



A Direct Estimation Approach to Sparse Linear Discriminant Analysis

Tony Cai & Weidong Liu

To cite this article: Tony Cai & Weidong Liu (2011) A Direct Estimation Approach to Sparse Linear Discriminant Analysis, Journal of the American Statistical Association, 106:496, 1566-1577, DOI: [10.1198/jasa.2011.tm11199](https://doi.org/10.1198/jasa.2011.tm11199)

To link to this article: <http://dx.doi.org/10.1198/jasa.2011.tm11199>



Published online: 24 Jan 2012.



Submit your article to this journal [↗](#)



Article views: 745



View related articles [↗](#)



Citing articles: 2 View citing articles [↗](#)

A Direct Estimation Approach to Sparse Linear Discriminant Analysis

Tony CAI and Weidong LIU

This article considers sparse linear discriminant analysis of high-dimensional data. In contrast to the existing methods which are based on separate estimation of the precision matrix Ω and the difference δ of the mean vectors, we introduce a simple and effective classifier by estimating the product $\Omega\delta$ directly through constrained ℓ_1 minimization. The estimator can be implemented efficiently using linear programming and the resulting classifier is called the linear programming discriminant (LPD) rule.

The LPD rule is shown to have desirable theoretical and numerical properties. It exploits the approximate sparsity of $\Omega\delta$ and as a consequence allows cases where it can still perform well even when Ω and/or δ cannot be estimated consistently. Asymptotic properties of the LPD rule are investigated and consistency and rate of convergence results are given. The LPD classifier has superior finite sample performance and significant computational advantages over the existing methods that require separate estimation of Ω and δ . The LPD rule is also applied to analyze real datasets from lung cancer and leukemia studies. The classifier performs favorably in comparison to existing methods.

KEY WORDS: Classification; Constrained l_1 minimization; Fisher's rule; Naive Bayes rule; Sparsity.

1. INTRODUCTION

Classification is an important problem which has been well studied in the classical low-dimensional setting. In particular, linear discriminant analysis (LDA), which uses a linear combination of features as the criterion for classification, has been shown to perform well and enjoy certain optimality as the sample size tends to infinity while the dimension is fixed. Consider two p -dimensional normal distributions $N(\mu_1, \Sigma)$ (class 1) and $N(\mu_2, \Sigma)$ (class 2) with the same covariance matrix. Let \mathbf{Z} be a random vector that is drawn from one of these two distributions with equal prior probabilities. The goal of classification is to determine from which class \mathbf{Z} is drawn. The problem is simple in the ideal setting where the parameters μ_1 , μ_2 , and Σ are known in advance. In this case, Fisher's linear discriminant rule

$$\psi_F(\mathbf{Z}) = I\{(\mathbf{Z} - \mu)' \Omega \delta \geq 0\}, \quad (1)$$

where $\mu = (\mu_1 + \mu_2)/2$, $\delta = \mu_1 - \mu_2$, and $\Omega = \Sigma^{-1}$, classifies \mathbf{Z} into class 1 if and only if $\psi_F(\mathbf{Z}) = 1$. This classifier is the Bayes rule with equal prior probabilities for the two classes and is thus optimal in such an ideal setting.

Fisher's rule can be used to serve as an oracle benchmark, but it is typically not directly applicable in real data analysis as the parameters are usually unknown and need to be estimated from the samples. It is a standard practice to separately estimate Ω and δ and then plug the estimates into (1) to construct a classifier. Let $\{\mathbf{X}_k; 1 \leq k \leq n_1\}$ and $\{\mathbf{Y}_k; 1 \leq k \leq n_2\}$ be independent and identically distributed random samples from $N(\mu_1, \Sigma)$ and $N(\mu_2, \Sigma)$, respectively. The classical estimates of μ_1 , μ_2 , and Ω in the low-dimensional setting are the sample means $\bar{\mathbf{X}}$ and $\bar{\mathbf{Y}}$ and the inverse sample covariance matrix $\hat{\Sigma}_n^{-1}$. Plugging these estimates into (1) results in $\hat{\psi}_F(\mathbf{Z})$, the empirical version

of $\psi_F(\mathbf{Z})$. Theoretical properties of $\hat{\psi}_F(\mathbf{Z})$ have been well studied when p is fixed and can be found, for example, in the book by Anderson (2003).

With dramatic advances in technology, high-dimensional data are now routinely collected in a wide range of applications and classification for these data has drawn considerable recent attention. Examples include genomics, functional magnetic resonance imaging, risk management, and web search problems. In the high-dimensional settings, the standard LDA performs poorly and can even fail completely. For example, Bickel and Levina (2004) showed that the LDA can be no better than random guessing when $p/(n_1 + n_2) \rightarrow \infty$. In such a setting, the sample covariance matrix $\hat{\Sigma}_n$ is singular and its inverse is not well defined. One natural remedy is to use instead the generalized inverse of the sample covariance matrix. However, such an estimate is highly biased and unstable and will lead to a classifier with poor performance when p is large. A naive method in this case is to simply ignore the dependence among the variables and replace Σ with the diagonal of the sample covariance matrix. This leads to the so-called naive Bayes rule, also called the independence rule; see the article by Bickel and Levina (2004). Assuming that the difference δ is sparse, Fan and Fan (2008) proposed the features annealed independence rule which applies the naive independence rule to a set of selected important features of δ that are chosen by thresholding. This rule ignores the correlations between the variables and can be inefficient. See Section 6 for further discussions.

In the high-dimensional setting, regularity conditions on Ω (or Σ) and δ are needed to ensure that they can be estimated consistently. The most commonly used structural assumptions are that Ω (or Σ) and δ are sparse. Under such assumptions, Ω and δ are estimated separately and are then plugged into the Fisher's rule (1). Assuming the covariance matrix Σ and the difference δ are sparse, Shao et al. (2011) used the thresholding procedures for estimating Σ and δ . More commonly in applications the sparsity assumption is on the precision matrix

Tony Cai is Professor of Statistics, Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104 (E-mail: tcai@wharton.upenn.edu). Weidong Liu is Faculty Member, Department of Mathematics and Institute of Natural Sciences, Shanghai Jiao Tong University, China and Postdoctoral Fellow, Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104. The research of Tony Cai and Weidong Liu was supported in part by NSF FRG grant DMS-0854973.

Ω instead of Σ . In such a setting, Rothman et al. (2008) used the Glasso estimator for Ω in (1). Witten and Tibshirani (2009) proposed the scout procedure for classification in which they replaced Ω with a shrunken estimate. See also the works of Friedman (1989), Tibshirani et al. (2002), Guo, Hastie, and Tibshirani (2007), Wu et al. (2009), and Hall, Titterington, and Xue (2009), and the references therein.

A simple but important observation is that the Fisher’s discriminant rule (1) depends on Ω and δ only through their product $\Omega\delta$. In the present article, we shall show that the product $\Omega\delta$ can be estimated directly and efficiently, even when Ω and/or δ cannot be well estimated individually. We introduce the following direct estimation method for sparse linear discriminant analysis by estimating $\Omega\delta$ through a constrained ℓ_1 minimization method. Specifically, we propose to estimate $\beta^* := \Omega\delta$ by

$$\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \{|\beta|_1 \text{ subject to } |\hat{\Sigma}_n \beta - (\bar{\mathbf{X}} - \bar{\mathbf{Y}})|_\infty \leq \lambda_n\},$$

where λ_n is a tuning parameter, and classify \mathbf{Z} to class 1 if and only if

$$(\mathbf{Z} - \hat{\mu})' \hat{\beta} \geq 0,$$

where $\hat{\mu} = (\bar{\mathbf{X}} + \bar{\mathbf{Y}})/2$. The estimator $\hat{\beta}$ can be implemented easily using linear programming. The resulting classification procedure is thus called the *linear programming discriminant (LPD) rule*. The LPD rule is data-driven and easy to implement. It has significant computational advantage over the existing methods that require separate estimation of Ω (or Σ) and δ , because it only requires the estimation of a p -dimensional vector via linear programming instead of the estimation of the inverse of a $p \times p$ covariance matrix.

Both the theoretical and numerical properties of the LPD rule are studied in this article. The LPD rule performs well when $\Omega\delta$ is approximately sparse, which is a weaker and more flexible assumption than that both Ω and δ are sparse. In particular, under this assumption the precision matrix Ω is not required to be sparse and may not be consistently estimable. The asymptotic properties of the LPD rule are investigated and consistency and rate of convergence results are given. In addition to the Gaussian case, extensions to the non-Gaussian distributions are also considered. Numerical performance of the LPD classifier is investigated using both simulated and real data. A simulation study is carried out and the numerical results show that the LPD rule has superior finite sample performance in comparison to several other classifiers. It significantly outperforms the alternative methods in terms of the average misclassification rate. The LPD rule is also applied to the analysis of real datasets from lung cancer and leukemia studies and performs favorably in comparison to existing methods.

The rest of the article is organized as follows. Section 2 introduces a constrained ℓ_1 minimization method for the direct estimation of $\Omega\delta$ which leads to the LPD classification rule. Section 3 investigates the asymptotic properties of the LPD rule in the Gaussian setting. Extensions to non-Gaussian distributions are given in Section 4. Section 5 first discusses the linear programming implementation of the LPD classifier, and then investigates the numerical performance of the LPD rule by simulations and by applications to lung cancer and leukemia datasets. Discussions of our results and other related work are given in Section 6. The main results are proved in Section 7.

2. CLASSIFICATION VIA DIRECT ESTIMATION OF $\Omega\delta$

In this section we introduce a constrained ℓ_1 minimization method for estimating the product $\Omega\delta$ directly. It will be shown in Sections 3–5 that the resulting classification rule enjoys desirable properties theoretically, computationally, and numerically. For ease of presentation, we shall focus on the Gaussian case in this section and Section 3. The non-Gaussian case is considered in Section 4. We begin by reviewing basic notation and definitions.

For a vector $\beta = (\beta_1, \dots, \beta_p)' \in \mathbb{R}^p$, define the ℓ_0 norm by $|\beta|_0 = \sum_{j=1}^p I\{\beta_j \neq 0\}$; the ℓ_q norm by $|\beta|_q = (\sum_{j=1}^p |\beta_j|^q)^{1/q}$ for $1 \leq q \leq \infty$ with the usual modification for $q = \infty$. The vector β is called k -sparse if it has at most k nonzero entries. For a matrix $\Omega = (\omega_{ij})_{p \times p}$, the matrix 1-norm is defined to be the maximum absolute column sum, $\|\Omega\|_{L_1} = \max_{1 \leq j \leq p} \sum_{i=1}^p |\omega_{ij}|$. For a matrix Ω , we say Ω is k -sparse if each row/column has at most k nonzero entries. For two sequences of real numbers $\{a_n\}$ and $\{b_n\}$, write $a_n = O(b_n)$ for $n \geq 1$ if there exists a constant C such that $|a_n| \leq C|b_n|$, write $a_n = o(b_n)$ if $\lim_{n \rightarrow \infty} a_n/b_n = 0$, and write $a_n \asymp b_n$ if there are positive constants c and C such that $c \leq a_n/b_n \leq C$ for all $n \geq 1$.

