_2 + 0.8X_2X_3 + \epsilon$ ,

$$\text{Model 4: } Y = 0.5X_1^2 + 0.9X_5X_8 + \epsilon,$$

$$\text{Model 5: } Y = X_1^2 + X_5X_8 + X_9^2 + \epsilon,$$

$$\text{Model 6: } Y = X_1 + X_5 + X_1X_5 + \epsilon,$$

$$\text{Model 7: } Y = 0.1X_1 + 0.1X_5 + X_1X_5 + \epsilon,$$

$$\text{Model 8: } Y = X_1X_5 + X_2X_3 \cdot \epsilon,$$

$$\text{Model 9: } Y = \sum_{j=1}^9 X_jX_{j+1} + \epsilon.$$

In Model 1, we consider the case where only the main effects are present. This is a benchmarking case in the sense that there should be no false inclusion of the interactions for a valid detection method. In the next four models, we consider the cases where only interaction terms are presented under different scenarios. With the interaction only models, we intend to show the advantage of the proposed one-step procedure. Then, we consider two models where some hierarchical structures exist in Model 6 and Model 7 with different signal strength. Finally, we consider an example with the heteroscedasticity in Model 8, in which the conditional variance of the response variable is not homogeneous. Model 9 is a less sparse case.

We choose the tuning parameter by 10-fold cross-validation. We point out that after extensive numerical investigations, we find that our method is insensitive to the tuning parameter selection. We report the empirical true positive rate (TPR) and the false positive rate (FPR) by (6) for each data generating scheme after repeating each scheme 200 times. In particular, let  $\mathcal{I}$  be the true set of interaction terms, and let  $\hat{\mathcal{I}}$  be the interaction terms selected by the estimator  $\hat{\Psi}$ . TPR and FPR are defined as

$$\text{TPR} = \frac{|\mathcal{I} \cap \hat{\mathcal{I}}|}{|\hat{\mathcal{I}}|}, \text{ and } \text{FPR} = \frac{|\hat{\mathcal{I}} \setminus \mathcal{I}|}{\binom{d}{2} + d - |\mathcal{I}|}$$

respectively.

We compare our method with the interaction pursuit (IP) method (Fan et al., 2016), which is two-step method with a screening procedure as the first step, and the RAMP method (Hao et al., 2018), which is based on a hierarchical structure. We report the results in Tables 1-5.

Firstly, as seen in Table 1 for Model 1, our one-step method performs very well with very little false inclusion when no interactions are present. Other competing methods also are performing very well in the sense of little of no false inclusions. In all other models for detecting the contributing interactions, we see that the proposed method outperforms the other two methods by large margins in all settings except Model 6, where a hierarchical structure is presented with good signal strength, which favors the RAMP method when the main effect is correctly detected first. Nevertheless, we see that in Model 7, though a hierarchical model is presented, due to the low signal strength of the main effect of  $X_1$  and  $X_5$ , our method still outperforms the RAMP method substantially. This clearly demonstrates the advantage of our one-step estimation based interaction detection method.

In all tables, we report the averaged running time in seconds, where all experiments are conducted on an iMac with 3.2 GHz Intel Core i5 Processor and 16 GB memory. It is seen

that the IP method is most computationally efficient, which is expected as it is a screening based method. Meanwhile, from the computational perspective, our method’s efficiency is similar to the RAMP method. However, we point out that our method directly deals with the principal Hessian matrix and do not impose any hierarchical structure. In comparison, the RAMP method imposes a hierarchical structure, and essentially deals with vectors.

## 4.2. Real Data

We further apply the proposed method to analyze the GPL96 microarray dataset analyzed in McCall et al. (2010); Wu et al. (2013); Fang et al. (2017). This dataset contains 13,182 samples of more than 2,000 biological contexts generated from Affymetrix Human 133A (GPL96) arrays, and each sample has 12,704 genes. The data set was preprocessed and normalized using frozen-RMA (McCall et al., 2010) to reduce batch effects. We use  $n_1 = 235$  samples of breast tumor and  $n_2 = 537$  normal samples in the analysis. To improve the efficiency in our demonstration, we conduct some pre-screening. In particular, we use a plug-in estimator  $\tilde{\Psi} = S^-QS^-$  to estimate  $\Psi$ , where  $S^-$  denotes the Moore-Penrose pseudo-inverse of  $S$ . We then screen out the variables where the  $\ell_1$ -norm of the corresponding columns of  $\tilde{\Psi}$  are small. In our analysis, we screen the genes and keep  $p = 1,000$  genes.

We treat the disease status as responses, and randomly split the dataset into a training set and a testing set. Each training set contains 100 samples from the breast tumor group and 150 samples from the normal group. We repeat the random split 100 times. We first report the averaged testing errors and their standard errors in Table 6. It is seen that the proposed method achieves better testing errors than the other two methods. We further provide the five most selected interaction terms in Table 7. It is seen that BRCA1 and BRCA2 are frequently selected by our proposed method including the interaction between these two genes, while they are less frequently selected by the other two methods. It is well known in literature that BRCA1 and BRCA2 are of fundamental importance for breast tumor as shown in King et al. (2003), and this demonstrates the potential usefulness of the proposed method in analyzing real data.

Furthermore, we consider another disease Ewing Sarcoma, where we have  $n_1 = 57$  samples. Following same procedure, each time, we randomly select 33 (57.9%) of these samples and 150 samples from the normal group as training samples and the rest as testing samples. We repeat the random splitting 100 times and report the average classification error rates in Table 8. It is seen that the RAMP method does not perform well in this case, and our method performs better than the IP method. This demonstrates the importance of including interaction terms in the model.

## 5. Some Theoretical Analysis

### 5.1. Non-asymptotic Results

We now perform some theoretical analysis of the estimator (5). Outline of the proofs is provided in the Appendix of the paper. Denote by  $\Psi^*$  the unknown truth of the principal Hessian matrix, and  $\mathcal{U} = \text{supp}(\Psi^*)$  the support of  $\Psi$ , and  $q = |\mathcal{U}|$  be the cardinality of  $\mathcal{U}$ . Denote by  $\|A\|_\infty$  the max-norm of the matrix  $A$ .

The following lemma establishes a non-asymptotic result on the support of  $\hat{\Psi}$ .

**Lemma 4** *Let  $\mathcal{C}(\theta^*) = \{\Psi : \|\Psi_{\mathcal{U}^c}\|_1 \leq 3\|\Psi_{\mathcal{U}} - \Psi_{\mathcal{U}}^*\|_1\}$  be the cone depending on the truth  $\Psi^*$ . We have that  $\widehat{\Psi} \in \mathcal{C}(\theta^*)$  provided that the tuning parameter  $\lambda$  satisfies  $2\|S\Psi^*S - Q\|_\infty \leq \lambda/2$ .*

We also have the following lemma containing a non-asymptotic oracle's inequality.

**Lemma 5** *Let  $\widehat{\Delta} = \widehat{\Psi} - \Psi^*$ . If  $2\|S\Psi^*S - Q\|_\infty < \lambda/2$ , then*

$$2\text{tr}(\widehat{\Delta}S\widehat{\Delta}S) + \lambda\|\widehat{\Psi}_{\mathcal{U}} - \Psi_{\mathcal{U}}^*\|_1 + \lambda\|\widehat{\Psi}_{\mathcal{U}^c}\|_1 \leq 4\lambda\|\widehat{\Psi}_{\mathcal{U}} - \Psi_{\mathcal{U}}^*\|.$$

To establish an estimation error bound of  $\widehat{\Psi}$ , we need the following condition.

**Condition 6** (*Restricted eigenvalue condition*) *Let  $\Delta = \Psi - \Psi^*$ . For all  $\Psi \in \mathcal{C}(\Psi^*)$ ,*

$$\text{vec}^T(\Delta)(S \otimes S)\text{vec}(\Delta) \geq \gamma\|\text{vec}(\Delta)\|_2^2. \quad (9)$$

Then we have the following theorem, establishing a key non-asymptotic result on the estimating error bound of  $\widehat{\Psi}$ .

