

Prediction and estimation consistency of sparse multi-class penalized optimal scoring

IRINA GAYNANOVA

*Department of Statistics
Texas A&M University
MS 3143
College Station, TX 77843
E-mail: irinag@stat.tamu.edu*

Sparse linear discriminant analysis via penalized optimal scoring is a successful tool for classification in high-dimensional settings. While the variable selection consistency of sparse optimal scoring has been established, the corresponding prediction and estimation consistency results have been lacking. We bridge this gap by providing probabilistic bounds on out-of-sample prediction error and estimation error of multi-class penalized optimal scoring allowing for diverging number of classes.

Keywords: classification, high-dimensional regression, lasso, linear discriminant analysis.

1. Introduction

Sparse linear discriminant analysis has been proven to be a successful tool for classification in high-dimensional settings [6, 8, 23, 38]. While multiple formulations have been proposed, exploiting the connection between the linear discriminant analysis and optimal scoring problem [15, 16] leads to a particularly attractive regularization due to the least squares loss function.

Let $(\mathbf{x}_i, \mathbf{z}_i)$, $i = 1, \dots, n$, be independent pairs, where $\mathbf{x}_i \in \mathbb{R}^p$ is a vector of features, and $\mathbf{z}_i \in \{0, 1\}^K$ is a vector indicating class membership, $z_{ik} = 1$ if i th sample belongs to class $k \in \{1, \dots, K\}$ and $z_{ik} = 0$ otherwise. Let $\mathbf{X} \in \mathbb{R}^{n \times p}$ be a column-centered data matrix, $\mathbf{Z} \in \mathbb{R}^{n \times K}$ be the corresponding class indicator matrix and n_k be the number of samples in class k . Let $\mathbf{1} \in \{1\}^K$ be a vector of ones. The unpenalized optimal scoring problem [16] is formulated as

$$\begin{aligned} & \underset{\Theta, \mathbf{B}}{\text{minimize}} \|\mathbf{Z}\Theta - \mathbf{X}\mathbf{B}\|_F^2 \\ & \text{subject to } n^{-1}\Theta^\top \mathbf{Z}^\top \mathbf{Z}\Theta = \mathbf{I}_{K-1}, \quad \Theta^\top \mathbf{Z}^\top \mathbf{Z}\mathbf{1} = \mathbf{0}, \end{aligned}$$

where $\mathbf{B} = [\beta_1 \dots \beta_{K-1}] \in \mathbb{R}^{p \times (K-1)}$ is the matrix of feature coefficients, and $\Theta \in \mathbb{R}^{K \times (K-1)}$ is the matrix of scores. It is shown in [16] that linear discriminant analysis can be carried out by solving unpenalized optimal scoring problem.

In the special case of two classes, $K = 2$, the solution for the vector of scores is $\widehat{\theta} = (\sqrt{n_2/n_1}, -\sqrt{n_1/n_2})^\top$ up to a sign. Defining $\mathbf{Y} = \mathbf{Z}\widehat{\theta}$, the optimal scoring problem reduces to the linear regression problem. Given the success of lasso [34] in high-dimensional linear regression, [23] consider the penalized optimal scoring problem

$$\widehat{\beta} = \underset{\beta}{\text{argmin}} \{(2n)^{-1}\|\mathbf{Y} - \mathbf{X}\beta\|_2^2 + \lambda\|\beta\|_1\}. \quad (1)$$

Further generalizations to the copula models [14], tensor data [30] as well as the multi-class case [8, 10, 26, 39] have been considered.

While the prediction and estimation consistency of lasso estimator in linear regression has been well-studied, see for example [1, 5, 9, 13, 32, 41] and references therein, the theoretical analysis of (1) and related extensions have been primarily focused on variable selection consistency [10, 11, 14, 19, 23, 30]. The latter requires the use of irrepresentable condition [42], which significantly limits the amount of correlations allowed in \mathbf{X} , and is more restrictive than conditions needed for the prediction consistency [35].

There are several reasons for the gap between theoretical understanding of sparse optimal scoring and lasso. First, the theory for lasso relies on the linear model assumption

$$y_i = \mathbf{x}_i^\top \boldsymbol{\beta} + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, \sigma^2), \quad \varepsilon_i \text{ independent from } \mathbf{x}_i, \quad (2)$$

that is for each sample i , the continuous response y_i is generated conditionally on the covariates \mathbf{x}_i . Model (2) does not hold for optimal scoring. In particular, the covariates \mathbf{x}_i are generated conditionally on the class membership encoded by discrete y_i . Secondly, since the covariates are random, it is of interest to investigate the expected out-of-sample prediction risk rather than in-sample prediction risk typically considered in linear regression literature [1, 9, 17]. Specifically, let $\mathbf{x} \in \mathbb{R}^p$ be a new vector of covariates with the same distribution as \mathbf{x}_i , and for $K = 2$ let $\boldsymbol{\beta}^*$ be the population matrix of coefficients, then the expected out-of-sample prediction risk is defined as

$$R(\hat{\boldsymbol{\beta}}) = \mathbb{E}_{\mathbf{x}}\{\|\mathbf{x}^\top(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)\|_2^2\} = (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)^\top \mathbb{E}(\mathbf{x}\mathbf{x}^\top)(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*),$$

whereas the in-sample prediction risk is defined as $\|\mathbf{X}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)\|_2^2$. Finally, defining the residual terms as $\boldsymbol{\varepsilon} := \mathbf{Y} - \mathbf{X}\boldsymbol{\beta}^*$, the residuals in $\boldsymbol{\varepsilon}$ and the covariates in \mathbf{X} are not independent, again in contrast to linear model (2). These challenges prevent direct application of lasso results to (1).

In this work we address these challenges and bridge the existing gap in theoretical understanding of sparse optimal scoring. Specifically, our work makes the following contributions:

- Compared to existing research specific to $K = 2$ case [21], we consider a multi-class framework, and show that the matrix of optimal scores $\hat{\boldsymbol{\Theta}}$ can be expressed in a closed form up to an orthogonal transformation (Lemma 2). This allows us to formulate a coordinate-sparse multi-class optimal scoring problem as the penalized multiple response linear regression problem, thus enabling the subsequent theoretical analysis. We believe this result is of independent interest.

- We derive the concentration bound for the maximal row ℓ_2 norm of $n^{-1}\mathbf{X}^\top \mathbf{E}$ (Theorem 3), where $\mathbf{E} := \mathbf{Y} - \mathbf{X}\boldsymbol{\beta}^*$ is the matrix of residuals. The key difficulties in deriving this bound is the non-gaussianity of \mathbf{X} and \mathbf{E} , and the lack of independence between \mathbf{X} and \mathbf{E} . The corresponding proof is the key theoretical contribution of this work.

- We derive out-of-sample prediction and estimation bounds for sparse multi-class optimal scoring problem which allow both the number of features p and the number of classes K to grow with the sample size n . The corresponding results for the estimator in (1) follow as a special case when $K = 2$. We derive bounds of two types, that are typically called slow-rate bounds and fast-rate bounds in the literature, we refer to [2, 9] for the discussion. Slow-rate bounds make no assumptions on the correlation structure of \mathbf{X} or the sparsity of the population matrix of coefficients $\boldsymbol{\beta}^*$, whereas fast-rate bounds lead to faster convergence rates, but rely on exact sparsity of $\boldsymbol{\beta}^*$ and restricted eigenvalue condition [1].

- We prove that out-of-sample prediction consistency implies classification consistency, and derive finite-sample bounds on misclassification error rate of multi-class penalized optimal scoring problem in terms of corresponding error rate of population Fisher's linear discriminant analysis rule.

1.1. Relations to Existing Literature

The variable selection consistency of estimator in (1) has been established in [19, 23], whereas the variable selection consistency for the estimator in the multi-class case has been established in [10, 11]. While the estimation consistency can be established under the same conditions, the proofs rely on irrepresentability condition. To our knowledge the results on prediction and estimation consistency of (1) without irrepresentability condition are lacking, with the exception of a recent work by Li and Jia [21].

In [21], Li and Jia establish ℓ_2 estimation consistency of penalized optimal scoring when $K = 2$. Our results and analysis differ in several ways. Most importantly, we consider the multi-class case, and allow the number of classes K to grow with the sample size. This generalization is far from trivial, and requires establishing score invariance (Lemma 1), derivation of the explicit form of the scores (Lemma 2) as well as a new proof of the concentration bound for $n^{-1}\mathbf{X}^\top \mathbf{E}$ term (Theorem 3). Theorem 3 applies to the two-class case as well, but our proof allows to explicitly characterize the dependence of constants on model parameters

and is significantly reduced compared to the proof in [21]. Secondly, in addition to ℓ_2 consistency, we establish the bounds on expected out-of-sample prediction error, where expectation is taken with respect to a new vector of features $\mathbf{x} \in \mathbb{R}^p$. Similar distinction is made in [7], where the difference between mean squared prediction error and estimated mean squared prediction error is discussed. The out-of-sample prediction bounds are not present in [21], largely due to the latter focus on fast-rate bounds. In contrast, we derive both fast-rate and slow-rate bounds. The main advantage of the slow-rate bounds is that they do not require either sparsity assumption or the restricted eigenvalue condition, we refer to [2, 9] for the discussion of the two types of bounds. As part of the slow-rate bound derivation, we demonstrate that the norm of $\widehat{\mathbf{B}}$ can always be bounded by a constant times the norm of \mathbf{B}^* . While the proof is rather simple, we found that this fact was not explicitly stated in the literature, and therefore could be of independent interest. Finally, the ℓ_2 estimation consistency in [21] is established explicitly under the restricted eigenvalue condition on \mathbf{X} . While the authors state that “a few class of matrices have been proved to satisfy the restricted eigenvalue condition with high probability” and refer to [31] for corresponding results for gaussian designs, these results are neither incorporated into the analysis nor is \mathbf{X} gaussian in optimal scoring. We show that the entries of \mathbf{X} are marginally sub-gaussian with explicit characterization of sub-gaussian constant (Lemma 3), and correspondingly rely on results of [33, 43] to establish restricted eigenvalue condition with high probability. We also incorporate these bounds within the analysis.

1.2. Notation

For two scalars $a, b \in \mathbb{R}$, we let $a \vee b = \max(a, b)$. For a vector $\mathbf{v} \in \mathbb{R}^p$, we define ℓ_1 -norm as $\|\mathbf{v}\|_1 = \sum_{i=1}^p |v_i|$, ℓ_2 -norm as $\|\mathbf{v}\|_2 = (\sum_{i=1}^p v_i^2)^{1/2}$ and ℓ_∞ norm as $\|\mathbf{v}\|_\infty = \max_i |v_i|$. We use $\mathbf{1} \in \mathbb{R}^p$ to denote a vector of ones, $\mathbf{0} \in \mathbb{R}^p$ to denote a vector of zeros, and $\mathbf{e}_j \in \mathbb{R}^p$ to denote a unit-norm vector with j th coordinate equal to one. For scalar $a \in \mathbb{R}$, we write $\{a\}_l$ to denote a row-vector of length l with each element equal to a . For a matrix $\mathbf{A} \in \mathbb{R}^{n \times p}$, we let $\|\mathbf{A}\|_{\infty,2} = \max_i (\sum_{j=1}^p a_{ij}^2)^{1/2}$, $\|\mathbf{A}\|_{1,2} = \sum_{i=1}^n (\sum_{j=1}^p a_{ij}^2)^{1/2}$, $\|\mathbf{A}\|_1 = \sum_{i=1}^n \sum_{j=1}^p |a_{ij}|$, $\|\mathbf{A}\|_2 = \sup_{\mathbf{x}: \|\mathbf{x}\|_2=1} \|\mathbf{A}\mathbf{x}\|_2$, $\|\mathbf{A}\|_F = (\sum_{i=1}^n \sum_{j=1}^p a_{ij}^2)^{1/2}$ and $\|\mathbf{A}\|_\infty = \max_{i,j} |a_{ij}|$. We use \mathbf{I} to denote the identity matrix. For vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^p$, we use $\langle \mathbf{a}, \mathbf{b} \rangle = \mathbf{a}^\top \mathbf{b}$ to define vector inner product, and for matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times p}$ we use $\langle \mathbf{A}, \mathbf{B} \rangle = \text{Tr}(\mathbf{A}^\top \mathbf{B})$ to define matrix inner product. For a sequence of scalars b_1, \dots, b_n, \dots , we use $b_n = o(a_n)$ if $\lim_{n \rightarrow \infty} (b_n/a_n) = 0$ and $b_n = O(a_n)$ if $\lim_{n \rightarrow \infty} (b_n/a_n) = C$ for some finite constant C . For a sequence of random variables x_1, \dots, x_n, \dots , we use $x_n = O_p(a_n)$ if for any $\varepsilon > 0$ $P(|x_n|/a_n < \varepsilon) \rightarrow 0$ as $n \rightarrow \infty$, and $x_n = O_p(a_n)$ if for any $\varepsilon > 0$ there exists M_ε such that $P(|x_n|/a_n > M_\varepsilon) < \varepsilon$ for all n . For random variable t , we use $\|t\|_{\psi_2} = \sup_{p \geq 1} p^{-1/2} (\mathbb{E}|t|^p)^{1/p}$ for sub-gaussian norm of t , and $\|t\|_{\psi_1} = \sup_{p \geq 1} p^{-1} (\mathbb{E}|t|^p)^{1/p}$ for sub-exponential norm of t .

1.3. Paper organization

The rest of the manuscript is organized as follows. In Section 2 we consider penalized optimal scoring for the multi-class case, and demonstrate that coordinate-sparse multi-class optimal scoring problem can be formulated as a multiple response penalized linear regression problem. In Section 3 we derive deterministic bounds for expected out-of sample prediction error and ℓ_2 estimation error of sparse optimal scoring. In Section 4 we derive concentration bound for the maximal row ℓ_2 norm of $n^{-1} \mathbf{X}^\top \mathbf{E}$, which subsequently allows us to derive probabilistic slow-rate and fast-rate bounds. In Section 5, we prove classification consistency by deriving finite-sample bounds on misclassification error rate. In Section 6 we conclude with discussion. All the proofs are deferred to Appendix.

2. Multi-class penalized optimal scoring

We consider multi-class penalized optimal scoring problem

$$\begin{aligned} & \underset{\Theta, \mathbf{B}}{\text{minimize}} \{ (2n)^{-1} \|\mathbf{Z}\Theta - \mathbf{X}\mathbf{B}\|_F^2 + \lambda \text{Pen}(\mathbf{B}) \} \\ & \text{subject to } \Theta^\top \mathbf{Z}^\top \mathbf{Z} \Theta = n\mathbf{I}_{K-1}, \quad \Theta^\top \mathbf{Z}^\top \mathbf{Z} \mathbf{1} = \mathbf{0}, \end{aligned} \tag{3}$$

where $\text{Pen}(\mathbf{B}) : \mathbb{R}^{p \times (K-1)} \rightarrow [0, \infty)$ is a penalty function. For example, [16] uses $\text{Pen}(\mathbf{B}) = \text{Tr}(\mathbf{B}^\top \mathbf{A} \mathbf{B})$ for some positive definite matrix \mathbf{A} , [8] use $\text{Pen}(\mathbf{B}) = \|\mathbf{B}\|_1$, and for $K = 2$ [23] use $\text{Pen}(\boldsymbol{\beta}) = \|\boldsymbol{\beta}\|_1$. We first show that if the penalty function is invariant with respect to orthogonal transformation, then any matrix of scores within the constraint set will lead to global solution of (3).

Lemma 1. *Let $\text{Pen}(\mathbf{B}) = \text{Pen}(\mathbf{B}\mathbf{R})$ hold for any $\mathbf{B} \in \mathbb{R}^{p \times (K-1)}$ and any orthogonal matrix $\mathbf{R} \in \mathbb{R}^{(K-1) \times (K-1)}$. Let $\tilde{\Theta} \in \mathbb{R}^{K \times (K-1)}$ be such that $\tilde{\Theta}^\top \mathbf{Z}^\top \mathbf{Z} \tilde{\Theta} = n\mathbf{I}_{K-1}$, $\tilde{\Theta}^\top \mathbf{Z}^\top \mathbf{Z} \mathbf{1} = \mathbf{0}$, and let*

$$\mathbf{B}_{\tilde{\Theta}} = \operatorname{argmin}_{\mathbf{B}} \left\{ (2n)^{-1} \|\mathbf{Z} \tilde{\Theta} - \mathbf{X} \mathbf{B}\|_F^2 + \lambda \text{Pen}(\mathbf{B}) \right\}. \quad (4)$$

Then the pair $(\tilde{\Theta}, \mathbf{B}_{\tilde{\Theta}})$ attains global minimum of (3).

Since any matrix $\tilde{\Theta}$ that satisfies the constraints leads to the pair $(\tilde{\Theta}, \mathbf{B}_{\tilde{\Theta}})$ that minimizes the objective function, we next show that such a matrix can be constructed explicitly based on the sample sizes n_k .

Lemma 2. *Let $\tilde{\Theta} \in \mathbb{R}^{K \times (K-1)}$ have columns $\tilde{\Theta}_l \in \mathbb{R}^K$, $l = 1, \dots, K-1$, defined as*

$$\tilde{\Theta}_l = \left(\left\{ \sqrt{\frac{nn_{l+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}} \right\}_l, -\sqrt{\frac{n \sum_{i=1}^l n_i}{n_{l+1} \sum_{i=1}^{l+1} n_i}}, \{0\}_{K-1-l} \right)^\top.$$

Then $\tilde{\Theta}^\top \mathbf{Z}^\top \mathbf{Z} \tilde{\Theta} = n\mathbf{I}_{K-1}$ and $\tilde{\Theta}^\top \mathbf{Z}^\top \mathbf{Z} \mathbf{1} = \mathbf{0}$.

Lemmas 1 and 2 show that to solve a penalized optimal scoring problem with orthogonally invariant penalty function, it is sufficient to fix the scores according to Lemma 2, and only consider problem (4), which has the same form as a penalized multiple response linear regression problem. The condition on the penalty function is satisfied by many commonly-used penalties, for example by $\text{Pen}(\mathbf{B}) = \text{Tr}(\mathbf{B}^\top \mathbf{A} \mathbf{B})$ for any K and by $\text{Pen}(\boldsymbol{\beta}) = \|\boldsymbol{\beta}\|_1$ for $K = 2$. When $K > 2$, the element-wise sparsity penalty $\text{Pen}(\mathbf{B}) = \|\mathbf{B}\|_1$ does not satisfy the condition, however the coordinate-wise sparsity penalty $\text{Pen}(\mathbf{B}) = \|\mathbf{B}\|_{1,2}$ does. While the difference between element-wise and coordinate-wise sparsity may seem minor, we argue that coordinate-wise sparsity is preferable in the discriminant analysis context. Similar argument is made in [3] for the principal component analysis. When $\text{Pen}(\mathbf{B}) = \|\mathbf{B}\|_1$, each column of $\mathbf{B}_{\tilde{\Theta}}$ is sparse, however the rows are not necessarily sparse which means both that the individual features are not completely eliminated from the classification rule, and that the sparsity is not preserved under orthogonal transformation. In contrast, when $\text{Pen}(\mathbf{B}) = \|\mathbf{B}\|_{1,2}$, all columns of $\mathbf{B}_{\tilde{\Theta}}$ share the same sparsity pattern leading to sparse rows, and consequently feature elimination. Therefore, we let $\mathbf{Y} = \mathbf{Z} \tilde{\Theta}$ and define the solution to sparse multi-class penalized optimal scoring as

$$\hat{\mathbf{B}} = \operatorname{argmin}_{\mathbf{B}} \left\{ (2n)^{-1} \|\mathbf{Y} - \mathbf{X} \mathbf{B}\|_F^2 + \lambda \|\mathbf{B}\|_{1,2} \right\}. \quad (5)$$

When $K = 2$, (5) reduces to (1). When $K > 2$, (5) can be rewritten in a form equivalent to sparse linear discriminant analysis proposal of [10], although the latter does not draw connections to optimal scoring. In the rest of the paper, we derive bounds on expected out-of-sample prediction risk and estimation error of estimator in (5).

