

Statistica Sinica Preprint No: SS-2015-0387R2

Title	A New Reduced-Rank Linear Discriminant Analysis Method and Its Applications
Manuscript ID	SS-2015-0387R2
URL	http://www.stat.sinica.edu.tw/statistica/
DOI	10.5705/ss.202015.0387
Complete List of Authors	Yue Selena Niu Ning Hao and Bin Dong
Corresponding Author	Ning Hao
E-mail	nhao@math.arizona.edu
Notice: Accepted version subject to English editing.	

A New Reduced-Rank Linear Discriminant Analysis Method and Its Applications

Yue Selena Niu, Ning Hao, and Bin Dong
University of Arizona and Peking University

December 2, 2016

Abstract

We consider multi-class classification problems for high dimensional data. Following the idea of reduced-rank linear discriminant analysis (LDA), we introduce a new dimension reduction tool with a flavor of supervised principal component analysis (PCA). The proposed method is computationally efficient and can incorporate the correlation structure among the features. Besides the theoretical insights, we show that our method is a competitive classification tool by simulated and real data examples.

Keywords: Dimension reduction, Gene expression data, High dimensional data, Multi-class classification, Supervised principal component analysis.

1 Introduction

Targeting on cancer classification and other modern applications, a great number of high dimensional classification techniques have been invented and studied recently; see Hastie et al. (2009) for an extensive introduction, and Witten & Tibshirani (2011); Cai & Liu (2011); Fan et al. (2012); Mai et al. (2012) for some recent developments. Although these

contemporary classification tools can be applied to the high dimensional data, most of them rely on some strong assumptions. For example, many methods assume that the features are independent to each other; some other methods assume various sparsity conditions. On one hand, these assumptions make the model simple and robust against growing dimensionality, so classification accuracy and computational efficiency can be achieved. On the other hand, they may be too restrictive, and when violated, lead to information loss in data analysis. Moreover, lots of methods target on the binary classification case and are not straightforward to use if more than two classes are present. Convenient and efficient classification tools for multi-class data are quite limited. Therefore, it is desirable to develop new classification techniques that can handle high dimensional, multi-class data and also take into account the correlation among the features.

Recall that many linear classification rules depend on the Mahalanobis distance. However, it cannot be well-estimated for high dimensional data when the number of features is greater than the sample size as the sample covariance is singular. Nevertheless, under the assumption that features are independent, the sample covariance matrix is diagonal and strictly positive definite, so the Mahalanobis distance can be calculated. That is one of the main reasons that the independence assumption is crucial in many popular classification methods. For example, the nearest shrunken centroids (NSC, Tibshirani et al., 2002), independence rule (IR, Bickel & Levina, 2004), features annealed independence rule (FAIR, Fan & Fan, 2008) all assume that the features are independent to each other. Moreover, some other methods such as regularized discriminant analysis (RDA, Guo et al., 2007) use a covariance estimator which is the sample covariance regularized towards a diagonal matrix. Recently, many new classification tools have been developed including penalized linear discriminant analysis (PLDA, Witten & Tibshirani, 2011), linear programming discriminant rule (LPD, Cai & Liu, 2011), regularized optimal affine discriminant rule (ROAD, Fan et al., 2012), direct sparse discriminant analysis (DSDA, Mai et al., 2012), sparse discriminant analysis (SDA, Clemmensen et al., 2012), multi-class sparse discriminant analysis (MSDA, Mai et al., 2015). Roughly speaking, these sparse methods obtain sparse models by solving penalized or constrained optimization problems, and their efficiency relies on the sparsity level of the normal vectors to the optimal discriminant boundaries.

Reduced-rank LDA is a classical approach to classification. It conducts dimension reduction by projecting the data to the centroid-spanning space and classifies the data based on nearest centroid. Another commonly used dimension reduction tool is PCA, which projects

the data to the space spanned by the top principal components of the total sample covariance matrix. Note that reduced-rank LDA makes use of the label information (through centroids) but ignores the (within class) covariance information. On the other hand, PCA relies on the covariance information only and is mainly regarded as an unsupervised learning tool in the literature.

In this paper, we propose a new reduced-rank LDA method combining the advantages of the classical reduced-rank LDA and PCA. The principal components of a weighted sum of the sample within class and between class covariance matrices are used for dimension reduction, and standard LDA is employed to the projected data for classification. In this dimension reduction process, both label and covariance information can be taken into account, through between class and within class covariance, respectively. Therefore, we regard it as a version of supervised PCA. This new method does not rely on the aforementioned sparsity or independence assumptions and offers an alternative classification tool for various applications, especially, when these technical assumptions are not satisfied.

To shed theoretical insight of our method, we consider the spiked structure of the covariance (Johnstone, 2001). Roughly speaking, a symmetric positive definite matrix is called spiked if all of its eigenvalues are equal except a few large ones. In other words, it is a sum of a scalar matrix and a low rank matrix. Intuitively, the spiked structure might be a better model than the diagonal one to approximate the true covariance, as it can take into account strong correlation among the features, which is not uncommon in applications. We will explain why the spiked structure can help for dimensional reduction. We should remark that, although the spiked structure plays a key role in the theoretical analysis, our proposed procedure does not rely on the spiked structure explicitly and can be applied to general data sets.

The main contribution of this paper is threefold. First, it proposes a novel dimension reduction and classification tool. The new method incorporates covariance among features and works well for high dimensional multi-class data. Second, it illustrates a new supervised way to conduct PCA. This technique is generally applicable to both classification and regression models. Last but not least, the proposed method is computationally efficient and can be applied directly to the real data such as gene expression data in cancer research.

The rest of the paper is organized as follows. Section 2 introduces notations and gives a

selective review on some popular linear classification methods from an aspect of dimension reduction. In Section 3, we study a new reduced-rank LDA method for classification and offer some theoretical insights. Numerical studies and real data applications are illustrated in Section 4, which is followed by a short discussion in Section 5. All technical proofs are given in the supplementary material.

