
Distributed Sparse Multicategory Discriminant Analysis

Hengchao Chen
University of Toronto

Qiang Sun
University of Toronto

Abstract

This paper proposes a convex formulation for sparse multicategory linear discriminant analysis and then extend it to the distributed setting when data are stored across multiple sites. The key observation is that for the purpose of classification it suffices to recover the discriminant subspace which is invariant to orthogonal transformations. Theoretically, we establish statistical properties ensuring that the distributed sparse multicategory linear discriminant analysis performs as good as the centralized version after a few rounds of communications. Numerical studies lend strong support to our methodology and theory.

1 INTRODUCTION

Classification aims to assign data points to the correct classes. One popular multicategory classification method is the Fisher’s linear discriminant analysis (LDA), which has appealing performances in many applications (Hand, 2006; Michie et al., 1994). Similar to the principal component analysis (PCA), a well-known unsupervised dimensionality reduction technique, LDA can also serve as an efficient supervised dimensionality reduction tool. However, when the number of features is larger than the number of observations, which prevails in modern datasets, LDA performs poorly due to the diverging spectra (Bickel and Levina, 2008) and the noise accumulation (Fan and Fan, 2008). In the moderately high dimensional setting, the naive Bayes, which assumes conditional independence between features, is shown to outperform the Fisher’s rule (Bickel and Levina, 2004), while in the ultra-high dimensional setting, any classifier using all the predictors will be no better than random

guessing due to the noise accumulation, even if the true covariance matrix is an identity matrix (Fan and Fan, 2008). Moreover, as pointed out by Witten and Tibshirani (2011), the Fisher’s LDA classifier is not interpretable when the discriminant vectors have no particular structure.

In recent years, many high-dimensional extensions of LDA have been proposed. For binary classification problems, the linear programming discriminant (Cai and Liu, 2011), the regularized optimal affine discriminant (Fan et al., 2012), and the direct sparse discriminant analysis (Mai et al., 2012) are three popular sparse linear discriminant analysis (SLDA) methods. All these three methods assume some sparsity assumptions on the discriminant directions. Another SLDA method is the thresholding linear discriminant analysis (Shao et al., 2011), which assumes sparsity conditions on the common covariance matrix and the mean difference vector. It is non-trivial to extend these binary SLDA methods to the multicategory case. For multicategory classification problems, Qiao et al. (2009) proposed to solve a penalized least squares problem with the lasso penalty, Clemmensen et al. (2011) proposed the sparse optimal scoring, and Witten and Tibshirani (2011) proposed the ℓ_1 penalized LDA. All of these methods involve solving a non-convex optimization problem and thus can be computationally expensive. Moreover, it is unclear how to extend these methods to a distributed setting when data are possibly stored in multiple sites. Recently, Mai et al. (2019) proposed to directly estimate the Bayes rule via a composite-type loss function. Safo and Ahn (2016) proposed to estimate the discriminant vectors by applying a basis approach to the methods studied in Cai and Liu (2011) and Shao et al. (2011). Gaynanova et al. (2016) proposed to extract the sparse discriminant vectors simultaneously through a convex programming problem. The underlying idea of these three methods are the same: for the purpose of classification, it suffices to accurately estimate the discriminant subspace, i.e., the subspace spanned by the discriminant vectors. In this paper, we will exploit this observation and propose an $\ell_{1,1}$ regularized multicategory linear discriminant analysis method, which can also esti-

mate the sparse discriminant vectors simultaneously. It is worth mentioning that we choose the $\ell_{1,1}$ penalty instead of the $\ell_{2,1}$ penalty in Mai et al. (2019) and Gaynanova et al. (2016), so that we can easily extend the multicategory LDA method to a distributed setting and establish the corresponding statistical properties.

In addition to challenges posed by the high dimensionality of many datasets, with the rapid developments of science and technology, we have seen more and more datasets that are often scattered across distant servers, possibly due to the limitation of storage resources. The difficulties of fusing or aggregating these datasets due to the communication cost and privacy concerns have inspired many works in communication-efficient statistical learning (Lee et al., 2017; Wang et al., 2017; Fan et al., 2019; Jordan et al., 2019). One popular distributed estimation framework is the one-shot divide-and-conquer algorithm, in which one first computes the local estimators and then obtains the distributed estimator as an average of the local ones (Lee et al., 2017; Fan et al., 2019). These averaging-based approaches suffer from at least three drawbacks: non-diminishing bias when local estimators are biased, allowing only a limited number of local machines (much smaller than \sqrt{N} , where N is the total sample size), and poor performance in nonlinear problems (Jordan et al., 2019). To overcome these issues, Jordan et al. (2019) and Wang et al. (2017) studied a multi-round framework for distributed learning, which is referred to as the communication-efficient surrogate likelihood (CSL) framework by Jordan et al. (2019).

Although distributed classification has attracted some attention (Kokiopoulou and Frossard, 2010; Wang et al., 2019; Lian and Fan, 2018), few research has been done on the distributed sparse discriminant analysis. For binary classification, Tian and Gu (2017) proposed a divide-and-conquer sparse LDA method, in which they averaged the debiased local estimators of the discriminant direction and then sparsified the aggregated estimator. Moreover, they showed that their distributed sparse LDA method can achieve the same performance as the centralized one which used all samples. However, it is not straightforward to generalize their approach to the multicategory case. To handle the multicategory case, in this paper, we propose an $\ell_{1,1}$ regularized multicategory SLDA formulation and then extend it to the distributed setting via the CSL framework. Theoretically, we show that the distributed multicategory SLDA (dmSLDA) method performs as good as the centralized SLDA after a few rounds of communications. Moreover, we conduct numerical experiments to further support our methodology and theory.

The rest of this paper proceeds as follows. In Section

2, we review the multicategory sparse discriminant analysis and propose an $\ell_{1,1}$ regularized multicategory SLDA method. We then extend it to the distributed setting and propose the dmSLDA method in Section 3. Section 4 establishes the $\ell_{2,2}$ and $\ell_{1,1}$ estimation error bounds of dmSLDA under certain conditions and explain how the proposed distributed estimator achieves the same performance as the centralized version. To further back up the methodology and theory, we conduct several numerical experiments in Section 5. In Section 6, we conclude this paper with some discussions.

Notation. We summarize some notations that will be used throughout this paper. By convention, we use regular letters for scalars and bold letters for both vectors and matrices. We employ $[M]$ to represent the set $\{1, \dots, M\}$ for any positive integer M and $[(a, b)]$ to abbreviate the set $\{(1, 1), (1, 2), \dots, (a, b)\}$ for any positive integers a and b . For a vector $\mathbf{u} = (u_1, \dots, u_d)^\top \in \mathbb{R}^d$, we denote its ℓ_q norm by $\|\mathbf{u}\|_q = (\sum_{i=1}^d |u_i|^q)^{1/q}$, its ℓ_0 norm by $\|\mathbf{u}\|_0 = \sum_{i=1}^d 1_{u_i \neq 0}$, and its ℓ_∞ norm by $\|\mathbf{u}\|_\infty = \max_{1 \leq i \leq d} |u_i|$. Moreover, for any two vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^d$, we use $\langle \mathbf{u}, \mathbf{v} \rangle = \mathbf{u}^\top \mathbf{v}$ to denote the Euclidean inner product of \mathbf{u} and \mathbf{v} . For a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, we define the $\ell_{p,q}$ norm of \mathbf{A} as $\|\mathbf{A}\|_{p,q} = \|(\|\mathbf{A}_1\|_p, \dots, \|\mathbf{A}_m\|_p)\|_q$, where \mathbf{A}_i is the i th row of \mathbf{A} . For any two matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{m \times n}$, we denote $\langle \mathbf{A}, \mathbf{B} \rangle = \text{Tr}(\mathbf{A}^\top \mathbf{B})$. For a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, we use $\text{span}(\mathbf{A})$ to denote the subspace spanned by the columns of \mathbf{A} . For two sequences of real numbers $\{a_n\}_{n \geq 1}$ and $\{b_n\}_{n \geq 1}$, we write $a_n \lesssim b_n$ if $a_n \leq C b_n$ for some constant $C > 0$ independent of n .

2 MULTICATEGORY SLDA

Let $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$ be a collection of N i.i.d. samples with $y_i \in [K]$ drawn from a discrete distribution $\mathbb{P}(y_i = k) = \pi_k$ and $\mathbf{x}_i \in \mathbb{R}^d$ drawn from a distribution satisfying $\mathbb{E}[\mathbf{x}_i | y_i = k] = \boldsymbol{\mu}_k$ and $\text{Var}(\mathbf{x}_i | y_i = k) = \boldsymbol{\Sigma}$, where K is the number of classes, π_k is the prior probability of class k such that $\pi_k \geq 0$ and $\sum_{k=1}^K \pi_k = 1$, and $\boldsymbol{\mu}_k$ and $\boldsymbol{\Sigma}$ are the class-conditional mean and covariance matrix of \mathbf{x}_i in class k , respectively. The K centroids lie in an affine space of dimension $r_0 \leq K - 1$, which can be much smaller than d (Hastie et al., 2009). To exploit the underlying low dimensional structure, Fisher proposed to solve a set of discriminant vectors $\{\mathbf{v}_k^*\}_{k=1}^q \subset \mathbb{R}^d$ by maximizing the following sequence of Rayleigh quotients under orthogonal constraints,

$$\begin{aligned} \mathbf{v}_k^* &= \underset{\mathbf{v}_k \in \mathbb{R}^d}{\text{argmax}} \frac{\mathbf{v}_k^\top \mathbf{B} \mathbf{v}_k}{\mathbf{v}_k^\top \boldsymbol{\Sigma} \mathbf{v}_k}, \\ \text{s.t. } \mathbf{v}_k^{*\top} \boldsymbol{\Sigma} \mathbf{v}_j^* &= 0, \quad \forall 1 \leq j < k \leq q, \end{aligned} \quad (2.1)$$

where $\mathbf{B} = \text{Var}(\mathbb{E}[\mathbf{x}|y])$ is the between-class covariance and $\mathbf{\Sigma} = \mathbb{E}[\text{Var}(\mathbf{x}|y)]$ is the within-class covariance. In other words, Fisher's proposal seeks a low dimensional projection of samples that maximizes the between-class variation relative to the within-class variation. Classification can then be performed on the projected data $\mathbf{V}_q^\top \mathbf{x} \in \mathbb{R}^q$, where $\mathbf{V}_q = (\mathbf{v}_1^*, \dots, \mathbf{v}_q^*) \in \mathbb{R}^{d \times q}$. In the classical setting where $N \gg d$, the unknown between-class covariance \mathbf{B} and within-class covariance $\mathbf{\Sigma}$ in (2.1) are substituted in practice by their corresponding sample versions $\widehat{\mathbf{B}}$ and $\widehat{\mathbf{\Sigma}}$ respectively, where

$$\begin{aligned} \widehat{\mathbf{B}} &= \sum_{k=1}^K \widehat{\pi}_k (\widehat{\boldsymbol{\mu}}_k - \widehat{\boldsymbol{\mu}}) (\widehat{\boldsymbol{\mu}}_k - \widehat{\boldsymbol{\mu}})^\top, \\ \widehat{\mathbf{\Sigma}} &= \frac{1}{N-K} \sum_{k=1}^K \sum_{y_i=k} (\mathbf{x}_i - \widehat{\boldsymbol{\mu}}_k) (\mathbf{x}_i - \widehat{\boldsymbol{\mu}}_k)^\top, \end{aligned} \quad (2.2)$$

in which $\widehat{\boldsymbol{\mu}}_k = \frac{1}{b_k} \sum_{y_i=k} \mathbf{x}_i$ is the sample mean of class k , $b_k = \sum_{i=1}^N 1_{y_i=k}$ is the sample size of class k , $\widehat{\pi}_k = b_k/N$ is the sample proportion of class k , and $\widehat{\boldsymbol{\mu}} = \sum_{k=1}^K \widehat{\pi}_k \widehat{\boldsymbol{\mu}}_k = \frac{1}{N} \sum_i \mathbf{x}_i$ is the overall sample mean.

In high dimensions when $d > N$, the Fisher's approach (2.1) suffers from the singularity of $\widehat{\mathbf{\Sigma}}$, the noise accumulation, and the lack of interpretability (Bickel and Levina, 2004; Fan and Fan, 2008; Witten and Tibshirani, 2011). To overcome these issues, Witten and Tibshirani (2011) and Clemmensen et al. (2011) proposed sparse discriminant analysis, which requires solving a sequence of non-convex optimization problems and thus can be computationally intractable. On the other hand, Mai et al. (2019), Safo and Ahn (2016) and Gaynanova et al. (2016) proposed to compute any set of basis vectors of the discriminant subspace, i.e., $\mathcal{V}_q := \text{span}(\mathbf{V}_q)$ and then sparsified these basis vectors. Though not directly reported, the core of these ideas is the following proposition.

Proposition 2.1. *Suppose $\mathbf{V}, \mathbf{W} \in \mathbb{R}^{d \times q}$ span the same subspace, then the LDA classification rule based on $\mathbf{V}^\top \mathbf{x}$ is equivalent to the LDA classification rule based on $\mathbf{W}^\top \mathbf{x}$.*

Proposition 2.1 is a generalization of Proposition 4 in Gaynanova et al. (2016), implying that for the purpose of classification, it suffices to find any set of basis vectors of the discriminant subspace.

Proposition 2.2. *Assume that the within-class covariance $\mathbf{\Sigma}$ is nonsingular and define $\mathbf{W}^* \in \mathbb{R}^{d \times (K-1)}$ as the solution to the following optimization problem*

$$\mathbf{W}^* = \underset{\mathbf{W} \in \mathbb{R}^{d \times (K-1)}}{\text{argmin}} \left\{ \frac{1}{2} \langle \mathbf{W}, \mathbf{\Sigma} \mathbf{W} \rangle - \langle \mathbf{W}, \mathbf{U} \rangle \right\}, \quad (2.3)$$

where $\mathbf{U} = (\boldsymbol{\mu}_1 - \bar{\boldsymbol{\mu}}, \dots, \boldsymbol{\mu}_{K-1} - \bar{\boldsymbol{\mu}}) \in \mathbb{R}^{d \times (K-1)}$ and $\bar{\boldsymbol{\mu}} = \sum_{k=1}^K \pi_k \boldsymbol{\mu}_k$. Then we have $\mathcal{W} := \text{span}(\mathbf{W}^*) = \mathcal{V}_{\text{rank}(\mathbf{B})}$, where \mathbf{B} is the between-class covariance.