3. Deterministic bounds

In this section we derive out-of-sample prediction and ℓ_2 estimation bounds for (5) that hold deterministically under certain conditions on \mathbf{X} and λ . We first review the explicit form of the matrix of population discriminant vectors \mathbf{B}^* .

Let $\mathbf{x}_i \in \mathcal{C}_k$ denote that sample i belongs to class $k \in \{1, \dots, K\}$, and let $\pi_k = \text{pr}(\mathbf{x}_i \in \mathcal{C}_k)$, $\boldsymbol{\mu}_k = \mathbb{E}(\mathbf{x}_i | \mathbf{x}_i \in \mathcal{C}_k)$, $\boldsymbol{\Sigma}_W = \text{Cov}(\mathbf{x}_i | \mathbf{x}_i \in \mathcal{C}_k)$. Let $\boldsymbol{\Sigma}_T$ be the marginal covariance matrix such that $\text{Cov}(\mathbf{x}_i) = \boldsymbol{\Sigma}_T$. [10] show that the population matrix of canonical vectors can be expressed as $\boldsymbol{\Sigma}_T^{-1} \boldsymbol{\Delta} \mathbf{R}$, where $\boldsymbol{\Delta} \in \mathbb{R}^{K \times (K-1)}$ is the matrix of orthogonal mean contrasts between K classes with r th column defined as

$$\boldsymbol{\Delta}_r = \frac{\sqrt{\pi_{r+1}} \sum_{k=1}^r \pi_k (\boldsymbol{\mu}_k - \boldsymbol{\mu}_{r+1})}{\sqrt{\sum_{k=1}^r \pi_k \sum_{k=1}^{r+1} \pi_k}}, \quad (6)$$

and \mathbf{R} is the $(K-1) \times (K-1)$ orthogonal matrix of eigenvectors of $\boldsymbol{\Delta}^\top \boldsymbol{\Sigma}_T^{-1} \boldsymbol{\Delta}$. Moreover, since the classification rule is invariant to orthogonal transformations, any orthogonal matrix \mathbf{R} will lead to equivalent classification rule. The orthogonal invariance of \mathbf{B}^* mimics the orthogonal invariance explored in Lemma 1, which is not by chance. Our choice of $\tilde{\boldsymbol{\Theta}}$ in Lemma 2 is such that $\mathbb{E}(n^{-1} \mathbf{X}^\top \mathbf{Z} \tilde{\boldsymbol{\Theta}}) = \boldsymbol{\Delta} + o(1)$ (see Lemma 4 in the Appendix), and we fix $\mathbf{B}^* = \boldsymbol{\Sigma}_T^{-1} \boldsymbol{\Delta}$ throughout the manuscript. A different choice of $\tilde{\boldsymbol{\Theta}}$ leads to equivalent conclusions by applying corresponding orthogonal transformation to \mathbf{B}^* . In a special case of two classes, $\beta^* = \boldsymbol{\Sigma}_T^{-1} \boldsymbol{\Delta}_1 = \sqrt{\pi_1 \pi_2} \boldsymbol{\Sigma}_T^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)$, which coincides with discriminant analysis direction considered in the literature [6, 23].

Let $\mathbf{x} \in \mathbb{R}^p$ be a new vector of covariates with the same distribution as \mathbf{x}_i . Given $\mathbf{B}^* = \boldsymbol{\Sigma}_T^{-1} \boldsymbol{\Delta}$, we aim to derive bounds on expected out-of-sample prediction error defined as

$$R(\hat{\mathbf{B}}) := \mathbb{E}_{\mathbf{x}} \|\mathbf{x}^\top (\mathbf{B}^* - \hat{\mathbf{B}})\|_F^2 = \text{Tr}\{(\mathbf{B}^* - \hat{\mathbf{B}})^\top \boldsymbol{\Sigma}_T (\mathbf{B}^* - \hat{\mathbf{B}})\},$$

and the estimation error $\|\hat{\mathbf{B}} - \mathbf{B}^*\|_F^2$. Throughout, we define residual terms as $\mathbf{E} = [\varepsilon_1 \dots \varepsilon_{K-1}] := \mathbf{Y} - \mathbf{X} \mathbf{B}^*$, which allows direct comparison with lasso bounds. Since (5) is of the same form as penalized multiple-response linear regression problem with group-lasso penalty, the proofs of deterministic bounds for (5) follow the proofs of deterministic bounds for lasso with extra triangle inequality to handle out-of-sample rather than in-sample prediction error. Because these bounds are deterministic, we follow the terminology in [20] by differentiating *penalty* and *sparsity* deterministic bounds. In Section 4, we use these bound to derive probabilistic *slow-rate* and *fast-rate* bounds correspondingly.

3.1. Penalty bounds

We start by providing a deterministic bound for in-sample prediction error $\|\mathbf{X}(\hat{\mathbf{B}} - \mathbf{B}^*)\|_F^2$. This bound makes no assumption on \mathbf{X} or on the sparsity of \mathbf{B}^* .

Theorem 1. *If $\lambda \geq \frac{1}{n} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$, then*

$$\frac{1}{n} \|\mathbf{X}(\hat{\mathbf{B}} - \mathbf{B}^*)\|_F^2 \leq \left(\lambda + \frac{1}{n} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2} \right) \|\mathbf{B}^*\|_{1,2}.$$

If, in addition, $\lambda \geq \frac{2}{n} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$, then $\|\hat{\mathbf{B}}\|_{1,2} \leq 3 \|\mathbf{B}^\|_{1,2}$.*

The first part of Theorem 1 is a well-known deterministic bound for lasso and group-lasso, see for example [20, 32]. The second part of Theorem 1 shows that the mixed ℓ_1/ℓ_2 norm of $\hat{\mathbf{B}}$ can be bounded by the ℓ_1/ℓ_2 norm of \mathbf{B}^* up to a constant. The latter, in particular, allows to bound expected out-of-sample prediction error following the proof similar to [7] for the constrained lasso case. In the constrained formulation, the bound on the norm of $\hat{\mathbf{B}}$ is immediate by choosing a constraint parameter that is as large as $\|\mathbf{B}^*\|_{1,2}$. We show that by choosing the tuning parameter λ large enough, similar bound holds for penalized formulation. Combining the norm bound with the in-sample prediction error bound leads to the deterministic bound on expected out-of-sample prediction error. The bound on estimation error follows by assuming positive definiteness of population marginal covariance matrix $\boldsymbol{\Sigma}_T$.

Corollary 1. Let $\lambda \geq \frac{2}{n} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$. Then

$$\text{Tr}\{(\widehat{\mathbf{B}} - \mathbf{B}^*)^\top \Sigma_T (\widehat{\mathbf{B}} - \mathbf{B}^*)\} \leq \frac{3}{2} \lambda \|\mathbf{B}^*\|_{1,2} + 16 \|\mathbf{B}^*\|_{1,2}^2 \|\frac{1}{n} \mathbf{X}^\top \mathbf{X} - \Sigma_T\|_\infty.$$

If, in addition, $\lambda_{\min}(\Sigma_T) > 0$, then

$$\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F^2 \leq \frac{3}{2\lambda_{\min}(\Sigma_T)} \lambda \|\mathbf{B}^*\|_{1,2} + \frac{16}{\lambda_{\min}(\Sigma_T)} \|\mathbf{B}^*\|_{1,2}^2 \|\frac{1}{n} \mathbf{X}^\top \mathbf{X} - \Sigma_T\|_\infty.$$

The bounds of Corollary 1 are deterministic, and therefore depend on \mathbf{X} via $\|n^{-1} \mathbf{X}^\top \mathbf{X} - \Sigma_T\|_\infty$. In Section 4.2, Theorem 4, we provide the corresponding probabilistic bounds which rely on concentration inequality for $n^{-1} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$ (Theorem 3) and concentration inequality for $\|n^{-1} \mathbf{X}^\top \mathbf{X} - \Sigma_T\|_\infty$ (Lemma 7).

3.2. Sparsity bounds

To derive the sparsity bound, we make additional assumption on \mathbf{B}^* .

Assumption 1 (Sparsity). \mathbf{B}^* is row-sparse with the support $S = \{j : \|e_j^\top \mathbf{B}^*\|_2 \neq 0\}$ with $s = \text{card}(S)$.

As in lasso, we also use restricted eigenvalue condition on the design matrix [1].

Definition 1 (Restricted eigenvalue condition). A $q \times p$ matrix \mathbf{Q} satisfies restricted eigenvalue condition $RE(s, c)$ with parameter $\gamma_{\mathbf{Q}} = \gamma(s, c, \mathbf{Q})$ if for all sets $S \subset \{1, \dots, p\}$ with $\text{card}(S) \leq s$, and for all $\mathbf{a} \in \mathcal{C}(S, c) = \{\mathbf{a} \in \mathbb{R}^p : \|\mathbf{a}_{S^c}\|_1 \leq c \|\mathbf{a}_S\|_1\}$ it holds that

$$\|\mathbf{Q}\mathbf{a}\|_2^2 \geq \frac{\|\mathbf{a}_S\|_2^2}{\gamma_{\mathbf{Q}}}.$$

In the group-lasso case, this condition is generalized to allow for the mixed ℓ_1/ℓ_2 norms, see for example [22]. In penalized optimal scoring, the generalization is needed when the number of classes $K > 2$.

Definition 2 (Group restricted eigenvalue condition). A $q \times p$ matrix \mathbf{Q} satisfies restricted eigenvalue condition $RE(s, c, K)$ with parameter $\gamma_{\mathbf{Q}} = \gamma(s, c, K, \mathbf{Q})$ if for all sets $S \subset \{1, \dots, p\}$ with $\text{card}(S) \leq s$, and for all $\mathbf{A} \in \mathcal{C}(S, c, K) = \{\mathbf{A} \in \mathbb{R}^{p \times (K-1)} : \|\mathbf{A}_{S^c}\|_{1,2} \leq c \|\mathbf{A}_S\|_{1,2}\}$ it holds that

$$\|\mathbf{Q}\mathbf{A}\|_F^2 \geq \frac{\|\mathbf{A}_S\|_F^2}{\gamma_{\mathbf{Q}}}.$$

When $K = 2$, Definitions 1 and 2 coincide. We next state deterministic sparsity bounds that hold whenever \mathbf{X} satisfies restricted eigenvalue condition.

Theorem 2. Under Assumption 1, if $\lambda \geq \frac{2}{n} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$ and $n^{-1/2} \mathbf{X}$ satisfies $RE(s, 3, K)$ with parameter $\gamma_{\mathbf{X}} = \gamma(s, 3, K, n^{-1/2} \mathbf{X})$, then

$$\frac{1}{n} \|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F^2 \leq \frac{9}{4} \gamma_{\mathbf{X}} s \lambda^2; \quad \|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F \leq \frac{15}{2} \gamma_{\mathbf{X}} \sqrt{s} \lambda \quad \text{and} \quad \|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{1,2} \leq 6 \gamma_{\mathbf{X}} s \lambda.$$

The bounds of Theorem 2 are well-known for lasso and group-lasso, see for example [1, 17, 27]. Compared to these results, our interest is in expected out-of-sample prediction error, and we provide corresponding bound in Corollary 2. This bound can be obtained in two ways. On the one hand, we can use triangle inequality as in Corollary 1. On the other hand, since restricted eigenvalue condition allows to directly bound estimation error, we can use that bound through the maximal eigenvalue of Σ_T . If the maximal eigenvalue of Σ_T can be treated as constant, the second approach leads to tighter probabilistic bounds.

Corollary 2. Under Assumption 1, if $\lambda \geq \frac{2}{n} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$ and $n^{-1/2} \mathbf{X}$ satisfies $RE(s, 3, K)$ with parameter $\gamma_{\mathbf{X}} = \gamma(s, 3, K, n^{-1/2} \mathbf{X})$, then

$$\text{Tr}\{(\widehat{\mathbf{B}} - \mathbf{B}^*)^\top \Sigma_T (\widehat{\mathbf{B}} - \mathbf{B}^*)\} \leq \min \left\{ \frac{9}{4} \gamma_{\mathbf{X}} s \lambda^2 + 36 \gamma_{\mathbf{X}}^2 s^2 \lambda^2 \|\frac{1}{n} \mathbf{X}^\top \mathbf{X} - \Sigma_T\|_\infty, \lambda_{\max}(\Sigma_T) 57 \gamma_{\mathbf{X}}^2 s \lambda^2 \right\}.$$

4. Probabilistic bounds

Both Corollary 1 and 2 rely on the deterministic condition for λ , that is $\lambda \geq \frac{1}{n} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$. Therefore, to derive corresponding probabilistic bounds, we need to derive a concentration bound for $\frac{1}{n} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$. This bound is provided in Section 4.1 and is the central result of the paper. The corresponding probabilistic slow-rate and fast-rate bounds are stated in Sections 4.2 and 4.3.

4.1. Concentration bound for $n^{-1} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$

There are several difficulties in deriving the concentration bound for $n^{-1} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$ in the context of penalized optimal scoring. First, both \mathbf{Y} and \mathbf{X} are random, and the linear model (2) for \mathbf{Y} doesn't hold. Secondly, \mathbf{X} and \mathbf{E} are not independent. These challenges prevent application of lasso results, and therefore require new derivations. For this, we make the following assumptions.

Assumption 2 (Class probabilities). $\text{pr}(\mathbf{x}_i \in \mathcal{C}_k) = \pi_k$ for $k = 1, \dots, K$ with $0 < \pi_{\min} \leq \pi_k \leq \pi_{\max} < 1$.

Assumption 3 (Normality). $\mathbf{x}_i | \mathbf{x}_i \in \mathcal{C}_k \sim \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_W)$ for all $k = 1, \dots, K$ with $\boldsymbol{\mu} = \sum_{k=1}^K \pi_k \boldsymbol{\mu}_k = 0$.

Assumption 4 (Sample size). $\log p = o(n)$

Assumption 2 requires prior group probabilities to be of the same order so that n_k grows with n for each k . Assumption 3 is typical in linear discriminant analysis, however it can be relaxed to sub-gaussianity without affecting the rates. The normality allows to express the constants in the rates through the variance terms rather than sub-gaussian parameters, which we find easier to interpret. Without loss of generality, we assume that the overall mean $\boldsymbol{\mu}$ is zero. In practice, we always column-center data matrix \mathbf{X} . Assumption 4 is a typical scaling for n and p in high-dimensional statistics.

Throughout, we use σ_j^2 to denote the diagonal elements of within-class covariance matrix $\boldsymbol{\Sigma}_W$, and define

$$\tau := \max_{j=1, \dots, p} \sqrt{\sigma_j^2 + \max_k \mu_{kj}^2}.$$

Theorem 3. Let $\lambda_0 = C\tau \sqrt{\frac{(K-1)\log(p\eta^{-1})}{n}}$ for some $\eta \in (0, 1)$ and constant $C > 0$. Under Assumptions 2–4

$$\text{pr}\left(\frac{1}{n} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2} \leq \lambda_0\right) \geq 1 - \eta.$$

Theorem 3 provides a scaling of tuning parameter with respect to the number of classes K , the sample size n and the number of variables p . While \mathbf{X} is random, and \mathbf{X} and \mathbf{E} are not independent, the scaling is the same up to constants as in lasso with fixed design, see for example [4, 17] and references therein.

We provide the sketch of the proof here to emphasize the new ideas. Since the linear model (2) for \mathbf{Y} does not hold, we explicitly take $\mathbf{E} = \mathbf{Y} - \mathbf{X}\mathbf{B}^* = \mathbf{Y} - \mathbf{X}\boldsymbol{\Sigma}_T^{-1}\boldsymbol{\Delta}$ and use triangle inequality:

$$\frac{1}{n} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2} = \left\| \frac{1}{n} \mathbf{X}^\top \mathbf{Y} - \frac{1}{n} \mathbf{X}^\top \mathbf{X} \boldsymbol{\Sigma}_T^{-1} \boldsymbol{\Delta} \right\|_{\infty,2} \leq \left\| \frac{1}{n} \mathbf{X}^\top \mathbf{Y} - \boldsymbol{\Delta} \right\|_{\infty,2} + \left\| \boldsymbol{\Delta} - \frac{1}{n} \mathbf{X}^\top \mathbf{X} \boldsymbol{\Sigma}_T^{-1} \boldsymbol{\Delta} \right\|_{\infty,2}.$$

For the first term, we take advantage of the exact form of the optimal scores derived in Lemma 2 as well as tail inequality for quadratic forms of gaussian random vectors [18]. For the second term, we prove that under Assumption 2–3, the elements of \mathbf{X} are marginally sub-gaussian with parameter at most τ (Lemma 3) and derive element-wise concentration bound for the covariance matrix of random vector with sub-gaussian elements (Lemma 7). A particular feature of the bound in Theorem 3 is that the constant $C > 0$ is not dependent on the model parameters $\boldsymbol{\Sigma}_W$, $\boldsymbol{\mu}_k$ or π_k . The bound depends on the model parameters only through τ , the dimension p and the number of classes K . This is in contrast with the results of [21] for $K = 2$ case, where the constant implicitly depends on the model parameters and $\lambda_{\max}(\boldsymbol{\Sigma}_W)$ in particular. To derive the explicit dependence through τ , we exploit the matrix decomposition of total covariance matrix $\boldsymbol{\Sigma}_T$ from [10], the Woodbury matrix identity, and the bound on $\|\boldsymbol{\Sigma}_T^{1/2}\|_\infty$ through $\|\boldsymbol{\Sigma}_T\|_\infty$ based on concavity arguments. The full proof of Theorem 3 is in the Appendix.

4.2. Slow rate bound

In this section, we derive the slow rate bounds for out-of-sample prediction and ℓ_2 estimation consistency of (5) by combining the penalty bounds of Section 3.1 with concentration bound of Theorem 3.

Theorem 4. *If $\lambda \geq C\tau\sqrt{\frac{(K-1)\log(p)}{n}}$ for some constant $C > 0$, then under Assumptions 2-4*

$$\text{Tr}\{(\widehat{\mathbf{B}} - \mathbf{B}^*)^\top \Sigma_T (\widehat{\mathbf{B}} - \mathbf{B}^*)\} = O_p \left[\left\{ 1 \vee \tau \|\mathbf{B}^*\|_{1,2} \right\} \tau \|\mathbf{B}^*\|_{1,2} \sqrt{\frac{(K-1)\log p}{n}} \right].$$

If, in addition, $\lambda_{\min}(\Sigma_T) > 0$, then

$$\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F^2 = O_p \left[\left\{ 1 \vee \tau \|\mathbf{B}^*\|_{1,2} \right\} \frac{\tau \|\mathbf{B}^*\|_{1,2}}{\lambda_{\min}(\Sigma_T)} \sqrt{\frac{(K-1)\log p}{n}} \right].$$

When $K = 2$, these results mimic Theorem 1 in [7] for the constrained lasso. Here we analyze the penalized formulation, and rely heavily on Theorem 3, which required separate derivations for optimal scoring problem. The bound allows both the number of features p and the number of classes K to grow with n . If $\|\mathbf{B}^*\|_{1,2}$ is a constant, the prediction consistency is achieved as long as $n \gg (K-1)\log p$. Otherwise, $\|\mathbf{B}^*\|_{1,2}$ is allowed to grow at a rate no faster than $\{n/(K-1)\log p\}^{1/4}$. This scaling is suboptimal compared to what would be expected in lasso with fixed design, $\{n/(K-1)\log p\}^{1/2}$, and this discrepancy is a result of considering out-of-sample rather than in-sample prediction error. We refer to [7] for further discussion.

