

Neyman-Pearson classification: parametrics and sample size requirement

Xin Tong

XINT@MARSHALL.USC.EDU

*Department of Data Sciences and Operations
Marshall Business School
University of Southern California*

Lucy Xia

LUCYXIA@UST.HK

*Department of ISOM
School of Business and Management
Hong Kong University of Science and Technology*

Jiacheng Wang

JIACHENGWANG@GALTON.UCHICAGO.EDU

*Department of Statistics
University of Chicago*

Yang Feng

YANG.FENG@NYU.EDU

*Department of Biostatistics
School of Global Public Health
New York University*

Editor: Xiaotong Shen

Abstract

The Neyman-Pearson (NP) paradigm in binary classification seeks classifiers that achieve a minimal type II error while enforcing the prioritized type I error controlled under some user-specified level α . This paradigm serves naturally in applications such as severe disease diagnosis and spam detection, where people have clear priorities among the two error types. Recently, Tong et al. (2018) proposed a nonparametric umbrella algorithm that adapts all scoring-type classification methods (e.g., logistic regression, support vector machines, random forest) to respect the given type I error (i.e., conditional probability of classifying a class 0 observation as class 1 under the 0-1 coding) upper bound α with high probability, without specific distributional assumptions on the features and the responses. Universal the umbrella algorithm is, it demands an explicit minimum sample size requirement on class 0, which is often the more scarce class, such as in rare disease diagnosis applications. In this work, we employ the parametric linear discriminant analysis (LDA) model and propose a new parametric thresholding algorithm, which does not need the minimum sample size requirements on class 0 observations and thus is suitable for small sample applications such as rare disease diagnosis. Leveraging both the existing nonparametric and the newly proposed parametric thresholding rules, we propose four LDA-based NP classifiers, for both low- and high-dimensional settings. On the theoretical front, we prove NP oracle inequalities for one proposed classifier, where the rate for excess type II error benefits from the explicit parametric model assumption. Furthermore, as NP classifiers involve a sample splitting step of class 0 observations, we construct a new adaptive sample splitting scheme that can be applied universally to NP classifiers, and this adaptive strategy reduces the

type II error of these classifiers. The proposed NP classifiers are implemented in the R package `nproc`.

Keywords: classification, asymmetric error, Neyman-Pearson (NP) paradigm, NP oracle inequalities, minimum sample size requirement, linear discriminant analysis (LDA), NP umbrella algorithm, adaptive splitting

1. Introduction

Classification aims to predict discrete outcomes (i.e., class labels) for new observations, using algorithms trained on labeled data. It is one of the most studied machine learning problems with applications including automatic disease diagnosis, email spam filters, and image classification. Binary classification, where the outcomes belong to one of two classes and the class labels are usually coded as $\{0, 1\}$ (or $\{-1, 1\}$ or $\{1, 2\}$), is the most common type. Most binary classifiers are constructed to minimize the overall classification error (i.e., *risk*), which is a weighted sum of type I and type II errors. Here, *type I error* is defined as the conditional probability of misclassifying a class 0 observation as class 1, and *type II error* is the conditional probability of misclassifying a class 1 observation as class 0¹. In the following, we refer to this paradigm as the *classical classification paradigm*. Along this line, numerous methods have been proposed, including linear discriminant analysis (LDA) in both low and high dimensions (Guo et al., 2005; Cai and Liu, 2011; Shao et al., 2011; Witten and Tibshirani, 2012; Fan et al., 2012; Mai et al., 2012), logistic regression, support vector machine (SVM) (Vapnik, 1999), random forest (Breiman, 2001), among others.

In contrast, the *Neyman-Pearson (NP) classification paradigm* (Cannon et al., 2002; Scott and Nowak, 2005; Rigollet and Tong, 2011; Tong, 2013; Zhao et al., 2016; Tong et al., 2016) was developed to seek a classifier that minimizes the type II error while maintaining the type I error below a user-specified level α , usually a small value (e.g., 5%). We call this target classifier the NP oracle classifier. The NP paradigm is appropriate in applications such as cancer diagnosis, where a type I error (i.e., misdiagnosing a cancer patient to be healthy) has more severe consequences than a type II error (i.e., misdiagnosing a healthy patient as with cancer). The latter incurs extra medical costs and patients' anxiety but will not result in the tragic loss of life, so it is appropriate to have type I error control as the priority. Cost-sensitive learning, which assigns different costs as weights of type I and type II errors (Elkan, 2001; Zadrozny et al., 2003) is a popular alternative paradigm to address asymmetric errors. This approach has merits and many practical values. However, when there is no consensus to assign costs to errors, or in applications such as medical diagnosis, where it is morally unacceptable to do a cost and benefit analysis, the NP paradigm is a more natural choice. Previous NP classification literature use both empirical risk minimization (ERM) (Cannon et al., 2002; Casasent and Chen, 2003; Scott, 2005; Scott and Nowak, 2005; Han et al., 2008; Rigollet and Tong, 2011) and plug-in approaches (Tong, 2013; Zhao et al., 2016), and its genetic application is suggested in Li and Tong (2016). More recently, Tong et al. (2018) took a different route, and proposed an NP umbrella algorithm that adapts scoring-type classification algorithms (e.g., logistic regression, support vector machines,

1. In verbal discussion, with a slight abuse of language, we also refer to the action of assigning a class 0 observation to class 1 as type I error, and that of assigning a class 1 observation to class 0 as type II error.

random forest, etc.) to the NP paradigm, by setting nonparametric order statistics based thresholds on the classification scores. To implement the NP paradigm, it is very tempting to simply tune the empirical type I error to (no more than) α . Nevertheless, as argued extensively in Tong et al. (2018), doing so would not lead to classifiers whose type I errors are bounded from above by α with high probability.

Universal the NP umbrella algorithm is, it demands an explicit minimum sample size requirement on class 0, which is usually the more scarce class. While this requirement is not too stringent, it does cause a problem when the smaller class 0 has an insufficient sample size. For instance, a commonly-used lung cancer diagnosis example (Gordon et al., 2002) in the high-dimensional statistics literature has 181 subjects and 12,533 features. These subjects carry two different types of lung cancer, namely, adenocarcinoma (150 subjects) and mesothelioma (31 subjects). If we were to treat mesothelioma as the class 0, and would like to control the type I error under $\alpha = .05$ with high probability $1 - \delta_0 = .95$, then the sample size requirement of the NP umbrella algorithm for the left-out class 0 observations is 59 (or $\log \delta_0 / \log(1 - \alpha)$ to be exact), which is almost twice as much as the total available class 0 sample size 31. In this work, we employ the linear discriminant analysis (LDA) model and propose a new thresholding algorithm, which does not have the minimum sample size requirements on class 0 observations and thus applies to small sample applications such as rare disease diagnosis.

In total, we develop four NP classifiers in this paper: combinations of two different ways to create scoring functions in low- and high-dimensional settings under the LDA model, and two thresholding methods that include the nonparametric order statistics based NP umbrella algorithm and the newly proposed parametric thresholding rule. We will denote these classifiers by NP-LDA, NP-sLDA, pNP-LDA and pNP-sLDA, where **p** means “parametric thresholding” and **s** means “sparse”. Extensive numerical experiments will suggest the recommended application domains of these methods.

On the theoretical front, NP oracle inequalities, a core theoretical criterion to evaluate classifiers under the NP paradigm, were only established for nonparametric plug-in classifiers (Tong, 2013; Zhao et al., 2016). The current work is the first to establish NP oracle inequalities under parametric models. Concretely, we will show that NP-sLDA satisfies the NP oracle inequalities. Another major contribution of this work is that we design an adaptive sample splitting scheme that can be applied universally to existing NP classifiers. This adaptive strategy enhances the power (i.e., reduces type II error) and therefore raises the practicality of the NP algorithms.

The rest of this paper is organized as follows. Section 2 introduces the notations and model setup. Section 3 constructs the new parametric thresholding rule based on the LDA model and introduces four new LDA based NP classifiers. Section 4 formulates theoretical conditions and derives NP oracle inequalities for NP-sLDA. Section 5 describes the new data-adaptive sample splitting scheme. Numerical results are presented in Section 6, followed by a short discussion in Section 7. Longer proofs, additional numerical results, and other supporting intermediate results are relegated to the appendices.

2. Notations and model setup

A few standard notations are introduced to facilitate our discussion. Let (X, Y) be a random pair where $X \in \mathcal{X} \subset \mathbb{R}^d$ is a d -dimensional vector of features, and $Y \in \{0, 1\}$ indicates X 's class label. Denote respectively by \mathbb{P} and \mathbb{E} generic probability distribution and expectation. A *classifier* $\phi : \mathcal{X} \rightarrow \{0, 1\}$ is a data-dependent mapping from \mathcal{X} to $\{0, 1\}$ that assigns X to one of the classes. The overall classification error of ϕ is $R(\phi) = \mathbb{E}\mathbb{I}\{\phi(X) \neq Y\} = \mathbb{P}\{\phi(X) \neq Y\}$, where $\mathbb{I}(\cdot)$ denotes the indicator function. By the law of total probability, $R(\phi)$ can be decomposed into a weighted average of type I error $R_0(\phi) = \mathbb{P}\{\phi(X) \neq Y|Y = 0\}$ and type II error $R_1(\phi) = \mathbb{P}\{\phi(X) \neq Y|Y = 1\}$ as

$$R(\phi) = \pi_0 R_0(\phi) + \pi_1 R_1(\phi), \quad (1)$$

where $\pi_0 = \mathbb{P}(Y = 0)$ and $\pi_1 = \mathbb{P}(Y = 1)$. While the classical paradigm minimizes $R(\cdot)$, the Neyman-Pearson (NP) paradigm seeks to minimize R_1 while controlling R_0 under a user-specified level α . The (level- α) *NP oracle classifier* is thus

$$\phi_\alpha^* \in \arg \min_{R_0(\phi) \leq \alpha} R_1(\phi), \quad (2)$$

where the *significance level* α reflects the level of conservativeness towards type I error.

In this paper, we assume that $(X|Y = 0)$ and $(X|Y = 1)$ follow multivariate Gaussian distributions with a common covariance matrix. That is, their probability density functions f_0 and f_1 are

$$f_0 \sim \mathcal{N}(\mu^0, \Sigma) \text{ and } f_1 \sim \mathcal{N}(\mu^1, \Sigma),$$

where the mean vectors $\mu^0, \mu^1 \in \mathbb{R}^d$ ($\mu^0 \neq \mu^1$) and the common positive definite covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$. This model is frequently referred to as the linear discriminant analysis (LDA) model. Despite its simplicity, the LDA model has been proved to be effective in many applications and benchmark datasets. Moreover, in the last ten years, several papers (Shao et al., 2011; Cai and Liu, 2011; Fan et al., 2012; Witten and Tibshirani, 2012; Mai et al., 2012) have developed LDA based algorithms under high-dimensional settings where the dimensionality of features is comparable to or larger than the sample size.

It is well known that the Bayes classifier (i.e., oracle classifier) of the classical paradigm is $\phi^*(x) = \mathbb{I}(\eta(x) > 1/2)$, where $\eta(x) = \mathbb{E}(Y|X = x) = \mathbb{P}(Y = 1|X = x)$ is the regression function. Since

$$\eta(x) = \frac{\pi_1 \cdot f_1(x)/f_0(x)}{\pi_1 \cdot f_1(x)/f_0(x) + \pi_0},$$

the oracle classifier can be written alternatively as $\mathbb{I}(f_1(x)/f_0(x) > \pi_0/\pi_1)$. When f_0 and f_1 follow the LDA model, the oracle classifier of the classical paradigm is

$$\phi^*(x) = \mathbb{I} \left\{ (x - \mu_a)^\top \Sigma^{-1} \mu_d + \log \frac{\pi_1}{\pi_0} > 0 \right\} = \mathbb{I} \left\{ (\Sigma^{-1} \mu_d)^\top x > \mu_a^\top \Sigma^{-1} \mu_d - \log \frac{\pi_1}{\pi_0} \right\}, \quad (3)$$

where $\mu_a = \frac{1}{2}(\mu^0 + \mu^1)$, $\mu_d = \mu^1 - \mu^0$, and $(\cdot)^\top$ denotes the transpose of a vector. In contrast, motivated by the famous *Neyman-Pearson Lemma* in hypothesis testing (attached in the Appendix A for readers' convenience), the NP oracle classifier is

$$\phi_\alpha^*(x) = \mathbb{I} \left\{ \frac{f_1(x)}{f_0(x)} > C_\alpha \right\}, \quad (4)$$

for some threshold C_α such that $P_0\{f_1(X)/f_0(X) > C_\alpha\} \leq \alpha$ and $P_0\{f_1(X)/f_0(X) \geq C_\alpha\} \geq \alpha$, where P_0 is the conditional probability distribution of X given $Y = 0$ (P_1 is defined similarly).

Under the LDA model assumption, the NP oracle classifier is $\phi_\alpha^*(x) = \mathbb{I}((\Sigma^{-1}\mu_d)^\top x > C_\alpha^{**})$, where $C_\alpha^{**} = \log C_\alpha + \mu_d^\top \Sigma^{-1} \mu_d$. Denote by $\beta^{\text{Bayes}} = \Sigma^{-1} \mu_d$ and $s^*(x) = (\Sigma^{-1} \mu_d)^\top x = (\beta^{\text{Bayes}})^\top x$, then the NP oracle classifier (4) can be written as

$$\phi_\alpha^*(x) = \mathbb{I}(s^*(x) > C_\alpha^{**}). \quad (5)$$

In fact, leveraging further the LDA model assumption, C_α^{**} can be written out explicitly. Note that when $X^j \sim \mathcal{N}(\mu^j, \Sigma)$, $(\Sigma^{-1} \mu_d)^\top X^j \sim \mathcal{N}(\mu_d^\top \Sigma^{-1} \mu^j, \mu_d^\top \Sigma^{-1} \mu_d)$, where $j = 0, 1$. Let $Z = \{(\Sigma^{-1} \mu_d)^\top X^0 - (\Sigma^{-1} \mu_d)^\top \mu^0\} (\mu_d^\top \Sigma^{-1} \mu_d)^{-1/2}$, then $Z \sim \mathcal{N}(0, 1)$. As Σ^{-1} is positive definite and $\mu^0 \neq \mu^1$, we have $\mu_d^\top \Sigma^{-1} \mu^0 < \mu_d^\top \Sigma^{-1} \mu^1$, and thus the level- α NP oracle is

$$\phi_\alpha^*(x) = \mathbb{I}(s^*(x) > C_\alpha^{**}) = \mathbb{I}\left((\Sigma^{-1} \mu_d)^\top x > \sqrt{\mu_d^\top \Sigma^{-1} \mu_d} \Phi^{-1}(1 - \alpha) + \mu_d^\top \Sigma^{-1} \mu^0\right), \quad (6)$$

where $\Phi(\cdot)$ is the CDF of the univariate standard normal distribution $\mathcal{N}(0, 1)$. This oracle holds independent of the feature dimensionality. In reality, we cannot expect that the type I error bound holds almost surely; instead, we can only hope that a classifier $\hat{\phi}_\alpha$ trained on a finite sample have $R_0(\hat{\phi}_\alpha) \leq \alpha$ with high probability. We will construct multiple versions of LDA-based $\hat{\phi}_\alpha$ to suit different application domains in the next section.

Other mathematical notations are introduced as follows. For a general $m_1 \times m_2$ matrix M , $\|M\|_\infty = \max_{i=1, \dots, m_1} \sum_{j=1}^{m_2} |M_{ij}|$, and $\|M\|$ denotes the operator norm. For a vector b , $\|b\|_\infty = \max_j |b_j|$, $|b|_{\min} = \min_j |b_j|$, and $\|b\|$ denotes the L_2 norm. Let $A = \{j : (\Sigma^{-1} \mu_d)_j \neq 0\}$, and μ_A^1 be a sub-vector of μ^1 of length $s := \text{cardinality}(A)$ that consists of the coordinates of μ^1 in A (similarly for μ_A^0). Up to permutation, the Σ matrix can be written as

$$\Sigma = \begin{bmatrix} \Sigma_{AA} & \Sigma_{AA^c} \\ \Sigma_{A^cA} & \Sigma_{A^cA^c} \end{bmatrix}.$$

We use $\lambda_{\max}(\cdot)$ and $\lambda_{\min}(\cdot)$ to denote maximum and minimum eigenvalues of a matrix, respectively.

3. Constructing LDA-based NP classifiers $\hat{\phi}_\alpha(\cdot)$

In this section, we construct two estimates of $s^*(\cdot)$ in Sections 3.1 and 3.2, and two estimates of C_α^{**} in Sections 3.3 and 3.4. Hence in total, we present in Section 3.5 four LDA-based NP classifiers $\hat{\phi}_\alpha(\cdot)$ for different application domains. We assume the following sampling scheme for the rest of this section and in the theoretical analysis. Let $\mathcal{S}_0 = \{X_1^0, \dots, X_{n_0}^0\}$ be an i.i.d. class 0 sample of size n_0 , $\mathcal{S}'_0 = \{X_{n_0+1}^0, \dots, X_{n_0+n'_0}^0\}$ be an i.i.d. class 0 sample of size n'_0 and $\mathcal{S}_1 = \{X_1^1, \dots, X_{n_1}^1\}$ be an i.i.d. class 1 sample of size n_1 . The example sizes n_0 , n'_0 and n_1 are considered as fixed numbers. Moreover, we assume that samples \mathcal{S}_0 , \mathcal{S}'_0 and \mathcal{S}_1 are independent of each other.

3.1. Estimating $s^*(\cdot)$ in low-dimensional settings

To estimate $s^*(x) = (\Sigma^{-1}\mu_d)^\top(x)$, we divide the situation into low- and high-dimensional feature dimensionality d . In low-dimensional settings, that is, when feature dimensionality d is small compared to the sample sizes, we use \mathcal{S}_0 and \mathcal{S}_1 to get the sample means $\hat{\mu}^0$ and $\hat{\mu}^1$ that estimate μ^0 and μ^1 respectively, and get the pooled sample covariance matrix $\hat{\Sigma}$ that estimates Σ . Precisely, we have

$$\hat{\Sigma} = \frac{1}{n_0 + n_1 - 2} \left(\sum_{i=1}^{n_0} (X_i^0 - \hat{\mu}^0)(X_i^0 - \hat{\mu}^0)^\top + \sum_{i=1}^{n_1} (X_i^1 - \hat{\mu}^1)(X_i^1 - \hat{\mu}^1)^\top \right), \quad (7)$$

$$\hat{\mu}^0 = \frac{1}{n_0} (X_1^0 + \dots + X_{n_0}^0), \quad (8)$$

$$\hat{\mu}^1 = \frac{1}{n_1} (X_1^1 + \dots + X_{n_1}^1). \quad (9)$$

Let $\hat{\mu}_d = \hat{\mu}^1 - \hat{\mu}^0$, then we can estimate $s^*(\cdot)$ by $\hat{s}(x) = (\hat{\Sigma}^{-1}\hat{\mu}_d)^\top x$.