Proposition 2.2 gives an explicit choice of the basis vectors of the discriminant subspace, simpler than that in Mai et al. (2019) or Gaynanova et al. (2016). Since $\mathbf{\Sigma}$ and \mathbf{U} in (2.3) are unknown in practice, we will substitute them by their corresponding sample versions $\widehat{\mathbf{\Sigma}}$ given by (2.2) and $\widehat{\mathbf{U}} = (\widehat{\boldsymbol{\mu}}_1 - \widehat{\boldsymbol{\mu}}, \dots, \widehat{\boldsymbol{\mu}}_{K-1} - \widehat{\boldsymbol{\mu}})$, respectively. In high dimensions, we assume element-wise sparsity on \mathbf{W}^* and propose the following $\ell_{1,1}$ regularized estimator

$$\begin{aligned} \widehat{\mathbf{W}} &= \underset{\mathbf{W} \in \mathbb{R}^{d \times (K-1)}}{\text{argmin}} \left\{ \frac{1}{2} \langle \mathbf{W}, \widehat{\mathbf{\Sigma}} \mathbf{W} \rangle - \langle \mathbf{W}, \widehat{\mathbf{U}} \rangle \right. \\ &\quad \left. + \lambda \|\mathbf{W}\|_{1,1} \right\}, \end{aligned} \quad (2.4)$$

where $\lambda > 0$ is a tuning parameter. We choose the $\ell_{1,1}$ penalty (or the element-wise sparsity) rather than the $\ell_{2,1}$ penalty (or the row sparsity) used in Mai et al. (2019) and Gaynanova et al. (2016) because only with the $\ell_{1,1}$ penalty shall we extend the multicategory SLDA method to the distributed setting easily. Moreover, the assumption of element-wise sparsity is a weaker assumption than the row sparsity. Finally, to perform classification, one can apply LDA to the projected data $\widehat{\mathbf{W}}^\top \mathbf{x} \in \mathbb{R}^{K-1}$.

Remark 2.3. *The optimization problem (2.4) and similarly (3.1) and (3.2) in the next section can be solved efficiently using the fast iterative shrinkage-thresholding algorithm (FISTA) proposed by Beck and Teboulle (2009).*

3 DISTRIBUTED MULTICATEGORY SLDA

In this section, we extend the proposed $\ell_{1,1}$ regularized multicategory SLDA method to the distributed setting. Suppose there are total of M machines and on machine $m \in [M]$, n samples¹ $\{(\mathbf{x}_i^m, y_i^m)\}_{i=1}^n$ are drawn independently from the distribution specified in Section 2. Moreover, samples on different machines are assumed to be independent. To recover the discriminant subspace, the most straightforward yet unrealistic approach is to aggregate data from all machines together and then apply the $\ell_{1,1}$ regularized multicategory SLDA method based on the full sample. Another approach is to gather the local statis-

¹Here we assume for simplicity that the sample sizes across different machines are equal. In general, our framework can allow different sample sizes on each machine as long as the sample sizes are approximately balanced; that is the proportion of any single class can not vanish to zero.

tics $\widehat{\Sigma}^m$ and \widehat{U}^m from all machines and then estimate \mathbf{W}^* using (2.4) with the aggregated statistics $\widehat{\Sigma} = \frac{1}{M} \sum_{m=1}^M \widehat{\Sigma}^m$ and $\widehat{U} = \frac{1}{M} \sum_{m=1}^M \widehat{U}^m$. Although these two centralized methods usually give the best estimation and prediction performance, they are communication expensive, with communication cost $\mathcal{O}(ndM)$ and $\mathcal{O}(d^2M)$, respectively. This motivates us to design a communication-efficient algorithm that achieves the same performance as the centralized SLDA.

We apply the communication-efficient surrogate likelihood framework (Jordan et al., 2019) to the multicategory SLDA method (2.4) and obtain the distributed multicategory SLDA method (dmSLDA). The full algorithm is collected in Algorithm 1. Specifically, we

Algorithm 1: Distributed multicategory SLDA (dmSLDA).

Input: $\{(\mathbf{x}_i^m, y_i^m)\}_{i=1}^{n_m}, T$

Output: $\widehat{\mathbf{W}}$

- 1 compute $\widehat{\Sigma}^m$ and \widehat{U}^m for all m ;
 - 2 compute $\widehat{\mathbf{W}}_0$ as the solution to (3.1);
 - 3 **for** t in $1 : T$ **do**
 - 4 send $\widehat{\mathbf{W}}_{t-1}$ to local machines;
 - 5 compute $\nabla \mathcal{L}_m(\widehat{\mathbf{W}}_{t-1})$ on local machines;
 - 6 send $\nabla \mathcal{L}_m(\widehat{\mathbf{W}}_{t-1})$ back to the master machine;
 - 7 compute $\widehat{\mathbf{W}}_t$ as the solution to (3.2);
 - 8 **end**
 - 9 **return** $\widehat{\mathbf{W}} = \widehat{\mathbf{W}}_{t^*}$ with $t^* = \operatorname{argmin}_{0 \leq t \leq T} \widehat{e}_t$;
-

first compute the local estimators $\widehat{\Sigma}^m$ and \widehat{U}^m . Then we compute the initial estimator $\widehat{\mathbf{W}}_0$ on the master machine, which is set to be the first local machine without loss of generality, as the solution to the following $\ell_{1,1}$ regularized optimization problem,

$$\widehat{\mathbf{W}}_0 = \operatorname{argmin}_{\mathbf{W} \in \mathbb{R}^{d \times (K-1)}} \{\mathcal{L}_1(\mathbf{W}) + \lambda_0 \|\mathbf{W}\|_{1,1}\}, \quad (3.1)$$

where $\lambda_0 > 0$ is a tuning parameter and $\mathcal{L}_m(\mathbf{W})$ is the loss function on machine m given by

$$\mathcal{L}_m(\mathbf{W}) = \frac{1}{2} \langle \mathbf{W}, \widehat{\Sigma}^m \mathbf{W} \rangle - \langle \mathbf{W}, \widehat{U}^m \rangle.$$

At the t th iteration, we send the $(t-1)$ th estimator $\widehat{\mathbf{W}}_{t-1}$ to local machines, compute the gradients $\nabla \mathcal{L}_m(\widehat{\mathbf{W}}_{t-1})$ of the m th loss function $\mathcal{L}_m(\cdot)$ at $\widehat{\mathbf{W}}_{t-1}$, and then send these gradients back to the master machine. Then on the master machine, we compute the t th estimator by minimizing the following shifted $\ell_{1,1}$

regularized objective function,

$$\begin{aligned} \widehat{\mathbf{W}}_t = & \operatorname{argmin}_{\mathbf{W} \in \mathbb{R}^{d \times (K-1)}} \left\{ \mathcal{L}_1(\mathbf{W}) + \lambda_t \|\mathbf{W}\|_{1,1} \right. \\ & \left. + \left\langle \frac{1}{M} \sum_{m=1}^M \nabla \mathcal{L}_m(\widehat{\mathbf{W}}_{t-1}) - \nabla \mathcal{L}_1(\widehat{\mathbf{W}}_{t-1}), \mathbf{W} \right\rangle \right\}, \quad (3.2) \end{aligned}$$

where $\lambda_t > 0$ is a tuning parameter. After repeating the procedure for T times, the dmSLDA algorithm returns $\widehat{\mathbf{W}} = \widehat{\mathbf{W}}_{t^*}$ with $t^* = \operatorname{argmin}_{0 \leq t \leq T} \widehat{e}_t$, where $\widehat{e}_t = \sum_{m=2}^M \mathcal{L}_m(\widehat{\mathbf{W}}_t)$. In other words, we choose the best estimator in $\{\widehat{\mathbf{W}}_t\}_{t=0}^T$ in the sense that it minimizes the validation loss. Finally, one can perform classification based on the projected data $\widehat{\mathbf{W}}^\top \mathbf{x} \in \mathbb{R}^{K-1}$.

Remark 3.1. *The communication cost of the dmSLDA method is $\mathcal{O}(dKM)$ since it only communicates the estimators $\widehat{\mathbf{W}}_t$ and the gradients $\nabla \mathcal{L}_m(\widehat{\mathbf{W}}_{t-1})$. This is significantly smaller than that of the centralized methods, i.e., $\mathcal{O}(ndM)$ (communicate all raw data $\{(\mathbf{x}_i^m, y_i^m)\}$) or $\mathcal{O}(d^2M)$ (communicate local estimators $\widehat{\Sigma}^m$ and \widehat{U}^m), when $K \ll d, n$.*

Remark 3.2. *In numerical experiments, we propose to choose λ_t from a candidate set Λ_t by minimizing the validation loss $\widehat{e}(\widehat{\mathbf{W}}_t) := \sum_{m=2}^M \mathcal{L}_m(\widehat{\mathbf{W}}_t)$, i.e., the loss on all other machines.*

Remark 3.3. *Since in general a few iterations suffice for the procedure to match the same accuracy of the centralized estimator (Jordan et al., 2019), we choose $T = 3$ or 5 in practice.*

4 STATISTICAL ANALYSIS

In this section, we establish statistical properties ensuring that dmSLDA can achieve the performance of the centralized SLDA after a few rounds of communication. We prove this in two steps: first, we will establish a recursive estimation error bound, i.e., we upper bound the estimation error $\|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|$ in terms of the error in the previous round $\|\widehat{\mathbf{W}}_t - \mathbf{W}^*\|$; second, by applying this recursive estimation error bound iteratively, we can obtain an estimation error bound indicating that dmSLDA matches the performance of the centralized SLDA after a few rounds.

To begin with, we make a few assumptions. For the sake of simplicity, we consider a balanced setting defined below². In the balanced setting, the prior probability and the subsample size of each class are equal

²If the balance assumption does not hold and some class has significantly small prior probability and subsample size, one may refer to literature related for imbalanced data, such as Krawczyk (2016) and the references therein.

across different classes and machines. Our framework does allow approximately balanced settings when the sample size proportions of any single class and any single machine do not vanish.

Assumption 4.1 (Balanced). *We assume all classes have equal prior probability, i.e., $\pi_k = \frac{1}{K}$. Moreover, the subsample size of class k on machine m is assumed to be constant, denoted by b . Thus, the subsample size on a single machine is $n = Kb$ and the full sample size is $N = Mn$.*

Furthermore, we assume that the conditional distribution of \mathbf{x} given $y = k$ is sub-Gaussian with mean $\boldsymbol{\mu}_k$ and covariance $\boldsymbol{\Sigma}$.

Assumption 4.2 (sub-Gaussian). *We assume that given $y = k$, the transformed variable $\mathbf{z}^k = \boldsymbol{\Sigma}^{-1/2}(\mathbf{x} - \boldsymbol{\mu}_k)$ follows a sub-Gaussian distribution with mean $\mathbf{0}$ and covariance \mathbf{I}_d . In particular, there exists some $\sigma^2 > 0$ such that the following inequality holds,*

$$\mathbb{E}[e^{\lambda \langle \mathbf{v}, \mathbf{z}^k \rangle} \mid y = k] \leq e^{\frac{\lambda^2 \sigma^2}{2}}, \forall \lambda \in \mathbb{R}, \mathbf{v} \in \mathbb{R}^d, \|\mathbf{v}\|_2 = 1.$$

Moreover, we assume a specific restricted eigenvalue (RE) condition on the population covariance matrix $\boldsymbol{\Sigma}$, which is a common assumption in high-dimensional statistics (Cai and Liu, 2011; Lee et al., 2017; Tian and Gu, 2017; Wang et al., 2017). First, given a subset $\mathcal{A} \subset [d]$, we define the cone $\mathcal{C}(\mathcal{A}, 3) \subset \mathbb{R}^d$ as

$$\mathcal{C}(\mathcal{A}, 3) = \{\boldsymbol{\zeta} \in \mathbb{R}^d \mid \|\boldsymbol{\zeta}_{\mathcal{A}^c}\|_1 \leq 3\|\boldsymbol{\zeta}_{\mathcal{A}}\|_1\},$$

where $\boldsymbol{\zeta}_{\mathcal{A}} \in \mathbb{R}^d$ is $\boldsymbol{\zeta}$ with nonzeros restricted on the set \mathcal{A} , i.e., $\boldsymbol{\zeta}_{\mathcal{A},i} = \boldsymbol{\zeta}_i$, if $i \in \mathcal{A}$, $\boldsymbol{\zeta}_{\mathcal{A},i} = 0$ otherwise; and $\boldsymbol{\zeta}_{\mathcal{A}^c}$ is analogously defined. Then we assume $\boldsymbol{\Sigma}$ satisfies the $RE(\kappa_0, s)$ condition for some constant $\kappa_0 > 0$ and positive integer s , which is defined below.

Assumption 4.3 (Restricted Eigenvalue). *We assume that the covariance matrix $\boldsymbol{\Sigma}$ satisfies the $RE(\kappa_0, s)$ condition, i.e., there exists some constant $\kappa_0 > 0$ such that*

$$\boldsymbol{\zeta}^\top \boldsymbol{\Sigma} \boldsymbol{\zeta} \geq \kappa_0 \|\boldsymbol{\zeta}\|_2^2, \forall \boldsymbol{\zeta} \in \mathcal{C}(\mathcal{A}, 3), |\mathcal{A}| \leq s.$$

Assumption 4.2 and 4.3 allow us to establish concentration inequalities for $\widehat{\boldsymbol{\Sigma}}$ and $\widehat{\mathbf{U}}$ as well as an RE condition on the sample covariance $\widehat{\boldsymbol{\Sigma}}^1$, which are stated in Lemmas A.1 and A.4. We are ready to present our main theoretical results. We need to define the shifted loss function $\widetilde{\mathcal{L}}_1(\mathbf{W}, \widehat{\mathbf{W}}_t)$ as follows,

$$\begin{aligned} \widetilde{\mathcal{L}}_1(\mathbf{W}, \widehat{\mathbf{W}}_t) &= \mathcal{L}_1(\mathbf{W}) + \left\langle \frac{1}{M} \sum_{m=1}^M \nabla \mathcal{L}_m(\widehat{\mathbf{W}}_t) \right. \\ &\quad \left. - \nabla \mathcal{L}_1(\widehat{\mathbf{W}}_t), \mathbf{W} \right\rangle. \end{aligned} \quad (4.1)$$

We establish in Lemma 4.4 an upper bound on the $\ell_{\infty, \infty}$ norm $\|\nabla \widetilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t)\|_{\infty, \infty}$ of the gradient of the shifted loss function at \mathbf{W}^* , which suggests an appropriate choice of the tuning parameter λ_{t+1} .