4.3. Fast rate bound

In this section, we derive the fast rate bounds for out-of-sample prediction and ℓ_2 estimation consistency of (5) by combining the sparsity bounds of Section 3.2 with concentration bound of Theorem 3 and restricted eigenvalue condition on the marginal covariance matrix Σ_T . The latter allows us to establish that restricted eigenvalue condition holds for random $n^{-1/2}\mathbf{X}$ with high probability. For clarity, we assume that $\lambda_{\max}(\Sigma_T)$ is a constant so that the minimum in Corollary 2 is achieved with the second bound.

We present the results for the case $K = 2$ and $K > 2$ separately. When $K = 2$, we use [33, 43] to show that $\text{RE}(s, c)$ holds with high probability for sub-gaussian matrices.

Theorem 5. *Under Assumptions 1-4, if $K = 2$, $\lambda = C\tau\sqrt{\frac{\log p}{n}}$ for some constant $C > 0$, $s\log p = o(n)$, $\Sigma_T^{1/2}$ satisfies $\text{RE}(s, 9)$ and $\gamma = \gamma(s, 3, \Sigma_T^{1/2})$ according to Definition 1, then*

$$\begin{aligned} (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)^\top \Sigma_T (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*) &= O_p \left\{ \lambda_{\max}(\Sigma_T) \tau^2 \gamma^2 \frac{s\log p}{n} \right\}; \\ \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2^2 &= O_p \left(\tau^2 \gamma^2 \frac{s\log p}{n} \right). \end{aligned}$$

When $K > 2$, we need to consider a more general condition $\text{RE}(s, c, K)$. We conjecture that the results of [33, 43] can be generalized to this condition, however the explicit proof is outside of the scope of this paper. For technical clarity, we instead bound $\gamma(s, c, K, n^{-1/2}\mathbf{X})$ through the element-wise maximum $\|n^{-1}\mathbf{X}^\top \mathbf{X} - \Sigma\|_\infty$ as in [35]. This approach, however, leads to sub-optimal scaling of s ($s^2\log p = o(n)$) compared to the $K = 2$ case ($s\log p = o(n)$). This scaling is not present directly in the bounds, but rather is needed to ensure that $\gamma_{\mathbf{X}}$ can be bounded by γ .

Theorem 6. Under Assumptions 1–4, if $\lambda = C\tau\sqrt{\frac{(K-1)\log p}{n}}$ for some constant $C > 0$, $s^2 \log p = o(n)$, $\Sigma_T^{1/2}$ satisfies $RE(s, 3, K)$ and $\gamma = \gamma(s, 3, K, \Sigma_T^{1/2})$ according to Definition 2, then

$$\begin{aligned}\text{Tr}\{(\widehat{\mathbf{B}} - \mathbf{B}^*)^\top \Sigma_T (\widehat{\mathbf{B}} - \mathbf{B}^*)\} &= O_p\left\{\lambda_{\max}(\Sigma_T)\tau^2\gamma^2\frac{(K-1)s\log p}{n}\right\}; \\ \|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F^2 &= O_p\left(\tau^2\gamma^2\frac{(K-1)s\log p}{n}\right).\end{aligned}$$

Comparing Theorem 4 with Theorem 5 reveals that the key differences are in the use $\|\mathbf{B}^*\|_{1,2}$ instead of cardinality s , and slower rate $\sqrt{\log p/n}$ compared to $\log p/n$ thus justifying commonly used slow-rate and fast-rate bounds terminology. The main advantage of Theorem 4 is the lack of sparsity assumption and restricted eigenvalue condition. For more discussion on the advantages and disadvantages of these two bounds we refer to [2, 9]. Our main goal here is to show that penalized multi-class optimal scoring achieves the same consistency guarantees as lasso despite the lack of linear model (2) for \mathbf{Y} and dependency between \mathbf{X} and residuals \mathbf{E} , and this is demonstrated via statements of Theorems 4 and 6. As with Theorem 4, Theorem 6 allows the number of classes K to grow with n .

5. Classification consistency

In this section we prove classification consistency of multi-class penalized optimal scoring. In Fisher's linear discriminant analysis, the population classification rule assigns a new observation $\mathbf{x} \in \mathbb{R}^p$ according to

$$h(\mathbf{x}) = \operatorname{argmin}_k \left\{ (\mathbf{x} - \boldsymbol{\mu}_k)^\top \mathbf{B}^* (\mathbf{B}^{*\top} \Sigma_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} (\mathbf{x} - \boldsymbol{\mu}_k) - 2 \log \pi_k \right\}. \quad (7)$$

Given the estimated $\widehat{\mathbf{B}}$ from (5), the multi-class penalized optimal scoring assigns $\mathbf{x} \in \mathbb{R}^p$ according to

$$\widehat{h}(\mathbf{x}) = \operatorname{argmin}_k \left\{ (\mathbf{x} - \bar{\mathbf{x}}_k)^\top \widehat{\mathbf{B}} (\widehat{\mathbf{B}}^\top \widehat{\Sigma}_W \widehat{\mathbf{B}})^{-1} \widehat{\mathbf{B}}^\top (\mathbf{x} - \bar{\mathbf{x}}_k) - 2 \log n_k/n \right\}, \quad (8)$$

where $\bar{\mathbf{x}}_k$ is the sample mean for class k , and $\widehat{\Sigma}_W$ is the pooled sample covariance matrix. We further establish the consistency of $\widehat{h}(\mathbf{x})$ with respect to $h(\mathbf{x})$.

We first provide a projection-based interpretation of the population rule $h(\mathbf{x})$. Let $\mathbf{x} \in \mathbb{R}^p$ follow Assumption 3, and let $\mathbf{z} = \mathbf{B}^{*\top} \mathbf{x} \in \mathbb{R}^{K-1}$. Then \mathbf{z} has normal distribution conditionally on the class membership, the rule $h(\mathbf{x})$ depends on \mathbf{x} only through \mathbf{z}

$$h(\mathbf{x}) = \operatorname{argmin}_k \left\{ (\mathbf{z} - \boldsymbol{\mu}_k^\top \mathbf{B}^*) (\mathbf{B}^{*\top} \Sigma_W \mathbf{B}^*)^{-1} (\mathbf{z}^\top - \mathbf{B}^{*\top} \boldsymbol{\mu}_k) - 2 \log \pi_k \right\} := g(\mathbf{z}),$$

and $g(\mathbf{z})$ is the Bayes rule for \mathbf{z} . When $K = 2$, the population rule (7) is Bayes for \mathbf{x} under normality Assumption 3 since $\mathbf{x} \rightarrow \mathbf{B}^{*\top} \mathbf{x}$ preserves all the discriminatory information between the classes. That is, the Bayes rule based on distribution of \mathbf{x} and the Bayes rule based on distribution of \mathbf{z} coincide. When $K > 2$, the Fisher's population rule (7) is Bayes when the population means are colinear [24, p. 330], but is not Bayes in general. That is, when $K > 2$, the Bayes rule based on distribution of \mathbf{x} and the Bayes rule based on distribution of \mathbf{z} do not generally coincide. Fisher's LDA is designed to seek projections of high-dimensional data onto the most discriminative low-dimensional subspace rather than projections that minimize the error rates. The former goal can be viewed as an approximation to the latter, with direct equivalence when $K = 2$. We refer the reader to [25, Chapter 3.9] for further discussion on the difference between likelihood-based discriminant analysis and Fisher's discriminant analysis for the case $K > 2$. Since in this work we consider an optimal scoring problem which corresponds to Fisher's formulation of linear discriminant analysis, we focus on establishing consistency of $\widehat{h}(\mathbf{x})$ with respect to Fisher's population rule (7).

Consider the estimated rule $\widehat{h}(\mathbf{x})$, which is the sample plug-in Bayes rule after the transformation $\mathbf{x} \rightarrow \widehat{\mathbf{B}}^\top \mathbf{x}$ is applied. The out-of-sample prediction error $R(\widehat{\mathbf{B}}) = \mathbb{E}_{\mathbf{x}} \|\mathbf{x}^\top \mathbf{B}^* - \mathbf{x}^\top \widehat{\mathbf{B}}\|_F^2$ thus coincides with the average distance between intended transformation $\mathbf{x}^\top \mathbf{B}^*$ and estimated transformation $\mathbf{x}^\top \widehat{\mathbf{B}}$ across \mathbf{x} . Section 4

establishes that $\mathbf{x}^\top \mathbf{B}^*$ and $\mathbf{x}^\top \widehat{\mathbf{B}}$ coincide with high probability, which intuitively is a stronger statement than the agreement of corresponding classification rules. Further, we formally show that the convergence of $R(\widehat{\mathbf{B}})$ to zero in probability is indeed sufficient for convergence of the corresponding misclassification error rates.

Consider the misclassification error probability associated with the rule (7)

$$e_h = \text{pr}(h(\mathbf{x}) \neq \mathcal{C}_\mathbf{x}) = \sum_{k=1}^K \text{pr}(h(\mathbf{x}) \neq k | \mathbf{x} \in \mathcal{C}_k) \text{pr}(\mathbf{x} \in \mathcal{C}_k) = \sum_{k=1}^K \pi_k \text{pr}(h(\mathbf{x}) \neq k | \mathbf{x} \in \mathcal{C}_k) := \sum_{k=1}^K \pi_k e_{hk},$$

where e_{hk} is conditional misclassification error rate for class k . Similarly, for rule (8)

$$\widehat{e}_h = \text{pr}(\widehat{h}(\mathbf{x}) \neq \mathcal{C}_\mathbf{x}) = \sum_{k=1}^K \pi_k \text{pr}(\widehat{h}(\mathbf{x}) \neq k | \mathbf{x} \in \mathcal{C}_k) := \sum_{k=1}^K \pi_k \widehat{e}_{hk}.$$

We further show that $R(\widehat{\mathbf{B}}) \rightarrow_p 0$ implies $\widehat{e}_h \rightarrow e_h$. For clarity, we additionally assume

Assumption 5 (Class separation). $c \leq \lambda_{\min}(\Delta^\top \Sigma_W \Delta) \leq \lambda_{\max}(\Delta^\top \Sigma_W \Delta) \leq C$ for some $c, C > 0$.

We now provide intuition for Assumption 5. Let $\delta_{kl} = \mu_k - \mu_l$ be the difference in means between classes k and l . Due to the form of Δ and equality of eigenvalues between $\Delta^\top \Sigma_W \Delta$ and $\mathbf{R}^\top \Delta^\top \Sigma_W \Delta \mathbf{R}$ for any orthogonal matrix \mathbf{R} , Assumption 5 implies that (i) matrix $\mathbf{B}^{*\top} \Sigma_W \mathbf{B}^*$ is full rank; (ii) $\delta_{kl}^\top \Sigma_W \delta_{kl}$ are constants for all class pairs k, l , which leads to fixed misclassification error probability e_h . Our proof can be adapted to the case of growing e_h , however the dependence of constants on eigenvalues of $\Delta^\top \Sigma_W \Delta$ is non-trivial. Assumption 5 thus allows us to present a simplified bound, however in the proofs we do not invoke the assumption until the end.

Theorem 7. Under conditions of Theorem 4 and Assumption 5, if $K = o(p)$, then

$$\widehat{e}_h \leq e_h + O\left\{\tau^{1/2} \|\mathbf{B}^*\|_{1,2}^{1/2} (1 \vee \tau^{1/2} \|\mathbf{B}^*\|_{1,2}^{1/2}) \left(\frac{(K-1) \log p}{n}\right)^{1/4}\right\}.$$

Theorem 8. Under conditions of Theorem 6 and Assumption 5, if $K = o(p)$, then

$$\widehat{e}_h \leq e_h + O\left\{\tau \gamma \|\Sigma_W^{1/2}\|_2 \left(\frac{(K-1)s \log p}{n}\right)^{1/2}\right\}.$$

The achieved classification consistency rates are square root of the prediction rates of Theorems 4 and 6, which is consistent with existing literature for binary case [21]. The assumption $K = o(p)$ is very mild, especially for high-dimensional settings, but it allows to explicitly state $O\{\log(pK)\} = O\{\log(p)\}$ in the proofs. The extra $\log(K)$ term that appears in the proofs is likely an artifact of using the different proof technique compared to Theorems 4 and 6. The appearance of $\|\Sigma_W^{1/2}\|_2$ in Theorem 8 is also due to the different proof. Since $\|\Sigma_W^{1/2}\|_2 \leq \sqrt{\lambda_{\max}(\Sigma_T)}$, the bound of Theorem 8 is sharper.

6. Discussion

There has been significant progress in understanding the consistency of lasso and group-lasso estimators in linear regression [1, 5, 9, 13, 22, 32, 40, 41]. These results can not be applied to penalized optimal scoring problem despite the similarity between corresponding optimization problems. The key difficulty is that linear model (2) for \mathbf{Y} does not hold, and the dependency between the random covariates in \mathbf{X} and the residual terms in \mathbf{E} . In this work we overcome these challenges by using sub-exponential concentration bounds, and exploiting the decomposition of marginal covariance matrix Σ_T . While for clarity we focus on the linear optimal scoring and penalties of group-lasso type, the underlying technique can be used as a building block

for investigating consistency of other problems, for example tensor discriminant analysis [30] or optimal scoring with weighted group-lasso penalty [26]. In our treatment of the fast rate bound for the multi-class case, we rely on group restricted eigenvalue condition for the random design matrix with sub-gaussian entries. When $K = 2$, the existing results of [33, 43] show the condition holds with high probability without affecting the rates. When $K > 2$, the results of [33, 43] do not strictly apply due to the different form of the cone constraint in Definition 2, but we conjecture that the same conclusions hold. It would be of interest to have a formal justification for this conjecture.

Appendix

A.1. Technical proofs

In this section, we prove the results stated in the main text. We use $C, C_1, C_2, C_3, \dots > 0$ to denote absolute constants that do not depend on model parameters. Their values may change from line to line.

Proof of Lemma 1. Let $f = f(\Theta, B)$ denote the objective function in (3) and let (Θ^*, B^*) be the global solution to (3), that is

$$f^* = f(\Theta^*, B^*) = \frac{1}{2n} \|\mathbf{Z}\Theta^* - \mathbf{X}B^*\|_F^2 + \lambda \text{Pen}(B^*) \leq f(\Theta, B)$$

for all B and all Θ that satisfy the constraints. Since both $\tilde{\Theta}$ and Θ^* satisfy the constraints, there exist orthogonal matrix $R \in \mathbb{R}^{(K-1) \times (K-1)}$ such that $\Theta^* R = \tilde{\Theta}$. Let $\tilde{B} = B^* R$, then using orthogonal invariance of the penalty function

$$\begin{aligned} f(\tilde{\Theta}, \tilde{B}) &= \frac{1}{2n} \text{Tr}\{(\mathbf{Z}\Theta^* - \mathbf{X}B^*)RR^\top(\mathbf{Z}\Theta^* - \mathbf{X}B^*)^\top\} + \lambda \text{Pen}(B^* R) \\ &= \frac{1}{2n} \text{Tr}\{(\mathbf{Z}\Theta^* - \mathbf{X}B^*)(\mathbf{Z}\Theta^* - \mathbf{X}B^*)^\top\} + \lambda \text{Pen}(B^*) = f^*, \end{aligned}$$

that is the pair $(\tilde{\Theta}, \tilde{B})$ also attains global minimum. Then $f^* = f(\tilde{\Theta}, \tilde{B}) \leq f(\tilde{\Theta}, B)$ for all B , that is

$$\tilde{B} = B_{\tilde{\Theta}} = \underset{B}{\operatorname{argmin}} \left\{ \frac{1}{2n} \|\mathbf{Z}\tilde{\Theta} - \mathbf{X}B\|_F^2 + \lambda \text{Pen}(B) \right\},$$

and the pair $(\tilde{\Theta}, B_{\tilde{\Theta}})$ attains global minimum of (3). \square

Proof of Lemma 2. For any l, j with $j > l$

$$\begin{aligned} \frac{1}{n} \Theta_l^\top \mathbf{Z}^\top \mathbf{Z} \Theta_l &= \left(\left\{ \sqrt{\frac{n_{l+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}} \right\}_l, -\sqrt{\frac{\sum_{i=1}^l n_i}{n_{l+1} \sum_{i=1}^{l+1} n_i}}, \{0\}_{K-1-l} \right) \begin{pmatrix} n_1 & 0 & \dots & 0 \\ 0 & n_2 & \dots & 0 \\ 0 & 0 & \dots & n_K \end{pmatrix} \\ &\quad \times \left(\left\{ \sqrt{\frac{n_{l+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}} \right\}_l, -\sqrt{\frac{\sum_{i=1}^l n_i}{n_{l+1} \sum_{i=1}^{l+1} n_i}}, \{0\}_{K-1-l} \right)^\top \\ &= \left(\sqrt{\frac{n_1^2 n_{l+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}}, \dots, \sqrt{\frac{n_l^2 n_{l+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}}, -\sqrt{\frac{n_{l+1} \sum_{i=1}^l n_i}{\sum_{i=1}^{l+1} n_i}} \right)^\top \\ &\quad \times \left(\left\{ \sqrt{\frac{n_{l+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}} \right\}_l, -\sqrt{\frac{\sum_{i=1}^l n_i}{n_{l+1} \sum_{i=1}^{l+1} n_i}} \right)^\top \\ &= \frac{n_{l+1}}{\sum_{i=1}^{l+1} n_i \sum_{i=1}^{l+1} n_i} \sum_{i=1}^l n_i + \frac{\sum_{i=1}^l n_i}{\sum_{i=1}^{l+1} n_i} = \frac{\sum_{i=1}^{l+1} n_i}{\sum_{i=1}^{l+1} n_i} = 1 \end{aligned}$$

and

$$\begin{aligned}
\frac{1}{n} \Theta_l^\top \mathbf{Z}^\top \mathbf{Z} \Theta_j &= \left(\left\{ \sqrt{\frac{n_{l+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}} \right\}_l, -\sqrt{\frac{\sum_{i=1}^l n_i}{n_{l+1} \sum_{i=1}^{l+1} n_i}}, \{0\}_{K-1-l} \right) \begin{pmatrix} n_1 & 0 & \dots & 0 \\ 0 & n_2 & \dots & 0 \\ 0 & 0 & \dots & n_K \end{pmatrix} \\
&\quad \times \left(\left\{ \sqrt{\frac{n_{j+1}}{\sum_{i=1}^j n_i \sum_{i=1}^{j+1} n_i}} \right\}_j, -\sqrt{\frac{\sum_{i=1}^j n_i}{n_{j+1} \sum_{i=1}^{j+1} n_i}}, \{0\}_{K-1-j} \right)^\top \\
&= \left(\sqrt{\frac{n_1^2 n_{l+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}}, \dots, \sqrt{\frac{n_l^2 n_{l+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}}, -\sqrt{\frac{n_{l+1} \sum_{i=1}^l n_i}{\sum_{i=1}^{l+1} n_i}}, \{0\}_{j-l} \right) \\
&\quad \times \left(\left\{ \sqrt{\frac{n_{j+1}}{\sum_{i=1}^j n_i \sum_{i=1}^{j+1} n_i}} \right\}_j, -\sqrt{\frac{\sum_{i=1}^j n_i}{n_{j+1} \sum_{i=1}^{j+1} n_i}} \right)^\top \\
&= \sum_{t=1}^l \sqrt{\frac{n_t^2 n_{l+1} n_{j+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i \sum_{i=1}^j n_i \sum_{i=1}^{j+1} n_i}} - \sqrt{\frac{n_{l+1} \sum_{i=1}^l n_i n_{j+1}}{\sum_{i=1}^{l+1} n_i \sum_{i=1}^j n_i \sum_{i=1}^{j+1} n_i}} \\
&= \sum_{t=1}^l n_t \sqrt{\frac{n_{l+1} n_{j+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i \sum_{i=1}^j n_i \sum_{i=1}^{j+1} n_i}} - \sum_{t=1}^l n_t \sqrt{\frac{n_{l+1} n_{j+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i \sum_{i=1}^j n_i \sum_{i=1}^{j+1} n_i}} \\
&= 0.
\end{aligned}$$