**Theorem 7** *Suppose that the restricted eigenvalue Condition 6 holds. If  $2\|S\Psi^*S - Q\|_\infty \leq \lambda/2$ , then  $\|\widehat{\Psi} - \Psi^*\|_1 \leq 4\lambda q/\gamma^2$ .*

The main implication from Theorem 7 is the good performance non-asymptotically in the sense of small error bound. The restricted eigenvalue Condition 6 essentially requires that on the set  $\mathcal{C}$ , the smallest eigenvalue of  $S \otimes S$  is strictly bounded away from 0. We remark that the condition (9) here is essentially not stronger than the analogous one with high-dimensional linear model  $Y = X\beta + \epsilon$  for sparse estimator of  $\beta$ . For example, as in Hastie et al. (2015) and Bühlmann and van de Geer (2011), the Lasso estimator requires  $(\beta - \beta^*)^T S(\beta - \beta^*) \geq \gamma_1 \|\beta - \beta^*\|_2^2$  for all  $\beta$  in an analogous restricted set depending on the unknown truth  $\beta^*$ . Then by observing the fact that the eigenvalues of  $H_1 = S \otimes S$  are the products of the eigenvalues of  $S$ , we see that Condition 6 is essentially of the same kind as those on the sample covariance matrix. This can be viewed as a merit of our method. In contrast, if one considers an alternative approach incorporating all interactions in a penalized regression, then the corresponding restricted eigenvalue condition is required on the matrix  $H_2 = n^{-1} \sum_{i=1}^n w_i w_i^T$  with  $w_i \in \mathbb{R}^{p(p+1)/2}$  containing all quadratic terms of the  $i$ th observation. Clearly, a much larger sample covariance matrix  $H_2$  is directly involved with fourth order moments; and conditions on its restricted eigenvalue are seen as more stringent. Additionally, comparing with hierarchical approaches using two-step procedures, the main difference is that our condition does not require condition on the model matrix to ensure correctly identifying the linear effects; see, for example, the irrepresentable condition of Hao et al. (2018).

## 5.2. Asymptotic Results

The restricted eigenvalue condition (9) is known important for establishing the error bounds of penalized estimators in regression problems; see, among others, Negahban et al. (2012) and Hastie et al. (2015). Since our estimator (5) is the minimizer of the sum a quadratic loss and  $\ell_1$  penalty, this type of the restricted eigenvalue condition is expected. It is known

in the literature that in an asymptotic setting assuming that the entries of  $X_i$  following some distributions satisfying some condition on the tail probabilistic behavior, then the condition (9) holds with probability tending to 1. We refer to the discussion in Negahban et al. (2012) and results in Raskutti et al. (2010), Raskutti et al. (2011), and Rudelson and Zhou (2011).

As in Lemma 4 and Theorem 7, the condition  $2\|S\Psi^*S - Q\|_\infty \leq \lambda/2$  imposes the requirement on the tuning parameter  $\lambda$ . By their definitions,  $Q$  is an estimator of  $\Sigma_{YXX}$ , and  $S$  is an estimator of  $\Sigma_X$ . As in an asymptotic setting with  $n \rightarrow \infty$  and  $p \rightarrow \infty$ , they all converge to the truths element-wise in probability under appropriate conditions on the tail distribution of  $Y$  and  $X$ . Thus, the tuning parameter is allowed going to 0 as  $n \rightarrow \infty$ , ensuring that  $\widehat{\Psi} \in \mathcal{C}(\Psi^*)$ , and  $\|\widehat{\Delta}\|_1 \rightarrow 0$  in probability, i.e. consistency of  $\widehat{\Psi}$ .

Formally, for showing the asymptotic properties of  $\widehat{\Psi}$ , we impose the following condition.

**Condition 8** *The random vectors  $\{(Y_i, X_i^T)^T\}_{i=1}^n$  are independent and identically distributed. The largest eigenvalue of the  $X_i$  is strictly bounded away from infinity. Both  $Y_i$  and  $X_{ij}$  ( $j = 1, \dots, p$ ) satisfy exponential tail property, i.e., there exist constants  $b_1, b_2, c_1$  and  $c_2$  such that*

$$\mathbb{P}(|Y_i| > t) \leq \exp\{-(t/b_1)^{c_1}\} \quad \text{and} \quad \mathbb{P}(|X_{ij}| > t) \leq \exp\{-(t/b_2)^{c_2}\}.$$

Condition 8 ensures the concentration property of  $\|Q - \Sigma_{YXX}\|_\infty$ . So that by choosing  $\lambda = C\sqrt{\frac{\log p}{n}}$ ,  $2\|Q - \Sigma_{YXX}\|_\infty \leq \lambda/2$  holds with probability at least  $1 - c_3 \exp(-c_4 n^{c_5} \lambda^2)$  when  $\log(p) = n^\eta$  with some constants  $C, c_3, c_4, c_5, \eta$ .

Then we have the following asymptotic result on the error bound of  $\widehat{\Psi}$ .

**Theorem 9** *Suppose that  $\Psi^*$  is sparse with support  $\mathcal{U}$ , and the cardinality  $|\mathcal{U}| = q$ . Suppose additionally that  $\|\Psi^*\|_\infty = O\left(\sqrt{n/\log p}\right)$ ,  $\|\Psi^*\Sigma_X\|_\infty = O(1)$ . Then under Condition 8, by choosing  $\lambda = C\sqrt{\frac{\log p}{n}}$ ,  $\|\widehat{\Psi} - \Psi^*\|_1 = \mathcal{O}_p(q\sqrt{\log p/n})$ .*

To ensure that zero components of  $\Psi^*$  are correctly estimated as zero by  $\widehat{\Psi}$ , we need the following irrepresentable condition.

**Condition 10** *(Irrepresentable condition) Let  $\Lambda = S \otimes S = \begin{pmatrix} \Lambda_{11} & \Lambda_{12} \\ \Lambda_{21} & \Lambda_{22} \end{pmatrix}$  with  $\Lambda_{11}$  and  $\Lambda_{22}$  corresponding to the partition of  $\text{vec}(\Psi_{\mathcal{U}}^*)$  and  $\text{vec}(\Psi_{\mathcal{U}^c}^*)$ . Then  $\max_j \|\Lambda_{11}^{-1}\Lambda_{12,j}\|_1 \leq 1 - \alpha$  for some  $\alpha > 0$ , where  $\Lambda_{12,j}$  is the  $j$ -th column of  $\Lambda_{12}$ .*

**Theorem 11** *By choosing  $\lambda = C\sqrt{\frac{\log p}{n}}$  and assuming Conditions 8 and 10,  $\widehat{\Psi}_{\mathcal{U}^c} = 0$  with probability tending to 1. Furthermore, if  $|\min(\Psi_{\mathcal{U}}^*)| > \|\lambda\Lambda_{11}^{-1}\|_\infty$ , then  $\widehat{\Psi}_{\mathcal{U}}$  is bounded away from zero with probability tending to 1.*

With Theorem 9 showing the consistence of  $\widehat{\Psi}$ , and Theorem 11 showing the correctness in identifying the nonzero component in  $\Psi$ , we demonstrate the validity of the our method in detecting the interactions.