Lemma 4.4. *Suppose Assumptions 4.1 and 4.2 hold and $\sup_j \boldsymbol{\Sigma}_{jj} \leq c_0$ for some constant c_0 , then there exist some universal constants c_1 and c_2 such that*

$$\|\nabla \widetilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t)\|_{\infty, \infty} \leq a_n + b_n \|\widehat{\mathbf{W}}_t - \mathbf{W}^*\|_{1,1},$$

holds for all t with probability at least $1 - 2\delta$, where

$$\begin{aligned} a_n &= c_1 \sqrt{\frac{\log(dM/\delta)}{N}} \|\mathbf{W}^*\|_{1,1} \\ &\quad + c_2 \sqrt{\frac{K \log(dMK/\delta)}{N}}, \end{aligned} \quad (4.2)$$

$$b_n = 2c_1 \sqrt{\frac{\log(dM/\delta)}{n}}. \quad (4.3)$$

The $\ell_{1,1}$ regularization parameter λ_{t+1} is set as twice the upper bound in Lemma 4.4, i.e.,

$$\lambda_{t+1} = 2a_n + 2b_n \|\widehat{\mathbf{W}}_t - \mathbf{W}^*\|_{1,1}. \quad (4.4)$$

Although the right hand side of (4.4) involves the oracle matrix \mathbf{W}^* and thus unknown, it however suggests how to tune the parameter in practice. For example, the parameter λ_{t+1} should decrease as the estimation error $\|\widehat{\mathbf{W}}_t - \mathbf{W}^*\|_{1,1}$ goes down during iterations. Next, with this specific choice of parameter λ_{t+1} , we establish in Theorem 4.5 a recursive estimation error bound connecting the estimation error $\|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|$ to that of the previous iteration $\|\widehat{\mathbf{W}}_t - \mathbf{W}^*\|$, where the norm is the $\ell_{1,1}$ or $\ell_{2,2}$ norm.

Theorem 4.5. *Suppose that Assumptions 4.1, 4.2, and 4.3 hold, $\sup_j \boldsymbol{\Sigma}_{jj} \leq c_0$ for some constant c_0 , and the regularization parameter λ_{t+1} is set as in (4.4), then there exist some universal constants c_1, c_2, c_3 such that for $n \geq c_3 s \log(d/\delta)$, the recursive estimation error bounds*

$$\begin{aligned} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2} &\leq \frac{48a_n \sqrt{s}}{\kappa_0} \\ &\quad + \frac{48b_n \sqrt{s}}{\kappa_0} \|\widehat{\mathbf{W}}_t - \mathbf{W}^*\|_{1,1}, \\ \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1} &\leq \frac{192a_n s}{\kappa_0} \\ &\quad + \frac{192b_n s}{\kappa_0} \|\widehat{\mathbf{W}}_t - \mathbf{W}^*\|_{1,1}, \end{aligned}$$

hold for all t with probability at least $1 - 3\delta$, where a_n and b_n are given by (4.2) and (4.3), respectively, and $s = |\mathcal{S}|$ is the cardinality of the support \mathcal{S} of \mathbf{W}^* .

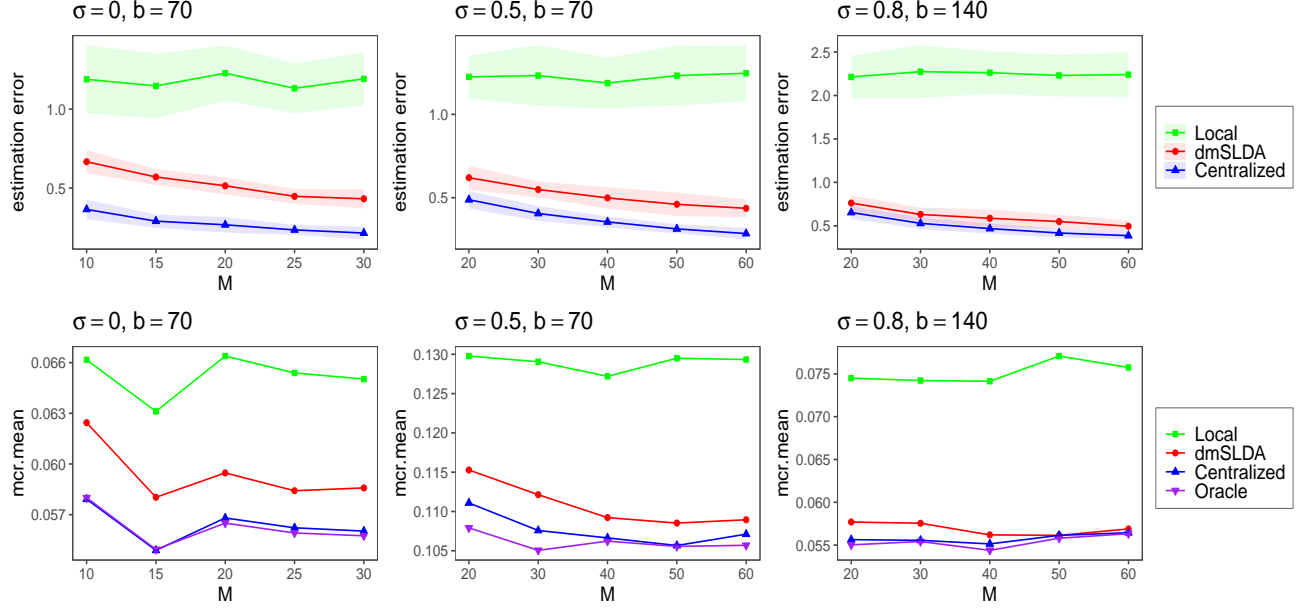


Figure 1: The top plots show the mean and the standard deviation of $\ell_{2,2}$ estimation error for dmSLDA, local SLDA, and centralized SLDA over 40 repetitions for different σ, b and M . The bottom plots show the averaged MCR on the testing data for all four methods over 40 repetitions.

Theorem 4.5 upper bounds the $(t + 1)$ th estimation error $\|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1}$ by a linear function of the t th estimation error $\|\widehat{\mathbf{W}}_t - \mathbf{W}^*\|_{1,1}$. Thus by applying Theorem 4.5 iteratively, we can bound $\|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|$ in terms of the size of the local estimation error $\|\widehat{\mathbf{W}}_0 - \mathbf{W}^*\|$.

Theorem 4.6. *Assume that the conditions in Theorem 4.5 hold and the regularization parameter λ_{t+1} is set as in (4.4), then there exist some universal constants c_1, c_2, c_3 such that for $n \geq c_3 s \log(d/\delta)$, the estimation error bounds*

$$\begin{aligned} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1} &\leq \frac{\mathbf{a}(1 - \mathbf{b}^{t+1})}{1 - \mathbf{b}} \\ &\quad + \mathbf{b}^{t+1} \|\widehat{\mathbf{W}}_0 - \mathbf{W}^*\|_{1,1}, \\ \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2} &\leq \frac{\mathbf{a}(1 - \mathbf{b}^{t+1})}{4\sqrt{s}(1 - \mathbf{b})} \\ &\quad + \frac{\mathbf{b}^{t+1}}{4\sqrt{s}} \|\widehat{\mathbf{W}}_0 - \mathbf{W}^*\|_{1,1}, \end{aligned}$$

hold with probability at least $1 - 3\delta$, where $\mathbf{a} = \frac{192a_n s}{\kappa_0}$ and $\mathbf{b} = \frac{192b_n s}{\kappa_0}$.

Theorem 4.6 provides theoretical guarantees that dmSLDA performs as well as the centralized SLDA in terms of the estimation accuracy. To see that, we observe that the coefficient $\mathbf{b} < 1$ when the local sample size n is sufficiently large and thus the terms involving \mathbf{b}^{t+1} are negligible after a few (logarithmic) iterations.

In other words, the estimation error for sufficiently large t satisfies

$$\|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2} \lesssim \sqrt{\frac{s \log(dM)}{N}},$$

which is the same as that of the centralized method.

5 NUMERICAL STUDIES

In this section, we conduct numerical studies to further back up our methodology and theory. The R code is released [here](#).

5.1 dmSLDA

In this subsection, we conduct numerical simulations to illustrate the theoretical results in Section 4. Four methods are implemented and compared: the dmSLDA, the local SLDA ($\widehat{\mathbf{W}}_0$), the centralized SLDA ($\widehat{\mathbf{W}}$ based on $\widehat{\Sigma} = \frac{1}{M} \sum_{m=1}^M \widehat{\Sigma}^m$ and $\widehat{\mathbf{U}} = \frac{1}{M} \sum_{m=1}^M \widehat{\mathbf{U}}^m$), and the oracle (\mathbf{W}^*). For the dmSLDA, we set $T = 3$ as the maximum rounds of communication. In the experiments, we fix the number of classes $K = 3$, the dimension $d = 400$, and the class-conditional means $\boldsymbol{\mu}_1 = (-2, -2, -2, 0, \dots, 0)^\top$, $\boldsymbol{\mu}_2 = (0, 0, 0, 2, 2, 2, 0, \dots, 0)^\top$, and $\boldsymbol{\mu}_3 = (0, \dots, 0)^\top$. The class-conditional covariance Σ is set as $\Sigma_{ij} = \sigma^{|i-j|}$ for all $1 \leq i \leq j \leq d$, where $\sigma \in \{0, 0.5, 0.8\}$ gives three different settings and $0^0 = 1$ by convention. We

set the subsample size $b = 70$ of each class on one machine when $\sigma \in \{0, 0.5\}$ and set $b = 140$ when $\sigma = 0.8$.³ Let the number of machines M vary from 20 to 60 by an increment of 10. For each setting, i.e., $(\sigma, b) \in \{(0, 70), (0.5, 70), (0.8, 140)\}$ and $M \in \{20, 30, \dots, 60\}$, we repeat the following procedure 40 times. First, we generate the training data $\{\mathbf{x}_{kj}^m\}_{j=1}^b$ of class k on machine m and the testing data $\{\mathbf{x}_{kj}\}_{j=1}^{300}$ of class k from $\mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma})$ independently. Once the data are generated, we estimate the transformation matrix \mathbf{W}^* based on the training data and then compute both the $\ell_{2,2}$ estimation error $\|\widehat{\mathbf{W}} - \mathbf{W}^*\|_{2,2}$ and the misclassification rate (MCR) based on the testing data. The results are displayed in Figure 1.

It can be seen from Figure 1 that in all settings, the $\ell_{2,2}$ estimation error of the dmSLDA is close to that of the centralized SLDA, which is much smaller than that of the local SLDA, as implied by Theorem 4.6. Moreover, the averaged MCR of the centralized SLDA is almost equal to that of the oracle SLDA, to which the averaged MCR of the dmSLDA is comparable. They all outperform the local SLDA in terms of the averaged MCR. The rationale behind this phenomenon is that a more accurate estimator of \mathbf{W}^* leads to a better linear discriminant classifier.

5.2 Binary Classification

In this subsection, we compare the dmSLDA method with other distributed classifiers in the literature. To the best of our knowledge, dmSLDA is the first distributed multicategorical classifier for high dimensional data. Thus, we restrict our experiments to the binary setting. Specifically, we compare dmSLDA with the Centralized SLDA (Cen.SLDA) and two other distributed binary classifiers, the divide-and-conquer SLDA (dSLDA) method proposed by Tian and Gu (2017) and the divide-and-conquer ℓ_1 SVM (dSVM) method proposed by Lian and Fan (2018). We adopt the following binary balanced setting: $K = 2$, $d = 200$, $b = 200$, $\boldsymbol{\Sigma}_{ij} = 0.8^{|i-j|}$, $\boldsymbol{\mu}_1 = (\text{rep}(0.5, 10), 0, \dots)$, $\boldsymbol{\mu}_2 = -\boldsymbol{\mu}_1$ and $M \in \{5, 10, 15, 20, 30\}$. For each setting, we generate the training data $\{\mathbf{x}_{kj}^m\}_{j=1}^b$ of class k on machine m and the testing data $\{\mathbf{x}_{kj}\}_{j=1}^{300}$ of class k from $\mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma})$ independently and then implement all four methods. The whole procedure is repeated 40 times and the averaged misclassification rates (MCR) are reported in Figure 2. It turns out that dmSLDA, dSLDA and Cen.SLDA are comparable and all these three methods outperform dSVM.

³When σ increases, κ_0 decreases and thus we need a larger subsample size.




Figure 2: Comparison between dSVM, dSLDA, dmSLDA, and Cen.SLDA. The averaged misclassification rate is computed over 40 repetitions.

6 DISCUSSION

In this paper, we propose an $\ell_{1,1}$ regularized multicategory SLDA method and extend it to the distributed setting when data are stored across multiple sites. We establish statistical properties ensuring that the distributed sparse multicategory LDA performs as good as the centralized sparse LDA after a few rounds of communications. Here we assume the data points follow sub-Gaussian distributions. It is possible to extend the current framework to deal with heavy-tailed data or data with adversarial contamination. We leave this to future work.

References

- Beck, A. and Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202.
- Bickel, P. J. and Levina, E. (2004). Some theory for Fisher’s linear discriminant function, ‘naive Bayes’, and some alternatives when there are many more variables than observations. *Bernoulli*, 10(6):989–1010.
- Bickel, P. J. and Levina, E. (2008). Regularized estimation of large covariance matrices. *The Annals of Statistics*, 36(1):199–227.
- Cai, T. and Liu, W. (2011). A direct estimation approach to sparse linear discriminant analysis. *Journal of the American statistical association*, 106(496):1566–1577.
- Clemmensen, L., Hastie, T., Witten, D., and Ersbøll, B. (2011). Sparse discriminant analysis. *Technometrics*, 53(4):406–413.
- Fan, J. and Fan, Y. (2008). High-dimensional classification using features annealed independence rules. *The Annals of Statistics*, 36(6):2605–2637.