Finally,

$$\begin{aligned}
\frac{1}{\sqrt{n}} \Theta_l^\top \mathbf{Z}^\top \mathbf{Z} \mathbf{1} &= \left(\left\{ \sqrt{\frac{n_{l+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}} \right\}_l, -\sqrt{\frac{\sum_{i=1}^l n_i}{n_{l+1} \sum_{i=1}^{l+1} n_i}}, \{0\}_{K-1-l} \right) \begin{pmatrix} n_1 & 0 & \dots & 0 \\ 0 & n_2 & \dots & 0 \\ 0 & 0 & \dots & n_K \end{pmatrix} \\
&\quad \times (\{1\}_l, 1, \{1\}_{K-1-l})^\top \\
&= \left(\sqrt{\frac{n_1^2 n_{l+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}}, \dots, \sqrt{\frac{n_l^2 n_{l+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}}, -\sqrt{\frac{n_{l+1} \sum_{i=1}^l n_i}{\sum_{i=1}^{l+1} n_i}} \right) \times (\{1\}_l, 1)^\top \\
&= \sum_{i=1}^l \sqrt{\frac{n_i^2 n_{l+1}}{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}} - \sqrt{\frac{n_{l+1} \sum_{i=1}^l n_i}{\sum_{i=1}^{l+1} n_i}} \\
&= \sqrt{\frac{\sum_{i=1}^l n_i n_{l+1}}{\sum_{i=1}^{l+1} n_i}} - \sqrt{\frac{n_{l+1} \sum_{i=1}^l n_i}{\sum_{i=1}^{l+1} n_i}} = 0.
\end{aligned}$$

□

Proof of Theorem 1. The first part of the proof follows the proof of the “slow-rate” bound for lasso, see for example [20]. We reproduce the proof for completeness. Consider the KKT conditions for (1):

$$-\frac{1}{n} \mathbf{X}^\top (\mathbf{Y} - \mathbf{X} \hat{\mathbf{B}}) + \lambda \hat{\boldsymbol{\Psi}} = 0,$$

where $\hat{\boldsymbol{\Psi}}$ is the subgradient of $\|\mathbf{B}\|_{1,2}$ evaluated at $\hat{\mathbf{B}}$. It follows that

$$\text{Tr}[(\hat{\mathbf{B}} - \mathbf{B}^*)^\top \left\{ -\frac{1}{n} \mathbf{X}^\top (\mathbf{Y} - \mathbf{X} \hat{\mathbf{B}}) + \lambda \hat{\boldsymbol{\Psi}} \right\}] = 0,$$

and using $\mathbf{Y} = \mathbf{X}\mathbf{B}^* + \mathbf{Y} - \mathbf{X}\mathbf{B}^* = \mathbf{X}\mathbf{B}^* + \mathbf{E}$

$$\frac{1}{n}\|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F^2 - \langle \frac{1}{n}\mathbf{X}^\top \mathbf{E}, \widehat{\mathbf{B}} - \widehat{\mathbf{B}}^* \rangle + \lambda \langle \widehat{\Psi}, \widehat{\mathbf{B}} - \mathbf{B}^* \rangle = 0.$$

Since $\widehat{\Psi}$ is the subgradient of the convex function $\|\mathbf{B}\|_{1,2}$ evaluated at $\widehat{\mathbf{B}}$, it follows that

$$\|\widehat{\mathbf{B}}\|_{1,2} \leq \|\mathbf{B}^*\|_{1,2} + \langle \widehat{\Psi}, \widehat{\mathbf{B}} - \mathbf{B}^* \rangle.$$

Combining the above two displays leads to

$$\frac{1}{n}\|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F^2 - \langle \frac{1}{n}\mathbf{X}^\top \mathbf{E}, \widehat{\mathbf{B}} - \widehat{\mathbf{B}}^* \rangle + \lambda(\|\widehat{\mathbf{B}}\|_{1,2} - \|\mathbf{B}^*\|_{1,2}) \leq 0.$$

Rearranging the terms gives

$$\frac{1}{n}\|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F^2 \leq \langle \frac{1}{n}\mathbf{X}^\top \mathbf{E}, \widehat{\mathbf{B}} - \widehat{\mathbf{B}}^* \rangle + \lambda\|\mathbf{B}^*\|_{1,2} - \lambda\|\widehat{\mathbf{B}}\|_{1,2}.$$

Using Hölder's inequality and triangle inequality gives

$$\begin{aligned} \frac{1}{n}\|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F^2 &\leq \frac{1}{n}\|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}\|\widehat{\mathbf{B}} - \widehat{\mathbf{B}}^*\|_{1,2} + \lambda\|\mathbf{B}^*\|_1 - \lambda\|\widehat{\mathbf{B}}\|_{1,2} \\ &\leq (\frac{1}{n}\|\mathbf{X}^\top \mathbf{E}\|_{\infty,2} + \lambda)\|\mathbf{B}^*\|_{1,2} + (\frac{1}{n}\|\mathbf{X}^\top \mathbf{E}\|_{\infty,2} - \lambda)\|\widehat{\mathbf{B}}\|_{1,2}. \end{aligned}$$

Using $\lambda \geq \frac{1}{n}\|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$ completes the proof of the first part.

For the second part of the proof, using $\lambda \geq \frac{2}{n}\|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$ and the above display gives

$$\frac{1}{n}\|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F^2 \leq (\lambda/2 + \lambda)\|\mathbf{B}^*\|_{1,2} + (\lambda/2 - \lambda)\|\widehat{\mathbf{B}}\|_{1,2}.$$

Since the left-hand side is non-negative, rearranging the terms gives

$$\frac{\lambda}{2}\|\widehat{\mathbf{B}}\|_{1,2} \leq 3\frac{\lambda}{2}\|\mathbf{B}^*\|_{1,2}.$$

Since $\lambda > 0$, the result follows. \square

Proof of Corollary 1. By triangle and Hölder's inequalities

$$\text{Tr}\{(\widehat{\mathbf{B}} - \mathbf{B}^*)^\top \Sigma_T(\widehat{\mathbf{B}} - \mathbf{B}^*)\} \leq \frac{1}{n}\|\mathbf{X}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_F^2 + \|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{1,2}^2 \|\frac{1}{n}\mathbf{X}^\top \mathbf{X} - \Sigma_T\|_\infty.$$

Applying Theorem 1 leads to stated bound. \square

Proof of Theorem 2. From the proof of Theorem 1

$$\frac{1}{n}\|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F^2 \leq \langle \frac{1}{n}\mathbf{X}^\top \mathbf{E}, \widehat{\mathbf{B}} - \widehat{\mathbf{B}}^* \rangle + \lambda\|\mathbf{B}^*\|_{1,2} - \lambda\|\widehat{\mathbf{B}}\|_{1,2}.$$

Using $\lambda \geq \frac{2}{n}\|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$, Hölder's inequality and Assumption 1,

$$\begin{aligned} \frac{1}{n}\|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F^2 &\leq \frac{\lambda}{2}\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{1,2} + \lambda\|\mathbf{B}^*\|_{1,2} - \lambda\|\widehat{\mathbf{B}}\|_{1,2} \\ &\leq \frac{\lambda}{2}\|\widehat{\mathbf{B}}_S - \mathbf{B}_S^*\|_{1,2} + \frac{\lambda}{2}\|\widehat{\mathbf{B}}_{S^c}\|_{1,2} + \lambda\|\mathbf{B}^*\|_{1,2} - \lambda\|\widehat{\mathbf{B}}_S\|_{1,2} - \lambda\|\widehat{\mathbf{B}}_{S^c}\|_{1,2} \\ &\leq \frac{\lambda}{2}\|\widehat{\mathbf{B}}_S - \mathbf{B}_S^*\|_{1,2} - \frac{\lambda}{2}\|\widehat{\mathbf{B}}_{S^c}\|_{1,2} + \lambda\|\mathbf{B}^*\|_{1,2} - \lambda\|\mathbf{B}^*\|_{1,2} + \lambda\|\widehat{\mathbf{B}}_S^* - \widehat{\mathbf{B}}_S\|_{1,2} \\ &\leq \frac{3\lambda}{2}\|\widehat{\mathbf{B}}_S - \mathbf{B}_S^*\|_{1,2} - \frac{\lambda}{2}\|\widehat{\mathbf{B}}_{S^c}\|_{1,2}, \end{aligned}$$

where in the third step we use triangle inequality $\|\widehat{\mathbf{B}}_S\|_{1,2} \geq \|\mathbf{B}^*\|_{1,2} - \|\widehat{\mathbf{B}}_S - \mathbf{B}_S^*\|_{1,2}$. Since $\|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F^2 \geq 0$, it follows that $\lambda\|\widehat{\mathbf{B}}_{S^c}\|_{1,2} \leq 3\lambda\|\widehat{\mathbf{B}}_S - \mathbf{B}_S^*\|_{1,2}$. Since $\lambda > 0$ and $\mathbf{B}_{S^c}^* = 0$ by Assumption 1, it follows that $\mathbf{A} := \widehat{\mathbf{B}} - \mathbf{B}^*$ belongs to the cone $C(S, 3, K)$ from Definition 2.

Since $\|\widehat{\mathbf{B}}_S - \mathbf{B}_S^*\|_{1,2} \leq \sqrt{s}\|\widehat{\mathbf{B}}_S - \mathbf{B}_S^*\|_F$, and $\frac{1}{n}\mathbf{X}^\top \mathbf{X}$ satisfies RE($s, 3, K$), from the above display

$$\frac{1}{n}\|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F^2 \leq \frac{3\lambda}{2}\sqrt{s}\|\widehat{\mathbf{B}}_S - \mathbf{B}_S^*\|_F \leq \frac{3\lambda}{2}\sqrt{s}\sqrt{\gamma_{\mathbf{X}}}\frac{1}{\sqrt{n}}\|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F.$$

If $\|\mathbf{X}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_F = 0$, the statement of the Theorem holds trivially. Otherwise dividing both sides by $\frac{1}{\sqrt{n}}\|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F$ gives

$$\frac{1}{\sqrt{n}}\|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F \leq \frac{3}{2}\sqrt{\gamma_{\mathbf{X}}}\sqrt{s}\lambda,$$

which leads to

$$\frac{1}{n}\|\mathbf{X}(\mathbf{B}^* - \widehat{\mathbf{B}})\|_F^2 \leq \frac{9}{4}\gamma_{\mathbf{X}}s\lambda^2.$$

Since $\frac{1}{n}\mathbf{X}^\top \mathbf{X}$ satisfies RE($s, 3, K$) and $\mathbf{A} = \widehat{\mathbf{B}} - \mathbf{B}^*$ belongs to the cone $C(S, 3, K)$,

$$\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{1,2} = \|\mathbf{A}\|_{1,2} = \|\mathbf{A}_S\|_{1,2} + \|\mathbf{A}_{S^c}\|_{1,2} \leq 4\|\mathbf{A}_S\|_{1,2} \leq 4\sqrt{s}\|\mathbf{A}_S\|_F \leq 4\sqrt{s}\sqrt{\gamma_{\mathbf{X}}}\sqrt{\|\mathbf{X}\mathbf{A}\|_F^2/n} \leq 6s\lambda\gamma_{\mathbf{X}}.$$

Finally, to prove the bound on $\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F$, we follow derivations in Appendix A.2 of [43]. Let T_0 correspond to the location of s largest in euclidean norm rows of \mathbf{A} , T_1 to the location of s largest in euclidean norm rows of $\mathbf{A}_{T_0^c}$, and so on for T_2, T_3, \dots . Then $\text{card}(T_j) = s$, and

$$\|\mathbf{A}_{T_0^c}\|_{1,2} = \|\mathbf{A}\|_{1,2} - \|\mathbf{A}_{T_0}\|_{1,2} \leq \|\mathbf{A}\|_{1,2} - \|\mathbf{A}_S\|_{1,2} = \|\mathbf{A}_{S^c}\|_{1,2} \leq 3\|\mathbf{A}_S\|_{1,2} \leq 3\|\mathbf{A}_{T_0}\|_{1,2} \leq 3\sqrt{s}\|\mathbf{A}_{T_0}\|_F.$$

Therefore

$$\begin{aligned} \|\mathbf{A}\|_F &\leq \|\mathbf{A}_{T_0}\|_F + \sum_{j \geq 1} \|\mathbf{A}_{T_j}\|_F \leq \|\mathbf{A}_{T_0}\|_F + \sum_{j \geq 1} \sqrt{s}\|\mathbf{A}_{T_j}\|_{\infty,2} \\ &\leq \|\mathbf{A}_{T_0}\|_F + \sum_{j \geq 0} \sqrt{s}\frac{1}{s}\|\mathbf{A}_{T_j}\|_{1,2} \leq \|\mathbf{A}_{T_0}\|_F + \frac{1}{\sqrt{s}}\|\mathbf{A}\|_{1,2}. \end{aligned}$$

Using that $\frac{1}{n}\mathbf{X}^\top \mathbf{X}$ satisfies RE($s, 3, K$) and derived bounds leads to

$$\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F = \|\mathbf{A}\|_F \leq \sqrt{\gamma_{\mathbf{X}}}\sqrt{\|\mathbf{X}\mathbf{A}\|_F^2/n} + \frac{1}{\sqrt{s}}6s\lambda\gamma_{\mathbf{X}} \leq \frac{3}{2}\gamma_{\mathbf{X}}\sqrt{s}\lambda + 6\gamma_{\mathbf{X}}\sqrt{s}\lambda = \frac{15}{2}\gamma_{\mathbf{X}}\sqrt{s}\lambda.$$

□

Proof of Corollary 2. Using triangle and Hölder's inequalities

$$\begin{aligned} \text{Tr}\{(\widehat{\mathbf{B}} - \mathbf{B}^*)^\top \Sigma_T(\widehat{\mathbf{B}} - \mathbf{B}^*)\} \\ = \text{Tr}\{(\widehat{\mathbf{B}} - \mathbf{B}^*)^\top \frac{1}{n}\mathbf{X}^\top \mathbf{X}(\widehat{\mathbf{B}} - \mathbf{B}^*)\} + \text{Tr}\{(\widehat{\mathbf{B}} - \mathbf{B}^*)^\top (\frac{1}{n}\mathbf{X}^\top \mathbf{X} - \Sigma_T)(\widehat{\mathbf{B}} - \mathbf{B}^*)\} \\ \leq \frac{1}{n}\|\mathbf{X}(\widehat{\mathbf{B}} - \mathbf{B}^*)\|_F^2 + \|\widehat{\mathbf{B}} - \mathbf{B}^*\|_{1,2}^2 \|\frac{1}{n}\mathbf{X}^\top \mathbf{X} - \Sigma_T\|_\infty. \end{aligned}$$

Applying Theorem 2 for $\lambda \geq \frac{2}{n}\|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$ gives

$$\text{Tr}\{(\widehat{\mathbf{B}} - \mathbf{B}^*)^\top \Sigma_T(\widehat{\mathbf{B}} - \mathbf{B}^*)\} \leq \frac{9}{4}\gamma_{\mathbf{X}}s\lambda^2 + 36\gamma_{\mathbf{X}}^2s^2\lambda^2 \|\frac{1}{n}\mathbf{X}^\top \mathbf{X} - \Sigma_T\|_\infty.$$

On the other hand, using Theorem 2 gives

$$\text{Tr}\{(\widehat{\mathbf{B}} - \mathbf{B}^*)^\top \Sigma_T(\widehat{\mathbf{B}} - \mathbf{B}^*)\} \leq \lambda_{\max}(\Sigma_T)\|\widehat{\mathbf{B}} - \mathbf{B}^*\|_F^2 \leq \lambda_{\max}(\Sigma_T)57\gamma_{\mathbf{X}}^2s\lambda^2.$$

□

Proof of Theorem 3. Consider

$$\frac{1}{n} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2} = \left\| \frac{1}{n} \mathbf{X}^\top \mathbf{Y} - \frac{1}{n} \mathbf{X}^\top \mathbf{X} \mathbf{B}^* \right\|_{\infty,2} \leq \underbrace{\left\| \frac{1}{n} \mathbf{X}^\top \mathbf{Y} - \mathbf{\Delta} \right\|_{\infty,2}}_{:=I_1} + \underbrace{\left\| \mathbf{\Delta} - \frac{1}{n} \mathbf{X}^\top \mathbf{X} \mathbf{\Sigma}_T^{-1} \mathbf{\Delta} \right\|_{\infty,2}}_{:=I_2}.$$

Consider I_1 . From Lemma 4, with probability at least $1 - \eta$ for some constant $C > 0$

$$I_1 = \left\| \frac{1}{n} \mathbf{X}^\top \mathbf{Y} - \mathbf{\Delta} \right\|_{\infty,2} \leq C \max_j \sigma_j \sqrt{\frac{(K-1) \log(p\eta^{-1})}{n}}.$$

Consider I_2 . Using $\|\mathbf{ABC}\|_{\infty,2} \leq \|\mathbf{A}\|_\infty \|\mathbf{B}\|_\infty \|\mathbf{C}\|_{\infty,2}$ [29, Lemma 8] gives

$$\begin{aligned} I_2 &= \left\| \mathbf{\Delta} - \frac{1}{n} \mathbf{X}^\top \mathbf{X} \mathbf{\Sigma}_T^{-1} \mathbf{\Delta} \right\|_{\infty,2} = \left\| \mathbf{\Sigma}_T^{1/2} \left(\mathbf{I} - \frac{1}{n} \mathbf{\Sigma}_T^{-1/2} \mathbf{X}^\top \mathbf{X} \mathbf{\Sigma}_T^{-1/2} \right) \mathbf{\Sigma}_T^{-1/2} \mathbf{\Delta} \right\|_{\infty,2} \\ &\leq \left\| \mathbf{\Sigma}_T^{1/2} \right\|_\infty \left\| \mathbf{I} - \frac{1}{n} \mathbf{\Sigma}_T^{-1/2} \mathbf{X}^\top \mathbf{X} \mathbf{\Sigma}_T^{-1/2} \right\|_\infty \left\| \mathbf{\Sigma}_T^{-1/2} \mathbf{\Delta} \right\|_{\infty,2}. \end{aligned} \quad (\text{A.1})$$

Consider $\|\mathbf{\Sigma}_T^{-1/2} \mathbf{\Delta}\|_{\infty,2}$. Since $\mathbf{\Sigma}_T = \mathbf{\Sigma}_W + \mathbf{\Delta} \mathbf{\Delta}^\top$ [10, Proposition 2], by Woodbury matrix identity $\mathbf{\Delta}^\top \mathbf{\Sigma}_T^{-1} \mathbf{\Delta} = \mathbf{\Delta}^\top \mathbf{\Sigma}_W^{-1} \mathbf{\Delta} (\mathbf{I} + \mathbf{\Delta}^\top \mathbf{\Sigma}_W^{-1} \mathbf{\Delta})^{-1}$. Therefore,

$$\begin{aligned} \|\mathbf{\Sigma}_T^{-1/2} \mathbf{\Delta}\|_{\infty,2} &= \max_j \|\mathbf{e}_j^\top \mathbf{\Sigma}_T^{-1/2} \mathbf{\Delta}\|_2 = \max_j \sqrt{\mathbf{e}_j^\top \mathbf{\Sigma}_T^{-1/2} \mathbf{\Delta} \mathbf{\Delta}^\top \mathbf{\Sigma}_T^{-1/2} \mathbf{e}_j} \\ &\leq \|\mathbf{\Sigma}_T^{-1/2} \mathbf{\Delta} \mathbf{\Delta}^\top \mathbf{\Sigma}_T^{-1/2}\|_2 = \|\mathbf{\Delta}^\top \mathbf{\Sigma}_T^{-1} \mathbf{\Delta}\|_2 = \|\mathbf{\Delta}^\top \mathbf{\Sigma}_W^{-1} \mathbf{\Delta} (\mathbf{I} + \mathbf{\Delta}^\top \mathbf{\Sigma}_W^{-1} \mathbf{\Delta})^{-1}\|_2 \leq 1. \end{aligned}$$

Consider $\|\mathbf{\Sigma}_T^{1/2}\|_\infty$. Let $\mathbf{\Sigma}_T = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^\top$ be the eigendecomposition of $\mathbf{\Sigma}_T$, then $\mathbf{\Sigma}_T^{1/2} = \mathbf{U} \mathbf{\Lambda}^{1/2} \mathbf{U}^\top$ is positive definite and $\|\mathbf{\Sigma}_T^{1/2}\|_\infty = \max_j |\sum_{i=1}^p \sqrt{\lambda_i} u_{ji}^2|$. Since $f(x) = \sqrt{x}$ is concave and $\sum_{i=1}^p u_{ji}^2 = 1$ for all j , it follows that

$$\|\mathbf{\Sigma}_T^{1/2}\|_\infty = \max_j \left| \sum_{i=1}^p \sqrt{\lambda_i} u_{ji}^2 \right| \leq \max_j \sqrt{\sum_{i=1}^p \lambda_i u_{ji}^2} \leq \sqrt{\|\mathbf{\Sigma}_T\|_\infty} \leq \max_j \sqrt{\sigma_{jj}^2 + \max_k \mu_{kj}^2} = \tau,$$

where the last inequality holds since $\mathbf{\Sigma}_T = \mathbf{\Sigma}_W + \sum_{k=1}^K \pi_k \boldsymbol{\mu}_k \boldsymbol{\mu}_k^\top$ for $\boldsymbol{\mu} = 0$.