Recall that $\{\mathbf{X}_k; 1 \leq k \leq n_1\}$ and $\{\mathbf{Y}_k; 1 \leq k \leq n_2\}$ are independent and identically distributed random samples from $N(\mu_1, \Sigma)$ and $N(\mu_2, \Sigma)$, respectively. Set

$$\begin{aligned} \bar{\mathbf{X}} &= \frac{1}{n_1} \sum_{i=1}^{n_1} \mathbf{X}_i, & \bar{\mathbf{Y}} &= \frac{1}{n_2} \sum_{i=1}^{n_2} \mathbf{Y}_i, \\ \hat{\delta} &= \bar{\mathbf{X}} - \bar{\mathbf{Y}}, & \hat{\mu} &= (\bar{\mathbf{X}} + \bar{\mathbf{Y}})/2. \end{aligned} \tag{2}$$

Denote the sample covariance matrices by

$$\begin{aligned} \hat{\Sigma}_{\mathbf{X}} &= \frac{1}{n_1} \sum_{i=1}^{n_1} (\mathbf{X}_i - \bar{\mathbf{X}})(\mathbf{X}_i - \bar{\mathbf{X}})', \\ \hat{\Sigma}_{\mathbf{Y}} &= \frac{1}{n_2} \sum_{i=1}^{n_2} (\mathbf{Y}_i - \bar{\mathbf{Y}})(\mathbf{Y}_i - \bar{\mathbf{Y}})', \end{aligned}$$

and set

$$\hat{\Sigma}_n = \frac{1}{n} (n_1 \hat{\Sigma}_{\mathbf{X}} + n_2 \hat{\Sigma}_{\mathbf{Y}}),$$

where $n = n_1 + n_2$.

As mentioned in the [Introduction](#), most of the classification methods in the literature involve separate estimation of the unknown precision matrix $\Omega = \Sigma^{-1}$ and the difference of the means δ in the Fisher’s rule (1). In the high-dimensional setting, the sample covariance matrix $\hat{\Sigma}_n$ is typically not invertible and regularity conditions are needed in order to be possible to construct good estimators. It should be noted that accurate estimation of a large sparse precision matrix is a difficult and computationally costly problem itself. See, for example, the works of Ravikumar et al. (2008), Yuan (2009), and Cai, Liu, and Luo (2011).

It is clear that the Fisher’s rule (1) depends on Ω and δ only through their product $\Omega\delta$. We now introduce a constrained ℓ_1 minimization method to directly estimate the product $\Omega\delta$ by exploiting the (approximate) sparsity of $\Omega\delta$. We should note here that the sparsity of $\Omega\delta$ is a weaker and more flexible condition than the sparsity of both Ω and δ . In particular, it does not

require the precision matrix Ω to be sparse. See Remark 1 below for more discussion. Specifically, we propose to estimate $\beta^* := \Omega\delta$ by the solution to the following optimization problem:

$$\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \{|\beta|_1 \text{ subject to } |\hat{\Sigma}_n \beta - (\bar{X} - \bar{Y})|_\infty \leq \lambda_n\}, \quad (3)$$

where λ_n is a tuning parameter which will be specified later. The constrained ℓ_1 minimization method (3) is known to be an effective way for reconstructing sparse signals. The readers are referred to the articles of Donoho, Elad, and Temlyakov (2006) and Candès and Tao (2007) for more details on the ℓ_1 minimization methods for sparse signal recovery. We shall show that the direct estimate leads to a classifier that is more effective and efficient than those based on estimating Ω and δ separately.

Given the solution $\hat{\beta}$ to (3), we propose the following classification rule: classify Z to class 1 if and only if

$$(Z - \hat{\mu})' \hat{\beta} \geq 0. \quad (4)$$

The optimization problem (3) can be cast as a linear program. We shall call the discriminant in (4) the *Linear Programming Discriminant (LPD)* and the classification rule (4) the LPD rule.

The motivation behind the constrained ℓ_1 minimization method (3) for estimating $\beta^* = \Omega\delta$ directly can be easily seen as follows. Note that β^* is the solution to the equation $\Sigma\beta - \delta = 0$. When Σ and δ are unknown, they are replaced by their respective sample versions $\hat{\Sigma}_n$ and $\hat{\delta} = \bar{X} - \bar{Y}$. We then seek the most sparse solution within the feasible set

$$\{\beta : |\hat{\Sigma}_n \beta - (\bar{X} - \bar{Y})|_\infty \leq \lambda_n\}$$

to account for the variability in $\hat{\Sigma}_n$ and $\hat{\delta}$. The convex relaxation of using ℓ_1 minimization in place of ℓ_0 minimization is a standard technique in sparse signal recovery. We shall show in the next sections that the resulting classification rule (4) has desirable properties both asymptotically and numerically. The ℓ_1 minimization method (3) works well when $\Omega\delta$ is approximately sparse. It thus allows the case where Ω itself is not sparse. In other words, it is possible to classify Z with accuracy using the classifier (4) even when Ω cannot be estimated consistently.

In addition to its good performance in terms of classification accuracy, the classifier given in (4) also enjoys significant computational advantages over existing methods that require separate estimation of Ω and δ . This can be seen at an intuitive level. There are only p parameters in $\Omega\delta$, while one needs to estimate $p^2/2$ parameters if Ω and δ are estimated separately. More discussions on the computational issues will be given in Section 5.

Remark 1. It is easy to see that if δ is k_1 -sparse and Ω is k_2 -sparse, then $\Omega\delta$ is at most $k_1 k_2$ -sparse. Furthermore, the sparsity of $\Omega\delta$ does not require Ω being sparse. Suppose δ is k_1 -sparse and without loss of generality assume the nonzeros are among the first k_1 coordinates. (In general we can always reorder the rows/columns of Ω accordingly.) So, δ can be written as

$$\delta = \begin{pmatrix} \delta_1 \\ \mathbf{0} \end{pmatrix},$$

where δ_1 is a k_1 -dimensional vector. Write Ω as

$$\Omega = \begin{pmatrix} \Omega_{11} & \Omega'_{21} \\ \Omega_{21} & \Omega_{22} \end{pmatrix},$$

where Ω_{11} is $k_1 \times k_1$, Ω_{21} is $(p - k_1) \times k_1$, and Ω_{22} is $(p - k_1) \times (p - k_1)$. Then the sparsity of $\Omega\delta$ does not depend on the submatrix Ω_{22} at all. $\Omega\delta = \begin{pmatrix} \Omega_{11}\delta_1 \\ \Omega_{21}\delta_1 \end{pmatrix}$ is sparse if Ω_{21} is sparse. In particular, if there are at most k_2 nonzero elements on each column of Ω_{21} , then $\Omega\delta$ is $k_1(k_2 + 1)$ -sparse. No condition on Ω_{22} is needed. In general, it is not possible to consistently estimate Ω under the spectral norm without regularity conditions on Ω_{22} . The consistency of $\hat{\Omega}$ was required by Shao et al. (2011) through the invertibility of the estimated covariance matrix $\hat{\Sigma}$ and for the good asymptotic performance of the resulting classification rule. In fact, even when Ω is the identity matrix, joint estimation of $\Omega\delta$ by (3) may lead to a better misclassification rate than estimating Ω and δ separately as in the work of Shao et al. (2011). See Remark 5 for more details.

Finally we note that there are also cases where neither Ω nor δ is sparse, but $\Omega\delta$ is. For example, if $\delta = (\sigma_{11}, \dots, \sigma_{p1})'$, the first column of Σ , then $\Omega\delta = (1, 0, \dots, 0)'$. Hence the sparsity on $\Omega\delta$ is more flexible than assuming both Ω and δ are sparse.

3. ASYMPTOTIC PROPERTIES

We now turn to the theoretical properties of the LPD classifier given in (4). Both consistency and convergence rate results are given. We shall focus on the Gaussian case in this section. Extensions to the non-Gaussian case are discussed in Section 4 and numerical performance of the classifier will be considered in Section 5.

The misclassification rate of Fisher's rule (see, e.g., Anderson 2003) is

$$R := 1 - \Phi\left(\frac{1}{2}\Delta_p^{1/2}\right) \quad \text{with } \Delta_p = \delta'\Omega\delta, \quad (5)$$

which is the best possible performance in the ideal setting where all the parameters μ_1 , μ_2 , and Σ are known in advance. This can serve as an oracle benchmark for the performance of any data-driven classifier based on the samples $\{\mathbf{X}_k\}$ and $\{\mathbf{Y}_k\}$.

It is not difficult to calculate that, given the samples $\{\mathbf{X}_k\}$ and $\{\mathbf{Y}_k\}$, the conditional misclassification rate of the LPD rule is

$$R_n := 1 - \frac{1}{2}\Phi\left(-\frac{(\hat{\mu} - \mu_1)'\hat{\beta}}{(\hat{\beta}'\hat{\Sigma}\hat{\beta})^{1/2}}\right) - \frac{1}{2}\Phi\left(\frac{(\hat{\mu} - \mu_2)'\hat{\beta}}{(\hat{\beta}'\hat{\Sigma}\hat{\beta})^{1/2}}\right),$$

where $\hat{\beta}$ is given in (3). The performance of the LPD rule can then be naturally measured by the difference (or ratio) between R_n and the Bayes misclassification rate R . In this section we will study the difference and ratio between R_n and R . To this end, we need to introduce some conditions.

$$(C1) \quad n_1 \asymp n_2, \log p \leq n, c_0^{-1} \leq \lambda_{\min}(\Sigma) \leq \lambda_{\max}(\Sigma) \leq c_0 \text{ for some constant } c_0 > 0 \text{ and } \Delta_p \geq c_1 \text{ for some } c_1 > 0.$$

Here we assume that the two samples are of comparable sizes and the eigenvalues of the covariance matrix Σ are bounded from below and above. These are commonly used conditions in the high-dimensional setting. In addition, we also assume Δ_p is bounded away from zero. If $\Delta_p \rightarrow 0$, then it can be seen easily from (5) that even the oracle rule is no better than random guessing.

Our first result is on the consistency of R_n .

Theorem 1. Let $\lambda_n = C\sqrt{\Delta_p \log p/n}$ with $C > 0$ being a sufficiently large constant. Suppose (C1) holds and

$$|\mathbf{\Omega}\delta|_0 = o\left(\sqrt{\frac{n}{\log p}}\right). \quad (6)$$

Then we have as $n \rightarrow \infty$ and $p \rightarrow \infty$,

$$R_n - R \rightarrow 0 \quad (7)$$

in probability.

This theorem shows that the LPD rule is consistent when $\mathbf{\Omega}\delta$ is sparse. In practice, the value of the tuning parameter λ_n is chosen by cross-validation. See Section 5 for further discussions on the implementation of the LPD rule.

Remark 2. As mentioned earlier, the condition $|\mathbf{\Omega}\delta|_0 = o(\sqrt{n/\log p})$ does not require $\mathbf{\Omega}$ to be sparse. Therefore, by estimating $\mathbf{\Omega}\delta$ directly, we do not need a consistent estimate for $\mathbf{\Omega}$ or $\mathbf{\Sigma}$ under the spectral norm to get the asymptotically optimal misclassification rate. In contrast, consistent estimation of $\mathbf{\Omega}$ was required by Shao et al. (2011). A basic condition in the work of Shao et al. (2011) is that $\mathbf{\Sigma} =: (\sigma_{ij})_{p \times p}$ is (approximately) sparse with the sparsity $s_0(p)$ of $\mathbf{\Sigma}$ satisfying $s_0(p)(\log p/n)^{(1-q)/2} = o(1)$ and $\max_{1 \leq i \leq p} \sum_{j=1}^p |\sigma_{ij}|^q \leq s_0(p)$ for $0 \leq q < 1$. It follows from the work of Cai and Zhou (2010) on the minimax rate of convergence for estimating sparse covariance matrices, this condition is necessary for the consistency under the spectral norm.