3.2. Estimating $s^*(\cdot)$ in high-dimensional settings

In high-dimensional settings where d is larger than the sample sizes, $\hat{\Sigma}$ in (7) is not invertible, and we need to resort to more sophisticated methods to estimate $s^*(\cdot)$. First, we note that although the decision thresholds are different, the NP oracle ϕ_α^* in (5) and the classical oracle ϕ^* in (3) both project an observation x to the $\beta^{\text{Bayes}} = \Sigma^{-1}\mu_d$ direction. Hence one can leverage existing works on sparse LDA methods under the classical paradigm to find a β^{Bayes} estimate, using samples \mathcal{S}_0 and \mathcal{S}_1 . In particular, we adopt $\hat{\beta}^{\text{lasso}}$, the lassoed (sparse) discriminant analysis (sLDA) direction in Mai et al. (2012), which is defined by

$$(\hat{\beta}^{\text{lasso}}, \hat{\beta}_0^\lambda) = \arg \min_{(\beta, \beta_0)} \left\{ n^{-1} \sum_{i=1}^n (y_i - \beta_0 - x_i^\top \beta)^2 + \lambda \sum_{j=1}^d |\beta_j| \right\}, \quad (10)$$

where $n = n_0 + n_1$, $y_i = -n/n_0$ if the i th observation is from class 0, and $y_i = n/n_1$ if the i th observation is from class 1. Then we estimate $s^*(\cdot)$ by

$$\hat{s}(x) = (\hat{\beta}^{\text{lasso}})^\top x.$$

Note that although the optimization program (10) is the same as in Mai et al. (2012), our sampling scheme is different from that in Mai et al. (2012), where they assumed i.i.d. observations from the joint distribution of (X, Y) . As a consequence, when analyzing theoretical properties for $\hat{\beta}^{\text{lasso}}$ in (10), it is necessary to establish results that are counterparts to those in Mai et al. (2012).

3.3. Estimating C_α^{**} via the nonparametric NP umbrella algorithm

Tong et al. (2018) provides a nonparametric order statistics based method, the NP umbrella algorithm, to estimate the threshold C_α^{**} . This algorithm leverages the following proposition.

Proposition 1 *Suppose that we use \mathcal{S}_0 and \mathcal{S}_1 to train a base algorithm (e.g., $sLDA$), and obtain a scoring function f (e.g., an estimate of s^*). Applying f to \mathcal{S}'_0 , we denote the resulting classification scores as $T_1, \dots, T_{n'_0}$, which are real-valued random variables. Then, denote by $T_{(k)}$ the k -th order statistic (i.e., $T_{(1)} \leq \dots \leq T_{(n'_0)}$). For a new observation X , if we denote its classification score $f(X)$ as T , we can construct classifiers $\hat{\phi}_k(X) = \mathbb{I}(T > T_{(k)})$, $k \in \{1, \dots, n'_0\}$. Then, the population type I error of $\hat{\phi}_k$, denoted by $R_0(\hat{\phi}_k)$, is a function of $T_{(k)}$ and hence a random variable, and it holds that*

$$\mathbb{P} \left[R_0(\hat{\phi}_k) > \alpha \right] \leq \sum_{j=k}^{n'_0} \binom{n'_0}{j} (1-\alpha)^j \alpha^{n'_0-j}. \quad (11)$$

That is, the probability that the type I error of $\hat{\phi}_k$ exceeds α is under a constant that only depends on k , α and n'_0 . We call this probability the violation rate of $\hat{\phi}_k$ and denote its upper bound by $v(k) = \sum_{j=k}^{n'_0} \binom{n'_0}{j} (1-\alpha)^j \alpha^{n'_0-j}$. When T_i 's are continuous, this bound is tight.

Proposition 1 is the key step towards the NP umbrella algorithm proposed in Tong et al. (2018), which applies to all scoring-type classification methods (base algorithms), including logistic regression, support vector machines, random forest, etc. In the theoretical analysis part of this paper, we always assume the continuity of scoring functions. Under this mild assumption, $v(k)$ is the violation rate of type I error for $\hat{\phi}_k$. An essential step of the proof of the proposition hinges on the symmetry property of permutation. It is obvious that $v(k)$ decreases as k increases. To choose from $\hat{\phi}_1, \dots, \hat{\phi}_{n'_0}$ such that a classifier achieves minimal type II error with type I error violation rate less than or equal to a user's specified δ_0 , the right order is

$$k^* = \min \{k \in \{1, \dots, n'_0\} : v(k) \leq \delta_0\}. \quad (12)$$

In our current setting, $\hat{s}(\cdot)$, constructed as in Section 3.1 or Section 3.2, plays the role of the scoring function. Let $\hat{s}_{(k^*)}$ be the k^* -th order statistic among the set $\{\hat{s}(X_{n_0+1}^0), \dots, \hat{s}(X_{n_0+n'_0}^0)\}$, then the NP umbrella algorithm sets

$$\hat{C}_\alpha = \hat{s}_{(k^*)}.$$

To achieve $\mathbb{P} \left[R_0(\hat{\phi}_k) > \alpha \right] \leq \delta_0$ for some $\hat{\phi}_k$, we need to control the violation rate under δ_0 at least in the extreme case when $k = n'_0$; that is, it is necessary to ensure $v(n'_0) = (1-\alpha)^{n'_0} \leq \delta_0$. Clearly, if the (n'_0) -th order statistic cannot guarantee the violation rate control, other order statistics certainly cannot. Therefore, for a given α and δ_0 , there exists a minimum left-out class 0 sample size requirement

$$n'_0 \geq \log \delta_0 / \log(1-\alpha),$$

for the type I error violation rate control. Note that the control on type I error violation rate does not demand any sample size requirements on \mathcal{S}_0 and \mathcal{S}_1 . But these two parts have an impact on the quality of scoring functions, and hence on the type II error performance.

3.4. Estimating C_{α}^{**} by leveraging the parametric assumption

We explicitly leverage the LDA model assumption to create an estimate of C_{α}^{**} . For simplicity, let \hat{A} be either $\hat{\Sigma}^{-1}\hat{\mu}_d$ in Section 3.1 or $\hat{\beta}^{\text{lasso}}$ in Section 3.2, corresponding to the low- and high-dimensional settings, respectively. First, we consider \hat{A} as fixed (i.e., fix \mathcal{S}_0 and \mathcal{S}_1). When $X^0 \sim \mathcal{N}(\mu^0, \Sigma)$, we have

$$\hat{A}^\top X^0 \sim \mathcal{N}(\hat{A}^\top \mu^0, \hat{A}^\top \Sigma \hat{A}).$$

Let

$$\tilde{\phi}_\alpha(x) = \mathbb{I} \left(\hat{A}^\top x > \sqrt{\hat{A}^\top \Sigma \hat{A}} \Phi^{-1}(1 - \alpha) + \hat{A}^\top \mu^0 \right).$$

For every fixed \mathcal{S}_0 and \mathcal{S}_1 , $\tilde{\phi}_\alpha$ is the NP oracle at level α . Note that $\tilde{\phi}_\alpha$ is not an accessible classifier because Σ and μ^0 are unknown. Plugging in the estimates of these parameters is not a good idea because this will not achieve a high probability control of the type I error under α . We plan to construct a statistic \hat{C}_α^p (the super index p stands for ‘‘parametric thresholding’’) such that with high probability,

$$\hat{C}_\alpha^p \geq \sqrt{\hat{A}^\top \Sigma \hat{A}} \Phi^{-1}(1 - \alpha) + \hat{A}^\top \mu^0.$$

Let

$$\hat{\phi}_\alpha(x) = \mathbb{I} \left(\hat{A}^\top x > \hat{C}_\alpha^p \right).$$

Then $R_0(\hat{\phi}_\alpha) \leq R_0(\tilde{\phi}_\alpha) = \alpha$ with high probability. Naturally, we want such a \hat{C}_α^p as small as possible, so that the resulting classifier has good type II error performance. Towards this end, we build tight sample-based upper bounds for $\hat{A}^\top \mu^0$ (Lemma 2) and $\hat{A}^\top \Sigma \hat{A}$ (Lemma 3), and then combine these bounds to get \hat{C}_α^p (Proposition 4).

3.4.1. UPPER BOUND FOR $\hat{A}^\top \mu^0$

Lemma 2 *Let $W_1 = \hat{A}^\top X_{n_0+1}^0$, $W_2 = \hat{A}^\top X_{n_0+2}^0, \dots$, and $W_{n'_0} = \hat{A}^\top X_{n_0+n'_0}^0$. Denote by $F_{t(n'_0-1)}(\cdot)$ the cumulative distribution function of the t -distribution with degrees of freedom $(n'_0 - 1)$. Then with probability at least $1 - \delta_0$, it holds that*

$$\hat{A}^\top \mu^0 \leq \bar{W} - F_{t(n'_0-1)}^{-1}(\delta_0) \frac{S}{\sqrt{n'_0}}, \quad (13)$$

where $\bar{W} = (W_1 + \dots + W_{n'_0})/n'_0$ and $S = \left\{ \left[(W_1 - \bar{W})^2 + \dots + (W_{n'_0} - \bar{W})^2 \right] / (n'_0 - 1) \right\}^{1/2}$.

Proof We invoke the following classic result. Suppose W_1, \dots, W_m are i.i.d. from one-dimensional $\mathcal{N}(\mu, \sigma^2)$. Let $\bar{W} = (W_1 + \dots + W_m)/m$ and $S = \left\{ \sum_{j=1}^m (W_j - \bar{W})^2 / (m - 1) \right\}^{1/2}$, then we have

$$\frac{\bar{W} - \mu}{S/\sqrt{m}} \sim t_{(m-1)} \text{ (} t \text{ distribution with df = } m - 1 \text{)}.$$

Take $W = \hat{A}^\top X$ and $m = n'_0$, then for fixed \hat{A} , $W_1 = \hat{A}^\top X_{n_0+1}^0, \dots, W_{n'_0} = \hat{A}^\top X_{n_0+n'_0}^0$ are i.i.d. from one-dimensional normally distributed variable with mean $\hat{A}^\top \mu^0$. Then it

follows that

$$\frac{\bar{W} - \hat{A}^\top \mu^0}{S/\sqrt{n'_0}} \sim t_{(n'_0-1)}.$$

Thus with probability $1 - \delta_0$ (randomness comes from \mathcal{S}'_0 while keeping \mathcal{S}_0 and \mathcal{S}_1 fixed),

$$\frac{\bar{W} - \hat{A}^\top \mu^0}{S/\sqrt{n'_0}} \geq F_{t_{(n'_0-1)}}^{-1}(\delta_0) \iff \hat{A}^\top \mu^0 \leq \bar{W} - F_{t_{(n'_0-1)}}^{-1}(\delta_0) \frac{S}{\sqrt{n'_0}}. \quad (14)$$

Since the above inequality is true for every realization of \mathcal{S}_0 and \mathcal{S}_1 , and \mathcal{S}'_0 is independent of \mathcal{S}_0 and \mathcal{S}_1 , the above upper bound for $\hat{A}^\top \mu^0$ holds with probability $1 - \delta_0$ regarding all sampling randomness. \blacksquare

3.4.2. UPPER BOUND FOR $\hat{A}^\top \Sigma \hat{A}$

Lemma 3 *Let $n = n_0 + n_1$. Suppose there exist some positive constants c and C such that $(n-2)^{1/C} < d < (n-2)$ and $d/(n-2) \leq 1 - c$. Then for any positive ϵ and D , there exists an $N(\epsilon, D)$ such that for all $n \geq N(\epsilon, D)$, it holds with probability at least $1 - (n-2)^{-D}$ that,*

$$\hat{A}^\top \Sigma \hat{A} \leq \frac{1}{\left(1 - \sqrt{\frac{d}{n-2}}\right)^2 - \frac{(n-2)^\epsilon}{\sqrt{n-2}d^{\frac{1}{6}}}} \lambda_{\max}(\hat{\Sigma}) \hat{A}^\top \hat{A}. \quad (15)$$

Proof An obvious upper bound for $\hat{A}^\top \Sigma \hat{A}$ is $\lambda_{\max}(\Sigma) \hat{A}^\top \hat{A}$, but $\lambda_{\max}(\Sigma)$ is not accessible from samples. Instead, $\lambda_{\max}(\hat{\Sigma})$ is accessible. In the following, we explore the relations between $\lambda_{\max}(\Sigma)$ and $\lambda_{\max}(\hat{\Sigma})$ to derive a sample-based upper bound for $\hat{A}^\top \Sigma \hat{A}$.

Let Z_1, \dots, Z_{n_0} be i.i.d. from d -dimensional Gaussian $\mathcal{N}(\mu^0, I_{d \times d})$, $Z_{n_0+1}, \dots, Z_{n_0+n_1}$ be i.i.d. from $\mathcal{N}(\mu^1, I_{d \times d})$, and $n = n_0 + n_1$. Further assume that all the Z_i 's, $i = 1, \dots, n$, are independent of each other. Let

$$\hat{U} = \frac{1}{n-2} \left(\sum_{i=1}^{n_0} (Z_i - \bar{Z}_0)(Z_i - \bar{Z}_0)^\top + \sum_{i=n_0+1}^n (Z_i - \bar{Z}_1)(Z_i - \bar{Z}_1)^\top \right),$$

where $\bar{Z}_0 = (Z_1 + \dots + Z_{n_0})/n_0$ and $\bar{Z}_1 = (Z_{n_0+1} + \dots + Z_n)/n_1$. Define $G_i = Z_i - \mu^0$ for $i = 1, \dots, n_0$ and $G_i = Z_i - \mu^1$ for $i = n_0 + 1, \dots, n$. Then $G_i \sim \mathcal{N}(0, I_{d \times d})$ for all $i = 1, \dots, n$. Let

$$\hat{V} = \frac{1}{n-2} \left(\sum_{i=1}^{n_0} (G_i - \bar{G}_0)(G_i - \bar{G}_0)^\top + \sum_{i=n_0+1}^n (G_i - \bar{G}_1)(G_i - \bar{G}_1)^\top \right),$$

where $\bar{G}_0 = (G_1 + \dots + G_{n_0})/n_0$ and $\bar{G}_1 = (G_{n_0+1} + \dots + G_n)/n_1$. Clearly $\hat{U} = \hat{V}$. Let

$$\mathbf{P} = \text{diag} \left(I_{n_0 \times n_0} - \frac{1}{n_0} \mathbf{e}_0 \mathbf{e}_0^\top, I_{n_1 \times n_1} - \frac{1}{n_1} \mathbf{e}_1 \mathbf{e}_1^\top \right),$$

where the entries of \mathbf{e}_0 and \mathbf{e}_1 are all equal to 1, and the lengths are n_0 and n_1 respectively. Then \mathbf{P} is a projection matrix with rank $n-2$. Note that a projection matrix has eigenvalues

all equal to 1 and 0, with the number of 1's equals to its rank. Hence, we can decompose \mathbf{P} by $\mathbf{P} = \mathbf{U}\mathbf{U}^\top$, where \mathbf{U} is an $n \times (n-2)$ orthogonal matrix with $\mathbf{U}^\top\mathbf{U} = I_{(n-2) \times (n-2)}$.

Let $\mathbf{G} = (G_1, \dots, G_n)$. \mathbf{G} is a $d \times n$ matrix and its columns are i.i.d. d -dimensional standard multivariate Gaussian. Let $\tilde{\mathbf{G}} = \mathbf{G}\mathbf{U}$, then $\tilde{\mathbf{G}}$ is a $d \times (n-2)$ matrix in which the columns are i.i.d. d -dimensional standard multivariate Gaussian. Therefore, we have

$$\hat{U} = \hat{V} = \frac{1}{n-2} \mathbf{G}\mathbf{P}\mathbf{G}^\top = \frac{1}{n-2} \mathbf{G}\mathbf{U}\mathbf{U}^\top\mathbf{G}^\top = \frac{1}{n-2} \tilde{\mathbf{G}}\tilde{\mathbf{G}}^\top.$$

By Bloemendal et al. (2015), we have the following concentration result on minimum eigenvalues: suppose there exist some positive constants c and C such that $(n-2)^{1/C} < d < (n-2)$ and $d/(n-2) \leq 1 - c$. Then for any positive ϵ and D , there exists an $N(\epsilon, D)$ such that for all $n \geq N(\epsilon, D)$, we have

$$\mathbb{P} \left(\left| \lambda_{\min} \left(\frac{1}{n-2} \tilde{\mathbf{G}}\tilde{\mathbf{G}}^\top \right) - \left(1 - \sqrt{\frac{d}{n-2}} \right)^2 \right| > \frac{(n-2)^\epsilon}{\sqrt{n-2}d^{\frac{1}{6}}} \right) \leq (n-2)^{-D}. \quad (16)$$

This result implies that for $n \geq N(\epsilon, D)$, we have with probability at least $1 - (n-2)^{-D}$,

$$I_{d \times d} \leq \frac{1}{\left(1 - \sqrt{\frac{d}{n-2}} \right)^2 - \frac{(n-2)^\epsilon}{\sqrt{n-2}d^{\frac{1}{6}}}} \left(\frac{1}{n-2} \tilde{\mathbf{G}}\tilde{\mathbf{G}}^\top \right),$$

where the inequality means “ $A \leq B$ iff $B - A$ is positive semi-definite”. This further implies

$$\hat{\Sigma} \stackrel{d}{=} \Sigma^{1/2} \hat{U} \Sigma^{1/2} = \Sigma^{1/2} \left(\frac{1}{n-2} \tilde{\mathbf{G}}\tilde{\mathbf{G}}^\top \right) \Sigma^{1/2} \geq \left(\left(1 - \sqrt{\frac{d}{n-2}} \right)^2 - \frac{(n-2)^\epsilon}{\sqrt{n-2}d^{\frac{1}{6}}} \right) \Sigma,$$

where the notation “ $\stackrel{d}{=}$ ” means equal in distribution. Therefore, for $n \geq N(\epsilon, D)$, we have with probability at least $1 - (n-2)^{-D}$,

$$\lambda_{\max}(\Sigma) \leq \frac{1}{\left(1 - \sqrt{\frac{d}{n-2}} \right)^2 - \frac{(n-2)^\epsilon}{\sqrt{n-2}d^{\frac{1}{6}}}} \lambda_{\max}(\hat{\Sigma}). \quad (17)$$

Hence we can bound $\hat{A}^\top \Sigma \hat{A}$ by

$$\hat{A}^\top \Sigma \hat{A} \leq \lambda_{\max}(\Sigma) \hat{A}^\top \hat{A} \leq \frac{1}{\left(1 - \sqrt{\frac{d}{n-2}} \right)^2 - \frac{(n-2)^\epsilon}{\sqrt{n-2}d^{\frac{1}{6}}}} \lambda_{\max}(\hat{\Sigma}) \hat{A}^\top \hat{A}. \quad (18)$$

■

Lemma 3 can be improved for large d and sparse \hat{A} . To bound $\hat{A}^\top \Sigma \hat{A} = (\hat{\beta}^{\text{lasso}})^\top \Sigma \hat{\beta}^{\text{lasso}}$ using results in Bloemendal et al. (2015), we need the condition $d < n - 2$. Even when d is moderate, the bound on the right-hand side of (15) can be loose. A remedy exists

when $\hat{\beta}^{\text{lasso}}$ is sparse, e.g., $\|\hat{\beta}^{\text{lasso}}\|_0 = s \ll d$. Concretely, we can replace $\hat{\beta}^{\text{lasso}}$ in the first inequality of (18) by its s -dimensional sub-vector that consists of the nonzero coordinates, and Σ by its $s \times s$ sub-matrix that corresponds to the nonzero elements of $\hat{\beta}^{\text{lasso}}$. Then in the multiplicative factor on the right-hand side of (18), we replace d by a much smaller s and replace the maximum eigenvalue of the $d \times d$ pooled sample covariance matrix by that of the $s \times s$ pooled sample covariance matrix. To summarize, we achieve a much tighter high probability bound of $\hat{A}^\top \Sigma \hat{A}$ by exploring the sparsity of \hat{A} . This variant also allows us to handle the situation when d is bigger than the sample sizes. In numerical implementation, when we use $\hat{\beta}^{\text{lasso}}$ for \hat{A} , this improved bound is what we always use, although for notational simplicity we still denote the threshold of scoring function as \hat{C}_α^p .