Finally, from Lemma 3, all elements of $\mathbf{X} \mathbf{\Sigma}_T^{-1/2}$ are sub-gaussian with parameter C that does not depend on $\mathbf{\Sigma}_W$ or $\boldsymbol{\mu}_k$. Therefore, from Lemma 7 with probability at least $1 - \eta$

$$\left\| \mathbf{I} - \frac{1}{n} \mathbf{\Sigma}_T^{-1/2} \mathbf{X}^\top \mathbf{X} \mathbf{\Sigma}_T^{-1/2} \right\|_\infty \leq C_1 \sqrt{\frac{\log(p\eta^{-1})}{n}}.$$

Combining the above displays with (A.1) gives

$$I_2 \leq C_2 \tau \sqrt{\frac{\log(p\eta^{-1})}{n}}.$$

Combining results for I_1 and I_2 gives

$$\left\| \frac{1}{n} \mathbf{X}^\top (\mathbf{Y} - \mathbf{X} \mathbf{B}^*) \right\|_{\infty,2} \leq C_3 \tau \sqrt{\frac{(K-1) \log(p\eta^{-1})}{n}}$$

with probability at least $1 - \eta$ for some constant $C_3 > 0$. \square

Proof of Theorem 4. From Corollary 1, if $\lambda \geq \frac{2}{n} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$,

$$\text{Tr}\{(\widehat{\mathbf{B}} - \mathbf{B}^*)^\top \mathbf{\Sigma}_T (\widehat{\mathbf{B}} - \mathbf{B}^*)\} \leq \frac{3}{2} \lambda \|\mathbf{B}^*\|_{1,2} + 16 \|\mathbf{B}^*\|_{1,2}^2 \left\| \frac{1}{n} \mathbf{X}^\top \mathbf{X} - \mathbf{\Sigma}_T \right\|_\infty.$$

Applying Theorem 3 for $\frac{1}{n} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$ and Lemmas 3 and 7 for $\|\frac{1}{n} \mathbf{X}^\top \mathbf{X} - \mathbf{\Sigma}_T\|_\infty$ leads to the desired statement. \square

Proof of Theorem 5. By Lemma 5, for $n \geq Cs\log(p/s)$, $n^{-1/2}\mathbf{X}$ satisfies RE($s, 3$) with probability at least $1 - O(e^{-n})$ with

$$\gamma_{\mathbf{X}} = \gamma(s, 3, n^{-1/2}\mathbf{X}) \leq 2\gamma(s, 3, \Sigma_T^{1/2}) = 2\gamma.$$

The first bound follows by combining this with Corollary 2 and Theorem 3. The second bound follows by combining this with the results of Theorem 2 and Theorem 3. \square

Proof of Theorem 6. By Lemmas 3 and 7, with probability at least $1 - \eta$

$$\|n^{-1}\mathbf{X}^\top \mathbf{X} - \Sigma_T\|_\infty \leq C\tau^2 \sqrt{\frac{\log(p\eta^{-1})}{n}}.$$

By Lemma 6, if $s \leq (32\gamma\|\Sigma_T - n^{-1}\mathbf{X}^\top \mathbf{X}\|_\infty)^{-1}$, then $\gamma_{\mathbf{X}} \leq 2\gamma$. Therefore, using $s = o(\sqrt{n/\log p})$, Corollary 2 and Theorem 3 gives that for $\lambda \geq C\tau\sqrt{\frac{(K-1)\log p}{n}}$

$$\text{Tr}\{(\widehat{\mathbf{B}} - \mathbf{B}^*)^\top \Sigma_T(\widehat{\mathbf{B}} - \mathbf{B}^*)\} = O_p\left\{\lambda_{\max}(\Sigma_T)\tau^2\gamma^2\frac{(K-1)s\log p}{n}\right\}.$$

The second bound follows by combining the results of Theorem 2 and Theorem 3. \square

Proof of Theorem 7. Consider

$$\widehat{e}_h - e_h = \sum_{k=1}^K \pi_k(\widehat{e}_{hk} - e_{hk}) \leq \max_k(\widehat{e}_{hk} - e_{hk}).$$

We further derive bound on $\max_k(\widehat{e}_{hk} - e_{hk})$. Let

$$\begin{aligned} h_k(\mathbf{x}) &:= (\mathbf{x} - \boldsymbol{\mu}_k)^\top \mathbf{B}^* (\mathbf{B}^{*\top} \Sigma_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} (\mathbf{x} - \boldsymbol{\mu}_k) - 2\log \pi_k, \\ \widehat{h}_k(\mathbf{x}) &:= (\mathbf{x} - \bar{\mathbf{x}}_k)^\top \widehat{\mathbf{B}} (\widehat{\mathbf{B}}^\top \widehat{\Sigma}_W \widehat{\mathbf{B}})^{-1} \widehat{\mathbf{B}}^\top (\mathbf{x} - \bar{\mathbf{x}}_k) - 2\log n_k/n. \end{aligned}$$

Then

$$\begin{aligned} e_{hk} &= \text{pr}(h(\mathbf{x}) \neq k | \mathbf{x} \in \mathcal{C}_k) = 1 - \text{pr}(h(\mathbf{x}) = k | \mathbf{x} \in \mathcal{C}_k) = 1 - \text{pr}(h_k(\mathbf{x}) < h_l(\mathbf{x}) \text{ for all } l \neq k | \mathbf{x} \in \mathcal{C}_k) \\ &= 1 - \text{pr}(h_k(\mathbf{x}) - h_l(\mathbf{x}) < 0 \text{ for all } l \neq k | \mathbf{x} \in \mathcal{C}_k). \end{aligned}$$

Similarly, $\widehat{e}_{hk} = 1 - \text{pr}(\widehat{h}_k(\mathbf{x}) - \widehat{h}_l(\mathbf{x}) < 0 \text{ for all } l \neq k | \mathbf{x} \in \mathcal{C}_k)$. Consider $a_{kl}(\mathbf{x}) = h_k(\mathbf{x}) - h_l(\mathbf{x})$, and $\widehat{a}_{kl}(\mathbf{x}) = \widehat{h}_k(\mathbf{x}) - \widehat{h}_l(\mathbf{x})$. Then for any $\varepsilon > 0$

$$\begin{aligned} \widehat{e}_{hk} &= \text{pr}\left(\bigcup_l \{\widehat{a}_{kl}(\mathbf{x}) > 0\} | \mathbf{x} \in \mathcal{C}_1\right) = \text{pr}\left(\bigcup_l \left\{\frac{a_{kl}(\mathbf{x}) + (\widehat{a}_{kl}(\mathbf{x}) - a_{kl}(\mathbf{x}))}{\text{sd}\{a_{kl}(\mathbf{x})\}} > 0\right\} | \mathbf{x} \in \mathcal{C}_1\right) \\ &\leq \text{pr}\left(\max_l \left\{\frac{a_{kl}(\mathbf{x})}{\text{sd}\{a_{kl}(\mathbf{x})\}} + \frac{|\widehat{a}_{kl}(\mathbf{x}) - a_{kl}(\mathbf{x})|}{\text{sd}\{a_{kl}(\mathbf{x})\}}\right\} > 0 | \mathbf{x} \in \mathcal{C}_1\right) \\ &\leq \text{pr}\left(\max_l \frac{a_{kl}(\mathbf{x})}{\text{sd}\{a_{kl}(\mathbf{x})\}} > 0 | \mathbf{x} \in \mathcal{C}_1\right) + \text{pr}\left(-\varepsilon < \max_l \frac{a_{kl}(\mathbf{x})}{\text{sd}\{a_{kl}(\mathbf{x})\}} \leq 0 | \mathbf{x} \in \mathcal{C}_1\right) \\ &\quad + \text{pr}\left(\max_l \frac{|\widehat{a}_{kl}(\mathbf{x}) - a_{kl}(\mathbf{x})|}{\text{sd}\{a_{kl}(\mathbf{x})\}} \geq \varepsilon | \mathbf{x} \in \mathcal{C}_1\right) \\ &\leq e_{hk} + \text{pr}\left(\max_l \frac{|a_{kl}(\mathbf{x})|}{\text{sd}\{a_{kl}(\mathbf{x})\}} < \varepsilon | \mathbf{x} \in \mathcal{C}_1\right) + \text{pr}\left(\max_l \frac{|\widehat{a}_{kl}(\mathbf{x}) - a_{kl}(\mathbf{x})|}{\text{sd}\{a_{kl}(\mathbf{x})\}} \geq \varepsilon | \mathbf{x} \in \mathcal{C}_1\right) \end{aligned} \tag{A.2}$$

Therefore,

$$\begin{aligned} \max_k(\widehat{e}_{hk} - e_{hk}) &\leq \max_k \text{pr}\left(\max_l \frac{|a_{kl}(\mathbf{x})|}{\text{sd}\{a_{kl}(\mathbf{x})\}} < \varepsilon | \mathbf{x} \in \mathcal{C}_k\right) + \max_k \text{pr}\left(\max_l \frac{|\widehat{a}_{kl}(\mathbf{x}) - a_{kl}(\mathbf{x})|}{\text{sd}\{a_{kl}(\mathbf{x})\}} \geq \varepsilon | \mathbf{x} \in \mathcal{C}_k\right) \\ &:= I_1 + I_2. \end{aligned} \tag{A.3}$$

We further provide bounds on I_1 and I_2 .

Let $\delta_{kl} = \mu_k - \mu_l$, then

$$\begin{aligned} a_{kl}(\mathbf{x}) &= -2\delta_{kl}^\top \mathbf{B}^* (\mathbf{B}^{*\top} \Sigma_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} \mathbf{x} + \mu_k^\top \mathbf{B}^* (\mathbf{B}^{*\top} \Sigma_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} \mu_k \\ &\quad - \mu_k^\top \mathbf{B}^* (\mathbf{B}^{*\top} \Sigma_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} \mu_k - 2 \log(\pi_k / \pi_l) \\ &= -2\delta_{kl}^\top \mathbf{B}^* (\mathbf{B}^{*\top} \Sigma_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} \mathbf{x} - \delta_{kl}^\top \mathbf{B}^* (\mathbf{B}^{*\top} \Sigma_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} \delta_{kl} \\ &\quad + 2\delta_{kl}^\top \mathbf{B}^* (\mathbf{B}^{*\top} \Sigma_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} \mu_k - 2 \log(\pi_k / \pi_l) \\ &:= u_{kl}(\mathbf{x}) - 2 \log(\pi_k / \pi_l). \end{aligned}$$

By Assumption 3, for fixed k , $a_{kl}(\mathbf{x}) / \text{sd}\{a_{kl}(\mathbf{x})\} | \mathbf{x} \in \mathcal{C}_k, l \neq k$, are jointly normally distributed with mean μ_{ak} and the covariance matrix Σ_{ak} . Then

$$\begin{aligned} \text{pr} \left(\max_l \frac{|a_{kl}(\mathbf{x})|}{\text{sd}\{a_{kl}(\mathbf{x})\}} < \varepsilon | \mathbf{x} \in \mathcal{C}_1 \right) &= \text{pr} \left(\bigcap_{l \neq k} \left\{ \frac{|a_{kl}(\mathbf{x})|}{\text{sd}\{a_{kl}(\mathbf{x})\}} < \varepsilon \right\} | \mathbf{x} \in \mathcal{C}_1 \right) \\ &= \int_{-\varepsilon}^{\varepsilon} \dots \int_{-\varepsilon}^{\varepsilon} \frac{1}{\sqrt{(2\pi)^{K-1} |\Sigma_{ak}|}} \exp[-(\mathbf{z} - \mu_{ak})^\top \Sigma_{1a}^{-1} (\mathbf{z} - \mu_{ak})] dz_1 \dots dz_{K-1} \\ &\leq \frac{1}{\sqrt{(2\pi)^{K-1} |\Sigma_{ak}|}} (2\varepsilon)^{K-1}. \end{aligned}$$

By Assumption 5, $|\Sigma_{ak}|$ is constant for each k , therefore for small $\varepsilon > 0$ there exists $C > 0$ such that

$$I_1 \leq C\varepsilon^{K-1}.$$

Consider next I_2 . Let $\mathbf{d}_{kl} = \bar{\mathbf{x}}_k - \bar{\mathbf{x}}_l$, then

$$\begin{aligned} \hat{a}_{kl}(\mathbf{x}) &= -2\mathbf{d}_{kl}^\top \hat{\mathbf{B}} (\hat{\mathbf{B}}^\top \hat{\Sigma}_W \hat{\mathbf{B}})^{-1} \hat{\mathbf{B}}^\top \mathbf{x} - \mathbf{d}_{kl}^\top \hat{\mathbf{B}} (\hat{\mathbf{B}}^\top \hat{\Sigma}_W \hat{\mathbf{B}})^{-1} \hat{\mathbf{B}}^\top \mathbf{d}_{kl} \\ &\quad + 2\mathbf{d}_{kl}^\top \hat{\mathbf{B}} (\hat{\mathbf{B}}^\top \hat{\Sigma}_W \hat{\mathbf{B}})^{-1} \hat{\mathbf{B}}^\top \bar{\mathbf{x}}_k - 2 \log(n_k / n_l) \\ &= \hat{u}_{kl}(\mathbf{x}) - 2 \log(n_k / n_l). \end{aligned}$$

Therefore, $a_{kl}(\mathbf{x}) - \hat{a}_{kl}(\mathbf{x}) = u_{kl}(\mathbf{x}) - \hat{u}_{kl}(\mathbf{x}) - 2 \log(\pi_k / \pi_l) + 2 \log(n_k / n_l)$. By Lemma 11, $|\log(\pi_k / \pi_l) - \log(n_k / n_l)| = O_p(\log(K-1)n^{-1/2})$, therefore we further focus on $u_{kl}(\mathbf{x}) - \hat{u}_{kl}(\mathbf{x})$. Conditionally on $\mathbf{x} \in \mathcal{C}_k$,

$$\frac{u_{kl}(\mathbf{x}) - \hat{u}_{kl}(\mathbf{x})}{\text{sd}\{a_{kl}(\mathbf{x})\}} = \frac{m_{kl}}{\text{sd}\{a_{kl}(\mathbf{x})\}} + \frac{\sqrt{v_{kl}}}{\text{sd}\{a_{kl}(\mathbf{x})\}} \mathbf{z},$$

where $\mathbf{z} | \mathbf{x} \in \mathcal{C}_k \sim \mathcal{N}(0, 1)$. Applying Lemma 8 and Slutsky, it follows that under conditions of Theorem 4

$$\max_{k,l} \frac{|u_{kl}(\mathbf{x}) - \hat{u}_{kl}(\mathbf{x})|}{\text{sd}\{a_{kl}(\mathbf{x})\}} = O_p \left\{ \tau^{1/2} \|\mathbf{B}^*\|_{1,2}^{1/2} (1 \vee \tau^{1/2} \|\mathbf{B}^*\|_{1,2}^{1/2}) \left(\frac{(K-1) \log p}{n} \right)^{1/4} \right\}.$$

Setting $\varepsilon = M_\eta \tau^{1/2} \|\mathbf{B}^*\|_{1,2}^{1/2} (1 \vee \tau^{1/2} \|\mathbf{B}^*\|_{1,2}^{1/2}) \left(\frac{(K-1) \log p}{n} \right)^{1/4}$ in (A.2) and invoking the bound for I_1 completes the proof. \square

Proof of Theorem 8. The proof follows the proof of Theorem 7 by substituting the appropriate bounds from Lemma 8. \square

A.2. Additional lemmas

Lemma 3. *Under Assumptions 2–3, all elements of \mathbf{X} are sub-gaussian, that is*

$$\mathbb{E} e^{\lambda x_{ij}} \leq e^{\lambda^2 \tau^2 / 2} \quad \text{for all } \lambda \in \mathbb{R}; \quad i = 1, \dots, n; \quad j = 1, \dots, p;$$

where $\tau = \max_{j=1,\dots,p} \sqrt{\sigma_j^2 + \max_k \mu_{kj}^2}$ with σ_j^2 being the diagonal elements of Σ_W . Similarly, all elements of $\mathbf{X}\Sigma_T^{-1/2}$ are subgaussian with parameter $C > 0$ that does not depend on Σ_W or μ_k .

Proof. Since $\mathbf{x}_i | \mathbf{x}_i \in \mathcal{C}_k \sim N(\mu_k, \Sigma_W)$,

$$\mathbf{x}_i = \sum_{k=1}^K \mu_k \mathbb{1}\{\mathbf{x}_i \in \mathcal{G}_k\} + \Sigma_W^{1/2} \zeta_i = \mathbf{t}_{1i} + \mathbf{t}_{2i},$$

where $\zeta_i \sim \mathcal{N}(0, \mathbf{I})$ and $\mathbf{t}_{1i}, \mathbf{t}_{2i}$ are independent random vectors.

Since $\mu = 0$, t_{1ij} is mean zero random variable with $|t_{1ij}| \leq \max_k |\mu_{kj}|$, hence t_{1ij} is sub-gaussian with parameter at most $\max_k |\mu_{kj}|$. On the other hand, \mathbf{t}_{2i} is mean zero gaussian random vector with $\text{Cov}(\mathbf{t}_{2i}) = \Sigma_W$. Hence, $\text{Var}(t_{2ij}) = \sigma_j^2$ for all j , and t_{2ij} is sub-gaussian with parameter σ_j . Since t_{1ij} and t_{2ij} are independent,

$$\mathbb{E}(e^{\lambda x_{ij}}) = \mathbb{E}\{e^{\lambda(t_{1ij} + t_{2ij})}\} = \mathbb{E}(e^{\lambda t_{1ij}})\mathbb{E}(e^{\lambda t_{2ij}}) \leq e^{\lambda^2\{\sigma_j^2 + \max_k \mu_{kj}^2\}/2}.$$

Therefore, x_{ij} is sub-gaussian with parameter $\tau_j = \sqrt{\sigma_j^2 + \max_k \mu_{kj}^2}$. Letting $\tau = \max_j \tau_j$, all elements of \mathbf{X} are sub-gaussian with parameter at most τ .

