Sparsifying the Fisher Linear Discriminant by Rotation

Ning Hao, Bin Dong, and Jianqing Fan

University of Arizona, University of Arizona, and Princeton University

August 4, 2014

Abstract

Many high dimensional classification techniques have been proposed in the literature based on sparse linear discriminant analysis (LDA). To efficiently use them, sparsity of linear classifiers is a prerequisite. However, this might not be readily available in many applications, and rotations of data are required to create the needed sparsity. In this paper, we propose a family of rotations to create the required sparsity. The basic idea is to use the principal components of the sample covariance matrix of the pooled samples and its variants to rotate the data first and to then apply an existing high dimensional classifier. This rotate-and-solve procedure can be combined with any existing classifiers, and is robust against the sparsity level of the true model. We show that these rotations do create the sparsity needed for high dimensional classifications and provide theoretical understanding why such a rotation works empirically. The effectiveness of the proposed method is demonstrated by a number of simulated and real data examples, and the improvements of our method over some popular high dimensional classification rules are clearly shown.

Keywords: Classification, Equivariance, Principal Components, High Dimensional Data, Linear Discriminant Analysis, Rotate-and-Solve.

1 Introduction

Linear discriminant analysis (LDA) is a useful classical tool for classification. Consider two p-dimensional normal distributions with the same covariance matrix, $N(\boldsymbol{\mu}_1, \boldsymbol{\Sigma})$ for class 1 and $N(\boldsymbol{\mu}_2, \boldsymbol{\Sigma})$ for class 2. Given a random vector \mathbf{X} which is from one of these distributions with equal prior probabilities, a *linear discriminant rule*

$$\psi_{\boldsymbol{\omega},\boldsymbol{\nu}}(\mathbf{X}) = I\{(\mathbf{X} - \boldsymbol{\nu})^{\top}\boldsymbol{\omega} \ge 0\}, \quad \boldsymbol{\omega}, \boldsymbol{\nu} \in \mathbb{R}^{p},$$
(1.1)

assigns **X** to class 1 when $\psi_{\omega,\nu}(\mathbf{X}) = 1$ and class 2 otherwise. Geometrically, the equation $(\mathbf{x} - \boldsymbol{\nu})^{\top} \boldsymbol{\omega} = 0$ defines an affine space passing through a point $\boldsymbol{\nu}$ with a normal vector $\boldsymbol{\omega}$, which is the discriminant boundary of the classification rule.

When μ_1 , μ_2 and Σ are known, the optimal classifier, namely the Fisher linear discriminant rule, is

$$\psi_F(\mathbf{X}) = I\{(\mathbf{X} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\delta} \ge 0\},\tag{1.2}$$

where $\boldsymbol{\mu} = \frac{1}{2}(\boldsymbol{\mu}_1 + \boldsymbol{\mu}_2), \, \boldsymbol{\delta} = \boldsymbol{\mu}_1 - \boldsymbol{\mu}_2$. In practice, these parameters are unknown and replaced by their estimates. Let $\{\mathbf{X}_i^{(1)} : 1 \leq i \leq n_1\}$ and $\{\mathbf{X}_i^{(2)} : 1 \leq i \leq n_2\}$ be independent and identically distributed (IID) observations from $N(\boldsymbol{\mu}_1, \boldsymbol{\Sigma})$ and $N(\boldsymbol{\mu}_2, \boldsymbol{\Sigma})$, respectively. In the classical setting with $n_1, n_2 \gg p, \, \boldsymbol{\mu}_1, \, \boldsymbol{\mu}_2$ and $\boldsymbol{\Sigma}^{-1}$ are usually estimated by sample means $\hat{\boldsymbol{\mu}}_1 = \bar{\mathbf{X}}^{(1)}, \, \hat{\boldsymbol{\mu}}_2 = \bar{\mathbf{X}}^{(2)}$ and the inverse pooled sample covariance matrix $\hat{\boldsymbol{\Sigma}}^{-1}$. The standard linear discriminant analysis (LDA) uses an empirical version of (1.2)

$$\psi_{\hat{F}}(\mathbf{X}) = I\{(\mathbf{X} - \hat{\boldsymbol{\mu}})^{\top} \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\delta}} \ge 0\},$$
(1.3)

where $\hat{\boldsymbol{\mu}} = \frac{1}{2}(\hat{\boldsymbol{\mu}}_1 + \hat{\boldsymbol{\mu}}_2), \ \hat{\boldsymbol{\delta}} = \hat{\boldsymbol{\mu}}_1 - \hat{\boldsymbol{\mu}}_2.$

Although the standard LDA has been widely used in applications, it does not work well for high dimensional data when p is comparable to or larger than the sample size. The reason is that, with limited number of observations, it is impossible to estimate too many parameters simultaneously and accurately. In particular, $\hat{\Sigma}$ is singular and not invertible when $n_1 + n_2 . One may use pseudo-inverse <math>\hat{\Sigma}^-$, but Bickel & Levina (2004) showed the LDA performs as poorly as random guessing when $p/(n_1 + n_2) \to \infty$. Since the work of Bickel & Levina (2004), a series of LDA-based methods have been proposed for the high dimensional classification problem. The main idea is to find methods which work well when the original classification problem is (nearly) sparse so that μ or $\beta = \Sigma^{-1}\delta$ in the optimal rule (1.2) can be well estimated. Ignoring the covariances among the features, Bickel & Levina (2004) proposed an independence rule (IR) which outperforms standard LDA in the high dimensional setting. Fan & Fan (2008) proposed the features annealed independence rule (FAIR) that selects a subset of features before applying the independence rule. In spite of the clear interpretations of the sparsity of the covariance matrix Σ and difference of centroids δ , in practice, it might be more efficient to find the sparse discriminant affine space directly (see Trendafilov & Jolliffe (2007); Wu et al. (2009); Cai & Liu (2011); Fan et al. (2012); Mai et al. (2012) among others). Here, a sparse discriminant affine space is an affine space with a sparse normal vector. In particular, Fan et al. (2012) and Cai & Liu (2011) clearly illustrated the advantages of their direct approaches over IR and FAIR, which oversimplify the problem in many scenarios.

For all aforementioned LDA-based high dimensional classification rules, various explicit sparsity conditions on one or some of Σ , Σ^{-1} , δ and β are crucial to the classification accuracy. For example, IR (Bickel & Levina, 2004) works well only when Σ is nearly diagonal; FAIR (Fan & Fan, 2008) needs ideally diagonal Σ and sparse δ ; ROAD (Fan et al., 2012) and LPD (Cai & Liu, 2011) need β to be sparse to achieve optimal classification. We shall refer to all of these methods as sparse LDA methods. They are efficient when the corresponding sparsity conditions are granted. However, they may not work well when the sparsity conditions are violated. Although these sparse assumptions make sense in some applications, they can be too restrictive in many scenarios (see Hall et al. (2009) and reference therein). It is a natural and challenging question how and to what extent we can sparsify a possibly non-sparse problem.

To solve a non-sparse model, a natural idea is to rotate the data to a nearly sparse setting before applying sparse LDA methods. For example, the classification problem can be easily solved by ROAD and LPD if the normal vector of the optimal discriminant affine space, β , is sparse after a rotation. In order to do this, we need an oracle that can rotate the data to such a sparse setting. For the ideal case when β is known, there are infinitely many orthogonal matrices which can rotate β to a sparse vector $(||\beta||_2, 0, ..., 0)^{\top}$. However, it is not realistic to approximate such rotations before estimating β itself. An alternative way might be to make Σ diagonal after a rotation, which is related to principal component analysis (PCA). However, such a rotation does not combine the information of the centroids and tends to get wrong directions with small variances, which may actually be crucial for classification.

In this paper, we propose a class of rotations which balance both mean and variance information. Intuitively, both δ and Σ should play essential roles in a rotation to make β sparse. In particular, if Σ is spiked (Johnstone, 2001), its principal components and δ span a linear space, which contains key information on the rotation. Following this intuition, we define $\Sigma_{\rho}^{tot} = \Sigma + \rho \delta \delta^{\top}$ for $\rho > 0$, whose principal components are determined by the ones of Σ as well as δ . Consider an orthogonal matrix \mathbf{U}_{ρ} , formed by the eigenvectors of Σ_{ρ}^{tot} , which diagonalizes Σ_{ρ}^{tot} . We shall show that $\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}$ is sparse when the covariance matrix $\boldsymbol{\Sigma}$ is spiked. In other words, the eigenvectors of Σ_{ρ}^{tot} are good directions to rotate. Similarly, we can define the empirical version $\hat{\mathbf{U}}_{\rho}$ which diagonalizes $\hat{\boldsymbol{\Sigma}}_{\rho}^{tot} = \hat{\boldsymbol{\Sigma}} + \rho \hat{\boldsymbol{\delta}} \hat{\boldsymbol{\delta}}^{\top}$. The rotation $\hat{\mathbf{U}}_{\rho}$ is a reasonably good approximation to \mathbf{U}_{ρ} when $p \ll n$ (Johnstone & Lu, 2009) or p > nwith some additional conditions (Zou et al., 2006; Fan et al., 2013). In other words, under some conditions on Σ , $\hat{\mathbf{U}}_{\rho}^{\top}\boldsymbol{\beta}$ is nearly sparse, regardless of the sparsity level of the original $\boldsymbol{\beta}$. Therefore, we propose to rotate the data by $\hat{\mathbf{U}}_{\rho}^{\top}$ first before applying ROAD or LPD, when the sparsity level of β is unknown. While our original motivation is to make β sparse by rotation, we find that our procedure is equivariant with respect to orthogonal transformation group O(p) consisting of all rotations. This feature makes our method robust against the sparsity level of β . The advantage of our method is illustrated by numerous simulated and real data examples. An implementation of the proposed rotations using MATLAB language can be found on the web http://math.arizona.edu/~dongbin/Data/HDF_Rotation.m.

The rest of our paper is organized as follows. Section 2 introduces a family of ideal rotations and analyzes their theoretical properties. In Section 3, we study a rotate-and-solve procedure for classification. Numerical studies on both simulated and real data are demonstrated in Section 4. All proofs are given in the appendix. Various norms of vectors and matrices appear frequently in the paper. For a vector \mathbf{a} , $||\mathbf{a}||_p$ denote the standard ℓ_p -norm. For a matrix \mathbf{A} , $||\mathbf{A}||$ is the spectral norm.