Theorem 1 can be extended to a more general setting where $\mathbf{\Omega}\delta$ is only approximately sparse. To state this result, we first relax the condition (C1) as follows.

$$(C2) \quad n_1 \asymp n_2, \log p \leq n, \max_{1 \leq i \leq p} \sigma_{ii} \leq K \text{ for some constant } K > 0 \text{ and } \Delta_p \geq c_1 \text{ for some constant } c_1 > 0.$$

Theorem 2. Let $\lambda_n = C\sqrt{\Delta_p \log p/n}$ with C being a sufficiently large constant. Suppose (C2) holds and

$$\frac{|\mathbf{\Omega}\delta|_1}{\Delta_p^{1/2}} + \frac{|\mathbf{\Omega}\delta|_1^2}{\Delta_p^2} = o\left(\sqrt{\frac{n}{\log p}}\right). \quad (8)$$

Then we have as $n \rightarrow \infty$ and $p \rightarrow \infty$,

$$R_n - R \rightarrow 0 \quad (9)$$

in probability.

Remark 3. It follows from the Cauchy–Schwarz inequality and (C1),

$$|\mathbf{\Omega}\delta|_1^2 \leq |\mathbf{\Omega}\delta|_0 |\mathbf{\Omega}\delta|_2^2 \leq c_0^2 |\mathbf{\Omega}\delta|_0 |\delta|_2^2$$

and $\Delta_p \geq c_0^{-1} |\delta|_2^2$. Thus (6) implies (8). The condition (6) can be further relaxed if the minimum magnitude of the nonzero elements of $\mathbf{\Omega}\delta$ is relatively large. Let $S = \{i: (\mathbf{\Omega}\delta)_i \neq 0\}$. If $\min_{i \in S} |(\mathbf{\Omega}\delta)_i| \geq C(\log p/n)^{1/4}$, then a sufficient condition of (8) is $|\mathbf{\Omega}\delta|_0 = o(n/\log p)$. Condition (8) allows the case where $\mathbf{\Omega}\delta$ is only approximately sparse with many small entries.

Remark 4. The condition $\max_{1 \leq i \leq p} \sigma_{ii} \leq K$ can be relaxed. Let $K_p := \max_{1 \leq i \leq p} \sigma_{ii}$ and $\lambda_n = C\sqrt{K_p \Delta_p \log p/n}$. Theorem 2 still holds under the condition

$$\frac{|\mathbf{\Omega}\delta|_1}{\Delta_p^{1/2}} + \frac{|\mathbf{\Omega}\delta|_1^2}{\Delta_p^2} = o\left(\sqrt{\frac{n}{K_p \log p}}\right).$$

Here K_p can grow and may tend to infinity as $p \rightarrow \infty$.

Theorems 1 and 2 provide the consistency results for the LPD rule. Consistency is important, but the fact $R_n - R \rightarrow 0$ does not give a detailed description of the properties of a classifier. For example, when the Bayes misclassification rate $R \rightarrow 0$, any classifier with $R_n \rightarrow 0$ is consistent. Stronger results on the rate of convergence can be obtained.

Theorem 3. Let $\lambda_n = C\sqrt{\Delta_p \log p/n}$ with C being a sufficiently large constant. Suppose (C2) holds and

$$|\mathbf{\Omega}\delta|_1 \Delta_p^{1/2} + |\mathbf{\Omega}\delta|_1^2 = o\left(\sqrt{\frac{n}{\log p}}\right).$$

Then

$$\frac{R_n}{R} - 1 = O\left(\left(|\mathbf{\Omega}\delta|_1 \Delta_p^{1/2} + |\mathbf{\Omega}\delta|_1^2\right) \sqrt{\frac{\log p}{n}}\right)$$

with probability greater than $1 - O(p^{-1})$. In particular, if (C1) holds and

$$|\mathbf{\Omega}\delta|_0 \Delta_p = o\left(\sqrt{\frac{n}{\log p}}\right),$$

then

$$\frac{R_n}{R} - 1 = O\left(|\mathbf{\Omega}\delta|_0 \Delta_p \sqrt{\frac{\log p}{n}}\right)$$

with probability greater than $1 - O(p^{-1})$.

Theorem 3 shows that a larger Δ_p implies a worse convergence rate for the relative classification error R_n/R . This is in fact to be expected. When Δ_p is large, the classification problem is easy and the Bayes misclassification rate R can be very small. It then becomes harder for any data-driven classification rule to mimic the performance of the oracle rule.

Remark 5. Due to the differences in setting, our results and the results of Shao et al. (2011) are not directly comparable. To make them comparable, it is necessary to assume both $\mathbf{\Sigma}$ and $\mathbf{\Omega}$ are sparse. For simplicity, we consider the case $\mathbf{\Sigma} = I_{p \times p}$. Suppose that $|\delta|_1 \leq K$ for some constant K and $\log p = o(n)$. Theorem 3 shows that $R_n/R - 1 = O_P(\sqrt{\log p/n})$. Let R_n^* be the conditional misclassification rate of the SLDA rule proposed by Shao et al. (2011). Their results show that $R_n^*/R - 1 = O_P(b_n)$ with $b_n = (\log p/n)^{\alpha(1-g)} |\delta|_{2g}^g$ for some $0 < \alpha < 1/2$ and $0 \leq g < 1$. It is easy to see that $b_n/\sqrt{\log p/n} \rightarrow \infty$. So the LPD rule outperforms the SLDA rule in this case.

The convergence rate in Theorem 3 can be further improved under stronger conditions.

Theorem 4. Let $\lambda_n = C\sqrt{\Delta_p \log p/n}$ with C being a sufficiently large constant. Suppose (C2) holds and

$$\|\Omega\|_{L_1} |\Omega\delta|_0 + |\Omega\delta|_1 \Delta_p^{1/2} = o\left(\sqrt{\frac{n}{\log p}}\right). \quad (10)$$

Then

$$\frac{R_n}{R} - 1 = O\left(|\Omega\delta|_1 \Delta_p^{1/2} \sqrt{\frac{\log p}{n}}\right)$$

with probability greater than $1 - O(p^{-1})$. In particular, if (C1) holds and

$$\|\Omega\|_{L_1} |\Omega\delta|_0 + |\Omega\delta|_0^{1/2} \Delta_p = o\left(\sqrt{\frac{n}{\log p}}\right),$$

then

$$\frac{R_n}{R} - 1 = O\left(|\Omega\delta|_0^{1/2} \Delta_p \sqrt{\frac{\log p}{n}}\right)$$

with probability greater than $1 - O(p^{-1})$.

4. EXTENSIONS

Section 3 establishes the theoretical properties of the LPD classifier in the Gaussian setting. The results can be extended to a class of non-Gaussian distributions satisfying certain moment conditions.

Let \mathbf{X} and \mathbf{Y} be two p -dimensional random vectors satisfying

$$\mathbf{X} = \boldsymbol{\mu}_1 + \mathbf{U}_1 \quad \text{and} \quad \mathbf{Y} = \boldsymbol{\mu}_2 + \mathbf{U}_2,$$

where \mathbf{U}_1 and \mathbf{U}_2 are independent and identically distributed random vectors with mean zero and covariance matrix $\boldsymbol{\Sigma} = (\sigma_{ij})_{p \times p}$. Fang and Anderson (1990) showed that Fisher's rule is still optimal when \mathbf{U}_1 has an elliptical distribution with zero mean and density

$$c_p |\boldsymbol{\Sigma}|^{-1/2} f(\mathbf{u}' \boldsymbol{\Sigma}^{-1} \mathbf{u}), \quad (11)$$

where f is a monotone function on $[0, \infty)$ and c_p is a normalizing constant. The optimal misclassification rate in this case is

$$R = \frac{1}{2} \mathbb{P}(\mathbf{U}'_1 \Omega \delta < -\frac{1}{2} \delta' \Omega \delta) + \frac{1}{2} \mathbb{P}(\mathbf{U}'_2 \Omega \delta \geq \frac{1}{2} \delta' \Omega \delta).$$

As in the work of Shao et al. (2011), we relax the normality of \mathbf{U}_1 to that, for any p -dimensional non-random vector \mathbf{l} with $\|\mathbf{l}\|_2 = 1$ and any $t \in R$,

$$\mathbb{P}(\mathbf{l}' \Omega^{1/2} \mathbf{U}_1 \leq t) =: \Psi(t)$$

is a continuous distribution function symmetric about 0 and does not depend on \mathbf{l} . The elliptical distributions [such as (11)] and the multivariate scale mixture of normals satisfy this condition. The conditional classification error of the LPD rule (4) given $\{\mathbf{X}_k\}$ and $\{\mathbf{Y}_k\}$ is

$$R_n := 1 - \frac{1}{2} \Psi\left(-\frac{(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_1)' \hat{\boldsymbol{\beta}}}{(\hat{\boldsymbol{\beta}}' \boldsymbol{\Sigma} \hat{\boldsymbol{\beta}})^{1/2}}\right) - \frac{1}{2} \Psi\left(\frac{(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_2)' \hat{\boldsymbol{\beta}}}{(\hat{\boldsymbol{\beta}}' \boldsymbol{\Sigma} \hat{\boldsymbol{\beta}})^{1/2}}\right).$$

To obtain the convergence rate for R_n , we shall impose an additional condition: for any $x < 0$ and $|\delta| \leq 1$,

$$\left| \frac{\Psi(x + \delta)}{\Psi(x)} - 1 \right| \leq c_1 |\delta| (|x| + 1) e^{c_2 |x\delta|} \quad (12)$$

for some positive constants c_1, c_2 which do not depend on x and δ . Note that the distribution with density function $p(x) = c_3(1 + |x|)^{-w} e^{-c_4 |x|^\varphi}$ satisfies (12), where c_3 and c_4 are positive constants, φ and w are constants with $0 < \varphi \leq 2$, $w \geq 0$, or $\varphi = 0$, $w > 1$.

The moment conditions are divided into two cases: the sub-Gaussian-type tails and the polynomial-type tails. Let $\mathbf{U}_1 =: (U_1, \dots, U_p)'$ and $U_\delta = \mathbf{U}'_1 \Omega \delta / \Delta_p^{1/2}$. Note that U_δ is a standardized random variable with zero mean and unit variance.

(C3) (*Sub-Gaussian-type tails*) Suppose that $\log p \leq n$ and there exist some constants $\eta > 0$ and $K_1 > 0$ such that

$$\mathbb{E} \exp(\eta U_\delta^2) \leq K_1, \quad \text{and} \quad (13)$$

$$\mathbb{E} \exp(\eta U_i^2 / \sigma_{ii}) \leq K_1 \quad \text{for all } i.$$

(C4) (*Polynomial-type tails*) Suppose that for some $\gamma, c_1 > 0$, $p \leq c_1 n^\gamma$, and for some $\epsilon > 0$

$$\mathbb{E} |U_\delta|^{4\gamma+4+\epsilon} \leq K_1 \quad \text{and} \quad (14)$$

$$\mathbb{E} |U_i / \sigma_{ii}^{1/2}|^{4\gamma+4+\epsilon} \leq K_1 \quad \text{for all } i.$$

Theorem 5. (i) Assume that the conditions in Theorem 2 hold. By replacing the normality with elliptical distributions satisfying (C3) or (C4), we have as $n \rightarrow \infty$ and $p \rightarrow \infty$,

$$R_n - R \rightarrow 0 \quad \text{in probability.} \quad (15)$$

(ii) Under the conditions in Theorem 3, (12), and (C3) [(or (C4))],

$$\frac{R_n}{R} - 1 = O\left((|\Omega\delta|_1 \Delta_p^{1/2} + |\Omega\delta|_1^2) \sqrt{\frac{\log p}{n}}\right)$$

with probability greater than $1 - O(p^{-1} + n^{-\epsilon/8})$.