3.4.3. COMBINING BOUNDS FOR $\hat{A}^\top \mu^0$ AND $\hat{A}^\top \Sigma \hat{A}$

The arguments at the beginning of Section 3.4 together with the derived upper bounds for $\hat{A}^\top \mu^0$ and $\hat{A}^\top \Sigma \hat{A}$ imply the following proposition.

Proposition 4 *Under definitions and conditions in Lemmas 2 and 3, if we set*

$$\hat{C}_\alpha^p = \left(\frac{1}{\left(1 - \sqrt{\frac{d}{n_0+n_1-2}}\right)^2 - \frac{(n_0+n_1-2)^\epsilon}{\sqrt{n_0+n_1-2}d^{\frac{1}{6}}}} \lambda_{\max}(\hat{\Sigma}) \hat{A}^\top \hat{A} \right)^{\frac{1}{2}} \cdot \Phi^{-1}(1-\alpha) + \bar{W} - F_{t_{(n'_0-1)}^{-1}}^{-1}(\delta_0) \frac{S}{\sqrt{n'_0}}, \quad (19)$$

then $\hat{\phi}_\alpha(x) = \mathbb{I}(\hat{A}^\top x > \hat{C}_\alpha^p)$ satisfies

$$\mathbb{P}(R_0(\hat{\phi}_\alpha) \leq \alpha) > 1 - \delta_0 - (n-2)^{-D}. \quad (20)$$

In many applications, n_1 is large, so the sample size requirement of (16) (and hence of (17) and (20)) on n , although impossible to check, can be comfortably assumed. Furthermore, simulation studies in Appendix C show that the inequality (17) is often satisfied with probability very close to 1 even when n is moderate, with the choice of $\epsilon = 1e - 3$. Hence the right-hand side of (20) is often almost $1 - \delta_0$ in practice.

3.5. The resulting four LDA-based classifiers $\hat{\phi}_\alpha$

Having obtained two estimates for s^* and two for C_α^{**} , we construct four classifiers in total: NP-LDA: $\mathbb{I}\left((\hat{\Sigma}^{-1} \hat{\mu}_d)^\top x > \hat{C}_\alpha\right)$, NP-sLDA: $\mathbb{I}\left((\hat{\beta}^{\text{lasso}})^\top x > \hat{C}_\alpha\right)$, pNP-LDA: $\mathbb{I}\left((\hat{\Sigma}^{-1} \hat{\mu}_d)^\top x > \hat{C}_\alpha^p\right)$, pNP-sLDA: $\mathbb{I}\left((\hat{\beta}^{\text{lasso}})^\top x > \hat{C}_\alpha^p\right)$. We will suggest the application domains of these four classifiers in the numerical analysis section.

4. Theoretical analysis

In the theoretical analysis, we focus on the NP-sLDA classifier, which we denote by $\hat{\phi}_{k^*}(x) = \mathbb{I}\left((\hat{\beta}^{\text{lasso}})^\top x > \hat{C}_\alpha\right)$ in this section. We will establish NP oracle inequalities for $\hat{\phi}_{k^*}$. The *NP oracle inequalities* were formulated for classifiers under the NP paradigm to reckon the spirit of oracle inequalities in the classical paradigm. They require two properties to hold

simultaneously with high probability: i). type I error $R_0(\hat{\phi}_{k^*})$ is bounded from above by α , and ii). excess type II error, that is $R_1(\hat{\phi}_{k^*}) - R_1(\phi_\alpha^*)$, diminishes as sample sizes increase. *By construction of the order k^* in NP-sLDA $\hat{\phi}_{k^*}$, the first property is already fulfilled, so in the following we bound the excess type II error.*

In the NP classification literature, nonparametric plug-in NP classifiers constructed in Tong (2013) and Zhao et al. (2016) were shown to satisfy the NP oracle inequalities. Both papers assume bounded feature support $[-1, 1]^d$. Under this assumption, uniform deviation bounds between f_1/f_0 and its nonparametric estimate \hat{f}_1/\hat{f}_0 were derived, and such uniform deviation bounds were crucial in bounding the excess type II error. However, as canonical parametric models in classification (such as LDA and QDA) have unbounded feature support, the development of NP theory under parametric settings cannot bypass the challenges arisen from the unboundedness of feature support. To address these challenges, we follow a conditioning-on-a-high-probability-set strategy and formulate conditional marginal assumption and conditional detection condition. For the LDA model, we elaborate on these high-level conditions in terms of specific parameters. In fact, the same conditioning-on-a-high-probability-set strategy can work under the nonparametric model assumptions, such as with a mild finite moment condition. Thus, we can also relax the bounded feature support assumption for the nonparametric methods. Before presenting the new assumptions and main theorem, we need a few technical lemmas to make the ‘‘conditioning’’ work.

4.1. A few technical lemmas

With kernel density estimates \hat{f}_1 , \hat{f}_0 , and an estimate of the threshold level \tilde{C}_α based on VC inequality, Tong (2013) constructed a plug-in classifier $\mathbb{I}\{\hat{f}_1(x)/\hat{f}_0(x) \geq \tilde{C}_\alpha\}$ that is of limited practical value unless the feature dimension is small and sample size is large. Zhao et al. (2016) analyzed high-dimensional Naive Bayes models under the NP paradigm and innovated the threshold estimate by invoking order statistics with an explicit analytic formula for the chosen order. We denote that order by k' . The order k^* derived in Tong et al. (2018) is a refinement of the order statistics approach to estimate the threshold. However, although the order k^* is optimal, it does not take an explicit formula and thus is not helpful in bounding the excess type II error. Interestingly, efforts to approximate k^* analytically for type II error control leads to k' , and so k' will be employed as a bridge in establishing NP oracle inequalities for $\hat{\phi}_{k^*}$.

To derive an upper bound for the excess type II error, it is essential to bound the deviation between type I error of $\hat{\phi}_{k^*}$ and that of the NP oracle ϕ_α^* . To achieve this, we first quote the next proposition from Zhao et al. (2016) and then derive a corollary.

Proposition 5 *Given $\delta_0 \in (0, 1)$, suppose $n'_0 \geq 4/(\alpha\delta_0)$, let the order k' be defined as follows*

$$k' = \lceil (n'_0 + 1)A_{\alpha, \delta_0}(n'_0) \rceil, \tag{21}$$

where $\lceil z \rceil$ denotes the smallest integer larger than or equal to z , and

$$A_{\alpha, \delta_0}(n'_0) = \frac{1 + 2\delta_0(n'_0 + 2)(1 - \alpha) + \sqrt{1 + 4\delta_0(1 - \alpha)\alpha(n'_0 + 2)}}{2\{\delta_0(n'_0 + 2) + 1\}}.$$

Then we have

$$\mathbb{P}\left(R_0(\hat{\phi}_{k'}) > \alpha\right) \leq \delta_0.$$

In other words, the type I error of classifier $\hat{\phi}_{k'}$ ($\hat{\phi}_k$ was defined in Proposition 1) is bounded from above by α with probability at least $1 - \delta_0$.

Corollary 6 Under continuity assumption of the classification scores T_i 's (which we always assume in this paper), the order k^* is smaller than or equal to the order k' .

Proof Under the continuity assumption of T_i 's, $v(k)$ is the exact violation rate of classifier $\hat{\phi}_k$. By construction, both $v(k')$ and $v(k^*)$ are smaller than or equal to δ_0 . Since k^* is the smallest k that satisfies $v(k) \leq \delta_0$, we have $k^* \leq k'$. \blacksquare

Lemma 7 Let $\alpha, \delta_0 \in (0, 1)$ and $n'_0 \geq 4/(\alpha\delta_0)$. For any $\delta'_0 \in (0, 1)$, the distance between $R_0(\hat{\phi}_{k'})$ and $R_0(\phi_\alpha^*)$ can be bounded as

$$\mathbb{P}\{|R_0(\hat{\phi}_{k'}) - R_0(\phi_\alpha^*)| > \xi_{\alpha, \delta_0, n'_0}(\delta'_0)\} \leq \delta'_0,$$

where

$$\xi_{\alpha, \delta_0, n'_0}(\delta'_0) = \sqrt{\frac{k'(n'_0 + 1 - k')}{(n'_0 + 2)(n'_0 + 1)^2\delta'_0}} + A_{\alpha, \delta_0}(n'_0) - (1 - \alpha) + \frac{1}{n'_0 + 1},$$

in which k' and $A_{\alpha, \delta_0}(n'_0)$ are the same as in Proposition 5. Moreover, if $n'_0 \geq \max(\delta_0^{-2}, \delta_0'^{-2})$, we have $\xi_{\alpha, \delta_0, n'_0}(\delta'_0) \leq (5/2)n_0'^{-1/4}$.

Lemma 7 is borrowed from Zhao et al. (2016), so its proof is omitted. Based on Lemma 7 and Corollary 6, we can derive the following result whose proof is in the Appendix.

Lemma 8 Under the same assumptions as in Lemma 7, the distance between $R_0(\hat{\phi}_{k^*})$ and $R_0(\phi_\alpha^*)$ can be bounded as

$$\mathbb{P}\{|R_0(\hat{\phi}_{k^*}) - R_0(\phi_\alpha^*)| > \xi_{\alpha, \delta_0, n'_0}(\delta'_0)\} \leq \delta_0 + \delta'_0.$$

Lemma 8 serves as an intermediate step towards the final “conditional” version to be elaborated in Lemma 11. Moving towards Lemma 11, we construct a set $\mathcal{C} \in \mathbb{R}^d$, such that \mathcal{C}^c is “small”. We also show that the uniform deviation between \hat{s} and s^* on \mathcal{C} is controllable (Lemma 10). To achieve that, we digress to introduce some more notations. Suppose the lassoed linear discriminant analysis (sLDA) finds the set A , which is the support of the Bayes rule direction β^{Bayes} , we have $\hat{\beta}_{A^c}^{\text{lasso}} = 0$ and $\hat{\beta}_A^{\text{lasso}} = \hat{\beta}_A$, where $\hat{\beta}_A$ is defined by

$$(\hat{\beta}_A, \tilde{\beta}_0) = \arg \min_{(\beta, \beta_0)} \left\{ n^{-1} \sum_{i=1}^n (y_i - \beta_0 - \sum_{j \in A} x_{ij} \beta_j)^2 + \sum_{j \in A} \lambda |\beta_j| \right\}.$$

The quantity $\hat{\beta}_A$ is only for theoretical analysis, as the definition assumes knowledge of the true support set A . The next proposition is a counterpart of Theorem 1 in Mai et al. (2012), but due to different sampling schemes, it differs from that theorem and a proof is attached in the Appendix.

Proposition 9 *Assume $\kappa := \|\Sigma_{A^c A}(\Sigma_{AA})^{-1}\|_\infty < 1$ and choose λ in the optimization program (10) such that $\lambda < \min\{|\beta^*|_{\min}/(2\varphi), \Delta\}$, where $\beta^* = (\Sigma_{AA})^{-1}(\mu_A^1 - \mu_A^0)$, $\varphi = \|(\Sigma_{AA})^{-1}\|_\infty$ and $\Delta = \|\mu_A^1 - \mu_A^0\|_\infty$, then it holds that*

1. *With probability at least $1 - \delta_1^*$, $\hat{\beta}_A^{lasso} = \hat{\beta}_A$ and $\hat{\beta}_{A^c}^{lasso} = 0$, where*

$$\delta_1^* = \sum_{l=0}^1 2d \exp\left(-c_2 n_l \frac{\lambda^2(1-\kappa-2\varepsilon\varphi)^2}{16(1+\kappa)^2}\right) + f(d, s, n_0, n_1, (\kappa+1)\varepsilon\varphi(1-\varphi\varepsilon)^{-1}),$$

in which ε is any positive constant less than $\min[\varepsilon_0, \lambda(1-\kappa)(4\varphi)^{-1}(\lambda/2 + (1+\kappa)\Delta)^{-1}]$ and ε_0 is some positive constant, and in which

$$f(d, s, n_0, n_1, \varepsilon) = (d+s)s \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4s^2 n_0}\right) + (d+s)s \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4s^2 n_1}\right),$$

for some constants c_1 and c_2 .

2. *With probability at least $1 - \delta_2^*$, none of the elements of $\hat{\beta}_A$ is zero, where*

$$\delta_2^* = \sum_{l=0}^1 2s \exp(-n_l \varepsilon^2 c_2) + \sum_{l=0}^1 2s^2 \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4n_l s^2}\right).$$

in which ε is any positive constant less than $\min[\varepsilon_0, \xi(3+\xi)^{-1}/\varphi, \Delta\xi(6+2\xi)^{-1}]$, where $\xi = |\beta^|_{\min}/(\Delta\varphi)$.*

3. *For any positive ε satisfying $\varepsilon < \min\{\varepsilon_0, \lambda(2\varphi\Delta)^{-1}, \lambda\}$, we have*

$$\mathbb{P}\left(\|\hat{\beta}_A - \beta^*\|_\infty \leq 4\varphi\lambda\right) \geq 1 - \delta_2^*.$$

Aided by Proposition 9, the next lemma constructs \mathcal{C} , a high probability set under both P_0 and P_1 . Moreover, a high probability bound is derived for the uniform deviation between \hat{s} and s^* on this set.

Lemma 10 *Suppose $\max\{tr(\Sigma_{AA}), tr(\Sigma_{AA}^2), \|\Sigma_{AA}\|, \|\mu_A^0\|^2, \|\mu_A^1\|^2\} \leq c_0 s$ for some constant c_0 , where $s = \text{cardinality}(A)$. For $\delta_3 = \exp\{-(n_0 \wedge n_1)^{1/2}\}$, there exists some constant $c'_1 > 0$, such that $\mathcal{C} = \{X \in \mathbb{R}^d : \|X_A\| \leq c'_1 s^{1/2} (n_0 \wedge n_1)^{1/4}\}$ satisfies $P_0(X \in \mathcal{C}) \geq 1 - \delta_3$ and $P_1(X \in \mathcal{C}) \geq 1 - \delta_3$. Moreover, let $\|\hat{s} - s^*\|_{\infty, \mathcal{C}} := \max_{x \in \mathcal{C}} |\hat{s}(x) - s^*(x)|$. Then for $\delta_1 \geq \delta_1^*$ and $\delta_2 \geq \delta_2^*$, where δ_1^* and δ_2^* are defined as in Proposition 9, it holds that with probability at least $1 - \delta_1 - \delta_2$,*

$$\|\hat{s} - s^*\|_{\infty, \mathcal{C}} \leq 4c'_1 \varphi \lambda s (n_0 \wedge n_1)^{1/4}.$$

The set \mathcal{C} was constructed with two opposing missions in mind. On one hand, we want to restrict the feature space \mathbb{R}^d to \mathcal{C} so that the restricted uniform deviation of \hat{s} from s^* is controlled. On the other hand, we also want \mathcal{C} to be sufficiently large, so that $P_0(X \in \mathcal{C}^c)$ and $P_1(X \in \mathcal{C}^c)$ diminish as sample sizes increase. The next lemma is implied by Lemma 8 and Lemma 10.

Lemma 11 *Let \mathcal{C} be defined as in Lemma 10. Then under the same conditions as in Lemma 7, the distance between $R_0(\hat{\phi}_{k^*}|\mathcal{C}) := P_0(\hat{s}(X) \geq \hat{C}_\alpha|X \in \mathcal{C})$ and $R_0(\phi_\alpha^*|\mathcal{C}) := P_0(s^*(X) \geq C_\alpha^{**}|X \in \mathcal{C})$ can be bounded as*

$$\mathbb{P}\{|R_0(\hat{\phi}_{k^*}|\mathcal{C}) - R_0(\phi_\alpha^*|\mathcal{C})| > 2[\xi_{\alpha, \delta_0, n'_0}(\delta'_0) + \exp\{-(n_0 \wedge n_1)^{1/2}\}]\} \leq \delta_0 + \delta'_0,$$

where $\xi_{\alpha, \delta_0, n'_0}(\delta'_0)$ is defined in Lemma 7.

4.2. Margin assumption and detection condition

Margin assumption and detection condition are critical theoretical assumptions in Tong (2013) and Zhao et al. (2016) for bounding excess type II error of the nonparametric NP classifiers constructed in those papers. To assist our proof strategy that divides the \mathbb{R}^d space into a high probability set (e.g., \mathcal{C} defined in Lemma 10) and its complement, we introduce conditional versions of these assumptions.

Definition 12 (conditional margin assumption) *A function $f(\cdot)$ is said to satisfy conditional margin assumption restricted to \mathcal{C}^* of order $\bar{\gamma}$ with respect to probability distribution P (i.e., $X \sim P$) at the level C^* if there exists a positive constant M_0 , such that for any $\delta \geq 0$,*

$$P\{|f(X) - C^*| \leq \delta | X \in \mathcal{C}^*\} \leq M_0 \delta^{\bar{\gamma}}.$$

The unconditional version of such an assumption was first introduced in Polonik (1995). In the classical binary classification framework, Mammen and Tsybakov (1999) proposed a similar condition named “margin condition” by requiring most data to be away from the optimal decision boundary and this condition has become a common assumption in classification literature. In the classical classification paradigm, Definition 12 reduces to the margin condition by taking $f = \eta$, $\mathcal{C}^* = \text{support}(X)$ and $C^* = 1/2$, with $\{x : |f(x) - C^*| = 0\} = \{x : \eta(x) = 1/2\}$ giving the decision boundary of the classical Bayes classifier. For a \mathcal{C}^* with nontrivial probability, the conditional margin assumption is weaker than the unconditional version. For example, suppose $P(X \in \mathcal{C}^*) \geq 1/2$, then the condition $P\{|f(X) - C^*| \leq \delta\} \leq 2M_0 \delta^{\bar{\gamma}}$ would imply the conditional margin assumption in view of the Bayes Theorem.