- Fan, J., Feng, Y., and Tong, X. (2012). A road to classification in high dimensional space: the regularized optimal affine discriminant. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(4):745–771.
- Fan, J., Wang, D., Wang, K., and Zhu, Z. (2019). Distributed estimation of principal eigenspaces. *The Annals of Statistics*, 47(6):3009–3031.
- Gaynanova, I., Booth, J. G., and Wells, M. T. (2016). Simultaneous sparse estimation of canonical vectors in the $p \gg n$ setting. *Journal of the American Statistical Association*, 111(514):696–706.
- Hand, D. J. (2006). Classifier technology and the illusion of progress. *Statistical science*, 21(1):1–14.
- Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, Berlin, 2 edition.
- Jordan, M. I., Lee, J. D., and Yang, Y. (2019). Communication-efficient distributed statistical inference. *Journal of the American Statistical Association*, 114(526):668–681.
- Kokiopoulou, E. and Frossard, P. (2010). Distributed classification of multiple observation sets by consensus. *IEEE Transactions on Signal Processing*, 59(1):104–114.
- Krawczyk, B. (2016). Learning from imbalanced data: open challenges and future directions. *Progress in Artificial Intelligence*, 5(4):221–232.
- Lee, J. D., Liu, Q., Sun, Y., and Taylor, J. E. (2017). Communication-efficient sparse regression. *The Journal of Machine Learning Research*, 18(1):115–144.
- Lian, H. and Fan, Z. (2018). Divide-and-conquer for debiased l_1 -norm support vector machine in ultra-high dimensions. *Journal of Machine Learning Research*, 18:1–26.
- Mai, Q., Yang, Y., and Zou, H. (2019). Multi-class sparse discriminant analysis. *Statistica Sinica*, 29(1):97–111.
- Mai, Q., Zou, H., and Yuan, M. (2012). A direct approach to sparse discriminant analysis in ultra-high dimensions. *Biometrika*, 99(1):29–42.
- Michie, D., Spiegelhalter, D. J., Taylor, C. C., and Campbell, J., editors (1994). *Machine Learning, Neural and Statistical Classification*. Ellis Horwood, Chichester.
- Pilanci, M. and Wainwright, M. J. (2015). Randomized sketches of convex programs with sharp guarantees. *IEEE Transactions on Information Theory*, 61(9):5096–5115.
- Qiao, Z., Zhou, L., and Huang, J. Z. (2009). Sparse linear discriminant analysis with applications to high dimensional low sample size data. *International Journal of Applied Mathematics*, 39(1):48–60.
- Rudelson, M. and Zhou, S. (2013). Reconstruction from anisotropic random measurements. *IEEE transactions on information theory*, 59(6):3434–3447.
- Safo, S. E. and Ahn, J. (2016). General sparse multi-class linear discriminant analysis. *Computational Statistics & Data Analysis*, 99:81–90.
- Shao, J., Wang, Y., Deng, X., and Wang, S. (2011). Sparse linear discriminant analysis by thresholding for high dimensional data. *The Annals of statistics*, 39(2):1241–1265.
- Tian, L. and Gu, Q. (2017). Communication-efficient distributed sparse linear discriminant analysis. In *Artificial Intelligence and Statistics*, volume 54, pages 1178–1187. PMLR.
- Wainwright, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press.
- Wang, J., Kolar, M., Srebro, N., and Zhang, T. (2017). Efficient distributed learning with sparsity. In *International Conference on Machine Learning*, volume 70, pages 3636–3645. PMLR.
- Wang, X., Yang, Z., Chen, X., and Liu, W. (2019). Distributed inference for linear support vector machine. *Journal of Machine Learning Research*, 20:1–41.
- Witten, D. M. and Tibshirani, R. (2011). Penalized classification using fisher’s linear discriminant. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(5):753–772.

Supplementary Material: Distributed Sparse Multicategory Discriminant Analysis

A PROOF OF MAIN RESULTS

A.1 Proof of Proposition 2.1

Proof of Proposition 2.1. Since $\mathbf{V}, \mathbf{W} \in \mathbb{R}^{d \times q}$ span the same subspace, there exists an invertible matrix $\mathbf{R} \in \mathbb{R}^{q \times q}$ such that $\mathbf{V} = \mathbf{W}\mathbf{R}$. The remaining proof follows the same argument of Proposition 4 in Gaynanova et al. (2016), while we note that the generalization of the orthogonal matrix \mathbf{R} to the general invertible matrix \mathbf{R} does not affect the proof. \square

A.2 Proof of Proposition 2.2

Proof of Proposition 2.2. Since $\mathbf{\Sigma}$ is invertible, the optimization problem (2.3) has a closed form solution $\mathbf{W}^* = \mathbf{\Sigma}^{-1}\mathbf{U}$. The core of the proof lies in the decomposition of the between-class covariance matrix $\mathbf{B} = \mathbf{U}\mathbf{P}\mathbf{U}^\top$ for some symmetric positive definite matrix $\mathbf{P} \in \mathbb{R}^{(K-1) \times (K-1)}$. By definition,

$$\begin{aligned} \mathbf{B} &= \sum_{k=1}^K \pi_k (\boldsymbol{\mu}_k - \bar{\boldsymbol{\mu}})(\boldsymbol{\mu}_k - \bar{\boldsymbol{\mu}})^\top \\ &= \sum_{k=1}^{K-1} \pi_k (\boldsymbol{\mu}_k - \bar{\boldsymbol{\mu}})(\boldsymbol{\mu}_k - \bar{\boldsymbol{\mu}})^\top + \frac{1}{\pi_K} \left(\sum_{k=1}^{K-1} \pi_k (\boldsymbol{\mu}_k - \bar{\boldsymbol{\mu}}) \right) \left(\sum_{k=1}^{K-1} \pi_k (\boldsymbol{\mu}_k - \bar{\boldsymbol{\mu}}) \right)^\top, \end{aligned}$$

where the second equality holds since $\pi_K (\boldsymbol{\mu}_K - \bar{\boldsymbol{\mu}}) = -\sum_{k=1}^{K-1} \pi_k (\boldsymbol{\mu}_k - \bar{\boldsymbol{\mu}})$. Then we can write $\mathbf{B} = \mathbf{U}\mathbf{P}\mathbf{U}^\top$, with $\mathbf{P} \in \mathbb{R}^{(K-1) \times (K-1)}$ given by

$$\mathbf{P} = \text{diag}(\pi_1, \dots, \pi_{K-1}) + \frac{1}{\pi_K} (\pi_1, \dots, \pi_{K-1})^\top (\pi_1, \dots, \pi_{K-1}).$$

Note that \mathbf{P} is a symmetric positive definite matrix. Then we study the relationship between \mathbf{W}^* and discriminant vectors $\{\mathbf{v}_k^*\}_{k=1}^q$ solved by Fisher's reduced rank LDA problem (2.1), where $q = \text{rank}(\mathbf{B})$. Let $\mathbf{A} = \mathbf{\Sigma}^{-1/2}\mathbf{U}\mathbf{P}\mathbf{U}^\top\mathbf{\Sigma}^{-1/2}$, then the Fisher's reduced rank LDA problem can be transformed into a generalized eigenvalue problem

$$\begin{aligned} \tilde{\mathbf{v}}_k^* &= \underset{\tilde{\mathbf{v}}_k \in \mathbb{R}^d}{\text{argmax}} \frac{\tilde{\mathbf{v}}_k^\top \mathbf{A} \tilde{\mathbf{v}}_k}{\tilde{\mathbf{v}}_k^\top \tilde{\mathbf{v}}_k}, \\ \text{s.t. } &\tilde{\mathbf{v}}_k^\top \tilde{\mathbf{v}}_j = 0, \quad \forall 1 \leq j \leq k \leq q, \end{aligned} \tag{A.1}$$

where $\tilde{\mathbf{v}}_k^* = \mathbf{\Sigma}^{1/2}\mathbf{v}_k^*$. Since $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{B}) = q$, the matrix \mathbf{A} can be orthogonally diagonalized as

$$\mathbf{A} = \mathbf{Q} \text{diag}(\lambda_1, \dots, \lambda_q, 0, \dots, 0) \mathbf{Q}^\top,$$

where $\mathbf{Q} = (\mathbf{a}_1, \dots, \mathbf{a}_d) \in \mathbb{R}^{d \times d}$ is an orthogonal matrix and $\lambda_1 \geq \dots \geq \lambda_q > 0$. Then the solution to the problem (A.1) is $(\tilde{\mathbf{v}}_1^*, \dots, \tilde{\mathbf{v}}_q^*) = (\mathbf{a}_1, \dots, \mathbf{a}_q)$ and thus the solution to the Fisher's reduced rank problem (2.1) is

$\mathbf{V}_q = (\mathbf{v}_1^*, \dots, \mathbf{v}_q^*) = \Sigma^{-1/2}(\mathbf{a}_1, \dots, \mathbf{a}_q)$. As a result, we can express the linear subspace $\mathcal{V}_q := \text{span}(\mathbf{V}_q)$ as

$$\begin{aligned} \mathcal{V}_q &= \text{span}(\Sigma^{-1/2}(\mathbf{a}_1, \dots, \mathbf{a}_q)) \\ &\stackrel{(i)}{=} \text{span}(\Sigma^{-1/2}(\sqrt{\lambda_1}\mathbf{a}_1, \dots, \sqrt{\lambda_q}\mathbf{a}_q)) \\ &\stackrel{(ii)}{=} \text{span}(\Sigma^{-1/2}\mathbf{A}) \\ &\stackrel{(iii)}{=} \text{span}(\Sigma^{-1/2}(\Sigma^{-1/2}\mathbf{U}\mathbf{P}^{1/2})) \\ &\stackrel{(iv)}{=} \text{span}(\Sigma^{-1}\mathbf{U}) = \text{span}(\mathbf{W}^*), \end{aligned}$$

where (i) holds since $\lambda_1 \geq \dots \geq \lambda_q > 0$, (ii) holds since $\text{span}(\sqrt{\lambda_1}\mathbf{a}_1, \dots, \sqrt{\lambda_q}\mathbf{a}_q) = \text{span}(\mathbf{A})$, (iii) holds since $\text{span}(\mathbf{A}) = \text{span}(\Sigma^{-1/2}\mathbf{U}\mathbf{P}^{1/2})$, and (iv) holds since $\mathbf{P}^{1/2}$ is invertible. \square

A.3 Proof of Lemma 4.4

Proof of Lemma 4.4. Taking derivative of the shifted loss function $\tilde{\mathcal{L}}_1(\mathbf{W}, \widehat{\mathbf{W}}_t)$ with respect to the first entity at \mathbf{W}^* , we get

$$\begin{aligned} \nabla \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t) &= \nabla \mathcal{L}_1(\mathbf{W}^*) + \frac{1}{M} \sum_{m=1}^M \nabla \mathcal{L}_m(\widehat{\mathbf{W}}_t) - \nabla \mathcal{L}_1(\widehat{\mathbf{W}}_t) \\ &= \frac{1}{M} \sum_{m=1}^M \nabla \mathcal{L}_m(\mathbf{W}^*) + \nabla \mathcal{L}_1(\mathbf{W}^*) - \nabla \mathcal{L}_1(\widehat{\mathbf{W}}_t) + \frac{1}{M} \sum_{m=1}^M (\nabla \mathcal{L}_m(\widehat{\mathbf{W}}_t) - \nabla \mathcal{L}_m(\mathbf{W}^*)), \end{aligned}$$

where the second equality is obtained by adding and subtracting a term $\frac{1}{M} \nabla \mathcal{L}_m(\mathbf{W}^*)$. By definition of $\mathcal{L}_m(\cdot)$, we have

$$\nabla \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t) = \frac{1}{M} \sum_{m=1}^M \nabla \mathcal{L}_m(\mathbf{W}^*) + (\widehat{\Sigma}^1 - \Sigma)(\mathbf{W}^* - \widehat{\mathbf{W}}_t) + \frac{1}{M} \sum_{m=1}^M (\widehat{\Sigma}^m - \Sigma)(\widehat{\mathbf{W}}_t - \mathbf{W}^*).$$

Thus, by the triangle inequality, we have

$$\begin{aligned} \|\nabla \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t)\|_{\infty, \infty} &\leq \left\| \frac{1}{M} \sum_{m=1}^M \nabla \mathcal{L}_m(\mathbf{W}^*) \right\|_{\infty, \infty} \\ &\quad + \left(\|\widehat{\Sigma}^1 - \Sigma\|_{\infty, \infty} + \frac{1}{M} \sum_{m=1}^M \|\widehat{\Sigma}^m - \Sigma\|_{\infty, \infty} \right) \|\widehat{\mathbf{W}}_t - \mathbf{W}^*\|_{1,1}. \end{aligned} \quad (\text{A.2})$$

Then we proceed to finish the proof by establishing upper bounds on the $\ell_{\infty, \infty}$ norm of $\frac{1}{M} \sum_{m=1}^M \nabla \mathcal{L}_m(\mathbf{W}^*)$ and $\widehat{\Sigma}^m - \Sigma$. First, we establish in Lemma A.1 upper bounds on the $\ell_{\infty, \infty}$ norm of the error matrix $\widehat{\Sigma}^m - \Sigma$. Then in Lemma A.2, we give upper bounds on the $\ell_{\infty, \infty}$ norm of the gradient term $\frac{1}{M} \sum_{m=1}^M \nabla \mathcal{L}_m(\mathbf{W}^*)$.

Lemma A.1. *Suppose the assumptions in Lemma 4.4 hold, then there exist some universal constants c_1 and c_2 such that the following bounds*

$$\|\widehat{\Sigma}^m - \Sigma\|_{\infty, \infty} \leq c_1 \sqrt{\frac{\log(dM/\delta)}{n}}, \quad \forall m \in [M], \quad (\text{A.3})$$

$$\left\| \frac{1}{M} \sum_{m=1}^M \widehat{\Sigma}^m - \Sigma \right\|_{\infty, \infty} \leq c_1 \sqrt{\frac{\log(dM/\delta)}{N}}, \quad (\text{A.4})$$

$$\|\widehat{\mathbf{U}}^m - \mathbf{U}\|_{\infty, \infty} \leq c_2 \sqrt{\frac{K \log(dMK/\delta)}{n}}, \quad \forall m \in [M], \quad (\text{A.5})$$

$$\left\| \frac{1}{M} \sum_{m=1}^M \widehat{\mathbf{U}}^m - \mathbf{U} \right\|_{\infty, \infty} \leq c_2 \sqrt{\frac{K \log(dMK/\delta)}{N}}, \quad (\text{A.6})$$

hold with probability at least $1 - 2\delta$.