Similarly, Theorem 4 remains valid if the normality assumption is replaced by elliptical distributions satisfying (C3) or (C4). For reasons of space, we do not restate the results here.

5. NUMERICAL INVESTIGATION

We now turn to the numerical performance of the LPD rule using both simulated and real data. We begin in Section 5.1 with a discussion on the implementation of the classifier using linear programming and the selection of the tuning parameter λ_n through cross-validation. Section 5.2 presents simulation results and comparisons with other methods including oracle features annealed independence rule (OFAIR) (the support of δ is assumed to be known), the nearest shrunken centroids method (NSC) proposed by Tibshirani et al. (2002), the sparse linear discriminant (SLD) introduced by Shao et al. (2011), the Naive-Bayes rule (Naive-LDA), the LDA rule with a generalized inverse (GLDA) as well as the oracle Fisher's rule (Oracle). The applications of the LPD rule to the analysis of a lung cancer dataset and a leukemia dataset are given in Section 5.3.

5.1 Implementation of LPD

Recall that the estimate of $\beta^* = \Omega \delta$ is obtained by solving the constrained ℓ_1 minimization problem

$$\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \{ \|\beta\|_1 \text{ subject to } |\hat{\Sigma}_n \beta - (\bar{\mathbf{X}} - \bar{\mathbf{Y}})|_\infty \leq \lambda_n \}.$$

This optimization problem is convex, and can easily be recast as the following linear program:

$$\begin{aligned} \min \quad & \sum_{j=1}^p u_j \\ \text{subject to:} \quad & -\beta_j \leq u_j \quad \text{for all } 1 \leq j \leq p, \\ & +\beta_j \leq u_j \quad \text{for all } 1 \leq j \leq p, \\ & -\hat{\sigma}'_k \beta + \hat{\delta}_k \leq \lambda_n \quad \text{for all } 1 \leq k \leq p, \\ & +\hat{\sigma}'_k \beta - \hat{\delta}_k \leq \lambda_n \quad \text{for all } 1 \leq k \leq p, \end{aligned} \quad (16)$$

where $(\hat{\delta}_1, \dots, \hat{\delta}_p)' := \hat{\delta}$ and $(\hat{\sigma}_1, \dots, \hat{\sigma}_p) := \hat{\Sigma}_n$.

This linear programming implementation is similar to that of the Dantzig selector in high-dimensional linear regression. See the article by Candès and Tao (2007). We then apply the primal-dual interior-point method to solve (16). See, for example, the book by Boyd and Vandenberghe (2004) for more details on the primal-dual interior-point method. We should note that there are other stable algorithms based on first-order method that may be used to implement the optimization problem (3); see the work of Becker, Candès, and Grant (2011). Similarly to many iterative methods, one needs to specify a feasible initialization. To this end, we replace $\hat{\Sigma}_n$ in (3) by $\hat{\Sigma}_\rho = \hat{\Sigma}_n + \rho I_{p \times p}$ with a small positive number ρ (e.g., $\rho = \sqrt{\log p/n}$) and take the initial value to be $\hat{\Sigma}_\rho^{-1} \hat{\delta}$. Such a perturbation does not noticeably affect the computational accuracy of the final solution in our numerical experiments. All the theoretical properties in Sections 3 and 4 still hold for $\rho \leq \sqrt{\log p/n}$ with the additional condition $\lambda_{\max}(\Sigma) \leq K$ for some constant $K > 0$.

The computational cost of estimating $\Omega \delta$ directly through linear programming as described above is much smaller than that of estimating the precision matrix Ω . For example, if one estimates Ω using the method of Yuan (2009) or the CLIME method in the article by Cai, Liu, and Luo (2011), the computational cost is p times that of estimating $\Omega \delta$ directly by (3).

There is a tuning parameter $\lambda = \lambda_n$ in the algorithm. As mentioned before, λ can be chosen empirically by cross-validation (CV). This can be done as follows. Divide the sets $\{1, 2, \dots, n_1\}$ and $\{1, 2, \dots, n_2\}$ into $2N$ subgroups H_{11}, \dots, H_{1N} and H_{21}, \dots, H_{2N} . Thus the samples $\{\mathbf{X}_i, \mathbf{Y}_j; 1 \leq i \leq n_1, 1 \leq j \leq n_2\}$ are divided into $\mathcal{X}_k := \{\mathbf{X}_i, \mathbf{Y}_j; i \in H_{1k}, j \in H_{2k}\}$, $1 \leq k \leq N$. Let $\hat{\mu}^{(k)}$, $\hat{\delta}^{(k)}$, and $\hat{\Sigma}^{(k)}$ be defined as in (2), based on $\{\mathbf{X}_i, \mathbf{Y}_j; 1 \leq i \leq n_1, 1 \leq j \leq n_2\} \setminus \mathcal{X}_k$. For any given choice of λ , calculate $\hat{\beta}^{(k)}(\lambda)$ based on $\hat{\delta}^{(k)}$ and $\hat{\Sigma}^{(k)}$ by (3). Let $I_{j1}^{(k)} = 1$ if $(\mathbf{X}_j - \hat{\mu}^{(k)})' \hat{\beta}^{(k)}(\lambda) \geq 0$ for $\mathbf{X}_j \in \mathcal{X}_k$; else $I_{j1}^{(k)} = 0$. Similarly, define $I_{j2}^{(k)} = 1$ if $(\mathbf{Y}_j - \hat{\mu}^{(k)})' \hat{\beta}^{(k)}(\lambda) < 0$ for $\mathbf{Y}_j \in \mathcal{X}_k$; else $I_{j2}^{(k)} = 0$. Then

$$\text{CV}(\lambda) = \sum_{k=1}^N \left(\sum_{i \in H_{1k}} I_{i1}^{(k)} + \sum_{j \in H_{2k}} I_{j2}^{(k)} \right)$$

is the total number of correctly classified cases among the validation sets for the classifier with a given choice of λ . The final choice of λ is $\hat{\lambda} = \max_\lambda \text{CV}(\lambda)$. If the maximum is attained at several λ 's, the minimum value of these λ 's is selected.

5.2 Simulation Results

We now present simulation results and compare the numerical performance of the LPD classifier with the oracle features annealed independence rule (OFAIR) (Fan and Fan 2008) where the support of the difference δ is assumed to be known, the nearest shrunken centroids method (NSC) (Tibshirani et al. 2002), the sparse linear discriminant (SLD) (Shao et al. 2011), the Naive-Bayes rule (Naive-LDA), the LDA rule with a generalized inverse (GLDA), and the oracle Fisher's rule (Oracle). The oracle rule is included as a benchmark.

The setup in the simulation study is as follows. We fix the sample sizes $n_1 = n_2 = 200$ and set $\mu_1 = 0$ and $\mu_2 = (1, \dots, 1, 0, \dots, 0)'$, where the number of 1's is $s_0 = 10$. Three models are considered:

- Model 1. $\Omega = (\sigma_{ij})_{p \times p}^{-1}$ with $\sigma_{ii} = 1$ for $1 \leq i \leq p$ and $\sigma_{ij} = \rho$ with $\rho = 0.5$ for $i \neq j$.
- Model 2. $\Omega = (\mathbf{B} + \delta \mathbf{I}) / (1 + \delta)$, where $\mathbf{B} = (b_{ij})_{p \times p}$ with independent $b_{ij} = b_{ji} = 0.5 \times \text{Ber}(1, 0.2)$ for $1 \leq i \leq s_0$, $i < j \leq p$; $b_{ij} = b_{ji} = 0.5$ for $s_0 + 1 \leq i < j \leq p$; $b_{ii} = 1$ for $1 \leq i \leq p$. Here $\text{Ber}(1, 0.2)$ is a Bernoulli random variable which takes value 1 with probability 0.2 and 0 with probability 0.8; and $\delta = \max(-\lambda_{\min}(\mathbf{B}), 0) + 0.05$ to ensure that Ω is positive definite. Finally, the matrix is standardized to have unit diagonals.
- Model 3. $\Omega = \Sigma^{-1}$, where $\Sigma = (\sigma_{ij})_{p \times p}$ with $\sigma_{ij} = 0.8^{|i-j|}$ for $1 \leq i, j \leq p$.

Ω in Model 1 is an approximately sparse matrix. It is diagonally dominant with the off-diagonal entries of order p^{-1} . In Model 1 $\Omega \delta$ is also approximately sparse. In Model 2, only the first s_0 rows and columns of Ω are sparse and the rest of the matrix is not sparse. In Model 3, Σ can be well approximated by a sparse matrix and the inverse Ω is a 3-sparse matrix. Model 3 satisfies the conditions in both the article by Shao et al. (2011) and the present article. This enables a fair comparison between the SLD of Shao et al. (2011) and the LPD rule.

In the simulation, we generate $n_1 = n_2 = 200$ training and test samples of the same size according to Models 1–3 with the multivariate normal distribution and the multivariate t distribution with five degrees of freedom. The tuning parameter λ_n is chosen by fivefold cross-validation as described in Section 5.1. Note that the covariance matrix Σ in Models 1 and 2 is not sparse. So the thresholding estimator for Σ used by Shao et al. (2011) may be not invertible. The generalized inverse of the thresholding estimator is used when the estimator itself is not invertible. The SLD rule in the work of Shao et al. (2011) requires to choose two tuning parameters by cross-validation. To reduce the computational cost, when implementing the SLD rule we assume the support of δ is known so that only the tuning parameter for estimating the covariance matrix is needed. The average classification errors for the test samples and the standard deviations based on 100 replications are stated in Tables 1 and 2.

Table 1. Average classification error for the test samples in percentage in the normal distribution case. Standard deviations are given in parentheses

p	LPD	OFAIR	NSC	SLD	Naive-LDA	GLDA	Oracle
Model 1							
100	2.42 (0.78)	25.07 (2.01)	18.58 (8.27)	3.20 (0.89)	21.39 (11.53)	3.54 (1.00)	1.60 (0.07)
200	2.45 (0.75)	24.80 (1.85)	17.70 (9.18)	6.23 (1.35)	25.51 (11.94)	7.28 (1.53)	1.51 (0.60)
400	2.27 (0.83)	24.28 (2.28)	19.35 (8.35)	41.45 (4.32)	32.12 (12.16)	41.95 (4.18)	1.41 (0.57)
800	2.51 (1.08)	24.51 (2.03)	19.40 (8.31)	13.28 (1.97)	39.57 (9.43)	17.24 (2.20)	1.30 (0.61)
Model 2							
100	3.23 (0.99)	13.88 (3.10)	13.38 (4.90)	10.73 (5.15)	19.08 (6.57)	3.53 (0.98)	1.62 (0.64)
200	5.12 (1.24)	25.75 (4.07)	26.13 (5.88)	18.92 (7.61)	36.15 (5.05)	8.21 (1.37)	1.83 (0.64)
400	8.18 (1.59)	18.04 (3.24)	21.01 (5.61)	20.87 (8.32)	37.85 (4.70)	43.90 (3.70)	2.64 (0.78)
800	14.87 (2.27)	23.52 (2.69)	30.40 (4.49)	26.48 (4.82)	45.59 (3.38)	32.12 (2.92)	3.12 (0.80)
Model 3							
100	18.93 (2.08)	24.92 (2.00)	25.06 (2.04)	25.09 (2.61)	25.65 (2.22)	22.63 (2.25)	16.55 (1.74)
200	19.42 (2.14)	24.81 (1.95)	25.02 (2.07)	25.40 (4.96)	26.23 (2.18)	29.31 (2.42)	16.47 (1.94)
400	19.64 (2.47)	24.50 (2.31)	24.73 (2.47)	24.60 (2.49)	27.56 (2.39)	47.46 (3.21)	16.44 (2.21)
800	19.90 (2.34)	24.94 (2.26)	25.24 (2.43)	25.37 (3.25)	29.70 (2.36)	34.41 (2.61)	16.61 (2.04)

Table 1 displays the numerical results of the six classifiers as well as the oracle Fisher’s rule in the Gaussian case. For Model 1, the performance of the LPD rule is similar to that of the oracle Fisher’s rule, and is better by a large margin than those of the other five classifiers OFAIR, NSC, SLD, Naive-LDA, and GLDA. Comparing to these methods, LPD has the smallest classification errors with the smallest standard deviations. The classification error is also quite stable as p increases from 100 to 800. The performance of SLD is not stable in Model 1 because Σ is not sparse and the generalized inverse of the thresholding estimator is used. For Models 2 and 3, the LPD rule again significantly outperforms the other five classifiers. The misclassification rate of the LPD rule in Model 2 is less than half of those of the other five methods.