Definition 12 is a high level assumption. In view of explicit Gaussian modeling assumptions, it is preferable to derive it based on more elementary assumptions on μ^0 , μ^1 and Σ , for our choices of f , P , C^* and \mathcal{C}^* . Recall that the NP oracle classifier can be written as

$$\phi_\alpha^*(x) = \mathbb{I}\{(\Sigma^{-1}\mu_d)^\top x > C_\alpha^{**}\}.$$

Here we take $f(x) = s^*(x) = (\Sigma^{-1}\mu_d)^\top x$, $\mathcal{C}^* = C_\alpha^{**}$, $P = P_0$, and $\mathcal{C} = \mathcal{C}$ in Lemma 10. When $X \sim \mathcal{N}(\mu^0, \Sigma)$, $(\Sigma^{-1}\mu_d)^\top X \sim \mathcal{N}(\mu_d^\top \Sigma^{-1}\mu^0, \mu_d^\top \Sigma^{-1}\mu_d)$. Lemma 10 guarantees that for $\delta_3 = \exp\{-(n_0 \wedge n_1)^{1/2}\}$, $P_0(X \in \mathcal{C}) \geq 1 - \delta_3$. Moreover,

$$\begin{aligned} & P_0(|s^*(X) - C_\alpha^{**}| \leq \delta | X \in \mathcal{C}) \\ & \leq P_0\left(C_\alpha^{**} - \delta \leq (\Sigma^{-1}\mu_d)^\top X \leq C_\alpha^{**} + \delta\right) / (1 - \delta_3) \\ & = [\Phi(U) - \Phi(L)] / (1 - \delta_3), \end{aligned}$$

where Φ is the cumulative distribution function of the standard normal distribution, $U = (C_\alpha^{**} + \delta - \mu_d^\top \Sigma^{-1} \mu^0) / \sqrt{\mu_d^\top \Sigma^{-1} \mu_d}$, and $L = (C_\alpha^{**} - \delta - \mu_d^\top \Sigma^{-1} \mu^0) / \sqrt{\mu_d^\top \Sigma^{-1} \mu_d}$. By the mean value theorem, we have

$$\Phi(U) - \Phi(L) = \phi(z)(U - L) = \phi(z) \frac{2\delta}{\sqrt{\mu_d^\top \Sigma^{-1} \mu_d}},$$

where ϕ is the probability distribution function of the standard normal distribution, and z is in $[L, U]$. Clearly ϕ is bounded from above by $\phi(0)$. Hence, under the assumptions of Lemma 10, if we additionally assume that $\mu_d^\top \Sigma^{-1} \mu_d \geq C$ for some universal positive constant C , the conditional margin assumption is met with the restricted set \mathcal{C} , the constant $M_0 = 2\phi(0)/(\sqrt{C}(1 - \delta_3))$ and $\bar{\gamma} = 1$. Since $\delta_3 < 1/2$, we can take $M_0 = 4\phi(0)/\sqrt{C}$.

Assumption 1 *i). $\max\{\text{tr}(\Sigma_{AA}), \text{tr}(\Sigma_{AA}^2), \|\Sigma_{AA}\|, \|\mu_A^0\|^2, \|\mu_A^1\|^2\} \leq c_0 s$ for some constant c_0 , where $s = \text{cardinality}(A)$, and $A = \{j : \{\Sigma^{-1} \mu_d\}_j \neq 0\}$; ii). $\mu_d^\top \Sigma^{-1} \mu_d \geq C$ for some universal positive constant C ; iii). the set \mathcal{C} is defined as in Lemma 10 .*

Remark 13 *Under Assumption 1, the function $s^*(\cdot)$ satisfies the conditional margin assumption restricted to \mathcal{C} of order $\bar{\gamma} = 1$ with respect to probability distribution P_0 at the level C_α^{**} . In addition, the constant M_0 can be taken as $M_0 = 4\phi(0)/\sqrt{C}$.*

Unlike the classical paradigm where the optimal threshold $1/2$ on regression function is known, the optimal threshold level in the NP paradigm is unknown and needs to be estimated, suggesting the necessity of having sufficient data around the decision boundary to detect it. This concern motivated Tong (2013) to formulate a detection condition that works as an opposite force to the margin assumption, and Zhao et al. (2016) improved upon it and proved its necessity in bounding excess type II error of an NP classifier. However, formulating a transparent detection condition for feature spaces of unbounded support is subtle: to generalize the detection condition in the same way as we generalize the margin assumption to a conditional version, it is not obvious what elementary general assumptions one should impose on the μ^0 , μ^1 and Σ . The good side is that we are able to establish explicit conditions for $s \leq 2$, aided by the literature on the truncated normal distribution. Also, we need a two-sided detection condition as in Tong (2013), because the technique in Zhao et al. (2016) to get rid of one side does not apply in the unbounded feature support situation.

Definition 14 (conditional detection condition) *A function $f(\cdot)$ is said to satisfy conditional detection condition restricted to \mathcal{C}^* of order $\underline{\gamma}$ with respect to P (i.e., $X \sim P$) at level (C^*, δ^*) if there exists a positive constant M_1 , such that for any $\delta \in (0, \delta^*)$,*

$$P\{C^* \leq f(X) \leq C^* + \delta | X \in \mathcal{C}^*\} \wedge P\{C^* - \delta \leq f(X) \leq C^* | X \in \mathcal{C}^*\} \geq M_1 \delta^{\underline{\gamma}}.$$

Unlike the conditional margin assumption, the conditional detection condition is stronger than its unconditional counterpart, in view of the Bayes Theorem. Although we do not have a proof of the necessity for the conditional detection condition, much efforts to bound excess type II error without it failed.

Assumption 2 *The function $s^*(\cdot)$ satisfies conditional detection condition restricted to \mathcal{C} (defined in Lemma 10) of order $\underline{\gamma} \geq 1$ with respect to P_0 at the level $(C_\alpha^{**}, \delta^*)$.*

Proposition 23 in the Appendix shows that under restrictive settings ($s \leq 2$), Assumption 2 can be implied by more elementary assumptions on the parameters of the LDA model.

4.3. NP oracle inequalities

Having introduced the technical assumptions and lemmas, we present the main theorem.

Theorem 15 *Suppose Assumptions 1 and 2, and the assumptions for Lemmas 1-4 hold. Further suppose $n'_0 \geq \max\{4/(\alpha\alpha_0), \delta_0^{-2}, (\delta'_0)^{-2}, (\frac{1}{10}M_1\delta^{*\underline{\gamma}}/4)\}$, $n_0 \wedge n_1 \geq [-\log(M_1\delta^{*\underline{\gamma}}/4)]^2$, and C_α and $\mu_a^\top \Sigma^{-1} \mu_d$ are bounded from above and below. For $\delta_0, \delta'_0 > 0$, $\delta_1 \geq \delta_1^*$ and $\delta_2 \geq \delta_2^*$, there exist constants \bar{c}_1 , \bar{c}_2 and \bar{c}_3 such that, with probability at least $1 - \delta_0 - \delta'_0 - \delta_1 - \delta_2$, it holds that*

$$(I) \quad R_0(\hat{\phi}_{k^*}) \leq \alpha,$$

$$(II) \quad R_1(\hat{\phi}_{k^*}) - R_1(\phi_\alpha^*) \leq \bar{c}_1(n'_0)^{-\frac{1}{4}(\frac{1+\bar{\gamma}}{\underline{\gamma}} \wedge 1)} + \bar{c}_2(\lambda s)^{1+\bar{\gamma}}(n_0 \wedge n_1)^{\frac{1+\bar{\gamma}}{4}} + \bar{c}_3 \exp \left\{ -(n_0 \wedge n_1)^{\frac{1}{2}} \left(\frac{1+\bar{\gamma}}{\underline{\gamma}} \wedge 1 \right) \right\}.$$

Theorem 15 establishes the NP oracle inequalities for the NP-sLDA classifier $\hat{\phi}_{k^*}$. Note that the upper bound for excess type II error does not contain the overall feature dimensionality d explicitly. However, the indirect dependency is two-fold: first, the choice of λ might depend on d ; second, the minimum requirements (i.e., lower bounds) for δ_1 and δ_2 , which are δ_1^* and δ_2^* defined in Proposition 9, depend on d .

By Assumption 1, $\bar{\gamma} = 1$ and by Proposition 23 in the Appendix, $\underline{\gamma} = 1$ under certain conditions. Take the special case $\bar{\gamma} = \underline{\gamma} = 1$ and take $\lambda \sim \sqrt{\frac{\log d}{(n_0+n_1)}}$, the upper bound on the excess type II error can be simplified to $B_1 = \bar{c}_1(n'_0)^{-1/4} + \bar{c}_2 s^2 \frac{\log d}{\sqrt{n_0+n_1}} + \bar{c}_3 \exp\{-(n_0 \wedge n_1)^{1/2}\}$, where $\bar{c}_1, \bar{c}_2, \bar{c}_3$ are generic constants. The upper bound on excess type II errors when $\bar{\gamma} = \underline{\gamma} = 1$ in Zhao et al. (2016) [nonparametric plug-in estimators for densities, high-dimensional settings, feature independence assumption] is $B_2 = \bar{c}_1(n'_0)^{-1/4} + \bar{c}_2 s^2 \left(\frac{\log n_0}{n_0}\right)^{2\beta/(2\beta+1)} + \bar{c}_3 s^2 \left(\frac{\log n_1}{n_1}\right)^{2\beta/(2\beta+1)}$, where β is a smoothness parameter of the kernel function and the densities. Note that the rate about the left-out class 0 sample size n'_0 (for threshold estimate) is the same. This is because although NP-sLDA relies on an optimal order k^* from the NP umbrella algorithm while the classifier in Zhao et al. (2016) uses the order k' specified in equation (19) of the current paper, analytic approximation of k^* falls back to k' . However, we should note that B_1 and B_2 are not directly comparable due to the simplifying feature independence assumption and a screening stage in Zhao et al. (2016). Concretely, there is a marginal screening step before constructing the scoring function and threshold, and that

requires reserving some class 0 and class 1 observations. The sample sizes of these reserved observations, as well as the full feature space dimensionality d , do not enter the bound B_2 because the theory part of Zhao et al. (2016) assumes a minimum sample size requirement on the observations for screening in terms of d . So the effect of the screening only enters as a probability compromise as opposed to an extra term in the upper bound. Moreover, after the marginal screening step, Zhao et al. (2016) effectively dealt with s one-dimensional problems due to the feature independence assumption. This explains the appearance of the typical exponent $2\beta/(2\beta + 1)$ in one-dimensional nonparametric estimation. Without the feature independence assumption, we would see the exponent as $2\beta/(2\beta + s)$. For a typical β value $\beta = 2$ and a moderate s value $s = 20$, we have $2\beta/(2\beta + s) = 1/6$ which is smaller than $1/2$. Hence without the feature independence assumption, the second and third terms in B_2 could represent slower rates in terms of the sample sizes n_0 and n_1 compared to the counterparts in B_1 . Overall, if $n_0, n'_0, n_1 \sim n$, taking $\beta = 2$ and $s = 20$ without the feature independence assumption, we have $B_2 \sim (\log n/n)^{1/6}$ while $B_1 \sim n^{-1/4}$; the parametric assumptions result in a better rate in this case. Moreover, in important applications such as severe disease diagnosis, the sample size n_0 is much smaller than n_1 . In B_1 , these sample sizes appear together as $n_0 + n_1$, but in B_2 , n_0 appears alone as in $(\log n_0/n_0)^{2\beta/(2\beta+1)}$, which is likely much larger than both $(\log n_1/n_1)^{2\beta/(2\beta+1)}$ and $1/\sqrt{n_0 + n_1}$, considering the $n_0 \ll n_1$ situation.

5. Data-adaptive sample splitting scheme

In practice, researchers and practitioners are not given data as separate sets \mathcal{S}_0 , \mathcal{S}'_0 and \mathcal{S}_1 . Instead, they have a single dataset \mathcal{S} that consists of mixed class 0 and class 1 observations. More class 0 observations to better train the base algorithm and more class 0 observations to provide more candidates for threshold estimate each has its own merits. Hence how to split the class 0 observations into two parts, one to train the base algorithm and the other to estimate the score threshold, is far from intuitive. Tong et al. (2018) adopts a half-half split for class 0 and ignores this issue in the development of the NP umbrella algorithm, as that paper focuses on the type I error violation rate control.

Now switching the focus to type II error, we propose a data-adaptive splitting scheme that universally enhances the power (i.e., reduces type II error) of NP classifiers, as demonstrated in the subsequent numerical studies. The procedure is to choose a split proportion τ according to rankings of K -fold cross-validated type II errors. Concretely, for each split proportion candidate $\tau \in \{.1, .2, \dots, .9\}$, the following steps are implemented.

1. Randomly split class 1 observations into K -folds.
2. Use all class 0 observations and $K - 1$ folds of class 1 observations to train an NP classifier. For class 0 observations, τ proportion is used to train the base algorithm, and $1 - \tau$ proportion for threshold estimate.
3. For this classifier, calculate its classification error on the validation fold of the class 1 observations (type II error).

4. Repeat steps (2) and (3) for K times, with each of the K folds used exactly once as the validation data. Compute the mean of type II errors in step (3), and denote it by $e(\tau)$.

Our choice of the split proportion is

$$\tau_{\min} = \underset{\tau \in \{.1, \dots, .9\}}{\operatorname{arg\,min}} e(\tau).$$

Note that τ_{\min} not only depends on the dataset \mathcal{S} , but also on the base algorithm and the thresholding algorithm one uses, as well as on the user-specified α and δ_0 . Merits of this adaptive splitting scheme will be revealed in the simulation section. Here we elaborate on how to reconcile this adaptive scheme with the violation rate control objective. The type I error violation rate control was proved based on a fixed split proportion of class 0 observations, so will the adaptive splitting scheme be overly aggressive on type II error such that we can no longer keep the type I error violation rate under control? If for each realization (among infinite realizations) of the mixed sample \mathcal{S} , we do adaptive splitting on class 0 observations before implementing NP-sLDA $\hat{\phi}_{k^*}$ (or other NP classifiers), then the overall procedure indeed does not lead to a classifier with type I error violation rate controlled under δ_0 . However, this is not how we think about this process; instead, we only adaptively split for one realization of \mathcal{S} , getting a split proportion $\hat{\tau}$, and then fix $\hat{\tau}$ in all rest realizations of \mathcal{S} . This implementation of the overall procedure keeps the type I error violation rate under control.

6. Numerical analysis

Through extensive simulations and real data analysis, we study the performance of NP-LDA, NP-sLDA, pNP-LDA and pNP-sLDA. In addition, we will study the new adaptive sample splitting scheme. In this section, N_0 denotes the total class 0 training sample size (we do not use n_0 and n'_0 here, as class 0 observations are not assumed to be pre-divided into two parts), n_1 denotes the class 1 training sample size, and $N = N_0 + n_1$ denotes the total sample size.

6.1. Simulation studies under low-dimensional settings

In this subsection, we consider the low-dimensional settings with two examples. In particular, we would like to compare the performance of four NP based methods: NP-LDA, NP-sLDA, pNP-LDA and pNP-sLDA. In all NP methods, τ , the class 0 split proportion, is fixed at 0.5. In every simulation setting, the experiments are repeated 1,000 times.

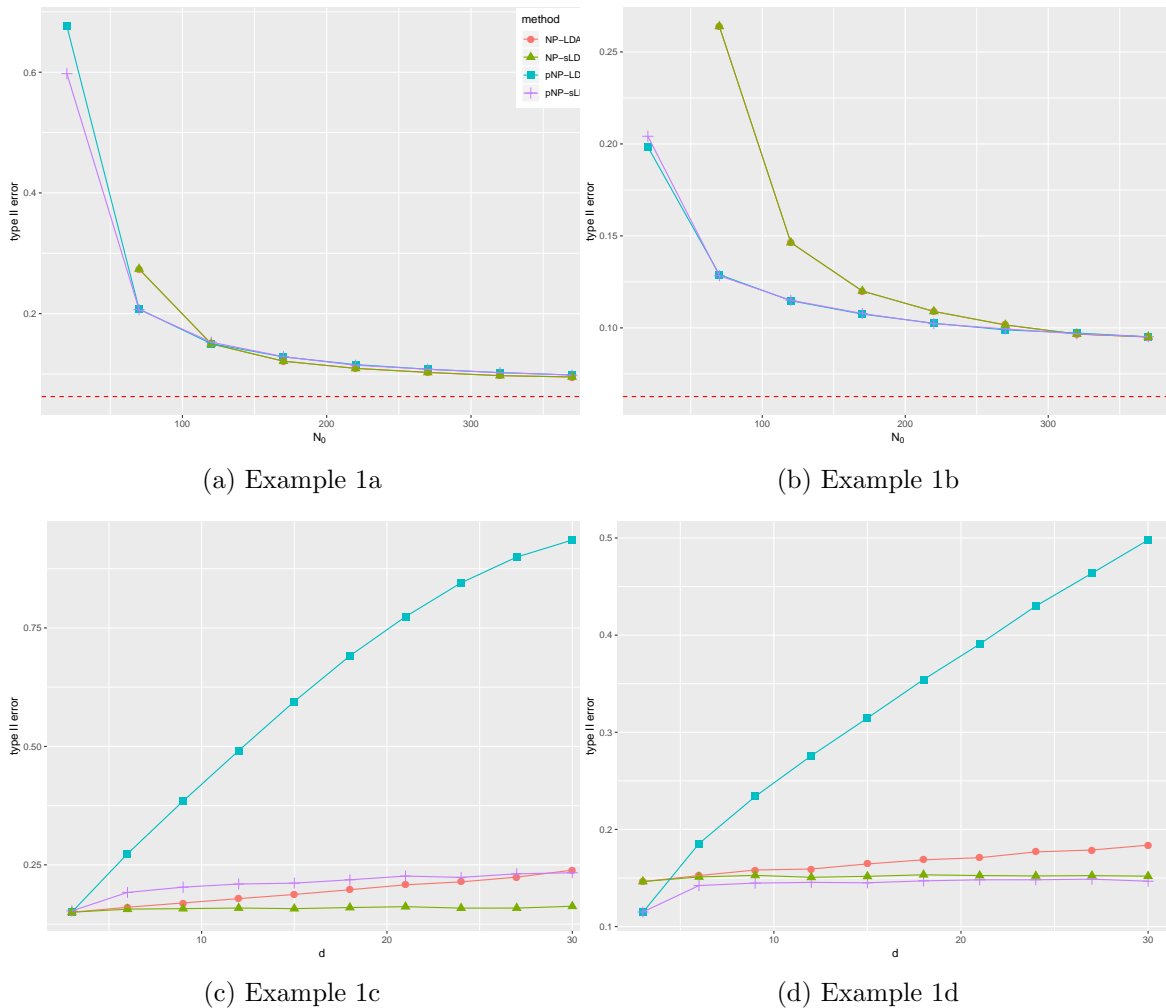
Example 1 *The data are generated from an LDA model with common covariance matrix Σ , where Σ is set to be an AR(1) covariance matrix with $\Sigma_{ij} = 0.5^{|i-j|}$ for all i and j . The $\beta^{\text{Bayes}} = \Sigma^{-1}\mu_d = 1.2 \times (1_{d_0}, 0_{d-d_0})^\top$, $\mu^0 = 0_d$, $d_0 = 3$. We set $\pi_0 = \pi_1 = 0.5$ and $\alpha = \delta_0 = 0.1$.*

(1a). $d = 3$, varying $N_0 = n_1 \in \{20, 70, 120, 170, 220, 270, 320, 370\}$.

(1b). $d = 3$, $n_1 = 500$, varying $N_0 \in \{20, 70, 120, 170, 220, 270, 320, 370\}$.

(1c). $N_0 = n_1 = 125$, varying $d \in \{3, 6, 9, 12, 15, 18, 21, 24, 26, 30\}$.

Figure 1: Example 1. Type II error of NP-LDA, NP-sLDA, pNP-LDA and pNP-sLDA vs. N_0 or d . The dashed red line represent the NP oracle.



(1d). $N_0 = 125$, $n_1 = 500$, varying $d \in \{3, 6, 9, 12, 15, 18, 21, 24, 26, 30\}$.