Table 1: Quantitative performance of the interaction detection methods for Model 1 and Model 2. We report the averaged true positive and false positive discovery rates after repeating each simulation setup 200 times with  $n = 100$ . Note that  $\rho = 0$  refers to the case where  $\Sigma$  is the identity matrix.

| Model | $(\rho, \sigma)$ | Method | $d = 100$ |       |            | $d = 200$ |       |            | $d = 300$ |       |            |
|-------|------------------|--------|-----------|-------|------------|-----------|-------|------------|-----------|-------|------------|
|       |                  |        | TPR       | FPR   | Time(s)    | TPR       | FPR   | Time(s)    | TPR       | FPR   | Time(s)    |
| 1     | (0, 0.1)         | ADMM   | NA        | 0.03% | 0.41(0.05) | NA        | 0.01% | 1.27(0.13) | NA        | 0.00% | 2.13(0.25) |
|       |                  | IP     | NA        | 0.15% | 0.12(0.02) | NA        | 0.11% | 0.20(0.03) | NA        | 0.03% | 0.34(0.04) |
|       |                  | RAMP   | NA        | 0.00% | 0.39(0.04) | NA        | 0.00% | 1.15(0.10) | NA        | 0.00% | 1.95(0.19) |
|       | (0, 2.0)         | ADMM   | NA        | 0.04% | 0.38(0.04) | NA        | 0.01% | 1.15(0.07) | NA        | 0.00% | 1.99(0.14) |
|       |                  | IP     | NA        | 0.14% | 0.09(0.01) | NA        | 0.02% | 0.17(0.03) | NA        | 0.00% | 0.29(0.07) |
|       |                  | RAMP   | NA        | 0.00% | 0.37(0.06) | NA        | 0.00% | 1.20(0.13) | NA        | 0.00% | 2.04(0.22) |
|       | (0.2, 1.0)       | ADMM   | NA        | 0.05% | 0.44(0.03) | NA        | 0.03% | 1.33(0.17) | NA        | 0.01% | 2.25(0.29) |
|       |                  | IP     | NA        | 0.08% | 0.14(0.01) | NA        | 0.02% | 0.23(0.03) | NA        | 0.00% | 0.36(0.03) |
|       |                  | RAMP   | NA        | 0.00% | 0.35(0.05) | NA        | 0.00% | 1.05(0.13) | NA        | 0.00% | 1.92(0.31) |
|       | (0.2, 2.0)       | ADMM   | NA        | 0.05% | 0.46(0.03) | NA        | 0.02% | 1.23(0.18) | NA        | 0.00% | 2.07(0.30) |
|       |                  | IP     | NA        | 0.12% | 0.10(0.01) | NA        | 0.03% | 0.18(0.02) | NA        | 0.00% | 0.31(0.02) |
|       |                  | RAMP   | NA        | 0.00% | 0.50(0.10) | NA        | 0.00% | 1.15(0.20) | NA        | 0.00% | 2.23(0.22) |
| 2     | (0, 0.1)         | ADMM   | 99.0%     | 0.08% | 0.44(0.08) | 96.0%     | 0.09% | 0.95(0.12) | 92.5%     | 0.13% | 2.31(0.32) |
|       |                  | IP     | 82.0%     | 0.06% | 0.08(0.01) | 67.0%     | 0.08% | 0.23(0.04) | 66.0%     | 0.02% | 0.34(0.04) |
|       |                  | RAMP   | 1.00%     | 0.03% | 0.48(0.10) | 0.00%     | 0.06% | 0.99(0.13) | 0.00%     | 0.01% | 2.03(0.24) |
|       | (0, 1.0)         | ADMM   | 98.5%     | 0.12% | 0.40(0.05) | 92.0%     | 0.13% | 1.22(0.10) | 90.0%     | 0.17% | 2.18(0.29) |
|       |                  | IP     | 54.0%     | 0.01% | 0.07(0.02) | 47.5%     | 0.04% | 0.18(0.03) | 42.5%     | 0.04% | 0.41(0.05) |
|       |                  | RAMP   | 2.00%     | 0.22% | 0.38(0.02) | 0.00%     | 0.06% | 1.15(0.09) | 0.00%     | 0.00% | 2.30(0.25) |
|       | (0.1, 0.1)       | ADMM   | 96.5%     | 0.14% | 0.35(0.09) | 93.5%     | 0.18% | 0.93(0.14) | 91.0%     | 0.11% | 2.20(0.28) |
|       |                  | IP     | 73.0%     | 0.19% | 0.09(0.02) | 68.5%     | 0.11% | 0.20(0.03) | 65.0%     | 0.08% | 0.38(0.05) |
|       |                  | RAMP   | 0.00%     | 0.00% | 0.29(0.11) | 0.00%     | 0.00% | 1.23(0.12) | 0.00%     | 0.00% | 2.37(0.30) |
|       | (0.1, 1.0)       | ADMM   | 95.0%     | 0.30% | 0.37(0.03) | 91.5%     | 0.26% | 1.20(0.18) | 89.0%     | 0.13% | 2.34(0.31) |
|       |                  | IP     | 49.5%     | 0.18% | 0.10(0.03) | 40.5%     | 0.10% | 0.25(0.03) | 28.0%     | 0.06% | 0.31(0.03) |
|       |                  | RAMP   | 0.00%     | 0.00% | 0.45(0.02) | 0.00%     | 0.00% | 1.27(0.23) | 0.00%     | 0.00% | 2.48(0.43) |

Table 2: Quantitative performance of the interaction detection methods for Model 3 and Model 4. We report the averaged true positive and false positive discovery rates after repeating each simulation setup 200 times with  $n = 100$ . Note that  $\rho = 0$  refers to the case where  $\Sigma$  is the identity matrix.