2 A family of oracle rotations and their properties

As mentioned in the introduction, the performance of the sparse LDA methods depend highly on the sparsity of β , which is unknown and hard to verify in practice. High dimensional classifiers will work more efficiently if an oracle rotates the data to a sparse setting before applying sparse LDA methods. If β is known, we can easily rotate β to a sparse vector $(||\beta||_2, 0, ..., 0)^{\top}$. Of course, it is meaningless to mimic such oracle, which motivates us to find other ideal rotations that can be estimated more easily.

Recall that the distributions of two classes are $N(\boldsymbol{\mu}_1, \boldsymbol{\Sigma})$ for class 1 and $N(\boldsymbol{\mu}_2, \boldsymbol{\Sigma})$ for class 2. Let $\boldsymbol{\mu} = \frac{1}{2}(\boldsymbol{\mu}_1 + \boldsymbol{\mu}_2), \, \boldsymbol{\delta} = \boldsymbol{\mu}_1 - \boldsymbol{\mu}_2$, and

$$\Sigma_{\rho}^{tot} = \Sigma + \rho \delta \delta^{\top}, \quad \text{for a given } \rho > 0.$$

Consider an orthogonal matrix \mathbf{U}_{ρ} , formed by the eigenvectors of $\boldsymbol{\Sigma}_{\rho}^{tot}$, which diagonalizes $\boldsymbol{\Sigma}_{\rho}^{tot}$. Then, without loss of generality by rearranging columns in \mathbf{U}_{ρ} , we assume that $\mathbf{U}_{\rho}^{\top} \boldsymbol{\Sigma}_{\rho}^{tot} \mathbf{U}_{\rho} = \mathbf{D}_{\rho}$ where $\mathbf{D}_{\rho} = diag(\eta_1, ..., \eta_p)$ is the diagonal matrix, consisting of eigenvalues in descending order.

Let $\{\lambda_j\}_{j=1}^p$ be eigenvalues of Σ , arranged from the largest to the smallest, and $\{\boldsymbol{\xi}_j\}_{j=1}^p$ be their corresponding eigenvectors. Note that, for repeated eigenvalues, say $\lambda_r = \lambda_{r+1} = \cdots = \lambda_s$, $\{\boldsymbol{\xi}_j\}_{j=r}^s$ can be chosen as any orthonormal basis of the corresponding eigenspace. Johnstone (2001) considered a spiked covariance model, where a few large eigenvalues clearly standing out of the rest.

Condition 1 (Spiked Covariance Structure): Assume that $\lambda_1 \geq \cdots \geq \lambda_k > \lambda_{k+1} = \cdots = \lambda_p$ for some integer k < p.

Theorem 1 Under Condition 1, we have $||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_{0} \leq k+1$.

Theorem 1 shows the sparsity property of $\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}$ when $\boldsymbol{\Sigma}$ is spiked and k + 1 < p. In particular, it implies that $||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_{1}/||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_{2} \leq \sqrt{k+1}$ by Cauchy-Schwarz inequality. The boundedness of the ℓ_{0} or ℓ_{1} norm is crucial for sparse LDA methods such as ROAD and LPD to be efficient. For a vector randomly picked on the unit sphere in \mathbb{R}^{p} , the expectation of its ℓ_{1} norm is of order \sqrt{p} . Therefore, both ℓ_{0} and ℓ_{1} norms of $\boldsymbol{\beta}$ have been greatly reduced after rotation when $k \ll p$.

The condition of Theorem 1 can still be relaxed somehow while keeping $||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_{1}/||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_{2}$ bounded. This is shown in Theorem 2 below.

Condition 2 (Quasi-Spiked Covariance Structure): Assume that $\lambda_k \geq \lambda_{k+1} + d$ and $\lambda_{k+1} - \lambda_p \leq \epsilon$ for some integer k < p, where $d, \epsilon > 0$.

Let \mathbf{W}_1 and \mathbf{W}_2 be two linear spaces spanned by $\{\boldsymbol{\xi}_j\}_{1 \leq j \leq k}$ and $\{\boldsymbol{\xi}_j\}_{k+1 \leq j \leq p}$, respectively. Then, we have $\mathbb{R}^p = \mathbf{W}_1 \oplus \mathbf{W}_2$ and the mean difference vector $\boldsymbol{\delta}$ can be decomposed as $\boldsymbol{\delta} = \boldsymbol{\delta}_1 + \boldsymbol{\delta}_2$ with $\boldsymbol{\delta}_1 \in \mathbf{W}_1$ and $\boldsymbol{\delta}_2 \in \mathbf{W}_2$.

Theorem 2 If $\boldsymbol{\delta} \in \mathbf{W}_1$ and $\lambda_k > \lambda_{k+1}$, then $||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_0 \leq k$ and

$$||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_{1}/||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_{2} \leq \sqrt{k}.$$

If $\boldsymbol{\delta} \notin \mathbf{W}_1$ and Condition 2 holds, then

$$||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_{1}/||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_{2} \leq \sqrt{k+1} + \sqrt{p-k-1}\frac{\lambda_{p}+\epsilon}{\lambda_{p}}\left(\frac{\epsilon}{\lambda_{p}} + \sqrt{\frac{\epsilon}{\tilde{d}-2\epsilon}}\right)$$

provided $\epsilon < \tilde{d}/2$, where $\tilde{d} = d \frac{\rho || \delta_2 ||_2^2}{d + \rho || \delta ||_2^2}$.

Theorem 1 and the first part of Theorem 2 demonstrate that the sparsity can be achieved after rotation even measured by the strong notion ℓ_0 -norm. However, the weaker measure of sparsity using ℓ_1 -norm is needed in order to obtain more general results, as shown in the second part of Theorem 2.

As a direct consequence of Theorem 2, we have the following corollary.

Corollary 1 If
$$\frac{\epsilon}{\lambda_p} = O(\sqrt{k/p})$$
 and $\frac{\epsilon ||\boldsymbol{\delta}||_2^2}{d||\boldsymbol{\delta}_2||_2^2} = O(k/p)$, then $||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_1/||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_2 = O(\sqrt{k})$.

Note that the construction of \mathbf{U}_{ρ} is independent of k, and conclusions of Theorem 2 hold for any k satisfying the technical conditions. Define $d_k = \lambda_k - \lambda_{k+1}$, $\epsilon_k = \lambda_{k+1} - \lambda_p$, and $\mathbf{W}_1^k = \operatorname{span}\{\boldsymbol{\xi}_j\}_{1 \leq j \leq k}$, $\mathbf{W}_2^k = \operatorname{span}\{\boldsymbol{\xi}_j\}_{k+1 \leq j \leq p}$, $\boldsymbol{\delta} = \boldsymbol{\delta}_1^k + \boldsymbol{\delta}_2^k$ with $\boldsymbol{\delta}_m^k \in \mathbf{W}_m^k$, m = 1, 2. Let $\tilde{d}_k = d_k \frac{\rho ||\boldsymbol{\delta}_2^k||_2^2}{d_k + \rho ||\boldsymbol{\delta}||_2^2}$. Define $C_k = \sqrt{k+1} + \sqrt{p-k-1} \frac{\lambda_p + \epsilon_k}{\lambda_p} (\frac{\epsilon_k}{\lambda_p} + \sqrt{\frac{\epsilon_k}{d_k - 2\epsilon_k}})$ if $\tilde{d}_k - 2\epsilon_k > 0$, and $C_k = \infty$ otherwise. Theorem 2 implies the following corollary.

Corollary 2 If K is the least integer such that $\boldsymbol{\delta} \in \mathbf{W}_1^K$, then $||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_1/||\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}||_2 \leq \min\{C, \sqrt{K}\}$, where $C = \min_{1 \leq k < K}\{C_k\}$.

Theorems 1 and 2 show that the classification problem is reduced to a sparse one after rotation by \mathbf{U}_{ρ}^{\top} when the covariance structure is spiked. And the sparsity level of $\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}$ can be controlled by the spiked covariance structure (k and eigenvalue distribution in Conditions 1 and 2).

Moreover, the procedure is invariant under orthonormal transformations. In other words, the normal vector of the optimal discriminant affine space after rotation, i.e., $\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}$, is invariant with respect to any rotation. Indeed, when the data are rotated by an arbitrary orthogonal matrix \mathbf{V} , then the new mean vectors and common covariance matrix are $\mathbf{V}\boldsymbol{\mu}_1$, $\mathbf{V}\boldsymbol{\mu}_2$ and $\mathbf{V}\boldsymbol{\Sigma}\mathbf{V}^{\top}$. Since

$$\mathbf{D}_{\rho} = \mathbf{U}_{\rho}^{\top} \boldsymbol{\Sigma}_{\rho}^{tot} \mathbf{U}_{\rho} = (\mathbf{V} \mathbf{U}_{\rho})^{\top} \mathbf{V} \boldsymbol{\Sigma}_{\rho}^{tot} \mathbf{V}^{\top} (\mathbf{V} \mathbf{U}_{\rho}),$$

the rotation matrix should be $(\mathbf{V}\mathbf{U}_{\rho})^{\top}$, and the rotated normal vector $(\mathbf{V}\mathbf{U}_{\rho})^{\top}\mathbf{V}\boldsymbol{\beta} = \mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}$, which is independent of \mathbf{V} .

3 A Rotate-and-Solve Procedure

In this section, we introduce a two-stage rotate-and-solve (RS) procedure for classification. The idea is to mimic the oracle rotations in the previous section and rotate the data such that $\boldsymbol{\beta}$ is nearly sparse. Namely, we first use the orthogonal matrix $\hat{\mathbf{U}}_{\rho}$, consisting of the eigenvectors of the empirical total covariance $\hat{\boldsymbol{\Sigma}}_{\rho}^{tot} = \hat{\boldsymbol{\Sigma}} + \rho \hat{\boldsymbol{\delta}} \hat{\boldsymbol{\delta}}^{\top}$ to rotate the data and then apply sparse LDA methods such as ROAD and LPD to the rotated data.

