

Sharp Threshold for Multivariate Multi-Response Linear Regression via Block Regularized Lasso ^{1 2}

Weiguang Wang, Yingbin Liang, Eric P. Xing ³

Abstract

The multivariate multi-response (MVMR) linear regression problem is investigated, in which design matrices are Gaussian with covariance matrices $\Sigma^{(1:K)} = (\Sigma^{(1)}, \dots, \Sigma^{(K)})$ for K linear regressions. The support union of K p -dimensional regression vectors (collected as columns of matrix B^*) is recovered using l_1/l_2 -regularized Lasso. Sufficient and necessary conditions on sample complexity are characterized as a sharp threshold to guarantee successful recovery of the support union. This model has been previously studied via l_1/l_∞ -regularized Lasso in [1] and via $l_1/l_1 + l_1/l_\infty$ -regularized Lasso in [2], in which sharp threshold on sample complexity is characterized only for $K = 2$ and under special conditions. In this work, using l_1/l_2 -regularized Lasso, sharp threshold on sample complexity is characterized under standard regularization conditions. Namely, if $n > c_{p1}\psi(B^*, \Sigma^{(1:K)}) \log(p - s)$ where c_{p1} is a constant, and s is the size of the support set, then l_1/l_2 -regularized Lasso correctly recovers the support union; and if $n < c_{p2}\psi(B^*, \Sigma^{(1:K)}) \log(p - s)$ where c_{p2} is a constant, then l_1/l_2 -regularized Lasso fails to recover the support union. In particular, the function $\psi(B^*, \Sigma^{(1:K)})$ captures the impact of the sparsity of K regression vectors and the statistical properties of the design matrices on the threshold on sample complexity. Therefore, such threshold function also demonstrates the advantages of joint support union recovery using multi-task Lasso over individual support recovery using single-task Lasso.

1 Introduction

Linear regression is a simple but practically very useful statistical model, in which an n -sample response vector \vec{Y} can be modeled as

$$\vec{Y} = X\vec{\beta} + \vec{W}$$

¹The material in this paper was presented in part at the International Conference on Artificial Intelligence and Statistics (AISTATS), Scottsdale, AZ, USA, April 2013.

²The work of W. Wang and Y. Liang was supported by NSF CAREER Award CCF-10-26565 and NSF CCF-10-26566. The work of Eric P. Xing was supported by NIH1R01GM087694, FA9550010247, and NIH1R01GM093156.

³Weiguang Wang and Yingbin Liang are with the Department of Electrical Engineering and Computer Science, Syracuse University, Syracuse, NY 13244 USA (email: {wwang23,yliang06}@syr.edu). Eric P. Xing is with the Department of Machine Learning, Carnegie Mellon University, Pittsburgh, PA 15213 USA (email: epxing@cs.cmu.edu).

where $X \in \mathbb{R}^{n \times p}$ is the design matrix containing n samples of feature vectors, $\vec{\beta} = (\beta_1, \dots, \beta_p) \in \mathbb{R}^p$ contains regression coefficients, and $\vec{W} \in \mathbb{R}^n$ is the noise vector. The goal is to find the regression coefficients $\vec{\beta}$ such that the linear relationship is as accurate as possible with regard to a certain performance criterion. The problem is more interesting in high dimensional regime with a sparse regression vector, in which the sample size n can be much smaller than the dimension p of the regression vector.

In order to estimate the sparse regression vector, it is natural to construct an optimization problem with an l_0 -constraint on $\vec{\beta}$, i.e., the number of nonzero components of $\vec{\beta}$. However, such an optimization problem is nonconvex and in general very difficult to solve in an efficient manner as commented in [3]. More recently, the convex relaxation (referred to as Lasso) has been studied with an l_1 -constraint on $\vec{\beta}$ based on the idea in some seminal works ([4–6]). More specifically, the regression problem can be formulated as:

$$\min_{\vec{\beta} \in \mathbb{R}^p} \frac{1}{n} \|\vec{Y} - X \vec{\beta}\|_{l_2}^2 + \lambda_n \|\vec{\beta}\|_{l_1}.$$

The l_1 -regularized estimator has been proved in [7] to have similar behavior to Dantzig Selector, which was proposed in [8]. Various efficient algorithms have been developed to solve the above convex problem efficiently (see a review monograph [9]), although the objective function is not differentiable everywhere due to l_1 -regularization. Moreover, the l_1 -regularization is critical to force the minimizer to have sparse components as shown in [4–6].

A vast amount of recent work has studied the high dimensional linear regression problem via l_1 -regularized Lasso under various assumptions. For example, the studies [5, 10–14] investigated the noiseless scenario and showed that recovery of true coefficients could be guaranteed with certain conditions on design matrices and sparsity. A number of studies focused on using l_1 -regularization to achieve sparsity recovery for noisy scenarios. Some work (e.g., [15–17]) focused on the problem with deterministic design matrices, whereas other work (e.g., [18, 19]) studied the problem with random design matrices. The work [20] investigated linear regression model via trace norm. [21] and [22] studied linear regression model using a fusion penalty (known as the total variational penalty).

Generalized from the l_1 -regularized linear regression problem which aims at selecting variables individually, group Lasso is applied to regression vector $\vec{\beta}$ in the linear regression model to select grouped variables (e.g., [23, 24]). The work [25] and [26] applied group Lasso for studying empirical risk minimization problems. The work [27] studied the least square optimization problem with group Lasso.

This line of research is further generalized to block-regularization for high-dimensional multi-response (i.e., multi-task) linear regression problem, (see, e.g., [28, 29] and references therein). For a multi-task regression problem, we have the following model:

$$Y = XB^* + W \tag{1}$$

where $Y \in \mathbb{R}^{n \times K}$ of which each column corresponds to the output of one task, $X \in \mathbb{R}^{n \times p}$

is the design matrix, the regression matrix $B^* \in \mathbb{R}^{p \times K}$ has each column corresponding to the regression vector for one task, and $W \in \mathbb{R}^{n \times K}$ has each column corresponding to the noise vector of one task. For each column $\vec{Y}^{(k)}$ of the matrix Y , it is clear that $\vec{Y}^{(k)} = X \vec{\beta}^{*(k)} + \vec{W}^{(k)}$, where $\vec{\beta}^{*(k)}$ and $\vec{W}^{(k)}$ are the corresponding columns in B^* and W . Then each column is a single-task linear regression problem and can be solved individually. However, the K individual problems (i.e., tasks) can also be coupled together via a block regularized Lasso and solved jointly in one problem.

Various types of block regularization have been proposed and studied. In the work [29], the l_1/l_2 -regularization was adopted to recover the support union of the regression vectors. More specifically, the following problem was studied

$$\min_{B \in \mathbb{R}^{p \times K}} \frac{1}{2n} \| \| Y - XB \| \|_F^2 + \lambda_n \| B \|_{l_1/l_2}, \quad (2)$$

where $\| \cdot \|_{l_a/l_b}$ is defined in (7) in section 2.1. Sufficient and necessary conditions for correct recovery of the support union (i.e., the union of the supports of all columns of B^*) have been characterized. Block regularized Lasso (as well as group Lasso) has also been applied to study various other models. For example, the l_1/l_q -regularized Lasso was adopted for learning structured linear regression model in [30]. The l_1/l_∞ -regularized Lasso was used to investigate a multi-response regression model in [31]. The l_1/l_2 -regularization was used for studying empirical risk minimization problems in [32], and to study multi-task feature problems in [33]. The l_1/l_q -regularized Lasso was adopted to analyze normal means model in [34]. Blockwise sparse regression was used to study a general loss function in [35].

In the multi-response linear regression problem given in (1), the design matrix is identical for all tasks, i.e., X is the same for all column vectors of Y and B^* . However, in many applications, it is often the case that different output variables may depend on design variables that are different or distributed differently. Thus, the resulting model includes K linear regression models with different design matrices and is given by:

$$\vec{Y}^{(k)} = X^{(k)} \vec{\beta}^{*(k)} + \vec{W}^{(k)} \quad (3)$$

for $k = 1, \dots, K$, where $\vec{Y}^{(k)} \in \mathbb{R}^n$, $X^{(k)} \in \mathbb{R}^{n \times p}$, $\vec{\beta}^{*(k)} \in \mathbb{R}^p$, and $\vec{W}^{(k)} \in \mathbb{R}^n$. We refer to the above problem as the *multivariate multi-response (MVMR) linear regression model*, and the goal is to recover $\vec{\beta}^{*(k)}$ for $k = 1, \dots, K$ jointly. This problem has been studied in [36] via the l_1/l_2 -regularized Lasso for fixed matrices $X^{(1)}, \dots, X^{(K)}$. For random design matrices, this model has been studied via l_1/l_∞ -regularized Lasso in [1] and via $l_1/l_1 + l_1/l_\infty$ -regularized Lasso in [2] for incorporating both row sparsity and individual sparsity.

In this paper, we study the MVMR problem for random design matrices via l_1/l_2 -regularized Lasso. Although this may seem to only likely offer expected results similar to those in [29], [1], and [2], our exploration turns out to provide more insights which were not captured in previous studies. We discuss these in depth in Section 1.2. In our model, it is assumed that the design matrices are Gaussian distributed, and are independent but not identical across tasks.

For each task k , the row vector of $X^{(k)}$ is Gaussian with mean zero and the covariance matrix $\Sigma^{(k)}$ for $k = 1, \dots, K$. The noise vectors and hence the output vectors are also Gaussian distributed and independent across tasks. We are interested in joint recovery of the union of the support sets (i.e., the support union) of regression vectors $\vec{\beta}^{*(1)}, \dots, \vec{\beta}^{*(K)}$. We collect these vectors together as a matrix $B^* = [\vec{\beta}^{*(1)}, \dots, \vec{\beta}^{*(K)}]$.

We adopt the l_1/l_2 -regularized Lasso problem for recovery of the support union via the following optimization problem:

$$\min_{B \in \mathcal{R}^{p \times K}} \frac{1}{2n} \sum_{k=1}^K \left\| \vec{Y}^{(k)} - X^{(k)} \vec{\beta}^{(k)} \right\|_2^2 + \lambda_n \|B\|_{l_1/l_2} \quad (4)$$

where $B = [\vec{\beta}^{(1)}, \dots, \vec{\beta}^{(K)}]$. In this way, the K linear regression problems are coupled together via the regularization constraint. We show that this approach is advantageous as opposed to individual recovery of the support set for each linear regression problem. This is because the K regression models may share their samples in joint support recovery so that the total number of samples needed can be significantly reduced compared to performing each task individually.

1.1 Main Contributions

In the following, we summarize the main contributions of this work. Our results contain two parts: the achievability and the converse, corresponding respectively to sufficient and necessary conditions under which the l_1/l_2 -regularized Lasso problem recovers the support union for the MVMR linear regression problem. Our proof adapts the techniques developed in [18] and in [29], but involves nontrivial development to deal with the differently distributed design matrices across tasks. This also leads to interesting generalization of the results in [29] as we articulate in Section 1.2.

More specifically, we show that under certain conditions that the distributions of the design matrices satisfy, if $n > c_{p1} \psi(B^*, \Sigma^{(1:K)}) \log(p - s)$, where $\psi(\cdot)$ is defined in (8) in Section 2.1 and c_{p1} is a constant, then the l_1/l_2 -regularized Lasso recovers the support union for the MVMR linear regression problem; and if $n < c_{p2} \psi(B^*, \Sigma^{(1:K)}) \log(p - s)$, where c_{p2} is a constant, then the l_1/l_2 -regularized Lasso fails to recover the support union. Thus, $\psi(B^*, \Sigma^{(1:K)}) \log(p - s)$ serves as a sharp threshold on the sample size.

In particular, $\psi(B^*, \Sigma^{(1:K)})$ captures the sparsity of B^* and the statistical properties of the design matrices, which are important in determining the sufficient and necessary conditions for successful recovery of the support union. The property of $\psi(B^*, \Sigma^{(1:K)})$ also captures the advantages of the multi-task Lasso over solving each problem individually via the single-task Lasso. We show that when the K tasks share the same support sets (although the design matrices can be differently distributed), $\psi(B^*, \Sigma^{(1:K)}) = \frac{1}{K} \max_{1 \leq k \leq K} \psi(\vec{\beta}_k^*, \Sigma^{(k)})$.

This means that the number of samples needed per task for multi-task Lasso to jointly recover the support union is reduced by K compared to that of single-task Lasso to recover each support set individually. On the other hand, if the K tasks have disjoint support sets, then $\psi(B^*, \Sigma^{(1:K)}) = \max_{1 \leq k \leq K} \psi(\vec{\beta}^{*(k)}, \Sigma^{(k)})$. This implies that the number of samples needed per task to correctly recover the support union is almost the same as that of single-task Lasso to recover each support individually. Between these two extreme cases, tasks can have overlapped support sets with different overlapping levels, and the impact of these properties on the sample size for recovery of the support union is quantitatively captured by $\psi(B^*, \Sigma^{(1:K)})$.

1.2 Comparison to Previous Results

As we mentioned before, the MVMR model has also been studied in [1] and [2], in which l_1/l_∞ and $l_1/l_1 + l_1/l_\infty$ -regularization were adopted for support union recovery, respectively. In these studies, sharp threshold on sample complexity is characterized only for $K = 2$ and under special conditions on $\frac{1}{n} X_{S_k}^{(k)T} X_{S_k}^{(k)}$. In our work, using l_1/l_2 -regularized Lasso, we are able to characterize the sharp threshold under standard regularization conditions.

The MVMR model (with differently distributed design matrices across tasks) can be viewed as generalization of the multi-response model (with an identical design matrix across tasks) studied in [29]. It is thus interesting to compare our results to the results in [29]. For the scenario when the tasks share the same regression vector, it is shown in [29] that the major advantage of jointly solving a multi-task Lasso problem over solving each single-task Lasso problem individually is reduction of effective noise variance by the factor K . But the sample size needed per task for recovery of the support union via multi-task Lasso is the same as that needed for recovery of each support set individually via single-task Lasso. This implies that multi-task Lasso does not offer benefit in reducing the sample size (in the order sense) for this case. Our result, on the other hand, shows that the benefit in sample complexity by using multi-task Lasso for recovery of support union arises when the design matrices are differently distributed across tasks. For such a case, the sample size needed per task is reduced by K via multi-task Lasso compared to recovery of each support set individually via single-task Lasso. Consequently, our result is a nontrivial generalization of the result in [29]. For the scenario when the tasks have disjoint support sets, our result is consistent with the result in [29], which suggests that there is no advantage of performing multi-task Lasso as opposed to performing single-task Lasso for each task.

1.3 Relationship to Jointly Learning Multiple Markov Networks

One application of the MVMR linear regression model is to jointly learning multiple Gaussian Markov network structures. In this context, it solves a multi-task neighbor selection problem. This is also a natural scenario, in which features and their distributions vary across tasks.

We consider K Gaussian Markov networks, each with $p + 1$ nodes represented by $X_1^{(k)}, \dots, X_{p+1}^{(k)}$ for $k = 1, \dots, K$. The distribution of the Gaussian vector for graph k is given by $\mathcal{N}\left(0, \Sigma_{p+1}^{(k)}\right)$, where $\Sigma_{p+1}^{(k)} \in \mathbb{R}^{(p+1) \times (p+1)}$. Assume for each graph, there are n i.i.d. samples generated based on the joint distribution of the nodes. The objective is to estimate the connection relationship of nodes based on the samples. We denote n samples of each variable $X_j^{(k)}$ by a column vector $\vec{X}_j^{(k)} \in \mathbb{R}^n$ for $j = 1, \dots, p + 1$ and $k = 1, \dots, K$. For each graph k and each node with index a , the sample vector $\vec{X}_a^{(k)}$ can be expressed as:

$$\vec{X}_a^{(k)} = X_{-a}^{(k)} \vec{\beta}^{(k)} + \vec{W}_a^{(k)} \quad (5)$$

where $X_{-a}^{(k)}$ is an $n \times p$ matrix that contains all column vectors $\vec{X}_j^{(k)}$ for $j \neq a$, $\vec{\beta}^{(k)}$ is a p -dimensional vector consisting of the estimation parameters of $X_a^{(k)}$ given $X_j^{(k)}$ with $j \neq a$, and $\vec{W}_a^{(k)}$ is the n -dimensional Gaussian vector containing i.i.d. components with zero mean and variance given by

$$\sigma_W^{(k)2} = \text{Var}(X_{1a}) - \text{Cov}(X_{1a}, X_{1,-a}) \text{Cov}^{-1}(X_{1,-a}) \text{Cov}(X_{1,-a}, X_{1a}).$$

It has been shown that the nonzero components of the vector $\vec{\beta}^{(k)}$ represent existence of the edges between the corresponding nodes and node a in graph k . Hence, estimation of the support set of $\vec{\beta}^{(k)}$ provides an estimation of the graph structure, which is referred to as the *neighbor selection problem* in [37].