Table 2 shows the corresponding numerical results in the case of the multivariate t_5 distribution. In comparison to the results for the Gaussian case given in Table 1, it can be seen from Table 2 that the classification errors of all methods including the oracle rule increase when the tail of the distribution becomes

heavier. In this case the performance of the LPD rule remains close to that of the oracle rule in Model 1 and the LPD classifier again significantly outperforms OFAIR, NSC, SLD, Naive-LDA, and GLDA in all of the three models.

Support recovery by $\hat{\beta}$ is also considered in the simulation. We only consider Model 3, in which $\Omega\delta$ has 11 nonzero elements. Note that in Model 1 all of the elements of $\Omega\delta$ are nonzero and most of the elements of $\Omega\delta$ (more than 88%) in Model 2 generated by our simulation are nonzero. We thus do not consider support recovery for Models 1 and 2. In Model 3, the number of nonzero elements (POS) and true nonzero elements (TPOS) in $\hat{\beta}$ are calculated. The ability to recover the support is evaluated via the true positive rate (TPR) in combination with the false positive rate (FPR), defined respectively as

$$TPR = \frac{\#\{i : \hat{\beta}_i \neq 0 \text{ and } (\Omega\delta)_i \neq 0\}}{\#\{i : (\Omega\delta)_i \neq 0\}} \quad \text{and}$$

Table 2. Average classification error for the test samples in percentage in the t_5 distribution case. Standard deviations are given in parentheses

p	LPD	OFAIR	NSC	SLD	Naive-LDA	GLDA	Oracle
Model 1							
100	6.70 (1.20)	30.12 (2.26)	23.84 (8.33)	7.79 (1.36)	27.22 (10.82)	8.50 (1.42)	5.02 (1.17)
200	6.64 (1.28)	30.12 (2.07)	24.98 (8.26)	11.76 (1.80)	31.23 (12.37)	13.59 (2.07)	4.61 (1.06)
400	6.29 (1.38)	29.90 (2.47)	25.40 (8.29)	43.05 (4.95)	39.37 (8.85)	44.19 (3.71)	4.38 (0.85)
800	5.86 (1.09)	30.02 (2.21)	26.60 (7.81)	19.73 (2.40)	41.83 (10.24)	25.25 (2.56)	4.06 (1.09)
Model 2							
100	8.02 (1.35)	19.90 (3.51)	20.60 (6.04)	17.25 (8.80)	27.34 (6.85)	8.46 (1.47)	4.91 (0.92)
200	11.06 (1.89)	30.40 (4.15)	31.67 (5.30)	42.73 (13.97)	40.80 (4.08)	14.57 (2.09)	5.15 (1.04)
400	15.15 (1.96)	25.10 (3.80)	30.13 (5.71)	39.12 (8.33)	42.63 (3.63)	45.20 (3.40)	6.33 (1.28)
800	23.19 (2.12)	30.60 (3.73)	36.40 (4.08)	31.17 (4.08)	47.25 (2.55)	38.00 (2.75)	7.28 (1.33)
Model 3							
100	24.04 (2.34)	29.30 (2.13)	29.55 (2.24)	29.80 (3.19)	30.52 (2.23)	28.81 (2.55)	21.46 (1.95)
200	25.01 (2.18)	29.23 (2.02)	29.39 (2.12)	29.74 (4.00)	31.80 (2.20)	34.88 (2.49)	21.76 (1.98)
400	25.61 (3.08)	29.27 (2.29)	29.57 (2.30)	29.79 (4.41)	33.26 (2.77)	48.00 (2.79)	21.70 (2.25)
800	25.92 (2.59)	28.88 (2.12)	29.08 (2.27)	29.11 (2.09)	35.50 (2.51)	39.43 (2.61)	21.60 (2.52)

Table 3. Support recovery of $\Omega\delta$ for Model 3. Standard deviations are given in parentheses

p	100	200	400	800
	Normal distribution			
POS	21.92 (9.75)	26.39 (17.88)	23.06 (14.14)	25.67 (16.96)
TPOS	8.47 (0.36)	8.14 (0.36)	8.25 (0.36)	8.36 (0.36)
TPR	0.77 (0.11)	0.74 (0.11)	0.75 (0.11)	0.76 (0.11)
FPR	0.15 (0.11)	0.10 (0.10)	0.04 (0.04)	0.02 (0.02)
	t_5 distribution			
POS	18.69 (10.17)	23.97 (17.65)	22.98 (16.96)	21.48 (14.90)
TPOS	7.15 (0.33)	7.15 (0.40)	7.15 (0.36)	7.05 (0.43)
TPR	0.65 (0.10)	0.65 (0.12)	0.65 (0.11)	0.64 (0.13)
FPR	0.13 (0.11)	0.09 (0.09)	0.04 (0.04)	0.02 (0.02)

$$\text{FPR} = \frac{\#\{i: \hat{\beta}_i \neq 0 \text{ and } (\Omega\delta)_i = 0\}}{\#\{i: (\Omega\delta)_i = 0\}}.$$

The simulation results are summarized in Table 3. We can see that our method leads to a sparse solution $\hat{\beta}$. It correctly recovers more than eight nonzero elements in the normal distribution case and more than seven nonzero elements in the t_5 distribution case on average. Note that FPR is low, and thus most of zero positions can be recovered by $\hat{\beta}$.

Finally, we carry out a simulation study to investigate the accuracy between the tuning parameter $\hat{\lambda}$ chosen by CV and the optimal value λ_{opt} which minimizes the misclassification rate for the test samples. The results are stated in Table 4 for Models 1–3 with the multivariate normal distribution. It can be seen that the value $\hat{\lambda}$ chosen by CV and the optimal choice λ_{opt} are quite close. Additional simulation results show that the performance of the LPD rule using $\hat{\lambda}$ is similar to that using the optimal choice λ_{opt} .

5.3 Real Data Analysis

In addition to the simulation results presented above, we also apply the LPD classifier to the analysis of two real datasets, one from a lung cancer study (Gordon et al. 2002) and another from a leukemia study (Golub et al. 1999) to further examine the performance of the LPD rule. The lung cancer dataset is available at <http://www.chest Surg.org> and the leukemia dataset is available at <http://www.broadinstitute.org/cgi-bin/cancer/datasets.cgi>.

5.3.1 Lung Cancer Data. The lung cancer dataset in the article by Gordon et al. (2002) consists of 181 tissue samples and each sample is described by 12,533 genes. Among the 181 tissue samples, there are two classes of tissue samples including

31 malignant pleural mesothelioma (MPM) and 150 adenocarcinoma (ADCA). Distinguishing MPM from ADCA is important and challenging from both clinical and pathological perspectives. This dataset has been analyzed by Fan and Fan (2008) using FAIR and NSC. In this section we apply the LPD rule to this dataset for disease classification.

The sample variances of the genes range over a wide interval. After rescaling by a factor of 10^4 , there are 165 genes with the sample variances larger than 10^2 and 41 genes with the sample variances smaller than 10^{-2} . See Figure 1 for a plot of the sorted sample variances. To ensure numerical stability, we drop these 206 genes to control the condition number of $\hat{\Sigma}_\rho$ so that the numerical solution of $\hat{\Sigma}_\rho^{-1}$ is accurate. We use 32 training samples with 16 from MPM and 16 from ADCA. The rest, 149 samples with 15 from MPM and 134 from ADCA, are used for testing. To reduce the computational costs, only 3000 genes with the largest absolute values of the two sample t statistics are used. The classification result is satisfactory, although only 3000 genes are used. Twofold cross-validation method is used for choosing the tuning parameter λ_n . The resulting estimate of $\Omega\delta$ contains 369 nonzero elements, which is about 12.3% of all elements. Classification results are summarized in Table 5. The LPD rule classifies all of 149 testing samples correctly. In contrast, the naive-Bayes rule misclassifies 12 of 149 testing samples and GLDA misclassifies 7 of 149 testing samples. Fan and Fan (2008) reported a test error rate of 7/149 for FAIR and a test error rate of 11/149 for NSC proposed by Tibshirani et al. (2002).

5.3.2 Leukemia Data. The leukemia dataset in the work of Golub et al. (1999) consists of 72 tissue samples, which were all from acute leukemia patients, either acute lymphoblastic leukemia (ALL) or acute myelogenous leukemia (AML). Each sample is described by 7129 genes. Distinguishing ALL from

Table 4. Average of $\hat{\lambda}$ and λ_{opt} (SD)

p	Model 1		Model 2		Model 3	
	$\hat{\lambda}$	λ_{opt}	$\hat{\lambda}$	λ_{opt}	$\hat{\lambda}$	λ_{opt}
100	0.13 (0.05)	0.12 (0.05)	0.14 (0.09)	0.11 (0.07)	0.18 (0.02)	0.14 (0.06)
200	0.15 (0.05)	0.14 (0.04)	0.13 (0.04)	0.11 (0.04)	0.19 (0.02)	0.17 (0.06)
400	0.20 (0.05)	0.18 (0.05)	0.19 (0.06)	0.18 (0.08)	0.24 (0.05)	0.21 (0.05)
800	0.23 (0.05)	0.20 (0.04)	0.17 (0.06)	0.15 (0.03)	0.26 (0.05)	0.24 (0.04)

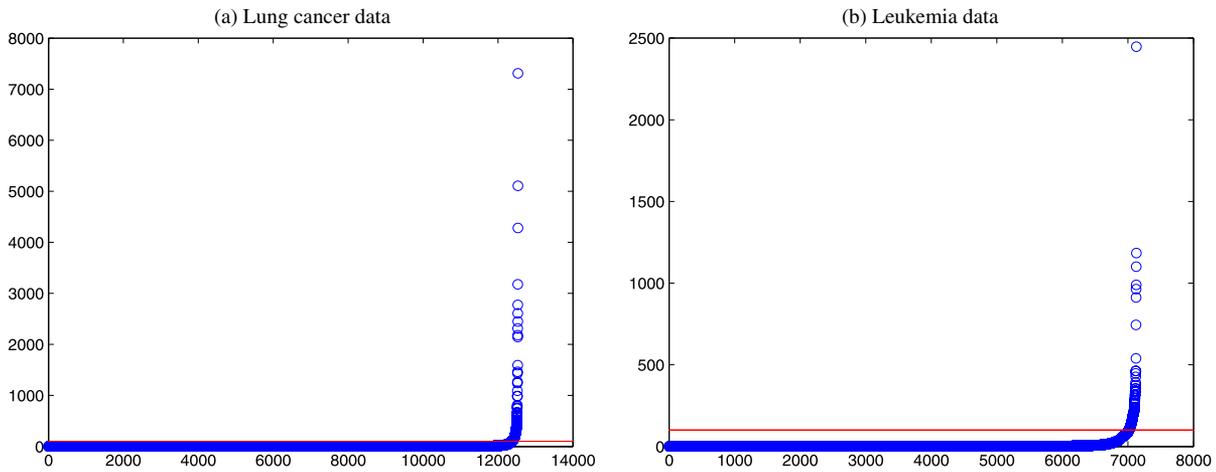


Figure 1. The diagonal of the sample covariance matrix. The online version of this figure is in color.