Let $\hat{\mu}_1$ and $\hat{\mu}_2$ be the sample mean vectors of classes 1 and 2 respectively. Set

$$\hat{\boldsymbol{\mu}} = (\hat{\boldsymbol{\mu}}_1 + \hat{\boldsymbol{\mu}}_2)/2, \text{ and } \hat{\boldsymbol{\delta}} = \hat{\boldsymbol{\mu}}_1 - \hat{\boldsymbol{\mu}}_2$$

Similarly, let $\hat{\Sigma}^{(1)}$ and $\hat{\Sigma}^{(2)}$ be their sample covariance matrices and

$$\hat{\Sigma} = \frac{1}{n_1 + n_2} (n_1 \hat{\Sigma}^{(1)} + n_2 \hat{\Sigma}^{(2)})$$

be the pooled sample covariance matrix. The degree of freedom can be adjusted, but the version of the maximum likelihood estimate (MLE) is used here to facilitate the expression in Remark 1 below. We then estimate Σ_{ρ}^{tot} by

$$\hat{\boldsymbol{\Sigma}}_{\boldsymbol{\rho}}^{tot} = \hat{\boldsymbol{\Sigma}} + \boldsymbol{\rho} \hat{\boldsymbol{\delta}} \hat{\boldsymbol{\delta}}^{\top}.$$

Perform singular-value decomposition

$$\hat{\mathbf{U}}_{\rho}^{\top} \hat{\boldsymbol{\Sigma}}_{\rho}^{tot} \hat{\mathbf{U}}_{\rho} = \hat{\mathbf{D}}_{\rho}, \qquad (3.1)$$

where $\hat{\mathbf{D}}_{\rho} = \text{diag}(\hat{\eta}_1, ..., \hat{\eta}_p)$ is the diagonal matrix with sorted eigenvalues.

The two-stage rotate-and-solve procedure can be implemented as follows.

Stage one: Calculate $\hat{\mathbf{U}}_{\rho}$ and rotate the data to get $\{\hat{\mathbf{U}}_{\rho}^{\top}\mathbf{X}_{i}^{(m)}\}_{i=1}^{n_{m}}$ for m = 1 and 2.

Stage two: apply ROAD, LPD or other sparse LDA methods to the rotated data $\mathcal{X}\hat{\mathbf{U}}_{\rho}$ to get a prediction rule.

Remark 1: Define

$$\bar{\mathbf{X}} = \frac{1}{n_1 + n_2} (n_1 \bar{\mathbf{X}}^{(1)} + n_2 \bar{\mathbf{X}}^{(2)}),$$
$$\hat{\mathbf{\Sigma}}_{sample}^{tot} = \frac{1}{n_1 + n_2} \sum_{m=1}^{2} \sum_{i=1}^{n_m} (\mathbf{X}_i^{(k)} - \bar{\mathbf{X}}) (\mathbf{X}_i^{(k)} - \bar{\mathbf{X}})^{\mathsf{T}}$$

which is the sample total covariance (ignoring the classes). It is straightforward to see $\hat{\Sigma}_{sample}^{tot} = \hat{\Sigma} + \frac{n_1 n_2}{(n_1 + n_2)^2} \hat{\delta} \hat{\delta}^{\top}$.

When $p \ll n = n_1 + n_2$, $\hat{\mathbf{U}}_{\rho}$ and \mathbf{U}_{ρ} are similar when the eigenvalues are separated from each other, and hence $\hat{\mathbf{U}}_{\rho}^{\top}\boldsymbol{\beta}$ is similar to $\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}$. The property of $\hat{\mathbf{U}}_{\rho}^{\top}\boldsymbol{\beta}$ is much more complicated when $p \sim n$ or $p \gg n$. In this case, it is hard to guarantee all estimated eigenvectors are close to the true ones. However, the eigenvectors that correspond to spiked eigenvalues can be consistently estimated. See for example Zou et al. (2006); Karoui (2008); Johnstone & Lu (2009); Agarwal et al. (2012); Fan et al. (2013); Shen et al. (2013). As these eigenvectors point at most important directions, the consistent estimation of these directions ensures the correct rotations in these important directions. This explains our empirical results that the RS procedure performs very well compared to several state-of-the-art methods, even when $p \gg n$.

To understand better the mathematics behind the excellent performance of RS procedure, the classification error of the idealized Fisher classifier depends on $\gamma \equiv \delta^{\top} \Sigma^{-1} \delta$. Let $\mathbf{U}_{\rho 1}$ be a $(k+1) \times p$ matrix, consisting of the eigenvectors of Σ_{ρ}^{tot} that correspond to the largest k+1eigenvalues $\{\eta_j\}_{j=1}^{k+1}$. If we restrict the information to the first k+1 dimensions of the rotated data $\mathbf{U}_{\rho 1}^{\top} \mathbf{X} |_m \sim N(\mathbf{U}_{\rho 1}^{\top} \boldsymbol{\mu}_m, \mathbf{U}_{\rho 1}^{\top} \Sigma \mathbf{U}_{\rho 1}), m = 1, 2$, then the classification error depends on

$$\gamma_1 \equiv (\mathbf{U}_{\rho_1}^{\top} \boldsymbol{\delta})^T (\mathbf{U}_{\rho_1}^{\top} \boldsymbol{\Sigma} \mathbf{U}_{\rho_1})^{-1} (\mathbf{U}_{\rho_1}^{\top} \boldsymbol{\delta}).$$

Clearly, $\gamma_1 \leq \gamma$. How much is the information loss when $\{\eta_j\}_{j=1}^{k+1}$ are spiked? Under Conditions in Theorem 1, there is no information loss if the first k + 1 most important features are used. Furthermore, the cited literatures above give the conditions under which \mathbf{U}_{ρ_1} can be consistently estimated.

The above argument is based on the fact that $\mathbf{U}_{\rho 1}^{\top} \boldsymbol{\delta}$ preserves the energy of $\boldsymbol{\delta}$. The result holds more generally for the covariance matrix $\boldsymbol{\Sigma}$ admitting spiked eigenvalues, including covariance matrices derived from approximate factor models (Fan et al., 2013) or admitting low rank plus sparse matrix decomposition (Agarwal et al., 2012). Recall that $\boldsymbol{\Sigma} = \sum_{i=1}^{p} \lambda_i \boldsymbol{\xi}_i \boldsymbol{\xi}_i^{\top}$ with $\boldsymbol{\xi}_i$ being the eigenvector of $\boldsymbol{\Sigma}$. Let $\lambda_i(\mathbf{B})$ be the i^{th} largest eigenvalue of a symmetric matrix \mathbf{B} .

Theorem 3 If $\lambda_{k+1}(\sum_{i=1}^k \lambda_i \boldsymbol{\xi}_i \boldsymbol{\xi}_i^\top + \rho \boldsymbol{\delta} \boldsymbol{\delta}^\top) > a\lambda_{k+1}$ for some a > 2, then $||\mathbf{U}_{\rho 1}\boldsymbol{\delta}||_2 \ge \frac{a-2}{a-1} ||\boldsymbol{\delta}||_2$ and $\gamma_1 \ge \frac{(a-2)^2}{(a-1)^2\lambda_1} ||\boldsymbol{\delta}||_2^2$.

The condition of Theorem 3 holds relatively easily. We can take k = 0 when $\rho \|\boldsymbol{\delta}\|_2^2 \ge a \|\boldsymbol{\Sigma}\|^2$. This holds easily by taking a sufficiently large ρ .

Note that $\gamma \leq \lambda_p^{-1} \|\boldsymbol{\delta}\|_2^2$ and γ is usually significantly smaller than this upper bound. Therefore, when λ_1/λ_p is bounded, the loss of information by using rotated data is limited. Yet, we reduce significantly the noise accumulation in classification (Fan & Fan, 2008). As noted above, the rotation $\mathbf{U}_{\rho 1}$ can be consistently estimated by regularization. These together provide theoretical endorsement of the advantages of using rotation.

Remark 2: (Dimensionality reduction) When p > n, $\hat{\mathbf{U}}_{\rho}$ is not unique since $\hat{\boldsymbol{\Sigma}}_{\rho}^{tot}$ is singular. (The null space of $\hat{\boldsymbol{\Sigma}}_{\rho}^{tot}$ is large and we can choose arbitrary basis of the null space as the columns of $\hat{\mathbf{U}}_{\rho}$.) Since the last p - n columns in $\hat{\mathbf{U}}_{\rho}$ are arbitrary and can not be controlled, we define $\tilde{\mathbf{U}}_{\rho}$ as the first n columns (or even fewer) of $\hat{\mathbf{U}}_{\rho}$ and conduct classification on the rotated data $\{\tilde{\mathbf{U}}_{\rho}^{\top}\mathbf{X}_{i}^{(m)}\}_{i=1}^{n_{m}}$ for m = 1 and 2. From the theoretical analysis in the last section, we see that, under ideal conditions, $\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}$ is sparse with nonvanishing part concentrated on the first k + 1 components. This implies that only first k + 1columns of the rotated data are useful to estimate $\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}$, which motivates us to use $\tilde{\mathbf{U}}_{\rho}$ instead of $\hat{\mathbf{U}}_{\rho}$ as a practical approach with reduced dimensionality. Theorem 3 further shows that the loss of classification power due to this dimensionality reduction is limited. Let $\tilde{\psi}$ be a classification rule constructed by some (fixed) sparse LDA method based on $\tilde{\mathcal{X}} = \mathcal{X}\tilde{\mathbf{U}}_{\rho}$. It is straightforward to see that $\tilde{\psi}$ is equivariant.

Remark 3 (Computation of transform) When p > n, the computation of $\tilde{\mathbf{U}}_{\rho}$ can be performed as follows. First of all, $\hat{\mathbf{\Sigma}}_{\rho}^{tot}$ can be written as $\mathbf{Y}^{\top}\mathbf{Y}$ for a given $(n+1) \times p$ matrix \mathbf{Y} (suitable scaling of centered observations and sample mean). Note that $\mathbf{Y}^{\top}\mathbf{Y}$ and $\mathbf{Y}\mathbf{Y}^{\top}$ have the same non-vanishing eigenvalues. Let $\tilde{\mathbf{U}}_{\rho} = \mathbf{Y}^{\top}\hat{\mathbf{V}}$, where $\hat{\mathbf{V}}$ is the orthogonal matrix consisting of eigenvectors of non-vanishing eigenvalues of the $(n+1) \times (n+1)$ matrix $\mathbf{Y}\mathbf{Y}^{\top}$. Then, the columns of $\tilde{\mathbf{U}}_{\rho}$ contain the eigenvectors of nonvanishing eigenvalues of $\mathbf{Y}^{\top}\mathbf{Y}$ and are orthogonal. In other words, $\tilde{\mathbf{U}}_{\rho}$ can be used to transform the data. The reduction of computation cost is significant when $p \gg n$, since the singular value decomposition of $\mathbf{Y}\mathbf{Y}^{\top}$ is much faster.