Therefore, multi-task Lasso for the MVMR linear regression problem provides an useful approach for joint neighbor selection over K graphs. It is clear that in this case, the design matrices $X_{-a}^{(k)}$ in general have different distributions across k , and hence the MVMR model is well justified. We note that jointly learning multiple graphs has also been studied in [38] and [39], which adopted a different objective function of the precision matrix Σ^{-1} . Via the MVMR linear regression model, we characterize the threshold-based sufficient and necessary conditions for joint recovery of the graphs.

2 Problem Formulation and Notations

In this paper, we study the MVMR linear regression problem given by (3), which contains K linear regressions. Here, the design matrices $X^{(1)}, \dots, X^{(K)}$ and noise vectors $\vec{W}^{(1)}, \dots, \vec{W}^{(K)}$ are Gaussian distributed, and are independent but not identical across k . For each task k , $X^{(k)}$ has independent and identically distributed (i.i.d.) row vectors with each being Gaussian with mean zero and covariance matrix $\Sigma^{(k)}$, and the noise vector $\vec{W}^{(k)}$ has i.i.d. components with each being Gaussian with mean zero and variance $\sigma_W^{(k)2}$. We let $\sigma_{max} = \max_{1 \leq k \leq K} \sigma_W^{(k)2}$.

In (3), $\vec{\beta}^{*(k)}$ denotes the true regression vector for each task k . We define the support set for each $\vec{\beta}^{*(k)}$ as $S_k := \{j \in \{1, \dots, p\} | \vec{\beta}_j^{*(k)} \neq 0\}$. The support union over K tasks is

defined to be $S := \cup_{k=1}^K S_k$. In this paper, we are interested in estimating the support union jointly for K tasks.

We adopt the l_1/l_2 -regularized Lasso to recover the support union for the MVMR linear regression model. More specifically, we solve the multi-task Lasso given in (4) and rewritten below:

$$\min_{B \in \mathbb{R}^{p \times K}} \frac{1}{2n} \sum_{k=1}^K \left\| \vec{Y}^{(k)} - X^{(k)} \vec{\beta}^{(k)} \right\|_2^2 + \lambda_n \|B\|_{l_1/l_2} \quad (6)$$

where $B = [\vec{\beta}^{(1)}, \dots, \vec{\beta}^{(K)}]$. In this way, the K linear regression problems are coupled together via the regularization constraint. In this paper, we characterize conditions under which the solution to the above multi-task Lasso problem correctly recover the support union of the true regression vectors for K tasks.

2.1 Notations

We introduce some notations that we use in this paper. For a matrix $A \in \mathbb{R}^{p \times K}$, we define the l_a/l_b block norm as

$$\|A\|_{l_a/l_b} := \left[\sum_{i=1}^p \left(\sum_{j=1}^K |A_{ij}|^b \right)^{a/b} \right]^{1/a}. \quad (7)$$

We also define the operator norm for a matrix as

$$\|A\|_{a,b} := \sup_{\|x\|_b=1} \|Ax\|_a.$$

In particular, we define the spectral norm as $\|A\|_2 = \|A\|_{2,2}$ and the l_∞ -operator norm as $\|A\|_\infty = \|A\|_{\infty,\infty} = \max_{j=1,\dots,p} \sum_{k=1}^K |A_{jk}|$, which are special cases of the operator norm.

For matrix $B = [\vec{\beta}^{(1)}, \dots, \vec{\beta}^{(K)}]$ that appears in (6), $\vec{\beta}^{(k)}$ denotes its k th columns for $k = 1, \dots, K$. We further let B_i to be the i th row of B . Similarly, for $B^* = [\vec{\beta}^{*(1)}, \dots, \vec{\beta}^{*(K)}]$ that contains true regression vectors, its k th column is denoted by $\vec{\beta}^{*(k)}$ and the i th row is denoted by B_i^* . We next define the normalized row vectors of B^* as

$$Z_i^* = \begin{cases} \frac{B_i^*}{\|B_i^*\|_{l_2}} & \text{if } B_i^* \neq 0 \\ 0 & \text{otherwise,} \end{cases}$$

and define the matrix Z^* to contain Z_i^* as its i th row for $i = 1, \dots, p$. To avoid confusion, we use \hat{B} to denote the solution to the multi-task Lasso problem (6).

The support union $S(B)$ for a matrix $B \in \mathbb{R}^{p \times K}$ is denoted as $S(B) = \{i \in \{1, \dots, p\} | B_i \neq 0\}$, which includes indices of the nonzero rows of the matrix B . We use S to represent $S(B^*)$ (i.e., the true support union) for convenience and use S^c to denote the complement of the set S . We let $s = |S|$ denote the size of the set S . For any matrix $X^{(k)} \in \mathbb{R}^{n \times p}$, the matrix $X_S^{(k)}$ contains the columns of matrix $X^{(k)}$ with column indices in the set S , and $X_{S^c}^{(k)}$ contains the columns of matrix $X^{(k)}$ with column indices in the set S^c . Similarly, B_S^* and Z_S^* respectively contain rows of B^* and Z^* with indices in S .

As each row of matrix $X^{(k)}$ is Gaussian distributed as $\mathcal{N}(0, \Sigma^{(k)})$, we use $\Sigma_{SS}^{(k)}$ to denote the covariance matrix for each row of $X_S^{(k)}$, and use $\Sigma_{S^c S}^{(k)}$ to denote the cross covariance between rows of $X_{S^c}^{(k)}$ and $X_S^{(k)}$.

For convenience, we use $\Sigma^{(1:K)}$ to denote a set of matrices $\Sigma^{(1)}, \dots, \Sigma^{(K)}$. We also define the following functions of matrices $Q^{(1:K)}$ to simplify our notations:

$$\begin{aligned} \rho_u(Q^{(1:K)}) &:= \max_{j \in S^c} \max_{1 \leq k \leq K} Q_{jj}^{(k)}, \\ \rho_l(Q^{(1:K)}) &:= \min_{i, j \in S^c, j \neq i} \min_{1 \leq k \leq K} \left[Q_{jj}^{(k)} + Q_{ii}^{(k)} - 2Q_{ji}^{(k)} \right]. \end{aligned}$$

In particular, our results contain the functions $\rho_u\left(\Sigma_{S^c S^c|S}^{(1:K)}\right)$ and $\rho_l\left(\Sigma_{S^c S^c|S}^{(1:K)}\right)$, where $\Sigma_{S^c S^c|S}^{(k)}$ is the covariance matrix of each row of $X_{S^c}^{(k)}$ with $X_S^{(k)}$ given.

For matrix B^* , we define $b_{min}^* = \min_{j \in S} \|B_j^*\|_{l_2}$. We define the following function that captures the sparsity of B^* and the statistical properties of the design matrices $X_S^{(1:K)}$:

$$\psi(B^*, \Sigma^{(1:K)}) := \max_{1 \leq k \leq K} \overrightarrow{Z}_{S^c}^{*T} \left(\Sigma_{SS}^{(k)} \right)^{-1} \overrightarrow{Z}_{S^c}^*, \quad (8)$$

where $\overrightarrow{Z}_{S^c}^*$ is the k th column of $Z_{S^c}^*$. We note that this definition of $\psi(\cdot)$ function is different from the previous work [29] with the same design matrix for all tasks. Here, due to different design matrices across the K tasks, $\psi(\cdot)$ depends on K quantities $\overrightarrow{Z}_{S^c}^{*T} \left(\Sigma_{SS}^{(k)} \right)^{-1} \overrightarrow{Z}_{S^c}^*$ with each depending on a column vector $\overrightarrow{Z}_{S^c}^*$.

We denote $g(\cdot) = o(f(\cdot))$ if $\lim_{n \rightarrow \infty} \frac{g(\cdot)}{f(\cdot)} \rightarrow 0$, and $g(\cdot) = O(f(\cdot))$ if $\lim_{n \rightarrow \infty} \frac{g(\cdot)}{f(\cdot)} \rightarrow c_o$, where the constant $0 < c_o < \infty$.

3 Main Results

In this section, we provide our main results on using l_1/l_2 -regularized Lasso to recover the support union for the MVMR linear regression model. Our results contain two parts: one is the achievability, i.e., sufficient conditions for the l_1/l_2 -regularized Lasso to recover the support union; and the other is the converse, i.e., conditions under which the l_1/l_2 -regularized

Lasso fails to recover the support union. We then discuss implications of our results by considering a few representative scenarios, and compare our results with those for the multivariate linear regression with an identical design matrix across tasks.

3.1 Achievability and Converse

We first introduce a number of conditions on covariance matrices $\Sigma^{(k)}$ for $k = 1, \dots, K$, which are useful for the statements of our results.

(C1). There exists a real number $\gamma \in (0, 1]$ such that $\|A\|_\infty \leq 1 - \gamma$, where $A_{js} = \max_{1 \leq k \leq K} \left| \left(\Sigma_{S^c S}^{(k)} \left(\Sigma_{SS}^{(k)} \right)^{-1} \right)_{js} \right|$ for $j \in S^c$ and $s \in S$.

(C2). There exist constants $0 < C_{min} \leq C_{max} < +\infty$ such that all eigenvalues of the matrix $\Sigma_{SS}^{(k)}$ are contained in the interval $[C_{min}, C_{max}]$ for $k = 1, \dots, K$.

(C3). There exists a constant $D_{max} < +\infty$ such that $\max_{1 \leq k \leq K} \left\| \left(\Sigma_{SS}^{(k)} \right)^{-1} \right\|_\infty \leq D_{max}$.

In this paper, we consider the asymptotic regime, in which $p \rightarrow \infty$, $s \rightarrow \infty$, and $\log(p - s) \rightarrow +\infty$. In such a regime, we introduce the conditions on the regularization parameter and the sample size n as follows:

(P1). Regularization parameter $\lambda_n = \sqrt{\frac{f(p) \log p}{n}}$, where the function $f(p)$ is chosen such that $f(p) \rightarrow +\infty$ as $p \rightarrow +\infty$, and $\frac{f(p) \log p}{n} \rightarrow 0$ as $n \rightarrow \infty$, i.e., $\lambda_n \rightarrow 0$ as $n \rightarrow +\infty$.

(P2). Define $\rho(n, s, \lambda_n)$ as

$$\rho(n, s, \lambda_n) := \sqrt{\frac{8\sigma_{max}^2 s \log s}{nC_{min}}} + \lambda_n \left(D_{max} + \frac{12s}{C_{min}\sqrt{n}} \right)$$

and require $\frac{\rho(n, s, \lambda_n)}{b_{min}^*} = o(1)$.

The following theorem characterizes sufficient conditions for recovery of the support union via l_1/l_2 -regularized Lasso.

Theorem 1. *Consider the MVMR problem in the asymptotic regime, in which $p \rightarrow \infty$, $s \rightarrow \infty$ and $\log(p - s) \rightarrow \infty$. We assume that the parameters $(n, p, s, B^*, \Sigma^{(1:K)})$ satisfy the conditions (C1)-(C3), and (P1)-(P2). If for some small constant $v > 0$,*

$$n > 2(1 + v)\psi(B^*, \Sigma^{(1:K)}) \log(p - s) \frac{\rho_u \left(\Sigma_{S^c S^c | S}^{(1:K)} \right)}{\gamma^2}, \quad (9)$$

then the multi-task Lasso problem (6) has a unique solution \widehat{B} , the support union $S(\widehat{B})$ is the same as the true support union $S(B^)$, and $\|\widehat{B} - B^*\|_{l_\infty/l_2} = o(b_{min}^*)$ with the probability*

greater than

$$1 - K \exp(-c_0 \log s) - \exp(-c_1 \log(p - s)) \quad (10)$$

where c_0 and c_1 are constants.

Theorem 1 provides sufficient conditions on the sample size such that the solution to the l_1/l_2 -regularized Lasso problem correctly recovers the support union of the MVMR linear regression model. We next provides a theorem about the conditions on the sample size under which the solution to the l_1/l_2 -regularized Lasso problem fails to recover the support union.

Theorem 2. *Consider the MVMR problem in the asymptotic regime, in which $p \rightarrow \infty$, $s \rightarrow \infty$ and $\log(p - s) \rightarrow \infty$. We assume that the parameters $(n, p, s, B^*, \Sigma^{(1:K)})$ satisfy the conditions (C1)-(C2) and the conditions: $s/n = o(1)$ and $\frac{1}{\lambda_n^2 s} \rightarrow 0$. If for some small constant $v > 0$,*

$$n < 2(1 - v)\psi(B^*, \Sigma^{(1:K)}) \log(p - s) \frac{\rho_l \left(\Sigma_{(S^c S^c | S)}^{(1:K)} \right)}{(2 - \gamma)^2}, \quad (11)$$

then with the probability greater than

$$1 - \exp(-c_2 s) - c_3 \exp\left(-c_4 \frac{n}{s}\right) \quad (12)$$

for some positive constants c_2, c_3 and c_4 , no solution \widehat{B} to the multi-task Lasso problem (6) recovers the true support union and achieves $\|\widehat{B} - B^*\|_{l_\infty/l_2} = o(b_{min}^*)$.

The proofs of Theorems 1 and 2 are provided in Sections 5 and 6, respectively. Combining Theorems 1 and 2, it is clear that the quantity $\psi(B^*, \Sigma^{(1:K)}) \log(p - s)$ serves as a threshold on the sample size n , which is tight in the order sense. As the sample size is above the threshold, the multi-task Lasso recovers the true support union, and as the sample size is below the threshold, the multi-task Lasso fails to recover the true support union. The following proposition provides bounds on the scaling behavior of the function $\psi(B^*, \Sigma^{(1:K)})$ in the asymptotic regime.

Proposition 1. *Consider the MVMR linear regression model with the regression matrix B^* and the covariance matrices $\Sigma^{(1:K)}$ satisfying the condition (C2), the function $\psi(B^*, \Sigma^{(1:K)})$ is bounded as*

$$\frac{s}{KC_{min}} \leq \psi(B^*, \Sigma^{(1:K)}) \leq \frac{s}{C_{min}}.$$

The proof of the Proposition 1 is provided in Appendix A.

In the next subsection, we explore the properties of the quantity $\psi(B^*, \Sigma^{(1:K)})$ in order to understand the impact of sparsity of B^* and covariance matrices $\Sigma^{(1:K)}$ on sample complexity for recovering the support union.

3.2 Implications

The quantity $\psi(B^*, \Sigma^{(1:K)})$ captures sparsity of B^* and statistical properties of design matrices $\Sigma^{(1:K)}$, and hence plays an important role in determining the conditions on the sample size for recovery of the support union as shown in Theorems 1 and 2. In this section, we analyze $\psi(B^*, \Sigma^{(1:K)})$ for a number of representative cases in order to understand advantages of multi-task Lasso which solves multiple linear regression problems jointly over single-task Lasso which solves each linear regression problem individually.