AML is critical for a successful treatment. The dataset has been analyzed by Fan and Fan (2008). In this section, we apply the LPD rule to this dataset and compare the classification results with those obtained by Fan and Fan (2008) using FAIR and NSC.

As in the analysis of the lung cancer data, we first drop 129 genes with extreme sample variances, either larger than 10^2 or smaller than 10^{-2} after rescale by a factor of 10^5 . See Figure 1. Among the 72 tissue samples, there are 38 training samples (27 in class ALL and 11 in class AML) and 34 test samples (20 in class ALL and 14 in class AML). Similarly to the analysis of the lung cancer data, to control the computational costs, we only use 3000 genes with the largest absolute values of the two sample t statistics. The estimate of $\Omega\delta$ contains 206 nonzero elements, which is about 6.87% of all elements. The classification results are summarized in Table 6. The LPD rule only misclassifies 1 of the 34 testing samples and makes 0 training error. In comparison, the naive-Bayes rule misclassifies 7 of 34 testing samples and 1 of 38 training samples. GLDA misclassifies 3 of 34 testing samples and 1 of 38 training samples. From the work of Fan and Fan (2008), FAIR makes 1/34 test error rate and 1/38 training error rate, and NSC makes 3/34 test error rate and 1/38 training error rate.

6. DISCUSSIONS

In this article we introduced the LPD rule for sparse linear discriminant analysis of high-dimensional data. The LPD classifier is based on the direct estimation of the product $\Omega\delta$ through constrained ℓ_1 minimization which can be implemented efficiently using linear programming. The classifier has desirable theoretical and numerical properties and performs well in the real data analysis.

The LPD rule exploits the approximate sparsity of $\Omega\delta$ which can be estimated more efficiently than Ω can. The sparsity of $\Omega\delta$

can be viewed as a relaxation of the conventional assumption on the sparsity of both Ω and δ . In certain settings, $\Omega\delta$ can be well estimated even when Ω is not estimable consistently which leads to the failure of some conventional classification methods. An interesting consequence is that the LPD classifier can still perform well even when Ω cannot be estimated well. This is a major advantage of the LPD rule over classification methods that are based on separate and good estimates of Ω (or Σ) and δ . Furthermore, as shown both in the theoretical results (Theorems 2 and 3 only require conditions on $|\Omega\delta|_1$ and not on $|\Omega\delta|_0$) and in the simulation (in Models 1 and 2 nearly all of the elements of $\Omega\delta$ are nonzero), only approximate sparsity of $\Omega\delta$ is needed in order for the LPD rule to perform well.

In this article we have focused on the case where the new observation \mathbf{Z} has equal prior probabilities of belonging to either class 1 or class 2. The procedure can be extended easily to the case of unequal prior probabilities π_1 and π_2 . In this case, we can define the LPD rule by classifying \mathbf{Z} to class 1 if and only if

$$(\mathbf{Z} - \hat{\boldsymbol{\mu}})' \hat{\boldsymbol{\beta}} \geq \log(\pi_2/\pi_1).$$

When the unequal prior probabilities π_1 and π_2 are unknown, we can simply estimate them by $\hat{\pi}_1 = n_1/n$ and $\hat{\pi}_2 = n_2/n$, respectively. The LPD rule also can be directly extended to multi-group classification problems. Suppose there are K classes with distributions $N(\boldsymbol{\mu}_k, \boldsymbol{\Sigma})$ for $1 \leq k \leq K$. In the ideal setting where all the parameters are known, the oracle rule classifies \mathbf{Z} to class k if and only if

$$(\mathbf{Z} - \boldsymbol{\mu}_{kl})' \boldsymbol{\Omega} \boldsymbol{\delta}_{kl} \geq 0 \quad \text{for all } l \neq k,$$

where $\boldsymbol{\delta}_{kl} = \boldsymbol{\mu}_k - \boldsymbol{\mu}_l$ and $\boldsymbol{\mu}_{kl} = (\boldsymbol{\mu}_k + \boldsymbol{\mu}_l)/2$. When the parameters are unknown and random samples from the distributions are available, the products $\boldsymbol{\Omega} \boldsymbol{\delta}_{kl}$ can then be estimated by solving a similar linear programming as in (3) and an LPD classifier can be constructed accordingly.

Table 5. Classification error of lung cancer data by various methods

	LPAD	FAIR	NSC	Naive-LDA	GLDA
Training error	0/32	0/32	0/32	0/32	0/32
Testing error	0/149	7/149	11/149	12/149	7/149

Table 6. Classification error of leukemia data by various methods

	LPAD	FAIR	NSC	Naive-LDA	GLDA
Training error	0/38	1/38	1/38	1/38	1/38
Testing error	1/34	1/34	3/34	7/34	3/34

6.1 Consequence of Feature Selection on Classification

In the high-dimensional setting, it is a common practice to exploit sparsity by first selecting a small number of important features (typically by thresholding) and then making inference based only on the selected features. A common aspect of these methods is that the correlations between the variables are ignored. It should be noted that these methods are inefficient in general even when the zero features are known in advance and all the important features are selected correctly. The main reason is that those “unimportant” features are in fact useful and even potentially important for classification because of the correlations. This can be seen as follows by considering the oracle rules in various settings.

Let us first consider the oracle independence rule which classifies \mathbf{Z} into class 1 if and only if $(\mathbf{Z} - \boldsymbol{\mu})'\mathbf{D}^{-1}\boldsymbol{\delta} \geq 0$, where $\mathbf{D} = \text{diag}(\boldsymbol{\Sigma})$. It is easy to see that the misclassification rate of the oracle independence rule is $1 - \Phi(\frac{1}{2}\Upsilon_p)$, where

$$\Upsilon_p = \frac{\boldsymbol{\delta}'\mathbf{D}^{-1}\boldsymbol{\delta}}{(\boldsymbol{\delta}'\mathbf{D}^{-1}\boldsymbol{\Sigma}\mathbf{D}^{-1}\boldsymbol{\delta})^{1/2}} \quad \text{with } \mathbf{D} = \text{diag}(\boldsymbol{\Sigma}). \quad (17)$$

It can be shown that Fisher’s rule based only on the important features outperforms the oracle independence rule. Write

$$\boldsymbol{\delta} = \begin{pmatrix} \boldsymbol{\delta}_1 \\ \boldsymbol{\delta}_2 \end{pmatrix} \quad \text{and} \quad \boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}'_{12} \\ \boldsymbol{\Sigma}_{12} & \boldsymbol{\Sigma}_{22} \end{pmatrix}, \quad (18)$$

where $\boldsymbol{\delta}_1$ is a k_1 -dimensional vector, $\boldsymbol{\Sigma}_{11}$ is $k_1 \times k_1$, $\boldsymbol{\Sigma}_{12}$ is $(p - k_1) \times k_1$, and $\boldsymbol{\Sigma}_{22}$ is $(p - k_1) \times (p - k_1)$. Suppose $\boldsymbol{\delta}_2$ is known to be $\mathbf{0}$. Then it follows from (5) that the oracle Fisher’s rule based only on the first k_1 variables has misclassification rate $1 - \Phi(\frac{1}{2}(\boldsymbol{\delta}'_1\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\delta}_1)^{1/2})$. Note that in this case the oracle independence rule only depends on the important features and Υ_p in (17) can be re-expressed as

$$\Upsilon_p = \frac{\boldsymbol{\delta}'_1\mathbf{D}_{11}^{-1}\boldsymbol{\delta}_1}{(\boldsymbol{\delta}'_1\mathbf{D}_{11}^{-1}\boldsymbol{\Sigma}_{11}\mathbf{D}_{11}^{-1}\boldsymbol{\delta}_1)^{1/2}} \quad \text{with } \mathbf{D}_{11} = \text{diag}(\boldsymbol{\Sigma}_{11}).$$

It is easy to verify that

$$\boldsymbol{\delta}'_1\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\delta}_1 = \max_{\mathbf{x} \in \mathbb{R}^{k_1}} \frac{\mathbf{x}'\boldsymbol{\delta}_1\boldsymbol{\delta}'_1\mathbf{x}}{\mathbf{x}'\boldsymbol{\Sigma}_{11}\mathbf{x}}.$$

Thus we have $\boldsymbol{\delta}'_1\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\delta}_1 \geq \Upsilon_p^2$, which implies that the oracle Fisher’s rule based only on the important features (the first k_1 variables) outperforms the independence rule. This shows that, given only the important features are used for classification, the independence rule can be inefficient and correlations among the features should be taken into account.

Although the oracle Fisher’s rule based on the important features is better than the independence rule, it is not an efficient rule itself because ignoring the zero (or “unimportant”) features also leads to inefficiency. Write $\boldsymbol{\delta}$ as in (18) and suppose the fact that $\boldsymbol{\delta}_1 \neq \mathbf{0}$ and $\boldsymbol{\delta}_2 = \mathbf{0}$ is known. We next show that the oracle Fisher’s rule based on all the features outperforms the Fisher’s rule based only on the important features. Note that $\Delta_p = \boldsymbol{\delta}'\boldsymbol{\Omega}\boldsymbol{\delta}$ can be decomposed as follows:

$$\boldsymbol{\delta}'\boldsymbol{\Omega}\boldsymbol{\delta} = \boldsymbol{\delta}'_1\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\delta}_1 + (\boldsymbol{\delta}_2 - \mathbf{B}\boldsymbol{\delta}_1)'\mathbf{W}^{-1}(\boldsymbol{\delta}_2 - \mathbf{B}\boldsymbol{\delta}_1), \quad (19)$$

where $\mathbf{B} = \boldsymbol{\Sigma}_{22}^{-1}\boldsymbol{\Sigma}_{12}$. Note that $\mathbf{W} = \boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{12}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}'_{12}$ is positive definite. Consequently, if $\mathbf{B}\boldsymbol{\delta}_1 \neq \mathbf{0}$, then the last term in (19) is positive and hence $\Delta_p \geq \boldsymbol{\delta}'_1\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\delta}_1$. This means that even

if the fact that $\boldsymbol{\delta}_1 \neq \mathbf{0}$ and $\boldsymbol{\delta}_2 = \mathbf{0}$ is known in advance, dropping the zero features would lead to inefficiency because of the correlations among all the features. Therefore classifiers based only on the important features are not efficient in general.

The above analysis shows that ignoring the correlations and feature selections in general lead to inefficient classifiers. A better alternative is to construct a classification rule taking into account all the features and their correlations. This analysis makes the LPD rule even more attractive in the ultra-high-dimensional case where p is very large. In such a setting estimating the full precision matrix $\boldsymbol{\Omega}$ well is very difficult if not impossible. In contrast, it is relatively easy to estimate the vector $\boldsymbol{\Omega}\boldsymbol{\delta}$ directly.

7. PROOFS

Throughout this section, we denote by C, C_1, C_2, \dots generic constants which may vary from place to place. We shall omit the proof of Theorem 1 as it follows directly from Theorem 2; see Remark 3 in Section 3. Before proving the other main theorems, we first collect some technical lemmas. The following lemma is an exponential inequality from the article of Cai and Liu (2011a). The proof also can be found in the report by Cai and Liu (2011b).