Remark 4 (Sensitivity of ρ). Our empirical studies show that the rotate-and-solve procedure is not sensitive to ρ in a broad range. For a large range of choices of ρ , the classification errors are significantly improved over the existing LDA algorithms, as will be shown by our numerical experiments in the next section. Ideally, ρ can be estimated using data-adaptive methods such as cross-validation. However, cross-validation on ρ may be computationally intractable for high dimensional data where p is huge. As noted from Remark 3, we may use $\tilde{\mathbf{U}}_{\rho}$ to rotate the data which reduces the dimension from p to n. Thus cross-validation on ρ is more tractable using the modified rotate-and-solve procedure, and the classification quality can be noticeably improved as to be shown by our numerical experiments.

4 Numerical Studies

In this section, we compare the rotate-and-solve (RS) procedure with a number of popular LDA-based methods including standard LDA (1.3) (using Moore-Penrose pseudoinverse when $\hat{\Sigma}$ is singular), IR, nearest shrunken centroids (NSC) (Tibshirani et al., 2002), ROAD and LPD, via simulation and real data examples. For the RS procedure, two variants RS-ROAD and RS-LPD are included. For simulated examples using the toy models, we also consider the oracle RS methods (O-RS-ROAD and O-RS-LPD) where the oracle rotation shown in Section 2 are used to rotate the data. Moreover, the oracle Fisher's rule (1.2) is used as a benchmark method. In all RS-related methods, the parameter ρ is fixed to $\frac{1}{2}$ unless explicitly defined. The same number of observations are generated for both classes for all simulated data in Section 4.1, i.e. $n_1 = n_2$. All simulation settings have been repeated 100 times unless noted otherwise.

4.1 Simulated Data

4.1.1 Toy Models

We begin with several toy models with relatively small n and p to illustrate the performance of the RS procedures versus aforementioned LDA methods. We consider the following three toy models:

- Toy Model 1. $\Sigma = \mathbf{I}_p$; $\mu_1 = \mathbf{0}_p$ and $\mu_2 = a_1 \mathbf{1}_p$.
- Toy Model 2. $\Sigma = (\sigma_{i,j})$ with $\sigma_{i,i} = 1$ and $\sigma_{i,j} = 0.5$ for $i \neq j$; $\mu_1 = \mathbf{0}_p$ and $\mu_2 = (a_2 \mathbf{1}_{\ell}^{\top}, \mathbf{0}_{p-\ell}^{\top})^{\top}$, where $\ell = 5$.
- Toy Model 3. The setting is the same as 2 except $\ell = p/2$, $\boldsymbol{\mu}_2 = (a_3 \mathbf{1}_{\ell}^{\top}, \mathbf{0}_{p-\ell}^{\top})^{\top}$.

The values of a_1 , a_2 and a_3 in each of the toy models are chosen such that the expected classification errors of the oracle Fisher's rule (1.2) are 1%, 5% and 10%. For each model,

we take p = 50 and $n_1 = 20$ or 30. The same number of observations have been collected independently as the testing set.

We apply IR, Standard LDA, NSC, ROAD, LPD, RS-ROAD, RS-LPD, O-RS-ROAD and O-RS-LPD to 100 replicates of every simulation scenario. Simulation results are presented in Figure 1 (for $n_1 = 20$) and Figure 2 (for $n_1 = 30$). The Oracle rule always performs best and gives a benchmark for other methods. The O-RS methods perform very well and are comparable with oracle rule. For toy model 1, the features are independent, so IR performs best besides the oracle rule. But RS methods are comparable with IR. For model 2, the true β is nearly sparse. Therefore, ROAD and LPD perform well but RS methods still improve their performance. For model 3, neither the covariance matrix nor true β is sparse. RS methods work significantly better than their competitors. We observe that the RS methods are uniformly good in all the three models.



Figure 1: Simulation results for the three toy models with $n_1 = 20$ and p = 50. 1=ROAD, 2=RS-ROAD, 3=O-RS-ROAD, 4=LPD, 5=RS-LPD, 6=O-RS-LPD, 7=IR, 8=Standard LDA, 9=NSC, 10=Oracle.



Figure 2: Simulation results for the three toy models with $n_1 = 30$ and p = 50. 1=ROAD, 2=RS-ROAD, 3=O-RS-ROAD, 4=LPD, 5=RS-LPD, 6=O-RS-LPD, 7=IR, 8=Standard LDA, 9=NSC, 10=Oracle.

To see why RS methods outperform their direct sparse competitors, we plot the percentages of sum squares of the first several largest components of true $\boldsymbol{\beta}$ before and after rotation. For a rotation $R = \mathbf{U}_{\rho}$ or $\hat{\mathbf{U}}_{\rho}$, define $\boldsymbol{\beta}^{R} = R^{\top}\boldsymbol{\beta}$. Denote by $|\boldsymbol{\beta}|_{(1)}, \cdots, |\boldsymbol{\beta}|_{(p)}$ and $|\boldsymbol{\beta}^{R}|_{(1)}, \cdots, |\boldsymbol{\beta}^{R}|_{(p)}$ the reversed order statistics (from largest to smallest) of $\{|\boldsymbol{\beta}_{j}|\}_{j=1}^{p}$ and $\{|\boldsymbol{\beta}_{j}^{R}|\}_{j=1}^{p}$, respectively. For each setting, we plot $\sum_{i=1}^{k} |\boldsymbol{\beta}|_{(i)}^{2}/||\boldsymbol{\beta}||_{2}^{2}, \sum_{i=1}^{k} |\boldsymbol{\beta}^{\mathbf{U}_{\rho}}|_{(i)}^{2}/||\boldsymbol{\beta}^{\mathbf{U}_{\rho}}||_{2}^{2}$ and $\frac{1}{100} \sum_{j=1}^{100} \sum_{i=1}^{k} |\boldsymbol{\beta}^{\hat{\mathbf{U}}_{\rho j}}|_{(i)}^{2}/||\boldsymbol{\beta}^{\hat{\mathbf{U}}_{\rho j}}||_{2}^{2}$ for k = 1,..., p, where $\hat{\mathbf{U}}_{\rho j}$ is the rotation matrix for *j*th replicate and \mathbf{U}_{ρ} is the oracle rotation matrix. In Figure 3, we see, after rotation, $\boldsymbol{\beta}$ is more concentrated in its largest components. $\mathbf{U}_{\rho}^{\top}\boldsymbol{\beta}$ is extremely sparse, and $\hat{\mathbf{U}}_{\rho}^{\top}\boldsymbol{\beta}$ is sparser than the original $\boldsymbol{\beta}$. Obviously, ROAD/LDP is more efficient after the rotation.



Figure 3: Sparsity levels of β before (black) and after rotation using \mathbf{U}_{ρ} (blue) and \mathbf{U}_{ρ} (red). The top row corresponds to the case $n_1 = 20$, p = 50; while the bottom row corresponds to $n_1 = 30$, p = 50.

4.1.2 More Simulations

In our next numerical simulations, we consider the following three covariance structures:

Model 1: $\Sigma = (\sigma_{i,j})$ with $\sigma_{i,i} = 1$ and $\sigma_{i,j} = 0.5$ for $i \neq j$.

Model 2: $\Sigma = (\sigma_{i,j})$ with $\sigma_{i,j} = 0.7^{|i-j|}$.

Model 3: $\Sigma = \mathbf{I} + \mathbf{A}\mathbf{A}^{\top}$ where \mathbf{I} is the identity matrix and \mathbf{A} is $p \times 5$ matrix with entries generated independently from $\mathcal{N}(0, 1)$.

Without loss of generality, we set $\mu_1 = \mathbf{0}$ and $\mu_2 = (a\mathbf{1}_{p/2}^{\top}, \mathbf{0}_{p/2}^{\top})^{\top}$, where *a* is chosen specifically for each model such that the expected classification error of the oracle rule is 2%. Similar as before, for each simulation, we generate $2n_1$ independent observations for each class, where n_1 observations are used as training data and the other n_1 observations are used for testing. Results of Models 1-3 are presented in Figures 4-6 respectively, with various sample sizes and dimensionality. For Models 1 and 3 where the covariance structure is spiked, the improvement of the RS methods over the ROAD/LPD is remarkable. For Model 2 where the covariance structure is far from being spiked, the RS methods still generally outperform their counterparts.



Figure 4: Simulation results of Model 1: the boxplots from left to right correspond to the cases $(n_1, p) = (30, 200), (50, 200), (50, 400), (100, 200)$ and (100, 400).



Figure 5: Simulation results of Model 2: the boxplots from left to right correspond to the cases $(n_1, p) = (30, 200), (50, 200), (50, 400), (100, 200)$ and (100, 400).

In order to show that the improvement by applying RS is relatively general, we consider the following two scenarios with randomly generated covariance matrices



Figure 6: Simulation results of Model 3: the boxplots from upper-left to lower-right correspond to the cases $(n_1, p) = (30, 200), (50, 200), (50, 400), (100, 200)$ and (100, 400).

Random Model 1: $\Sigma = \left(\frac{M}{\|M\|}\right)^{\top} \left(\frac{M}{\|M\|}\right) + \operatorname{diag}(v)$ with each entry of $p \times p$ matrix M being generated independently from $\mathcal{N}(0,1)$ and v from $\mathcal{U}(0,1)$, where $\|M\|$ is the operator norm of M.

Random Model 2: $\Sigma = 4 \left(\frac{M}{\|M\|} \right)^{\top} \left(\frac{M}{\|M\|} \right)$ with each entry of M being generated independently from $\mathcal{N}(0, 1)$.

We fix $n_1 = 30$ and p = 300 and consider different sparsity levels of β with $||\beta||_0/p = 5\%, 10\%, \ldots, 95\%, 100\%$. We randomly generate β with a given sparsity level, whose nonzero entries are IID from $\mathcal{N}(0, 1)$. We then normalize β such that $\beta^{\top}\Sigma\beta = 12$. We fix $\mu_1 = 0$ and let $\mu_2 = -\Sigma\beta$. We repeat our data generation and classification 100 times for each scenario and record the average classification errors and their standard deviations.

We compare the results of ROAD and RS-ROAD which are shown in Figure 7. As one can see when β is very sparse, ROAD outperforms RS-ROAD as expected. However, the performance of ROAD highly depends on the sparsity level. On the other hand, RS-ROAD has significantly smaller overall error rates, and has the same qualitative behavior as the ORACLE. In particular, RS-ROAD is robust against the sparsity level of the true data generating procedure.