2 Linear methods for classification

2.1 Notations

In this section, we review several popular classification tools from the aspect of dimension reduction. We consider a standard setup. Let $\mathcal{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)^\top$ be a $n \times p$ matrix with n observations $\mathbf{X}_1, \dots, \mathbf{X}_n$, each of which is a p dimensional vector. Let $\mathbf{Y} = (Y_1, \dots, Y_n)^\top$ be a response vector with $Y_i \in \{1, \dots, K\}$, $1 \leq i \leq n$. That is, \mathbf{X}_i belongs to group k if $Y_i = k$. Define the index set of group k as $C_k = \{i : Y_i = k\}$ and its cardinality $n_k = |C_k|$, where $1 \leq k \leq K$. The goal of classification is to establish a classification rule and label a new observation \mathbf{X}^* based on the training data.

In the literature, the Gaussian assumption is often used to facilitate statistical analysis of various methods. In the simplest setting, it is assumed that data from all groups share a common covariance matrix Σ_w , that is, $(\mathbf{X}|Y = k) \sim \mathcal{N}(\boldsymbol{\mu}_k, \Sigma_w)$, $1 \leq k \leq K$. For easy presentation, we also assume that the prior probabilities $\pi_k = \mathbf{P}(Y = k)$ are equal for all k . In practice, the prior probability can be always estimated and taken into account for most methods considered in this paper.

Recall that, for a specified strictly positive definite symmetric matrix \mathbf{S} , the Mahalanobis distance between two vectors \mathbf{u} and \mathbf{v} is

$$d_M(\mathbf{u}, \mathbf{v}) = \{(\mathbf{u} - \mathbf{v})^\top \mathbf{S}^{-1}(\mathbf{u} - \mathbf{v})\}^{\frac{1}{2}}.$$

In fact, under the Gaussian assumption, the optimal classification rule minimizing the expected classification error is called the Bayes rule, which simply classifies a data point to a

group with the nearest centroid in terms of Mahalanobis distance with $\mathbf{S} = \Sigma_w$. That is,

$$Y = \operatorname{argmin}_{1 \leq k \leq K} (\mathbf{X} - \boldsymbol{\mu}_k)^\top \Sigma_w^{-1} (\mathbf{X} - \boldsymbol{\mu}_k) = \operatorname{argmin}_{1 \leq k \leq K} \left\| \Sigma_w^{-1/2} (\mathbf{X} - \boldsymbol{\mu}_k) \right\|_2^2. \quad (2.1)$$

A key observation from (2.1) is that if we rotate the sample space by $\Sigma_w^{-1/2}$ first, then the Bayes rule is equivalent to a nearest centroid classifier with standard Euclidian distance. Moreover, it follows (2.1) that the optimal decision boundary separating groups k and ℓ is an affine space defined by $\{\mathbf{X} - (\boldsymbol{\mu}_k + \boldsymbol{\mu}_\ell)/2\} \Sigma_w^{-1} (\boldsymbol{\mu}_k - \boldsymbol{\mu}_\ell) = 0$. Note that the normal vector to this affine space is $\Sigma_w^{-1} (\boldsymbol{\mu}_k - \boldsymbol{\mu}_\ell)$. Therefore, the decision boundary of the Bayes rule to the whole classification problem is a subset of the union of these affine spaces, whose normal vectors span a vector space $\Sigma_w^{-1} \mathbf{C}$, where \mathbf{C} is the vector space spanned by $\{\boldsymbol{\mu}_k - \boldsymbol{\mu}_\ell\}_{1 \leq k < \ell \leq K}$. Note that $\dim \mathbf{C} = \dim(\operatorname{span}\{\boldsymbol{\mu}_k - \boldsymbol{\mu}_\ell\}_{k=1}^{K-1}) \leq K - 1$, and the equal sign holds when the set of centroids $\{\boldsymbol{\mu}_k\}_{k=1}^K$ is in linear general position. By Lemma 2 in the supplementary material, when p is larger than K , we will lose no information to project the data from \mathbb{R}^p to a small subspace $\Sigma_w^{-1} \mathbf{C}$ for classification. That is, it is equivalent to apply the Bayes rule to the projected data and the original data. In practice, when p is large, it is extremely helpful to find a reasonable approximation subspace to $\Sigma_w^{-1} \mathbf{C}$ to reduce the dimensionality.

Let us introduce some notations of important statistics. Without loss of generality, we assume that the features, i.e., columns of \mathcal{X} , are centered to have mean zero since all methods considered in this paper are translation invariant. The within-class sample covariance matrix is

$$\mathbf{W} = \frac{1}{n} \sum_{k=1}^K \sum_{i \in C_k} (\mathbf{X}_i - \hat{\boldsymbol{\mu}}_k)(\mathbf{X}_i - \hat{\boldsymbol{\mu}}_k)^\top,$$

where $\hat{\boldsymbol{\mu}}_k = n_k^{-1} \sum_{i \in C_k} \mathbf{X}_i$. The between-class sample covariance matrix is

$$\mathbf{B} = \frac{1}{n} \sum_{k=1}^K n_k (\hat{\boldsymbol{\mu}}_k - \hat{\boldsymbol{\mu}})(\hat{\boldsymbol{\mu}}_k - \hat{\boldsymbol{\mu}})^\top = \frac{1}{n} \sum_{k=1}^K n_k \hat{\boldsymbol{\mu}}_k \hat{\boldsymbol{\mu}}_k^\top,$$

where $\hat{\boldsymbol{\mu}} = n^{-1} \sum_{k=1}^K n_k \hat{\boldsymbol{\mu}}_k = 0$. The total sample covariance matrix

$$\mathbf{T} = n^{-1} \mathcal{X}^\top \mathcal{X} = \mathbf{W} + \mathbf{B}.$$

2.2 Simple reduced-rank linear discriminant analysis

A simple reduced-rank LDA (Hastie et al., 2009) projects the data to the centroid-spanning subspace $\hat{\mathbf{C}} = \text{span}\{\hat{\boldsymbol{\mu}}_k - \hat{\boldsymbol{\mu}}_\ell\}_{1 \leq k < \ell \leq K}$. The idea is that, when calculating the (Euclidean) distances to find the closest centroid, one can ignore the distances orthogonal to $\hat{\mathbf{C}}$ which contribute equally to all groups. This simple method reduces the dimensionality remarkably. The main drawback is that it does not incorporate the covariance structure. We may lose much information if $\boldsymbol{\Sigma}_w$ is far from a scalar matrix.