The results are summarized in the Figure 1. Several observations are made in order. First, from the first row of the figure (1a and 1b), we observe that when N_0 is very small, the implementable methods that can achieve the desired type I error control are pNP-LDA and pNP-sLDA; the NP umbrella algorithm based methods fail its minimum class 0 sample size requirement. Second, the type II error of all methods decreases as N_0 increases from the first row of the figure (1a and 1b), and increases when d increases from the second row of the figure (1c and 1d). Third, from the first row of the figure, we see pNP-LDA and pNP-sLDA have advantages over NP-LDA and NP-sLDA when the sample sizes are small. Finally, from the second row of the figure, we see that the nonparametric NP umbrella algorithm gains more and more advantages over the parametric thresholding rule (pNP-LDA) as d increases, since \widehat{C}_α^p specified in (19) can become loose when d is large. It is worth to mention that by

taking advantage of the sparse solution generated by **sLDA**, the performance of **pNP-sLDA** does not deteriorates as d increases and performs the best for Example 1d.

Example 2 *The data are generated from an LDA model with common covariance matrix Σ , where Σ is set to be AR(1) covariance matrix with $\Sigma_{ij} = 0.5^{|i-j|}$ for all i and j . The $\beta^{Bayes} = \Sigma^{-1}\mu_d = C_d \cdot \mathbf{1}_d^\top$, $\mu^0 = \mathbf{0}_d$. We set $\pi_0 = \pi_1 = 0.5$ and $\alpha = \delta_0 = 0.1$. Here C_d is a constant depending on d such that the oracle classifier always has type II error 0.112 for any choice of d .*

(2a). $d = 3$, varying $N_0 = n_1 \in \{20, 70, 120, 170, 220, 270, 320, 370\}$.

(2b). $d = 6$, varying $N_0 = n_1 \in \{20, 70, 120, 170, 220, 270, 320, 370\}$.

(2c). $N_0 = n_1 = 125$, varying $d \in \{3, 6, 9, 12, 15, 18, 21, 24, 26, 30\}$.

(2d). $N_0 = 125$, $n_1 = 500$, varying $d \in \{3, 6, 9, 12, 15, 18, 21, 24, 26, 30\}$.

Example 2 is a more challenging scenario where the oracle rule depends on all features. Similar observations as in Example 1 can be made for Example 2 from Figure 2. It is worth to mention that in this case, although we still see improvement of **pNP-sLDA** over **pNP-LDA** throughout all d , the performance of **pNP-sLDA** is dominated by **NP-LDA** and **NP-sLDA** for moderate sample sizes as all features are important.

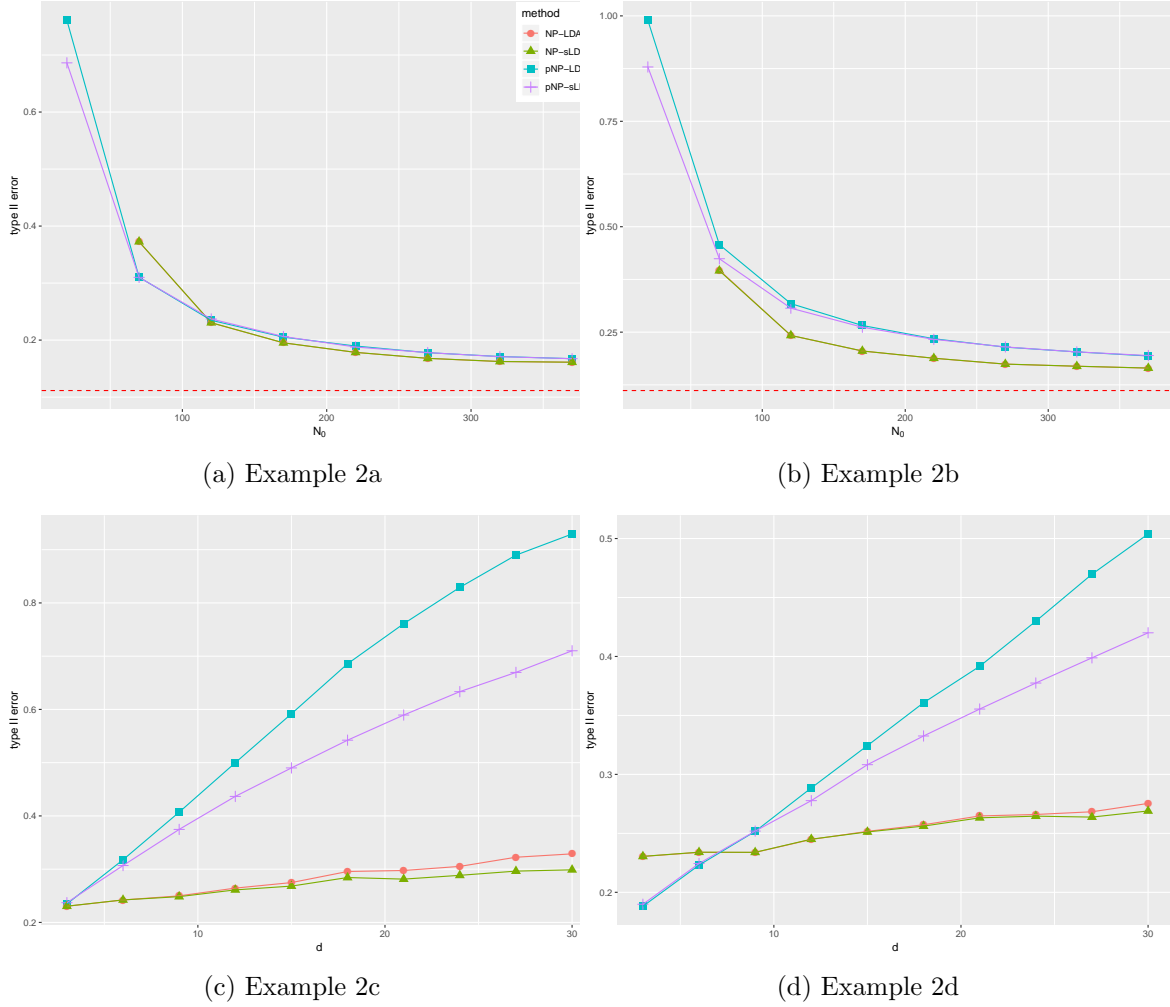
From Figure 1 (a)(b) and Figure 2(a)(b), we see that as sample sizes increase, the performance of **NP-sLDA** gets better and eventually dominates **pNP-sLDA** even though the latter takes advantage of the parametric model assumption in both training the scoring function and constructing the threshold. This might seem a little counter-intuitive at first glance. The reason is that the construction of threshold estimate \widehat{C}_α in **pNP-sLDA** relies on a high probability upper bound on the inaccessible model-specific oracle threshold (see Lemma 3). The construction of this upper bound involves bounding a quadratic form of Σ by $\lambda_{\max}(\Sigma)$ and studying the relations between $\lambda_{\max}(\Sigma)$ and $\lambda_{\max}(\widehat{\Sigma})$ without structural assumptions on Σ . Thus, the upper bound can be on the conservative side (i.e., larger than what is necessary) for a specific covariance structure. On the other hand, when the sample sizes are large, the number of candidate thresholds in **NP-sLDA** becomes large, then one can choose an order k in the nonparametric NP umbrella algorithm to make the violation rate $v(k)$ very close to δ_0 in equation (12). Thus the loss due to universal handling of covariance structure in **pNP-sLDA** may outweigh the loss due to discretization in **NP-sLDA**.

6.2. Simulation studies under high-dimensional settings

In Examples 3–5, we conduct simulations to compare the empirical performance of the proposed **NP-sLDA** and **pNP-sLDA** with other non-LDA based NP classifiers as well as the **sLDA** (Mai et al., 2012). In every simulation setting, the experiments are repeated 1,000 times.

Example 3 *The data are generated from an LDA model with common covariance matrix Σ , where Σ is set to be an AR(1) covariance matrix with $\Sigma_{ij} = 0.5^{|i-j|}$ for all i and j . The $\beta^{Bayes} = \Sigma^{-1}\mu_d = 0.556 \times (3, 1.5, 0, 0, 2, 0, \dots, 0)^\top$, $\mu^0 = \mathbf{0}^\top$, $d = 1,000$, and $N_0 = n_1 = 200$. Bayes error = 10% under $\pi_0 = \pi_1 = 0.5$. $\alpha = \delta_0 = 0.1$.*

Figure 2: Example 2. Type II error of NP-LDA, NP-sLDA, pNP-LDA and pNP-sLDA vs. N_0 or d . The dashed red line represents the NP oracle.



Example 4 The data are generated from an LDA model with common covariance matrix Σ , where Σ is set to be a compound symmetric covariance matrix with $\Sigma_{ij} = 0.5$ for all $i \neq j$ and $\Sigma_{ii} = 1$ for all i . The $\beta^{Bayes} = \Sigma^{-1}\mu_d = 0.551 \times (3, 1.7, -2.2, -2.1, 2.55, 0, \dots, 0)^\top$, $\mu^0 = 0^\top$, $d = 2,000$, and $N_0 = n_1 = 300$. Bayes error = 10% under $\pi_0 = \pi_1 = 0.5$. $\alpha = \delta_0 = 0.1$.

Example 5 Same as in Example 4, except $d = 3,000$, $N_0 = n_1 = 400$, and the $\beta^{Bayes} = \Sigma^{-1}\mu_d = 0.362 \times (3, 1.7, -2.2, -2.1, 2.55, 0, \dots, 0)^\top$. Bayes error = 20% under $\pi_0 = \pi_1 = 0.5$. $\alpha = 0.2$ and $\delta_0 = 0.1$.

In Examples 3–5, we compare the empirical type I/II error performance of NP-sLDA, NP-penlog (*penlog* stands for penalized logistic regression), NP-svm, pNP-sLDA, and sLDA on a large test data set of size 20,000 that consist of 10,000 observations from each class.

Table 1: Violation rate and type II error for Examples 3, 4 and 5 over 1,000 repetitions.

		NP-sLDA	NP-penlog	NP-svm	pNP-sLDA	sLDA
Ex 3	violation rate	.068	.055	.054	.001	.764
	type II error (mean)	.189	.205	.621	.220	.104
	type II error (sd)	.057	.063	.077	.052	.010
Ex 4	violation rate	.073	.081	.081	.000	1.000
	type II error (mean)	.246	.255	.615	.824	.129
	type II error (sd)	.051	.053	.070	.121	.010
Ex 5	violation rate	.079	.088	.099	.000	.997
	type II error (mean)	.332	.334	.584	.748	.231
	type II error (sd)	.044	.044	.045	.128	.012

In all NP methods, τ , the class 0 split proportion, is fixed at 0.5. The choices for α in these examples match the corresponding Bayes errors, so that the comparison between NP and classical methods does not obviously favor the former.

Table 1 indicates that while the classical sLDA method cannot control the type I error violation rate under δ_0 , all the NP classifiers are able to do so. In addition, among the four NP classifiers, NP-sLDA gives the smallest mean type II error. pNP-sLDA performs reasonably well in Example 3, where the NP oracle rule is extremely sparse. In Examples 4 and 5, however, the threshold \widehat{C}_α^p becomes overly large as a result of more selected features due to the less sparse oracle, thus leading to an overly conservative classifier with 0 violation rate ².

6.3. Adaptive sample splitting

By explanations in the last paragraph of Section 5, the adaptive splitting scheme does not affect the type I error control objective. Examples 6 and 7 investigate the power enhancement as a result of the adaptive splitting scheme over the default half-half choice. These examples include an array of situations, including low- and high-dimensional settings ($d = 20$ and 1,000), balanced and imbalanced classes ($N_0 : n_1 = 1 : 1$ to $1 : 256$), and small to medium sample sizes ($N_0 = 100$ to 500).

Example 6 *Same as in Example 3, except taking the following sample sizes.*

(6a). $N_0 = 100$ and varying $n_1/N_0 \in \{1, 2, 4, 8, 16, 32, 64, 128, 256\}$.

(6b). Varying $n_1 = N_0 \in \{100, 150, 200, 250, 300, 350, 400, 450, 500\}$.

Example 7 *Same as in Example 3, except that $d = 20$, $N_0 = 100$ and varying $n_1/N_0 \in \{1, 2, 4, 8, 16\}$.*

2. Strictly speaking, the observed type I error violation rate is only an approximation to the real violation rate. The approximation is two-fold: i). in each repetition of an experiment, the population type I error is approximated by empirical type I error on a large test set; ii). the violation rate should be calculated based on infinite repetitions of the experiment, but we only calculate it based on 1,000 repetitions. However, such approximation is unavoidable in numerical studies.

Note that Examples 6a and 6b each includes 9 different simulation settings, and Example 7 includes 5. For each simulation setting, we generate 1,000 (training) datasets and a common test set of size 100,000 from class 1. Only class 1 test data are needed because only type II error is investigated in these examples. In each simulation setting, we train 10 NP classifiers of the same base algorithm using each of the 1,000 datasets. Nine of these 10 NP classifiers use fixed split proportions in $\{.1, \dots, .9\}$, and the last one uses adaptive split proportion using $K = 5$. Overall in Examples 6 and 7, we set $\alpha = \delta_0 = 0.1$, and train an enormous number of NP classifiers. For instance, in Example 6a, we train $9 \times 1,000 \times 10 = 90,000$ NP-sLDA classifiers, and the same number of NP classifiers for any other base algorithm under investigation. We fix the thresholding rule as the NP umbrella algorithm in this subsection.

For each simulation setting, denote by $\tilde{R}_1(\cdot)$ the empirical type II error on the test set. We fix a simulation setting so that we do not need to have overly complex sub or sup indexes in the following discussion. Denote by $\hat{h}_{i,b,\tau}$ an NP classifier with base algorithm b , trained on the i th dataset ($i \in \{1, \dots, 1000\}$) using split proportion τ . This classifier also depends on users' choices of α and δ_0 , but we suppress these dependencies here to highlight our focus. In fixed proportion scenarios, $\tau \in \{.1, \dots, .9\}$. Let $\tau^{\text{ada}}(j, b)$ represent the adaptive split proportion trained on the j -th dataset with base algorithm b using adaptive splitting scheme described in Section 5. Therefore, $\hat{h}_{i,b,\tau^{\text{ada}}(j,b)}$ refers to the NP classifier with base algorithm b , trained on the i -th dataset using the split proportion $\tau^{\text{ada}}(j, b)$ pre-determined in the j -th dataset, where $i, j \in \{1, \dots, 1000\}$. Let $\text{Ave}_{b,\tau}$ and $\text{Ave}_{b,\hat{\tau}}$ be our performance measures for fix proportion and adaptive proportion respectively, which are defined by,

$$\text{Ave}_{b,\tau} = \frac{1}{1000} \sum_{i=1}^{1000} \tilde{R}_1(\hat{h}_{i,b,\tau}), \text{ and } \text{Ave}_{b,\hat{\tau}} = \text{median}_{j=1, \dots, 1000} \left(\frac{1}{1000} \sum_{i=1}^{1000} \tilde{R}_1(\hat{h}_{i,b,\tau^{\text{ada}}(j,b)}) \right).$$

While the meaning of the measure $\text{Ave}_{b,\tau}$ is almost self-evident, $\text{Ave}_{b,\hat{\tau}}$ deserves some elaboration. As we explained in the last paragraph of Section 5, the adaptive splitting scheme returns a proportion based on one realization of \mathcal{S} , and then we adopt it in all subsequent realizations. Let

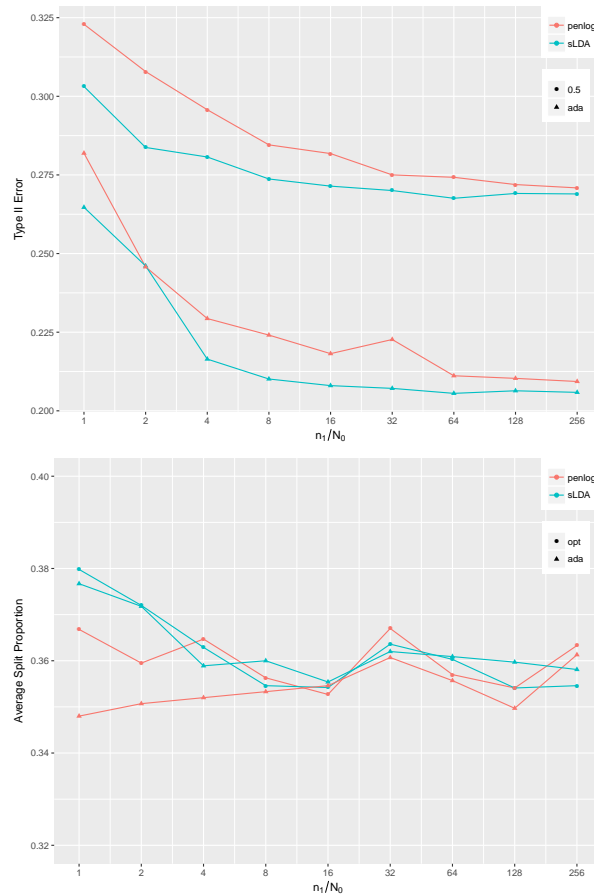
$$w_b(j) = \frac{1}{1000} \sum_{i=1}^{1000} \tilde{R}_1(\hat{h}_{i,b,\tau^{\text{ada}}(j,b)}),$$

then $w_b(j)$ is a performance measure of the adaptive scheme if the proportion is returned from training on the j -th dataset. To account for the variation among $w_b(j)$'s for different choices of j , we take the median over $w_b(j)$'s as our final measure. Also, we denote the average of adaptively selected proportions by $\tau_{b,\text{ada}} = \frac{1}{1000} \sum_{j=1}^{1000} \tau^{\text{ada}}(j, b)$, and define the average optimal split proportion $\tau_{b,\text{opt}}$ by

$$\tau_{b,\text{opt}} = \frac{1}{1000} \sum_{i=1}^{1000} \arg \min_{\tau \in \{.1, \dots, .9\}} \tilde{R}_1(\hat{h}_{i,b,\tau}).$$

With Example 6, we investigate i). the effectiveness (in terms of type II error) of the adaptive splitting strategy compared to a fixed half-half split, illustrated by the left panels of Figures 3 and 4; ii). how close is $\tau_{b,\text{ada}}$ compared to $\tau_{b,\text{opt}}$, illustrated by the right panels

Figure 3: Example 6a. Left panel: type II error ($Ave_{b,.5}$ and $Ave_{b,\hat{\tau}}$) of NP-sLDA and NP-penlog vs. n_1/N_0 ; Right panel: average split proportion ($\tau_{b,ada}$ and $\tau_{b,opt}$) vs. n_1/N_0 . N_0 is fixed to be 100 for both panels.




of Figures 3 and 4; iii). how the class imbalance affects NP-sLDA and NP-penlog, illustrated by both panels of Figure 3; and iv). how the absolute class 0 sample size affects NP-sLDA and NP-penlog, illustrated by both panels of Figure 4.

In Figure 3 (Example 6a), the left panel presents the trend of type II errors ($Ave_{b,.5}$ and $Ave_{b,\hat{\tau}}$) as the sample size ratio n_1/N_0 increases from 1 to 256 for fixed $N_0 = 100$. For both NP-penlog and NP-sLDA, type II error decreases as n_1/N_0 increases from 1 to 16 and gradually stabilizes afterwards. Neither NP-penlog nor NP-sLDA suffers from training on imbalanced classes. In terms of type II error performance, the adaptive splitting strategy significantly improves over the fixed split proportion 0.5. The right panel of Figure 3 shows that, on average the adaptive split proportion is very close to the optimal one throughout all sample size ratios.