| Model | $(\rho, \sigma)$ | Method | $d = 100$ |       |            | $d = 200$ |       |            | $d = 300$ |       |            |
|-------|------------------|--------|-----------|-------|------------|-----------|-------|------------|-----------|-------|------------|
|       |                  |        | TPR       | FPR   | Time(s)    | TPR       | FPR   | Time(s)    | TPR       | FPR   | Time(s)    |
| 3     | (0.1, 0.1)       | ADMM   | 99.5%     | 0.10% | 0.36(0.03) | 95.0%     | 0.08% | 1.03(0.05) | 94.5%     | 0.13% | 2.04(0.10) |
|       |                  | IP     | 93.5%     | 0.19% | 0.08(0.01) | 82.5%     | 0.05% | 0.17(0.02) | 80.5%     | 0.03% | 0.28(0.02) |
|       |                  | RAMP   | 2.00%     | 0.02% | 0.38(0.02) | 2.00%     | 0.01% | 1.09(0.06) | 2.00%     | 0.00% | 2.23(0.15) |
|       | (0.1, 1.0)       | ADMM   | 96.5%     | 0.29% | 0.46(0.05) | 93.0%     | 0.13% | 0.96(0.17) | 91.5%     | 0.08% | 2.10(0.14) |
|       |                  | IP     | 76.0%     | 0.09% | 0.15(0.02) | 68.5%     | 0.03% | 0.25(0.05) | 62.0%     | 0.02% | 0.37(0.04) |
|       |                  | RAMP   | 0.00%     | 0.00% | 0.43(0.06) | 1.00%     | 0.00% | 0.99(0.10) | 0.00%     | 0.00% | 2.39(0.18) |
|       | (0.4, 0.1)       | ADMM   | 100.0%    | 0.20% | 0.40(0.03) | 99.0%     | 0.15% | 1.35(0.12) | 96.0%     | 0.07% | 2.36(0.28) |
|       |                  | IP     | 97.0%     | 0.10% | 0.07(0.01) | 91.5%     | 0.08% | 0.15(0.04) | 89.0%     | 0.05% | 0.24(0.06) |
|       |                  | RAMP   | 2.00%     | 0.00% | 0.49(0.05) | 0.00%     | 0.00% | 1.34(0.14) | 3.00%     | 0.00% | 2.59(0.30) |
|       | (0.4, 1.0)       | ADMM   | 99.0%     | 0.40% | 0.43(0.07) | 99.0%     | 0.19% | 1.31(0.16) | 97.0%     | 0.10% | 2.39(0.21) |
|       |                  | IP     | 92.0%     | 0.10% | 0.06(0.02) | 89.5%     | 0.08% | 0.16(0.03) | 86.0%     | 0.02% | 0.25(0.04) |
|       |                  | RAMP   | 2.00%     | 0.00% | 0.37(0.13) | 0.50%     | 0.00% | 0.98(0.10) | 0.50%     | 0.00% | 2.27(0.18) |
| 4     | (0, 1.0)         | ADMM   | 98.5%     | 0.12% | 0.24(0.05) | 96.0%     | 0.09% | 1.17(0.11) | 95.5%     | 0.03% | 2.19(0.15) |
|       |                  | IP     | 89.0%     | 0.08% | 0.07(0.02) | 82.5%     | 0.04% | 0.20(0.03) | 73.0%     | 0.02% | 0.38(0.03) |
|       |                  | RAMP   | 3.50%     | 0.00% | 0.25(0.03) | 1.00%     | 0.06% | 1.29(0.24) | 0.00%     | 0.01% | 2.10(0.37) |
|       | (0, 1.5)         | ADMM   | 98.0%     | 0.13% | 0.33(0.06) | 94.0%     | 0.09% | 1.21(0.17) | 91.5%     | 0.04% | 2.15(0.20) |
|       |                  | IP     | 74.0%     | 0.09% | 0.06(0.04) | 62.0%     | 0.05% | 0.18(0.03) | 52.0%     | 0.03% | 0.45(0.04) |
|       |                  | RAMP   | 0.50%     | 0.00% | 0.27(0.05) | 2.00%     | 0.00% | 1.36(0.15) | 0.00%     | 0.00% | 2.32(0.26) |
|       | (0.05, 0.5)      | ADMM   | 98.0%     | 0.15% | 0.31(0.07) | 95.0%     | 0.09% | 1.08(0.15) | 93.0%     | 0.13% | 2.51(0.24) |
|       |                  | IP     | 94.5%     | 0.12% | 0.06(0.01) | 91.5%     | 0.10% | 0.15(0.02) | 89.0%     | 0.06% | 0.33(0.03) |
|       |                  | RAMP   | 9.00%     | 0.00% | 0.28(0.03) | 6.00%     | 0.00% | 1.23(0.12) | 0.00%     | 0.00% | 2.40(0.16) |
|       | (0.15, 1.0)      | ADMM   | 98.5%     | 0.18% | 0.24(0.04) | 94.0%     | 0.13% | 1.19(0.10) | 93.0%     | 0.06% | 2.20(0.27) |
|       |                  | IP     | 90.5%     | 0.13% | 0.08(0.02) | 81.5%     | 0.06% | 0.16(0.03) | 76.5%     | 0.02% | 0.29(0.03) |
|       |                  | RAMP   | 8.00%     | 0.00% | 0.19(0.03) | 3.00%     | 0.00% | 1.05(0.24) | 2.50%     | 0.00% | 2.35(0.23) |

Table 3: Quantitative performance of the interaction detection methods for Model 5 and Model 6. We report the averaged true positive and false positive discovery rates after repeating each simulation setup 200 times with  $n = 100$ . Note that  $\rho = 0$  refers to the case where  $\Sigma$  is the identity matrix.

| Model | $(\rho, \sigma)$ | Method | $d = 100$ |       |            | $d = 200$ |       |            | $d = 300$ |       |            |
|-------|------------------|--------|-----------|-------|------------|-----------|-------|------------|-----------|-------|------------|
|       |                  |        | TPR       | FPR   | Time(s)    | TPR       | FPR   | Time(s)    | TPR       | FPR   | Time(s)    |
| 5     | (0, 0.1)         | ADMM   | 94.7%     | 0.17% | 0.33(0.02) | 92.0%     | 0.09% | 1.14(0.04) | 90.3%     | 0.05% | 2.27(0.13) |
|       |                  | IP     | 80.7%     | 0.14% | 0.05(0.01) | 91.5%     | 0.10% | 0.18(0.02) | 70.3%     | 0.04% | 0.30(0.04) |
|       |                  | RAMP   | 9.33%     | 0.00% | 0.35(0.03) | 9.67%     | 0.00% | 1.19(0.10) | 10.7%     | 0.00% | 2.15(0.18) |
|       | (0, 1.0)         | ADMM   | 92.0%     | 0.12% | 0.29(0.04) | 91.3%     | 0.07% | 1.24(0.14) | 89.7%     | 0.05% | 2.38(0.21) |
|       |                  | IP     | 78.3%     | 0.11% | 0.08(0.01) | 75.0%     | 0.06% | 0.15(0.02) | 68.7%     | 0.03% | 0.26(0.03) |
|       |                  | RAMP   | 8.00%     | 0.00% | 0.24(0.02) | 3.00%     | 0.00% | 1.32(0.13) | 6.67%     | 0.00% | 2.41(0.24) |
|       | (0.2, 0.1)       | ADMM   | 95.3%     | 0.31% | 0.42(0.02) | 94.0%     | 0.10% | 1.20(0.17) | 93.7%     | 0.04% | 2.48(0.26) |
|       |                  | IP     | 81.3%     | 0.28% | 0.11(0.01) | 72.3%     | 0.09% | 0.16(0.02) | 71.0%     | 0.03% | 0.34(0.02) |
|       |                  | RAMP   | 16.0%     | 0.00% | 0.39(0.05) | 11.3%     | 0.00% | 1.37(0.16) | 0.00%     | 0.01% | 2.66(0.47) |
|       | (0.2, 1.0)       | ADMM   | 93.7%     | 0.11% | 0.36(0.04) | 92.7%     | 0.08% | 1.24(0.15) | 91.5%     | 0.04% | 2.53(0.31) |
|       |                  | IP     | 79.7%     | 0.08% | 0.07(0.02) | 77.3%     | 0.06% | 0.19(0.02) | 69.7%     | 0.04% | 0.33(0.03) |
|       |                  | RAMP   | 10.7%     | 0.00% | 0.31(0.04) | 8.33%     | 0.00% | 1.40(0.18) | 7.33%     | 0.00% | 2.45(0.39) |
| 6     | (0, 1.0)         | ADMM   | 98.5%     | 0.23% | 0.24(0.03) | 93.0%     | 0.19% | 1.18(0.13) | 91.5%     | 0.10% | 2.34(0.32) |
|       |                  | IP     | 98.5%     | 0.20% | 0.08(0.02) | 92.5%     | 0.12% | 0.21(0.03) | 92.0%     | 0.08% | 0.31(0.04) |
|       |                  | RAMP   | 99.5%     | 0.00% | 0.32(0.03) | 100.0%    | 0.06% | 1.50(0.22) | 100.0%    | 0.00% | 2.24(0.19) |
|       | (0, 1.5)         | ADMM   | 92.0%     | 0.31% | 0.43(0.08) | 88.5%     | 0.16% | 1.52(0.18) | 82.0%     | 0.04% | 2.40(0.24) |
|       |                  | IP     | 91.5%     | 0.15% | 0.06(0.02) | 86.0%     | 0.11% | 0.17(0.03) | 80.0%     | 0.03% | 0.34(0.05) |
|       |                  | RAMP   | 96.5%     | 0.00% | 0.29(0.04) | 95.5%     | 0.00% | 1.34(0.27) | 94.0%     | 0.00% | 2.39(0.25) |
|       | (0.2, 1.0)       | ADMM   | 96.0%     | 0.23% | 0.36(0.05) | 95.0%     | 0.14% | 1.49(0.20) | 92.5%     | 0.08% | 2.61(0.38) |
|       |                  | IP     | 97.0%     | 0.12% | 0.09(0.02) | 94.5%     | 0.08% | 0.18(0.03) | 91.0%     | 0.05% | 0.29(0.03) |
|       |                  | RAMP   | 100.0%    | 0.00% | 0.41(0.06) | 100.0%    | 0.00% | 1.33(0.24) | 99.0%     | 0.00% | 2.45(0.36) |
|       | (0.4, 1.0)       | ADMM   | 97.5%     | 0.19% | 0.42(0.06) | 95.5%     | 0.10% | 1.38(0.22) | 95.0%     | 0.04% | 2.65(0.31) |
|       |                  | IP     | 98.0%     | 0.11% | 0.07(0.02) | 96.0%     | 0.06% | 0.20(0.03) | 93.5%     | 0.01% | 0.32(0.04) |
|       |                  | RAMP   | 100.0%    | 0.00% | 0.35(0.05) | 100.0%    | 0.00% | 1.39(0.30) | 100.0%    | 0.00% | 2.50(0.39) |