We denote $\psi(\vec{\beta}^{*(k)}, \Sigma^{(k)})$ as the function corresponding to a single linear regression problem, where $\vec{\beta}^{*(k)}$ represents the k th column of B^* . It is clear that $\psi(\vec{\beta}^{*(k)}, \Sigma^{(1:K)})$ captures the threshold on the sample size for the single-task Lasso problem. Comparison of $\psi(B^*, \Sigma^{(1:K)})$ and $\psi(\vec{\beta}^{*(k)}, \Sigma^{(k)})$ provides comparison between multi-task Lasso and single-task Lasso in terms of the number of samples needed for recovery of the support union/set. We explicitly express $\psi(B^*, \Sigma^{(1:K)})$ and $\psi(\vec{\beta}^{*(k)}, \Sigma^{(k)})$ as follows:

$$\psi(B^*, \Sigma^{(1:K)}) = \max_{1 \leq k \leq K} \sum_{i \in S} \sum_{j \in S} \frac{B_{ik}^* B_{jk}^*}{\|B_i^*\|_{l_2} \|B_j^*\|_{l_2}} \left(\left(\Sigma_{SS}^{(k)} \right)^{-1} \right)_{ij} \quad (13)$$

$$\psi(\vec{\beta}^{*(k)}, \Sigma^{(k)}) = \sum_{i \in S} \sum_{j \in S} \frac{\vec{\beta}_i^{*(k)} \vec{\beta}_j^{*(k)}}{|\vec{\beta}_i^{*(k)}| |\vec{\beta}_j^{*(k)}|} \left(\left(\Sigma_{SS}^{(k)} \right)^{-1} \right)_{ij} \quad (14)$$

where B_{ik}^* denotes the (i, k) th entry of the matrix B^* and $\vec{\beta}_i^{*(k)}$ denotes the i th entry of the vector $\vec{\beta}^{*(k)}$.

We first study the scenario, in which all K tasks have the same regression vectors, and hence have the same support sets.

Corollary 1. (*Identical Regression Vectors*) If B^* has identical column vectors, i.e., $\vec{\beta}^{*(k)} = \vec{\beta}^*$ for $k = 1, \dots, K$, then

$$\psi(B^*, \Sigma^{(1:K)}) = \frac{1}{K} \max_{1 \leq k \leq K} \psi(\vec{\beta}^*, \Sigma^{(k)}). \quad (15)$$

Proof. Under the assumption of the corollary, $B^* = \vec{\beta}^* \mathbf{1}_K^T$, where $\vec{\beta}^* \in \mathbb{R}^p$. Hence, $\vec{Z}_{S_k}^* = \frac{\text{sign}(\vec{\beta}_S^*)}{\sqrt{K}}$, where the vector $\vec{\beta}_S^*$ contains components in the support S .

$$\begin{aligned} \psi(B^*, \Sigma^{(1:K)}) &= \max_{1 \leq k \leq K} \vec{Z}_{S_k}^{*T} \left(\Sigma_{SS}^{(k)} \right)^{-1} \vec{Z}_{S_k}^* \\ &= \max_{1 \leq k \leq K} \frac{\text{sign}(\vec{\beta}_S^*)^T}{\sqrt{K}} \left(\Sigma_{SS}^{(k)} \right)^{-1} \frac{\text{sign}(\vec{\beta}_S^*)}{\sqrt{K}} \\ &= \frac{1}{K} \max_{1 \leq k \leq K} \psi(\vec{\beta}^*, \Sigma^{(k)}). \end{aligned} \quad (16)$$

□

Remark 1. *Corollary 1 implies that the number of samples per task needed to correctly recover the support union via multi-task Lasso is reduced by a factor of K compared to single-task Lasso that recovers each support set individually.*

It can be seen that although the K tasks involve design matrices that have different covariances, as long as dependence of the output variables on the feature variables is the same for all tasks, the tasks share samples in multi-task Lasso to recover the support union so that the sample size needed per task is reduced by a factor of K . Hence, there is a significant advantage of grouping tasks with similar regression vectors together for multi-task learning.

Corollary 1 can be viewed as a generalization of the result in [29], in which the design matrices for the tasks are the same. The result in [29] suggests that if the tasks share the same regression vector, there is no benefit in terms of the number of samples needed for support recovery using multi-task Lasso compared to single-task Lasso. Our result suggests that the benefit of multi-task Lasso in fact arises when the design matrices are differently distributed. For such a case, we show that the sample size needed per design matrix (i.e., per task) is reduced by the factor K .

Moreover, compared to recovery of each support set individually via single-task Lasso, multi-task Lasso also reduces sample size per task by the factor K . However, such an advantage does not appear if the K tasks have the same design matrix and regression vectors as in [29].

We next study a more general case when regression vectors are also different across tasks (but the support sets of tasks are the same) in addition to varying design matrices across tasks.

Corollary 2. *(Varying Regression Vectors with Same Supports) Suppose all entries $B_{jk}^* > 0$ for $j \in S$ and $k = 1, \dots, K$, and all coefficients are bounded, i.e., $\bar{B}_k - \Delta_k \leq B_{jk}^* \leq \bar{B}_k + \Delta_k$, where $\Delta_k > 0$ is a small perturbation constant with $\bar{B}_k > \Delta_k$. Then,*

$$\frac{\psi(B^*, \Sigma^{(1:K)})}{\max_{1 \leq k \leq K} \psi(\vec{\beta}^{*(k)}, \Sigma^{(k)})} \leq \frac{1}{K} \max_{1 \leq k \leq K} \frac{(\bar{B}_k + \Delta_k)^2}{(\bar{B}_k - \Delta_k)^2}.$$

Proof. Based on the assumption for B^* , we obtain the following upper bound on $\psi(B^*, \Sigma^{(1:K)})$ and lower bound on $\psi(\vec{\beta}^{*(k)}, \Sigma^{(k)})$:

$$\psi(B^*, \Sigma^{(1:K)}) \leq \frac{1}{K} \max_{1 \leq k \leq K} \frac{1}{(\bar{B}_k - \Delta_k)^2} \sum_{i \in S} \sum_{j \in S} B_{ik}^* B_{jk}^* \left(\left(\Sigma_{SS}^{(k)} \right)^{-1} \right)_{ij}; \quad (17)$$

$$\psi(\vec{\beta}^{*(k)}, \Sigma^{(k)}) \geq \frac{1}{(\bar{B}_k + \Delta_k)^2} \sum_{i \in S} \sum_{j \in S} B_{ik}^* B_{jk}^* \left(\left(\Sigma_{SS}^{(k)} \right)^{-1} \right)_{ij}. \quad (18)$$

Combining the above bounds, we obtain

$$\frac{\psi(B^*, \Sigma^{(1:K)})}{\max_{1 \leq k \leq K} \psi(\vec{\beta}^{*(k)}, \Sigma^{(k)})} \leq \frac{1}{K} \max_{1 \leq k \leq K} \frac{(\bar{B}_k + \Delta_k)^2}{(\bar{B}_k - \Delta_k)^2}.$$

□

Corollary 2 is a strengthened version of Corollary 1 in that Corollary 2 allows both the regression vectors and design matrices to be different across tasks and still shows that the number of samples needed is reduced by a factor of K compared to single-task Lasso, as long as the support sets across tasks are the same.

Corollary 3. (*Disjoint Support Sets*) Suppose the distribution of all design matrices are the same, i.e., $\Sigma^{(k)} = \Sigma$ for $k = 1, \dots, K$, and suppose that the support sets S_k of all tasks are disjoint. Let $s_k = |S_k|$, and hence $s = \sum_{k=1}^K s_k$. Then,

$$\psi(B^*, \Sigma^{(1:K)}) = \max_{1 \leq k \leq K} \psi(\vec{\beta}^{*(k)}, \Sigma^{(k)}),$$

where $\Sigma^{(1:K)} = (\Sigma, \dots, \Sigma)$.

Proof. By the assumption of the corollary, we obtain:

$$\psi(B^*, \Sigma^{(1:K)}) = \max_{1 \leq k \leq K} \text{sign} \left(\vec{\beta}_S^{*(k)} \right)^T \Sigma_{SS}^{-1} \text{sign} \left(\vec{\beta}_S^{*(k)} \right) = \max_{1 \leq k \leq K} \psi(\vec{\beta}^{*(k)}, \Sigma^{(k)}). \quad (19)$$

□

We note that

$$\max_{1 \leq k \leq K} \psi(\vec{\beta}^{*(k)}, \Sigma^{(k)}) \log(p - s) \leq \max_{1 \leq k \leq K} \psi(\vec{\beta}^{*(k)}, \Sigma^{(k)}) \log(p - s_k).$$

Since the number of samples needed per task for multi-task Lasso is proportional to $\max_{1 \leq k \leq K} \psi(\vec{\beta}^{*(k)}, \Sigma^{(k)}) \log(p - s)$, and the number of samples needed for single-task Lasso for task k is proportional to $\psi(\vec{\beta}^{*(k)}, \Sigma^{(k)}) \log(p - s_k)$, the above equation implies that the required number of samples for multi-task Lasso is smaller than (in fact almost the same as) that for single-task Lasso.

Corollary 3 suggests that if the tasks have disjoint support sets for regression vectors, the advantage of the multi-task Lasso vanishes. This is reasonable because the tasks do not benefit from sharing the samples for recovering the supports if their support sets are disjoint. The essential message of Corollary 3 should not change if the tasks have different design matrices and/or different regression vectors. The critical assumption in Corollary 3 is the disjoint support sets.

Corollaries 1 and 3 provide two extreme cases when the tasks share the same support sets and have disjoint support sets, respectively. The number of samples needed per task for recovery of the support union goes from $1/K$ of to the same as the sample size needed for single-task Lasso. Between these two extreme cases, tasks may have overlapped support sets with various overlapping levels. Correspondingly, the number of samples needed for recovering the support union should depend on the overlapping levels of the support sets and is captured precisely by the quantity $\psi(B^*, \Sigma^{(1:K)})$. We demonstrate such behavior via our numerically results in the next section.

4 Numerical Results

In this section, we provide numerical simulations to demonstrate our theoretical results on using block-regularized multi-task Lasso for recovery of the support union for the MVMR linear regression model. We study how the sample size needed for correct recovery of the support union depends on sparsity of the regression vectors, on the distributions of the design matrices, and on the number of tasks.

We first study the scenario considered in Corollary 1 when the K tasks have the same regression vectors, i.e., $B^* = \vec{\beta}^* \mathbf{1}_K^T$. We set $\vec{\beta}^* = \frac{1}{\sqrt{K}} \vec{1}_S$, where S is the common support set across K tasks. We set the covariance matrix $\Sigma^{(k)}$ to be different across K tasks as follows. For $k = 1, \dots, K$, we set $Cov(X_a, X_b) > 0$ (where $a, b \in \{1, 2, \dots, p\}$) if $a = b \pm 1$, and otherwise $Cov(X_a, X_b) = 0$. In particular, $Cov(X_a, X_b) = 1 + 1/k$ if $a = b \pm 1$ and a is odd, and $Cov(X_a, X_b) = 1 - 0.8/k$ if $a = b \pm 1$ and a is even. The sparsity of linear regression vectors is linearly proportional to the dimension p , i.e., $s = \alpha p$, with the parameter α controlling the sparsity of the model. We set $\alpha = 1/8$. We choose the dimension $p = 128, 256, 512$. We set the regularization parameter $\lambda_n = 3.5 \times \sqrt{\log(p-s) \log s/n}$. We solve the l_1/l_2 -regularized multi-task Lasso problem (6) for recovery of the support union for $K = 2, 4, 6, 8$.

Fig. 1 plots the probability of correct recovery of the support union as a function of the scaled sample size. It can be seen that the sample size for guaranteeing correct recovery scales in the order of $s \log(p-s)$ for all plots. Moreover, as the number of tasks K increases, the sample size (per task) needed for correct recovery decreases inversely proportionally with K , which is consistent with Corollary 1. These results demonstrate that when the regression vectors are the same across tasks, multi-task Lasso has a great advantage compared to single-task Lasso in terms of reduction in the sample size needed per task.

We are also interested in the influence of non-equal regression values on the sample size for correct recovery. Our next experiment is taken for the scenario in which all tasks share the same support sets but have non-equal regression values across tasks. For $k = 1, \dots, K$, $\vec{\beta}_j^{*(k)} = \frac{1}{\sqrt{K}} \times (1 + \frac{k}{16})$ for $j = 16t_{pe}$, and $\vec{\beta}_j^{*(k)} = \frac{1}{\sqrt{K}} \times (1 - \frac{k}{16})$ for $j = 16t_{pe} + 8$, where t_{pe} is any nonnegative integer such that $j \leq p$. The covariance matrices $\Sigma^{(k)}$ are set to be identical




Figure 1: Impact of number of tasks on the sample size for scenarios with identical regression vectors and varying distributions for design matrices across tasks




Figure 2: Impact of number of tasks on the sample size for scenarios with non-equal regression values and identical design matrix distribution across tasks

across all tasks. We set $Cov(X_a, X_b) = 1$ (where $a, b \in \{1, 2, \dots, p\}$) if $a = b \pm 1$, and otherwise $Cov(X_a, X_b) = 0$. Other parameters are chosen to be the same as the experiment in Fig. 1. Fig. 2 plots how the probability of correct recovery changes with the sample size for $p = 128, 256, 512$. It exhibits the same behavior as Fig. 1, although now the regression vectors have unequal values across tasks. In particular, it can be seen that the sample size needed for correct recovery decreases as the number of tasks increased, demonstrating the advantage of multi-task Lasso.

We next study how the overlapping levels of the support sets across tasks affect the sample size for correct recovery of the support union. We set $K = 2$, i.e., two tasks, and study three overlapping models for the two tasks: (1) same support sets $S_1 = S_2 = \{j \leq p : 8t_{pe} + 1, \text{ where integer } t_{pe} \geq 0\}$; (2) disjoint support sets $S_1 \cap S_2 = \phi$ in which $S_1 = \{j \leq p : 16t_{pe} + 1, \text{ where integer } t_{pe} \geq 0\}$ and $S_2 = \{j \leq p : 16t_{pe} + 2, \text{ where integer } t_{pe} \geq 0\}$; (3) overlapping support sets in which $S_1 = \{j \leq p : j = 24t_{pe} + 1 \text{ or } j = 24t_{pe} + 2 \text{ where integer } t_{pe} \geq 0\}$, and $S_2 = \{j \leq p : j = 24t_{pe} + 2 \text{ or } j = 24t_{pe} + 3 \text{ where integer } t_{pe} \geq 0\}$. We choose the linear sparsity model with $\alpha = 1/8$. We set $p = 128, 256, 512$, and $\Sigma^{(k)} = I_p$ for $k = 1$ and 2 . We also set $\lambda_n = 3.5 \times \sqrt{\log(p-s) \log s/n}$.

Fig. 3 compares the probability of correct recovery as a function of the scaled sample size for the three overlapping models. It can be seen that the model with the same support set requires the smallest sample size, and the model with disjoint support sets requires the largest sample size. The model with overlapping support sets needs the sample size between the two extreme models. This is reasonable because as the support sets overlap more, tasks share more information in samples for support recovery and hence need less number of samples for correct recovery.

The preceding experiment is taken for the case when the design matrices of the two tasks have the same covariance matrix and the regression vectors are identical on overlapping entries. It is interesting to investigate how non-equal values in regression vectors and different covariance matrices across the two tasks affect the sample complexity. We first study the case when the regression vectors of the two tasks do not have the same values on the overlapping entries. For the case when the two tasks have the same support sets, we let $\vec{\beta}_j^{*(k)} = \frac{1}{\sqrt{K}} \times (1 + \frac{k}{16})$ for $j = 16t_{pe}$, and $\vec{\beta}_j^{*(k)} = \frac{1}{\sqrt{K}} \times (1 - \frac{k}{16})$ for $j = 16t_{pe} + 8$, where integer $t_{pe} \geq 0$ such that $j \leq p$ for $k = 1, 2$. For the overlapping model, S_1 and S_2 are the same as the preceding experiment. For $k = 1$, $\vec{\beta}_j^{*(k)} = 1$ if $j = 24t_{pe} + 1$, and $\vec{\beta}_j^{*(k)} = \frac{1}{\sqrt{K}} \times (1 + \frac{1}{16})$ if $j = 24t_{pe} + 2$, where integer $t_{pe} \geq 0$ such that $j \leq p$. For $k = 2$, $\vec{\beta}_j^{*(k)} = \frac{1}{\sqrt{K}} \times (1 - \frac{1}{16})$ if $j = 24t_{pe} + 2$, and $\vec{\beta}_j^{*(k)} = 1$ if $j = 24t_{pe} + 3$, where integer $t_{pe} \geq 0$ such that $j \leq p$. For the disjoint case, the regression vectors are the same as the preceding experiment since no overlapping exists in the disjoint model. Other parameters $(\Sigma^{(1:K)}, n, p, s, \lambda)$ are kept the same as the preceding experiment. Fig. 4 plots the probability of correct recovery of the support union versus the scaled sample size for this experiment. It can be observed that Fig. 4 exhibits same behavior as Fig. 3 and demonstrates that higher overlapping level




Figure 3: Impact of overlapping levels of support sets on the sample size with same regression values for overlapping entries and identical distributions for design matrices across tasks




Figure 4: Impact of overlapping levels of support sets on the sample size with non-equal regression values for overlapping entries and identical covariance matrices across tasks




Figure 5: Impact of overlapping levels of support sets on the sample size with same regression values for overlapping entries and varying covariance matrices across tasks

across two tasks leads to smaller sample size needed for recovery, although the regression vectors do not match values for the overlapping entries. We also denote that more careful comparison of Fig. 4 and Fig. 3 suggests that the model with perturbation on overlapping entries in regression vectors requires a slightly larger sample size than the model without perturbation.