Figure 7: Average classification errors of "Random Model 1" (left) and "Random Model 2" (right) for $||\beta||_0/p = 5\%, 10\%, \ldots, 95\%, 100\%$, with ROAD in blue, RS-ROAD in red and ORACLE in green. Bars indicate the standard deviations of classification errors across 100 simulations.

4.2 Real Data: Leukemia and Lung Cancer

We now evaluate the performance of our proposed RS procedure on two popular gene expression data set: Leukemia (Golub et al., 1999) and lung cancer (Gordon et al., 2002). The two data sets come with separate training and testing sets of data vectors. The Leukemia data set contains p = 7129 genes with $n_1 = 27$ acute lymphoblastic leukemia (ALL) and $n_2 = 11$ acute myeloid leukemia (AML) vectors in the training set. The testing set includes 20 ALL and 14 AML vectors. The Lung Cancer data set contains p = 12533 genes with $n_1 = 16$ adenocarcinoma (ADCA) and $n_2 = 16$ mesothelioma training vectors. The testing set has 134 ADCA and 15 mesothelioma vectors. These two data sets can be downloaded from the web sites http://www.bioconductor.org/packages/release/data/experiment/html/golubEsets.html and http://www.chestsurg.org/publications/2002-microarray.aspx, respectively.

In our experiments, we put all the 47 (27 training + 20 testing data) ALL vectors and 25 (11 training + 14 testing data) AML vectors together and randomly select 23 ALL and 12 AML as training and the rest as testing. We repeat the experiments 20 times. We conduct a similar experiment on Lung cancer data by randomly select 75 ADCA and 15 mesothelioma data vector as training and the rest as testing, and repeat 20 times. The classification results of the aforementioned experiments using IR, NSC, ROAD and RS-ROAD are presented in Table 1, where RS-ROAD has the best overall performance.

Errors % (std %)	IR	NSC	ROAD	RS-ROAD
Leukemia	4.2708(2.9998)	8.5135(8.4232)	6.3514(5.9650)	4.4595(3.0721)
Lung Cancer	3.4669(1.4381)	$10.4396\ (7.2675)$	1.3736(1.0621)	$0.9341 \ (0.8931)$

Table 1: Classification errors for cancer data

4.3 Real Data: Shape Classifications

We also evaluate the performance of RS on shape classification, which is one of the most fundamental and important problems in computer vision and machine learning. All the shapes are represented by 2D binary images. We downloaded the MPEG-7 CE Shape-1 Part-B data set (Thakoor et al., 2007) from web site http://visionlab.uta.edu/shape_data.htm and selected a subset of it for our tests. Since the images in the dataset generally have different sizes, we resized them to the same size 50×50 (i.e., p = 2500) using the Matlab command imresize with bi-cubic interpolation. All the selected and resized shape images are shown in Figure 8.

There are 20 images for each shape class. After being loaded, each image is a matrix, with elements taking integer values in [0, 255]. In order to test the robustness of the classifiers, we also added Gaussian noise $N(0, 50^2)$ to all the selected images. For every pair of shapes, we randomly select 10 from each class as testing data and the rest as training data (i.e., $n_1 = n_2 = 10$). We repeat this 50 times for each of the shape pairs. The average classification errors by IR, NSC, ROAD and RS-ROAD are summarized in Table 2. We observe that RS-ROAD has the best overall performance, and it consistently improves ROAD in all scenarios.

Errors (%)	IR		NSC		ROAD		RS-ROAD	
Shape Pairs:	mean	std	mean	std	mean	std	mean	std
Apple & Bell	7.9	3.0	7.7	3.1	8.3	4.5	7.8	3.4
Pencil & Watch	19.2	6.1	20.4	7.1	18.2	6.9	16.0	7.1
Personal Car & Shoe	7.7	5.1	11.3	6.6	13.2	6.9	6.1	4.2
Camel & Elephant	8.8	6.6	12.1	9.1	20.3	11.0	6.9	4.0
Camel & Horse	9.6	7.1	11.8	9.5	22.0	10.6	7.7	5.6
Elephant & Horse	8.8	6.4	11.9	8.4	15.1	10.1	6.9	5.9

Table 2: Classification errors for shapes.



Figure 8: Selected shape images: resized to 50×50 with additive Gaussian noise.

4.4 Choice of ρ

Here we shall mainly discuss two issues related to the choice of ρ in Σ_{ρ}^{tot} : (1) the sensitivity of the classification results to the choices of ρ ; (2) data-adaptive selection of ρ by cross-validation.

4.4.1 Sensitivity to ρ

In the following simulations, we take the toy models 1-3 with a_i 's chosen such that the oracle error rate is 10%, and use the method RS-ROAD as an example. Let $\hat{\mathbf{U}}_{\rho}$ be the eigenvectors of $\hat{\boldsymbol{\Sigma}} + \rho \hat{\boldsymbol{\delta}} \hat{\boldsymbol{\delta}}^{\top}$ with various values of ρ . The average classification errors (among 100 replicates) of RS-ROAD with various ρ are shown in Figure 9, where the blue curves

show the errors associated to ρ and the red horizontal lines indicate the errors of ROAD. As we can see, the best choice of ρ depends on the scenario. Although it seems that choosing ρ optimally is a complicated issue, the plots in Figure 9 do indicate that for a large range of ρ , the classification results have significant improvements over a non-rotated classifier such as ROAD. This also indicates the robustness of the RS procedure to the choices of the parameter ρ . In general, any reasonable positive value of ρ should work well in most applications (Figure 9 shows the workable range of $\log \rho \in [-1, 10]$), if one does not have the resources or time to perform cross-validation.



Figure 9: Classification errors of RS-ROAD with various ρ (blue curves) v.s. ROAD (red lines). Plots in the first row correspond to the case $n_1 = n_2 = 30$ and plots in the second row correspond to $n_1 = 30$ and $n_2 = 45$. Columns 1-3 correspond to the Toy Models 1-3.

4.4.2 Cross-Validation choice of ρ

Cross-validation on ρ is computationally expensive when p is large. See Remark 4 for reduction of computation. When Σ has a (quasi)-spiked covariance structure, i.e. there are k eigenvalues that are significantly larger than the rest p - k eigenvalues, and if k is much less than the number of observations n, then we may use $\tilde{\mathbf{U}}_{\rho}$ to rotate the data instead of using $\hat{\mathbf{U}}_{\rho}$. Recall that $\tilde{\mathbf{U}}_{\rho}$ is the collections of the n eigenvectors of $\hat{\Sigma}^{tot}$ corresponding to the n largest eigenvalues. Then after rotating the data using $\tilde{\mathbf{U}}_{\rho}$, we reduce the dimension of the problem from p to n which will be significant reduction when $n \ll p$ (e.g. the real data considered in the previous two sections). We can also take $\tilde{\mathbf{U}}_{\rho}$ to be principal components, with dimensionality much less than n.

Our first simulations show that using $\tilde{\mathbf{U}}_{\rho}$ instead of $\hat{\mathbf{U}}_{\rho}$ does not hurt the classification error. We take the toy model 1-3 with a_i 's chosen such that the oracle error rate is 10%, and use the method RS-ROAD as an example. We set $n_1 = n_2 = 10$ (i.e. n = 20) and p = 50. The results are summarized in Table 3.

Errors % (std %)	Toy Model 1	Toy Model 2	Toy Model 3
Using $\hat{\mathbf{U}}_{\rho}$	24.9500 (11.3817)	26.8500(12.6861)	26.7000 (12.5171)
Using $\tilde{\mathbf{U}}_{ ho}$	25.0000 (11.5470)	26.5000(12.5831)	26.9500(12.5508)

Table 3: Classification errors and their standard deviations.

The previous simulation shows that we can reduce the size of the problem from p to n without sacrificing much of the classification quality. Since the computation cost can be greatly reduced in this way, cross-validation on ρ is now a computationally viable approach. In our next experiments, we take the data of Leukemia and Lung cancer in Section 4.2, and conduct a similar experiment as we did before, except that we use $\tilde{\mathbf{U}}_{\rho}$ and choose ρ using 5-folds cross-validation. The classification results are summarized in Table 4, where we also reproduce the results in Table 1 for comparison. We also presented therein the average values of ρ chosen by cross-validation along with their standard deviations. We repeat the same simulation to the shape data we presented in Section 4.3 and present comparisons and the estimated values of ρ in Table 5. As one can see that the choice of ρ is generally different for different type of data, and using cross-validation to select ρ , we can further reduce the classification errors.

Table 4: Classification errors and their standard deviations for Leukemia and Lung cancer.

Errors $\%$ (std $\%$)	Leukemia Cancer	Lung cancer
W/O Cross-Validation	4.4595(3.0721)	$0.9341 \ (0.8931)$
Cross-Validation	4.0541 (3.9702)	0.6593 (0.7479)
Estimated ρ	$0.2241 \ (0.2630)$	$0.0848 \ (0.0890)$

Acknowledgements

The authors are grateful to the Joint Editor, the Associate Editor and two referees for helpful comments, which have led to many improvements. Hao's research is supported by

Shape Pairs	W/O Cross-Validation	Cross-Validation	Estimated ρ	
	error mean (std)	error mean(std)	mean(std)	
Apple & Bell	7.8(3.4)	7.4(3.4)	$0.0806\ (0.0810)$	
Pencil & Watch	16.0(7.1)	15.0(6.5)	$0.2561 \ (0.2757)$	
Personal Car & Shoe	6.1 (4.2)	5.5(3.4)	0.0639(0.0808)	
Camel & Elephant	6.9(4.0)	7.1(4.1)	0.0723(0.0763)	
Camel & Horse	7.7(5.6)	6.0(6.1)	0.1196(0.1829)	
Elephant & Horse	6.9(5.9)	6.5(4.9)	$0.0667 \ (0.0725)$	

Table 5: Classification errors for shapes and the average values of ρ .

National Science Foundation grant DMS-1309507 and AMS-Simons Travel Grant. Fan's research is supported by the National Institute of General Medical Sciences of the National Institutes of Health through Grant Numbers R01-GM072611-9 and R01-GMR01GM100474 and National Science Foundation grants DMS-1206464 and DMS-1406266.

5 Appendix

The parameter ρ is a fixed positive constant through the appendix. So for easy presentation, we drop it from the notations Σ_{ρ}^{tot} , \mathbf{U}_{ρ} , etc.