2.3 Fisher's approach and the standard LDA

Fisher's approach is to find a subspace so the projected centroids are spread out as much as possible with respect to the covariance. It finds the first direction by solving

$$\mathbf{v}_1 = \underset{\mathbf{v}}{\operatorname{argmax}} \frac{\mathbf{v}^\top \mathbf{B} \mathbf{v}}{\mathbf{v}^\top \mathbf{W} \mathbf{v}}, \quad (2.2)$$

which is equivalent to

$$\mathbf{v}_1 = \underset{\mathbf{v}}{\operatorname{argmax}} \mathbf{v}^\top \mathbf{B} \mathbf{v} \text{ subject to } \mathbf{v}^\top \mathbf{W} \mathbf{v} = 1, \quad (2.3)$$

provided \mathbf{W} is not singular. It is well known that when $K = 2$, \mathbf{v}_1 is the same as the normal vector (up to a scalar) of the decision boundary separating two groups obtained by standard LDA. When $K > 2$, one can continue to solve this generalized eigenvalue problem until step $\text{rank}(\mathbf{B})$, as follows.

$$\begin{aligned} \mathbf{v}_2 &= \underset{\mathbf{v}}{\operatorname{argmax}} \mathbf{v}^\top \mathbf{B} \mathbf{v} \text{ subject to } \mathbf{v}^\top \mathbf{W} \mathbf{v} = 1; \mathbf{v}^\top \mathbf{W} \mathbf{v}_1 = 0, \\ &\dots \\ \mathbf{v}_k &= \underset{\mathbf{v}}{\operatorname{argmax}} \mathbf{v}^\top \mathbf{B} \mathbf{v} \text{ subject to } \mathbf{v}^\top \mathbf{W} \mathbf{v} = 1; \mathbf{v}^\top \mathbf{W} \mathbf{v}_\ell = 0, \ell = 1, 2, \dots, k-1, \\ &\dots \end{aligned} \quad (2.4)$$

Obviously, the covariance plays a role here through the sample pooled covariance \mathbf{W} . Moreover, there is an order for these directions so the dimension of the subspace can be pre-

specified or chosen data-adaptively. In fact, the simple reduced-rank LDA is a special case by restricting \mathbf{W} to a scalar matrix. In this sense, Fisher's approach is also a reduced-rank LDA method.

The standard LDA can be viewed as a dimension reduction technique as well. Roughly speaking, it mimics the Bayes rule by plugging in estimators of the common covariance and centroids. It labels an observation \mathbf{X} by $\hat{Y} = \operatorname{argmin}_{1 \leq k \leq K} \|\mathbf{W}^{-1/2}(\mathbf{X} - \hat{\boldsymbol{\mu}}_k)\|_2$. Similar to the analysis of the Bayes rule, the normal vectors of the decision boundaries of standard LDA span a subspace $\mathbf{W}^{-1}\hat{\mathbf{C}} \subset \mathbb{R}^p$. It is equivalent to apply standard LDA to, instead of the original data, the projected data onto subspace $\mathbf{W}^{-1}\hat{\mathbf{C}}$. The following proposition indicates the equivalence between Fisher's approach and the standard LDA in the context of dimension reduction. It seems that this well-known result is not established formally anywhere so we include it here.

Proposition 1 If \mathbf{W} is nonsingular, then $\dim \mathbf{W}^{-1}\hat{\mathbf{C}} = \operatorname{rank}(\mathbf{B})$, and $\mathbf{W}^{-1}\hat{\mathbf{C}} = \operatorname{span}\{\mathbf{v}_k\}_{k=1}^r$ where $r = \dim \hat{\mathbf{C}}$, \mathbf{v}_k is defined in (2.4).

The original idea of the reduced-rank LDA is to find and project the data to a small subspace with high discriminant power before applying LDA. We see that, as a by product of the standard LDA, it gives automatically such a subspace, i.e., $\mathbf{W}^{-1}\hat{\mathbf{C}}$. A quick remark here is that Fisher's approach or the standard LDA performs well only when the sample size is large enough so $\mathbf{W}^{-1}\hat{\mathbf{C}}$ is a good approximation to $\boldsymbol{\Sigma}_w^{-1}\mathbf{C}$.

2.4 The independence rule and related approaches

It is well known that LDA does not work well when $p \sim n$ and $p > n$. In the context of dimension reduction reviewed in Section 2.3, the reason is that $\mathbf{W}^{-1}\hat{\mathbf{C}}$ or $\mathbf{W}^-\hat{\mathbf{C}}$ is no longer a good approximation to $\boldsymbol{\Sigma}_w^{-1}\mathbf{C}$ for high dimension data, where \mathbf{W}^- is a pseudo-inverse of \mathbf{W} . The main problem is that the sample pooled covariance \mathbf{W} is singular and a poor estimate for $\boldsymbol{\Sigma}_w$. One remedy is to assume that the features are independent, i.e., $\boldsymbol{\Sigma}_w$ is diagonal. This leads to the independence rule or diagonal LDA. To apply diagonal LDA, one just uses the diagonal part $\hat{\mathbf{D}}_w = \operatorname{diag}(\mathbf{W})$ instead of \mathbf{W} in the standard LDA. Although the features are rarely independent in applications, the IR or diagonal LDA usually outplays standard LDA when $p > n$ (Bickel & Levina, 2004).