We finally study how the varying covariance matrices across the two tasks influence the result. We set the covariance matrices $\Sigma^{(k)}$ for $k = 1, 2$ as follows. We let $Cov(X_a, X_b) > 0$ ($a, b \in \{1, 2, \dots, p\}$) if $a = b \pm 1$, and otherwise $Cov(X_a, X_b) = 0$. More specifically, we let $Cov(X_a, X_b) = 1 + 1/k$ if $a = b \pm 1$ and a is odd, and $Cov(X_a, X_b) = 1 - 0.8/k$ if $a = b \pm 1$ and a is even. Other parameters (B^*, n, p, s, λ) are the same as the experiment in Fig. 3. Fig. 5 compares the probability of correct recovery versus the scaled sample size for the three overlapping models under the varying covariance matrices but the same values for overlapping regression entries across the two tasks. The behavior is similar to that in Fig. 3 and Fig. 4. More careful comparison of Fig. 5 and Fig. 3 suggests that the varying covariance matrices across the two tasks require larger sample size than the case with identical covariance matrices.

5 Proof of Theorem 1

Our proof applies the framework developed in [18] and in [29] based on the idea of primal-dual witness. However, for the MVMR model, we need to develop novel adaption due to varying design matrices across tasks. In [29], since the model can be expressed by a matrix operation on regression matrix, the proof involves many operations for matrices, for which properties/bounds for matrices can be applied. However, the MVMR model is expressed by K operations on individual regression vectors. The proof mostly involves first manipulating/bounding individual regression vectors and then integrating these manipulations/bounds together for conditions across all tasks. Our adaption needs to make bounds in both steps as tight as possible in order to develop sharp threshold conditions. We next present our proof in detail.

The objective function in the multi-task Lasso problem given in (6) is convex, and hence the following Karush-Kuhn-Tucker (KKT) condition is sufficient and necessary to characterize an optimal solution:

$$\nabla_B f(B) + \lambda_n Z = 0 \quad (20)$$

where $f(B) = \frac{1}{2n} \sum_{k=1}^K \left\| \vec{Y}^{(k)} - X^{(k)} \vec{\beta}^{(k)} \right\|^2$, and $Z \in \partial \|B\|_{l_1/l_2}$.

Before introducing the sufficient conditions, we first present the following lemma which provides an important property about the optimal solution to the above problem.

Lemma 1. *Suppose there exists an optimal solution \hat{B} to the multi-task Lasso problem given in (6). Suppose \hat{Z} is in the subdifferential of $\|B\|_{l_1/l_2}$ at \hat{B} , and satisfies the KKT condition in (20) jointly with \hat{B} . Suppose that \hat{Z} satisfies $\left\| \hat{Z}_\Omega \right\|_{l_\infty/l_2} < 1$, where \hat{Z}_Ω denotes the submatrix that contains rows of \hat{Z} with indices in the set Ω . Then any optimal solution \tilde{B} to (6) must satisfy $\tilde{B}_\Omega = 0$.*

The proof of Lemma 1 is similar to that of Lemma 1 in [18]. For completeness of our paper, we provide the proof of Lemma 1 in Appendix B.

We now construct a pair (\hat{B}, \hat{Z}) that satisfy the KKT condition in (20). We first let \hat{B}_S be an optimal solution to the following optimization problem:

$$\hat{B}_S = \operatorname{argmin}_{B_S} \left[f(B)|_{B_{S^c}=0} + \lambda_n \|B_S\|_{l_1/l_2} \right] \quad (21)$$

and let \hat{Z}_S be the associated element in the subdifferential of $\|B_S\|_{l_1/l_2}$ such that (\hat{B}_S, \hat{Z}_S) satisfy the KKT condition for the optimization problem given in (21). We then let $\hat{B}_{S^c} = 0$, and let \hat{Z}_{S^c} be an element in the subdifferential of $\|B_{S^c}\|_{l_1/l_2}$ that satisfies the KKT condition jointly with $\hat{B}_{S^c} = 0$ for the following problem

$$\operatorname{argmin}_{B_{S^c}} \left[f(B)|_{B_S=\hat{B}_S} + \lambda_n \|B_{S^c}\|_{l_1/l_2} \right]. \quad (22)$$

Such \widehat{Z}_{S^c} must exist if the KKT condition for the optimization problem (22) implies $\left\| \widehat{Z}_{S^c} \right\|_{l_\infty/l_2} \leq$

1. Now it is easy to see that $(\widehat{B}, \widehat{Z})$ obtained above satisfies the KKT condition in (20) and is hence an optimal solution to the problem (6). Furthermore, following Lemma 1, if $\left\| \widehat{Z}_{S^c} \right\|_{l_\infty/l_2} < 1$, then any optimal solution \widetilde{B} to (6) satisfies $\widetilde{B}_{S^c} = 0$. Therefore, condition $\left\| \widehat{Z}_{S^c} \right\|_{l_\infty/l_2} < 1$ guarantees both that there exists an optimal solution with the structure described as above and that all optimal solutions \widetilde{B} satisfies $\widetilde{B}_{S^c} = 0$. Furthermore, the condition $\left\| \widehat{Z}_{S^c} \right\|_{l_\infty/l_2} < 1$ guarantees uniqueness of the optimal solution. The arguments follow from the proof of Lemma 2 in [18].

We next proceed to characterize the conditions that guarantee $\left\| \widehat{Z}_{S^c} \right\|_{l_\infty/l_2} < 1$. For $j \in S^c$ and $k = 1, \dots, K$, we have

$$\widehat{Z}_{jk} = -\frac{1}{\lambda_n n} \overrightarrow{X}_j^{(k)T} \left(\Pi_S^{(k)} - I_n \right) \overrightarrow{W}^{(k)} + \frac{1}{n} \overrightarrow{X}_j^{(k)T} X_S^{(k)} \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \widehat{Z}_{S^c k}, \quad (23)$$

where $\overrightarrow{X}_j^{(k)}$ denotes the j th column of the matrix $X^{(k)}$, $\widehat{\Sigma}_{SS}^{(k)} = \frac{1}{n} X_S^{(k)T} X_S^{(k)}$, and $\Pi_S^{(k)} = \frac{X_S^{(k)} \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} X_S^{(k)T}}{n}$. The steps to obtain the above \widehat{Z}_{jk} is provided in Appendix C for completeness.

Analysis of V_{S^c} : We let $V_j = \left(\widehat{Z}_{j1}, \dots, \widehat{Z}_{jK} \right)$. We need to characterize the conditions so that $\|V_j\|_{l_2} < 1$ for all $j \in S^c$ with high probability. We write V_j into three terms as follows

$$\begin{aligned} V_j &= \underbrace{\mathbb{E} \left(V_j \mid X_S^{(1:K)} \right)}_{T_{j1}} \\ &\quad + \underbrace{\mathbb{E} \left(V_j \mid X_S^{(1:K)}, \overrightarrow{W}^{(1:K)} \right) - \mathbb{E} \left(V_j \mid X_S^{(1:K)} \right)}_{T_{j2}} \\ &\quad + \underbrace{V_j - \mathbb{E} \left(V_j \mid X_S^{(1:K)}, \overrightarrow{W}^{(1:K)} \right)}_{T_{j3}} \end{aligned} \quad (24)$$

where $X_S^{(1:K)} = \left(X_S^{(1)}, \dots, X_S^{(K)} \right)$ and $\overrightarrow{W}^{(1:K)} = \left(\overrightarrow{W}^{(1)}, \dots, \overrightarrow{W}^{(K)} \right)$. We next evaluate T_{j1}, T_{j2} , and T_{j3} one by one.

Evaluation of T_{j1} : By the definition of \widehat{Z}_S , we have the following conditional independencies:

$$\left(\overrightarrow{W}^{(k)} \perp \overrightarrow{X}_j^{(k)} \mid X_S^{(1:K)} \right), \quad \left(\widehat{Z}_{S^c k} \perp \overrightarrow{X}_j^{(k)} \mid X_S^{(1:K)} \right), \quad \left(\widehat{Z}_{S^c k} \perp \overrightarrow{X}_j^{(k)} \mid X_S^{(1:K)}, \overrightarrow{W}^{(1:K)} \right). \quad (25)$$

Given the above independence properties, we first derive

$$\begin{aligned}
\mathbb{E} \left(\widehat{Z}_{jk} \middle| X_S^{(1:K)} \right) &= -\frac{1}{\lambda_n n} \mathbb{E} \left(\overrightarrow{X}_j^{(k)T} \middle| X_S^{(1:K)} \right) \left(\Pi_S^{(k)} - I_n \right) \mathbb{E} \left(\overrightarrow{W}^{(k)} \right) \\
&\quad + \frac{1}{n} \mathbb{E} \left(\overrightarrow{X}_j^{(k)T} \middle| X_S^{(1:K)} \right) X_S^{(k)} \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \mathbb{E} \left(\widehat{Z}_{Sk} \middle| X_S^{(1:K)} \right) \\
&= \Sigma_{jS}^{(k)} \left(\Sigma_{SS}^{(k)} \right)^{-1} \mathbb{E} \left(\widehat{Z}_{Sk} \middle| X_S^{(1:K)} \right)
\end{aligned} \tag{26}$$

for $j \in S^c$, where $\Sigma_{jS}^{(k)}$ represents the covariance between a component in $\overrightarrow{X}_j^{(k)}$ and a row in $X_S^{(k)}$. We then obtain the following bound on $\|T_{j1}\|_{l_2}$ with the proof provided in Appendix D:

$$\|T_{j1}\|_{l_2} \leq \sum_{a=1}^{|S|} A_{ja}, \tag{27}$$

where $A_{ja} = \max_k \left| \left(\Sigma_{S^c S}^{(k)} \left(\Sigma_{SS}^{(k)} \right)^{-1} \right)_{ja} \right|$ for $j \in S^c$ and $a \in S$. We hence obtain

$$\max_{j \in S^c} \|T_{j1}\|_{l_2} \leq \max_{j \in S^c} \sum_{a=1}^{|S|} A_{ja} = \|A\|_{\infty} \leq 1 - \gamma.$$

Evaluation of T_{j2} : Due to the independency $\left(\widehat{Z}_{Sk} \perp \overrightarrow{X}_j^{(k)} \middle| X_S^{(1:K)}, \overrightarrow{W}^{(1:K)} \right)$, we obtain

$$\begin{aligned}
&\mathbb{E} \left(\widehat{Z}_{jk} \middle| X_S^{(1:K)}, \overrightarrow{W}^{(1:K)} \right) \\
&= -\frac{1}{\lambda_n n} \mathbb{E} \left(X_j^{(k)T} \middle| X_S^{(1:K)}, \overrightarrow{W}^{(1:K)} \right) \left(\Pi_S^{(k)} - I_n \right) \overrightarrow{W}^{(k)} \\
&\quad + \frac{1}{n} \mathbb{E} \left(X_j^{(k)T} \middle| X_S^{(1:K)}, \overrightarrow{W}^{(1:K)} \right) X_S^{(k)} \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \mathbb{E} \left(\widehat{Z}_{Sk} \middle| X_S^{(1:K)}, \overrightarrow{W}^{(1:K)} \right) \\
&= \Sigma_{jS}^{(k)} \left(\Sigma_{SS}^{(k)} \right)^{-1} \widehat{Z}_{Sk}
\end{aligned} \tag{28}$$

where the second equality follows because \widehat{Z}_{Sk} is a function of $X_S^{(1:K)}$ and $\overrightarrow{W}^{(1:K)}$. We then obtain

$$\mathbb{E} \left(\widehat{Z}_{jk} \middle| X_S^{(1:K)}, \overrightarrow{W}^{(1:K)} \right) - \mathbb{E} \left(\widehat{Z}_{jk} \middle| X_S^{(1:K)} \right) = \Sigma_{jS}^{(k)} \left(\Sigma_{SS}^{(k)} \right)^{-1} \left(\widehat{Z}_{Sk} - \mathbb{E} \left(\widehat{Z}_{Sk} \middle| X_S^{(1:K)} \right) \right). \tag{29}$$

Thus, following from steps similar to those in Appendix D, we obtain

$$\|T_{j2}\|_{l_2} \leq \sum_{a=1}^{|S|} A_{ja} \left\| \widehat{Z}_S - \mathbb{E}(\widehat{Z}_S | X_S^{(1:K)}) \right\|_{l_{\infty}/l_2}, \tag{30}$$

and hence

$$\begin{aligned}
\max_{j \in S^c} \|T_{j2}\|_{l_2} &\leq \max_{j \in S^c} \sum_{a=1}^{|S|} A_{ja} \left\| \widehat{Z}_S - \mathbb{E}(\widehat{Z}_S | X_S^{(1:K)}) \right\|_{l_\infty/l_2} \\
&= \|A\|_\infty \left\| \widehat{Z}_S - \mathbb{E}(\widehat{Z}_S | X_S^{(1:K)}) \right\|_{l_\infty/l_2} \\
&\leq (1 - \gamma) \left\| \widehat{Z}_S - \mathbb{E}(\widehat{Z}_S | X_S^{(1:K)}) \right\|_{l_\infty/l_2} \\
&\leq (1 - \gamma) \left\| \widehat{Z}_S - Z_S^* \right\|_{l_\infty/l_2} + (1 - \gamma) \mathbb{E} \left[\left\| \widehat{Z}_S - Z_S^* \right\|_{l_\infty/l_2} \middle| X_S^{(1:K)} \right]. \quad (31)
\end{aligned}$$

We next provide the following lemma given in [29], which is useful for our proof.

Lemma 2. [29] Consider the matrix $\Delta \in R^{S \times K}$ with rows $\Delta_i := \frac{\widehat{B}_i - B_i^*}{\|B_i^*\|_2}$. If $\|\Delta\|_{l_\infty/l_2} < \frac{1}{2}$, then

$$\|\widehat{Z}_S - Z_S^*\|_{l_\infty/l_2} \leq 4\|\Delta\|_{l_\infty/l_2}.$$

By applying the above lemma, given the condition $\|\Delta\|_{l_\infty/l_2} < \frac{1}{2}$ that we will show later, we obtain

$$\max_{j \in S^c} \|T_{j2}\|_{l_2} \leq 4(1 - \gamma) \left(\|\Delta\|_{l_\infty/l_2} + \mathbb{E} \left[\|\Delta\|_{l_\infty/l_2} \middle| X_S^{(1:K)} \right] \right)$$

We will show later in the analysis of U_S that $\|\Delta\|_{l_\infty/l_2}$ is of order $o(1)$ with high probability, and hence the above inequality holds with high probability.