Table 4: Quantitative performance of the interaction detection methods for Model 7 and Model 8. We report the averaged true positive and false positive discovery rates after repeating each simulation setup 200 times with  $n = 100$ . Note that  $\rho = 0$  refers to the case where  $\Sigma$  is the identity matrix.

| Model | $(\rho, \sigma)$ | Method | $d = 100$ |       |            | $d = 200$ |       |            | $d = 300$ |       |            |
|-------|------------------|--------|-----------|-------|------------|-----------|-------|------------|-----------|-------|------------|
|       |                  |        | TPR       | FPR   | Time(s)    | TPR       | FPR   | Time(s)    | TPR       | FPR   | Time(s)    |
| 7     | (0, 1.0)         | ADMM   | 100.0%    | 0.29% | 0.31(0.04) | 100.0%    | 0.19% | 1.22(0.12) | 97.5%     | 0.08% | 2.38(0.35) |
|       |                  | IP     | 98.5%     | 0.27% | 0.08(0.02) | 95.5%     | 0.13% | 0.17(0.02) | 90.0%     | 0.04% | 0.36(0.04) |
|       |                  | RAMP   | 1.00%     | 0.00% | 0.36(0.06) | 2.50%     | 0.00% | 1.31(0.11) | 2.00%     | 0.00% | 2.50(0.42) |
|       | (0, 1.5)         | ADMM   | 97.0%     | 0.34% | 0.41(0.06) | 95.5%     | 0.17% | 1.30(0.21) | 94.5%     | 0.07% | 2.43(0.28) |
|       |                  | IP     | 94.0%     | 0.29% | 0.08(0.02) | 92.0%     | 0.16% | 0.21(0.03) | 89.5%     | 0.06% | 0.37(0.06) |
|       |                  | RAMP   | 1.50%     | 0.00% | 0.37(0.04) | 2.00%     | 0.00% | 1.43(0.32) | 1.00%     | 0.00% | 2.62(0.38) |
|       | (0.2, 1.0)       | ADMM   | 100.0%    | 0.23% | 0.39(0.03) | 99.5%     | 0.13% | 1.40(0.18) | 97.0%     | 0.06% | 2.44(0.31) |
|       |                  | IP     | 97.0%     | 0.24% | 0.07(0.02) | 94.0%     | 0.12% | 0.19(0.03) | 90.5%     | 0.05% | 0.32(0.04) |
|       |                  | RAMP   | 1.50%     | 0.00% | 0.40(0.03) | 2.50%     | 0.00% | 1.35(0.25) | 2.00%     | 0.00% | 2.23(0.40) |
|       | (0.2, 1.5)       | ADMM   | 96.5%     | 0.25% | 0.33(0.04) | 95.0%     | 0.17% | 1.24(0.18) | 94.0%     | 0.06% | 2.48(0.29) |
|       |                  | IP     | 93.0%     | 0.27% | 0.09(0.02) | 91.5%     | 0.13% | 0.16(0.03) | 88.0%     | 0.04% | 0.31(0.04) |
|       |                  | RAMP   | 1.00%     | 0.00% | 0.38(0.06) | 1.50%     | 0.00% | 1.32(0.24) | 1.00%     | 0.00% | 2.35(0.33) |
| 8     | (0, 1.0)         | ADMM   | 97.0%     | 0.31% | 0.27(0.03) | 96.5%     | 0.18% | 1.40(0.18) | 94.5%     | 0.05% | 2.51(0.32) |
|       |                  | IP     | 75.5%     | 0.29% | 0.06(0.02) | 60.0%     | 0.17% | 0.18(0.03) | 57.5%     | 0.08% | 0.39(0.05) |
|       |                  | RAMP   | 0.00%     | 0.00% | 0.25(0.04) | 1.00%     | 0.00% | 1.37(0.29) | 0.00%     | 0.00% | 2.40(0.34) |
|       | (0, 1.25)        | ADMM   | 95.0%     | 0.28% | 0.40(0.06) | 93.5%     | 0.16% | 1.28(0.26) | 91.0%     | 0.08% | 2.48(0.41) |
|       |                  | IP     | 48.5%     | 0.23% | 0.07(0.02) | 39.0%     | 0.13% | 0.17(0.02) | 35.0%     | 0.06% | 0.37(0.06) |
|       |                  | RAMP   | 1.00%     | 0.00% | 0.36(0.05) | 2.50%     | 0.00% | 1.29(0.17) | 1.00%     | 0.00% | 2.36(0.38) |
|       | (0.2, 1.0)       | ADMM   | 95.5%     | 0.35% | 0.32(0.04) | 92.5%     | 0.23% | 1.19(0.15) | 91.5%     | 0.12% | 2.58(0.31) |
|       |                  | IP     | 67.5%     | 0.28% | 0.08(0.02) | 54.0%     | 0.20% | 0.15(0.03) | 52.5%     | 0.13% | 0.33(0.05) |
|       |                  | RAMP   | 1.00%     | 0.00% | 0.36(0.06) | 1.00%     | 0.00% | 1.41(0.16) | 1.50%     | 0.00% | 2.43(0.41) |
|       | (0.2, 1.25)      | ADMM   | 94.0%     | 0.29% | 0.42(0.06) | 92.0%     | 0.17% | 1.38(0.19) | 90.5%     | 0.09% | 2.43(0.44) |
|       |                  | IP     | 64.0%     | 0.25% | 0.09(0.02) | 52.0%     | 0.03% | 0.20(0.04) | 50.5%     | 0.08% | 0.33(0.05) |
|       |                  | RAMP   | 1.00%     | 0.00% | 0.30(0.04) | 2.50%     | 0.00% | 1.24(0.26) | 2.00%     | 0.00% | 2.39(0.35) |

Table 5: Quantitative performance of the interaction detection methods for Model 9. We report the averaged true positive and false positive discovery rates after repeating each simulation setup 200 times with  $n = 100$ . Note that  $\rho = 0$  refers to the case where  $\Sigma$  is the identity matrix.