By the same spirit of Proposition 1, one can see the equivalence between the IR and Fisher's approach with \mathbf{W} replaced by $\hat{\mathbf{D}}_w$ in (2.4), as stated in Corollary 1 in the supplementary material. Witten & Tibshirani (2011) further imposed some sparse assumptions on \mathbf{v}_k 's to derive a penalized LDA (PLDA). One can also conduct dimension reduction based on the rank of marginal discriminant power. Two well-known approaches are the NSC and FAIR, studied by Tibshirani et al. (2002) and Fan & Fan (2008), respectively.

2.5 Principal component analysis

As a dominant approach to dimension reduction in the context of unsupervised learning, PCA has been also used to solve supervised learning problems, e.g., principal component regression (Jolliffe, 2002), supervised PCA (Bair et al., 2006), etc. Roughly speaking, PCA extracts the mutually orthogonal directions with largest variances of the data and discards the rest. In our context, standard PCA would ignore the label information and keep the eigenvectors corresponding to q top eigenvalues of \mathbf{T} , where q can be pre-specified or chosen data adaptively. People have been concerned with this approach as there is no guarantee that the top principal components have good discriminant power. Bair et al. (2006) proposed a variant of supervised PCA. It is a two-stage procedure in which marginal statistics are used to reduce dimension before applying standard PCA. It seems that the label information is used only in the first stage. It is interesting to find alternative ways to conduct supervised PCA.

2.6 Summary

Reduced-rank LDA can be viewed as a general strategy which finds and projects the data into a linear subspace before applying standard LDA. Using Fisher's framework (2.4) with different within-class covariance estimates, we unified simple reduced-rank LDA, standard LDA and IR as special cases. One main drawback of these methods is that none of them can compromise degeneracy of the sample covariance \mathbf{W} and non-triviality of the true covariance Σ_w . PCA would be a potentially good approach which requires neither nondegeneracy of \mathbf{W} nor independence assumption. Next, we will show a new reduced-rank LDA method which uses PCA in a supervised way.

3 A New Reduced-Rank Linear Discriminant Analysis Method

3.1 Method

To take advantage of existing methods described in the last section and study the multi-class classification problem in a unified manner, we consider $\mathbf{T}_\gamma = \mathbf{W} + \gamma\mathbf{B}$ with $\gamma > 0$. Recall that \mathbf{W} and \mathbf{B} are within class and between class sample covariance matrices, respectively. Here γ is a tuning parameter. As we will see, if $\gamma = 1$, our proposed procedure is equivalent to the standard PCA; and if $\gamma \rightarrow \infty$, the procedure is equivalent to the simple reduced-rank LDA illustrated in Section 2.2.

Consider eigenvalue decomposition

$$\mathbf{U}_\gamma^\top \mathbf{T}_\gamma \mathbf{U}_\gamma = \mathbf{D}_\gamma \quad (3.1)$$

where \mathbf{D}_γ is a diagonal matrix with diagonal entries ranked in a descending order and \mathbf{U}_γ is an orthogonal matrix. We propose a new reduced-rank LDA procedure based on the first q principal components of \mathbf{T}_γ . It can be conducted as follows.

1. Calculate \mathbf{T}_γ and \mathbf{U}_γ and project the data from \mathbb{R}^p to the linear subspace spanned by the first q columns of \mathbf{U}_γ . In practice, the parameters γ and q can be chosen data adaptively.
2. Apply the standard LDA to the projected data.

This procedure generalized classical methods such as PCA and simple reduced-rank LDA by allowing a varying parameter γ . The label information is taken into account through γ in dimension reduction. In this sense, our procedure is a version of supervised PCA, so we call it supervised PCA-based LDA (SPCALDA). Although the parameter γ plays a minor role in the later theoretical analysis, it makes our procedure more flexible. For example, the qualities of \mathbf{W} and \mathbf{B} to approximate their population counterparts are usually not equally good. Therefore, γ can serve as a weight to balance them.

3.2 Theoretical insights

It is not uncommon that PCA is used to solve supervised learning problems, especially, in the analysis of high dimensional or highly correlated data. However, it is usually used as an ad-hoc method without any theoretical justification. We now illustrate some theoretical insights why our procedure (including PCA as a special case) works well under certain conditions.

To understand the proposed method, we consider a population version. Denote by Σ_b and Σ_t the population versions of between-class and total covariance matrix, respectively. Define $\Sigma_\gamma = \Sigma_w + \gamma\Sigma_b$, $\gamma > 0$ with eigenvalue decomposition

$$\mathbf{U}_O^\top \Sigma_\gamma \mathbf{U}_O = \mathbf{D}_O,$$

where \mathbf{D}_O is a diagonal matrix with diagonal entries ranked in a descending order and \mathbf{U}_O is orthogonal. Because γ plays only a minor role in this section, we drop it from notations \mathbf{U}_O , \mathbf{D}_O , etc. For an oracle procedure assuming parameters are known, we can project the data to the linear subspace spanned by top principal components of \mathbf{U}_O . A natural question is when we can project the data by oracle procedure without information loss.

Let $\{\lambda_j\}_{j=1}^p$ be eigenvalues of Σ_w in a descending order. We consider a spiked covariance structure Johnstone (2001).

Spiked Condition: Assume that $\lambda_1 \geq \dots \geq \lambda_s > \lambda_{s+1} = \dots = \lambda_p$ for some integer $s < p$.