Evaluation of T_{j3} : We introduce the vector $\vec{D}^{(k)}$ such that

$$\begin{aligned}
\widehat{Z}_{jk} &= -\frac{1}{\lambda_n n} \vec{X}_j^{(k)T} \left(\Pi_S^{(k)} - I_n \right) \vec{W}^{(k)} + \frac{1}{n} \vec{X}_j^{(k)T} X_S^{(k)} \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \widehat{Z}_{Sk} \\
&= \vec{X}_j^{(k)T} \vec{D}^{(k)}. \quad (32)
\end{aligned}$$

It is clear that for $j \in S^c$,

$$\text{Cov} \left(\vec{X}_j^{(k)} | X_S^{(1:K)}, \vec{W}^{(1:K)} \right) = \left(\Sigma_{S^c S^c | S}^{(k)} \right)_{jj} I_n.$$

Under the condition that $X_S^{(1:K)}$ and $\vec{W}^{(1:K)}$ are given, we have

$$\left(\widehat{Z}_{jk} | X_S^{(1:K)}, \vec{W}^{(1:K)} \right) - \mathbb{E} \left[\widehat{Z}_{jk} | X_S^{(1:K)}, \vec{W}^{(1:K)} \right] \sim \mathcal{N}(0, \sigma_{jk}^2) \quad (33)$$

where

$$\sigma_{jk}^2 = \frac{1}{n} \left(\Sigma_{S^c S^c | S}^{(k)} \right)_{jj} \widehat{Z}_{Sk}^T \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \widehat{Z}_{Sk} - \frac{1}{n^2 \lambda_n^2} \left(\Sigma_{S^c S^c | S}^{(k)} \right)_{jj} \vec{W}^{(k)T} \left(\Pi_S^{(k)} - I_n \right) \vec{W}^{(k)}. \quad (34)$$

Given $(X_S^{(1:K)}, \vec{W}^{(1:K)})$, \widehat{Z}_{jk} is independently distributed across k for $k = 1, \dots, K$. Hence,

$$\widehat{Z}_{jk} - \mathbb{E} \left[\widehat{Z}_{jk} | X_S^{(1:K)}, \vec{W}^{(1:K)} \right] \stackrel{d.}{=} \sigma_{jk} \xi_{jk} \quad \text{given} \quad \left(X_S^{(1:K)}, \vec{W}^{(1:K)} \right) \quad (35)$$

where $\xi_{jk} \sim \mathcal{N}(0, 1)$ is independently distributed across k for $k = 1, \dots, K$. Thus,

$$\left\| V_j - \mathbb{E} \left[V_j | X_S^{(1:K)}, \vec{W}^{(1:K)} \right] \right\|_{l_2}^2 \stackrel{d.}{=} \sum_{k=1}^K \sigma_{jk}^2 \xi_{jk}^2 \quad \text{given} \quad \left(X_S^{(1:K)}, \vec{W}^{(1:K)} \right). \quad (36)$$

We hence obtain

$$\begin{aligned} \max_{j \in S^c} \|T_{j3}\|_{l_2}^2 &\stackrel{d.}{=} \max_{j \in S^c} \sum_{k=1}^K \sigma_{jk}^2 \xi_{jk}^2 \\ &\leq \max_{j \in S^c} \max_{1 \leq k \leq K} \sigma_{jk}^2 \max_{j \in S^c} \left(\sum_{k=1}^K \xi_{jk}^2 \right) \quad \text{given} \quad \left(X_S^{(1:K)}, \vec{W}^{(1:K)} \right) \end{aligned} \quad (37)$$

We next provide a useful bound for χ^2 random variable, which was given in [29].

Lemma 3. [29] *Let Z be a central χ^2 distributed random variable with the degree d . Then for all $t > d$, we have*

$$P(Z \geq 2t) \leq \exp \left(-t \left[1 - 2\sqrt{\frac{d}{t}} \right] \right).$$

Applying the above lemma, we obtain for all $t > K$,

$$\begin{aligned} P \left(\max_{j \in S^c} \left(\sum_{k=1}^K \xi_{jk}^2 \right) > 2t \right) &\leq (p-s) P \left(\left(\sum_{k=1}^K \xi_{jk}^2 \right) > 2t \right) \\ &\leq (p-s) \exp \left(-t \left[1 - 2\sqrt{\frac{K}{t}} \right] \right) \end{aligned} \quad (38)$$

By applying the bound on σ_{jk}^2 derived in appendix E together with (38), we further have

$$\max_{j \in S^c} \|T_{j3}\|_{l_2}^2 \leq 2t \rho_u \left(\Sigma_{S^c S^c}^{(1:K)} \right) \left(\frac{\psi(B^*, \Sigma^{(1:K)})}{n} + \Gamma \right) \quad (39)$$

with the probability larger than

$$\begin{aligned} &1 - 2(K+1) \exp \left(-\frac{s}{2} \right) - 4(K+1) \exp \left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}} \right)_+^2 \right) \\ &- K \exp \left(-\log s + 2\sqrt{2 \log s} \right) - (p-s) \exp \left(-t \left[1 - 2\sqrt{\frac{K}{t}} \right] \right) \\ &- \exp \left(-5(n-s) \left[1 - 2\sqrt{\frac{1}{5}} \right] \right) \end{aligned} \quad (40)$$

for $t > K$, where

$$\Gamma = \frac{16s \|\Delta\|_{l_\infty/l_2}}{nC_{min}} (1 + 2 \|\Delta\|_{l_\infty/l_2}) + \frac{12}{C_{min}} \left(\frac{s}{n}\right)^{\frac{3}{2}} + \frac{10(n-s)\sigma_W^{(k)^2}}{n^2\lambda_n^2}. \quad (41)$$

For n large enough, Γ converges to zero with an order $o\left(\frac{s}{n}\right)$. We also note that $\psi(B^*, \Sigma^{(1:K)})$ has an order $O(s)$ based on Proposition 1. In (39), we set $t = \frac{1+v}{1+\delta} \log(p-s)$ where $v > 0$ and $\delta = v/(3v+4)$. We can then show that if

$$n > 2(1+v)\psi(B^*, \Sigma^{(1:K)}) \log(p-s) \frac{\rho_u\left(\Sigma_{S^c S^c | S}^{(1:K)}\right)}{\gamma^2},$$

then

$$\max_{j \in S^c} \|T_{j3}\|_{l_2} < \gamma \quad (42)$$

with the probability larger than

$$\begin{aligned} & 1 - 2(K+1) \exp\left(-\frac{s}{2}\right) - 4(K+1) \exp\left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}}\right)_+^2\right) \\ & - K \exp\left(-\log s + 2\sqrt{2 \log s}\right) - \exp\left(-\frac{v}{2} \log(p-s)\right) \\ & - \exp\left(-5(n-s) \left[1 - 2\sqrt{\frac{1}{5}}\right]\right). \end{aligned} \quad (43)$$

It follows from (24) that

$$\|V_j\|_{l_2} \leq \|V_{j1}\|_{l_2} + \|V_{j2}\|_{l_2} + \|V_{j3}\|_{l_2}.$$

Combining the above equation with the evaluation for T_{j1}, T_{j2}, T_{j3} , we conclude that $\|V_j\|_{l_2} < 1$.

Analysis of U_S : We have obtained the sufficient conditions for the existence and uniqueness of an optimal solution to the problem given in (6), which guarantees $\widehat{B}_{S^c} = 0$. It remains to characterize conditions such that all rows of \widehat{B}_S are nonzero and hence $S(\widehat{B})$ recovers the true support union.

In order to guarantee that every row of \widehat{B}_S is nonzero, it suffices to guarantee that

$$\|U_S\|_{l_\infty/l_2} \leq \frac{1}{2} b_{min}^*$$

where

$$U_S = \widehat{B}_S - B_S^* = \left[\vec{U}_S^{(1)} \dots \vec{U}_S^{(K)} \right].$$

Each column $\vec{U}_S^{(k)}$ is given by

$$\vec{U}_S^{(k)} := \widehat{\beta}_S^{(k)} - \vec{\beta}_S^{*(k)} = \left(\widehat{\Sigma}_{SS}^{(k)}\right)^{-1} \left(\frac{1}{n} X_S^{(k)T} \vec{W}^{(k)} - \lambda_n \widehat{Z}_{Sk}\right). \quad (44)$$

It suffices to guarantee that

$$\left\| \vec{U}_S^{(k)} \right\|_{l_\infty} \leq \frac{1}{2K} b_{min}^*,$$

for $k = 1, \dots, K$. In order to bound $\left\| \vec{U}_S^{(k)} \right\|_{l_\infty}$, we define

$$\widetilde{W}^{(k)} = \frac{1}{\sqrt{n}} \left(\widehat{\Sigma}_{SS}^{(k)}\right)^{-\frac{1}{2}} X_S^{(k)T} \vec{W}^{(k)},$$

and hence

$$\vec{U}_S^{(k)} = \frac{1}{\sqrt{n}} \left(\widehat{\Sigma}_{SS}^{(k)}\right)^{-\frac{1}{2}} \widetilde{W}^{(k)} - \lambda_n \left(\widehat{\Sigma}_{SS}^{(k)}\right)^{-1} \widehat{Z}_{Sk}.$$

We then obtain the following bound

$$\left\| \vec{U}_S^{(k)} \right\|_{l_\infty} \leq \underbrace{\left\| \frac{1}{\sqrt{n}} \left(\widehat{\Sigma}_{SS}^{(k)}\right)^{-\frac{1}{2}} \widetilde{W}^{(k)} \right\|_{l_\infty}}_{T'_{k1}} + \lambda_n \underbrace{\left\| \left(\widehat{\Sigma}_{SS}^{(k)}\right)^{-1} \widehat{Z}_{Sk} \right\|_{l_\infty}}_{T'_{k2}}. \quad (45)$$

We next evaluate the bounds on the two terms T'_{k2} and T'_{k1} , respectively.

Evaluation of T'_{k2} : We first derive the following bound

$$\begin{aligned} \left\| \left(\widehat{\Sigma}_{SS}^{(k)}\right)^{-1} \widehat{Z}_{Sk} \right\|_{l_\infty} &\leq \max_{i \in S} \sum_{j \in S} \left| \left(\widehat{\Sigma}_{SS}^{(k)}\right)^{-1} \right|_{ij} \\ &= \left\| \left(\widehat{\Sigma}_{SS}^{(k)}\right)^{-1} \right\|_{\infty} \\ &\leq \left\| \left(\Sigma_{SS}^{(k)}\right)^{-1} \right\|_{\infty} + \left\| \left(\widehat{\Sigma}_{SS}^{(k)}\right)^{-1} - \left(\Sigma_{SS}^{(k)}\right)^{-1} \right\|_{\infty} \\ &\stackrel{(a)}{\leq} D_{max} + \sqrt{s} \left\| \left(\widehat{\Sigma}_{SS}^{(k)}\right)^{-1} - \left(\Sigma_{SS}^{(k)}\right)^{-1} \right\|_2 \\ &\stackrel{(b)}{\leq} D_{max} + \frac{12s}{C_{min} \sqrt{n}} \end{aligned} \quad (46)$$

with probability larger than $1 - 2 \exp\left(-\frac{s}{2}\right) - 3 \exp\left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}}\right)_+^2\right)$. In the above derivation, step (a) follows from the assumption of the theorem and $\|A\|_{\infty} \leq \sqrt{s} \|A\|_2$ for $A \in \mathbb{R}^{s \times n}$, and step (b) applies the bound given in (95) in Appendix F. Therefore,

$$T'_{k2} \leq \lambda_n \left(D_{max} + \frac{12s}{C_{min} \sqrt{n}} \right) \quad (47)$$

with probability larger than $1 - 2 \exp\left(-\frac{s}{2}\right) - 3 \exp\left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}}\right)_+^2\right)$.

Evaluation of T'_{k1} : We first have

$$\begin{aligned} \mathbb{E} \left(\widetilde{\vec{W}}^{(k)} \widetilde{\vec{W}}^{(k)T} \middle| X_S^{(1:K)} \right) &= \mathbb{E} \left(\left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-\frac{1}{2}} \frac{1}{n} X_S^{(k)T} \widetilde{\vec{W}}^{(k)} \widetilde{\vec{W}}^{(k)T} X_S^{(k)} \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-\frac{1}{2}} \middle| X_S^{(1:K)} \right) \\ &= \sigma_W^{(k)2} I_S \end{aligned} \quad (48)$$

which implies that given $X_S^{(1:K)}$, $\widetilde{\vec{W}}^{(k)}$ has i.i.d. components with each being Gaussian distributed as $N\left(0, \sigma_w^{(k)2}\right)$. Hence, given $X_S^{(1:K)}$, we have

$$T'_{k1} \leq \left\| \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-\frac{1}{2}} \right\|_{\infty} \left\| \frac{\widetilde{\vec{W}}^{(k)}}{\sqrt{n}} \right\|_{l_{\infty}} \leq \sigma_{max} \sqrt{\frac{2s}{C_{min}}} \max_{j \in S} \sqrt{\frac{1}{n} \xi_j^2} \quad (49)$$

with probability larger than $1 - \exp\left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}}\right)_+^2\right)$, where $\sigma_{max} = \max_{1 \leq k \leq K} \sigma_W^{(k)}$, and ξ_j is the standard Gaussian random variable. The second inequality in the preceding derivation follows because $\|A\|_{\infty} \leq \sqrt{s} \|A\|_2$ for $A \in \mathbb{R}^{s \times n}$, and from the bound (93) provided in Appendix F. By applying Lemma 3 with $d = 1$, we have

$$P \left(\frac{1}{n} \max_{j \in S} \xi_j^2 \geq \frac{2t}{n} \right) \leq s \cdot \exp \left(-t \left[1 - 2\sqrt{\frac{1}{t}} \right] \right). \quad (50)$$

By setting $t = 2 \log s$ in the above bound, we then obtain

$$T'_{k1} \leq \sigma_{max} \sqrt{\frac{2s}{C_{min}}} \cdot \sqrt{\frac{2t}{n}} \leq \sqrt{\frac{8s \log(s) \sigma_{max}^2}{nC_{min}}} \quad (51)$$

with the probability larger than

$$1 - \exp \left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}} \right)_+^2 \right) - \exp \left(-\log s + 2\sqrt{2 \log s} \right). \quad (52)$$

Combining the bounds on T'_{k1} and T'_{k2} , we obtain

$$\begin{aligned} \left\| \vec{U}_S^{(k)} \right\|_{l_{\infty}} &\leq \sqrt{\frac{8s \log(s) \sigma_{max}^2}{nC_{min}}} + \lambda_n \left(D_{max} + \frac{12s}{C_{min} \sqrt{n}} \right) \\ &= \rho(n, s, \lambda_n) \end{aligned} \quad (53)$$

with the probability larger than

$$1 - 2 \exp \left(-\frac{s}{2} \right) - 4 \exp \left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}} \right)_+^2 \right) - \exp \left(-\log s + 2\sqrt{2 \log s} \right). \quad (54)$$

Thus, the assumption $\frac{\rho(n,s,\lambda_n)}{b_{min}^*} = o(1)$ guarantees that $\left\| \vec{U}_S^{(k)} \right\|_{l_\infty} \leq \frac{1}{2K} b_{min}^*$ for sufficiently large n .

Furthermore, we derive the following bound

$$\begin{aligned} \|\Delta\|_{l_\infty/l_2} &\leq \frac{\left\| \widehat{B}_S - B_S^* \right\|_{l_\infty/l_2}}{\min_{j \in S} \|B_j^*\|_2} = \frac{\|U_S\|_{l_\infty/l_2}}{b_{min}^*} \\ &\leq \frac{\max_{j \in S} \sum_{k=1}^K |U_{jk}|}{b_{min}^*} \leq \sum_{k=1}^K \frac{\max_{j \in S} |U_{jk}|}{b_{min}^*} \\ &= \sum_{k=1}^K \frac{\left\| \vec{U}_S^{(k)} \right\|_{l_\infty}}{b_{min}^*} \leq \frac{K\rho(n,s,\lambda_n)}{b_{min}^*} = o(1) \end{aligned} \quad (55)$$

with the probability larger than

$$1 - 2K \exp\left(-\frac{s}{2}\right) - 4K \exp\left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}}\right)_+^2\right) - K \exp\left(-\log s + 2\sqrt{2 \log s}\right). \quad (56)$$

Summarizing the analysis of V_{S^c} and U_S , we conclude that the multi-task Lasso problem given in (6) has a unique solution \widehat{B} , whose support union recovers the true support union $S(B^*)$ with high probability under the assumption of the theorem.