| Model | $(\rho, \sigma)$ | Method | $d = 100$ |       |            | $d = 200$ |       |            | $d = 300$ |       |            |
|-------|------------------|--------|-----------|-------|------------|-----------|-------|------------|-----------|-------|------------|
|       |                  |        | TPR       | FPR   | Time(s)    | TPR       | FPR   | Time(s)    | TPR       | FPR   | Time(s)    |
| 9     | (0, 1.0)         | ADMM   | 100.0%    | 0.33% | 0.28(0.05) | 100.0%    | 0.23% | 1.32(0.20) | 98.4%     | 0.10% | 2.46(0.33) |
|       |                  | IP     | 99.2%     | 0.26% | 0.07(0.01) | 96.6%     | 0.21% | 0.15(0.02) | 91.4%     | 0.06% | 0.31(0.04) |
|       |                  | RAMP   | 0.84%     | 0.02% | 0.30(0.04) | 2.01%     | 0.01% | 1.41(0.23) | 1.12%     | 0.00% | 2.53(0.36) |
|       | (0, 1.5)         | ADMM   | 98.4%     | 0.25% | 0.37(0.04) | 96.7%     | 0.23% | 1.44(0.26) | 95.2%     | 0.09% | 2.60(0.37) |
|       |                  | IP     | 96.5%     | 0.22% | 0.07(0.01) | 95.3%     | 0.13% | 0.25(0.03) | 93.8%     | 0.08% | 0.30(0.04) |
|       |                  | RAMP   | 1.39%     | 0.01% | 0.34(0.05) | 1.69%     | 0.03% | 1.58(0.41) | 0.96%     | 0.00% | 2.43(0.49) |
|       | (0.2, 1.0)       | ADMM   | 99.9%     | 0.18% | 0.35(0.04) | 99.7%     | 0.09% | 1.34(0.21) | 98.2%     | 0.03% | 2.61(0.40) |
|       |                  | IP     | 98.6%     | 0.13% | 0.06(0.02) | 98.3%     | 0.05% | 0.16(0.02) | 95.0%     | 0.02% | 0.38(0.04) |
|       |                  | RAMP   | 1.08%     | 0.00% | 0.46(0.04) | 1.35%     | 0.00% | 1.49(0.31) | 1.29%     | 0.00% | 2.51(0.35) |
|       | (0.2, 1.5)       | ADMM   | 98.9%     | 0.21% | 0.42(0.05) | 96.9%     | 0.14% | 1.40(0.22) | 95.8%     | 0.07% | 2.59(0.43) |
|       |                  | IP     | 96.1%     | 0.24% | 0.08(0.01) | 94.2%     | 0.10% | 0.13(0.04) | 90.3%     | 0.03% | 0.27(0.05) |
|       |                  | RAMP   | 0.82%     | 0.00% | 0.49(0.05) | 0.79%     | 0.00% | 1.44(0.30) | 0.66%     | 0.00% | 2.43(0.39) |

Table 6: The means and standard errors (in parentheses) of testing errors and median model sizes in GPL96 data analysis for breast cancer.

| Method | Classification Error | Median Model Size |
|--------|----------------------|-------------------|
| ADMM   | 6.13% (0.24%)        | 97                |
| IP     | 7.38% (0.22%)        | 85                |
| RAMP   | 8.15% (0.35%)        | 68                |

Table 7: The means and standard errors (in parentheses) of testing errors and median model sizes in GPL96 data analysis.

| ADMM                   |           | IP                           |           | RAMP                  |           |
|------------------------|-----------|------------------------------|-----------|-----------------------|-----------|
| Interaction            | Frequency | Interaction                  | Frequency | Interaction           | Frequency |
| BRCA1 $\times$ LMO3    | 75        | c-Myc $\times$ KLK3          | 69        | ESR1 $\times$ IRF4    | 64        |
| BRCA1 $\times$ PALB2   | 71        | FOXA1 $\times$ NF $\kappa$ B | 62        | SERINC5 $\times$ CHPF | 61        |
| TP53 $\times$ RXRA     | 64        | JUND $\times$ GSTP1          | 56        | EPCAM $\times$ RBP1   | 55        |
| BRCA2 $\times$ HNF4A   | 59        | BRCA1 $\times$ ALCAM         | 52        | MAF $\times$ LMOD1    | 51        |
| PDLIM5 $\times$ SETDB1 | 58        | DPT $\times$ SETDB1          | 51        | BRCA2 $\times$ TARP   | 51        |

Table 8: The means and standard errors (in parentheses) of testing errors and median model sizes in GPL96 data analysis for Ewing Sarcomar.

| Method | Classification Error | Median Model Size |
|--------|----------------------|-------------------|
| ADMM   | 5.58% (1.81%)        | 79                |
| IP     | 8.90% (1.69%)        | 71                |
| RAMP   | 19.72% (4.55%)       | 95                |

## Acknowledgement

We are grateful to the Editor, Action Editor and two reviewers for their helpful comments, which lead to a significant improvement of the earlier version of this paper. Tang is partially supported by NSF grant IIS-1546087 and ES-1533956. Fang is partially supported by NSF grant DMS 1820702.

## References

- P. J. Bickel and E. Levina. Regularized estimation of large covariance matrices. *Annals of Statistics*, 36:199–227, 2008.
- Jacob Bien, Jonathan Taylor, and Robert Tibshirani. A lasso for hierarchical interactions. *The Annals of Statistics*, 41(3):1111–1141, jun 2013. doi: 10.1214/13-aos1096.
- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- P. Bühlmann and S. van de Geer. *Statistics for High-dimensional Data: Methods, Theory and Applications*. New York: Springer, 2011.
- Nam Hee Choi, William Li, and Ji Zhu. Variable selection with the strong heredity constraint and its oracle property. *Journal of the American Statistical Association*, 105(489): 354–364, mar 2010. doi: 10.1198/jasa.2010.tm08281.
- R. Dennis Cook. Principal hessian directions revisited. *Journal of the American Statistical Association*, 93(441):84–94, mar 1998. doi: 10.1080/01621459.1998.10474090.
- Heather J Cordell. Detecting gene-gene interactions that underlie human diseases. *Nature Reviews Genetics*, 10(6):392–404, jun 2009. doi: 10.1038/nrg2579.
- Serena Dato, Mette Soerensen, Francesco De Rango, Giuseppina Rose, Kaare Christensen, Lene Christiansen, and Giuseppe Passarino. The genetic component of human longevity: New insights from the analysis of pathway-based SNP-SNP interactions. *Aging Cell*, 17(3):e12755, mar 2018. doi: 10.1111/accel.12755.
- Chao Ding, Defeng Sun, and Kim-Chuan Toh. An introduction to a class of matrix cone programming. *Mathematical Programming*, 144(1-2):141–179, 2014.
- Jonathan Eckstein and W Yao. Augmented lagrangian and alternating direction methods for convex optimization: A tutorial and some illustrative computational results. *RUTCOR Research Reports*, 32:3, 2012.
- J. Fan and J. Lv. A selective overview of variable selection in high dimensional feature space. *Statistica Sinica*, 20:101–148, 2010.
- Yingying Fan, Yinfei Kong, Daoji Li, and Jinchi Lv. Interaction pursuit with feature screening and selection. *Manuscript. arXiv:1605.08933.*, 2016.

- Ethan X Fang, Bingsheng He, Han Liu, and Xiaoming Yuan. Generalized alternating direction method of multipliers: new theoretical insights and applications. *Mathematical Programming Computation*, 7(2):149–187, 2015.
- Ethan X Fang, Min-Dian Li, Michael I Jordan, and Han Liu. Mining massive amounts of genomic data: a semiparametric topic modeling approach. *Journal of the American Statistical Association*, 112(519):921–932, 2017.
- Roland Glowinski and A Marroco. Sur l’approximation, par éléments finis d’ordre un, et la résolution, par pénalisation-dualité d’une classe de problèmes de dirichlet non linéaires. *Revue française d’automatique, informatique, recherche opérationnelle. Analyse numérique*, 9(R2):41–76, 1975.
- Ning Hao and Hao Helen Zhang. Interaction screening for ultrahigh-dimensional data. *Journal of the American Statistical Association*, 109(507):1285–1301, jul 2014. doi: 10.1080/01621459.2014.881741.
- Ning Hao, Yang Feng, and Hao Helen Zhang. Model selection for high-dimensional quadratic regression via regularization. *Journal of the American Statistical Association*, 113:615–625, 2018.
- Asad Haris, Daniela Witten, and Noah Simon. Convex modeling of interactions with strong heredity. *Journal of Computational and Graphical Statistics*, 25(4):981–1004, oct 2016. doi: 10.1080/10618600.2015.1067217.
- Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical Learning with Sparsity: The Lasso and Generalizations*. Taylor & Francis Inc, 2015. ISBN 1498712169. URL [http://www.ebook.de/de/product/23642172/trevor\\_hastie\\_robert\\_tibshirani\\_martin\\_wainwright\\_statistical\\_learning\\_with\\_sparsity.html](http://www.ebook.de/de/product/23642172/trevor_hastie_robert_tibshirani_martin_wainwright_statistical_learning_with_sparsity.html).
- J. J. Jaccard and R. Turrisi. *Interaction Effects in Multiple Regression*. SAGE, 2003. ISBN 0761927425. URL [https://www.ebook.de/de/product/4324119/james\\_j\\_jaccard\\_robert\\_turrisi\\_interaction\\_effects\\_in\\_multiple\\_regression.html](https://www.ebook.de/de/product/4324119/james_j_jaccard_robert_turrisi_interaction_effects_in_multiple_regression.html).
- Binyan Jiang, Xiangyu Wang, and Chenlei Leng. A direct approach for sparse quadratic discriminant analysis. *Journal of Machine Learning Research*, 19:1–37, 2018.
- Bo Jiang and Jun S. Liu. Variable selection for general index models via sliced inverse regression. *The Annals of Statistics*, 42(5):1751–1786, oct 2014. doi: 10.1214/14-aos1233.
- Mary-Claire King, Joan H Marks, Jessica B Mandell, et al. Breast and ovarian cancer risks due to inherited mutations in BRCA1 and BRCA2. *Science*, 302(5645):643–646, 2003.
- Yinfei Kong, Daoji Li, Yingying Fan, and Jinchi Lv. Interaction pursuit in high-dimensional multi-response regression via distance correlation. *The Annals of Statistics*, 45(2):897–922, apr 2017. doi: 10.1214/16-aos1474.
- Ker-Chau Li. Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414):316–327, Jun 1991. ISSN 1537-274X. doi: 10.1080/01621459.1991.10475035. URL <http://dx.doi.org/10.1080/01621459.1991.10475035>.