Lemma A.2. *Assume that the covariance matrix Σ is nonsingular and the upper bounds (A.3), (A.4), (A.5), and (A.6) in Lemma A.1 hold, then we have for all $m \in [M]$,*

$$\|\nabla \mathcal{L}_m(\mathbf{W}^*)\|_{\infty, \infty} \leq c_1 \sqrt{\frac{\log(dM/\delta)}{n}} \|\mathbf{W}^*\|_{1,1} + c_2 \sqrt{\frac{K \log(dMK/\delta)}{n}}, \quad (\text{A.7})$$

$$\left\| \frac{1}{M} \sum_{m=1}^M \nabla \mathcal{L}_m(\mathbf{W}^*) \right\|_{\infty, \infty} \leq c_1 \sqrt{\frac{\log(dM/\delta)}{N}} \|\mathbf{W}^*\|_{1,1} + c_2 \sqrt{\frac{K \log(dMK/\delta)}{N}}, \quad (\text{A.8})$$

where the constants c_1, c_2 are defined in Lemma A.1.

By Lemma A.1 and A.2, we have with probability at least $1 - 2\delta$ that inequalities (A.3), (A.4), (A.5), (A.6), (A.7) and (A.8) hold for some universal constants c_1 and c_2 . Substituting these inequalities into the inequality (A.2), we have

$$\begin{aligned} \|\nabla \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t)\|_{\infty, \infty} &\leq c_1 \sqrt{\frac{\log(dM/\delta)}{N}} \|\mathbf{W}^*\|_{1,1} + c_2 \sqrt{\frac{K \log(dMK/\delta)}{N}} \\ &\quad + 2c_1 \sqrt{\frac{\log(dM/\delta)}{n}} \|\mathbf{W}^* - \widehat{\mathbf{W}}_t\|_{1,1}, \end{aligned}$$

with probability at least $1 - 2\delta$, which concludes the proof. \square

A.4 Proof of Theorem 4.5

Before giving a proof of Theorem 4.5, we first state two technical lemmas that will be used. Let us begin with the definition of a series of subsets in $\mathbb{R}^{d \times (K-1)}$. For any integer s , we define the subset $\mathcal{H}(s) \subset \mathbb{R}^{d \times (K-1)}$ by

$$\mathcal{H}(s) = \left\{ \Delta \in \mathbb{R}^{d \times (K-1)} \mid \Delta_{\cdot k} \in \mathcal{C}(\mathcal{A}, 3), \text{ for some } |\mathcal{A}| \leq s, \forall 1 \leq k \leq K-1 \right\},$$

where $\Delta_{\cdot k}$ is the k th column of Δ . Furthermore, given an index set $\mathcal{S} \subset [(d, K-1)]$, we define the cone $\tilde{\mathcal{C}}(\mathcal{S}, 3)$ in $\mathbb{R}^{d \times (K-1)}$ analogous to $\mathcal{C}(\mathcal{A}, 3)$ in \mathbb{R}^d by

$$\tilde{\mathcal{C}}(\mathcal{S}, 3) = \left\{ \Delta \in \mathbb{R}^{d \times (K-1)} \mid \|\Delta_{\mathcal{S}^c}\|_{1,1} \leq 3 \|\Delta_{\mathcal{S}}\|_{1,1} \right\},$$

where $\Delta_{\mathcal{S}} \in \mathbb{R}^{d \times (K-1)}$ is the restriction of Δ on the set \mathcal{S} and $\Delta_{\mathcal{S}^c}$ is analogously defined. It is shown in Lemma A.3 that the t th error $\widehat{\mathbf{W}}_t - \mathbf{W}^*$ is in the subset $\mathcal{H}(s) \cap \tilde{\mathcal{C}}(\mathcal{S}, 3)$ under certain conditions, where $s = |\mathcal{S}|$ and $\mathcal{S} = \text{supp}(\mathbf{W}^*)$ is the support of \mathbf{W}^* . In addition, we establish in Lemma A.4 a strong restricted convexity property of $\mathcal{L}_1(\cdot)$, which proves to be useful.

Lemma A.3. *Assume that the upper bound in Lemma 4.4 holds, λ_{t+1} is chosen by (4.4), and the cardinality $|\mathcal{S}| = s$, where $\mathcal{S} = \text{supp}(\mathbf{W}^*)$ is the support of \mathbf{W}^* . Then for any t , the $(t+1)$ th error $\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*$ is in the subset $\mathcal{H}(s) \cap \tilde{\mathcal{C}}(\mathcal{S}, 3)$.*

Lemma A.4 (Strong Restricted Convexity). *Suppose the conditions in Theorem 4.5 hold, then there exists some constant c_3 such that for $n \geq c_3 s \log(d/\delta)$, the empirical loss function \mathcal{L}_1 satisfies the strong restricted convexity property, i.e.,*

$$\mathcal{L}_1(\mathbf{W}^* + \Delta) - \mathcal{L}_1(\mathbf{W}^*) - \langle \nabla \mathcal{L}_1(\mathbf{W}^*), \Delta \rangle \geq \frac{\kappa_0}{4} \|\Delta\|_{2,2}^2, \quad \forall \Delta \in \mathcal{H}(s), \quad (\text{A.9})$$

with probability at least $1 - \delta$.

Remark A.5. *By definition of $\mathcal{L}_1(\cdot)$, the strong restricted convexity is reduced to the following inequality,*

$$\langle \Delta, \widehat{\Sigma}^1 \Delta \rangle \geq \frac{\kappa_0}{2} \|\Delta\|_{2,2}^2, \quad \forall \Delta \in \mathcal{H}(s).$$

As shown in the proof, we prove this inequality by establishing the restricted eigenvalue $RE(\kappa_0/2, s)$ property on the empirical covariance matrix $\widehat{\Sigma}^1$. When we are finishing the paper, we find that the restricted eigenvalue property has been established in Rudelson and Zhou (2013). They proved the RE property via a reduction principle while in this paper we prove the RE property directly by applying Proposition 1 in Pilanci and Wainwright (2015).

Proof of Theorem 4.5. By definition of $\tilde{\mathcal{L}}_1(\cdot, \cdot)$, we have

$$\begin{aligned} & \tilde{\mathcal{L}}_1(\widehat{\mathbf{W}}_{t+1}, \widehat{\mathbf{W}}_t) - \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t) \\ &= \mathcal{L}_1(\widehat{\mathbf{W}}_{t+1}) - \mathcal{L}_1(\mathbf{W}^*) + \left\langle \frac{1}{M} \sum_{m=1}^M \nabla \mathcal{L}_m(\widehat{\mathbf{W}}_t) - \nabla \mathcal{L}_1(\widehat{\mathbf{W}}_t), \widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \right\rangle. \end{aligned} \quad (\text{A.10})$$

Since Assumption 4.1 and 4.2 hold and $\sup \Sigma_{jj} \leq c_0$, the upper bound in Lemma 4.4 holds for some universal constants c_1 and c_2 with probability at least $1 - 2\delta$. Since λ_{t+1} is chosen by (4.4), the $(t+1)$ th error $\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \in \mathcal{H}(s) \cap \tilde{\mathcal{C}}(\mathcal{S}, 3)$ by Lemma A.3, where $s = |\mathcal{S}|$ and $\mathcal{S} = \text{supp}(\mathbf{W}^*)$ is the support of \mathbf{W}^* . In addition, since Assumption 4.3 also holds, the empirical loss function \mathcal{L}_1 satisfies the strong restricted convexity property (A.9) when $n \geq c_3 s \log(d/\delta)$ for some universal constant c_3 with probability at least $1 - \delta$ by Lemma A.4. In the remaining part of analysis, we first assume the conclusions in Lemma 4.4, A.3 and A.4 hold and we will go back to the high-probability language in the end.

Combining the fact that $\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \in \mathcal{H}(s)$ and the strong restricted convexity property (A.9) of \mathcal{L}_1 , we obtain the following inequality,

$$\mathcal{L}_1(\widehat{\mathbf{W}}_{t+1}) - \mathcal{L}_1(\mathbf{W}^*) \geq \left\langle \nabla \mathcal{L}_1(\mathbf{W}^*), \widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \right\rangle + \frac{\kappa_0}{4} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2}^2. \quad (\text{A.11})$$

Substitute this inequality into the equation (A.10), we get

$$\begin{aligned} & \tilde{\mathcal{L}}_1(\widehat{\mathbf{W}}_{t+1}, \widehat{\mathbf{W}}_t) - \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t) \\ & \geq \left\langle \nabla \mathcal{L}_1(\mathbf{W}^*) + \frac{1}{M} \sum_{m=1}^M \nabla \mathcal{L}_m(\widehat{\mathbf{W}}_t) - \nabla \mathcal{L}_1(\widehat{\mathbf{W}}_t), \widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \right\rangle + \frac{\kappa_0}{4} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2}^2 \\ & = \left\langle \nabla \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t), \widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \right\rangle + \frac{\kappa_0}{4} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2}^2, \end{aligned} \quad (\text{A.12})$$

where the equality follows from the definition of $\nabla \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t)$. By the optimality (3.2) of $\widehat{\mathbf{W}}_{t+1}$, we can derive the following inequality,

$$\tilde{\mathcal{L}}_1(\widehat{\mathbf{W}}_{t+1}, \widehat{\mathbf{W}}_t) - \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t) + \lambda_{t+1} \|\widehat{\mathbf{W}}_{t+1}\|_{1,1} - \lambda_{t+1} \|\mathbf{W}^*\|_{1,1} \leq 0. \quad (\text{A.13})$$

Combining the inequalities (A.12) and (A.13), the upper bound in Lemma 4.4, and the definition (4.4) of λ_{t+1} , we get

$$\begin{aligned} & \lambda_{t+1} \|\mathbf{W}^*\|_{1,1} - \lambda_{t+1} \|\widehat{\mathbf{W}}_{t+1}\|_{1,1} \geq \tilde{\mathcal{L}}_1(\widehat{\mathbf{W}}_{t+1}, \widehat{\mathbf{W}}_t) - \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t) \\ & \geq \left\langle \nabla \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t), \widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \right\rangle + \frac{\kappa_0}{4} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2}^2 \\ & \geq - \|\nabla \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t)\|_{\infty, \infty} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1} + \frac{\kappa_0}{4} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2}^2 \\ & \geq - \frac{\lambda_{t+1}}{2} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1} + \frac{\kappa_0}{4} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2}^2, \end{aligned} \quad (\text{A.14})$$

where the third inequality follows from the Hölder's inequality. Combining this inequality with the fact that $\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \in \tilde{\mathcal{C}}(\mathcal{S}, 3)$, we obtain that

$$\begin{aligned} \frac{\kappa_0}{4} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2}^2 & \leq \lambda_{t+1} \|\mathbf{W}^*\|_{1,1} - \lambda_{t+1} \|\widehat{\mathbf{W}}_{t+1}\|_{1,1} + \frac{\lambda_{t+1}}{2} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1} \\ & \leq \frac{3\lambda_{t+1}}{2} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1} \\ & \leq 6\sqrt{s}\lambda_{t+1} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2}, \end{aligned} \quad (\text{A.15})$$

where the last inequality holds since $\|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1} \leq 4\sqrt{s} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2}$ when $\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \in \tilde{\mathcal{C}}(\mathcal{S}, 3)$. Dividing $\|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2}$ from both sides of the inequality (A.15), we have

$$\|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2} \leq \frac{24\sqrt{s}\lambda_{t+1}}{\kappa_0}. \quad (\text{A.16})$$

Again by the fact that $\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \in \widetilde{\mathcal{C}}(\mathcal{S}, 3)$, we have

$$\|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1} \leq 4\sqrt{s}\|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2} \leq \frac{96s\lambda_{t+1}}{\kappa_0}. \quad (\text{A.17})$$

Recall the definition (4.4) of λ_{t+1} , the inequality (A.16) and (A.17) are exactly what we want. Finally, let us conclude the proof by noting that this analysis holds for all t and thus the inequalities (A.16) and (A.17) hold for all t with probability $1 - 3\delta$ when $n \geq c_3s \log(d/\delta)$. \square

A.5 Proof of Theorem 4.6

Proof of Theorem 4.6. Setting $\mathbf{a} = \frac{192sa_n}{\kappa_0}$ and $\mathbf{b} = \frac{192sb_n}{\kappa_0}$, we can rewrite Theorem 4.5 in the following iteration form,

$$\|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1} \leq \mathbf{a} + \mathbf{b}\|\widehat{\mathbf{W}}_t - \mathbf{W}^*\|_{1,1},$$

Applying Theorem 4.5 iteratively, we can obtain the following upper bound on $\|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1}$,

$$\begin{aligned} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1} &\leq \mathbf{a} + \mathbf{b}\|\widehat{\mathbf{W}}_t - \mathbf{W}^*\|_{1,1} \\ &\leq \mathbf{a} + \mathbf{b}\left(\mathbf{a} + \mathbf{b}\|\widehat{\mathbf{W}}_{t-1} - \mathbf{W}^*\|_{1,1}\right) \\ &\leq \dots \\ &\leq \frac{\mathbf{a}(1 - \mathbf{b}^{t+1})}{1 - \mathbf{b}} + \mathbf{b}^{t+1}\|\widehat{\mathbf{W}}_0 - \mathbf{W}^*\|_{1,1}. \end{aligned} \quad (\text{A.18})$$

Then by Theorem 4.5 again, we can upper bound $\|\widehat{\mathbf{W}} - \mathbf{W}^*\|_{2,2}$ as follows,

$$\begin{aligned} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{2,2} &\leq \frac{\mathbf{a}}{4\sqrt{s}} + \frac{\mathbf{b}}{4\sqrt{s}}\|\widehat{\mathbf{W}}_t - \mathbf{W}^*\|_{1,1} \\ &\leq \frac{\mathbf{a}(1 - \mathbf{b}^{t+1})}{4\sqrt{s}(1 - \mathbf{b})} + \frac{\mathbf{b}^{t+1}}{4\sqrt{s}}\|\widehat{\mathbf{W}}_0 - \mathbf{W}^*\|_{1,1}, \end{aligned}$$

where the second inequality follows by (A.18). \square

B PROOF OF AUXILIARY LEMMAS

B.1 Proof of Lemma A.1

Proof of Lemma A.1. First, let us upper bound the $\ell_{\infty, \infty}$ norm of $\widehat{\Sigma}^m - \Sigma$ for a fixed m . By definition of $\widehat{\Sigma}^m$ and Assumption 4.1, we have

$$\begin{aligned} \widehat{\Sigma}^m - \Sigma &= \frac{1}{Kb - K} \sum_{k=1}^K \sum_{y_i^m = k} \{(\mathbf{x}_i^m - \boldsymbol{\mu}_k)(\mathbf{x}_i^m - \boldsymbol{\mu}_k)^\top - \Sigma\} \\ &\quad - \frac{1}{Kb - K} \sum_{k=1}^K \{b(\bar{\mathbf{x}}^{mk} - \boldsymbol{\mu}_k)(\bar{\mathbf{x}}^{mk} - \boldsymbol{\mu}_k)^\top - \Sigma\}, \end{aligned}$$

where $\bar{\mathbf{x}}^{mk} = \frac{1}{b} \sum_{y_i^m = k} \mathbf{x}_i^m$. For convenience, we introduce the following transformed random vectors,

$$\begin{aligned} \tilde{\mathbf{z}}_i^m &= \Sigma^{-1/2}(\mathbf{x}_i^m - \boldsymbol{\mu}_k), \quad \text{if } y_i^m = k, \quad \forall i, \\ \tilde{\mathbf{z}}^{mk} &= \sqrt{b}\Sigma^{-1/2}(\bar{\mathbf{x}}^{mk} - \boldsymbol{\mu}_k), \quad \forall k. \end{aligned}$$