6 Proof of Theorem 2

Our proof adapts and further develops the proof techniques established in [29] due to varying design matrices across tasks.

Following from the proof in Section 5, it can be shown that if either $\left\| \widehat{Z}_{S^c} \right\|_{l_\infty/l_2} > 1$ holds or $\left\| \widehat{B} - B^* \right\|_{l_\infty/l_2} = o(b_{min}^*)$ does not hold, no solution \tilde{B} to the multi-task Lasso problem given in (6) recovers the correct support union and satisfies $\left\| \tilde{B} - B^* \right\|_{l_\infty/l_2} = o(b_{min}^*)$. Hence, if $\left\| \widehat{B} - B^* \right\|_{l_\infty/l_2} = o(b_{min}^*)$ does not hold, it is already the case that the multi-task Lasso does not provide the desired solution. Then the following proof is to identify sufficient conditions such that $\|V_{S^c}\|_{l_\infty/l_2} > 1$ when $\left\| \widehat{B} - B^* \right\|_{l_\infty/l_2} = o(b_{min}^*)$ holds, where $V_j = \left(\widehat{Z}_{j1}, \dots, \widehat{Z}_{jK} \right)$ for $j \in S^c$.

We use the decomposition in (24), which is rewritten below:

$$V_j = T_{j1} + T_{j2} + T_{j3}.$$

However, we are now interested in lower bounding $\|V_{S^c}\|_{l_\infty/l_2}$. We first bound this quantity as follows:

$$\|V_{S^c}\|_{l_\infty/l_2} \geq \|T_{S^{c3}}\|_{l_\infty/l_2} - \|T_{S^{c1}}\|_{l_\infty/l_2} - \|T_{S^{c2}}\|_{l_\infty/l_2}.$$

By the assumption of the theorem, $\|T_{S^{c1}}\|_{l_\infty/l_2} \leq 1 - \gamma$. We next consider $T_{S^{c2}}$. Due to (31), we have

$$\|T_{S^{c2}}\|_{l_\infty/l_2} \leq (1 - \gamma) \left\| \widehat{Z}_S - Z_S^* \right\|_{l_\infty/l_2} + (1 - \gamma) \mathbb{E} \left[\left\| \widehat{Z}_S - Z_S^* \right\|_{l_\infty/l_2} \middle| X_S^{(1:K)} \right]. \quad (57)$$

By the assumption that $\left\| \widehat{B} - B^* \right\|_{l_\infty/l_2} = o(b_{min}^*)$ holds, following the proof in Section 5, $\|T_{S^{c2}}\|_{l_\infty/l_2} = o(1)$ holds.

It then suffices to guarantee that $\|T_{S^{c3}}\|_{l_\infty/l_2} > 2 - \gamma$. We recall from (37) that

$$\max_{j \in S^c} \|T_{j3}\|_{l_2} \stackrel{d.}{=} \max_{j \in S^c} \sqrt{\sum_{k=1}^K \sigma_{jk}^2 \xi_{jk}^2} \quad \text{given} \quad \left(X_S^{(1:K)}, \vec{W}^{(1:K)} \right), \quad (58)$$

where $\xi_{jk} \sim \mathcal{N}(0, 1)$ are independently distributed across k .

We let $V_{max} := \|T_{S^{c3}}\|_{l_\infty/l_2}$, and the remaining part of the proof is to derive a lower bound on V_{max} , which takes several steps. The first step is to show that V_{max} is concentrated around its expectation when $\left(X_S^{(1:K)}, \vec{W}^{(1:K)} \right)$ are given.

Lemma 4. *For any $\delta > 0$,*

$$P \left[|V_{max} - \mathbb{E}V_{max}| \geq \delta \middle| X_S^{(1:K)}, \vec{W}^{(1:K)} \right] \leq 4 \exp \left(- \frac{\delta^2}{2\rho_u \left(\sum_{S^c S^c | S}^{(1:K)} \right) \max_{1 \leq k \leq K} M_k} \right). \quad (59)$$

Proof. We first construct the following function $g : \mathbb{R}^{(p-s) \times K} \rightarrow \mathbb{R}$

$$g(\xi) := \max_{j \in S^c} \left(\sqrt{\sum_{k=1}^K \sigma_{jk}^2 \xi_{jk}^2} \right)$$

where ξ_{jk} is the entry of the matrix ξ with the index pair $\{j, k\}$.

To explore the continuity property of the constructed function g , we let $u = (u_{jk}, j \in S^c, k = 1, \dots, K)$ and $v = (v_{jk}, j \in S^c, k = 1, \dots, K)$ be two matrices. We derive the

following bound given $(X_S^{(1:K)}, \vec{W}^{(1:K)})$.

$$\begin{aligned}
|g(u) - g(v)| &= \left| \max_{j \in S^c} \left(\sqrt{\sum_{k=1}^K \sigma_{jk}^2 u_{jk}^2} \right) - \max_{n \in S^c} \left(\sqrt{\sum_{k=1}^K \sigma_{nk}^2 v_{nk}^2} \right) \right| \\
&\leq \max_{j \in S^c} \left| \sqrt{\sum_{k=1}^K \sigma_{jk}^2 u_{jk}^2} - \sqrt{\sum_{k=1}^K \sigma_{jk}^2 v_{jk}^2} \right| \\
&\stackrel{(a)}{\leq} \left(\max_{j \in S^c} \max_{1 \leq k \leq K} \sigma_{jk} \right) \left(\max_{j \in S^c} \|u_j - v_j\|_2 \right) \\
&\leq \sqrt{\rho_u \left(\Sigma_{S^c S^c | S}^{(1:K)} \right)} \max_{1 \leq k \leq K} M_k \|u - v\|_F, \tag{60}
\end{aligned}$$

where (a) follows by taking square on both sides and comparing various cross terms.

Therefore, the function g is Lipschitz continuous with constant $L = \sqrt{\rho_u \left(\Sigma_{S^c S^c | S}^{(1:K)} \right)} \max_{1 \leq k \leq K} M_k$. The proof completes by applying Gaussian concentration inequality given below for a standard Gaussian vector X and the Lipschitz function g with the constant L :

$$P[|g(X) - \mathbb{E}g(X)| \geq \delta] \leq 4 \exp(-\delta^2/(2L^2)).$$

□

The second step is to find a lower bound on $\mathbb{E}[V_{max}]$.

Lemma 5. *For any fixed δ' and sufficiently large $(p - s)$, the following inequality holds:*

$$\mathbb{E} \left[V_{max} \mid X_S^{(1:K)}, \vec{W}^{(1:K)} \right] \geq \max_{1 \leq k \leq K} \sqrt{M_k} \sqrt{(1 - \delta') \rho_l \left(\Sigma_{S^c S^c | S}^{(1:K)} \right) \log(p - s)/2}.$$

Proof. The proof is under the assumption that $(X_S^{(1:K)}, \vec{W}^{(1:K)})$ are given. Define $\eta_{jk} = \sqrt{(\Sigma_{S^c S^c | S}^{(k)})_{jj}} \xi_{jk}$ and therefore, $\eta_{jk} \sim \mathcal{N} \left(0, (\Sigma_{S^c S^c | S}^{(k)})_{jj} \right)$. We then have

$$\sqrt{\sum_{k=1}^K \sigma_{jk}^2 \xi_{jk}^2} = \sqrt{\sum_{k=1}^K M_k (\Sigma_{S^c S^c | S}^{(k)})_{jj} \xi_{jk}^2} = \sqrt{\sum_{k=1}^K M_k \eta_{jk}^2} \geq \sqrt{M_{k^*}} |\eta_{jk^*}| \tag{61}$$

where $k^* = \operatorname{argmax}_{1 \leq k \leq K} \sqrt{M_k}$. Without loss of generality, let $k^* = 1$.

$$\mathbb{E} \left[V_{max} \mid X_S^{(1:K)}, \vec{W}^{(1:K)} \right] \geq \sqrt{M_{k^*}} \cdot \mathbb{E} \left(\max_{j \in S^c} |\eta_{j1}| \right) \tag{62}$$

The proof completes by applying the lower bound of $\mathbb{E}(\max_{j \in S^c} |\eta_{j1}|)$. It can be shown that

$$\mathbb{E}[(\eta_{i1} - \eta_{j1})^2] \geq \rho_l \left(\Sigma_{S^c S^c | S}^{(1:K)} \right) \mathbb{E}[(\xi_{i1} - \xi_{j1})^2].$$

Using the result in [40], we have

$$\mathbb{E} \left(\max_{j \in S^c} |\eta_{j1}| \right) \geq \frac{1}{2} \sqrt{\rho_l \left(\Sigma_{S^c S^c | S}^{(1:K)} \right)} \mathbb{E} \left(\max_{j \in S^c} |\xi_{j1}| \right)$$

Furthermore, the standard Gaussian random vector has the following bound by the result in [41]:

$$\mathbb{E} \left(\max_{j \in S^c} |\xi_{j1}| \right) \geq \sqrt{2(1 - \delta') \log(p - s)}$$

if $(p - s)$ is large enough, where δ' is a small positive number. \square

In Appendix E, we obtain the following lower bound

$$\max_{1 \leq k \leq K} M_k \geq \frac{\psi(B^*, \Sigma^{(1:K)})}{n} - \Gamma$$

with the probability larger than

$$1 - 2 \exp\left(-\frac{s}{2}\right) - 4 \exp\left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}}\right)_+^2\right) - \exp\left(-5(n - s) \left[1 - 2\sqrt{\frac{1}{5}}\right]\right). \quad (63)$$

Since Γ converges to 0 with an order $o\left(\frac{s}{n}\right)$, $\max_{1 \leq k \leq K} M_k \geq \frac{\psi(B^*, \Sigma^{(1:K)})}{n} (1 - \delta'')$ holds for any small constant $\delta'' > 0$ and large enough n . We then have

$$\begin{aligned} \mathbb{E} \left[V_{max} \middle| X_S^{(1:K)}, \vec{W}^{(1:K)} \right] &\geq \sqrt{\frac{\psi(B^*, \Sigma^{(1:K)})}{n} (1 - \delta'')} \sqrt{(1 - \delta') \rho_l \left(\Sigma_{S^c S^c | S}^{(1:K)} \right) \log(p - s) / 2} \\ &\stackrel{(a)}{\geq} (2 - \gamma) \sqrt{\frac{(1 - \delta')(1 - \delta'')}{4(1 - v)}} \\ &\stackrel{(b)}{>} 2 - \gamma + \delta \end{aligned} \quad (64)$$

with high probability, where (a) follows from the assumption of the theorem on the sample size n , and (b) follows by choosing $v > 1 - \frac{(1 - \delta')(1 - \delta'')}{4[1 + (\delta/(2 - \gamma))]^2}$.

By applying lemma 4 and $\max_{1 \leq k \leq K} M_k \leq \frac{\psi(B^*, \Sigma^{(1:K)})}{n} (1 + \delta'')$, i.e., equation (89) in Appendix E, we obtain

$$\begin{aligned} P \left[|V_{max} - \mathbb{E}V_{max}| \geq \delta \middle| X_S^{(1:K)}, \vec{W}^{(1:K)} \right] &\leq 4 \exp \left(-\frac{n\delta^2}{2\rho_u \left(\Sigma_{S^c S^c | S}^{(1:K)} \right) \psi(B^*, \Sigma^{(1:K)}) (1 + \delta'')} \right) \\ &\leq 4 \exp \left(-\frac{n\delta^2 C_{min}}{2\rho_u \left(\Sigma_{S^c S^c | S}^{(1:K)} \right) s(1 + \delta'')} \right) \end{aligned} \quad (65)$$

which implies $V_{max} > 2 - \gamma$ with high probability.

Therefore, $\|V_{S^c}\|_{l_\infty/l_2} > 1$ holds with probability larger than

$$1 - 2 \exp\left(-\frac{s}{2}\right) - 4 \exp\left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}}\right)_+^2\right) - 4 \exp\left(-\frac{n\delta^2 C_{min}}{2\rho_u \left(\Sigma_{S^c S^c|S}^{(1:K)}\right) s(1 + \delta'')}\right) \\ - \exp\left(-5(n-s) \left[1 - 2\sqrt{\frac{1}{5}}\right]\right),$$

which concludes the proof.

7 Conclusions

In this paper, we have investigated the Gaussian MVMR linear regression model. We have characterized sufficient and necessary conditions under which the l_1/l_2 -regularized multi-task Lasso guarantees successful recovery of the support union of K linear regression vectors. The two conditions are characterized by a threshold and hence are tight in the order sense. Our numerical results have demonstrated the advantage of joint recovery of the support union compared to using single-task Lasso to recover the support set of each task individually. Further studying the MVMR model under other block-constraints is an interesting topic in the future. Applications of the approach here to structure learning problems based on real data sets such as social network data are also interesting.

Appendix

A Bounds on $\psi(B^*, \Sigma^{(1:K)})$

We first derive an upper bound on $\psi(B^*, \Sigma^{(1:K)})$ as follows:

$$\begin{aligned} \psi(B^*, \Sigma^{(1:K)}) &= \max_{1 \leq k \leq K} \vec{Z}_{S_k}^{*T} \left(\Sigma_{SS}^{(k)}\right)^{-1} \vec{Z}_{S_k}^* \\ &\leq \sum_{k=1}^K \vec{Z}_{S_k}^{*T} \left(\Sigma_{SS}^{(k)}\right)^{-1} \vec{Z}_{S_k}^* \\ &\leq \sum_{k=1}^K \|\vec{Z}_{S_k}^*\|_{l_2}^2 \left\| \left(\Sigma_{SS}^{(k)}\right)^{-1} \right\|_2 \\ &\leq \frac{s}{C_{min}}. \end{aligned} \tag{66}$$

We then derive a lower bound on $\psi(B^*, \Sigma^{(1:K)})$ as follows:

$$\begin{aligned}
\psi(B^*, \Sigma^{(1:K)}) &= \max_{1 \leq k \leq K} \vec{Z}_{Sk}^{*T} \left(\Sigma_{SS}^{(k)} \right)^{-1} \vec{Z}_{Sk}^* \\
&\geq \frac{1}{K} \sum_{k=1}^K \vec{Z}_{Sk}^{*T} \left(\Sigma_{SS}^{(k)} \right)^{-1} \vec{Z}_{Sk}^* \\
&\geq \frac{1}{K} \sum_{k=1}^K \|\vec{Z}_{Sk}^*\|_{l_2}^2 \cdot \min_{\vec{x}} \frac{\vec{x}^T \left(\Sigma_{SS}^{(k)} \right)^{-1} \vec{x}}{\|\vec{x}\|_{l_2}^2} \\
&\geq \frac{s}{KC_{max}}
\end{aligned} \tag{67}$$

Therefore, $\psi(B^*, \Sigma^{(1:K)})$ is of the order of $O(s)$.