Proof of Theorem 1: Let $a = \lambda_p > 0$ and $a_i = \lambda_i - \lambda_p > 0$. It then follows directly from Condition 1 and the singular value decomposition that

$$\boldsymbol{\Sigma} = a\mathbf{I} + \sum_{i=1}^{k} a_i \boldsymbol{\xi}_i \boldsymbol{\xi}_i^{\top}$$
(5.1)

and

$$\boldsymbol{\Sigma}^{tot} = a\mathbf{I} + \rho \boldsymbol{\delta} \boldsymbol{\delta}^{\top} + \sum_{i=1}^{k} a_i \boldsymbol{\xi}_i \boldsymbol{\xi}_i^{\top}.$$
 (5.2)

It can be shown that

$$(a\mathbf{I} + \sum_{i=1}^{k} a_i \boldsymbol{\xi}_i \boldsymbol{\xi}_i^{\top})^{-1} = a^{-1}\mathbf{I} - \sum_{i=1}^{k} \frac{a_i}{a(a+a_i)} \boldsymbol{\xi}_i \boldsymbol{\xi}_i^{\top}.$$
 (5.3)

This can be directly verified by

$$(a\mathbf{I} + \sum_{i=1}^{k} a_i \boldsymbol{\xi}_i \boldsymbol{\xi}_i^{\top})(a^{-1}\mathbf{I} - \sum_{i=1}^{k} \frac{a_i}{a(a+a_i)} \boldsymbol{\xi}_i \boldsymbol{\xi}_i^{\top})$$

= $\mathbf{I} + \sum_{i=1}^{k} a^{-1} a_i \boldsymbol{\xi}_i \boldsymbol{\xi}_i^{\top} - \sum_{i=1}^{k} \frac{a_i}{a+a_i} \boldsymbol{\xi}_i \boldsymbol{\xi}_i^{\top} - \sum_{i=1}^{k} \frac{a_i^2}{a(a+a_i)} \boldsymbol{\xi}_i \boldsymbol{\xi}_i^{\top}$
= \mathbf{I} ,

using the orthogonality

$$\boldsymbol{\xi}_i^{\top} \boldsymbol{\xi}_j = \begin{cases} 0, & i \neq j; \\ 1, & i = j. \end{cases}$$

By (5.3),

$$\boldsymbol{\beta} = \boldsymbol{\Sigma}^{-1} \boldsymbol{\delta} = a^{-1} \boldsymbol{\delta} - \sum_{i=1}^{k} \frac{a_i \boldsymbol{\xi}_i^{\top} \boldsymbol{\delta}}{a(a+a_i)} \boldsymbol{\xi}_i.$$
(5.4)

In other words, $\boldsymbol{\beta}$ is in the space spanned by $\{\boldsymbol{\delta}, \boldsymbol{\xi}_1, ..., \boldsymbol{\xi}_k\}$. On the other hand, by (5.2), it is easy to see that the space spanned by eigenvectors of $\boldsymbol{\Sigma}^{tot}$ corresponding to eigenvalues greater than a is exactly the space spanned by $\{\boldsymbol{\delta}, \boldsymbol{\xi}_1, ..., \boldsymbol{\xi}_k\}$. Therefore, $\boldsymbol{\beta}$ is perpendicular to the p - k - 1 dimensional eigenspace corresponding to eigenvalue a, i.e. $||\mathbf{U}^{\top}\boldsymbol{\beta}||_0 \leq k + 1$.

Before proving Theorem 2, we need a couple of results on the eigenvalues and eigenspaces of hermitian/symmetric matrices.

Lemma 1 (Weyl, 1912) If A and B are symmetric $p \times p$ matrices that differ by a matrix of rank at most r, then their eigenvalues (in descending order) $\{\alpha_j\}_{1 \leq j \leq p}$ and $\{\gamma_j\}_{1 \leq j \leq p}$ satisfy

$$\alpha_{j+r} \leq \gamma_j$$
 and $\gamma_{j+r} \leq \alpha_j$ for $1 \leq j, j+r \leq p$.

In particular, if r = 1 and $A \ge B$, it implies an interlacing property

$$\alpha_1 \ge \gamma_1 \ge \alpha_2 \ge \cdots \ge \alpha_p \ge \gamma_p$$

Lemma 2 (Davis & Kahan, 1970) Let A and B be symmetric matrices with A - B = Hand eigenvalues $\{\alpha_j\}_{1 \leq j \leq p}$ and $\{\gamma_j\}_{1 \leq j \leq p}$, respectively. If there exist a subset $S \subset \{1, ..., p\}$, an interval [s,t] and a positive constant z, such that $\alpha_j, \gamma_j \in [s,t]$ when $j \in S$ and $\alpha_j, \gamma_j \in$ $(-\infty, s - z] \cup [t + z, \infty)$ when $j \notin S$, then $||P - Q|| \leq ||H||/z$, where P and Q are projection matrices to the subspaces spanned by eigenvectors corresponding to $\{\alpha_j\}_{j \in S}$ and $\{\gamma_j\}_{j \in S}$, respectively.

The following lemmas are crucial in the proof of Theorem 2.

Lemma 3 Under Condition 2, if $\delta \in \mathbf{W}_1$, then the eigenvalues of Σ^{tot} satisfy

$$\eta_1 \ge \eta_2 \ge \dots \ge \eta_k \ge \eta_{k+1} + d > \eta_{k+1} \ge \dots \ge \eta_p; \tag{5.5}$$

otherwise,

$$\eta_1 \ge \eta_2 \ge \dots \ge \eta_{k+1} \ge \eta_{k+2} + d \frac{\rho || \delta_2 ||_2^2}{d + \rho || \delta ||_2^2} - \epsilon \ge \eta_{k+2} \ge \dots \ge \eta_p.$$
(5.6)

Proof of Lemma 3: Recall that $\{\lambda_j\}_{1 \le j \le p}$ are eigenvalues of Σ in descending order and $\boldsymbol{\xi}_j$ is the eigenvector corresponding to λ_j . \mathbf{W}_1 and \mathbf{W}_2 are linear spaces spanned by $\{\boldsymbol{\xi}_j\}_{1 \le j \le k}$ and $\{\boldsymbol{\xi}_j\}_{k+1 \le j \le p}$, respectively. $\boldsymbol{\delta} = \boldsymbol{\delta}_1 + \boldsymbol{\delta}_2$ with $\boldsymbol{\delta}_m \in \mathbf{W}_m$, m = 1, 2.

If $\boldsymbol{\delta} \in \mathbf{W}_1$, then $\boldsymbol{\delta} \perp \boldsymbol{\xi}_j$ for $k+1 \leq j \leq p$. Therefore, $\{\boldsymbol{\xi}_j\}_{k+1 \leq j \leq p}$ are eigenvectors of $\boldsymbol{\Sigma}^{tot} = \boldsymbol{\Sigma} + \rho \boldsymbol{\delta} \boldsymbol{\delta}^{\top}$ as well, and the corresponding eigenvalues satisfy $\eta_j = \lambda_j$ for $k+1 \leq j \leq p$. Moreover, by Lemma 1, $\eta_1 \geq \lambda_1 \geq \eta_2 \geq \cdots \geq \eta_k \geq \lambda_k$. Thus, Condition 2 implies

$$\eta_1 \ge \eta_2 \ge \cdots \ge \eta_k \ge \eta_{k+1} + d > \eta_{k+1} \ge \cdots \ge \eta_p.$$

If $\delta \notin \mathbf{W}_1$, i.e., $\delta_2 \neq 0$, define $\mathbf{W} = \mathbf{W}_1 \oplus \delta_2$. For all $\mathbf{w} \in \mathbf{W}$, with $||\mathbf{w}||_2 = 1$, we may

write $\mathbf{w} = \mathbf{w}_1 + \mathbf{w}_2$ where $\mathbf{w}_1 \in \mathbf{W}_1$ and $\mathbf{w}_2 = c\boldsymbol{\delta}_2 \in \mathbf{W}_2$. It follows that

$$\begin{split} \mathbf{w}^{\top} \mathbf{\Sigma}^{tot} \mathbf{w} &= \mathbf{w}^{\top} \mathbf{\Sigma} \mathbf{w} + \rho \mathbf{w}^{\top} \boldsymbol{\delta} \boldsymbol{\delta}^{\top} \mathbf{w} \\ &= \mathbf{w}_{1}^{\top} \mathbf{\Sigma} \mathbf{w}_{1} + \mathbf{w}_{2}^{\top} \mathbf{\Sigma} \mathbf{w}_{2} + \rho \left((\mathbf{w}_{1}^{\top} + \mathbf{w}_{2}^{\top}) (\boldsymbol{\delta}_{1} + \boldsymbol{\delta}_{2}) \right)^{2} \\ &= \mathbf{w}_{1}^{\top} \mathbf{\Sigma} \mathbf{w}_{1} + \mathbf{w}_{2}^{\top} \mathbf{\Sigma} \mathbf{w}_{2} + \rho \left(\mathbf{w}_{1}^{\top} \boldsymbol{\delta}_{1} + \mathbf{w}_{2}^{\top} \boldsymbol{\delta}_{2} \right)^{2} \\ &\geq \lambda_{k} || \mathbf{w}_{1} ||_{2}^{2} + \lambda_{p} || \mathbf{w}_{2} ||_{2}^{2} + \rho \left(\mathbf{w}_{1}^{\top} \boldsymbol{\delta}_{1} + \mathbf{w}_{2}^{\top} \boldsymbol{\delta}_{2} \right)^{2} \\ &\geq \lambda_{p} + d || \mathbf{w}_{1} ||_{2}^{2} + \rho \left(|\mathbf{w}_{1}^{\top} \boldsymbol{\delta}_{1} | - |\mathbf{w}_{2}^{\top} \boldsymbol{\delta}_{2} | \right)^{2} \end{split}$$

It is easy to see that

$$\inf \left\{ \rho \left(|\mathbf{w}_1^{\top} \boldsymbol{\delta}_1| - |\mathbf{w}_2^{\top} \boldsymbol{\delta}_2| \right)^2 \right\} = \left\{ \begin{array}{l} 0, & \text{if } ||\mathbf{w}_1||_2 \ge ||\boldsymbol{\delta}_2||_2/||\boldsymbol{\delta}||_2; \\ \rho (||\mathbf{w}_2||_2||\boldsymbol{\delta}_2||_2 - ||\mathbf{w}_1||_2||\boldsymbol{\delta}_1||_2)^2, & \text{if } ||\mathbf{w}_1||_2 < ||\boldsymbol{\delta}_2||_2/||\boldsymbol{\delta}||_2. \end{array} \right.$$