Theorem 1 *If $p > s + K - 1$, $s > 1$, write $\mathbf{U}_O = (\mathbf{U}_{O1}, \mathbf{U}_{O2})$ where \mathbf{U}_{O1} and \mathbf{U}_{O2} are $p \times (s + K - 1)$ and $p \times (p - s - K + 1)$ matrices respectively. Under spiked condition, we have*

$$\mathbf{U}_{O2}^\top \Sigma_w^{-1} (\boldsymbol{\mu}_k - \boldsymbol{\mu}_\ell) = 0, \text{ for all } 1 \leq k < \ell \leq K.$$

In other words, $\Sigma_w^{-1} \mathbf{C}$ are located in the subspace spanned by columns of \mathbf{U}_{O1} .

In short, we will lose no discriminant power to project the data to a $s + K - 1$ dimensional subspace spanned by columns of \mathbf{U}_{O1} .

We may generalize this result further as follows. Without loss of generality, assume $\boldsymbol{\mu} = \sum_{k=1}^K \pi_k \boldsymbol{\mu}_k = 0$. From the proof of Theorem 1 in the supplementary material, we see

the conclusion of Theorem 1 holds for principal components of $\Sigma_{\rho} = \Sigma_w + \sum_{k=1}^K \rho_k \boldsymbol{\mu}_k \boldsymbol{\mu}_k^{\top}$, where $\boldsymbol{\rho} = (\rho_1, \dots, \rho_K)^{\top}$ with $\rho_k > 0$, $1 \leq k \leq K$. Nevertheless, when K is more than three, it is technically more complicated to tune K parameters.

A more general model than the Gaussian one considered here is the mixture Gaussian model which allows each group to be distributed as mixture Gaussian with the same covariance; see e.g. Hastie et al. (2009), Section 12.7. It is useful in applications when the groups are inhomogeneous. Let $(\mathbf{X}|Y = k) \sim \sum_{t=1}^{R_k} \pi_{kt} \mathcal{N}(\boldsymbol{\mu}_{kt}, \Sigma_w)$, where $1 \leq k \leq K$, $1 \leq t \leq R_k$, $\sum_{t=1}^{R_k} \pi_{kt} = 1$. That is, there are R_k prototypes for the group k . We can generalize Theorem 1 to this case.

Theorem 2 Let $R = \sum_{k=1}^K R_k$. Then Theorem 1 still holds if we replace K by R everywhere.

Remark 1. The spiked condition is crucial in Theorems 1 and 2. It is a reasonable model to approximate the true covariance when strong correlation among the features is present. It is also employed by Hao et al. (2015), which aimed to sparsify the normal vector of optimal discriminant boundary for binary classification problems. In spite of the theoretical insight, the spiked condition plays no role in conducting our procedure. In real applications, it may not hold exactly. Nevertheless, our numerical studies show that our procedure performs very well.

Remark 2. In practice, we will work with \mathbf{U} instead of its population version \mathbf{U}_O . Although \mathbf{U}_O might be very different from \mathbf{U} when $n \ll p$, the left part of \mathbf{U} , i.e., \mathbf{U}_1 , can be similar to \mathbf{U}_{O1} under some conditions. For example, when the leading eigenvalues are large enough or their corresponding eigenvectors are sparse, \mathbf{U}_{O1} can be well-estimated by \mathbf{U}_1 or its sparse counterpart (Johnstone & Lu, 2009).

3.3 Computation

In many contemporary applications, p is much larger than n . For example, in some gene expression data sets, p is a few thousands or more, and n is a few hundreds or less. So it is time-consuming to calculate $p \times p$ matrix \mathbf{T}_{γ} and its eigenvalue decomposition directly. The following lemma offers a shortcut to find \mathbf{U}_1 .

Lemma 1 We can write $\mathbf{T}_\gamma = n^{-1}\mathbf{A}_\gamma^\top\mathbf{A}_\gamma$, where

$$\mathbf{A}_\gamma = (\mathbf{X}_1 - \hat{\boldsymbol{\mu}}_{Y_1}, \dots, \mathbf{X}_n - \hat{\boldsymbol{\mu}}_{Y_n}, (\gamma n_1)^{1/2}(\hat{\boldsymbol{\mu}}_1 - \hat{\boldsymbol{\mu}}), \dots, (\gamma n_G)^{1/2}(\hat{\boldsymbol{\mu}}_K - \hat{\boldsymbol{\mu}}))^\top$$

is an $(n + K) \times p$ matrix. Note that $\hat{\boldsymbol{\mu}}_{Y_i} = \hat{\boldsymbol{\mu}}_k$ when $Y_i = k$, $\hat{\boldsymbol{\mu}} = n^{-1} \sum_{k=1}^K n_k \hat{\boldsymbol{\mu}}_k = \mathbf{0}$ by our convention.

When $p > n + K$ we can conduct eigenvalue decomposition for $(n + K) \times (n + K)$ matrix $\mathbf{A}_\gamma\mathbf{A}_\gamma^\top$ instead of $p \times p$ matrix $\mathbf{A}_\gamma^\top\mathbf{A}_\gamma$, because the eigenvectors of $\mathbf{A}_\gamma^\top\mathbf{A}_\gamma$ and $\mathbf{A}_\gamma\mathbf{A}_\gamma^\top$ are up to a linear transformation. To elaborate, by singular decomposition $\mathbf{A}_\gamma = \mathbf{V}_\gamma\boldsymbol{\Gamma}\mathbf{U}_\gamma^\top$, where $\boldsymbol{\Gamma}$ is diagonal, \mathbf{V}_γ is $(n + K) \times (n + K)$ orthogonal matrix, and \mathbf{U}_γ is $p \times p$ orthogonal matrix identical to \mathbf{U}_γ in (3.1). When $n + K$ is small or moderate, it is easy to find \mathbf{V}_γ which consists of eigenvectors of $\mathbf{A}_\gamma\mathbf{A}_\gamma^\top$. And the first $(n + K)$ columns of \mathbf{U}_γ can be obtained by standardizing $\mathbf{A}_\gamma^\top\mathbf{V}_\gamma$ column-wisely, as $\mathbf{A}_\gamma^\top\mathbf{V}_\gamma = \mathbf{U}_\gamma\boldsymbol{\Gamma}^\top$. In practice, it is sufficient to consider only the first $(n + K)$ columns of \mathbf{U}_γ (i.e., top principal components of \mathbf{T}_γ), because the rest columns correspond to eigenvalue 0 and contains little information.