Similarly,

$$\Sigma_T^{-1/2} \mathbf{x}_i = \Sigma_T^{-1/2} \sum_{k=1}^K \mu_k \mathbb{1}\{\mathbf{x}_i \in \mathcal{G}_k\} + \Sigma_T^{-1/2} \Sigma_W^{1/2} \zeta_i = \mathbf{u}_{1i} + \mathbf{u}_{2i}.$$

Let $\mathbf{M} = [\mu_1 \dots \mu_k] \in \mathbb{R}^{p \times k}$, then

$$\|\mathbf{u}_{1i}\|_\infty = \|\Sigma_T^{-1/2} \sum_{k=1}^K \mu_k \mathbb{1}\{\mathbf{x}_i \in \mathcal{G}_k\}\|_\infty \leq \|\Sigma_T^{-1/2} \mathbf{M}\|_{\infty,2} \leq \|\mathbf{M}^\top \Sigma_T^{-1} \mathbf{M}\|_2.$$

Let $\mathbf{\Pi} = \text{diag}(\pi_1, \dots, \pi_K)$, then $\Sigma_T = \Sigma_W + \mathbf{M}\mathbf{\Pi}\mathbf{M}^\top$, and by Woodbury matrix identity

$$\|\mathbf{M}\Sigma_T^{-1}\mathbf{M}\|_2 = \|\mathbf{M}^\top \Sigma_W^{-1} \mathbf{M}\mathbf{\Pi}^{-1}(\mathbf{\Pi}^{-1} + \mathbf{M}^\top \Sigma_W^{-1} \mathbf{M})^{-1}\|_2 \leq C,$$

where the last inequality uses Assumption 2. It follows that all elements of \mathbf{u}_{1i} are bounded by at most C , hence are sub-gaussian with parameter at most C . On the other hand, using $\Sigma_T = \Sigma_W + \Delta\Delta^\top$ [10, Proposition 2],

$$\text{Cov}(\mathbf{u}_{2i}) = \Sigma_T^{-1/2} \Sigma_W \Sigma_T^{-1/2} = \Sigma_T^{-1/2} (\Sigma_T - \Delta\Delta^\top) \Sigma_T^{-1/2} = \mathbf{I} - \Sigma_T^{-1/2} \Delta\Delta^\top \Sigma_T^{-1/2},$$

therefore by Assumption 3 all elements of \mathbf{u}_{2i} are sub-gaussian with parameter at most one. Since $\Sigma_T^{-1/2} \mathbf{x}_i = \mathbf{u}_{1i} + \mathbf{u}_{2i}$, it follows that all elements of $\Sigma_T^{-1/2} \mathbf{x}_i$ are subgaussian with parameter at most C_1 independent of Σ_W and μ_k . \square

Lemma 4. Let $\mathbf{D} = \frac{1}{n} \mathbf{X}^\top \mathbf{Y} = \frac{1}{n} \mathbf{X}^\top \mathbf{Z} \tilde{\Theta}$, where $\tilde{\Theta}$ is from Lemma 2, and let Δ be as in (6). Under Assumptions 2–3, with probability at least $1 - \eta$

$$\|\mathbf{D} - \Delta\|_{\infty,2} \leq C \max_j \sigma_j \sqrt{\frac{(K-1) \log(p\eta^{-1})}{n}}.$$

Proof. Using the definition of $\tilde{\Theta}$ and \mathbf{Z} , it follows that the l th column of \mathbf{D} has the form

$$\mathbf{D}_l = \frac{1}{\sqrt{n}} \frac{\sqrt{n_{l+1}} \sum_{i=1}^l n_i (\bar{\mathbf{x}}_i - \bar{\mathbf{x}}_{l+1})}{\sqrt{\sum_{i=1}^l n_i \sum_{i=1}^{l+1} n_i}}.$$

Using Assumptions 2–3 and Lemma 8 in [11], \mathbf{D} has matrix-normal distribution with $\mathbb{E}(\mathbf{D}) = \mathbf{\Delta} + o(1)$ and $\text{Cov}(\mathbf{D}) = n^{-1}\mathbf{\Sigma}_W + o(1)$. Applying the tail inequality for quadratic form of the gaussian random vector [18, Proposition 1.1] gives for all $t > 0$

$$\text{pr}(\|\mathbf{e}_j^\top \mathbf{D} - \mathbf{e}_j^\top \mathbf{\Delta}\|_2^2 / \sigma_j^2 > (K-1) + 2\sqrt{(K-1)t} + 2t) \leq e^{-t}.$$

Applying union bound over all $j \in \{1, \dots, p\}$ and taking large t gives that with probability at least $1 - \eta$

$$\|\mathbf{D} - \mathbf{\Delta}\|_{\infty,2} \leq C \max_j \sigma_j \sqrt{\frac{(K-1) \log(p\eta^{-1})}{n}}.$$

□

Lemma 5. *Under Assumptions 1–4, if $\mathbf{\Sigma}_T^{1/2}$ satisfies RE($s, 9$), then for $n \geq Cs \log(p/s)$ with probability at least $1 - O(e^{-n}) n^{-1/2} \mathbf{X}$ satisfies RE($s, 3$) with parameter*

$$0 < \gamma(s, 3, n^{-1/2} \mathbf{X}) \leq 2\gamma(s, 3, \mathbf{\Sigma}_T^{1/2}).$$

Proof. Under Assumptions 2–3, all elements of \mathbf{X} are sub-gaussian with marginal covariance matrix $\mathbf{\Sigma}_T$. The result follows using the assumption on $\mathbf{\Sigma}_T^{1/2}$ and applying Theorem 6 in [33] with $\delta = 1/2$. □

Lemma 6. *Let $\mathbf{\Sigma}_T^{1/2}$ satisfy RE($s, 3, K$) with $\gamma = \gamma(s, 3, K, \mathbf{\Sigma}_T^{1/2})$. If $s \leq (32\gamma \|\mathbf{\Sigma}_T - n^{-1} \mathbf{X}^\top \mathbf{X}\|_\infty)^{-1}$, then $n^{-1/2} \mathbf{X}$ satisfies RE($s, 3, K$) with*

$$0 < \gamma(s, 3, K, n^{-1/2} \mathbf{X}) \leq 2\gamma(s, 3, K, \mathbf{\Sigma}_T^{1/2}).$$

Proof of Lemma 6. Since $\mathbf{\Sigma}_T^{1/2}$ satisfies RE($s, 3, K$), for all $\mathbf{A} \in \mathcal{C}(S, 3, K)$

$$\frac{1}{n} \text{Tr}(\mathbf{A}^\top \mathbf{X}^\top \mathbf{X} \mathbf{A}) = \text{Tr}(\mathbf{A}^\top \mathbf{\Sigma}_T \mathbf{A}) + \text{Tr}\{\mathbf{A}^\top (\mathbf{\Sigma}_T - n^{-1} \mathbf{X}^\top \mathbf{X}) \mathbf{A}\} \geq \frac{1}{\gamma} \|\mathbf{A}_S\|_F^2 - \|\mathbf{A}\|_{1,2} \|\mathbf{\Sigma}_T - n^{-1} \mathbf{X}^\top \mathbf{X}\|_\infty.$$

Since $\mathbf{A} \in \mathcal{C}(S, 3, K)$, $\|\mathbf{A}\|_{1,2} = \|\mathbf{A}_S\|_{1,2} + \|\mathbf{A}_{S^c}\|_{1,2} \leq 4\|\mathbf{A}_S\|_{1,2}$. Therefore

$$\begin{aligned} \frac{1}{n} \text{Tr}(\mathbf{A}^\top \mathbf{X}^\top \mathbf{X} \mathbf{A}) &\geq \frac{1}{\gamma} \|\mathbf{A}_S\|_F^2 - 16\|\mathbf{A}_S\|_{1,2}^2 \|\mathbf{\Sigma}_T - n^{-1} \mathbf{X}^\top \mathbf{X}\|_\infty \\ &\geq \frac{1}{\gamma} \|\mathbf{A}_S\|_F^2 - 16s \|\mathbf{A}_S\|_F^2 \|\mathbf{\Sigma}_T - n^{-1} \mathbf{X}^\top \mathbf{X}\|_\infty \\ &\geq \frac{1}{\gamma} \|\mathbf{A}_S\|_F^2 - \frac{1}{2\gamma} \|\mathbf{A}_S\|_F^2 = \frac{1}{2\gamma} \|\mathbf{A}_S\|_F^2, \end{aligned}$$

where we used the condition on s in the last inequality. □

Lemma 7. *Let $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^p$ be independent zero-mean random vectors with $\max_{j=1, \dots, p} \|x_{ij}\|_{\psi_2} \leq \tau$, $\text{Cov}(\mathbf{x}_i) = \mathbf{\Sigma}$, and let $\mathbf{X} = [\mathbf{x}_1 \dots \mathbf{x}_n]^\top$. Under Assumption 4, for some constant $C > 0$ and a fixed $\eta \in (0, 1)$*

$$\|n^{-1} \mathbf{X}^\top \mathbf{X} - \mathbf{\Sigma}\|_\infty \leq C\tau^2 \sqrt{\frac{\log(p\eta^{-1})}{n}}$$

with probability at least $1 - \eta$.

Proof. The statement is equivalent to Lemma F.2 in [28]. For completeness, we provide the full proof. Let $t_{ikj} = x_{ik} x_{ij}$. Since $\max_{j=1, \dots, p} \|x_{ij}\|_{\psi_2} \leq \tau$, applying Cauchy-Schwartz inequality together with [36, Lemma 5.14] leads to

$$\|t_{ikj}\|_{\psi_1} \leq \sqrt{\|x_{ik}^2\|_{\psi_1} \|x_{ij}^2\|_{\psi_1}} \leq \sqrt{2\|x_{ik}\|_{\psi_2}^2 2\|x_{ij}\|_{\psi_2}^2} = 2\|x_{ik}\|_{\psi_2} \|x_{ij}\|_{\psi_2} \leq 2\tau^2.$$

That is, t_{ikj} is sub-exponential with parameter $2\tau^2$. Moreover, $\|t_{ikj} - \sigma_{kj}\|_{\psi_2} = \|t_{ikj} - \mathbb{E}(t_{ikj})\|_{\psi_2} \leq 2\|t_{ikj}\|_{\psi_2} \leq 4\tau^2$ is also sub-exponential with parameter $4\tau^2$. Applying Bernstein's bound [36, Corollary 5.17] together with the union bound leads to

$$\text{pr}(\|n^{-1}\mathbf{X}^\top \mathbf{X} - \boldsymbol{\Sigma}\|_\infty \geq \varepsilon) \leq 2p^2 \exp[-C \min(\varepsilon^2/16\tau^4, \varepsilon/4\tau^2)n].$$

for some constant $C > 0$. Letting $\varepsilon = C_1\tau^2\sqrt{\frac{\log(p\eta^{-1})}{n}}$ for fixed $\eta \in (0, 1)$ and using Assumption 4 gives $\|n^{-1}\mathbf{X}^\top \mathbf{X} - \boldsymbol{\Sigma}\|_\infty \leq \varepsilon$ with probability at least $1 - \eta$. \square

Lemma 8. *Let $\boldsymbol{\delta}_{kl} = \boldsymbol{\mu}_k - \boldsymbol{\mu}_l$, $\mathbf{A} = \mathbf{B}^{*\top} \boldsymbol{\Sigma}_W \mathbf{B}^*$ and $\mathbf{q}_{kl} = \mathbf{A}^{-1/2} \mathbf{B}^{*\top} \boldsymbol{\delta}_{kl}$. Let m_{kl} and v_{kl} be conditional mean and variance of $u_{kl}(\mathbf{x}) - \hat{u}_{kl}(\mathbf{x})$ defined in the proof of Theorem 7, let Assumption 5 hold and let $K = o(p)$.*

1. Under conditions of Theorem 4

$$\begin{aligned} \max_{k,l} \frac{m_{kl}}{\text{sd}\{a_{kl}(\mathbf{x})\}} &= O_p\left\{(\tau\|\mathbf{B}^*\|_{1,2})^{1/2}\left(\frac{(K-1)\log p}{n}\right)^{1/4}\right\}, \\ \max_{k,l} \frac{\sqrt{v_{kl}}}{\text{sd}\{a_{kl}(\mathbf{x})\}} &= O_p\left\{\tau^{1/2}\|\mathbf{B}^*\|_{1,2}^{1/2}(1 \vee \tau^{1/2}\|\mathbf{B}^*\|_{1,2}^{1/2})\left(\frac{(K-1)\log p}{n}\right)^{1/4}\right\}. \end{aligned}$$

2. Under conditions of Theorem 6

$$\begin{aligned} \max_{k,l} \frac{m_{kl}}{\text{sd}\{a_{kl}(\mathbf{x})\}} &= O_p\left\{\tau\gamma\left(\frac{(K-1)s\log p}{n}\right)^{1/2}\right\}, \\ \max_{k,l} \frac{\sqrt{v_{kl}}}{\text{sd}\{a_{kl}(\mathbf{x})\}} &= O_p\left\{\tau\gamma\|\boldsymbol{\Sigma}_W^{1/2}\|_2\left(\frac{(K-1)s\log p}{n}\right)^{1/2}\right\}. \end{aligned}$$

Proof. By definition

$$\begin{aligned} u_{kl}(\mathbf{x}) - \hat{u}_{kl}(\mathbf{x}) &= -2\{\boldsymbol{\delta}_{kl}^\top \mathbf{B}^* (\mathbf{B}^{*\top} \boldsymbol{\Sigma}_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} - \mathbf{d}_{kl}^\top \hat{\mathbf{B}} (\hat{\mathbf{B}}^\top \hat{\boldsymbol{\Sigma}}_W \hat{\mathbf{B}})^{-1} \hat{\mathbf{B}}^\top\} \mathbf{x} \\ &\quad - \boldsymbol{\delta}_{kl}^\top \mathbf{B}^* (\mathbf{B}^{*\top} \boldsymbol{\Sigma}_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} \boldsymbol{\delta}_{kl} + \mathbf{d}_{kl}^\top \hat{\mathbf{B}} (\hat{\mathbf{B}}^\top \hat{\boldsymbol{\Sigma}}_W \hat{\mathbf{B}})^{-1} \hat{\mathbf{B}}^\top \mathbf{d}_{kl} \\ &\quad + 2\boldsymbol{\delta}_{kl}^\top \mathbf{B}^* (\mathbf{B}^{*\top} \boldsymbol{\Sigma}_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} \boldsymbol{\mu}_k - 2\mathbf{d}_{kl}^\top \hat{\mathbf{B}} (\hat{\mathbf{B}}^\top \hat{\boldsymbol{\Sigma}}_W \hat{\mathbf{B}})^{-1} \hat{\mathbf{B}}^\top \bar{\mathbf{x}}_k. \end{aligned}$$

Under Assumption 3, conditionally on the training data $(\mathbf{x}_i, \mathbf{z}_i)$, $i = 1, \dots, n$, and $\mathbf{x} \in \mathcal{C}_k$, $u_{kl}(\mathbf{x}) - \hat{u}_{kl}(\mathbf{x})$ is normal with mean

$$\begin{aligned} m_{kl} &= \mathbb{E}\{u_{kl}(\mathbf{x}) - \hat{u}_{kl}(\mathbf{x}) | \mathbf{x} \in \mathcal{C}_k, (\mathbf{x}_i, \mathbf{z}_i)\} \\ &= -\boldsymbol{\delta}_{kl}^\top \mathbf{B}^* (\mathbf{B}^{*\top} \boldsymbol{\Sigma}_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} \boldsymbol{\delta}_{kl} + \mathbf{d}_{kl}^\top \hat{\mathbf{B}} (\hat{\mathbf{B}}^\top \hat{\boldsymbol{\Sigma}}_W \hat{\mathbf{B}})^{-1} \hat{\mathbf{B}}^\top \mathbf{d}_{kl} + 2\mathbf{d}_{kl}^\top \hat{\mathbf{B}} (\hat{\mathbf{B}}^\top \hat{\boldsymbol{\Sigma}}_W \hat{\mathbf{B}})^{-1} \hat{\mathbf{B}}^\top (\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k) \end{aligned}$$

and variance

$$\begin{aligned} v_{kl} &= \text{Var}\{u_{kl}(\mathbf{x}) - \hat{u}_{kl}(\mathbf{x}) | \mathbf{x} \in \mathcal{C}_k, (\mathbf{x}_i, \mathbf{z}_i)\} \\ &= 4\boldsymbol{\delta}_{kl}^\top \mathbf{B}^* (\mathbf{B}^{*\top} \boldsymbol{\Sigma}_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} \boldsymbol{\delta}_{kl} + 4\mathbf{d}_{kl}^\top \hat{\mathbf{B}} (\hat{\mathbf{B}}^\top \hat{\boldsymbol{\Sigma}}_W \hat{\mathbf{B}})^{-1} (\hat{\mathbf{B}} \boldsymbol{\Sigma}_W \hat{\mathbf{B}}) (\hat{\mathbf{B}}^\top \hat{\boldsymbol{\Sigma}}_W \hat{\mathbf{B}})^{-1} \hat{\mathbf{B}}^\top \mathbf{d}_{kl} \\ &\quad - 8\boldsymbol{\delta}_{kl}^\top \mathbf{B}^* (\mathbf{B}^{*\top} \boldsymbol{\Sigma}_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} \boldsymbol{\Sigma}_W \hat{\mathbf{B}} (\hat{\mathbf{B}}^\top \hat{\boldsymbol{\Sigma}}_W \hat{\mathbf{B}})^{-1} \hat{\mathbf{B}}^\top \mathbf{d}_{kl}. \end{aligned}$$

Let $\mathbf{A} := \mathbf{B}^{*\top} \boldsymbol{\Sigma}_W \mathbf{B}^*$, $\hat{\mathbf{A}} := \hat{\mathbf{B}}^\top \hat{\boldsymbol{\Sigma}}_W \hat{\mathbf{B}}$, $\mathbf{q}_{kl} = \mathbf{A}^{-1/2} \mathbf{B}^{*\top} \boldsymbol{\delta}_{kl}$ and $\hat{\mathbf{q}}_{kl} = \hat{\mathbf{A}}^{-1/2} \hat{\mathbf{B}}^\top \mathbf{d}_{kl}$. Then we can simplify the expressions for m_{kl} and v_{kl} as

$$\begin{aligned} m_{kl} &= -\mathbf{q}_{kl}^\top \mathbf{q}_{kl} + \hat{\mathbf{q}}_{kl}^\top \hat{\mathbf{q}}_{kl} + 2\hat{\mathbf{q}}_{kl}^\top \hat{\mathbf{A}}^{-1/2} \hat{\mathbf{B}}^\top (\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k), \\ v_{kl} &= 4\mathbf{q}_{kl}^\top \mathbf{q}_{kl} + 4\hat{\mathbf{q}}_{kl}^\top \hat{\mathbf{A}}^{-1/2} (\hat{\mathbf{B}} \boldsymbol{\Sigma}_W \hat{\mathbf{B}}) \hat{\mathbf{A}}^{-1/2} \hat{\mathbf{q}}_{kl} - 8\mathbf{q}_{kl}^\top \mathbf{A}^{-1/2} (\mathbf{B}^{*\top} \boldsymbol{\Sigma}_W \hat{\mathbf{B}}) \hat{\mathbf{A}}^{-1/2} \hat{\mathbf{q}}_{kl}. \end{aligned}$$