In Figure 4 (Example 6b), the left panel presents the trend of type II errors ($Ave_{b,.5}$ and $Ave_{b,\hat{\tau}}$) as the class 0 sample size N_0 ($n_1 = N_0$) increases from 100 to 500, indicating that

Figure 4: Example 6b. Left panel: type II error ($Ave_{b,.5}$ and $Ave_{b,\hat{\tau}}$) of NP-sLDA and NP-penlog vs. N_0 ; Right panel: average split proportion ($\tau_{b,ada}$ and $\tau_{b,opt}$) vs. N_0 . $n_1 = N_0$ for both panels.



type II error clearly benefits from increasing training sample sizes of both classes. For the same base algorithm, the adaptive splitting strategy significantly improves over the fixed split proportion 0.5 for N_0 and n_1 small, although the improvement diminishes as both sample sizes become large. The right panel of Figure 4 shows that on average, the adaptive split proportion is very close to the optimal one throughout all sample sizes. Furthermore, the average optimal split proportion seems to increase as N_0 increases in general. An intuition is that when N_0 is smaller, a higher proportion of class 0 observations is needed for threshold estimate, to guarantee the type I error violation rate control.

With Example 7, we investigate the interaction between adaptive splitting strategy and multiple random splits on different NP classifiers. Multiple random splits of class 0 observations were proposed in the NP umbrella algorithm in Tong et al. (2018) to increase the stability of the type II error performance. When an NP classifier uses $M > 1$ multiple splits, each split will result in a classifier, and the final prediction rule is a majority

Figure 5: Example 7. Type II error ($Ave_{b, .5}$ and $Ave_{b, \hat{\tau}}$) vs. sample size ratios for four NP classifiers (NP-sLDA, NP-penlog, NP-svm, NP-randomforest), with both multiple random splits ($M = 11$) and single random split. $N_0 = 100$ for all sample size ratios.




Table 2: Example 7. Average computational cost (in seconds) for four NP-classifiers (NP-sLDA, NP-penlog, NP-randomforest, NP-svm) over 1,000 repetitions (standard deviation in parentheses).

n_1/N_0	NP-sLDA	NP-penlog	NP-randomforest	NP-svm
1	1.17(0.14)	3.58(0.58)	1.49(4.83)	33.81(2.08)
2	1.19(0.13)	5.24(1.03)	1.43(0.16)	36.09(1.98)
4	1.19(0.14)	7.63(1.79)	2.06(0.10)	41.44(2.38)
8	1.22(0.09)	11.11(2.31)	3.43(0.19)	53.25(6.21)
16	1.30(0.08)	16.08(3.79)	6.92(0.25)	84.64(5.28)

vote of these classifiers. Figure 5 shows the trend of type II error of NP-sLDA, NP-penlog, NP-randomforest, and NP-svm, as the sample size ratio n_1/N_0 increases from 1 to 16 while keeping $N_0 = 100$. For each base algorithm, four scenarios are considered: (fixed 0.5 split proportion, single split), (adaptive split proportion, single split), (fixed 0.5 split proportion, multiple splits), and (adaptive split proportion, multiple splits). Figure 5 suggests the following interesting findings: i). type II error decreases for NP-sLDA and NP-penlog but increases for NP-randomforest and NP-svm, as a function of n_1/N_0 while keeping N_0 constant; ii). with both fixed 0.5 split proportion and adaptive splitting strategy, performing multiple splits leads to a smaller type II error compared with their single split counterparts; iii). for both single split and multiple splits, the adaptive split always improves upon the fixed 0.5 split proportion; iv). NP-svm and NP-randomforest are affected by the imbalance scenario, and one might consider downsampling or upsampling methods before applying an NP algorithm; and v). adding multiple splits to the adaptive splitting strategy leads to a further reduction on the type II error. Nevertheless, the reduction in type II error from the adaptive splitting scheme alone is much larger than the marginal gain from adding multiple splits on top of it. Therefore, when computation power is limited, one should implement the adaptive splitting scheme before considering multiple splits.

Lastly, from Table 2, we would like to point out NP-sLDA is the fastest method to compute among the four NP classifiers with more evident advantages as the sample size increases.³

6.4. Real data analysis

We study two high-dimensional datasets in this subsection. The first is a neuroblastoma dataset containing $d = 43,827$ gene expression measurements from $N = 498$ neuroblastoma samples generated by the Sequencing Quality Control (SEQC) consortium (Wang et al., 2014). The samples fall into two classes: 176 high-risk (HR) samples and 322 non-HR samples. It is usually understood that misclassifying an HR sample as non-HR will have more severe consequences than the other way around. Formulating this problem under the NP classification framework, we label the HR samples as class 0 observations and the non-

3. All numerical experiments were performed on HP Enterprise XL170r with CPU E5-2650v4 (2.20 GHz) and 16 GB memory.

HR samples as class 1 observations and, use all gene expression measurements as features to perform classification. We set $\alpha = \delta_0 = 0.1$, and compare NP-sLDA with NP-penlog, NP-randomforest and NP-svm. We randomly split the dataset 1,000 times into a training set (70%) and a test set (30%), and then train the NP classifiers on each training data and compute their empirical type I and type II errors over the corresponding test data. We consider each fixed split proportion in $\{.1, .2, .3, .4, .5, .6, .7, .8\}$ as well as the adaptive splitting strategy. Here, the split proportion 0.9 is not considered since it leads to a left-out sample size which is too small to control the type I error at the given α and δ_0 values. Figure 6 indicates that the average type I error is less than α across different split proportions for all four methods considered. Regarding the average type II error, NP-sLDA has the smallest values for a wide range of split proportions. In particular, the smallest average type II error for NP-sLDA corresponds to split proportion 0.4. The average location of the split proportion chosen by the adaptive splitting scheme would lead to a type II error close to the minimum. This demonstrates that the adaptive splitting scheme works well for different NP classifiers. Lastly, we note that for a specific splitting proportion, the median computation time is 213.95 seconds for pNP-sLDA vs. 1717.69 seconds for NP-randomforest over 1,000 random splits.

The second dataset is a high-dimensional breast cancer dataset ($d = 22,215, N = 118$) (Chin et al., 2006) with gene expression measurements of subjects that fall into positive (class 0, $N = 75$) and negative (class 1, $n_1 = 43$) groups. We set $\delta_0 = 0.1$ and vary α from 0.05 to 0.15, and compare the performance of NP-sLDA with pNP-sLDA. We randomly split the dataset 1,000 times into a training set (2/3) and a test set (1/3), train the two methods on the training set, and compute the empirical type I and type II errors on the corresponding test set. Due to the limited sample size, it is clear from Table 3 that when α is small, the minimum sample size requirement for the NP umbrella algorithm is not satisfied. The pNP-sLDA, on the other hand, took advantage of the parametric assumption and the corresponding violation rate is under δ_0 throughout all choices of α .


Table 3: Violation rate and average type II errors over 1,000 replications for the breast cancer data. Here, $\delta_0 = 0.1$ and α varies from 0.05 to 0.15.

α	vio (NP)	vio (pNP)	type II (NP)	type II (pNP)
0.05	NA	0.021	NA	0.663
0.06	NA	0.025	NA	0.626
0.07	NA	0.029	NA	0.594
0.08	NA	0.015	NA	0.565
0.09	0.111	0.015	0.448	0.536
0.1	0.111	0.015	0.448	0.515

7. Discussion

This work develops Neyman-Pearson (NP) classification theory and methodology under parametric model assumptions. Most specifically, based on the linear discriminant analysis

Figure 6: The average type I and type II errors vs. splitting proportion on the neuroblastoma data set for NP-sLDA, NP-penlog, NP-randomforest and NP-svm over 1,000 random splits. The “*” point on each line represents the average split proportion chosen by adapting splitting.



(LDA) model, we develop a new parametric model-based thresholding rule for high probability type I error control, and this complements the nonparametric NP umbrella algorithm when the minimum sample size requirement of the latter is not met. In practice, when the minimum sample size requirement is met and the scoring function depends on more than a few features, the NP umbrella algorithm is still recommended based on better empirical performance. For future work, it would be interesting to investigate NP classifiers under other parametric settings, such as quadratic discriminant analysis (QDA) model and heavy-tailed distributions which are appropriate to model financial data. We expect that new model-specific thresholding rules will be developed for NP classification.

Acknowledgement

The authors would like to thank the Action Editor and three anonymous referees for many constructive comments which greatly improved the paper. This work was partially supported by National Science Foundation grants DMS-1554804 and DMS-1613338, and National Institutes of Health grant R01 GM120507.

Appendix A. Neyman-Pearson Lemma

The oracle classifier under the NP paradigm arises from its close connection to the Neyman-Pearson Lemma in statistical hypothesis testing. Hypothesis testing bears strong resemblance to binary classification if we assume the following model. Let P_1 and P_0 be two *known* probability distributions on $\mathcal{X} \subset \mathbb{R}^d$. Assume that $Y \sim \text{Bern}(\zeta)$ for some $\zeta \in (0, 1)$, and the conditional distribution of X given Y is P_Y . Given such a model, the goal of statistical hypothesis testing is to determine if we should reject the null hypothesis that X was generated from P_0 . To this end, we construct a randomized test $\phi : \mathcal{X} \rightarrow [0, 1]$ that rejects the null with probability $\phi(X)$. Two types of errors arise: type I error occurs when P_0 is rejected yet $X \sim P_0$, and type II error occurs when P_0 is not rejected yet $X \sim P_1$. The Neyman-Pearson paradigm in hypothesis testing amounts to choosing ϕ that solves the following constrained optimization problem

$$\text{maximize } \mathbb{E}[\phi(X)|Y = 1], \text{ subject to } \mathbb{E}[\phi(X)|Y = 0] \leq \alpha,$$

where $\alpha \in (0, 1)$ is the significance level of the test. A solution to this constrained optimization problem is called *a most powerful test* of level α . The Neyman-Pearson Lemma gives mild sufficient conditions for the existence of such a test.

Lemma 16 (Neyman-Pearson Lemma) *Let P_1 and P_0 be two probability measures with densities f_1 and f_0 respectively, and denote the density ratio as $r(x) = f_1(x)/f_0(x)$. For a given significance level α , let C_α be such that $P_0\{r(X) > C_\alpha\} \leq \alpha$ and $P_0\{r(X) \geq C_\alpha\} \geq \alpha$. Then, the most powerful test of level α is*

$$\phi_\alpha^*(X) = \begin{cases} 1 & \text{if } r(X) > C_\alpha, \\ 0 & \text{if } r(X) < C_\alpha, \\ \frac{\alpha - P_0\{r(X) > C_\alpha\}}{P_0\{r(X) = C_\alpha\}} & \text{if } r(X) = C_\alpha. \end{cases}$$

Under mild continuity assumption, we take the *NP oracle classifier*

$$\phi_\alpha^*(x) = \mathbb{I}\{f_1(x)/f_0(x) > C_\alpha\} = \mathbb{I}\{r(x) > C_\alpha\}, \quad (22)$$

as our plug-in target for NP classification.

Appendix B. Additional Lemmas and Propositions

Lemma 17 (Hsu et al. (2012)) *Let $A \in R^{m \times n}$ be a matrix, and let $\Sigma := A^\top A$. Let $x = (x_1, \dots, x_n)^\top$ be an isotropic multivariate Gaussian random vector with mean zero. For all $t > 0$,*

$$\mathbb{P}\left(\|Ax\|^2 > \text{tr}(\Sigma) + 2\sqrt{\text{tr}(\Sigma^2)t} + 2\|\Sigma\|t\right) \leq e^{-t}.$$

Lemma 18 *Recall that $\beta^{\text{Bayes}} = \Sigma^{-1}\mu_d = \Sigma^{-1}(\mu^1 - \mu^0)$ and $A = \{j : \{\tilde{\Sigma}^{-1}\mu_d\}_j \neq 0\}$. Denote by $\beta^* = (\Sigma_{AA})^{-1}(\mu_A^1 - \mu_A^0)$ and $\tilde{\beta}^{\text{Bayes}}$ by letting $\tilde{\beta}_A^{\text{Bayes}} = \beta^*$ and $\tilde{\beta}_{A^c}^{\text{Bayes}} = 0$. Then $\tilde{\beta}^{\text{Bayes}} = \beta^{\text{Bayes}}$.*

Recall these notations for the following lemma: let $\mathcal{S}_0 = \{x_1^0, \dots, x_{n_0}^0\}$ be an i.i.d. sample of class 0 of size n_0 and $\mathcal{S}_1 = \{x_1^1, \dots, x_{n_1}^1\}$ be an i.i.d. sample of class 1 of size n_1 , and $n = n_0 + n_1$. We use \mathcal{S}_0 and \mathcal{S}_1 to find an estimate of β^{Bayes} . Let \tilde{X} be the $(n \times d)$ centered predictor matrix, whose column-wise mean is zero, which can be decomposed into \tilde{X}^0 , the $(n_0 \times d)$ centered predictor matrix based on class 0 observations and \tilde{X}^1 , the $(n_1 \times d)$ centered predictor matrix based on class 1 observations. Let $C^{(n)} = (\tilde{X})^\top \tilde{X}/n$, then

$$C^{(n)} = \frac{n_0}{n} \hat{\Sigma}^0 + \frac{n_1}{n} \hat{\Sigma}^1,$$

where $\hat{\Sigma}^0 = (\tilde{X}^0)^\top \tilde{X}^0/n_0$, and $\hat{\Sigma}^1 = (\tilde{X}^1)^\top \tilde{X}^1/n_1$.

Lemma 19 *Suppose there exists $c > 0$ such that $\Sigma_{jj} \leq c$ for all $j = 1, \dots, d$. There exist constants ε_0 and c_1, c_2 such that for any $\varepsilon \leq \varepsilon_0$ we have,*

$$\mathbb{P}\left(|(\hat{\mu}_j^1 - \hat{\mu}_j^0) - (\mu_j^1 - \mu_j^0)| \geq \varepsilon\right) \leq 2 \exp(-n_0 \varepsilon^2 c_2) + 2 \exp(-n_1 \varepsilon^2 c_2), \text{ for } j = 1, \dots, d. \quad (23)$$

$$\mathbb{P}\left(|\hat{\Sigma}_{ij}^l - \Sigma_{ij}| \geq \varepsilon\right) \leq 2 \exp(-n_l \varepsilon^2 c_1), \text{ for } l = 0, 1, i, j = 1, \dots, d. \quad (24)$$

$$\mathbb{P}\left(\|\hat{\Sigma}_{AA}^l - \Sigma_{AA}\|_\infty \geq \varepsilon\right) \leq 2s^2 \exp(-n_l s^{-2} \varepsilon^2 c_1). \quad (25)$$

$$\mathbb{P}\left(\|\hat{\Sigma}_{A^c A}^l - \Sigma_{A^c A}\|_\infty \geq \varepsilon\right) \leq (d-s)s \exp(-n_l s^{-2} \varepsilon^2 c_1). \quad (26)$$

$$\mathbb{P}\left(\|(\hat{\mu}^1 - \hat{\mu}^0) - (\mu^1 - \mu^0)\|_\infty \geq \varepsilon\right) \leq 2d \exp(-n_0 \varepsilon^2 c_2) + 2d \exp(-n_1 \varepsilon^2 c_2). \quad (27)$$

$$\mathbb{P}\left(\|(\hat{\mu}_A^1 - \hat{\mu}_A^0) - (\mu_A^1 - \mu_A^0)\|_\infty \geq \varepsilon\right) \leq 2s \exp(-n_0 \varepsilon^2 c_2) + 2s \exp(-n_1 \varepsilon^2 c_2). \quad (28)$$

$$\mathbb{P}\left(|C_{ij}^{(n)} - \Sigma_{ij}| \geq \varepsilon\right) \leq 2 \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4n_0}\right) + 2 \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4n_1}\right). \quad (29)$$

$$\mathbb{P}\left(|C_{AA}^{(n)} - \Sigma_{AA}| \geq \varepsilon\right) \leq 2s^2 \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4s^2 n_0}\right) + 2s^2 \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4s^2 n_1}\right). \quad (30)$$

$$\mathbb{P}\left(|C_{A^c A}^{(n)} - \Sigma_{A^c A}| \geq \varepsilon\right) \leq (d-s)s \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4s^2 n_0}\right) + (d-s)s \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4s^2 n_1}\right). \quad (31)$$

Lemma 20 Recall that $\kappa = \|\Sigma_{A^c A}(\Sigma_{AA})^{-1}\|_\infty$, $\varphi = \|(\Sigma_{AA})^{-1}\|_\infty$ and $\Delta = \|\mu_A^1 - \mu_A^0\|_\infty$. Let $C_{A^c A}^{(n)} = \frac{n_0}{n}(\tilde{X}_{A^c}^0)^\top \tilde{X}_A^0 + \frac{n_1}{n}(\tilde{X}_{A^c}^1)^\top \tilde{X}_A^1 = \frac{n_0}{n}\tilde{\Sigma}_{A^c A}^0 + \frac{n_1}{n}\tilde{\Sigma}_{A^c A}^1$, and $C_{AA}^{(n)} = \frac{n_0}{n}\tilde{\Sigma}_{AA}^0 + \frac{n_1}{n}\tilde{\Sigma}_{AA}^1$. There exist constants c_1 and ε_0 such that for any $\varepsilon \leq \min(\varepsilon_0, 1/\varphi)$, we have

$$\mathbb{P}\left(\|C_{A^c A}^{(n)}(C_{AA}^{(n)})^{-1} - \Sigma_{A^c A}(\Sigma_{AA})^{-1}\|_\infty \geq (\kappa + 1)\varepsilon\varphi(1 - \varphi\varepsilon)^{-1}\right) \leq f(d, s, n_0, n_1, \varepsilon),$$

where $f(d, s, n_0, n_1, \varepsilon) = (d+s)s \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4s^2 n_0}\right) + (d+s)s \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4s^2 n_1}\right)$, and $n = n_0 + n_1$.

Lemma 21 Let $\tilde{\mathcal{C}}^0$ be in Equation (39), \mathcal{C} as in Lemma 10, and $\tilde{X}_A = \Sigma_{AA}^{-1/2} X_A$. Assume $\lambda_m = \lambda_{\min}(\Sigma_{AA}^{-1/2})$ is bounded from below, then we have

$$P_0\{C_\alpha^{**} \leq s^*(X) \leq C_\alpha^{**} + \delta | X \in \mathcal{C}\} \geq (1 - \delta_3)P_0(C_\alpha^{**} \leq (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1/2} \tilde{X}_A \leq C_\alpha^{**} + \delta | \tilde{\mathcal{C}}^0),$$

where $\delta_3 = \exp\{-(n_0 \wedge n_1)^{1/2}\}$.

Lemma 22 Let us denote $a = \Sigma_{AA}^{-1/2}(\mu_A^1 - \mu_A^0)$. Assume there exist $M > 0$ such that the following conditions hold:

- i) $C_\alpha^{**} - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0 \in (C^1, C^2)$ for some constants C^1, C^2 .
- ii) When $s = 1$, a is a scalar. $f_{\mathcal{N}(0, |a|)}$ is bounded below on interval $(C^1, C^2 + \delta^*)$ by M .
- iii) When $s = 2$, $a = (a_1, a_2)^\top$ is a vector.

$$\left(\frac{1}{\sqrt{2\pi}|a_1|(\frac{a_2^2}{a_1^2} + 1)} \exp\left\{-\frac{a_1^2 + a_2^2 - a_2^2 a_1^2 t^2}{a_1^2(a_1^2 + a_2^2)}\right\}\right) \left(2\Phi\left(\sqrt{\tilde{L}^2 - \frac{a_1^2 + a_2^2 - a_2^2 a_1^2 t^2}{a_1^2(a_1^2 + a_2^2)}}\right) - 1\right)$$

is bounded below on interval $t \in (C^1, C^2 + \delta^*)$ by M .