Note that conditional on $\{y_i^m\}_{i=1}^n$, $\{\tilde{\mathbf{z}}_i^m\}_{i=1}^n$ are i.i.d. zero mean σ -sub-Gaussian random vectors and $\{\tilde{\mathbf{z}}^{mk}\}_{k=1}^K$ are i.i.d. zero mean σ -sub-Gaussian random vectors, though $\{\tilde{\mathbf{z}}_i^m\}_{i=1}^n$ and $\{\tilde{\mathbf{z}}^{mk}\}_{k=1}^K$ can be dependent. Using these transformed random vectors, we can rewrite the error $\widehat{\Sigma}^m - \Sigma$ as

$$\widehat{\Sigma}^m - \Sigma = \frac{Kb}{Kb - K} \mathbf{A} - \frac{K}{Kb - K} \mathbf{A}',$$

where

$$\mathbf{A} = \Sigma^{1/2} \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) \Sigma^{1/2}, \quad (\text{B.1})$$

$$\mathbf{A}' = \Sigma^{1/2} \left(\frac{1}{K} \sum_{k=1}^K \tilde{\mathbf{z}}^{mk} \tilde{\mathbf{z}}^{mk\top} - \mathbf{I}_d \right) \Sigma^{1/2}. \quad (\text{B.2})$$

By the triangle inequality, we have

$$\|\widehat{\Sigma}^m - \Sigma\|_{\infty, \infty} \leq \frac{Kb}{Kb - K} \|\mathbf{A}\|_{\infty, \infty} + \frac{K}{Kb - K} \|\mathbf{A}'\|_{\infty, \infty}. \quad (\text{B.3})$$

In the remaining part of the proof, we will apply Proposition 1 in [Pilanci and Wainwright \(2015\)](#) to upper bound the $\ell_{\infty, \infty}$ norm of \mathbf{A} and \mathbf{A}' separately, which leads to an upper bound on the $\ell_{\infty, \infty}$ norm of $\widehat{\Sigma}^m - \Sigma$ by (B.3). For reader's convenience, we present this proposition below.

Proposition B.1. *Let $\{\mathbf{s}_i\}_{i=1}^n \subset \mathbb{R}^d$ be i.i.d. samples from a zero-mean σ -sub-Gaussian distribution with $\text{Cov}(\mathbf{s}_i) = \mathbf{I}_d$. Then there exist some universal constants c_1, c_2 such that for any subset $\mathcal{Y} \subset \mathbb{S}^{d-1}$, we have with probability at least $1 - e^{-\frac{c_2 n \delta^2}{\sigma^4}}$,*

$$\sup_{\boldsymbol{\eta} \in \mathcal{Y}} \left| \boldsymbol{\eta}^\top \left(\frac{\mathbf{S}^\top \mathbf{S}}{n} - \mathbf{I}_d \right) \boldsymbol{\eta} \right| \leq c_1 \frac{\mathbb{W}(\mathcal{Y})}{\sqrt{n}} + \delta,$$

where $\mathbf{S}^\top = (\mathbf{s}_1, \dots, \mathbf{s}_n)$ and $\mathbb{W}(\mathcal{Y})$ is the Gaussian width of the subset \mathcal{Y} . Specifically, $\mathbb{W}(\mathcal{Y})$ is defined by

$$\mathbb{W}(\mathcal{Y}) = \mathbb{E}[\sup_{\boldsymbol{\eta} \in \mathcal{Y}} |\langle \mathbf{g}, \boldsymbol{\eta} \rangle|],$$

where the expectation is taken on $\mathbf{g} \in \mathbb{R}^d$, which is a standard normal random vector.

Let us first bound the $\ell_{\infty, \infty}$ norm of \mathbf{A} by applying Proposition B.1 to $\{\mathbf{s}_i\}_{i=1}^n = \{\tilde{\mathbf{z}}_i^m\}_{i=1}^n$ and $\mathcal{Y} = \mathcal{Y}_1 \subset \mathbb{S}^{d-1}$ defined by

$$\mathcal{Y}_1 = \left\{ \boldsymbol{\beta}_j = \frac{\boldsymbol{\Sigma}_{\cdot j}^{1/2}}{\|\boldsymbol{\Sigma}_{\cdot j}^{1/2}\|_2} \right\}_{1 \leq j \leq d} \cup \left\{ \boldsymbol{\gamma}_{jk} = \frac{\boldsymbol{\beta}_j + \boldsymbol{\beta}_k}{\|\boldsymbol{\beta}_j + \boldsymbol{\beta}_k\|_2} \right\}_{1 \leq j < k \leq d},$$

where $\boldsymbol{\Sigma}_{\cdot j}^{1/2}$ is the j th column of the matrix $\boldsymbol{\Sigma}^{1/2}$. Note that the subset \mathcal{Y}_1 is a finite set with at most $\frac{d(d+1)}{2}$ elements. Since $\{\tilde{\mathbf{z}}_i^m\}_{i=1}^n$ are i.i.d. samples from a zero mean σ -sub-Gaussian distribution with $\text{Cov}(\tilde{\mathbf{z}}_i^m) = \mathbf{I}_d$, we have with probability at least $1 - e^{-\frac{c_2 n \delta^2}{\sigma^4}}$,

$$\sup_{\boldsymbol{\eta} \in \mathcal{Y}_1} \left| \boldsymbol{\eta}^\top \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) \boldsymbol{\eta} \right| \leq c_1 \frac{\mathbb{W}(\mathcal{Y}_1)}{\sqrt{n}} + \delta, \quad (\text{B.4})$$

for some universal constants c_1 and c_2 . By definition of \mathbf{A} , we have

$$\begin{aligned} \|\mathbf{A}\|_{\infty, \infty} &= \sup_{1 \leq j \leq k \leq d} |\mathbf{A}_{jk}| = \sup_{1 \leq j \leq k \leq d} \left| \boldsymbol{\Sigma}_{\cdot j}^{1/2\top} \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) \boldsymbol{\Sigma}_{\cdot k}^{1/2} \right| \\ &= \sup_{1 \leq j \leq k \leq d} \left| \boldsymbol{\beta}_j^\top \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) \boldsymbol{\beta}_k \right| \cdot (\boldsymbol{\Sigma}_{jj} \boldsymbol{\Sigma}_{kk})^{1/2} \\ &\stackrel{(i)}{\leq} c_0 \cdot \sup_{1 \leq j \leq k \leq d} \left| \boldsymbol{\beta}_j^\top \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) \boldsymbol{\beta}_k \right|, \end{aligned}$$

where the inequality (i) follows from the condition $\sup_{1 \leq j \leq d} \boldsymbol{\Sigma}_{jj} \leq c_0$. By the polarization identity, we have

$$\begin{aligned} & \boldsymbol{\beta}_j^\top \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) \boldsymbol{\beta}_k \\ &= \frac{1}{2} (\boldsymbol{\beta}_j + \boldsymbol{\beta}_k)^\top \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) (\boldsymbol{\beta}_j + \boldsymbol{\beta}_k) \\ & \quad - \frac{1}{2} \boldsymbol{\beta}_j^\top \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) \boldsymbol{\beta}_j - \frac{1}{2} \boldsymbol{\beta}_k^\top \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) \boldsymbol{\beta}_k. \end{aligned}$$

Then by the triangle inequality, we have

$$\begin{aligned} \|\mathbf{A}\|_{\infty, \infty} &\leq c_0 \cdot \sup_{1 \leq j \leq k \leq d} \left| \boldsymbol{\beta}_j^\top \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) \boldsymbol{\beta}_k \right| \\ &\leq c_0 \cdot \frac{1}{2} \cdot \|\boldsymbol{\beta}_j + \boldsymbol{\beta}_k\|_2^2 \cdot \sup_{1 \leq j \leq k \leq d} \left| \boldsymbol{\gamma}_{jk}^\top \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) \boldsymbol{\gamma}_{jk} \right| \\ &\quad + c_0 \cdot \sup_{1 \leq j \leq d} \left| \boldsymbol{\beta}_j^\top \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) \boldsymbol{\beta}_j \right| \\ &\leq 3c_0 \cdot \sup_{\boldsymbol{\eta} \in \mathcal{Y}_1} \left| \boldsymbol{\eta}^\top \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) \boldsymbol{\eta} \right|, \end{aligned}$$

where the last inequality holds since $\|\boldsymbol{\beta}_j + \boldsymbol{\beta}_k\|_2^2 \leq 4$ and $\{\boldsymbol{\beta}_j, \boldsymbol{\gamma}_{jk}\}_{1 \leq j \leq k \leq d} \subset \mathcal{Y}_1$. Combining this inequality with the upper bound in (B.4), we can upper bound the $\ell_{\infty, \infty}$ norm of \mathbf{A} in terms of the Gaussian width $\mathbb{W}(\mathcal{Y}_1)$,

$$\begin{aligned} \|\mathbf{A}\|_{\infty, \infty} &\leq 3c_0 \cdot \sup_{\boldsymbol{\eta} \in \mathcal{Y}_1} \left| \boldsymbol{\eta}^\top \left(\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{z}}_i^m \tilde{\mathbf{z}}_i^{m\top} - \mathbf{I}_d \right) \boldsymbol{\eta} \right| \\ &\leq 3c_0 \left(c_1 \frac{\mathbb{W}(\mathcal{Y}_1)}{\sqrt{n}} + \delta \right), \end{aligned}$$

with probability at least $1 - e^{-\frac{c_2 n \delta^2}{\sigma^4}}$. Now we give an upper bound on the Gaussian width $\mathbb{W}(\mathcal{Y}_1)$ using the maximal inequality [Wainwright \(2019\)](#). Since the random variable $\langle \mathbf{g}, \boldsymbol{\eta} \rangle$ is Gaussian with zero mean and variance 1 for any $\boldsymbol{\eta} \in \mathcal{Y}_1 \subset \mathbb{S}^{d-1}$, and the cardinality of \mathcal{Y}_1 is finite less than $\frac{d(d+1)}{2}$, we have

$$\mathbb{W}(\mathcal{Y}_1) = \mathbb{E} \left[\sup_{\boldsymbol{\eta} \in \mathcal{Y}_1} |\langle \mathbf{g}, \boldsymbol{\eta} \rangle| \right] \leq \sqrt{2 \log(d(d+1))}.$$

Therefore, we have with probability at least $1 - e^{-\frac{c_2 n \delta^2}{\sigma^4}}$,

$$\|\mathbf{A}\|_{\infty, \infty} \leq c_0 \cdot \left(c_1 \sqrt{\frac{\log(d)}{n}} + \delta \right), \quad (\text{B.5})$$

for some universal constants c_1 and c_2 . Applying the same argument to \mathbf{A}' , we obtain the following upper bound

$$\|\mathbf{A}'\|_{\infty, \infty} \leq c_0 \cdot \left(c_1 \sqrt{\frac{\log(d)}{K}} + \delta \right), \quad (\text{B.6})$$

with probability at least $1 - e^{-\frac{c_2 K \delta^2}{\sigma^2}}$. Combining inequalities (B.3), (B.5), and (B.6), and the observation that \mathbf{A}' -term is negligible in $\widehat{\boldsymbol{\Sigma}}^m - \boldsymbol{\Sigma}$ relative to \mathbf{A} -term, we have with probability at least $1 - e^{-\frac{c_2 n \delta^2}{\sigma^4}}$,

$$\|\widehat{\boldsymbol{\Sigma}}^m - \boldsymbol{\Sigma}\|_{\infty, \infty} \leq \frac{b}{b-1} \|\mathbf{A}\|_{\infty, \infty} + \frac{1}{b-1} \|\mathbf{A}'\|_{\infty, \infty} \leq c_0 \cdot c_1 \left(\sqrt{\frac{\log(d)}{n}} + \delta \right), \quad (\text{B.7})$$

for some universal constants c_1, c_2 . Furthermore, we can bound the $\ell_{\infty, \infty}$ norm of $\frac{1}{M} \sum_{m=1}^M \widehat{\Sigma}^m - \Sigma$ in the same way. In specific, with probability at least $1 - e^{-\frac{c_2 n \delta^2}{\sigma^4}}$, we have

$$\left\| \frac{1}{M} \sum_{m=1}^M \widehat{\Sigma}^m - \Sigma \right\|_{\infty, \infty} \leq c_0 \cdot c_1 \left(\sqrt{\frac{\log(d)}{N}} + \frac{\delta}{\sqrt{M}} \right), \quad (\text{B.8})$$

for some constants c_1, c_2 . Using the method of union bound for (B.7) for all $m \in [M]$ and (B.8) and the change of variable of δ , we have

$$\begin{aligned} \|\widehat{\Sigma}^m - \Sigma\|_{\infty, \infty} &\leq c_1 \sqrt{\frac{\log(dM/\delta)}{n}}, \quad \forall m, \\ \left\| \frac{1}{M} \sum_{m=1}^M \widehat{\Sigma}^m - \Sigma \right\|_{\infty, \infty} &\leq c_1 \sqrt{\frac{\log(dM/\delta)}{N}}, \end{aligned} \quad (\text{B.9})$$

for some universal constant c_1 , with probability at least $1 - \delta$.