B Proof of Lemma 1

Suppose \tilde{B} is another optimal solution to the problem given in (6), then we have

$$f(\hat{B}) + \lambda_n \|\hat{B}\|_{l_1/l_2} = f(\tilde{B}) + \lambda_n \|\tilde{B}\|_{l_1/l_2}, \tag{68}$$

where $f(B) = \frac{1}{2n} \sum_{k=1}^K \left\| \vec{Y}^{(k)} - X^{(k)} \vec{\beta}^{(k)} \right\|_2^2$. It is clear that

$$\|\hat{B}\|_{l_1/l_2} = \sum_{j=1}^p \hat{Z}_j \hat{B}_j^T, \tag{69}$$

where \hat{Z}_j is the j th row of \hat{Z} and \hat{B}_j is the j th row of \hat{B} . We substitute (69) into (68) and obtain

$$f(\hat{B}) + \lambda_n \sum_{j=1}^p \hat{Z}_j \hat{B}_j^T = f(\tilde{B}) + \lambda_n \|\tilde{B}\|_{l_1/l_2}.$$

We then subtract $\lambda_n \sum_{j=1}^p \hat{Z}_j \tilde{B}_j^T$ from both sides of the above equation, and move $f(\tilde{B})$ to the left-hand-side (LHS) to obtain

$$f(\hat{B}) + \lambda_n \sum_{j=1}^p \hat{Z}_j (\hat{B}_j^T - \tilde{B}_j^T) - f(\tilde{B}) = \lambda_n \|\tilde{B}\|_{l_1/l_2} - \lambda_n \sum_{j=1}^p \hat{Z}_j \tilde{B}_j^T. \tag{70}$$

We further substitute the KKT condition $\nabla_B f(\hat{B}) + \lambda_n \hat{Z} = 0$ into (70), and obtain

$$f(\hat{B}) + \sum_{j=1}^p \nabla_{B_j} f(\hat{B}) (\tilde{B}_j^T - \hat{B}_j^T) - f(\tilde{B}) = \lambda_n \|\tilde{B}\|_{l_1/l_2} - \lambda_n \sum_{j=1}^p \hat{Z}_j \tilde{B}_j^T \tag{71}$$

Due to the convexity of $f(B)$, the LHS of the above equation is less than or equal to 0. Hence, we have

$$\|\tilde{B}\|_{l_1/l_2} \leq \sum_{j=1}^p \hat{Z}_j \tilde{B}_j^T.$$

Since $\sum_{j=1}^p \|\tilde{B}_j\|_{l_2} \geq \sum_{j=1}^p \hat{Z}_j \tilde{B}_j^T$, we obtain

$$\sum_{j=1}^p \|\tilde{B}_j\|_{l_2} = \sum_{j=1}^p \hat{Z}_j \tilde{B}_j^T.$$

Based on the assumption of the lemma, $\|\hat{Z}_j\|_{l_2} < 1$ if $j \in \Omega$. Therefore, $\|\tilde{B}_j\|_{l_2} = 0$ for $j \in \Omega$.

C Derivation of \hat{Z}_{S^c}

We write the function $f(B)$ as

$$\begin{aligned} f(B) &= \frac{1}{2n} \sum_{k=1}^K \left\| \vec{Y}^{(k)} - (X_S^{(k)}, X_{S^c}^{(k)}) \begin{pmatrix} \vec{\beta}_S^{(k)} \\ \vec{\beta}_{S^c}^{(k)} \end{pmatrix} \right\|^2 \\ &= \frac{1}{2n} \sum_{k=1}^K \left\| X_S^{(k)} \vec{\beta}_S^{(k)*} + \vec{W}^{(k)} - X_S^{(k)} \vec{\beta}_S^{(k)} - X_{S^c}^{(k)} \vec{\beta}_{S^c}^{(k)} \right\|^2 \end{aligned} \quad (72)$$

and take partial derivative over components of B to obtain

$$\frac{\partial f(B)}{\partial B_{jk}} = -\frac{1}{n} \vec{X}_j^{(k)T} \left(X_S^{(k)} \vec{\beta}_S^{(k)*} + \vec{W}^{(k)} - X_S^{(k)} \vec{\beta}_S^{(k)} - X_{S^c}^{(k)} \vec{\beta}_{S^c}^{(k)} \right),$$

where \vec{X}_j denotes the j th column of the matrix X . Hence, \hat{B}_S satisfies

$$-\frac{1}{n} X_S^{(k)T} \left(X_S^{(k)} \vec{\beta}_S^{(k)*} + \vec{W}^{(k)} - X_S^{(k)} \widehat{\vec{\beta}}_S^{(k)} \right) + \lambda_n \widehat{Z}_{S^c k} = 0$$

for $k = 1, \dots, K$, where $\widehat{Z}_{S^c k}$ denotes the k th column of \widehat{Z} with row indices in the set S , and $\widehat{Z}_S \in \partial \|\widehat{B}_S\|_{l_1/l_2}$. Furthermore, \widehat{Z}_{S^c} satisfies

$$-\frac{1}{n} X_{S^c}^{(k)T} \left(X_S^{(k)} \vec{\beta}_S^{(k)*} + \vec{W}^{(k)} - X_S^{(k)} \widehat{\vec{\beta}}_S^{(k)} \right) + \lambda_n \widehat{Z}_{S^c k} = 0$$

for $k = 1, \dots, K$, where $\widehat{Z}_{S^c} \in \partial \left\| \widehat{B}_{S^c} \right\|_{l_1/l_2}$. As we introduce the notations $\widehat{\Sigma}_{SS}^{(k)} = \frac{1}{n} X_S^{(k)T} X_S^{(k)}$ and $\widehat{\Sigma}_{S^cS}^{(k)} = \frac{1}{n} X_{S^c}^{(k)T} X_S^{(k)}$, the above two equations become

$$\widehat{\Sigma}_{SS}^{(k)} \left(\widehat{\beta}_S^{(k)} - \vec{\beta}_S^{(k)*} \right) - \frac{1}{n} X_S^{(k)T} \vec{W}^{(k)} = -\lambda_n \widehat{Z}_{Sk}, \quad (73)$$

$$\widehat{\Sigma}_{S^cS}^{(k)} \left(\widehat{\beta}_S^{(k)} - \vec{\beta}_S^{(k)*} \right) - \frac{1}{n} X_{S^c}^{(k)T} \vec{W}^{(k)} = -\lambda_n \widehat{Z}_{S^c k}, \quad (74)$$

for $k = 1, \dots, K$. We now solve $\widehat{\beta}_S^{(k)} - \vec{\beta}_S^{(k)*}$ from (73), substitute it into (74), reorganize the terms, and obtain

$$\widehat{Z}_{S^c k} = -\frac{1}{\lambda_n n} X_{S^c}^{(k)T} \left(\Pi_S^{(k)} - I_n \right) \vec{W}^{(k)} + \frac{1}{n} X_{S^c}^{(k)T} X_S^{(k)T} \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \widehat{Z}_{Sk}, \quad (75)$$

where $\Pi_S^{(k)} = \frac{X_S^{(k)} \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} X_S^{(k)T}}{n}$.

Hence, for $j \in S^c$,

$$\widehat{Z}_{jk} = -\frac{1}{\lambda_n n} \vec{X}_j^{(k)T} \left(\Pi_S^{(k)} - I_n \right) \vec{W}^{(k)} + \frac{1}{n} \vec{X}_j^{(k)T} X_S^{(k)} \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \widehat{Z}_{Sk}. \quad (76)$$

D Bound on $\|T_{j1}\|_{l_2}$

We let $A^{(k)} = \Sigma_{S^c S}^{(k)} \left(\Sigma_{SS}^{(k)} \right)^{-1}$ and $\vec{C}_{Sk} = \mathbb{E}(\widehat{\vec{Z}}_{Sk} | X_S^{(1:K)})$, and derive

$$\begin{aligned}
\|T_{j1}\|_{l_2} &= \sqrt{\sum_{k=1}^K \mathbb{E}^2 \left(\widehat{Z}_{jk} | X_S \right)} \\
&= \sqrt{\sum_{k=1}^K \left(\Sigma_{jS}^{(k)} \left(\Sigma_{SS}^{(k)} \right)^{-1} \mathbb{E} \left(\widehat{\vec{Z}}_{Sk} | X_S^{(1:K)} \right) \right)^2} \\
&= \sqrt{\sum_{k=1}^K \left(A_{jS}^{(k)} \vec{C}_{Sk} \right)^2} \\
&= \sqrt{\sum_{k=1}^K \sum_{a=1}^{|S|} A_{ja}^{(k)} C_{ak} \sum_{a'=1}^{|S|} A_{ja'}^{(k)} C_{a'k}} \\
&\leq \sqrt{\sum_{a=1}^{|S|} \sum_{a'=1}^{|S|} \sum_{k=1}^K \left| A_{ja}^{(k)} \right| \left| A_{ja'}^{(k)} \right| |C_{ak} C_{a'k}|} \\
&\leq \sqrt{\sum_{a=1}^{|S|} \sum_{a'=1}^{|S|} \max_k \left| A_{ja}^{(k)} \right| \max_k \left| A_{ja'}^{(k)} \right| \sum_{k=1}^K |C_{ak} C_{a'k}|} \\
&\leq \sqrt{\sum_{a=1}^{|S|} \sum_{a'=1}^{|S|} \max_k \left| A_{ja}^{(k)} \right| \max_k \left| A_{ja'}^{(k)} \right| \sqrt{\sum_{k=1}^K C_{ak}^2} \sqrt{\sum_{k=1}^K C_{a'k}^2}} \\
&\leq \sqrt{\sum_{a=1}^{|S|} \sum_{a'=1}^{|S|} \max_k \left| A_{ja}^{(k)} \right| \max_k \left| A_{ja'}^{(k)} \right|} \\
&= \sum_{a=1}^{|S|} \max_k \left| A_{ja}^{(k)} \right| = \sum_{a=1}^{|S|} A_{ja} \tag{77}
\end{aligned}$$

where $A_{ja} = \max_k \left| A_{ja}^{(k)} \right| = \max_k \left| \left(\Sigma_{S^c S}^{(k)} \left(\Sigma_{SS}^{(k)} \right)^{-1} \right)_{ja} \right|$.

E Bound on σ_{jk}^2

We let $\sigma_{jk}^2 = \left(\Sigma_{S^c S^c | S}^{(k)} \right)_{jj} M_k$, where

$$M_k = \frac{1}{n} \widehat{\mathbf{Z}}_{Sk}^T \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \widehat{\mathbf{Z}}_{Sk} - \frac{1}{n^2 \lambda_n^2} \widehat{\mathbf{W}}^{(k)T} \left(\Pi_S^{(k)} - I_n \right) \widehat{\mathbf{W}}^{(k)}. \quad (78)$$

We derive bounds on the term $\max_{j \in S^c} \max_{1 \leq k \leq K} \sigma_{jk}^2$. We first define

$$M_k^* := \frac{1}{n} \overrightarrow{\mathbf{Z}}_{Sk}^{*T} \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \overrightarrow{\mathbf{Z}}_{Sk}^* - \frac{1}{n^2 \lambda_n^2} \overrightarrow{\mathbf{W}}^{(k)T} \left(\Pi_S^{(k)} - I_n \right) \overrightarrow{\mathbf{W}}^{(k)}. \quad (79)$$

We also define

$$\bar{M}^* = \frac{1}{n} \overrightarrow{\mathbf{Z}}_{Sk}^{*T} \left(\Sigma_{SS}^{(k)} \right)^{-1} \overrightarrow{\mathbf{Z}}_{Sk}^* + \frac{(n-s)\sigma_W^{(k)2}}{n^2 \lambda_n^2}.$$

We then have

$$|M_k - \bar{M}^*| \leq |M_k - M_k^*| + |M_k^* - \bar{M}^*|.$$

To find upper and lower bounds on M_k , we start with

$$\bar{M}^* - |M_k - M_k^*| - |M_k^* - \bar{M}^*| \leq M_k \leq \bar{M}^* + |M_k - M_k^*| + |M_k^* - \bar{M}^*|. \quad (80)$$

We first bound

$$\begin{aligned} & |M_k^* - M_k| \\ &= \frac{1}{n} \left| \overrightarrow{\mathbf{Z}}_{Sk}^{*T} \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \overrightarrow{\mathbf{Z}}_{Sk}^* - \widehat{\mathbf{Z}}_{Sk}^T \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \widehat{\mathbf{Z}}_{Sk} \right| \\ &= \frac{1}{n} \left| \overrightarrow{\mathbf{Z}}_{Sk}^{*T} \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \left(\overrightarrow{\mathbf{Z}}_{Sk}^* - \widehat{\mathbf{Z}}_{Sk} \right) + \left(\overrightarrow{\mathbf{Z}}_{Sk}^{*T} - \widehat{\mathbf{Z}}_{Sk}^T \right) \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \left(\overrightarrow{\mathbf{Z}}_{Sk}^* + \left(\widehat{\mathbf{Z}}_{Sk} - \overrightarrow{\mathbf{Z}}_{Sk}^* \right) \right) \right| \\ &\leq \frac{1}{n} \left\| \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \right\|_2 \left\| \overrightarrow{\mathbf{Z}}_{Sk}^* - \widehat{\mathbf{Z}}_{Sk} \right\|_{l_2} \left(\left\| \overrightarrow{\mathbf{Z}}_{Sk}^* \right\|_{l_2} + \left\| \overrightarrow{\mathbf{Z}}_{Sk}^* + \left(\widehat{\mathbf{Z}}_{Sk} - \overrightarrow{\mathbf{Z}}_{Sk}^* \right) \right\|_{l_2} \right) \\ &\leq \frac{1}{n} \left\| \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \right\|_2 \left\| \overrightarrow{\mathbf{Z}}_{Sk}^* - \widehat{\mathbf{Z}}_{Sk} \right\|_{l_2} \left(2 \left\| \overrightarrow{\mathbf{Z}}_{Sk}^* \right\|_{l_2} + \left\| \widehat{\mathbf{Z}}_{Sk} - \overrightarrow{\mathbf{Z}}_{Sk}^* \right\|_{l_2} \right). \end{aligned} \quad (81)$$

In the above equations,

$$\left\| \overrightarrow{\mathbf{Z}}_{Sk}^* \right\|_{l_2} \leq \sqrt{s}.$$

Following (93) in Appendix F, we have

$$\left\| \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} \right\|_2 \leq \frac{2}{C_{min}}$$

with probability larger than $1 - \exp\left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}}\right)_+^2\right)$.

We also derive:

$$\begin{aligned}
\max_{1 \leq k \leq K} \|\vec{Z}_{Sk}^* - \widehat{\vec{Z}}_{Sk}\|_{l_2} &= \max_{1 \leq k \leq K} \sqrt{\sum_{j=1}^s (Z_{jk}^* - \widehat{Z}_{jk})^2} \\
&= \sqrt{\max_{1 \leq k \leq K} \sum_{j=1}^s (Z_{jk}^* - \widehat{Z}_{jk})^2} \leq \sqrt{\sum_{k=1}^K \sum_{j=1}^s (Z_{jk}^* - \widehat{Z}_{jk})^2} \\
&= \sqrt{\sum_{j=1}^s \sum_{k=1}^K (Z_{jk}^* - \widehat{Z}_{jk})^2} \leq \sqrt{s \max_{j \in S} \sum_{k=1}^K (Z_{jk}^* - \widehat{Z}_{jk})^2} \\
&= \sqrt{s} \max_{j \in S} \sqrt{\sum_{k=1}^K (Z_{jk}^* - \widehat{Z}_{jk})^2} = \sqrt{s} \|Z_S^* - \widehat{Z}_S\|_{l_\infty/l_2} \quad (82)
\end{aligned}$$

Hence, following from Lemma 2, we have if $\|\Delta\|_{l_\infty/l_2} < \frac{1}{2}$, then

$$\max_{1 \leq k \leq K} \|\vec{Z}_{Sk}^* - \widehat{\vec{Z}}_{Sk}\|_{l_2} \leq 4\sqrt{s} \|\Delta\|_{l_\infty/l_2}.$$

Based on the above bound, we have

$$\begin{aligned}
|M_k^* - M_k| &\leq \frac{2}{nC_{min}} \left(4\sqrt{s} \|\Delta\|_{l_\infty/l_2}\right) \left(2\sqrt{s} + 4\sqrt{s} \|\Delta\|_{l_\infty/l_2}\right) \\
&= \frac{16s \|\Delta\|_{l_\infty/l_2}}{nC_{min}} (1 + 2\|\Delta\|_{l_\infty/l_2}) \quad (83)
\end{aligned}$$

with probability larger than

$$1 - 2K \exp\left(-\frac{s}{2}\right) - (4K + 1) \exp\left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}}\right)_+^2\right) - K \exp\left(-\log s + 2\sqrt{2 \log s}\right). \quad (84)$$

We next derive a bound on $|M_k^* - \bar{M}^*|$ as follows.

$$\begin{aligned}
&|M_k^* - \bar{M}^*| \\
&\leq \frac{1}{n} \left| \vec{Z}_{Sk}^{*T} \left(\left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} - \left(\Sigma_{SS}^{(k)} \right)^{-1} \right) \vec{Z}_{Sk}^* \right| + \frac{1}{n^2 \lambda_n^2} \left| \vec{W}^{(k)T} (I_n - \Pi_s^{(k)}) \vec{W}^{(k)} - (n-s) \sigma_W^{(k)2} \right| \\
&\leq \frac{1}{n} \left\| \vec{Z}_{Sk}^{*T} \right\|_{l_2}^2 \left\| \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} - \left(\Sigma_{SS}^{(k)} \right)^{-1} \right\|_2 + \frac{1}{n^2 \lambda_n^2} \left| \vec{W}^{(k)T} (I_n - \Pi_s^{(k)}) \vec{W}^{(k)} - (n-s) \sigma_W^{(k)2} \right|.
\end{aligned}$$

In the above equation,

$$\left\| \vec{Z}_{Sk}^{*T} \right\|_{l_2}^2 \leq s.$$

Following (95) in Appendix F, we have

$$\left\| \left(\widehat{\Sigma}_{SS}^{(k)} \right)^{-1} - \left(\Sigma_{SS}^{(k)} \right)^{-1} \right\|_2 \leq \frac{12}{C_{min}} \sqrt{\frac{s}{n}}$$

with probability larger than $1 - 2 \exp\left(-\frac{s}{2}\right) - 3 \exp\left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}}\right)_+^2\right)$.