Then, for $s \leq 2$, for any $\delta \in (0, \delta^*)$, there exists M_1 which is a constant depending on M , such that the following inequality holds

$$P_0(C_\alpha^{**} \leq (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1/2} \tilde{X}_A \leq C_\alpha^{**} + \delta | \tilde{\mathcal{C}}^0) \geq M_1 \delta.$$

Proposition 23 Suppose that $\lambda_{\min}(\Sigma_{AA}^{-1/2})$, the minimum eigenvalue of $\Sigma_{AA}^{-1/2}$, is bounded from below. Let us denote $a = \Sigma_{AA}^{-1/2}(\mu_A^1 - \mu_A^0)$. Let us also assume that there exists $M > 0$ such that the following conditions hold:

- i) $C_{\alpha}^{**} - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0 \in (C^1, C^2)$ for some constants C^1, C^2 .
- ii) When $s = 1$, a is a scalar. $f_{\mathcal{N}(0,|a|)}$ is bounded below on interval $(C^1 - \delta^*, C^2 + \delta^*)$ by M .
- iii) Let $\tilde{L} = \lambda_{\min}(\Sigma_{AA}^{-1/2}) c_1' s^{1/2} (n_0 \wedge n_1)^{1/4}$. When $s = 2$, $a = (a_1, a_2)$ is a vector.

$$\left(\frac{1}{\sqrt{2\pi}|a_1| \left(\frac{a_2^2}{a_1^2} + 1\right)} \exp\left\{-\frac{a_1^2 + a_2^2 - a_2^2 a_1^2}{a_1^2(a_1^2 + a_2^2)} t^2\right\} \right) \left(2\Phi\left(\sqrt{\tilde{L}^2 - \frac{a_1^2 + a_2^2 - a_2^2 a_1^2}{a_1^2(a_1^2 + a_2^2)} t^2}\right) - 1 \right)$$

is bounded below on interval $t \in (C^1 - \delta^*, C^2 + \delta^*)$ by M .

Then for $s \leq 2$, the function $s^*(\cdot)$ satisfies conditional detection condition restricted to \mathcal{C} of order $\underline{\gamma} = 1$ with respect to P_0 at the level $(C_{\alpha}^{**}, \delta^*)$. In other words, Assumption 2 is satisfied.

Appendix C. Simulation for the control of largest empirical eigenvalues

Here, we present a simulation to verify the inequality on the largest eigenvalue of the sample covariance matrix in relation to that of the population covariance matrix, as stated in (17).

Example 8 The data are generated from an LDA model with common covariance matrix Σ . We set $\epsilon = 1e - 3$ and consider the 8 combinations for the following three factors and vary $N_0 = n_1 \in \{20, 40, 60, 80, 100, 120, 140, 160, 180, 200\}$. $\mu^0 = 0_d$.

(8a). The dimension $d \in \{3, 10\}$. $\mu^1 = 1.16 \cdot 1_3$ when $d = 3$ and $\mu^1 = 0.75 \cdot 1_{10}$ when $d = 10$.

(8b). Σ is set to be an AR(1) covariance matrix with $\Sigma_{ij} = \rho^{|i-j|}$ or an CS (compound symmetry) covariance matrix with $\Sigma_{ii} = 1$ and $\Sigma_{ij} = \rho$ for all $i \neq j$.

(8c). The correlation parameter $\rho \in \{.5, .9\}$.

From Table 4, it is clear that the bound is satisfied with very high probability across all scenarios considered.

Appendix D. Proofs

Proof [Proof of Lemma 8] By Corollary 6, $k^* \leq k'$. This implies that $R_0(\hat{\phi}_{k^*}) \geq R_0(\hat{\phi}_{k'})$. Moreover, by Lemma 7, for any $\delta'_0 \in (0, 1)$ and $n'_0 \geq 4/(\alpha\delta_0)$,

$$\mathbb{P}\left(|R_0(\hat{\phi}_{k'}) - R_0(\phi_{\alpha}^*)| > \xi_{\alpha, \delta_0, n'_0}(\delta'_0)\right) \leq \delta'_0.$$

Let $\mathcal{E}_0 = \{R_0(\hat{\phi}_{k^*}) \leq \alpha\}$ and $\mathcal{E}_1 = \{|R_0(\hat{\phi}_{k'}) - R_0(\phi_{\alpha}^*)| \leq \xi_{\alpha, \delta_0, n'_0}(\delta'_0)\}$. On the event $\mathcal{E}_0 \cap \mathcal{E}_1$, we have

$$\alpha = R_0(\phi_{\alpha}^*) \geq R_0(\hat{\phi}_{k^*}) \geq R_0(\hat{\phi}_{k'}) \geq R_0(\phi_{\alpha}^*) - \xi_{\alpha, \delta_0, n'_0}(\delta'_0),$$

This implies that

$$|R_0(\hat{\phi}_{k^*}) - R_0(\phi_{\alpha}^*)| \leq \xi_{\alpha, \delta_0, n'_0}(\delta'_0).$$

Table 4: The probability that the bound stated in (17) is satisfied over 1000 replications.

	Covariance	N_0										
		ρ	20	40	60	80	100	120	140	160	180	200
$d = 3$	AR(1)	.5	1.0	1.0	1.0	1.0	1.0	.999	.999	1.0	.999	1.0
		.9	1.0	1.0	1.0	1.0	1.0	.999	.999	1.0	.999	1.0
	CS	.5	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
		.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
$d = 10$	AR(1)	.5	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
		.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	CS	.5	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
		.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0

■

Proof [Proof of Lemma 10] Note that $\Sigma_{AA}^{-1/2}(X_A - \mu_A^0) \sim \mathcal{N}(0, I_s)$. By Lemma 17, for all $t > 0$,

$$P_0 \left(\|X_A - \mu_A^0\|^2 > \text{tr}(\Sigma_{AA}) + 2\sqrt{\text{tr}(\Sigma_{AA}^2)t} + 2\|\Sigma_{AA}\|t \right) \leq e^{-t}.$$

For $t = (n_0 \wedge n_1)^{1/2} (> 1)$, the above inequality implies there exists some $c_1'' > 0$ such that

$$P_0(\|X_A - \mu_A^0\|^2 > c_1''st) \leq e^{-t}.$$

Similarly, $P_1(\|X_A - \mu_A^1\|^2 > c_1''st) \leq e^{-t}$. Let $\mathcal{C}^0 = \{X : \|X_A - \mu_A^0\|^2 \leq c_1''st\}$ and $\mathcal{C}^1 = \{X : \|X_A - \mu_A^1\|^2 \leq c_1''st\}$. There exists some $c_1' > 0$, such that both \mathcal{C}^0 and \mathcal{C}^1 are subsets of $\mathcal{C} = \{X : \|X_A\| \leq c_1's^{1/2}t^{1/2}\}$. Then $P_0(X \in \mathcal{C}) \geq 1 - \delta_3$ and $P_1(X \in \mathcal{C}) \geq 1 - \delta_3$, for $\delta_3 = \exp\{-(n_0 \wedge n_1)^{-1/2}\}$.

By Proposition 9, for $\delta_1 \geq \delta_1^*$ and $\delta_2 \geq \delta_2^*$, we have with probability at least $1 - \delta_1 - \delta_2$, $\hat{\beta}_{A^c}^{\text{lasso}} = \beta_{A^c}^{\text{Bayes}} = 0$. Moreover,

$$\begin{aligned} \|\hat{s} - s^*\|_{\infty, \mathcal{C}} &\leq \max_{x \in \mathcal{C}} |x_A^\top \hat{\beta}_A^{\text{lasso}} - x_A^\top \beta_A^{\text{Bayes}}| + \max_{x \in \mathcal{C}} |x_{A^c}^\top \hat{\beta}_{A^c}^{\text{lasso}} - x_{A^c}^\top \beta_{A^c}^{\text{Bayes}}| \\ &= \max_{x \in \mathcal{C}} |x_A^\top \hat{\beta}_A^{\text{lasso}} - x_A^\top \beta_A^{\text{Bayes}}| \\ &\leq \|\hat{\beta}_A^{\text{lasso}} - \beta_A^{\text{Bayes}}\|_\infty \cdot \max_{x \in \mathcal{C}} \|X_A\|_1 \\ &\leq \|\hat{\beta}_A^{\text{lasso}} - \beta_A^{\text{Bayes}}\|_\infty \cdot \sqrt{s} \max_{x \in \mathcal{C}} \|X_A\|_2 \\ &\leq 4\varphi\lambda \cdot c_1's(n_0 \wedge n_1)^{1/4}, \end{aligned}$$

where the last inequality uses a relation $\beta^* = \beta_A^{\text{Bayes}}$, which is derived in Lemma 18. ■

Proof [Proof of Lemma 11] Note that by Lemma 10, $P_0(X \in \mathcal{C}) \geq 1 - \exp\{-(n_0 \wedge n_1)^{1/2}\}$, so we have

$$\begin{aligned}
 & |R_0(\hat{\phi}_{k^*}) - R_0(\phi_\alpha^*)| \\
 &= |[R_0(\hat{\phi}_{k^*}|\mathcal{C}) - R_0(\phi_\alpha^*|\mathcal{C})]P_0(X \in \mathcal{C}) + [R_0(\hat{\phi}_{k^*}|\mathcal{C}^c) - R_0(\phi_\alpha^*|\mathcal{C}^c)]P_0(X \in \mathcal{C}^c)| \\
 &\geq |[R_0(\hat{\phi}_{k^*}|\mathcal{C}) - R_0(\phi_\alpha^*|\mathcal{C})]P_0(X \in \mathcal{C})| - |[R_0(\hat{\phi}_{k^*}|\mathcal{C}^c) - R_0(\phi_\alpha^*|\mathcal{C}^c)]P_0(X \in \mathcal{C}^c)| \\
 &\geq |R_0(\hat{\phi}_{k^*}|\mathcal{C}) - R_0(\phi_\alpha^*|\mathcal{C})|(1 - \exp\{-(n_0 \wedge n_1)^{1/2}\}) - 1 \cdot \exp\{-(n_0 \wedge n_1)^{1/2}\}.
 \end{aligned}$$

Lemma 8 says that

$$\mathbb{P}\{|R_0(\hat{\phi}_{k^*}) - R_0(\phi_\alpha^*)| > \xi_{\alpha, \delta_0, n'_0}(\delta'_0)\} \leq \delta_0 + \delta'_0.$$

This combined with the above inequality chain implies

$$\mathbb{P}\{|R_0(\hat{\phi}_{k^*}|\mathcal{C}) - R_0(\phi_\alpha^*|\mathcal{C})| > \frac{[\xi_{\alpha, \delta_0, n'_0}(\delta'_0) + \exp\{-(n_0 \wedge n_1)^{1/2}\}]}{1 - \exp\{-(n_0 \wedge n_1)^{1/2}\}}\} \leq \delta_0 + \delta'_0.$$

Since $\exp\{-(n_0 \wedge n_1)^{1/2}\} \leq 1/2$, the conclusion follows. \blacksquare

Proof [Proof of Lemma 18] Note that $\mu^1 - \mu^0 = \Sigma \beta^{\text{Bayes}}$. After shuffling the A coordinates to the front if necessary, we have

$$\mu^1 - \mu^0 = \begin{bmatrix} \Sigma_{AA} & \Sigma_{AA^c} \\ \Sigma_{A^cA} & \Sigma_{A^cA^c} \end{bmatrix} \begin{bmatrix} \beta_A^{\text{Bayes}} \\ \beta_{A^c}^{\text{Bayes}} \end{bmatrix}.$$

Then, $\mu_A^1 - \mu_A^0 = (\Sigma_{AA}) \beta_A^{\text{Bayes}}$ as $\beta_{A^c}^{\text{Bayes}} = 0 \in \mathbb{R}^{|A^c|}$ by definition. Therefore we have,

$$\beta^* = \Sigma_{AA}^{-1}(\mu_A^1 - \mu_A^0) = \beta_A^{\text{Bayes}},$$

which combined with $\tilde{\beta}_{A^c}^{\text{Bayes}} = \beta_{A^c}^{\text{Bayes}} = 0$ leads to $\tilde{\beta}^{\text{Bayes}} = \beta^{\text{Bayes}}$. \blacksquare

Proof [Proof of Lemma 19]

Inequalities (23)-(28) can be proved similarly as in Mai et al. (2012), so proof is omitted.

Inequalities (29)-(31) can be proved by applying (24)-(26) respectively and observe that $A + B \geq \varepsilon$ implies $A \geq \varepsilon/2$ or $B \geq \varepsilon/2$. More concretely, they are proven by the following arguments:

$$\begin{aligned}
 \mathbb{P}\left(|C_{ij}^{(n)} - \Sigma_{ij}| \geq \varepsilon\right) &= \mathbb{P}\left(\left|\frac{n_0}{n} \hat{\Sigma}_{ij}^0 + \frac{n_1}{n} \hat{\Sigma}_{ij}^1 - \Sigma_{ij}\right| \geq \varepsilon\right) \\
 &\leq \mathbb{P}\left(\frac{n_0}{n} |\hat{\Sigma}_{ij}^0 - \Sigma_{ij}| \geq \varepsilon/2\right) + \mathbb{P}\left(\frac{n_1}{n} |\hat{\Sigma}_{ij}^1 - \Sigma_{ij}| \geq \varepsilon/2\right) \\
 &= \mathbb{P}\left(|\hat{\Sigma}_{ij}^0 - \Sigma_{ij}| \geq \frac{n\varepsilon}{2n_0}\right) + \mathbb{P}\left(|\hat{\Sigma}_{ij}^1 - \Sigma_{ij}| \geq \frac{n\varepsilon}{2n_1}\right) \\
 &\leq 2 \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4n_0}\right) + 2 \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4n_1}\right).
 \end{aligned}$$

$$\begin{aligned}
 \mathbb{P}\left(|C_{AA}^{(n)} - \Sigma_{AA}| \geq \varepsilon\right) &= \mathbb{P}\left(\left|\frac{n_0}{n}\widehat{\Sigma}_{AA}^0 + \frac{n_1}{n}\widehat{\Sigma}_{AA}^1 - \Sigma_{AA}\right| \geq \varepsilon\right) \\
 &\leq \mathbb{P}\left(\frac{n_0}{n}|\widehat{\Sigma}_{AA}^0 - \Sigma_{AA}| \geq \varepsilon/2\right) + \mathbb{P}\left(\frac{n_1}{n}|\widehat{\Sigma}_{AA}^1 - \Sigma_{AA}| \geq \varepsilon/2\right) \\
 &\leq 2s^2 \exp\left(-\frac{c_1\varepsilon^2 n^2}{4n_0 s^2}\right) + 2s^2 \exp\left(-\frac{c_1\varepsilon^2 n^2}{4n_1 s^2}\right).
 \end{aligned}$$

$$\begin{aligned}
 \mathbb{P}\left(|C_{A^c A}^{(n)} - \Sigma_{A^c A}| \geq \varepsilon\right) &= \mathbb{P}\left(\left|\frac{n_0}{n}\widehat{\Sigma}_{A^c A}^0 + \frac{n_1}{n}\widehat{\Sigma}_{A^c A}^1 - \Sigma_{A^c A}\right| \geq \varepsilon\right) \\
 &\leq \mathbb{P}\left(\frac{n_0}{n}|\widehat{\Sigma}_{A^c A}^0 - \Sigma_{A^c A}| \geq \varepsilon/2\right) + \mathbb{P}\left(\frac{n_1}{n}|\widehat{\Sigma}_{A^c A}^1 - \Sigma_{A^c A}| \geq \varepsilon/2\right) \\
 &\leq (d-s)s \exp\left(-\frac{c_1 n^2 \varepsilon^2}{4s^2 n_0}\right) + (d-s)s \exp\left(-\frac{c_1 n^2 \varepsilon^2}{4s^2 n_1}\right).
 \end{aligned}$$

■

Proof [Proof of Lemma 20] Let $\eta_1 = \|\Sigma_{AA} - C_{AA}^{(n)}\|_\infty$, $\eta_2 = \|\Sigma_{A^c A} - C_{A^c A}^{(n)}\|_\infty$, and $\eta_3 = \|(C_{AA}^{(n)})^{-1} - (\Sigma_{AA})^{-1}\|_\infty$.

$$\begin{aligned}
 \|C_{A^c A}^{(n)}(C_{AA}^{(n)})^{-1} - \Sigma_{A^c A}(\Sigma_{AA})^{-1}\|_\infty &\leq \|C_{A^c A}^{(n)} - \Sigma_{A^c A}\|_\infty \times \|(C_{AA}^{(n)})^{-1} - (\Sigma_{AA})^{-1}\|_\infty \\
 &\quad + \|C_{A^c A}^{(n)} - \Sigma_{A^c A}\|_\infty \times \|(\Sigma_{AA})^{-1}\|_\infty \\
 &\quad + \|\Sigma_{A^c A}(\Sigma_{AA})^{-1}\|_\infty \times \|\Sigma_{AA} - C_{AA}^{(n)}\|_\infty \times \|(\Sigma_{AA})^{-1}\|_\infty \\
 &\quad + \|\Sigma_{A^c A}(\Sigma_{AA})^{-1}\|_\infty \times \|\Sigma_{AA} - C_{AA}^{(n)}\|_\infty \\
 &\quad \times \|(C_{AA}^{(n)})^{-1} - (\Sigma_{AA})^{-1}\|_\infty \\
 &\leq (\kappa\eta_1 + \eta_2)(\varphi + \eta_3).
 \end{aligned}$$

Moreover, $\eta_3 \leq \|(C_{AA}^{(n)})^{-1}\|_\infty \times \|C_{AA}^{(n)} - \Sigma_{AA}\|_\infty \times \|(\Sigma_{AA})^{-1}\|_\infty \leq (\varphi + \eta_3)\varphi\eta_1$. Hence, if $\varphi\eta_1 < 1$, we have $\eta_3 \leq \varphi^2\eta_1(1 - \varphi\eta_1)^{-1}$. Hence we have,

$$\|C_{A^c A}^{(n)}(C_{AA}^{(n)})^{-1} - \Sigma_{A^c A}(\Sigma_{AA})^{-1}\|_\infty \leq (\kappa\eta_1 + \eta_2)\varphi(1 - \varphi\eta_1)^{-1}.$$

Then we consider the event $\max(\eta_1, \eta_2) \leq \varepsilon$. Note that $\varepsilon < 1/\varphi$ ensures that $\varphi\eta_1 < 1$ on this event. The conclusion follows from inequalities (30) and (31). ■