Therefore, if $||\mathbf{w}_1||_2 \ge ||\boldsymbol{\delta}_2||_2/||\boldsymbol{\delta}||_2$,

$$\mathbf{w}^{\top} \boldsymbol{\Sigma}^{tot} \mathbf{w} \ge \lambda_p + d||\mathbf{w}_1||_2^2 \ge \lambda_p + d \frac{||\boldsymbol{\delta}_2||_2^2}{||\boldsymbol{\delta}||_2^2};$$

 $\text{if } ||\mathbf{w}_1||_2 < ||\boldsymbol{\delta}_2||_2/||\boldsymbol{\delta}||_2,$

$$\mathbf{w}^{\top} \mathbf{\Sigma}^{tot} \mathbf{w} \geq \lambda_{p} + d||\mathbf{w}_{1}||_{2}^{2} + \rho(||\mathbf{w}_{2}||_{2}||\boldsymbol{\delta}_{2}||_{2} - ||\mathbf{w}_{1}||_{2}||\boldsymbol{\delta}_{1}||_{2})^{2}$$

$$\geq \lambda_{p} + d||\mathbf{w}_{1}||_{2}^{2} + \rho(||\boldsymbol{\delta}_{2}||_{2} - ||\mathbf{w}_{1}||_{2}||\boldsymbol{\delta}||_{2})^{2}$$

$$\geq \lambda_{p} + d\frac{\rho||\boldsymbol{\delta}_{2}||_{2}^{2}}{d + \rho||\boldsymbol{\delta}||_{2}^{2}}.$$

Overall, we have

$$\mathbf{w}^{\top} \mathbf{\Sigma}^{tot} \mathbf{w} \ge \lambda_p + \tilde{d} \quad \text{for all} \quad \mathbf{w} \in \mathbf{W},$$
(5.7)

where $\tilde{d} = d \frac{\rho || \boldsymbol{\delta}_2 ||_2^2}{d + \rho || \boldsymbol{\delta} ||_2^2}$. Since dim $\mathbf{W} = k + 1$, (5.7) implies that there are k + 1 eigenvalues that

are greater than $\lambda_p + \tilde{d}$ for Σ^{tot} . Together with Lemma 1, we conclude

$$\eta_1 \ge \eta_2 \ge \cdots \ge \eta_k \ge \eta_{k+1} \ge \lambda_p + d > \lambda_{k+1} \ge \eta_{k+2} \ge \cdots \ge \eta_p.$$

which leads to (5.6).

Similarly, we have

Lemma 4 Under Condition 1, if $\boldsymbol{\delta} \in \mathbf{W}_1$, then the eigenvalues of $\boldsymbol{\Sigma}^{tot}$ satisfy

$$\eta_1 \ge \eta_2 \ge \dots \ge \eta_k \ge \eta_{k+1} + d > \eta_{k+1} = \dots = \eta_p; \tag{5.8}$$

otherwise,

$$\eta_1 \ge \eta_2 \ge \dots \ge \eta_k \ge \eta_{k+1} \ge \eta_{k+2} + d \frac{\rho ||\boldsymbol{\delta}_2||_2^2}{d + \rho ||\boldsymbol{\delta}||_2^2} \ge \eta_{k+2} = \dots = \eta_p.$$
(5.9)

Proof of Lemma 4: The only difference is that the last p - k - 1 eigenvalues are equal, which is implies by Lemma 1 and the fact that $\lambda_{k+1} = \lambda_{k+2} = \cdots = \lambda_p$.

Proof of Theorem 2: Again, let $\boldsymbol{\xi}_j$ be the eigenvector of $\boldsymbol{\Sigma}$ corresponding to λ_j for $1 \leq j \leq p$. $a = \lambda_p$ and $a_j = \lambda_j - \lambda_p$.

Part I: $\delta \in \mathbf{W}_1$ implies $\delta \perp \boldsymbol{\xi}_j$ for $k < j \leq p$, so the eigenvectors $\{\boldsymbol{\xi}_j\}_{k < j \leq p}$ are also eigenvectors of $\boldsymbol{\Sigma}^{tot}$. Write $\mathbf{U} = (\mathbf{U}_1 \mathbf{U}_2)$ where \mathbf{U}_2 is submatrix of \mathbf{U} , consisting of right p - k columns. Then $\mathbf{U}_2 = (\boldsymbol{\xi}_{k+1}, \cdots, \boldsymbol{\xi}_p)$. Therefore,

$$\mathbf{U}_{2}^{\top} \boldsymbol{\beta} = \mathbf{U}_{2}^{\top} \boldsymbol{\Sigma}^{-1} \boldsymbol{\delta} = \mathbf{D}_{2}^{-1} \mathbf{U}_{2}^{\top} \boldsymbol{\delta} = \mathbf{0},$$

where $\mathbf{D}_2 = diag(\lambda_{k+1}, ..., \lambda_p).$

Part II: Under Condition 2, we can write $\Sigma = \Sigma_0 + \Delta$ where $\Sigma_0 = a\mathbf{I} + \sum_{j=1}^k a_j \boldsymbol{\xi}_j \boldsymbol{\xi}_j^{\top}$ and $\Delta = \sum_{j=k+1}^p a_j \boldsymbol{\xi}_j \boldsymbol{\xi}_j^{\top}$. Thus, Σ_0 satisfies Condition 1, and Δ is a semipositive matrix with maximal eigenvalue less than ϵ . Define

$$\boldsymbol{\Sigma}_0^{tot} = \boldsymbol{\Sigma}_0 +
ho \boldsymbol{\delta} \boldsymbol{\delta}^{ op}$$
 and $\boldsymbol{\Sigma}^{tot} = \boldsymbol{\Sigma} +
ho \boldsymbol{\delta} \boldsymbol{\delta}^{ op}.$

And let $\{\eta_{0j}\}_{1 \leq j \leq p}$ and $\{\eta_j\}_{1 \leq j \leq p}$ be their eigenvalues, in the descending order, respectively. Moreover, let **V** and **U** be orthogonal matrices such that

$$\mathbf{V}^{\top} \mathbf{\Sigma}_{0}^{tot} \mathbf{V} = \mathbf{D}_{0} \quad \text{and} \quad \mathbf{U}^{\top} \mathbf{\Sigma}^{tot} \mathbf{U} = \mathbf{D}_{0}$$

where $\mathbf{D}_0 = diag(\eta_{01}, ..., \eta_{0p})$ and $\mathbf{D} = diag(\eta_1, ..., \eta_p)$.

Here is the strategy of the proof. By Theorem 1, $\mathbf{V}^{\top} \mathbf{\Sigma}_{0}^{-1} \boldsymbol{\delta}$ is sparse so its ℓ_{1} -norm can be well controlled. Because of the results on the separated eigenvalues (Lemmas 3 and 4), we can show $\mathbf{U}^{\top} \boldsymbol{\beta}$ is similar to $\mathbf{V}^{\top} \mathbf{\Sigma}_{0}^{-1} \boldsymbol{\delta}$ using Lemma 2. Therefore, the ℓ_{1} -norm can be controlled as well.

Write $\mathbf{U} = (\mathbf{U}_1 \ \mathbf{U}_2)$ and $\mathbf{V} = (\mathbf{V}_1 \ \mathbf{V}_2)$ where \mathbf{U}_2 and \mathbf{V}_2 are submatrices of \mathbf{U} and \mathbf{V} respectively, consisting of right p - k - 1 columns. Note that

$$||\mathbf{U}^{\top}\boldsymbol{\beta}||_{1} = ||\mathbf{U}_{1}^{\top}\boldsymbol{\beta}||_{1} + ||\mathbf{U}_{2}^{\top}\boldsymbol{\beta}||_{1},$$

where $||\mathbf{U}_1^{\top}\boldsymbol{\beta}||_1 \leq \sqrt{||\mathbf{U}_1^{\top}\boldsymbol{\beta}||_0 \cdot ||\mathbf{U}_1^{\top}\boldsymbol{\beta}||_2^2} \leq \sqrt{k+1}||\boldsymbol{\beta}||_2$. So it is crucial to control $||\mathbf{U}_2^{\top}\boldsymbol{\beta}||_1$. From the proof of Theorem 1, we see that $\mathbf{V}_2^{\top}\boldsymbol{\Sigma}_0^{-1}\boldsymbol{\delta} = \mathbf{0}$. Hence

$$\begin{aligned} ||\mathbf{U}_{2}^{\top}\boldsymbol{\beta}||_{2} &= ||\mathbf{U}_{2}^{\top}\boldsymbol{\Sigma}^{-1}\boldsymbol{\delta} - \mathbf{U}_{2}^{\top}\boldsymbol{\Sigma}_{0}^{-1}\boldsymbol{\delta} + \mathbf{U}_{2}^{\top}\boldsymbol{\Sigma}_{0}^{-1}\boldsymbol{\delta}||_{2} \\ &\leq ||\mathbf{U}_{2}^{\top}\boldsymbol{\Sigma}^{-1}\boldsymbol{\delta} - \mathbf{U}_{2}^{\top}\boldsymbol{\Sigma}_{0}^{-1}\boldsymbol{\delta}||_{2} + ||\mathbf{U}_{2}^{\top}\boldsymbol{\Sigma}_{0}^{-1}\boldsymbol{\delta}||_{2} \\ &\leq ||\mathbf{U}_{2}^{\top}\boldsymbol{\Sigma}^{-1}\boldsymbol{\delta} - \mathbf{U}_{2}^{\top}\boldsymbol{\Sigma}_{0}^{-1}\boldsymbol{\delta}||_{2} + \sqrt{||\mathbf{U}_{2}^{\top}\boldsymbol{\Sigma}_{0}^{-1}\boldsymbol{\delta}||_{2}^{2} - ||\mathbf{V}_{2}^{\top}\boldsymbol{\Sigma}_{0}^{-1}\boldsymbol{\delta}||_{2}^{2}} \\ &\leq ||\boldsymbol{\Sigma}^{-1} - \boldsymbol{\Sigma}_{0}^{-1}|| \cdot ||\boldsymbol{\delta}_{2}||_{2} + \sqrt{\boldsymbol{\delta}^{\top}(\boldsymbol{\Sigma}_{0}^{-1})^{\top}\left(\mathbf{U}_{2}\mathbf{U}_{2}^{\top} - \mathbf{V}_{2}\mathbf{V}_{2}^{\top}\right)\boldsymbol{\Sigma}_{0}^{-1}\boldsymbol{\delta}} \\ &= S_{1} + S_{2} \end{aligned}$$

and

$$||\boldsymbol{\Sigma}^{-1} - \boldsymbol{\Sigma}_0^{-1}|| = \lambda_p^{-1} - \lambda_{k+1}^{-1} = a^{-1} - (a + a_{k+1})^{-1} = \frac{a_{k+1}}{a(a + a_{k+1})} \le \frac{\epsilon}{a^2}.$$

Thus, $S_1 \leq \frac{\epsilon}{a^2} ||\boldsymbol{\delta}_2||_2 \leq \frac{\epsilon(a+\epsilon)}{a^2} ||\boldsymbol{\beta}||_2$.