For a fixed K , consider the scenario $n < p$. The computational complexities to find $\mathbf{A}_\gamma\mathbf{A}_\gamma^\top$ and conduct its singular decomposition are $O(n^2p)$ and $O(n^3)$, respectively. The computational complexity to find and standardize $\mathbf{A}_\gamma^\top\mathbf{V}_\gamma$ is $O(n^2p)$. Therefore, the overall computational complexity is $O(n^2p)$, which is linear with respect to p for a fixed n . Therefore, our method is computationally efficient to analyze high dimensional data.

4 Numerical studies

4.1 Simulated data examples

We compare our proposed new reduced-rank LDA method (SPCALDA) with some other popular classification tools by simulated data examples. In particular, we considered simple reduced-rank LDA (SRRLDA) reviewed in Section 2.2, LDA after standard PCA (PCALDA) that is a special case of SPCALDA with fixed $\gamma = 1$, and the independence rule (IR) reviewed in Section 2.4. Moreover, we added the Bayes rule as an oracle benchmark for comparison.

Six scenarios are considered here. For each scenario, 200 observations are generated and equally split between the four classes. Among 200 observations, 100 are assigned to the training set, and the rest 100 serve as test data. There are $p = 500$ features. For each class k , $\mathbf{X} \sim \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_w)$, where $\boldsymbol{\mu}_k$ and $\boldsymbol{\Sigma}_w$ are defined as follows.

Scenario 1. The common covariance $\boldsymbol{\Sigma}_w = \mathbf{I}_p$. The mean vectors are given by $\mu_{1j} = 0.3 * \mathbb{I}_{1 \leq j \leq 125}$, $\mu_{2j} = 0.3 * \mathbb{I}_{126 \leq j \leq 250}$, $\mu_{3j} = 0.3 * \mathbb{I}_{251 \leq j \leq 375}$, $\mu_{4j} = 0.3 * \mathbb{I}_{376 \leq j \leq 500}$, where $\mathbb{I}_{\mathcal{S}}$ is a vector with entries 1 for indices in \mathcal{S} and 0 elsewhere.

Scenario 2. $\boldsymbol{\Sigma}_w = \mathbf{I}_p$. $\mu_{1j} \sim \mathcal{N}(0, 0.3^2)$ when $1 \leq j \leq 125$, and $\mu_{1j} = 0$ otherwise, $\mu_{2j} \sim \mathcal{N}(0, 0.3^2)$ when $126 \leq j \leq 250$, and $\mu_{2j} = 0$ otherwise, $\mu_{3j} \sim \mathcal{N}(0, 0.3^2)$ when $251 \leq j \leq 375$, and $\mu_{3j} = 0$ otherwise, $\mu_{4j} \sim \mathcal{N}(0, 0.3^2)$ when $376 \leq j \leq 500$, and $\mu_{4j} = 0$ otherwise.

Scenario 3. $\boldsymbol{\Sigma}_w = (\sigma_{ij})$ with $\sigma_{ii} = 1$ and $\sigma_{ij} = 0.5$ for $i \neq j$. $\mu_{1j} = 0.21 * \mathbb{I}_{1 \leq j \leq 125}$, $\mu_{2j} = 0.21 * \mathbb{I}_{126 \leq j \leq 250}$, $\mu_{3j} = 0.21 * \mathbb{I}_{251 \leq j \leq 375}$, $\mu_{4j} = 0.21 * \mathbb{I}_{376 \leq j \leq 500}$.

Scenario 4. $\boldsymbol{\Sigma}_w$ is the same as in Scenario 3. $\mu_{1j} \sim \mathcal{N}(0, 0.21^2)$ when $1 \leq j \leq 125$, and $\mu_{1j} = 0$ otherwise, $\mu_{2j} \sim \mathcal{N}(0, 0.21^2)$ when $126 \leq j \leq 250$, and $\mu_{2j} = 0$ otherwise, $\mu_{3j} \sim \mathcal{N}(0, 0.21^2)$ when $251 \leq j \leq 375$, and $\mu_{3j} = 0$ otherwise, $\mu_{4j} \sim \mathcal{N}(0, 0.21^2)$ when $376 \leq j \leq 500$, and $\mu_{4j} = 0$ otherwise.

To test the robustness of the proposed method against departures from Gaussian and equal covariance assumptions, we include two more scenarios. In particular, Scenario 5 considers a case when the data are contaminated by a random heavy-tailed noise. Scenario 6 considers the situation when the observations from different classes do not share a common covariance structure.

Scenario 5. \mathbf{X} is generated as in Scenario 3. Let \mathbf{Z} is a p dimensional random vector with entries IID from t -distribution with degrees of freedom 3. We assume that realizations of $\tilde{\mathbf{X}} = \mathbf{X} + 0.2\mathbf{Z}$ are observed instead of \mathbf{X} .

Scenario 6. \mathbf{X} is generated as in Scenario 3. For each class k , we first generate a p dimensional vector \mathbf{d}_k with entries IID from standard uniform distribution, and define a diagonal covariance matrix $\boldsymbol{\Delta}_k = \text{diag}(\mathbf{d}_k^2)$. For each class k , let $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Delta}_k)$. We assume that realizations of $\tilde{\mathbf{X}} = \mathbf{X} + \mathbf{Z}$ are observed instead of \mathbf{X} .