Consider the mean term

$$\begin{aligned}
|m_{kl}| &\leq |\mathbf{q}_{kl}^\top \mathbf{q}_{kl} - \hat{\mathbf{q}}_{kl}^\top \hat{\mathbf{q}}_{kl}| + 2|\hat{\mathbf{q}}_{kl}^\top \hat{\mathbf{A}}^{-1/2} \hat{\mathbf{B}}^\top (\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k)| \\
&\leq 2|\mathbf{q}_{kl}^\top (\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl})| + \|\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl}\|_2^2 + 2\|\hat{\mathbf{A}}^{-1/2}\|_2 \|\hat{\mathbf{q}}_{kl}\|_2 \|\hat{\mathbf{B}}^\top (\boldsymbol{\mu}_1 - \bar{\mathbf{x}}_1)\|_2 \\
&\leq 2\|\mathbf{q}_{kl}\|_2 \|\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl}\|_2 + \|\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl}\|_2^2 + 2\|\hat{\mathbf{A}}^{-1/2}\|_2 \|\hat{\mathbf{B}}^\top (\boldsymbol{\mu}_1 - \bar{\mathbf{x}}_1)\|_2 (\|\mathbf{q}_{kl}\|_2 + \|\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl}\|_2) \\
&\leq 2\left(\|\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl}\|_2 + \|\hat{\mathbf{A}}^{-1/2}\|_2 \|\hat{\mathbf{B}}^\top (\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k)\|_2\right)\left(\|\mathbf{q}_{kl}\|_2 + \|\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl}\|_2\right)
\end{aligned}$$

Since $\text{sd}\{a_{kl}(\mathbf{x}) | \mathbf{x} \in \mathcal{C}_k\} = 2\sqrt{\delta_{kl}^\top \mathbf{B}^* (\mathbf{B}^{*\top} \boldsymbol{\Sigma}_W \mathbf{B}^*)^{-1} \mathbf{B}^{*\top} \delta_{kl}} = 2\sqrt{\mathbf{q}_{kl}^\top \mathbf{q}_{kl}} = 2\|\mathbf{q}_{kl}\|_2$, it follows that

$$\frac{m_{kl}}{\text{sd}\{a_{kl}(\mathbf{x})\}} \leq \left(\|\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl}\|_2 + \|\hat{\mathbf{A}}^{-1/2}\|_2 \|\hat{\mathbf{B}}^\top (\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k)\|_2\right)\left(1 + \|\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl}\|_2 / \|\mathbf{q}_{kl}\|_2\right).$$

We first focus on bounding $\|\mathbf{q}_{kl} - \hat{\mathbf{q}}_{kl}\|_2$. Let $\mathbf{E}_1 := \hat{\mathbf{A}} - \mathbf{A}$ and $\mathbf{e}_{kl} := \hat{\mathbf{B}}^\top \mathbf{d}_{kl} - \mathbf{B}^{*\top} \delta_{kl}$. Recall that $\mathbf{q}_{kl} = \mathbf{A}^{-1/2} \mathbf{B}^{*\top} \delta_{kl}$ and $\hat{\mathbf{q}}_{kl} = \hat{\mathbf{A}}^{-1/2} \hat{\mathbf{B}}^\top \mathbf{d}_{kl} = (\mathbf{A} + \mathbf{E}_1)^{-1/2} (\mathbf{B}^{*\top} \delta_{kl} + \mathbf{e}_{kl})$. Therefore

$$\begin{aligned}
\|\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl}\|_2 &\leq \|(\mathbf{A} + \mathbf{E})^{-1/2} \mathbf{B}^{*\top} \delta_{kl} - \mathbf{A}^{-1/2} \mathbf{B}^{*\top} \delta_{kl}\| + \|(\mathbf{A} + \mathbf{E}_1)^{-1/2} \mathbf{e}_{kl}\|_2 \\
&\leq \|(\mathbf{I} + \mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2})^{-1/2} - \mathbf{I}\|_2 \|\mathbf{A}^{-1/2} \mathbf{B}^{*\top} \delta_{kl}\|_2 + \|(\mathbf{I} + \mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2})^{-1/2}\|_2 \|\mathbf{A}^{-1/2} \mathbf{e}_{kl}\|_2.
\end{aligned}$$

If $\|\mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2}\|_2 \leq 1$, then for some constants $C_1, C_2 > 0$

$$\|\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl}\|_2 \leq C_1 \|\mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2}\|_2 \|\mathbf{q}_{kl}\|_2 + \|\mathbf{A}^{-1/2} \mathbf{e}_{kl}\|_2 (1 + C_2 \|\mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2}\|_2),$$

and

$$\frac{\|\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl}\|_2}{\|\mathbf{q}_{kl}\|_2} \leq C_1 \|\mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2}\|_2 + \frac{\|\mathbf{A}^{-1/2} \mathbf{e}_{kl}\|_2}{\|\mathbf{q}_{kl}\|_2} (1 + C_2 \|\mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2}\|_2).$$

Therefore, from Lemmas 9 and 10, under conditions of Theorem 4

$$\max_{k,l} \|\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl}\|_2 = O_p \left\{ (1 \vee \max_{k,l} \|\mathbf{q}_{kl}\|_2) (\tau \|\mathbf{B}^*\|_{1,2} \|\mathbf{A}^{-1}\|_2)^{1/2} \left(\frac{(K-1) \log p}{n} \right)^{1/4} \right\},$$

and

$$\max_{k,l} \frac{\|\hat{\mathbf{q}}_{kl} - \mathbf{q}_{kl}\|_2}{\|\mathbf{q}_{kl}\|_2} = O_p \left\{ (1 \vee \max_{k,l} \|\mathbf{q}_{kl}\|_2^{-1}) (\tau \|\mathbf{B}^*\|_{1,2} \|\mathbf{A}^{-1}\|_2)^{1/2} \left(\frac{(K-1) \log p}{n} \right)^{1/4} \right\}.$$

Similarly, consider

$$\max_k \|\hat{\mathbf{A}}^{-1/2}\|_2 \|\hat{\mathbf{B}}^\top (\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k)\|_2 \leq \|\mathbf{A}^{-1/2}\|_2 (1 + C_1 \|\mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2}\|_2) \|\hat{\mathbf{B}}\|_{1,2} \max_k \|\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k\|_\infty.$$

From Theorem 1, $\|\hat{\mathbf{B}}\|_{1,2} \leq 3\|\mathbf{B}^*\|_{1,2}$ when $\lambda \geq 2n^{-1} \|\mathbf{X}^\top \mathbf{E}\|_{\infty,2}$. Also, since $\bar{\mathbf{x}}_k \sim \mathcal{N}(\boldsymbol{\mu}_k, n_k^{-1} \boldsymbol{\Sigma}_W)$, applying Gaussian concentration inequality with union bound and Assumption 2 leads to $\max_k \|\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k\|_\infty = O_p(\tau \sqrt{\log(Kp)/n})$. Combining this with the above display and Lemma 9, under conditions of Theorem 4

$$\max_k \|\hat{\mathbf{A}}^{-1/2}\|_2 \|\hat{\mathbf{B}}^\top (\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k)\|_2 = O_p \left\{ \|\mathbf{A}^{-1/2}\|_2 \|\mathbf{B}^*\|_{1,2} \tau \sqrt{\frac{\log(Kp)}{n}} \right\}.$$

Combining the above displays, under conditions of Theorem 4, and using $K = o(p)$

$$\max_{k,l} \frac{m_{kl}}{\text{sd}\{a_{kl}(\mathbf{x})\}} = O_p \left\{ (\max_{k,l} \|\mathbf{q}_{kl}\|_2 \vee \max_{k,l} \|\mathbf{q}_{kl}\|_2^{-1}) (\tau \|\mathbf{B}^*\|_{1,2} \|\mathbf{A}^{-1}\|_2)^{1/2} \left(\frac{(K-1) \log p}{n} \right)^{1/4} \right\}.$$

If in addition Assumption 5 holds, then $\|\mathbf{A}^{-1/2}\|_2$ and $\|\mathbf{q}_{kl}\|_2$ are constants, therefore

$$\max_l \frac{m_{kl}}{\text{sd}\{a_{kl}(\mathbf{x})\}} = O_p\left\{(\tau\|\mathbf{B}^*\|_{1,2})^{1/2}\left(\frac{(K-1)\log p}{n}\right)^{1/4}\right\}.$$

Similarly, using Lemma 4 in [12],

$$\max_k \|\mathbf{B}^{*\top}(\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k)\|_2 = O_p\left(\sqrt{\frac{(K-1)\log K}{n}}\right),$$

therefore under conditions of Theorem 6 and using $K = o(p)$

$$\begin{aligned} \max_k \|\hat{\mathbf{B}}^\top(\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k)\|_2 &\leq \max_k \|(\hat{\mathbf{B}} - \mathbf{B}^*)^\top(\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k)\|_2 + \max_k \|\mathbf{B}^{*\top}(\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k)\|_2 \\ &\leq \max_k \|\hat{\mathbf{B}} - \mathbf{B}^*\|_{1,2} \|\boldsymbol{\mu}_k - \bar{\mathbf{x}}_k\|_\infty + O_p\left(\sqrt{\frac{(K-1)\log K}{n}}\right) \\ &\leq \|\hat{\mathbf{B}} - \mathbf{B}^*\|_F O_p\left(\tau\sqrt{\frac{\log(pK)}{n}}\right) + O_p\left(\sqrt{\frac{(K-1)\log K}{n}}\right) \\ &= O_p\left(\gamma\tau\sqrt{\frac{(K-1)s\log p}{n}}\right). \end{aligned}$$

From Lemmas 9 and 10, it follows that under conditions of Theorem 6

$$\max_{k,l} \frac{m_{kl}}{\text{sd}\{a_{kl}(\mathbf{x})\}} = O_p\left\{(\max_{k,l} \|\mathbf{q}_{kl}\|_2 \vee \max_{k,l} \|\mathbf{q}_{kl}\|_2^{-1})\tau\gamma\|\mathbf{A}^{-1}\|_2\left(\frac{(K-1)s\log p}{n}\right)^{1/2}\right\}.$$

If in addition Assumption 5 holds, then $\|\mathbf{A}^{-1/2}\|_2$ and $\|\mathbf{q}_{kl}\|_2$ are constants, therefore

$$\max_{k,l} \frac{m_{kl}}{\text{sd}\{a_{kl}(\mathbf{x})\}} = O_p\left\{\tau\gamma\left(\frac{(K-1)s\log p}{n}\right)^{1/2}\right\}.$$

Next, consider the standard deviation term

$$\begin{aligned} \sqrt{v_{kl}} &= 2\|\Sigma_W^{1/2}\mathbf{B}^*\mathbf{A}^{-1/2}\mathbf{q}_{kl} - \Sigma_W^{1/2}\hat{\mathbf{B}}\hat{\mathbf{A}}^{-1/2}\hat{\mathbf{q}}_{kl}\|_2 \\ &\leq 2\|\Sigma_W^{1/2}\mathbf{B}^*\mathbf{A}^{-1/2}\mathbf{q}_{kl} - \Sigma_W^{1/2}\mathbf{B}^*\mathbf{A}^{-1/2}\hat{\mathbf{q}}_{kl}\|_2 + \|\Sigma_W^{1/2}\mathbf{B}^*\mathbf{A}^{-1/2}\hat{\mathbf{q}}_{kl} - \Sigma_W^{1/2}\hat{\mathbf{B}}\hat{\mathbf{A}}^{-1/2}\hat{\mathbf{q}}_{kl}\|_2 \\ &\leq \|\Sigma_W^{1/2}\mathbf{B}^*\mathbf{A}^{-1/2}\|_2 \|\mathbf{q}_{kl} - \hat{\mathbf{q}}_{kl}\|_2 + \|\Sigma_W^{1/2}\mathbf{B}^*\mathbf{A}^{-1/2} - \Sigma_W^{1/2}\hat{\mathbf{B}}\hat{\mathbf{A}}^{-1/2}\|_2 \|\hat{\mathbf{q}}_{kl}\|_2 \\ &\leq \|\mathbf{q}_{kl} - \hat{\mathbf{q}}_{kl}\|_2 + \|\Sigma_W^{1/2}\mathbf{B}^*\mathbf{A}^{-1/2} - \Sigma_W^{1/2}\hat{\mathbf{B}}\hat{\mathbf{A}}^{-1/2}\|_2 (\|\mathbf{q}_{kl}\|_2 + \|\mathbf{q}_{kl} - \hat{\mathbf{q}}_{kl}\|_2), \end{aligned}$$

and correspondingly

$$\frac{\sqrt{v_{kl}}}{\text{sd}\{a_{kl}(\mathbf{x})\}} \leq \|\Sigma_W^{1/2}\mathbf{B}^*\mathbf{A}^{-1/2} - \Sigma_W^{1/2}\hat{\mathbf{B}}\hat{\mathbf{A}}^{-1/2}\|_2 + \frac{\|\mathbf{q}_{kl} - \hat{\mathbf{q}}_{kl}\|_2}{\|\mathbf{q}_{kl}\|_2} \left(1 + \|\Sigma_W^{1/2}\mathbf{B}^*\mathbf{A}^{-1/2} - \Sigma_W^{1/2}\hat{\mathbf{B}}\hat{\mathbf{A}}^{-1/2}\|_2\right).$$

Since

$$\begin{aligned} \|\Sigma_W^{1/2}\mathbf{B}^*\mathbf{A}^{-1/2} - \Sigma_W^{1/2}\hat{\mathbf{B}}\hat{\mathbf{A}}^{-1/2}\|_2 &= \|\Sigma_W^{1/2}\mathbf{B}^*\mathbf{A}^{-1/2} - \Sigma_W^{1/2}(\hat{\mathbf{B}} - \mathbf{B}^* + \mathbf{B}^*)(\mathbf{A} + \mathbf{E}_1)^{1/2}\|_2 \\ &\leq \|\Sigma_W^{1/2}\mathbf{B}^*\mathbf{A}^{-1/2} - \Sigma_W^{1/2}\mathbf{B}^*(\mathbf{A} + \mathbf{E}_1)^{-1/2}\|_2 + \|\Sigma_W^{1/2}(\hat{\mathbf{B}} - \mathbf{B}^*)\|_2 \\ &\leq C_1 \|\mathbf{A}^{-1/2}\mathbf{E}_1\mathbf{A}^{-1/2}\|_2 + \|\Sigma_W^{1/2}(\hat{\mathbf{B}} - \mathbf{B}^*)\|_2, \end{aligned}$$

under conditions of Theorem 4 and using Lemma 9

$$\begin{aligned} \|\Sigma_W^{1/2}\mathbf{B}^*\mathbf{A}^{-1/2} - \Sigma_W^{1/2}\hat{\mathbf{B}}\hat{\mathbf{A}}^{-1/2}\|_2 &\leq C_1 \|\mathbf{A}^{-1/2}\mathbf{E}_1\mathbf{A}^{-1/2}\|_2 + \|\Sigma_T^{1/2}(\hat{\mathbf{B}} - \mathbf{B}^*)\|_2 \\ &= O_p\left\{(\tau\|\mathbf{B}^*\|_{1,2}\|\mathbf{A}^{-1}\|_2)^{1/2}\left(\frac{(K-1)\log p}{n}\right)^{1/4}\right\} \\ &\quad + O_p\left\{(1 \vee \tau^{1/2}\|\mathbf{B}^*\|_{1,2}^{1/2})\tau^{1/2}\|\mathbf{B}^*\|_{1,2}^{1/2}\left(\frac{(K-1)\log p}{n}\right)^{1/4}\right\} \end{aligned}$$

Combining this with the above bounds on $\|\mathbf{q}_{kl} - \hat{\mathbf{q}}_{kl}\|_2$, under conditions of Theorem 4, using $K = o(p)$ and Assumption 5

$$\max_{k,l} \frac{\sqrt{v_{kl}}}{\text{sd}(a_{kl}(\mathbf{x}))} = O_p \left\{ \tau^{1/2} \|\mathbf{B}^*\|_{1,2}^{-1/2} (1 \vee \tau^{1/2} \|\mathbf{B}^*\|_{1,2}^{1/2}) \left(\frac{(K-1) \log p}{n} \right)^{1/4} \right\}.$$

Similarly, under conditions of Theorem 6 and using Lemma 9

$$\begin{aligned} \|\Sigma_W^{1/2} \mathbf{B}^* \mathbf{A}^{-1/2} - \Sigma_W^{1/2} \hat{\mathbf{B}} \hat{\mathbf{A}}^{-1/2}\|_2 &\leq C_1 \|\mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2}\|_2 + \|\Sigma_W^{1/2}\|_2 \|\hat{\mathbf{B}} - \mathbf{B}^*\|_F \\ &\leq O_p \left(\tau \gamma \|\mathbf{A}^{-1/2}\|_2 \sqrt{\frac{(K-1)s \log p}{n}} \right) + \|\Sigma_W^{1/2}\|_2 O_p \left(\tau \gamma \sqrt{\frac{(K-1)s \log p}{n}} \right). \end{aligned}$$

Combining with the bounds on $\|\mathbf{q}_{kl} - \hat{\mathbf{q}}_{kl}\|_2$, under conditions of Theorem 6, $K = o(p)$ and using Assumption 5

$$\max_{k,l} \frac{\sqrt{v_{kl}}}{\text{sd}(a_{kl}(\mathbf{x}))} = O_p \left\{ \tau \gamma \|\Sigma_W^{1/2}\|_2 \sqrt{\frac{(K-1)s \log p}{n}} \right\}.$$

□

Lemma 9. Let $\mathbf{A} := \mathbf{B}^{*\top} \Sigma_W \mathbf{B}^*$, $\hat{\mathbf{A}} := \hat{\mathbf{B}}^\top \hat{\Sigma}_W \hat{\mathbf{B}}$ and $\mathbf{E}_1 = \hat{\mathbf{A}} - \mathbf{A}$. Then under Assumption 3

$$\begin{aligned} &\|\mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2}\|_2 \\ &\leq \frac{\|\mathbf{X}^\top (\hat{\mathbf{B}} - \mathbf{B}^*)\|_F^2}{n-K} \|\mathbf{A}^{-1}\|_2 + 2 \frac{\|\mathbf{X}^\top (\hat{\mathbf{B}} - \mathbf{B}^*)\|_F}{\sqrt{n-K}} \|\mathbf{A}^{-1/2}\|_2 \left\{ 1 + O_p \left(\sqrt{\frac{K-1}{n-K}} \right) \right\} + O_p \left(\sqrt{\frac{K-1}{n-K}} \right). \end{aligned}$$

Moreover,

1. Under conditions of Theorem 4,

$$\|\mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2}\|_2 = O_p \left\{ (\tau \|\mathbf{B}^*\|_{1,2} \|\mathbf{A}^{-1}\|_2)^{1/2} \left(\frac{(K-1) \log p}{n} \right)^{1/4} \right\}.$$

2. Under conditions of Theorem 6,

$$\|\mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2}\|_2 = O_p \left\{ \tau \gamma \|\mathbf{A}^{-1/2}\|_2 \left(\frac{(K-1)s \log p}{n} \right)^{1/2} \right\}.$$

Proof. By definition

$$\mathbf{E}_1 = \hat{\mathbf{A}} - \mathbf{A} = (\hat{\mathbf{B}} - \mathbf{B}^*)^\top \hat{\Sigma}_W (\hat{\mathbf{B}} - \mathbf{B}^*) + 2 \mathbf{B}^{*\top} \hat{\Sigma}_W (\hat{\mathbf{B}} - \mathbf{B}^*) + \mathbf{B}^{*\top} \hat{\Sigma}_W \mathbf{B}^* - \mathbf{B}^{*\top} \Sigma_W \mathbf{B}^*,$$

therefore

$$\begin{aligned} \|\mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2}\|_2 &= \|\mathbf{A}^{-1/2} (\hat{\mathbf{B}} - \mathbf{B}^*)^\top \hat{\Sigma}_W (\hat{\mathbf{B}} - \mathbf{B}^*) \mathbf{A}^{-1/2}\|_2 \\ &\quad + 2 \|\mathbf{A}^{-1/2} \mathbf{B}^{*\top} \hat{\Sigma}_W (\hat{\mathbf{B}} - \mathbf{B}^*) \mathbf{A}^{-1/2}\|_2 + \|\mathbf{A}^{-1/2} \mathbf{B}^{*\top} \hat{\Sigma}_W \mathbf{B}^* \mathbf{A}^{-1/2} - \mathbf{I}\|_2 \\ &=: I_1 + 2I_2 + I_3. \end{aligned}$$