Lemma 1. Let ξ_1, \dots, ξ_n be independent random variables with mean zero. Suppose that there exist some $t > 0$ and \bar{B}_n such that

$$\sum_{k=1}^n \mathbb{E}\xi_k^2 e^{t|\xi_k|} \leq \bar{B}_n^2.$$

Then uniformly for $0 < x \leq \bar{B}_n$ and $n \geq 1$,

$$\mathbb{P}\left(\sum_{k=1}^n \xi_k \geq C_t \bar{B}_n x\right) \leq \exp(-x^2), \quad (20)$$

where $C_t = t + t^{-1}$.

The next lemma shows that the true $\boldsymbol{\Omega}\boldsymbol{\delta}$ belongs to the feasible set of (3) with high probability.

Lemma 2. (i) Under (C2) and (C3), we have with probability greater than $1 - O(p^{-1})$,

$$|\hat{\boldsymbol{\Sigma}}_n \boldsymbol{\Omega}\boldsymbol{\delta} - (\bar{\mathbf{X}} - \bar{\mathbf{Y}})|_\infty \leq \lambda_n. \quad (21)$$

(ii) Under (C2) and (C4), (21) holds with probability greater than $1 - O(p^{-1} + n^{-\epsilon/8})$.

Proof. We only prove the lemma under (C3). The proof under (C4) is similar by using a truncation technique as in the work of Cai, Liu, and Luo (2011). The details are given in the report by Cai and Liu (2011b). Write $\mathbf{X}_k = \boldsymbol{\mu}_1 + \mathbf{U}_{k1}$ and $\mathbf{Y}_k = \boldsymbol{\mu}_2 + \mathbf{U}_{k2}$. Set $\bar{\mathbf{X}} = (\bar{X}_1, \dots, \bar{X}_p)'$ and $\boldsymbol{\mu}_1 = (\mu_1, \dots, \mu_p)'$. By Lemma 1, we have for any $M > 0$, there exists some $C_1 > 0$ such that

$$\max_{1 \leq i \leq p} \mathbb{P}\left(|\bar{X}_i - \mu_i| \geq C_1 \sqrt{\frac{\sigma_{ii} \log p}{n}}\right) \leq 2p^{-M} \quad (22)$$

and

$$\mathbb{P}\left(|(\bar{\mathbf{X}} - \boldsymbol{\mu}_1)'\boldsymbol{\Omega}\boldsymbol{\delta}| \geq C_1 \sqrt{\frac{\Delta_p \log p}{n}}\right) \leq 2p^{-M}. \quad (23)$$

Similar inequalities hold for $\bar{\mathbf{Y}}$ and $\boldsymbol{\mu}_2$. Note that

$$\begin{aligned} \hat{\boldsymbol{\Sigma}}_n &= \frac{1}{n} \left(\sum_{i=1}^{n_1} \mathbf{U}_{i1} \mathbf{U}'_{i1} + \sum_{j=1}^{n_2} \mathbf{U}_{j2} \mathbf{U}'_{j2} \right) \\ &\quad - \frac{n_1}{n} (\bar{\mathbf{U}}_1 - \boldsymbol{\mu}_1) (\bar{\mathbf{U}}_1 - \boldsymbol{\mu}_1)' \\ &\quad - \frac{n_2}{n} (\bar{\mathbf{U}}_2 - \boldsymbol{\mu}_2) (\bar{\mathbf{U}}_2 - \boldsymbol{\mu}_2)' \\ &=: \tilde{\boldsymbol{\Sigma}} - \frac{n_1}{n} (\bar{\mathbf{U}}_1 - \boldsymbol{\mu}_1) (\bar{\mathbf{U}}_1 - \boldsymbol{\mu}_1)' \\ &\quad - \frac{n_2}{n} (\bar{\mathbf{U}}_2 - \boldsymbol{\mu}_2) (\bar{\mathbf{U}}_2 - \boldsymbol{\mu}_2)'. \end{aligned}$$

Thus by (22) and (23), it suffices to show that with probability greater than $1 - O(p^{-1})$,

$$|\tilde{\boldsymbol{\Sigma}} \boldsymbol{\Omega} \boldsymbol{\delta} - (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)|_\infty \leq C \sqrt{\frac{\max_i \sigma_{ii} \Delta_p \log p}{n}}.$$

For brevity, we set $\mathbf{Z}_i = \mathbf{U}_{i1}$ for $1 \leq i \leq n_1$ and $\mathbf{Z}_{i+n_1} = \mathbf{U}_{i2}$ for $1 \leq i \leq n_2$. Note that

$$\tilde{\boldsymbol{\Sigma}} \boldsymbol{\Omega} \boldsymbol{\delta} - (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) = (\tilde{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}) \boldsymbol{\Omega} \boldsymbol{\delta}$$

which can be further written as

$$\frac{1}{n} \sum_{i=1}^n (\mathbf{Z}_i \mathbf{Z}'_i \boldsymbol{\Omega} \boldsymbol{\delta} - \mathbf{E}(\mathbf{Z}_i \mathbf{Z}'_i \boldsymbol{\Omega} \boldsymbol{\delta})).$$

We use Lemma 1 to bound the above partial sums. Write $\mathbf{Z}_i = (Z_{i1}, \dots, Z_{ip})'$. By (C3), we have $\mathbf{E} \exp(t_0 |Z_{ij} \mathbf{Z}'_i \boldsymbol{\Omega} \boldsymbol{\delta}| / (\sigma_{ij} \boldsymbol{\delta}' \boldsymbol{\Omega} \boldsymbol{\delta})^{1/2}) \leq K_0$ for some bounded constants $t_0 > 0$ and $K_0 > 0$. For any constant $\tau > 0$, let $\xi_i = (Z_{ij} \mathbf{Z}'_i \boldsymbol{\Omega} \boldsymbol{\delta} - \mathbf{E} Z_{ij} \mathbf{Z}'_i \boldsymbol{\Omega} \boldsymbol{\delta}) / (\sigma_{ij} \boldsymbol{\delta}' \boldsymbol{\Omega} \boldsymbol{\delta})^{1/2}$ in Lemma 1 and $\bar{B}_n^2 = c\tau n$ with some large constant c depending on τ, t_0, K_0 . Then we can get that for any $\tau > 0$, there exists some constant $C_2 > 0$ depending only on c, τ, t_0 , and K_0 ,

$$\begin{aligned} &\max_{1 \leq j \leq p} \mathbf{P} \left(\left| \sum_{i=1}^n \xi_i \right| \geq C_2 \sqrt{n \log p} \right) \\ &\leq \max_{1 \leq j \leq p} \mathbf{P} \left(\left| \sum_{i=1}^{n_1} \xi_i \right| \geq 2^{-1} C_2 \sqrt{n \log p} \right) \\ &\quad + \max_{1 \leq j \leq p} \mathbf{P} \left(\left| \sum_{i=n_1+1}^n \xi_i \right| \geq 2^{-1} C_2 \sqrt{n \log p} \right) \\ &\leq 4p^{-\tau}. \end{aligned}$$

This implies that for any $\tau > 0$,

$$\mathbf{P} \left(|\tilde{\boldsymbol{\Sigma}} \boldsymbol{\Omega} \boldsymbol{\delta} - (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)|_\infty \geq C_2 \sqrt{\frac{\max_i \sigma_{ii} \Delta_p \log p}{n}} \right) \leq 4p^{-\tau+1}.$$

Lemma 2 is proved.

We are now ready to prove Theorems 2–5. Throughout the proof, we assume (21), $|\hat{\boldsymbol{\delta}} - \boldsymbol{\delta}|_\infty \leq C \sqrt{\log p/n}$, $|\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}|_\infty \leq C \sqrt{\log p/n}$, and $|\hat{\boldsymbol{\Sigma}}_n - \boldsymbol{\Sigma}|_\infty \leq C \sqrt{\log p/n}$ for some large constant $C > 0$. The above four inequalities hold with probability greater than $1 - O(p^{-1})$ or $1 - O(p^{-1} + n^{-\epsilon/8})$ under (C3) or (C4), respectively.

Proof of Theorems 2 and 5(i). By the definition of $\hat{\boldsymbol{\beta}}$, we have

$$\begin{aligned} |(\boldsymbol{\Omega} \boldsymbol{\delta})' \hat{\boldsymbol{\Sigma}}_n \hat{\boldsymbol{\beta}} - (\boldsymbol{\Omega} \boldsymbol{\delta})' \boldsymbol{\delta}| &\leq \lambda_n |\boldsymbol{\Omega} \boldsymbol{\delta}|_1 + |\hat{\boldsymbol{\delta}} - \boldsymbol{\delta}|_\infty |\boldsymbol{\Omega} \boldsymbol{\delta}|_1 \\ &\leq 2\lambda_n |\boldsymbol{\Omega} \boldsymbol{\delta}|_1. \end{aligned} \tag{24}$$

By (21), we have

$$|(\boldsymbol{\Omega} \boldsymbol{\delta})' \hat{\boldsymbol{\Sigma}}_n \hat{\boldsymbol{\beta}} - \boldsymbol{\delta}' \hat{\boldsymbol{\beta}}| \leq \lambda_n |\hat{\boldsymbol{\beta}}|_1 + |\hat{\boldsymbol{\delta}} - \boldsymbol{\delta}|_\infty |\hat{\boldsymbol{\beta}}|_1 \leq 2\lambda_n |\boldsymbol{\Omega} \boldsymbol{\delta}|_1,$$

which together with (24) implies that

$$|(\hat{\boldsymbol{\beta}} - \boldsymbol{\Omega} \boldsymbol{\delta})' \boldsymbol{\delta}| \leq 4\lambda_n |\boldsymbol{\Omega} \boldsymbol{\delta}|_1. \tag{25}$$

Thus we have

$$\begin{aligned} \left| (\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_1)' \hat{\boldsymbol{\beta}} + \frac{1}{2} \boldsymbol{\delta}' \boldsymbol{\Omega} \boldsymbol{\delta} \right| &\leq |(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu})' \hat{\boldsymbol{\beta}}| + \frac{1}{2} |\boldsymbol{\delta}' \hat{\boldsymbol{\beta}} - \boldsymbol{\delta}' \boldsymbol{\Omega} \boldsymbol{\delta}| \\ &\leq |(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu})' \hat{\boldsymbol{\beta}}| + 2\lambda_n |\boldsymbol{\Omega} \boldsymbol{\delta}|_1 \\ &\leq C \sqrt{\frac{\log p}{n}} |\boldsymbol{\Omega} \boldsymbol{\delta}|_1 + 2\lambda_n |\boldsymbol{\Omega} \boldsymbol{\delta}|_1. \end{aligned} \tag{26}$$

Similarly,

$$\left| (\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_2)' \hat{\boldsymbol{\beta}} - \frac{1}{2} \boldsymbol{\delta}' \boldsymbol{\Omega} \boldsymbol{\delta} \right| \leq C \sqrt{\frac{\log p}{n}} |\boldsymbol{\Omega} \boldsymbol{\delta}|_1 + 2\lambda_n |\boldsymbol{\Omega} \boldsymbol{\delta}|_1.$$

We next consider the denominator in R_n . We have

$$|\boldsymbol{\Sigma} \hat{\boldsymbol{\beta}} - \boldsymbol{\delta}|_\infty \leq |\boldsymbol{\Sigma} \hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\Sigma}}_n \hat{\boldsymbol{\beta}}|_\infty + 2\lambda_n \leq C |\boldsymbol{\Omega} \boldsymbol{\delta}|_1 \sqrt{\frac{\log p}{n}} + 2\lambda_n.$$