There are two tuning parameters for SPCALDA (i.e., γ and q) and one parameter for PCALDA method, which were chosen by five-fold cross validation. The fitted models were evaluated using the test set for all methods. We repeated each experiment 100 times. The average and standard deviation of classification error rates for each method are listed in Table 1. The SPCALDA method always outperformed PCALDA, which indicates that it is helpful to tune the parameter γ . In the independence cases, SPCALDA is comparable with SRRLDA and IR but much better than them in the correlated cases. In general, the performance of SRRLDA and IR highly depends on the true covariance structure and SPCALDA is robust over different covariance structures. Moreover, when the Gaussian assumption or equal covariance assumption is violated, we see that SPCALDA still performs reasonably well.

Table 1: Mean (and standard errors) of classification error rates (%).

	SPCALDA	PCALDA	SRRLDA	IR	Oracle
Scenario 1	18.93(4)	26.53(4.52)	19.33(3.94)	18.45(3.86)	2.69(1.6)
Scenario 2	19.96(3.91)	27.71(5.1)	20.46(4.7)	19.29(4.03)	2.8(1.75)
Scenario 3	20.73(4.32)	30(5.64)	36.61(10.75)	63.92(5.41)	2.73(1.63)
Scenario 4	22.78(4.4)	32.26(5.82)	38.61(10.31)	64.38(7.92)	3.07(1.69)
Scenario 5	28.8(4.82)	38.42(6.41)	43.52(9.66)	64.38(5.8)	NA
Scenario 6	38.29(5.35)	50.75(6.72)	49.44(8.85)	64.79(6.57)	NA

4.2 Real data examples

In this section, we illustrate the performance of our method using six popular gene expression data sets, which have been studied in the literature. In particular, we considered three binary data sets, **Chin** (Chin et al., 2006), **Chowdary** (Chowdary et al., 2006), **Gordon** (Gordon et al., 2002), and three multi-class data sets, **Golub** (Golub et al., 1999), **Nakayama** (Nakayama et al., 2007), and **Sun** (Sun et al., 2006). The three binary data sets are available in R package `datamicroarray`. The data set **Golub** is available in R package `golubEsets`. The original **Nakayama** data set contains 105 samples from 10 types of soft tissue tumors. We considered a subset of 86 samples belonging to 5 tumor types and ignored the rest tumor types for which less than 7 samples were available. **Nakayama** and **Sun** are available on Gene Expression Omnibus (Barrett et al., 2005) with accession numbers GDS2736 and GDS1962, respectively. We list in Table 2 the sample size, number of features, number of classes, data distribution among different classes and related disease for each data set.

Table 2: Data sets used in this study.

Data set	related disease	# samples	# features	# classes	data distribution
Chin	breast cancer	118	22,215	2	43, 75
Chowdary	breast cancer	104	22,283	2	42, 62
Gordon	lung cancer	181	12,533	2	87, 94
Golub	leukemia	72	7,129	3	9, 25, 38
Nakayama	soft tissue tumor	86	22,283	5	15, 15, 16, 19, 21
Sun	glioma	180	54,613	4	23, 26, 50, 81

Besides the methods considered in Section 4.1, we included some state-of-the-art multi-class classification tools including NSC (Tibshirani et al., 2002), RDA (Guo et al., 2007), PLDA (Witten & Tibshirani, 2011), and SDA (Clemmensen et al., 2012), which have been implemented by R packages `pamr`, `rda`, `penalizedLDA`, and `sparseLDA`, respectively. These methods are based on various sparsity assumptions.

For each of these data sets, we randomly split the data with a 3 to 1 ratio into a training set and a test set, in a balanced manner. Five-fold cross-validation was conducted on the training set to select the tuning parameters for all methods, and the classification error rates using test set were recorded. In Table 3 we list the average classification error rates and their standard deviations over 25 random training/test set splits. We omit the results of PCALDA and IR which are dominated by SPCALDA and NSC, respectively. We see that SPCALDA performed best for two data sets, and second best for four data sets. In particular, SPCALDA is clearly the best in pairwise comparisons with other methods. We list in Table 4 the computation time for each method. We find that all these methods are reasonably fast to handle contemporary high dimensional data sets. In summary, our proposed method SPCALDA offers a competitive classification tool for high dimensional gene expression data.

Table 3: Mean (and standard errors) of classification error rates (%).

	SPCALDA	SRRLDA	NSC	PLDA	RDA	SDA
Chin	11.57(6.57)	12.11(6.13)	12.41(7.42)	13.87(6.79)	12.25(5.32)	10.13(4.47)
Chowdary	4.13(3.62)	10.43(5.82)	5.19(5.03)	33.63(9.52)	4.75(3.9)	17.19(8.26)
Gordon	0.62(1.19)	2.29(2.99)	0.79(1.08)	0.53(0.96)	1.4(1.25)	6.02(3.41)
Golub	5.43(5.44)	24.2(12.93)	4.6(4.17)	7.41(5.91)	6.3(3.74)	15.89(11.62)
Nakayama	16.37(7.05)	20.6(8.94)	23.51(6.36)	27.6(8.84)	15.68(7.98)	33.73(6.89)
Sun	30.43(5.73)	31.63(6.89)	33.24(6.03)	33.21(5.89)	33.48(6.97)	33.33(8.78)

Table 4: Mean (and standard error) of computation time per replicate (in second).