Next, we derive an upper bound on the $\ell_{\infty, \infty}$ norm of the error $\widehat{U}^m - U$ for a fixed m . By Assumption 4.1, we can rewrite $U_{\cdot k}$ in the following form

$$U_{\cdot k} = (\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K) \mathbf{h}^k, \quad \forall k \in \{1, \dots, K-1\},$$

where the vector $\mathbf{h}^k \in \mathbb{R}^K$ satisfies

$$\mathbf{h}_k^k = 1 - \frac{1}{K}, \mathbf{h}_j^k = -\frac{1}{K}, \forall j \neq k.$$

Using \mathbf{h}^k , we can express elements of $\widehat{U}^m - U$ as

$$[\widehat{U}^m - U]_{jk} = \widehat{U}_{jk}^m - U_{jk} = \widetilde{\mathbf{X}}_{j\cdot}^m \mathbf{h}^k - U_{jk},$$

where $\widetilde{\mathbf{X}}^m = (\widetilde{\mathbf{x}}^{m1}, \dots, \widetilde{\mathbf{x}}^{mK}) \in \mathbb{R}^{d \times K}$ is the standard unbiased estimator of $(\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K)$. Note that the (j, k) th element of $\widetilde{\mathbf{X}}^m$ is a linear combination of coordinates of \mathbf{x}_i^m and thus $\widetilde{\mathbf{X}}_{j\cdot}^m$ is sub-Gaussian with parameter

$$\sigma_{jk}^2 = \frac{\sigma^2 \boldsymbol{\Sigma}_{jj}}{b} \leq \tilde{\sigma}^2 := \frac{\sigma^2 c_0}{b},$$

where the inequality follows from the condition $\sup_j \boldsymbol{\Sigma}_{jj} \leq c_0$. Furthermore, columns of $\widetilde{\mathbf{X}}^m$ are independent random vectors. Therefore, $[\widehat{U}^m - U]_{jk}$ is zero-mean sub-Gaussian variable with parameter

$$\|\mathbf{h}^k\|_2 \cdot \tilde{\sigma} = \frac{\sqrt{K-1}}{K} \tilde{\sigma} \leq \tilde{\sigma}.$$

Then by the maximal inequality, we have

$$\|\widehat{U}^m - U\|_{\infty, \infty} = \sup_{1 \leq j \leq d, 1 \leq k \leq K-1} \left| \widehat{U}_{jk}^m - U_{jk} \right| \leq \delta, \quad (\text{B.10})$$

with probability at least $1 - 2dKe^{-\frac{\delta^2 b}{2c_0 \sigma^2}}$. Applying the same argument to $\frac{1}{M} \sum_{m=1}^M \widehat{U}^m - U$, we have

$$\left\| \frac{1}{M} \sum_{m=1}^M \widehat{U}^m - U \right\|_{\infty, \infty} \leq \frac{\delta}{\sqrt{M}}, \quad (\text{B.11})$$

with probability at least $1 - 2dKe^{-\frac{\delta^2 b}{2c_0 \sigma^2}}$. Using the method of union bound for (B.10) for all m and (B.11) and the change of variable of δ , we have with probability at least $1 - \delta$,

$$\begin{aligned} \|\widehat{U}^m - U\|_{\infty, \infty} &\leq c_2 \sqrt{\frac{K \log(dMK/\delta)}{n}}, \quad \forall m, \\ \left\| \frac{1}{M} \sum_{m=1}^M \widehat{U}^m - U \right\|_{\infty, \infty} &\leq c_2 \sqrt{\frac{K \log(dMK/\delta)}{N}}, \end{aligned} \quad (\text{B.12})$$

for some universal constant c_2 .

Finally, we conclude our proof by using the method of union bound for (B.9) and (B.12). \square

B.2 Proof of Lemma A.2

Proof of Lemma A.2. By definition of $\mathcal{L}_m(\cdot)$ and $\mathbf{W}^* = \Sigma^{-1}\mathbf{U}$, the gradient of $\mathcal{L}_m(\cdot)$ at the point \mathbf{W}^* is

$$\nabla \mathcal{L}_m(\mathbf{W}^*) = \widehat{\Sigma}^m \mathbf{W}^* - \widehat{\mathbf{U}}^m = (\widehat{\Sigma}^m - \Sigma) \mathbf{W}^* + (\mathbf{U} - \widehat{\mathbf{U}}^m).$$

Then by the triangle inequality, we obtain the following inequalities,

$$\begin{aligned} \|\nabla \mathcal{L}_m(\mathbf{W}^*)\|_{\infty, \infty} &\leq \|\widehat{\Sigma}^m - \Sigma\|_{\infty, \infty} \|\mathbf{W}^*\|_{1,1} + \|\widehat{\mathbf{U}}^m - \mathbf{U}\|_{\infty, \infty}, \\ \left\| \frac{1}{M} \sum_{m=1}^M \nabla \mathcal{L}_m(\mathbf{W}^*) \right\|_{\infty, \infty} &\leq \left\| \frac{1}{M} \sum_{m=1}^M \widehat{\Sigma}^m - \Sigma \right\|_{\infty, \infty} \|\mathbf{W}^*\|_{1,1} + \left\| \frac{1}{M} \sum_{m=1}^M \widehat{\mathbf{U}}^m - \mathbf{U} \right\|_{\infty, \infty} \end{aligned}$$

Substituting the upper bounds (A.3), (A.4), (A.5), and (A.6) in Lemma A.1 into the previous two inequalities, we have for all $m \in [M]$,

$$\begin{aligned} \|\nabla \mathcal{L}_m(\mathbf{W}^*)\|_{\infty, \infty} &\leq c_1 \sqrt{\frac{\log(dM/\delta)}{n}} \|\mathbf{W}^*\|_{1,1} + c_2 \sqrt{\frac{K \log(dMK/\delta)}{n}}, \\ \left\| \frac{1}{M} \sum_{m=1}^M \nabla \mathcal{L}_m(\mathbf{W}^*) \right\|_{\infty, \infty} &\leq c_1 \sqrt{\frac{\log(dM/\delta)}{N}} \|\mathbf{W}^*\|_{1,1} + c_2 \sqrt{\frac{K \log(dMK/\delta)}{N}}, \end{aligned}$$

where the constants c_1, c_2 are defined in Lemma A.1. \square

B.3 Proof of Lemma A.3

Proof of Lemma A.3. We prove this lemma in two steps: first, we show that $\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \in \mathcal{H}(s)$; second, we show that $\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \in \widetilde{\mathcal{C}}(\mathcal{S}, 3)$, where $\mathcal{S} = \text{supp}(\mathbf{W}^*)$ and $s = |\mathcal{S}|$.

Step 1: $\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \in \mathcal{H}(s)$. Before we give an proof, let us introduce some notations first. We use $\widehat{\mathbf{w}}_k^t$ and \mathbf{w}_k^* to denote the k th column of the matrix $\widehat{\mathbf{W}}_t$ and \mathbf{W}^* , respectively. Then we use $\mathcal{S}_k = \text{supp}(\mathbf{w}_k^*)$ to denote the support of the k th column of \mathbf{W}^* . Note that the cardinality of \mathcal{S}_k is smaller than that of $\mathcal{S} = \text{supp}(\mathbf{W}^*)$, which is s .

For each k and t , we define the following transformed loss function

$$\widetilde{\ell}_{1k}(\mathbf{w}, \widehat{\mathbf{w}}_k^t) = \ell_{1k}(\mathbf{w}) + \left\langle \frac{1}{M} \sum_{m=1}^M \nabla \ell_{mk}(\widehat{\mathbf{w}}_k^t) - \nabla \ell_{1k}(\widehat{\mathbf{w}}_k^t), \mathbf{w} \right\rangle,$$

where

$$\ell_{mk}(\mathbf{w}) = \frac{1}{2} \left\langle \mathbf{w}, \widehat{\Sigma}^m \mathbf{w} \right\rangle - \left\langle \widehat{\mathbf{U}}_k^m, \mathbf{w} \right\rangle.$$

Since the loss function $\widetilde{\mathcal{L}}_1(\mathbf{W}, \widehat{\mathbf{W}}_t)$ and the $\ell_{1,1}$ norm of \mathbf{W} can be decomposed in the form of

$$\begin{aligned} \widetilde{\mathcal{L}}_1(\mathbf{W}, \widehat{\mathbf{W}}_t) &= \sum_{k=1}^{K-1} \widetilde{\ell}_{1k}(\mathbf{w}_k, \widehat{\mathbf{w}}_k^t), \\ \|\mathbf{W}\|_{1,1} &= \sum_{k=1}^{K-1} \|\mathbf{w}_k\|_1, \end{aligned}$$

where $\mathbf{w}_k = \mathbf{W}_{\cdot k}$ is the k th column of \mathbf{W} , the optimization problem

$$\widehat{\mathbf{W}}_{t+1} = \underset{\mathbf{W} \in \mathbb{R}^{d \times (K-1)}}{\text{argmin}} \left\{ \widetilde{\mathcal{L}}_1(\mathbf{W}, \widehat{\mathbf{W}}_t) + \lambda_{t+1} \|\mathbf{W}\|_{1,1} \right\}$$

is equivalent to the following $K - 1$ minimization problems

$$\widehat{\mathbf{w}}_k^{t+1} = \underset{\mathbf{w}_k \in \mathbb{R}^d}{\text{argmin}} \left\{ \widetilde{\ell}_{1k}(\mathbf{w}_k, \widehat{\mathbf{w}}_k^t) + \lambda_{t+1} \|\mathbf{w}_k\|_1 \right\}, \quad \forall 1 \leq k \leq K - 1. \quad (\text{B.13})$$

Moreover, we have the following inequality

$$\|\nabla\tilde{\ell}_{1k}(\mathbf{w}_k^*, \hat{\mathbf{w}}_k^t)\|_\infty \leq \|\nabla\tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t)\|_{\infty, \infty} \leq \frac{\lambda_{t+1}}{2},$$

since $\nabla\tilde{\ell}_{1k}(\mathbf{w}_k^*, \hat{\mathbf{w}}_k^t)$ is the k th column of $\nabla\tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t)$ and λ_{t+1} is defined by (4.4).

Now let us prove the lemma by analyzing the k th column of $\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*$ with k and t fixed. By the triangle inequality, we have

$$\begin{aligned} \|\hat{\mathbf{w}}_k^{t+1}\|_1 - \|\mathbf{w}_k^*\|_1 &= \|\mathbf{w}_k^* + (\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k^c} + (\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k}\|_1 - \|\mathbf{w}_k^*\|_1 \\ &\geq \|\mathbf{w}_k^* + (\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k^c}\|_1 - \|(\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k}\|_1 - \|\mathbf{w}_k^*\|_1 \\ &= \|(\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k^c}\|_1 - \|(\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k}\|_1, \end{aligned}$$

where $\mathcal{S}_k = \text{supp}(\mathbf{w}_k^*)$. By the optimality (B.13) of $\hat{\mathbf{w}}_k^{t+1}$, we have

$$\tilde{\ell}_{1k}(\hat{\mathbf{w}}_k^{t+1}, \hat{\mathbf{w}}_k^t) + \lambda_{t+1}\|\hat{\mathbf{w}}_k^{t+1}\|_1 - \tilde{\ell}_{1k}(\mathbf{w}_k^*, \hat{\mathbf{w}}_k^t) - \lambda_{t+1}\|\mathbf{w}_k^*\|_1 \leq 0.$$

Combining the above two inequalities, we get

$$\tilde{\ell}_{1k}(\hat{\mathbf{w}}_k^{t+1}, \hat{\mathbf{w}}_k^t) - \tilde{\ell}_{1k}(\mathbf{w}_k^*, \hat{\mathbf{w}}_k^t) + \lambda_{t+1}(\|(\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k^c}\|_1 - \|(\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k}\|_1) \leq 0.$$

By the convexity of $\tilde{\ell}_{1k}(\cdot, \hat{\mathbf{w}}_k^t)$ we have

$$\tilde{\ell}_{1k}(\hat{\mathbf{w}}_k^{t+1}, \hat{\mathbf{w}}_k^t) - \tilde{\ell}_{1k}(\mathbf{w}_k^*, \hat{\mathbf{w}}_k^t) \geq \left\langle \nabla\tilde{\ell}_{1k}(\mathbf{w}_k^*, \hat{\mathbf{w}}_k^t), \hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^* \right\rangle.$$

Therefore,

$$\begin{aligned} 0 &\geq \left\langle \nabla\tilde{\ell}_{1k}(\mathbf{w}_k^*, \hat{\mathbf{w}}_k^t), \hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^* \right\rangle + \lambda_{t+1}(\|(\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k^c}\|_1 - \|(\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k}\|_1) \\ &\geq -\|\nabla\tilde{\ell}_{1k}(\mathbf{w}_k^*, \hat{\mathbf{w}}_k^t)\|_\infty\|\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*\|_1 + \lambda_{t+1}(\|(\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k^c}\|_1 - \|(\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k}\|_1). \end{aligned}$$

Since $\frac{\lambda_{t+1}}{2} \geq \|\nabla\tilde{\ell}_{1k}(\mathbf{w}_k^*, \hat{\mathbf{w}}_k^t)\|_\infty$, we have

$$\begin{aligned} 0 &\geq -\frac{1}{2}\|\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*\|_1 + \|(\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k^c}\|_1 - \|(\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k}\|_1 \\ &= \frac{1}{2}\|(\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k^c}\|_1 - \frac{3}{2}\|(\hat{\mathbf{w}}_k^{t+1} - \mathbf{w}_k^*)_{\mathcal{S}_k}\|_1, \end{aligned}$$

which proves that $\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \in \mathcal{H}(s)$ since $|\mathcal{S}_k| \leq |\mathcal{S}| = s$.