- Ker-Chau Li. On principal hessian directions for data visualization and dimension reduction: Another application of stein’s lemma. *Journal of the American Statistical Association*, 87(420):1025–1039, dec 1992. doi: 10.1080/01621459.1992.10476258.
- Yang Li and Jun S. Liu. Robust variable and interaction selection for logistic regression and general index models. *Journal of the American Statistical Association*, pages 1–16, nov 2017. doi: 10.1080/01621459.2017.1401541.
- Matthew N McCall, Benjamin M Bolstad, and Rafael A Irizarry. Frozen robust multiarray analysis (fRMA). *Biostatistics*, 11(2):242–253, 2010.
- Sahand N. Negahban, Pradeep Ravikumar, Martin J. Wainwright, and Bin Yu. A unified framework for high-dimensional analysis of  $m$ -estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, nov 2012. doi: 10.1214/12-sts400.
- J. A. Nelder. A reformulation of linear models. *Journal of the Royal Statistical Society. Series A (General)*, 140(1):48, 1977. doi: 10.2307/2344517.
- Garvesh Raskutti, Martin J Wainwright, and Bin Yu. Restricted eigenvalue properties for correlated gaussian designs. *Journal of Machine Learning Research*, 11(Aug):2241–2259, 2010.
- Garvesh Raskutti, Martin J. Wainwright, and Bin Yu. Minimax rates of estimation for high-dimensional linear regression over  $\ell_q$ -balls. *IEEE Transactions on Information Theory*, 57(10):6976–6994, oct 2011. doi: 10.1109/tit.2011.2165799.
- Marylyn D. Ritchie and Kristel Van Steen. The search for gene-gene interactions in genome-wide association studies: challenges in abundance of methods, practical considerations, and biological interpretation. *Annals of Translational Medicine*, 6(8):157–157, apr 2018. doi: 10.21037/atm.2018.04.05.
- Mark Rudelson and Shuheng Zhou. Reconstruction from anisotropic random measurements. *Manuscript. arXiv:1106.1151.*, 2011.
- Rajen D Shah. Modelling interactions in high-dimensional data with backtracking. *Journal of Machine Learning Research*, 17(1):7225–7255, 2016.
- Yiyuan She, Zhifeng Wang, and He Jiang. Group regularized estimation under structural hierarchy. *Journal of the American Statistical Association*, 113(521):445–454, dec 2018. doi: 10.1080/01621459.2016.1260470.
- Charles M. Stein. Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics*, 9(6):1135–1151, nov 1981. doi: 10.1214/aos/1176345632.
- Defeng Sun, Kim-Chuan Toh, and Liuqin Yang. A convergent 3-block semiproximal alternating direction method of multipliers for conic programming with 4-type constraints. *SIAM Journal on Optimization*, 25(2):882–915, 2015.
- George Wu, Jason T Yustein, Matthew N McCall, Michael Zilliox, Rafael A Irizarry, Karen Zeller, Chi V Dang, and Hongkai Ji. ChIP-PED enhances the analysis of ChIP-seq and ChIP-chip data. *Bioinformatics*, 2013.

Ming Yuan, V. Roshan Joseph, and Hui Zou. Structured variable selection and estimation. *The Annals of Applied Statistics*, 3(4):1738–1757, dec 2009. doi: 10.1214/09-aos254.

Peng Zhao, Guilherme Rocha, and Bin Yu. The composite absolute penalties family for grouped and hierarchical variable selection. *The Annals of Statistics*, 37(6A):3468–3497, dec 2009. doi: 10.1214/07-aos584.

## Appendix

### Proof of Proposition 3

Since  $Z = \Sigma^{-1}X$ , we have  $\text{Var}(Z) = \Sigma^{-1}$ , and  $\text{Cov}(Z, X) = I$ . Under model (2),

$$Y - \mathbb{E}(Y) = \sum_{k=1}^p \sum_{\ell=k}^p \beta_{k\ell} \{X_k X_\ell - \mathbb{E}(X_k X_\ell)\} + \sum_{k=1}^p \theta_k X_k + \epsilon. \quad (10)$$

Then, we have from C1 of Condition 2

$$\Psi = \mathbb{E}\{Y - \mathbb{E}(Y)ZZ^T\} = \mathbb{E}\left[\sum_{k=1}^p \sum_{\ell=k}^p \beta_{k\ell} \{X_k X_\ell - \mathbb{E}(X_k X_\ell)\} ZZ^T\right].$$

This implies  $\psi_{ij} = \sum_{k=1}^p \sum_{\ell=k}^p \beta_{k\ell} \text{Cov}(X_k X_\ell, Z_i Z_j)$ . Then Proposition 3 follows from C2 of Condition 2.

### Proof of Lemma 4

We prove the main results following the framework of penalized M-estimator as in Negahban et al. (2012) with  $\ell_1$  penalty function.

Let  $\mathcal{L}(\Psi) = \text{tr}(\Psi^T S \Psi S) - 2\text{tr}(\Psi Q)$ , and denote by  $\Psi^*$  the truth of  $\Psi$ . By the definition

$$\mathcal{L}(\widehat{\Psi}) + \lambda \|\widehat{\Psi}\|_1 \leq \mathcal{L}(\Psi^*) + \lambda \|\widehat{\Psi}^*\|_1.$$

For ease in presentation, we take  $\widehat{\Psi}$  to be symmetric. Let  $\widehat{\Delta} = \widehat{\Psi} - \Psi^*$ , then we have

$$\mathcal{L}(\Psi^* + \widehat{\Delta}) - \mathcal{L}(\Psi^*) + \lambda \|\widehat{\Psi}\|_1 \leq \lambda \|\widehat{\Psi}^*\|_1. \quad (11)$$

We note that

$$\text{tr}\{(\Psi^* + \widehat{\Delta})^T S (\Psi^* + \widehat{\Delta}) S\} - \text{tr}(\Psi^* S \Psi^* S) = \text{tr}(\widehat{\Delta} S \widehat{\Delta} S) + 2\text{tr}\{\widehat{\Delta} S \Psi^* S\}.$$

Then (11) implies that

$$\text{tr}(\widehat{\Delta} S \widehat{\Delta} S) + \lambda \|\widehat{\Psi}\|_1 \leq -2\text{tr}\{\widehat{\Delta} (S \Psi^* S - Q)\} + \lambda \|\Psi^*\|_1. \quad (12)$$