We next bound the term $\left| \vec{W}^{(k)T} (I_n - \Pi_S^{(k)}) \vec{W}^{(k)} - (n-s) \sigma_W^{(k)2} \right|$. Since $\Pi_S^{(k)}$ is a projection matrix, eigenvalues of $I_n - \Pi_S^{(k)}$ can only be 1 or 0. Thus, $\text{Tr}(I_n - \Pi_S^{(k)}) = (n-s)$ implies that if we decompose $I_n - \Pi_S^{(k)}$ into $U^T \Lambda U$ with $U^T U = I$, then Λ has $(n-s)$ of “1” and s of “0”. Moreover, $U \vec{W}^{(k)}$ is a Gaussian vector with zero mean, and $\text{E}\left(U \vec{W}^{(k)} \vec{W}^{(k)T} U^T\right) = \sigma_W^{(k)2} I_n$. Therefore, we conclude that

$$\begin{aligned} U \vec{W}^{(k)} &\stackrel{d.}{=} \vec{W}^{(k)} \\ \vec{W}^{(k)T} U^T \Lambda U \vec{W}^{(k)} &\stackrel{d.}{=} \vec{W}^{(k)T} \Lambda \vec{W}^{(k)} \stackrel{d.}{=} H \sigma_W^{(k)2} \end{aligned}$$

where $H \sim \chi_{(n-s)}^2$. We now consider the term

$$\begin{aligned} &\left| \vec{W}^{(k)T} (I_n - \Pi_S^{(k)}) \vec{W}^{(k)} - (n-s) \sigma_W^{(k)2} \right| \\ &= \left| \vec{W}^{(k)T} U^T \Lambda U \vec{W}^{(k)} - (n-s) \sigma_W^{(k)2} \right| \\ &= \left| \vec{W}^{(k)T} \Lambda \vec{W}^{(k)} - (n-s) \sigma_W^{(k)2} \right| \\ &= \left| H \sigma_W^{(k)2} - (n-s) \sigma_W^{(k)2} \right|. \end{aligned} \tag{85}$$

We derive the probability of the following event:

$$\begin{aligned} &P\left(\left| H \sigma_W^{(k)2} - (n-s) \sigma_W^{(k)2} \right| \leq 9(n-s) \sigma_W^{(k)2}\right) \\ &= P(\{H < 10(n-s)\} \cap \{H > -8(n-s)\}) \\ &= P(H < 10(n-s)). \end{aligned} \tag{86}$$

Following from Lemma 3, we have

$$P(H \geq 10(n-s)) \leq \exp\left(-5(n-s) \left[1 - 2\sqrt{\frac{1}{5}}\right]\right).$$

It then follows that

$$\left| \vec{W}^{(k)T} (I_n - \Pi_S^{(k)}) \vec{W}^{(k)} - (n-s) \sigma_W^{(k)2} \right| \leq 9(n-s) \sigma_W^{(k)2}$$

with probability larger than

$$1 - \exp\left(-5(n-s) \left[1 - 2\sqrt{\frac{1}{5}}\right]\right).$$

To summarize,

$$|M_k^* - \bar{M}^*| \leq \frac{12}{C_{min}} \left(\frac{s}{n}\right)^{\frac{3}{2}} + \frac{9(n-s)\sigma_W^{(k)2}}{n^2\lambda_n^2}$$

with probability larger than

$$1 - 2 \exp\left(-\frac{s}{2}\right) - 3 \exp\left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}}\right)_+^2\right) - \exp\left(-5(n-s) \left[1 - 2\sqrt{\frac{1}{5}}\right]\right).$$

Therefore,

$$\begin{aligned} |M_k - \bar{M}^*| &\leq |M_k - M_k^*| + |M_k^* - \bar{M}^*| \\ &\leq \frac{16s \|\Delta\|_{l_\infty/l_2}}{nC_{min}} (1 + 2 \|\Delta\|_{l_\infty/l_2}) + \frac{12}{C_{min}} \left(\frac{s}{n}\right)^{\frac{3}{2}} + \frac{9(n-s)\sigma_W^2}{n^2\lambda_n^2} \end{aligned} \quad (87)$$

with high probability.

To simplify the result, we define the following quantity

$$\Gamma := \frac{16s \|\Delta\|_{l_\infty/l_2}}{nC_{min}} (1 + 2 \|\Delta\|_{l_\infty/l_2}) + \frac{12}{C_{min}} \left(\frac{s}{n}\right)^{\frac{3}{2}} + \frac{10(n-s)\sigma_W^2}{n^2\lambda_n^2} \quad (88)$$

and our bounds on M_k can be expressed as

$$\frac{1}{n} \vec{Z}_{s_k}^{*T} \left(\Sigma_{SS}^{(k)}\right)^{-1} \vec{Z}_{s_k}^* - \Gamma \leq M_k \leq \frac{1}{n} \vec{Z}_{s_k}^{*T} \left(\Sigma_{SS}^{(k)}\right)^{-1} \vec{Z}_{s_k}^* + \Gamma.$$

Using the definition of $\psi(B^*, \Sigma^{(1:K)})$, we have

$$\frac{\psi(B^*, \Sigma^{(1:K)})}{n} - \Gamma \leq \max_{1 \leq k \leq K} M_k \leq \frac{\psi(B^*, \Sigma^{(1:K)})}{n} + \Gamma \quad (89)$$

with probability larger than

$$\begin{aligned} &1 - 2(K+1) \exp\left(-\frac{s}{2}\right) - 4(K+1) \exp\left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}}\right)_+^2\right) \\ &- \exp\left(-5(n-s) \left[1 - 2\sqrt{\frac{1}{5}}\right]\right) - K \exp\left(-\log s + 2\sqrt{2 \log s}\right). \end{aligned} \quad (90)$$

F Bounds on Spectral Norms

In this section, we provide some useful bounds on spectral norms. Detailed proof can be found in [29].

Let $U \in R^{n \times s}$ be a random matrix with i.i.d. entries, and each entry has a Gaussian distribution with zero mean and unit variance.

The bound for $\left\| \frac{1}{n} U^T U \right\|_2$:

$$P \left(\left\| \frac{1}{n} U^T U \right\|_2 \leq \frac{1}{2} \right) \leq \exp \left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}} \right)_+^2 \right).$$

The bound for $\left\| \frac{1}{n} U^T U - I_{s \times s} \right\|_2$:

$$P \left(\left\| \frac{1}{n} U^T U - I_{s \times s} \right\|_2 \geq 6\sqrt{\frac{s}{n}} \right) \leq 2 \exp \left(-\frac{s}{2} \right) + 2 \exp \left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}} \right)_+^2 \right) \quad (91)$$

Let $X = U\sqrt{\Sigma}$ where $\Sigma \in R^{s \times s}$ is positive definite. Then $X \in R^{n \times s}$ has i.i.d. rows, and each row X_i is a Gaussian vector with the distribution $\mathcal{N}(0, \Sigma)$. Suppose the eigenvalues of Σ are in the interval $[C_{min}, C_{max}]$, where C_{min} and C_{max} are both positive. We next provide the bounds on several spectral norms.

The bound for $\left\| \left(\frac{U^T U}{n} \right)^{-1} \right\|_2$:

$$P \left(\left\| \left(\frac{U^T U}{n} \right)^{-1} \right\|_2 \leq 2 \right) \geq 1 - \exp \left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}} \right)_+^2 \right). \quad (92)$$

The bound for $\left\| \left(\frac{X^T X}{n} \right)^{-1} \right\|_2$:

$$P \left(\left\| \left(\frac{X^T X}{n} \right)^{-1} \right\|_2 \leq \frac{2}{C_{min}} \right) \geq 1 - \exp \left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}} \right)_+^2 \right). \quad (93)$$

The bound for $\left\| \frac{X^T X}{n} - \Sigma \right\|_2$:

$$P \left(\left\| \frac{X^T X}{n} - \Sigma \right\|_2 \leq 6C_{max}\sqrt{\frac{s}{n}} \right) \geq 1 - 2 \exp \left(-\frac{s}{2} \right) - 2 \exp \left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}} \right)_+^2 \right). \quad (94)$$

The bound for $\left\| \left(\frac{X^T X}{n} \right)^{-1} - \Sigma^{-1} \right\|_2$:

$$P \left(\left\| \left(\frac{X^T X}{n} \right)^{-1} - \Sigma^{-1} \right\|_2 \leq \frac{12}{C_{min}} \sqrt{\frac{s}{n}} \right) \geq 1 - 2 \exp \left(-\frac{s}{2} \right) - 3 \exp \left(-\frac{n}{2} \left(\frac{1}{4} - \sqrt{\frac{s}{n}} \right)_+^2 \right). \quad (95)$$

Acknowledgements

The authors would like to thank Dr. Junming Yin and Dr. Mladen Kolar at Carnegie Mellon University for their helpful discussions.

References

- [1] S. Negahban and M. J. Wainwright. Simultaneous support recovery in high dimensions: Benefits and perils of block ℓ_1/ℓ_∞ -regularization. *IEEE Trans. Inform. Theory*, 57(6):3841–3863, 2011.
- [2] A. Jalali, P. Ravikumara, S. Sanghavi, and C. Ruan. A dirty model for multi-task learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2010.
- [3] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24(2):227–234, 1995.
- [4] R. Tibshirani. Regression shrinkage and selection via the Lasso. *J. Roy. Statist. Soc. Ser. B*, 58:267–288, 1996.
- [5] S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, 1998.
- [6] D. Donoho and X. Huo. Uncertainty principles and ideal atomic decomposition. *IEEE Trans. Inform. Theory*, 47(7):2845–2862, 2001.
- [7] P. J. Bickel, Y. Ritov, and A. B. Tsybakov. Simultaneous analysis of Lasso and dantzig selector. *Annals of Statistics*, 37(4):1705–1732, 2009.
- [8] E. Candes and T. Tao. The dantzig selector: Statistical estimation when p is much larger than n . *Annals of Statistics*, 35(6):2313–2351, 2007.
- [9] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski. Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning*, 4(1):1–106, Now Publishers, Hanover, MA, USA, 2011.
- [10] E. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52(2):489–509, 2006.
- [11] M. Elad and A. M. Bruckstein. A generalized uncertainty principle and sparse representation in pairs of bases. *IEEE Trans. Inform. Theory*, 48(9):2558–2567, 2002.
- [12] A. Feuer and A. Nemirovski. On sparse representation in pairs of bases. *IEEE Trans. Inform. Theory*, 49(6):1579–1581, 2003.
- [13] D. M. Malioutov, M. Cetin, and A. S. Willsky. Optimal sparse representations in general overcomplete bases. In *Proc. Int. Conf. Acoustics, Speech, and Signal Processing*, Montreal, Canada, 2004.
- [14] J. Tropp. Greedy is good: Algorithm results for sparse approximation. *IEEE Trans. Inform. Theory*, 50(10):2231–2242, 2004.
- [15] J. J. Fuchs. Recovery of exact sparse representations in the presence of noise. *IEEE Trans. Inform. Theory*, 51(10):3601–3608, 2005.
- [16] P. Zhao and B. Yu. On model selection consistency of Lasso. *Journal of Machine Learning Research*, 7:2541–2567, 2006.

- [17] N. Meinshausen and B. Yu. Lasso-type recovery of sparse representations for high-dimensional data. *Annals of Statistics*, 37(1):246–270, 2009.
- [18] M. J. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using l_1 -constrained quadratic programming (Lasso). *IEEE Trans. Inform. Theory*, 55(5):2183–2202, 2009.
- [19] G. Raskutti, M. J. Wainwright, and B. Yu. Restricted eigenvalue properties for correlated Gaussian designs. *Journal of Machine Learning Research*, 11:2241–2259, 2010.
- [20] F. Bach. Consistency of trace norm minimization. *Journal of Machine Learning Research*, 9:1019–1048, 2008.
- [21] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused Lasso. *Journal of the Royal Statistical Society: Series B*, 67:91–108, 2005.
- [22] Y. Chen and A. Dalalyan. Fused sparsity and robust estimation for linear models with unknown variance. In *Advances in Neural Information Processing Systems (NIPS)*, volume 25, pages 1268–1276, 2012.
- [23] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B*, 68:49–67, 2006.
- [24] J. Huang and T. Zhang. The benefit of group sparsity. *Annals of Statistics*, 38(4):1978–2004, 2010.
- [25] L. Jacob, G. Obozinski, and J.-P. Vert. Group Lasso with overlaps and graph Lasso. In *International Conference on Machine Learning (ICML)*, 2009.
- [26] P. Zhao, G. Rocha, and B. Yu. The composite absolute penalties family for grouped and hierarchical variable selection. *Annals of Statistics*, 37(6A):3468–3497, 2009.
- [27] F. Bach. Consistency of the group Lasso and multiple kernel learning. *Journal of Machine Learning Research*, 9:1179–1225, 2008.
- [28] S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, 2012.
- [29] G. Obozinski, M. J. Wainwright, and M. I. Jordan. Support union recovery in high-dimensional multivariate regression. *Annals of Statistics*, 39(1):1–47, 2011.
- [30] H. Liu and J. Zhang. On the $l_1 - l_q$ regularized regression. *arXiv:0802.1517v1*, 2008.
- [31] B. A. Turlach, W. N. Venables, and S. J. Wright. Simultaneous variable selection. *Technometrics*, 47(3):349–363, 2005.
- [32] G. Obozinski, B. Tarskar, and M. I. Jordan. Joint covariate selection and joint subspace selection for multiple classification problems. *Statistics and Computing*, 20(2):231–252, 2010.
- [33] A. Argyriou, T. Evgeniou, and M. Pontil. Multi-task feature learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2006.

- [34] M. Kolar, J. Lafferty, and L. Wasserman. Union support recovery in multi-task learning. *Journal of Machine Learning Research*, 12:2415–2435, 2011.
- [35] Y. Kim, J. Kim, and Y. Kim. Blockwise sparse regression. *Statistica Sinica*, 16:375–390, 2006.
- [36] K. Lounici, M. Pontil, S. Geer, and A. B. Tsybakov. Oracle inequalities and optimal inference under group sparsity. *Annals of Statistics*, 39:2164–2204, 2011.
- [37] N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the Lasso. *Annals of Statistics*, 34(3):1436–1462, 2006.
- [38] P. Danaher, P. Wang, and D. M. Witten. The joint graphical Lasso for inverse covariance estimation across multiple classes. *arXiv1111.0324*, 2011.
- [39] J. Guo, E. Levina, G. Michailidis, and J. Zhu. Joint estimation of multiple graphical models. *Biometrika*, 98(1):1–15, 2011.
- [40] G. Samorodnitsky and M. S. Taqqu. Stochastic monotonicity and slepian-type inequalities for infinitely divisible and stable random vectors. *Annals of Statistics*, 21(1):143–160, 1993.
- [41] M. Ledoux and M. Talagrand. *Probability in Banach Spaces*. Springer-Verlag, Berlin Heidelberg, Germany, 1999.