Therefore

$$|\hat{\boldsymbol{\beta}}' \boldsymbol{\Sigma} \hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}' \boldsymbol{\delta}| \leq C |\boldsymbol{\Omega} \boldsymbol{\delta}|_1^2 \sqrt{\frac{\log p}{n}} + 2\lambda_n |\boldsymbol{\Omega} \boldsymbol{\delta}|_1.$$

By (25), we have

$$|\hat{\boldsymbol{\beta}}' \boldsymbol{\Sigma} \hat{\boldsymbol{\beta}} - \boldsymbol{\delta}' \boldsymbol{\Omega} \boldsymbol{\delta}| \leq C |\boldsymbol{\Omega} \boldsymbol{\delta}|_1^2 \sqrt{\frac{\log p}{n}} + 6\lambda_n |\boldsymbol{\Omega} \boldsymbol{\delta}|_1. \tag{27}$$

Suppose that $\boldsymbol{\delta}' \boldsymbol{\Omega} \boldsymbol{\delta} \geq M$ for some $M > 0$. By (8), (26), and (27), we have

$$\left| \frac{(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_1)' \hat{\boldsymbol{\beta}}}{\sqrt{\hat{\boldsymbol{\beta}}' \boldsymbol{\Sigma} \hat{\boldsymbol{\beta}}}} \right| \geq C \left| \frac{\boldsymbol{\delta}' \boldsymbol{\Omega} \boldsymbol{\delta}}{\sqrt{\hat{\boldsymbol{\beta}}' \boldsymbol{\Sigma} \hat{\boldsymbol{\beta}}}} \right| \geq C (\Delta_p^{-1} + o(1))^{-1/2} \geq CM^{1/2}.$$

This inequality implies that

$$|R_n - R| \leq \exp(-CM). \tag{28}$$

Suppose that $\boldsymbol{\delta}' \boldsymbol{\Omega} \boldsymbol{\delta} \leq M$. By (8) and (27), we have

$$\left| \frac{\hat{\boldsymbol{\beta}}' \boldsymbol{\Sigma} \hat{\boldsymbol{\beta}}}{\boldsymbol{\delta}' \boldsymbol{\Omega} \boldsymbol{\delta}} - 1 \right| = o(1). \tag{29}$$

This together with (26) yields that

$$\left| \frac{(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_1)' \hat{\boldsymbol{\beta}}}{\sqrt{\hat{\boldsymbol{\beta}}' \boldsymbol{\Sigma} \hat{\boldsymbol{\beta}}}} + \frac{(1/2) \boldsymbol{\delta}' \boldsymbol{\Omega} \boldsymbol{\delta}}{\sqrt{\hat{\boldsymbol{\beta}}' \boldsymbol{\Sigma} \hat{\boldsymbol{\beta}}}} \right| \leq C \frac{|\boldsymbol{\Omega} \boldsymbol{\delta}|_1}{(\boldsymbol{\delta}' \boldsymbol{\Omega} \boldsymbol{\delta})^{1/2}} \lambda_n. \tag{30}$$

By (27) and some simple calculations,

$$\begin{aligned} & \left| \frac{1}{\sqrt{\hat{\beta}'\Sigma\hat{\beta}}} - \frac{1}{\sqrt{\delta'\Omega\delta}} \right| \\ & \leq \frac{C|\Omega\delta|_1^2\sqrt{\log p/n} + 6|\Omega\delta|_1\lambda_n}{\sqrt{\hat{\beta}'\Sigma\hat{\beta}}\sqrt{\delta'\Omega\delta}(\sqrt{\hat{\beta}'\Sigma\hat{\beta}} + \sqrt{\delta'\Omega\delta})} \\ & \leq C(\delta'\Omega\delta)^{-3/2} \left(|\Omega\delta|_1^2\sqrt{\frac{\log p}{n}} + |\Omega\delta|_1\lambda_n \right) \end{aligned} \quad (31)$$

and

$$\begin{aligned} & \left| \frac{(1/2)\delta'\Omega\delta}{\sqrt{\hat{\beta}'\Sigma\hat{\beta}}} - \frac{1}{2}(\delta'\Omega\delta)^{1/2} \right| \\ & \leq C\frac{|\Omega\delta|_1^2}{(\delta'\Omega\delta)^{1/2}}\sqrt{\frac{\log p}{n}} + C\frac{|\Omega\delta|_1}{(\delta'\Omega\delta)^{1/2}}\lambda_n \\ & =: r_n. \end{aligned} \quad (32)$$

Note that by (12), (30), and (32),

$$R_n = R \times (1 + O(1)r_n(\delta'\Omega\delta)^{1/2} \exp(O(1)(\delta'\Omega\delta)^{1/2}r_n)). \quad (33)$$

By the condition (8) and the assumption $\delta'\Omega\delta \leq M$, we have $(|\Omega\delta|_1 + |\Omega\delta|_1^2)\sqrt{\log p/n} = o(1)$. Thus $R_n = (1 + o(1))R$. This together with (28) proves the theorems by letting $n, p \rightarrow \infty$ first and then $M \rightarrow \infty$.

Proof of Theorems 3 and 5(ii). Under the conditions of Theorem 3 or 5(ii), we have (29)–(32) hold without assuming $\delta'\Omega\delta \leq M$. So Theorems 3 and 5(ii) follow from (33) and the fact $(\delta'\Omega\delta)^{1/2}r_n = o(1)$ immediately.

Proof of Theorem 4. We shall prove a better rate for $|\hat{\beta}'\Sigma\hat{\beta} - \delta'\Omega\delta|$ under the condition (10). We have

$$\begin{aligned} |\Sigma(\hat{\beta} - \Omega\delta)|_\infty & \leq |\hat{\Sigma}_n(\hat{\beta} - \Omega\delta)|_\infty + |(\hat{\Sigma}_n - \Sigma)(\hat{\beta} - \Omega\delta)|_\infty \\ & \leq 2\lambda_n + C|\hat{\beta} - \Omega\delta|_1\sqrt{\frac{\log p}{n}} \\ & \leq 2\lambda_n + C|\Omega\delta|_0\sqrt{\frac{\log p}{n}}|\hat{\beta} - \Omega\delta|_\infty \\ & \leq 2\lambda_n + C\|\Omega\|_{L_1}|\Omega\delta|_0\sqrt{\frac{\log p}{n}}|\Sigma(\hat{\beta} - \Omega\delta)|_\infty. \end{aligned}$$

This together with $\|\Omega\|_{L_1}|\Omega\delta|_0\sqrt{\frac{\log p}{n}} = o(1)$ implies that $|\Sigma(\hat{\beta} - \Omega\delta)|_\infty \leq C\lambda_n$. Thus we have

$$|\hat{\beta}'\Sigma\hat{\beta} - \hat{\beta}'\Sigma\Omega\delta| \leq C|\Omega\delta|_1\lambda_n$$

and

$$|\hat{\beta}'\Sigma\Omega\delta - \delta'\Omega\delta| \leq C|\Omega\delta|_1\lambda_n.$$

The remaining steps follow from the proof of (33).

[Received March 2011. Revised June 2011.]

REFERENCES

Anderson, T. W. (2003), *An Introduction to Multivariate Statistical Analysis* (3rd ed.), New Jersey: Wiley-Interscience. [1566,1568]

Becker, S., Candès, E., and Grant, M. (2011), “Templates for Convex Cone Problems With Applications to Sparse Signal Recovery,” *Mathematical Programming Computation*, 3, 165–218. [1571]

Bickel, P., and Levina, L. (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. [1566]

Boyd, S., and Vandenberghe, L. (2004), *Convex Optimization*, Cambridge: Cambridge University Press. [1571]

Cai, T., and Liu, W. (2011a), “Adaptive Thresholding for Sparse Covariance Matrix Estimation,” *Journal of the American Statistical Association*, 106, 672–684. [1575]

——— (2011b), Supplementary Material of “A Direct Estimation Approach to Sparse Linear Discriminant Analysis,” technical report, University of Pennsylvania. [1575]

Cai, T., and Zhou, H. (2010), “Optimal Rates of Convergence for Sparse Covariance Matrix Estimation,” technical report, University of Pennsylvania. [1569]

Cai, T., Liu, W., and Luo, X. (2011), “A Constrained l_1 Minimization Approach to Sparse Precision Matrix Estimation,” *Journal of the American Statistical Association*, 106, 594–607. [1567,1571,1575]

Candès, E., and Tao, T. (2007), “The Dantzig Selector: Statistical Estimation When p Is Much Larger Than n ,” *The Annals of Statistics*, 35, 2313–2351. [1568,1571]

Donoho, D., Elad, M., and Temlyakov, V. (2006), “Stable Recovery of Sparse Overcomplete Representations in the Presence of Noise,” *IEEE Transactions on Information Theory*, 52, 6–18. [1568]

Fan, J., and Fan, Y. (2008), “High Dimensional Classification Using Features Annealed Independence Rules,” *The Annals of Statistics*, 36, 2605–2637. [1566,1571,1573,1574]

Fang, K. T., and Anderson, T. W. (1990), *Statistical Inference in Elliptically Contoured and Related Distributions*, New York: Allerton Press Inc. [1570]

Friedman, J. H. (1989), “Regularized Discriminant Analysis,” *Journal of the American Statistical Association*, 84, 165–175. [1567]

Golub, T. R., Slonim, D. K., Tamayo, P., Huard, C., Gaasenbeek, M., Mesirov, J. P., Coller, H., Loh, M. L., Downing, J. R., Caligiuri, M. A., Bloomfield, C. D., and Lander, E. S. (1999), “Molecular Classification of Cancer: Class Discovery and Class Prediction by Gene Expression Monitoring,” *Science*, 286, 531–537. [1573]

Gordon, G. J., Jensen, R. V., Hsiao, L. L., Gullans, S. R., Blumenstock, J. E., Ramaswamy, S., Richards, W. G., Sugarbaker, D. J., and 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. [1573]

Guo, Y., Hastie, T., and Tibshirani, R. (2007), “Regularized Linear Discriminant Analysis and Its Application in Microarrays,” *Biostatistics*, 8, 86–100. [1567]

Hall, P., Titterton, D. M., and Xue, J. H. (2009), “Median-Based Classifiers for High-Dimensional Data,” *Journal of the American Statistical Association*, 104, 1597–1608. [1567]

Ravikumar, P., Wainwright, M., Raskutti, G., and Yu, B. (2008), “High-Dimensional Covariance Estimation by Minimizing l_1 -Penalized Log-Determinant Divergence,” Technical Report 797, UC Berkeley, Statistics Dept. [1567]

Rothman, A., Bickel, P., Levina, E., and Zhu, J. (2008), “Sparse Permutation Invariant Covariance Estimation,” *Electronic Journal of Statistics*, 2, 494–515. [1567]

Shao, J., Wang, Y., Deng, X., and Wang, S. (2011), “Sparse Linear Discriminant Analysis With High Dimensional Data,” *The Annals of Statistics*, 39, 1241–1265. [1566,1568-1571]

Tibshirani, R., Hastie, T., Narasimhan, B., and Chu, G. (2002), “Diagnosis of Multiple Cancer Types by Shrunken Centroids of Gene Expression,” *Proceedings of the National Academy of Sciences of the United States of America*, 99, 6567–6572. [1567,1570,1571,1573]

Witten, D., and Tibshirani, R. (2009), “Covariance-Regularized Regression and Classification for High Dimensional Problems,” *Journal of the Royal Statistical Society, Ser. B*, 71, 615–636. [1567]

Wu, M. C., Zhang, L., Wang, Z., Christiani, D. C., and 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. [1567]

Yuan, M. (2009), “Sparse Inverse Covariance Matrix Estimation via Linear Programming,” *Journal of Machine Learning Research*, 11, 2261–2286. [1567,1571]