Step 2: $\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \in \tilde{\mathcal{C}}(\mathcal{S}, 3)$. The proof of this step is analogous to that of the first step. By the triangle inequality, we have

$$\begin{aligned} \|\widehat{\mathbf{W}}_{t+1}\|_{1,1} - \|\mathbf{W}^*\|_{1,1} &= \|\mathbf{W}^* + (\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}^c} + (\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}}\|_{1,1} - \|\mathbf{W}^*\|_{1,1} \\ &\geq \|\mathbf{W}^* + (\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}^c}\|_{1,1} - \|(\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}}\|_{1,1} - \|\mathbf{W}^*\|_{1,1} \\ &= \|(\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}^c}\|_{1,1} - \|(\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}}\|_{1,1}, \end{aligned}$$

where $\mathcal{S} = \text{supp}(\mathbf{W}^*)$. By the optimality of $\widehat{\mathbf{W}}_{t+1}$, we have

$$\tilde{\mathcal{L}}_1(\widehat{\mathbf{W}}_{t+1}, \widehat{\mathbf{W}}_t) + \lambda_{t+1}\|\widehat{\mathbf{W}}_{t+1}\|_{1,1} - \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t) - \lambda_{t+1}\|\mathbf{W}^*\|_{1,1} \leq 0.$$

Combining the above two inequalities, we get

$$\tilde{\mathcal{L}}_1(\widehat{\mathbf{W}}_{t+1}, \widehat{\mathbf{W}}_t) - \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t) + \lambda_{t+1}(\|(\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}^c}\|_{1,1} - \|(\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}}\|_{1,1}) \leq 0.$$

By the convexity of $\tilde{\mathcal{L}}_1(\cdot, \widehat{\mathbf{W}}_t)$ we have

$$\tilde{\mathcal{L}}_1(\widehat{\mathbf{W}}_{t+1}, \widehat{\mathbf{W}}_t) - \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t) \geq \left\langle \nabla\tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t), \widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \right\rangle.$$

Therefore,

$$\begin{aligned} 0 &\geq \left\langle \nabla \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t), \widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \right\rangle + \lambda_{t+1} (\|(\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}^c}\|_{1,1} - \|(\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}}\|_{1,1}) \\ &\geq -\|\nabla \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t)\|_{\infty, \infty} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1} + \lambda_{t+1} (\|(\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}^c}\|_{1,1} - \|(\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}}\|_{1,1}). \end{aligned}$$

Since $\frac{\lambda_{t+1}}{2} \geq \|\nabla \tilde{\mathcal{L}}_1(\mathbf{W}^*, \widehat{\mathbf{W}}_t)\|_{\infty, \infty}$, we have

$$\begin{aligned} 0 &\geq -\frac{1}{2} \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_{1,1} + \|(\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}^c}\|_{1,1} - \|(\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}}\|_{1,1} \\ &= \frac{1}{2} \|(\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}^c}\|_{1,1} - \frac{3}{2} \|(\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*)_{\mathcal{S}}\|_{1,1}, \end{aligned}$$

which proves that $\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^* \in \tilde{\mathcal{C}}(\mathcal{S}, 3)$. \square

B.4 Proof of Lemma A.4

Proof of Lemma A.4. By definition of $\mathcal{L}_1(\cdot)$, we have for each $\Delta \in \mathcal{H}(s)$,

$$\begin{aligned} \mathcal{L}_1(\mathbf{W}^* + \Delta) - \mathcal{L}_1(\mathbf{W}^*) - \langle \nabla \mathcal{L}_1(\mathbf{W}^*), \Delta \rangle &= \frac{1}{2} \left\langle \Delta, \widehat{\Sigma}^1 \Delta \right\rangle \\ &= \frac{1}{2} \sum_{j=1}^{K-1} \Delta_{\cdot j}^\top \Sigma \Delta_{\cdot j} \cdot \left\{ 1 + \frac{\Delta_{\cdot j}^\top (\widehat{\Sigma}^1 - \Sigma) \Delta_{\cdot j}}{\Delta_{\cdot j}^\top \Sigma \Delta_{\cdot j}} \right\}, \end{aligned} \quad (\text{B.14})$$

where $\Delta_{\cdot j}$ is the j th column of Δ . In the remain part of proof, we first establish an upper bound on the following term

$$\sup_{\Delta \in \mathcal{H}(s), 1 \leq j \leq K-1} \frac{|\Delta_{\cdot j}^\top (\widehat{\Sigma}^1 - \Sigma) \Delta_{\cdot j}|}{\Delta_{\cdot j}^\top \Sigma \Delta_{\cdot j}},$$

using Proposition B.1 and the $\text{RE}(\kappa_0, s)$ condition of Σ , then we prove the strong restricted convexity of \mathcal{L}_1 by the equation (B.14) and the $\text{RE}(\kappa_0, s)$ condition of Σ .

By definition of $\mathcal{H}(s)$, we have

$$\sup_{\Delta \in \mathcal{H}(s), 1 \leq j \leq K-1} \frac{|\Delta_{\cdot j}^\top (\widehat{\Sigma}^1 - \Sigma) \Delta_{\cdot j}|}{\Delta_{\cdot j}^\top \Sigma \Delta_{\cdot j}} \leq \sup_{\zeta \in \mathcal{C}(\mathcal{A}, 3), |\mathcal{A}| \leq s} \frac{|\zeta^\top (\widehat{\Sigma}^1 - \Sigma) \zeta|}{\zeta^\top \Sigma \zeta}.$$

Similar to the proof of Lemma A.1, by definition of $\widehat{\Sigma}^1$ and the triangle inequality, we have

$$\left| \zeta^\top (\widehat{\Sigma}^1 - \Sigma) \zeta \right| \leq \frac{Kb}{Kb - K} |\zeta^\top \mathbf{A} \zeta| + \frac{K}{Kb - K} |\zeta^\top \mathbf{A}' \zeta|, \quad (\text{B.15})$$

where \mathbf{A} and \mathbf{A}' are defined in (B.1) and (B.2) for $m = 1$. By Proposition B.1, for some constants c_1 and c_2 , we have

$$\begin{aligned} \sup_{\zeta \in \mathcal{C}(\mathcal{A}, 3), |\mathcal{A}| \leq s} \frac{|\zeta^\top \mathbf{A} \zeta|}{\zeta^\top \Sigma \zeta} &= \sup_{\eta = \frac{\Sigma^{1/2} \zeta}{\|\Sigma^{1/2} \zeta\|_2}, \zeta \in \mathcal{C}(\mathcal{A}, 3), |\mathcal{A}| \leq s} \left| \eta^\top \left(\frac{1}{n} \sum_{i=1}^n \tilde{z}_i^1 \tilde{z}_i^{1\top} - \mathbf{I}_d \right) \eta \right| \\ &\leq c_1 \frac{\mathbb{W}(\mathcal{Y}_2)}{\sqrt{n}} + \delta, \end{aligned} \quad (\text{B.16})$$

with probability at least $1 - e^{-\frac{c_2 n \delta^2}{\sigma^4}}$, where

$$\mathcal{Y}_2 = \left\{ \eta = \frac{\Sigma^{1/2} \zeta}{\|\Sigma^{1/2} \zeta\|_2} : \zeta \in \mathcal{C}(\mathcal{A}, 3), |\mathcal{A}| \leq s \right\} \subset \mathbb{S}^{d-1}.$$

Now let us upper bound the Gaussian width $\mathbb{W}(\mathcal{Y}_2)$. The Gaussian width $\mathbb{W}(\mathcal{Y}_2)$ is

$$\mathbb{W}(\mathcal{Y}_2) = \mathbb{E}[\sup_{\boldsymbol{\eta} \in \mathcal{Y}_2} |\langle \boldsymbol{g}, \boldsymbol{\eta} \rangle|] = \mathbb{E}[\sup_{\boldsymbol{\zeta} \in \mathcal{C}(\mathcal{A}, 3), |\mathcal{A}| \leq s, \|\boldsymbol{\Sigma}^{1/2} \boldsymbol{\zeta}\|_2 = 1} |\langle \boldsymbol{g}, \boldsymbol{\Sigma}^{1/2} \boldsymbol{\zeta} \rangle|].$$

Then by Hölder's inequality, we have

$$\begin{aligned} \mathbb{W}(\mathcal{Y}_2) &\leq \mathbb{E}[\sup_{\boldsymbol{\zeta} \in \mathcal{C}(\mathcal{A}, 3), |\mathcal{A}| \leq s, \|\boldsymbol{\Sigma}^{1/2} \boldsymbol{\zeta}\|_2 = 1} \|\boldsymbol{\Sigma}^{1/2} \boldsymbol{g}\|_\infty \|\boldsymbol{\zeta}\|_1] \\ &\stackrel{(i)}{\leq} \mathbb{E}[\sup_{\boldsymbol{\zeta} \in \mathcal{C}(\mathcal{A}, 3), |\mathcal{A}| \leq s, \|\boldsymbol{\Sigma}^{1/2} \boldsymbol{\zeta}\|_2 = 1} \|\boldsymbol{\Sigma}^{1/2} \boldsymbol{g}\|_\infty \|\boldsymbol{\zeta}\|_2] \cdot 4\sqrt{s} \\ &\stackrel{(ii)}{\leq} \mathbb{E}[\|\boldsymbol{\Sigma}^{1/2} \boldsymbol{g}\|_\infty] \cdot 4\sqrt{\frac{s}{\kappa_0}}, \end{aligned}$$

where the inequality (i) follows from the property

$$\|\boldsymbol{\zeta}\|_1 = \|\boldsymbol{\zeta}_{\mathcal{A}}\|_1 + \|\boldsymbol{\zeta}_{\mathcal{A}^c}\|_1 \leq 4\|\boldsymbol{\zeta}_{\mathcal{A}}\|_1 \leq 4\sqrt{|\mathcal{A}|}\|\boldsymbol{\zeta}_{\mathcal{A}}\|_2 \leq 4\sqrt{s}\|\boldsymbol{\zeta}\|_2,$$

and the inequality (ii) holds due to the RE(κ_0, s) condition of $\boldsymbol{\Sigma}$ and $\|\boldsymbol{\Sigma}^{1/2} \boldsymbol{\zeta}\|_2 = 1$. Then using the maximal inequality and the assumption $\sup_j \boldsymbol{\Sigma}_{jj} \leq c_0$, we can bound the Gaussian width $\mathbb{W}(\mathcal{Y}_2)$ as follows

$$\mathbb{W}(\mathcal{Y}_2) \leq \sqrt{\frac{32c_0 s \log(2d)}{\kappa_0}}. \quad (\text{B.17})$$

Combining the upper bound (B.17) of the Gaussian width $\mathbb{W}(\mathcal{Y}_2)$ with the bound (B.16), we have with probability at least $1 - e^{-\frac{c_2 n \delta^2}{\sigma^4}}$,

$$\sup_{\boldsymbol{\zeta} \in \mathcal{C}(\mathcal{A}, 3), |\mathcal{A}| \leq s} \frac{|\boldsymbol{\zeta}^\top \boldsymbol{A} \boldsymbol{\zeta}|}{\boldsymbol{\zeta}^\top \boldsymbol{\Sigma} \boldsymbol{\zeta}} \leq c_1 \sqrt{\frac{s \log(d)}{n}} + \delta,$$

for some constants c_1, c_2 . Apply the same argument to \boldsymbol{A}' and we get with probability at least $1 - e^{-\frac{c_2 K b \delta^2}{\sigma^4}}$,

$$\sup_{\boldsymbol{\zeta} \in \mathcal{C}(\mathcal{A}, 3), |\mathcal{A}| \leq s} \frac{|\boldsymbol{\zeta}^\top \boldsymbol{A}' \boldsymbol{\zeta}|}{\boldsymbol{\zeta}^\top \boldsymbol{\Sigma} \boldsymbol{\zeta}} \leq c_1 \sqrt{\frac{s \log(d)}{K}} + \sqrt{b} \delta,$$

for some constants c_1 and c_2 . Combining these two upper bounds with the inequality (B.15), we have with probability at least $1 - 2e^{-\frac{c_2 n \delta^2}{\sigma^4}}$,

$$\sup_{\boldsymbol{\zeta} \in \mathcal{C}(\mathcal{A}, 3), |\mathcal{A}| \leq s} \frac{|\boldsymbol{\zeta}^\top (\widehat{\boldsymbol{\Sigma}}^1 - \boldsymbol{\Sigma}) \boldsymbol{\zeta}|}{\boldsymbol{\zeta}^\top \boldsymbol{\Sigma} \boldsymbol{\zeta}} \leq c_1 \sqrt{\frac{s \log(d)}{n}} + \delta,$$

for some constants c_1 and c_2 . Note again that \boldsymbol{A}' -term is negligible relative to \boldsymbol{A} -term. By change of variable of δ , we have with probability at least $1 - \delta$,

$$\sup_{\boldsymbol{\zeta} \in \mathcal{C}(\mathcal{A}, 3), |\mathcal{A}| \leq s} \frac{|\boldsymbol{\zeta}^\top (\widehat{\boldsymbol{\Sigma}}^1 - \boldsymbol{\Sigma}) \boldsymbol{\zeta}|}{\boldsymbol{\zeta}^\top \boldsymbol{\Sigma} \boldsymbol{\zeta}} \leq c_1 \sqrt{\frac{s \log(d/\delta)}{n}}.$$

Then there exists some constant c_3 , if $n \geq c_3 s \log(d/\delta)$,

$$\begin{aligned} &\left| 1 + \frac{\boldsymbol{\Delta}_{\cdot j}^\top (\widehat{\boldsymbol{\Sigma}}^1 - \boldsymbol{\Sigma}) \boldsymbol{\Delta}_{\cdot j}}{\boldsymbol{\Delta}_{\cdot j}^\top \boldsymbol{\Sigma} \boldsymbol{\Delta}_{\cdot j}} \right| \\ &\geq 1 - \sup_{\boldsymbol{\zeta} \in \mathcal{C}(\mathcal{A}, 3), |\mathcal{A}| \leq s} \frac{|\boldsymbol{\zeta}^\top (\widehat{\boldsymbol{\Sigma}}^1 - \boldsymbol{\Sigma}) \boldsymbol{\zeta}|}{\boldsymbol{\zeta}^\top \boldsymbol{\Sigma} \boldsymbol{\zeta}} \\ &\geq 1/2, \end{aligned}$$

with probability at least $1 - \delta$.

By the RE(κ_0, s) condition of Σ and the property that $\Delta_{\cdot j} \in \mathcal{C}(\mathcal{A}, 3)$ for some $|\mathcal{A}| \leq s$, we have

$$\Delta_{\cdot j}^\top \Sigma \Delta_{\cdot j} \geq \kappa_0 \|\Delta_{\cdot j}\|_2^2.$$

Thus we have with probability at least $1 - \delta$,

$$\begin{aligned} \mathcal{L}_1(\mathbf{W}^* + \Delta) - \mathcal{L}_1(\mathbf{W}^*) - \langle \nabla \mathcal{L}_1(\mathbf{W}^*), \Delta \rangle &\geq \frac{1}{4} \sum_{j=1}^{K-1} \Delta_{\cdot j}^\top \Sigma \Delta_{\cdot j} \\ &\geq \frac{\kappa_0}{4} \sum_{j=1}^{K-1} \|\Delta_{\cdot j}\|_2^2 = \frac{\kappa_0}{4} \|\Delta\|_{2,2}^2, \end{aligned}$$

if $n \geq c_3 s \log(d/\delta)$ for some constant c_3 . □