We observe that with large probability  $2\|S \Psi^* S - Q\|_\infty < \lambda/2$ . Additionally, let  $\mathcal{S} = \text{supp}(\Psi^*)$ ,

$$\begin{aligned} \|\widehat{\Psi}\|_1 &= \|\widehat{\Psi}_U\|_1 + \|\widehat{\Psi}_{U^c}\|_1 = \|\widehat{\Psi}_U + \Psi_{\mathcal{U}}^* - \Psi_{\mathcal{U}}^*\|_1 + \|\widehat{\Psi}_{U^c}\|_1 \\ &\geq \|\Psi_{\mathcal{U}}^*\|_1 - \|\widehat{\Psi}_U - \Psi_{\mathcal{U}}^*\|_1 + \|\widehat{\Psi}_{\mathcal{S}^c}\|_1, \\ \|\widehat{\Delta}\|_1 &= \|\widehat{\Psi}_U - \Psi_{\mathcal{U}}^*\|_1 + \|\widehat{\Psi}_{U^c}\|_1, \text{ and } \|\Psi^*\|_1 = \|\Psi_U\|_1. \end{aligned}$$

These together with (12) imply that

$$2\text{tr}(\widehat{\Delta} S \widehat{\Delta} S) + \lambda \|\widehat{\Psi}_{U^c}\|_1 \leq 3\lambda \|\widehat{\Psi}_U - \Psi_{\mathcal{U}}^*\|_1. \quad (13)$$

Since  $\text{tr}(\widehat{\Delta} S \widehat{\Delta} S) \geq 0$ , then with large probability

$$\|\widehat{\Psi}_{U^c}\|_1 \leq 3\|\widehat{\Psi}_U - \Psi_{\mathcal{U}}^*\|_1.$$

This completes the proof of Lemma 4.

**Proof of Lemma 5**

From the proof of Lemma 4, Lemma 5 follows straightforwardly by noting that

$$\begin{aligned} 2\text{tr}(\widehat{\Delta}S\widehat{\Delta}S) + \lambda\|\widehat{\Delta}\|_1 &= 2\text{tr}(\widehat{\Delta}S\widehat{\Delta}S) + \lambda\|\widehat{\Psi}_{\mathcal{U}} - \Psi_{\mathcal{U}}^*\|_1 + \lambda\|\widehat{\Psi}_{\mathcal{U}^c}\|_1 \\ &\leq 4\lambda\|\widehat{\Psi}_{\mathcal{U}} - \Psi_{\mathcal{U}}^*\|. \end{aligned} \quad (14)$$

**Proof of Theorem 7**

We further note that

$$\text{tr}(\widehat{\Delta}S\widehat{\Delta}S) = \text{vec}^T(\widehat{\Delta})(S \otimes S)\text{vec}(\widehat{\Delta}).$$

Then the oracle inequality (14) and the restricted eigenvalue condition (9) imply that

$$\begin{aligned} 2\text{tr}(\widehat{\Delta}S\widehat{\Delta}S) + \lambda\|\widehat{\Psi}_{\mathcal{U}} - \Psi_{\mathcal{U}}^*\|_1 + \lambda\|\widehat{\Psi}_{\mathcal{U}^c}\|_1 &\leq 4\lambda\|\widehat{\Psi}_{\mathcal{U}} - \Psi_{\mathcal{U}}^*\|_1 \\ &\leq 4\lambda|\mathcal{U}|^{1/2}\{\text{vec}^T(\Delta)(S \otimes S)\text{vec}(\Delta)\}^{1/2}/\gamma \\ &\leq \text{vec}^T(\Delta)(S \otimes S)\text{vec}(\Delta) + 4\lambda^2|\mathcal{U}|/\gamma^2 \\ &= \text{tr}(\widehat{\Delta}S\widehat{\Delta}S) + 4\lambda^2|\mathcal{U}|/\gamma^2. \end{aligned}$$

That is

$$\text{tr}(\widehat{\Delta}S\widehat{\Delta}S) + \lambda\|\widehat{\Psi} - \Psi^*\|_1 \leq 4\lambda^2|\mathcal{U}|/\gamma^2 \quad (15)$$

This completes the proof of Theorem 7.

**5.3. Proof of Theorem 9**

Let  $\Delta_{\Sigma} = S - \Sigma_X$ . By its definition

$$\begin{aligned} S\Psi^*S - Q &= (\Sigma_X + \Delta_{\Sigma})\Psi^*(\Sigma_X + \Delta_{\Sigma}) - Q \\ &= \Sigma_X\Psi^*\Sigma_X - Q + \Delta_{\Sigma}\Psi^*\Sigma_X + \Sigma_X\Psi^*\Delta_{\Sigma} + \Delta_{\Sigma}\Psi^*\Delta_{\Sigma}. \end{aligned}$$

Then by Condition 8 and Lemma A.3 of Bickel and Levina (2008),  $\|\Delta_{\Sigma}\|_{\infty} = \mathcal{O}_p(\sqrt{\log p/n})$ . Additionally, since  $\Sigma_X\Psi^*\Sigma_X = \Sigma_{YXX}$ , and  $\|\Sigma_{YXX} - Q\|_{\infty} = \mathcal{O}_p(\sqrt{\log p/n})$  follows from Condition 8, similar to that of Lemma A.3 of Bickel and Levina (2008). Then Theorem 9 follows by establishing that  $\|S\Psi^*S - Q\|_{\infty} = \mathcal{O}_p(\sqrt{\log p/n})$  and applying the triangular inequality and the condition that  $\|\Psi^*\|_{\infty} = \mathcal{O}(\sqrt{n/\log p})$ ,  $\|\Psi^*\Sigma_X\|_{\infty} = \mathcal{O}(1)$ .

**5.4. Proof of Theorem 11**

The first order conditon of (5) implies that

$$0 = S\widehat{\Psi}S - Q + \lambda Z$$

where  $Z = \text{sign}(\widehat{\Psi})$ . Let  $\eta = \text{vec}(\widehat{\Psi})$ ,  $\eta^* = \text{vec}(\Psi^*)$ ,  $\xi = \text{vec}(Q)$ , and  $\tau = \text{vec}(Z)$ . By writing in vector form, we have

$$\begin{pmatrix} \Lambda_{11} & \Lambda_{12} \\ \Lambda_{21} & \Lambda_{22} \end{pmatrix} \begin{pmatrix} \eta_1 - \eta_1^* \\ \eta_2 \end{pmatrix} - \begin{pmatrix} \xi_1 - \Lambda_{11}\eta_1^* \\ \xi_2 - \Lambda_{21}\eta_1^* \end{pmatrix} + \lambda \begin{pmatrix} \tau_1 \\ \tau_2 \end{pmatrix} = 0$$



where  $(\eta_1^T, \eta_2^T)^T$  is the decomposition of  $\text{vec}(\widehat{\Psi})$  corresponding to the set  $\mathcal{U}$ . Then to ensure  $\eta_2 = 0$ , we must have

$$\begin{aligned}\eta_1 - \eta_1^* &= \Lambda_{11}^{-1}(\xi_1 - \Lambda_{11}\eta_1^*) - \lambda\Lambda_{11}^{-1}\tau_1, \\ \tau_2 &= \Lambda_{21}\Lambda_{11}^{-1}Z_1 + \lambda^{-1}(\xi_2 - \Lambda_{21}\Lambda_{11}^{-1}\xi_1)\end{aligned}$$

where  $Z_1$  is the component of  $Z$  corresponding to the blocking. Then the irrerepresentable condition (8) implies that  $\eta_2 = 0$  with probability tending to 1, and the minimal signal strength condition  $|\min(\Psi_{\mathcal{U}}^*)| > \|\lambda\Lambda_{11}^{-1}\|_{\infty}$  ensures that  $\eta_2$  is bounded around from 0 with probability tending to 1.