	SPCA-LDA	SRR-LDA	NSC	PLDA	RDA	SDA
Chin	14.5(1.03)	0.05(0.03)	3.93(0.61)	13.51(0.77)	51.17(3.27)	1.12(0.22)
Chowdary	12.09(0.25)	0.05(0.04)	3.63(0.21)	12.41(0.75)	49.52(0.43)	1.09(0.17)
Gordon	11.68(0.15)	0.04(0.02)	1.95(0.07)	6.82(0.1)	30.2(0.44)	0.53(0.04)
Golub	2.98(0.08)	0.01(0.01)	0.71(0.04)	3.9(0.34)	10.97(0.25)	0.22(0.03)
Nakayama	8.77(0.2)	0.06(0.02)	2.33(0.09)	27.46(0.73)	38.14(0.55)	1.75(0.22)
Sun	55.16(1.19)	0.24(0.02)	9.16(0.28)	56.3(0.94)	162.03(2.46)	6.42(0.75)

5 Discussion

Feature selection and feature extraction are two popular strategies in statistical machine learning. In the context of this paper, the sparse methods such as NSC and SDA can conduct variable selection and model estimation simultaneously, and belong to the first category. On the other hand, our methods, including classical reduced-rank LDA and PCA as special cases, belong to the latter. Both of these two approaches have their strength and weakness. For example, sparse methods can achieve model selection consistency and efficiency under various sparse assumptions. But they may fail when the true model is far from sparse. In contrast, our approach does not depend on sparse assumptions and is robust against the sparsity level of the true model. Our real data examples also verify the robustness of our method. In general, we can not expect a result on model selection consistency or efficiency. Nevertheless, we discuss a spiked covariance condition under which our method may achieve efficiency.

6 Supplementary Materials

The supplementary material is available on the journal web site. It contains Corollary 1, Lemma 2, and technical proofs for Theorems 1 and 2.

Acknowledgements

This research was partially supported by grants DMS 1309507 from National Science Foundation and 11671022 from National Science Foundation of China. The authors are grateful to Daniela Witten for helpful discussion, and to the editor, associate editor, and two referees for comments and suggestions that greatly improved the paper.

References

- BAIR, E., HASTIE, T., PAUL, D. & TIBSHIRANI, R. (2006). Prediction by supervised principal components. *Journal of the American Statistical Association* **101**.
- BARRETT, T., SUZEK, T. O., TROUP, D. B., WILHITE, S. E., NGAU, W.-C., LEDOUX, P., RUDNEV, D., LASH, A. E., FUJIBUCHI, W. & EDGAR, R. (2005). Ncbi geo: mining millions of expression profiles database and tools. *Nucleic acids research* **33**, D562–D566.
- 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**.
- CHIN, K., DEVRIES, S., FRIDLAND, J., SPELLMAN, P. T., ROYDASGUPTA, R., KUO, W.-L., LAPUK, A., NEVE, R. M., QIAN, Z., RYDER, T. et al. (2006). Genomic and transcriptional aberrations linked to breast cancer pathophysiologies. *Cancer cell* **10**, 529–541.
- CHOWDARY, D., LATHROP, J., SKELTON, J., CURTIN, K., BRIGGS, T., ZHANG, Y., YU, J., WANG, Y. & MAZUMDER, A. (2006). Prognostic gene expression signatures can be measured in tissues collected in rnalater preservative. *The journal of molecular diagnostics* **8**, 31–39.
- CLEMMENSEN, L., HASTIE, T., WITTEN, D. & ERSBØLL, B. (2012). Sparse discriminant analysis. *Technometrics* .

- FAN, J. & FAN, Y. (2008). High dimensional classification using features annealed independence rules. *Annals of statistics* **36**, 2605.
- FAN, J., FENG, Y. & TONG, X. (2012). A road to classification in high dimensional space: the regularized optimal affine discriminant. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **74**, 745–771.
- GOLUB, T. R., SLONIM, D. K., TAMAYO, P., HUARD, C., GAASENBEEK, M., MESIROV, J. P., COLLIER, H., LOH, M. L., DOWNING, J. R., CALIGIURI, M. A. et al. (1999). Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *science* **286**, 531–537.
- GORDON, G. J., JENSEN, R. V., HSIAO, L.-L., GULLANS, S. R., BLUMENSTOCK, J. E., RAMASWAMY, S., RICHARDS, W. G., SUGARBAKER, D. J. & 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**, 4963–4967.
- GUO, Y., HASTIE, T. & TIBSHIRANI, R. (2007). Regularized linear discriminant analysis and its application in microarrays. *Biostatistics* **8**, 86–100.
- HAO, N., DONG, B. & FAN, J. (2015). Sparsifying the fisher linear discriminant by rotation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* .
- HASTIE, T., TIBSHIRANI, R. & FRIEDMAN, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer series in statistics. Springer-Verlag New York.
- 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.
- JOLLIFFE, I. (2002). *Principal Component Analysis*. Springer Series in Statistics. Springer.
- MAI, Q., YANG, Y. & ZOU, H. (2015). Multiclass sparse discriminant analysis. *arXiv preprint arXiv:1504.05845* .
- MAI, Q., ZOU, H. & YUAN, M. (2012). A direct approach to sparse discriminant analysis in ultra-high dimensions. *Biometrika* **102**, 33–45.

NAKAYAMA, R., NEMOTO, T., TAKAHASHI, H., OHTA, T., KAWAI, A., SEKI, K., YOSHIDA, T., TOYAMA, Y., ICHIKAWA, H. & HASEGAWA, T. (2007). Gene expression analysis of soft tissue sarcomas: characterization and reclassification of malignant fibrous histiocytoma. *Modern pathology* **20**, 749–759.

SUN, L., HUI, A.-M., SU, Q., VORTMEYER, A., KOTLIAROV, Y., PASTORINO, S., PASANITI, A., MENON, J., WALLING, J., BAILEY, R. et al. (2006). Neuronal and glioma-derived stem cell factor induces angiogenesis within the brain. *Cancer cell* **9**, 287–300.

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.

WITTEN, D. M. & TIBSHIRANI, R. (2011). Penalized classification using fisher’s linear discriminant. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**, 753–772.

University of Arizona, Tucson, AZ, 85721, USA

E-mail: yueniu@math.arizona.edu

University of Arizona, Tucson, AZ, 85721, USA

E-mail: nhao@math.arizona.edu

Beijing International Center for Mathematical Research, Peking University, Beijing, China

E-mail: dongbin@math.pku.edu.cn