To control S_2 , we have to show that the spaces spanned by column vectors of \mathbf{V}_2 and \mathbf{U}_2 are close to each other. By Lemmas 3 and 4, we have

$$\eta_1 \ge \eta_2 \ge \cdots \ge \eta_k \ge \eta_{k+1} \ge \eta_{k+2} + d - \epsilon \ge \eta_{k+2} \ge \cdots \ge \eta_p,$$

$$\eta_{01} \ge \eta_{02} \ge \dots \ge \eta_{0k} \ge \eta_{0,k+1} \ge \eta_{0,k+2} + d \ge \eta_{0,k+2} = \dots = \eta_{0p},$$

where $\tilde{d} = d \frac{\rho ||\boldsymbol{\delta}_2||_2^2}{d+\rho ||\boldsymbol{\delta}||_2^2}$. Moreover, by Lemma 1, $\eta_{k+2} \leq \lambda_{k+1} \leq \lambda_p + \epsilon = a + \epsilon$, $\eta_{0,k+2} = \lambda_p = a$. On the other hand, $\eta_{k+1} \geq \eta_{k+2} + \tilde{d} - \epsilon \geq a + \tilde{d} - \epsilon$, $\eta_{0,k+1} \geq \eta_{0,k+2} + \tilde{d} = a + \tilde{d}$.

By Lemma 2, $||\mathbf{U}_2\mathbf{U}_2^\top - \mathbf{V}_2\mathbf{V}_2^\top|| \le ||\mathbf{\Delta}||/(\tilde{d} - 2\epsilon) \le \epsilon/(\tilde{d} - 2\epsilon).$

$$\begin{split} ||\boldsymbol{\Sigma}_{0}^{-1}\boldsymbol{\delta}||_{2} &= ||\boldsymbol{\Sigma}_{0}^{-1}\boldsymbol{\Sigma}\boldsymbol{\beta}||_{2} \leq ||\boldsymbol{\Sigma}_{0}^{-1}\boldsymbol{\Sigma}|| \cdot ||\boldsymbol{\beta}||_{2} \leq \frac{a+\epsilon}{a}||\boldsymbol{\beta}||_{2}. \text{ Thus, } S_{2} \leq \sqrt{\frac{\epsilon}{\tilde{d}-2\epsilon}}\frac{a+\epsilon}{a}||\boldsymbol{\beta}||_{2}.\\ \text{Therefore, } ||\mathbf{U}_{2}^{\top}\boldsymbol{\beta}||_{2} \leq \frac{a+\epsilon}{a}(\frac{\epsilon}{a}+\sqrt{\frac{\epsilon}{\tilde{d}-2\epsilon}})||\boldsymbol{\beta}||_{2}. \ ||\mathbf{U}_{2}^{\top}\boldsymbol{\beta}||_{1} \leq \sqrt{p-k-1}\frac{a+\epsilon}{a}(\frac{\epsilon}{a}+\sqrt{\frac{\epsilon}{\tilde{d}-2\epsilon}})||\boldsymbol{\beta}||_{2}.\\ \text{Finally, } ||\mathbf{U}^{\top}\boldsymbol{\beta}||_{1}/||\boldsymbol{\beta}||_{2} \leq \sqrt{k+1}+\sqrt{p-k-1}\frac{a+\epsilon}{a}(\frac{\epsilon}{a}+\sqrt{\frac{\epsilon}{\tilde{d}-2\epsilon}}). \end{split}$$

Proof of Theorem 3. Let \mathbf{V}_1 be a matrix whose columns vectors are the eigenvectors corresponding to the nonvanishing eigenvalues of the matrix $\mathbf{A} = \sum_{i=1}^k \lambda_i \boldsymbol{\xi}_i \boldsymbol{\xi}_i^\top + \rho \boldsymbol{\delta} \boldsymbol{\delta}^\top$. Recall $\lambda_i(\mathbf{B})$ be the *i*th largest eigenvalue of a symmetric matrix **B**. Then, by Lemma 2,

$$\|\mathbf{U}_1 - \mathbf{V}_1\| \le \frac{\|\mathbf{\Sigma}^{tot} - \mathbf{A}\|}{\lambda_{k+1}(\mathbf{A}) - \lambda_{k+2}(\mathbf{\Sigma}^{tot})} = \frac{\lambda_{k+1}}{\lambda_{k+1}(\mathbf{A}) - \lambda_{k+2}(\mathbf{\Sigma}^{tot})}.$$

By Lemma 1, $\lambda_{k+2}(\Sigma^{tot}) \leq \lambda_{k+1}$. Hence,

$$\|\mathbf{U}_1 - \mathbf{V}_1\| \le \frac{\lambda_{k+1}}{\lambda_{k+1}(\mathbf{A}) - \lambda_{k+1}} \le \frac{1}{a-1},$$

Let \mathbf{V}_2 be the eigenvectors that are orthogonal to \mathbf{V}_1 . Then, $\mathbf{V}_2^{\top}\boldsymbol{\delta} = 0$, since the columns of \mathbf{V}_1 are the linear combinations of $\boldsymbol{\delta}$ and $\{\boldsymbol{\xi}_i\}_{i=1}^k$. Consequently, $\|\mathbf{V}_1^{\top}\boldsymbol{\delta}\|_2 = \|\boldsymbol{\delta}\|_2$ and

$$\|\mathbf{U}_1^{\top}\boldsymbol{\delta}\|_2 = \|\mathbf{V}_1^{\top}\boldsymbol{\delta} + (\mathbf{U}_1 - \mathbf{V}_1)^{\top}\boldsymbol{\delta}\|_2 \ge \|\boldsymbol{\delta}\|_2 - \|\mathbf{U}_1 - \mathbf{V}_1\|\|\boldsymbol{\delta}\|_2 = \frac{a-2}{a-1}\|\boldsymbol{\delta}\|_2$$

The second conclusion follows directly from the fact that $\|\mathbf{U}_1^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{U}_1\| \leq \|\boldsymbol{\Sigma}\| = \lambda_1$.

References

AGARWAL, A., NEGAHBAN, S. & WAINWRIGHT, M. J. (2012). Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions. *The Annals of Statistics* 40, 1171–1197.

- BICKEL, P. & LEVINA, E. (2004). Some theory for fisher's linear discriminant function, naive bayes', and some alternatives when there are many more variables than observations. *Bernoulli* 10, 989–1010.
- CAI, T. & LIU, W. (2011). A direct estimation approach to sparse linear discriminant analysis. *Journal of the American Statistical Association* **106**, 1566–1577.
- DAVIS, C. & KAHAN, W. (1970). The rotation of eigenvectors by a perturbation. iii. SIAM Journal on Numerical Analysis 7, 1–46.
- FAN, J. & FAN, Y. (2008). High dimensional classification using features annealed independence rules. Annals of statistics 36, 2605–2637.
- FAN, J., FENG, Y. & TONG, X. (2012). A road to classification in high dimensional space. Journal of the Royal Statistical Society: Series B, 74, 745–771.
- FAN, J., LIAO, Y. & MINCHEVA, M. (2013). Large covariance estimation by thresholding principal orthogonal complements (with discussion). J. Roy. Statist. Soc. Ser. B.
- GOLUB, T., SLONIM, D., TAMAYO, P., HUARD, C., GAASENBEEK, M., MESIROV, J., COLLER, H., LOH, M., DOWNING, J., CALIGIURI, M. et al. (1999). Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science* 286, 531–537.
- GORDON, G., JENSEN, R., HSIAO, L., GULLANS, S., BLUMENSTOCK, J., RAMASWAMY, S., RICHARDS, W., SUGARBAKER, D. & BUENO, R. (2002). Translation of microarray data into clinically relevant cancer diagnostic tests using gene expression ratios in lung cancer and mesothelioma. *Cancer research* 62, 4963.
- HALL, P., JIN, J. & MILLER, H. (2009). Feature selection when there are many influential features. *Manuscript*.
- JOHNSTONE, I. (2001). On the distribution of the largest eigenvalue in principal components analysis. Ann. Statist **29**, 295–327.
- JOHNSTONE, I. & LU, A. (2009). On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association* **104**, 682–693.
- KAROUI, N. (2008). Operator norm consistent estimation of large-dimensional sparse covariance matrices. The Annals of Statistics, 2717–2756.

- MAI, Q., ZOU, H. & YUAN, M. (2012). A direct approach to sparse discriminant analysis in ultra-high dimensionals. *Biometrika* **99**, pp. 29–42.
- SHEN, D., SHEN, H., ZHU, H. & MARRON, J. (2013). Surprising asymptotic conical structure in critical sample eigen-directions. arXiv preprint arXiv:1303.6171.
- THAKOOR, N., GAO, J. & JUNG, S. (2007). Hidden markov model-based weighted likelihood discriminant for 2-d shape classification. *Image Processing, IEEE Transactions on* 16, 2707–2719.
- TIBSHIRANI, R., HASTIE, T., NARASIMHAN, B. & CHU, G. (2002). Diagnosis of multiple cancer types by shrunken centroids of gene expression. *Proceedings of the National Academy of Sciences* 99, 6567–6572.
- TRENDAFILOV, N. T. & JOLLIFFE, I. T. (2007). Dalass: Variable selection in discriminant analysis via the lasso. *Computational Statistics & Data Analysis* **51**, 3718 3736.
- WEYL, H. (1912). Das asymtotische verteilungsgesetz der eigenwerte lineare partieller differentialgleichungen. Math. Ann. **71**, 441–479.
- WU, M. C., ZHANG, L., WANG, Z., CHRISTIANI, D. C. & LIN, X. (2009). Sparse linear discriminant analysis for simultaneous testing for the significance of a gene set/pathway and gene selection. *Bioinformatics* 25, 1145–1151.
- ZOU, H., HASTIE, T. & TIBSHIRANI, R. (2006). Sparse principal component analysis. Journal of computational and graphical statistics 15, 265–286.