Since $\hat{\Sigma}_W$ is the pooled sample covariance matrix, $n^{-1} \mathbf{X}^\top \mathbf{X} = \frac{n-K}{n} \hat{\Sigma}_W + \hat{\Sigma}_B$, where Σ_B is between-class sample covariance matrix which is positive semi-definite. Therefore

$$\begin{aligned} I_1 &= \|\mathbf{A}^{-1/2} (\hat{\mathbf{B}} - \mathbf{B}^*)^\top \hat{\Sigma}_W (\hat{\mathbf{B}} - \mathbf{B}^*) \mathbf{A}^{-1/2}\|_2 \leq \|\mathbf{A}^{-1}\|_2 \|(\hat{\mathbf{B}} - \mathbf{B}^*)^\top (n-K)^{-1} \mathbf{X}^\top \mathbf{X} (\hat{\mathbf{B}} - \mathbf{B}^*)\|_2 \\ &\leq \|\mathbf{A}^{-1}\|_2 (n-K)^{-1} \|\mathbf{X}^\top (\hat{\mathbf{B}} - \mathbf{B}^*)\|_F^2. \end{aligned}$$

Similarly,

$$I_2 = \|\mathbf{A}^{-1/2} \mathbf{B}^{*\top} \hat{\Sigma}_W (\hat{\mathbf{B}} - \mathbf{B}^*) \mathbf{A}^{-1/2}\|_2 \leq \|\mathbf{A}^{-1/2} \mathbf{B}^{*\top} \hat{\Sigma}_W^{1/2}\|_2 (n-K)^{-1/2} \|\mathbf{X}^\top (\hat{\mathbf{B}} - \mathbf{B}^*)\|_F \|\mathbf{A}^{-1/2}\|_2.$$

By Assumption 3, $(n - K)\widehat{\Sigma}_W$ has Wishart distribution, $(n - K)\widehat{\Sigma}_W \sim \mathcal{W}_p(n - K, \Sigma_W)$, hence one can write $\Sigma_W = (n - K)^{-1} \sum_{i=1}^{n-K} \mathbf{z}_i \mathbf{z}_i^\top$, where $\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \Sigma_W)$ iid. Then $\mathbf{A}^{-1/2} \mathbf{B}^{*\top} \widehat{\Sigma}_W \mathbf{B}^* \mathbf{A}^{-1/2} = (n - K)^{-1} \sum_{i=1}^{n-K} \tilde{\mathbf{z}}_i \tilde{\mathbf{z}}_i^\top$ where $\tilde{\mathbf{z}}_i = \mathbf{A}^{-1/2} \mathbf{B}^{*\top} \mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{K-1})$ iid. By Lemma 9 in [37],

$$I_3 = \|\mathbf{A}^{-1/2} \mathbf{B}^{*\top} \widehat{\Sigma}_W \mathbf{B}^* \mathbf{A}^{-1/2} - \mathbf{I}\|_2 = O_p\left(\sqrt{\frac{K-1}{n-K}}\right)$$

and

$$\|\mathbf{A}^{-1/2} \mathbf{B}^{*\top} \widehat{\Sigma}_W^{1/2}\|_2 = \sqrt{\|\mathbf{A}^{-1/2} \mathbf{B}^{*\top} \widehat{\Sigma}_W \mathbf{B}^* \mathbf{A}^{-1/2}\|_2} \leq 1 + O_p\left(\sqrt{\frac{K-1}{n-K}}\right).$$

The first result follows by combining the bounds for I_1 , I_2 and I_3 .

Using Theorem 1 and conditions of Theorem 4, $n^{-1} \|\mathbf{X}^\top (\widehat{\mathbf{B}}^* - \mathbf{B})\|_F^2 = O_p\left\{\tau \|\mathbf{B}^*\|_{1,2} \left(\frac{(K-1) \log p}{n}\right)^{1/2}\right\}$.

Substituting this into bound for $\|\mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2}\|_2$ leads to desired statement.

Similarly, using Theorem 2 and conditions of Theorem 6, $n^{-1} \|\mathbf{X}^\top (\widehat{\mathbf{B}}^* - \mathbf{B})\|_F^2 = O_p\left\{\tau^2 \gamma^2 \frac{(K-1)s \log p}{n}\right\}$.

Substituting this into bound for $\|\mathbf{A}^{-1/2} \mathbf{E}_1 \mathbf{A}^{-1/2}\|_2$ leads to desired statement. \square

Lemma 10. *Let $\delta_{kl} = \mu_k - \mu_l$, $\mathbf{d}_{kl} = \bar{\mathbf{x}}_k - \bar{\mathbf{x}}_l$, $\mathbf{A} := \mathbf{B}^{*\top} \Sigma_W \mathbf{B}^*$, for some $k, l \in \{1, \dots, K\}$. Under Assumptions 2 and 3,*

$$\begin{aligned} \max_{k,l} \|\mathbf{A}^{-1/2} (\mathbf{B}^{*\top} \delta_{kl} - \widehat{\mathbf{B}}^\top \mathbf{d}_{kl})\|_2 \\ \leq \|\mathbf{A}^{-1/2}\|_2 n^{-1} \|\mathbf{X}^\top (\widehat{\mathbf{B}} - \mathbf{B}^*)\|_F \left\{1 + O_p\left(\sqrt{\frac{\log(K-1)}{n}}\right)\right\} + O_p\left(\sqrt{\frac{(K-1) \log(K-1)}{n}}\right). \end{aligned}$$

Moreover, if $K = o(p)$, then

1. under conditions of Theorem 4,

$$\max_{k,l} \|\mathbf{A}^{-1/2} (\mathbf{B}^{*\top} \delta_{kl} - \widehat{\mathbf{B}}^\top \mathbf{d}_{kl})\|_2 = O_p\left\{\left(\tau \|\mathbf{B}^*\|_{1,2} \|\mathbf{A}^{-1}\|_2\right)^{1/2} \left(\frac{(K-1) \log p}{n}\right)^{1/4}\right\}.$$

2. under conditions of Theorem 6,

$$\max_{k,l} \|\mathbf{A}^{-1/2} (\mathbf{B}^{*\top} \delta_{kl} - \widehat{\mathbf{B}}^\top \mathbf{d}_{kl})\|_2 = O_p\left\{\tau \gamma \|\mathbf{A}^{-1/2}\|_2 \left(\frac{(K-1)s \log p}{n}\right)^{1/2}\right\}.$$

Proof. Let $\mathbf{e}_{kl} = \mathbf{B}^{*\top} \delta_{kl} - \widehat{\mathbf{B}}^\top \mathbf{d}_{kl}$. Then by definition

$$\|\mathbf{A}^{-1/2} \mathbf{e}_{kl}\|_2 \leq \|\mathbf{A}^{-1/2}\|_2 \|(\widehat{\mathbf{B}} - \mathbf{B}^*)^\top \mathbf{d}_{kl}\|_2 + \|\mathbf{A}^{-1/2} \mathbf{B}^{*\top} (\mathbf{d}_{kl} - \delta_{kl})\|_2.$$

By Assumption 3, conditionally on the training data class assignments,

$$\mathbf{A}^{-1/2} \mathbf{B}^{*\top} (\mathbf{d}_{kl} - \delta_{kl}) \sim \mathcal{N}\left\{\mathbf{0}, \frac{n_k/n + n_l/n}{(n_k/n) \cdot (n_l/n)} n^{-1} \mathbf{I}\right\}.$$

Using Hoeffding's inequality $\text{pr}(|n_k/n - \pi_k| \geq \varepsilon) \leq C_1 \exp(-C_2 n \varepsilon^2)$, Assumption 2 and Lemma 4 in [12] together with union bound leads to

$$\max_{k,l} \|\mathbf{A}^{-1/2} \mathbf{B}^{*\top} (\mathbf{d}_{kl} - \delta_{kl})\|_2 = O_p\left(\sqrt{\frac{(K-1) \log(K-1)}{n}}\right).$$

On the other hand, since $\frac{1}{n} \mathbf{X}^\top \mathbf{X} = \frac{n-K}{n} \widehat{\Sigma}_W + \widehat{\Sigma}_B$, where $\widehat{\Sigma}_B$ is sample between-class covariance matrix, and for any k, l the decomposition

$$\widehat{\Sigma}_B = \sqrt{\frac{n_k n_l}{n^2}} \mathbf{d}_{kl} \mathbf{d}_{kl}^\top + \tilde{\Sigma}_B$$

holds for some positive semi-definite $\tilde{\Sigma}_B$ (Proposition 2 in [10]), it follows that for any k, l

$$\|(\hat{\mathbf{B}} - \mathbf{B}^*)^\top \mathbf{d}_{kl}\|_2 \leq \sqrt{\text{Tr}\{(\hat{\mathbf{B}} - \mathbf{B}^*)^\top \mathbf{d}_{kl} \mathbf{d}_{kl}^\top (\hat{\mathbf{B}} - \mathbf{B}^*)\}} \leq \sqrt{\|\mathbf{X}^\top (\hat{\mathbf{B}} - \mathbf{B}^*)\|_F^2 / \sqrt{n_k n_l}}.$$

Using Lemma 2 leads to

$$\max_{k,l} \|(\hat{\mathbf{B}} - \mathbf{B}^*)^\top \mathbf{d}_{kl}\|_2 \leq n^{-1} \|\mathbf{X}^\top (\hat{\mathbf{B}} - \mathbf{B}^*)\|_F \left\{ 1 + O_p\left(\sqrt{\frac{\log(K-1)}{n}}\right) \right\}.$$

Combining the above displays leads to

$$\max_{k,l} \|\mathbf{A}^{-1/2} \mathbf{e}_{kl}\|_2 \leq \|\mathbf{A}^{-1/2}\|_2 n^{-1} \|\mathbf{X}^\top (\hat{\mathbf{B}} - \mathbf{B}^*)\|_F \left\{ 1 + O_p\left(\sqrt{\frac{\log(K-1)}{n}}\right) \right\} + O_p\left(\sqrt{\frac{(K-1)\log(K-1)}{n}}\right).$$

The final statement follows by using corresponding bounds on $n^{-1} \|\mathbf{X}^\top (\hat{\mathbf{B}} - \mathbf{B}^*)\|_F$ following the proof of Lemma 9. \square

Lemma 11. *Under Assumption 2,*

$$\begin{aligned} \max_k |n_k/n - \pi_k| &= O_p\left\{\sqrt{\frac{\log(K-1)}{n}}\right\}. \\ \max_{k,l} |\log(n_l/n_k) - \log(\pi_l/\pi_k)| &= O_p\left\{\sqrt{\frac{\log(K-1)}{n}}\right\}. \end{aligned}$$

Proof. By Hoeffding's inequality, $\text{pr}(|n_k/n - \pi_k| \geq \varepsilon) \leq 2 \exp(-2n\varepsilon^2)$. Using $\pi_K = 1 - \sum_{k=1}^{K-1} \pi_k$, and applying the union bound gives the first result. Using the Taylor expansion for $f(n_k/n) = \log(n_k/n)$ around π_k , and Assumption 2 implies

$$\text{pr}(|\log(n_k/n) - \log(\pi_k)| \geq \varepsilon) \leq C_1 \exp(-C_2 n \varepsilon^2).$$

Using $\pi_K = 1 - \sum_{k=1}^{K-1} \pi_k$, and applying the union bound gives

$$\text{pr}\left(\bigcup_k \{|\log(n_k/n) - \log(\pi_k)| \geq \varepsilon\}\right) \leq C_1(K-1) \exp(-C_2 n \varepsilon^2).$$

The final claim follows since $\log(n_l/n_k) - \log(\pi_l/\pi_k) = \log(n_l/n) - \log(\pi_l) - \log(n_k/n) + \log(\pi_k)$. \square

Acknowledgement

The author is grateful to Jacob Bien and Mladen Kolar for valuable comments on the earlier draft of this work. The author also thanks the Editors and the Reviewer for the helpful suggestions that improved the manuscript. Gaynanova's research was supported by National Science Foundation grant DMS-1712943.

References

- [1] BICKEL, P. J. J., RITOV, Y. and TSYBAKOV, A. B. (2009). Simultaneous analysis of Lasso and Dantzig selector. *Annals of Statistics* **37** 1705–1732.
- [2] BIEN, J., GAYNANOVA, I., LEDERER, J. and MULLER, C. L. (2019). Prediction error bounds for linear regression with the TREX. *TEST* forthcoming.
- [3] BOUVEYRON, C., LATOUCHE, P. and MATTEI, P.-A. (2016). Bayesian Variable Selection for Globally Sparse Probabilistic PCA. *arXiv.org* arXiv:1605.05918.

[4] BÜHLMANN, P. and VAN DE GEER, S. A. (2011). *Statistics for High-Dimensional Data. Methods, Theory and Applications*. Springer-Verlag New York Inc.

[5] BUNEA, F., TSYBAKOV, A. B. and WEGKAMP, M. (2007). Sparsity oracle inequalities for the Lasso. *Electronic Journal of Statistics* **1** 169–194.

[6] CAI, T. and LIU, W. (2011). A direct estimation approach to sparse linear discriminant analysis. *Journal of the American Statistical Association* **106** 1566–1577.

[7] CHATTERJEE, S. (2013). Assumptionless consistency of the Lasso. *arXiv.org*.

[8] CLEMMENSEN, L., WITTEN, D. M., HASTIE, T. J. and ERSBØLL, B. (2011). Sparse Discriminant Analysis. *Technometrics* **53** 406–413.

[9] DALALYAN, A. S., HEBIRI, M. and LEDERER, J. (2017). On the prediction performance of the Lasso. *Bernoulli* **23** 552–581.

[10] GAYNANOVA, I., BOOTH, J. G. and WELLS, M. T. (2016). Simultaneous Sparse Estimation of Canonical Vectors in the $p >> N$ Setting. *Journal of the American Statistical Association* **111** 696–706.

[11] GAYNANOVA, I. and KOLAR, M. (2015). Optimal variable selection in multi-group sparse discriminant analysis. *Electronic Journal of Statistics* **9** 2007–2034.

[12] GAYNANOVA, I. and WANG, T. (2019). Sparse quadratic classification rules via linear dimension reduction. *Journal of Multivariate Analysis* **169** 278–299.

[13] GREENSHTEIN, E. and RITOV, Y. (2004). Persistence in high-dimensional linear predictor selection and the virtue of overparametrization. *Bernoulli* **10** 971–988.

[14] HAN, F., ZHAO, T. and LIU, H. (2013). CODA: High dimensional copula discriminant analysis. *Journal of Machine Learning Research* **14** 629–671.

[15] HASTIE, T. J., BUJA, A. and TIBSHIRANI, R. J. (1995). Penalized Discriminant Analysis. *Annals of Statistics* **23** 73–102.

[16] HASTIE, T. J., TIBSHIRANI, R. J. and BUJA, A. (1994). Flexible discriminant analysis by optimal scoring. *Journal of the American Statistical Association* **89** 1255–1270.

[17] HASTIE, T. J., TIBSHIRANI, R. J. and WAINWRIGHT, M. J. (2015). *Statistical Learning with Sparsity. The Lasso and Generalizations*. CRC Press.

[18] HSU, D., KAKADE, S. M. and ZHANG, T. (2012). A tail inequality for quadratic forms of subgaussian random vectors. *Electronic Communications in Probability* **17**.

[19] KOLAR, M. and LIU, H. (2015). Optimal feature selection in high-dimensional discriminant analysis. *IEEE Transactions on Information Theory* **61** 1063–1083.

[20] LEDERER, J., YU, L. and GAYNANOVA, I. (2019). Oracle Inequalities for High-dimensional Prediction. *Bernoulli* **25** 1225–1255.

[21] LI, Y. and JIA, J. (2017). L1 least squares for sparse high-dimensional LDA. *Electronic Journal of Statistics* **11** 2499–2518.

[22] LOUNICI, K., PONTIL, M., VAN DE GEER, S. A. and TSYBAKOV, A. B. (2011). Oracle inequalities and optimal inference under group sparsity. *Annals of Statistics* **39** 2164–2204.

[23] MAI, Q., ZOU, H. and YUAN, M. (2012). A direct approach to sparse discriminant analysis in ultra-high dimensions. *Biometrika* **99** 29–42.

[24] MARDIA, K. V., KENT, J. T. and BIBBY, J. M. (1979). *Multivariate Analysis*. Academic Press Inc.

[25] McLACHLAN, G. J. (1992). *Discriminant analysis and statistical pattern recognition*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley and Sons, Inc., New York.

[26] MERCHANTE, L. F. S., GRANDVALET, Y. and GOVAERT, G. (2012). An efficient approach to sparse linear discriminant analysis. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012* 1167–1174. Universite de Technologie de Compiegne, Compiegne, France.

[27] NEGAHBAN, S. N., RAVIKUMAR, P., WAINWRIGHT, M. J. and YU, B. (2012). A Unified Framework for High-Dimensional Analysis of M-Estimators with Decomposable Regularizers. *Statistical Science* **27** 538–557.

[28] NEYKOV, M., NING, Y., LIU, J. S. and LIU, H. (2018). A unified theory of confidence regions and testing for high dimensional estimating equations. *Statistical Science* **33** 427 - 443.

[29] OBOZINSKI, G., WAINWRIGHT, M. J. and JORDAN, M. I. (2011). Support union recovery in high-dimensional multivariate regression. *Annals of Statistics* **39** 1–47.

- [30] PAN, Y., MAI, Q. and ZHANG, X. (2018). Covariate-Adjusted Tensor Classification in High dimensions. *Journal of the American Statistical Association* 1–41.
- [31] RASKUTTI, G., WAINWRIGHT, M. J. and YU, B. (2010). Restricted eigenvalue properties for correlated Gaussian designs. *Journal of Machine Learning Research* **11** 2241–2259.
- [32] RIGOLLET, P. and TSYBAKOV, A. B. (2011). Exponential Screening and optimal rates of sparse estimation. *Annals of Statistics* **39** 731–771.
- [33] RUDELSON, M. and ZHOU, S. (2013). Reconstruction from anisotropic random measurements. *IEEE Transactions on Information Theory* **59** 3434–3447.
- [34] TIBSHIRANI, R. J. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Ser. B* **58** 267–288.
- [35] VAN DE GEER, S. A. and BÜHLMANN, P. (2009). On the conditions used to prove oracle results for the Lasso. *Electronic Journal of Statistics*.
- [36] VERSHYNIN, R. (2012). Introduction to the non-asymptotic analysis of random matrices. In *Compressed sensing* 210–268. Cambridge Univ. Press, Cambridge.
- [37] WAINWRIGHT, M. J. (2009). Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (Lasso). *IEEE Transactions on Information Theory* **55** 2183–2202.
- [38] WITTEN, D. M. and TIBSHIRANI, R. J. (2011). Penalized classification using Fisher’s linear discriminant. *Journal of the Royal Statistical Society, Ser. B* **73** 753–772.
- [39] WU, Y., WIPF, D. P. and YUN, J. M. (2015). Understanding and Evaluating Sparse Linear Discriminant Analysis. *AISTATS*.
- [40] YUAN, M. and LIN, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society, Ser. B* **68** 49–67.
- [41] ZHANG, T. (2009). Some sharp performance bounds for least squares regression with L 1 regularization. *Annals of Statistics* **37** 2109–2144.
- [42] ZHAO, P. and YU, B. (2006). On model selection consistency of Lasso. *Journal of Machine Learning Research* **7** 2541–2563.
- [43] ZHOU, S. (2009). Restricted Eigenvalue Conditions on Subgaussian Random Matrices.