Proof [Proof of Lemma 21] Since $(\Sigma^{-1}\mu_d)_A = \Sigma_{AA}^{-1}(\mu_A^1 - \mu_A^0)$ (by Lemma 18) and $\widetilde{\mathcal{C}}^0 \subset \mathcal{C}^0 \subset \mathcal{C}$, we have

$$\begin{aligned}
 &P_0\{C_\alpha^{**} \leq s^*(X) \leq C_\alpha^{**} + \delta | X \in \mathcal{C}\} \\
 &\geq P_0(\{C_\alpha^{**} \leq (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} X_A \leq C_\alpha^{**} + \delta\} \cap \mathcal{C}) \\
 &\geq P_0(\{C_\alpha^{**} \leq (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} X_A \leq C_\alpha^{**} + \delta\} \cap \widetilde{\mathcal{C}}^0) \\
 &= P_0(\{C_\alpha^{**} \leq (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1/2} \widetilde{X}_A \leq C_\alpha^{**} + \delta\} | \widetilde{\mathcal{C}}^0) P_0(\widetilde{\mathcal{C}}^0) \\
 &\geq (1 - \delta_3) P_0(C_\alpha^{**} \leq (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1/2} \widetilde{X}_A \leq C_\alpha^{**} + \delta | \widetilde{\mathcal{C}}^0),
 \end{aligned}$$

where the last inequality uses $P_0(\tilde{C}^0) \geq 1 - \delta_3$. To derive this inequality, let V^0 (defined in the proof of Proposition 23) play the role of x and take $A = I_s$ in Lemma 17, then we have

$$\mathbb{P}\left(\|V^0\|^2 \geq s + 2\sqrt{st} + 1 \cdot t\right) \leq e^{-t}, \text{ for all } t > 0.$$

For $s, t \in \mathbb{N}$, the above inequality clearly implies $\mathbb{P}(\|V^0\|^2 \geq 4st) \leq \exp(-t)$. Take $t = (n_0 \wedge n_1)^{1/2}$, then as long as $c'_1 \geq 2/\lambda_m$,

$$\{x : \|V^0\|^2 \leq 4st\} \subset \{x : \|V^0\|^2 \leq \lambda_m^2 (c'_1)^2 st\} = \tilde{C}^0.$$

Since λ_m is bounded from below, we can certainly take $c'_1 \geq 2/\lambda_m$ in the proof of Lemma 10 in constructing \tilde{C}^0 . Therefore, $\mathbb{P}(\|V^0\|^2 \leq s + 2\sqrt{st} + t) \geq 1 - \exp(-t)$ implies that $\mathbb{P}(\tilde{C}^0) \geq 1 - \exp(-t)$ for $t = (n_0 \wedge n_1)^{1/2}$. \blacksquare

Proof [Proof of Lemma 22] Since $V^0 = \tilde{X}_A - \Sigma_{AA}^{-1/2} \mu_A^0$, it follows that,

$$\begin{aligned} & P_0(C_\alpha^{**} \leq (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1/2} \tilde{X}_A \leq C_\alpha^{**} + \delta | \tilde{C}^0) \\ & = P_0(C_\alpha^{**} - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0 \leq a^\top V^0 \leq C_\alpha^{**} + \delta - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0 | \tilde{C}^0). \end{aligned}$$

By Mukerjee and Ong (2015), the probability density function of $V^0 | \tilde{C}^0$ is given by

$$f_{V^0 | \tilde{C}^0}(v) = \begin{cases} k_{\tilde{L},s} \prod_{i=1}^s \phi(v_i) & \text{if } \|v\| \leq \tilde{L} \\ 0, & \text{otherwise,} \end{cases} \quad (32)$$

where ϕ is the pdf for the standard normal random variable, \tilde{L} is defined in equation (39), and $k_{\tilde{L},s}$ is a normalizing constant. Note that $k_{\tilde{L},s}$ is a monotone decreasing function of \tilde{L} for each s , and when \tilde{L} goes to infinity, $k_{\tilde{L},s} = k_s^0$ is a positive constant. Therefore, $k_{\tilde{L},s}$ is bounded below by k_s^0 . Since we only consider $s \in \{1, 2\}$, we can take k^0 as a universal constant independent of s , and $k_{\tilde{L},s}$ is bounded below by k^0 universally.

Let $f_{a^\top V^0 | \tilde{C}^0}(z)$ be the density of $a^\top V^0 | \tilde{C}^0$. Thus, we want to lower bound

$$\begin{aligned} & P_0(C_\alpha^{**} - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0 \leq a^\top V^0 \leq C_\alpha^{**} + \delta - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0 | \tilde{C}^0) \\ & = \int_{C_\alpha^{**} - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0}^{C_\alpha^{**} + \delta - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0} f_{a^\top V^0 | \tilde{C}^0}(z) dz. \end{aligned}$$

Let us analyze $f_{a^\top V^0 | \tilde{C}^0}(z)$ when $s = 1$ and $s = 2$.

Case 1 ($s = 1$): a is a scalar. Hence

$$f_{aV^0 | \tilde{C}^0}(z) = \begin{cases} \frac{k_{\tilde{L},1}}{|a|} \phi\left(\frac{z}{a}\right), & \text{for } |z| \leq |a|\tilde{L} \\ 0, & \text{otherwise,} \end{cases}$$

which is the density function of a truncated Normal random variable with parent distribution $\mathcal{N}(0, |a|)$ symmetrically truncated to $-|a|\tilde{L}$ and $|a|\tilde{L}$, i.e. $TN(0, |a|, -|a|\tilde{L}, |a|\tilde{L})$. Here $|a|$ is the standard deviation of the parent Normal distribution. Therefore,

$$f_{aV^0 | \tilde{C}^0}(z) \geq f_{\mathcal{N}(0, |a|)}(z), \text{ for } |z| \leq |a|\tilde{L}.$$

This implies

$$\begin{aligned}
 & \int_{C_{\alpha}^{**} - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0}^{C_{\alpha}^{**} + \delta - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0} f_{aV^0|\tilde{\mathcal{C}}^0}(z) dz \\
 & \geq \delta \min\{f_{\mathcal{N}(0,|a|)}(C_{\alpha}^{**} - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0), f_{\mathcal{N}(0,|a|)}(C_{\alpha}^{**} - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0 + \delta^*)\} \\
 & \geq \delta M.
 \end{aligned}$$

where the inequality follows from the mean-value theorem and our assumption (ii).

Case 2 ($s = 2$): $a = (a_1, a_2)$ is a vector. Now let us do the following change of variable from $(V_1, V_2) = V^0|\tilde{\mathcal{C}}$ to $(Z_1, Z_2) = (a^\top V^0|\mathcal{C}, V_2)$.

$$\begin{cases} Z_1 = a_1 V_1 + a_2 V_2 \\ Z_2 = V_2 \end{cases}, \text{ and thus } \begin{cases} V_1 = \frac{Z_1 - a_2 Z_2}{a_1} \\ V_2 = Z_2 \end{cases} \quad (33)$$

The original event $S_{V_1, V_2} = \{V_1^2 + V_2^2 \leq \tilde{L}^2\}$ is equivalent to

$$\begin{aligned}
 S_{Z_1, Z_2} & = \left(\frac{Z_1 - a_2 Z_2}{a_1} \right)^2 + Z_2^2 \leq \tilde{L}^2 \\
 & \Leftrightarrow \left(\frac{a_2^2}{a_1^2} + 1 \right) \left(Z_2 - \left(\frac{a_1 a_2}{a_1^2 + a_2^2} \right) Z_1 \right)^2 + \frac{a_1^2 + a_2^2 - a_1^2 a_2^2}{a_1^2 (a_1^2 + a_2^2)} Z_1^2 \leq \tilde{L}^2.
 \end{aligned}$$

Now for any z_1 , the marginal density of $a^\top V^0$ can be carried out as

$$\begin{aligned}
 & f_{Z_1}(z_1) \\
 & = \int_{S_{z_1, z_2}} \frac{k_{\tilde{L}, 2}}{|a_1|} \phi\left(\frac{z_1 - a_2 z_2}{a_1}\right) \phi(z_2) dz_2 \\
 & = \int_{S_{z_1, z_2}} \frac{k_{\tilde{L}, 2}}{2\pi|a_1|} \exp\left\{-\frac{1}{2} \left(\frac{z_1 - a_2 z_2}{a_1}\right)^2 - \frac{z_2^2}{2}\right\} dz_2 \\
 & = \left(\frac{k_{\tilde{L}, 2}}{2\pi|a_1|} \exp\left\{-\frac{a_1^2 + a_2^2 - a_2^2 a_1^2}{2a_1^2(a_1^2 + a_2^2)} z_1^2\right\} \right) \int_{S_{z_1, z_2}} \exp\left\{-\frac{(z_2 - \left(\frac{a_1 a_2}{a_1^2 + a_2^2}\right) z_1)^2}{\frac{2}{a_2^2/a_1^2 + 1}}\right\} dz_2 \\
 & = \left(\frac{k_{\tilde{L}, 2}}{|a_1| \sqrt{2\pi\left(\frac{a_2^2}{a_1^2} + 1\right)}} \exp\left\{-\frac{a_1^2 + a_2^2 - a_2^2 a_1^2}{2a_1^2(a_1^2 + a_2^2)} z_1^2\right\} \right) \int_{S_{z_1, z_2}} \phi_{\mathcal{N}\left(\left(\frac{a_1 a_2}{a_1^2 + a_2^2}\right) z_1, \frac{1}{\sqrt{a_2^2/a_1^2 + 1}}\right)}(z_2) dz_2 \\
 & = \left(\frac{k_{\tilde{L}, 2}}{\sqrt{2\pi}|a_1|\left(\frac{a_2^2}{a_1^2} + 1\right)} \exp\left\{-\frac{a_1^2 + a_2^2 - a_2^2 a_1^2}{a_1^2(a_1^2 + a_2^2)} z_1^2\right\} \right) \int_{-\sqrt{\tilde{L}^2 - \frac{a_1^2 + a_2^2 - a_2^2 a_1^2}{a_1^2(a_1^2 + a_2^2)} z_1^2}}^{\sqrt{\tilde{L}^2 - \frac{a_1^2 + a_2^2 - a_2^2 a_1^2}{a_1^2(a_1^2 + a_2^2)} z_1^2}} \phi_{\mathcal{N}(0,1)}(z) dz \\
 & = \left(\frac{k_{\tilde{L}, 2}}{\sqrt{2\pi}|a_1|\left(\frac{a_2^2}{a_1^2} + 1\right)} \exp\left\{-\frac{a_1^2 + a_2^2 - a_2^2 a_1^2}{a_1^2(a_1^2 + a_2^2)} z_1^2\right\} \right) \left(2\Phi\left(\sqrt{\tilde{L}^2 - \frac{a_1^2 + a_2^2 - a_2^2 a_1^2}{a_1^2(a_1^2 + a_2^2)} z_1^2}\right) - 1 \right).
 \end{aligned}$$

This implies

$$\begin{aligned}
 & \int_{C_{\alpha}^{**} - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0}^{C_{\alpha}^{**} + \delta - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0} f_{a^\top V^0} \tilde{c}^0(z) dz \\
 & \geq \delta \min\{f_{Z_1}(C_{\alpha}^{**} - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0), f_{Z_1}(C_{\alpha}^{**} - (\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1} \mu_A^0 + \delta^*)\} \\
 & \geq \delta M k_0.
 \end{aligned}$$

where the inequality follows from the mean-value theorem and our assumption (iii).

We can safely conclude our proof by combining cases $s = 1$ and $s = 2$, and taking $M_1 = \min\{M, M k_0\}$. \blacksquare

Proof [Proof of Proposition 9] The proof is largely identical to that of Theorem 1 in Mai et al. (2012), except the differences due to a different sampling scheme.

Similarly to Mai et al. (2012), by the definition of $\hat{\beta}_A$, we can write $\hat{\beta}_A = (n^{-1} \tilde{X}_A^\top \tilde{X}_A)^{-1} \{(\hat{\mu}_A^1 - \hat{\mu}_A^0) - \lambda t_A / 2\}$, where t_A represents the subgradient such that $t_j = \text{sign}(\hat{\beta}_j)$ if $\hat{\beta}_j \neq 0$ and $-1 < t_j < 1$ if $\hat{\beta}_j = 0$. To show that $\hat{\beta}^{\text{lasso}} = (\hat{\beta}_A, 0)$, it suffices to verify that

$$\|n^{-1} \tilde{X}_{A^c}^\top \tilde{X}_A \hat{\beta}_A - (\hat{\mu}_{A^c}^1 - \hat{\mu}_{A^c}^0)\|_\infty \leq \lambda / 2. \quad (34)$$

The left-hand side of (34) is equal to

$$\|C_{A^c A}^{(n)} (C_{AA}^{(n)})^{-1} (\hat{\mu}_A^1 - \hat{\mu}_A^0) - C_{A^c A}^{(n)} (C_{AA}^{(n)})^{-1} \lambda t_A / 2 - (\hat{\mu}_{A^c}^1 - \hat{\mu}_{A^c}^0)\|_\infty. \quad (35)$$

Using $\Sigma_{A^c A} \Sigma_{AA}^{-1} (\mu_A^1 - \mu_A^0) = (\mu_{A^c}^1 - \mu_{A^c}^0)$, (35) is bounded from above by

$$\begin{aligned}
 U_1 &= \|C_{A^c A}^{(n)} (C_{AA}^{(n)})^{-1} - \Sigma_{A^c A} \Sigma_{AA}^{-1}\|_\infty \Delta + \|(\hat{\mu}_{A^c}^1 - \hat{\mu}_{A^c}^0) - (\mu_{A^c}^1 - \mu_{A^c}^0)\|_\infty \\
 &+ (\|C_{A^c A}^{(n)} (C_{AA}^{(n)})^{-1} - \Sigma_{A^c A} \Sigma_{AA}^{-1}\|_\infty + \kappa) \|(\hat{\mu}_A^1 - \hat{\mu}_A^0) - (\mu_A^1 - \mu_A^0)\|_\infty \\
 &+ (\|C_{A^c A}^{(n)} (C_{AA}^{(n)})^{-1} - \Sigma_{A^c A} \Sigma_{AA}^{-1}\|_\infty + \kappa) \lambda / 2.
 \end{aligned}$$

If $\|C_{A^c A}^{(n)} (C_{AA}^{(n)})^{-1} - \Sigma_{A^c A} \Sigma_{AA}^{-1}\|_\infty \leq (\kappa + 1) \varepsilon \varphi (1 - \varphi \varepsilon)^{-1}$ (invoke Lemma 20), and $\|(\hat{\mu}_A^1 - \hat{\mu}_A^0) - (\mu_A^1 - \mu_A^0)\|_\infty \leq 4^{-1} \lambda (1 - \kappa - 2\varepsilon \varphi) / (1 + \kappa)$, and given $\varepsilon \leq \min[\varepsilon_0, \lambda (1 - \kappa) (4\varphi)^{-1} (\lambda / 2 + (1 + \kappa) \Delta)^{-1}]$, then $U_1 \leq \lambda / 2$.

Therefore, by Lemmas 19 and 20, we have

$$\begin{aligned}
 & \mathbb{P}\{\|n^{-1} \tilde{X}_{A^c}^\top \tilde{X}_A \hat{\beta}_A - (\hat{\mu}_{A^c}^1 - \hat{\mu}_{A^c}^0)\|_\infty \leq \lambda / 2\} \\
 & \geq 1 - 2d \exp(-n_0 \varepsilon^{*2} c_2) - 2d \exp(-n_1 \varepsilon^{*2} c_2) - f(d, s, n_0, n_1, (\kappa + 1) \varepsilon \varphi (1 - \varphi \varepsilon)^{-1}),
 \end{aligned}$$

where $\varepsilon^* = 4^{-1} \lambda (1 - \kappa - 2\varepsilon \varphi) / (1 + \kappa)$, and f is the same as in Lemma 20. Tidy up the algebra a bit, we can write

$$\delta_1^* = \sum_{l=0}^1 2d \exp\left(-c_2 n_l \frac{\lambda^2 (1 - \kappa - 2\varepsilon \varphi)^2}{16(1 + \kappa)^2}\right) + f(d, s, n_0, n_1, (\kappa + 1) \varepsilon \varphi (1 - \varphi \varepsilon)^{-1}).$$

To prove the 2nd conclusion, note that

$$\hat{\beta}_A = (\Sigma_{AA})^{-1} (\mu_A^1 - \mu_A^0) + (C_{AA}^{(n)})^{-1} \{(\hat{\mu}_A^1 - \hat{\mu}_A^0) - (\mu_A^1 - \mu_A^0)\} \quad (36)$$

$$+ \{(C_{AA}^{(n)})^{-1} - (\Sigma_{AA})^{-1}\} (\mu_A^1 - \mu_A^0) - \lambda (C_{AA}^{(n)})^{-1} t_A / 2. \quad (37)$$

Let $\xi = |\beta^*|_{\min}/(\Delta\varphi)$. Write $\eta_1 = \|\Sigma_{AA} - C_{AA}^{(n)}\|_\infty$ and $\eta_3 = \|(C_{AA}^{(n)})^{-1} - \Sigma_{AA}^{-1}\|_\infty$. Then for any $j \in A$,

$$|\hat{\beta}_j| \geq \xi\Delta\varphi - (\eta_3 + \varphi)\{\lambda/2 + \|(\hat{\mu}_A^1 - \hat{\mu}_A^0) - (\mu_A^1 - \mu_A^0)\|_\infty\} - \eta_3\Delta.$$

When $\eta_1\varphi < 1$, we have shown that $\eta_3 < \varphi^2\eta_1(1 - \eta_1\varphi)^{-1}$ in Lemma 20. Therefore,

$$|\hat{\beta}_j| \geq \xi\Delta\varphi - (1 - \eta_1\varphi)^{-1}\{\lambda\varphi/2 + \|(\hat{\mu}_A^1 - \hat{\mu}_A^0) - (\mu_A^1 - \mu_A^0)\|_\infty\varphi + \varphi^2\eta_1\Delta\} \equiv L_1.$$

Because $\|\beta^*\|_\infty \leq \Delta\varphi$, $\xi \leq 1$. Hence $\lambda \leq |\beta^*|_{\min}/(2\varphi) \leq 2|\beta^*|_{\min}/\{(3 + \xi)\varphi\}$. Under the events $\eta_1 \leq \varepsilon$ and $\|(\hat{\mu}_A^1 - \hat{\mu}_A^0) - (\mu_A^1 - \mu_A^0)\|_\infty \leq \varepsilon$, together with restriction on ε , we have $L_1 > 0$. Therefore,

$$\mathbb{P}(L_1 > 0) \geq 1 - \sum_{l=0}^1 2s \exp(-n_l \varepsilon^2 c_2) - \sum_{l=0}^1 2s^2 \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4n_l s^2}\right).$$

To prove the 3rd conclusion, equation (36) and $\eta_1\varphi < 1$ imply that

$$\|\hat{\beta}_A - \beta^*\|_\infty \leq (1 - \eta_1\varphi)^{-1}\{\lambda\varphi/2 + \|(\hat{\mu}_A^1 - \hat{\mu}_A^0) - (\mu_A^1 - \mu_A^0)\|_\infty\varphi + \varphi^2\eta_1\Delta\}.$$

On the events $\{\eta_1 < \varepsilon\}$ and $\{\|(\hat{\mu}_A^1 - \hat{\mu}_A^0) - (\mu_A^1 - \mu_A^0)\|_\infty \leq \varepsilon\}$, and under restrictions for ε and λ in the assumption, we have $\|\hat{\beta}_A - \beta^*\|_\infty \leq 4\varphi\lambda$. Hence,

$$\mathbb{P}(\|\hat{\beta}_A - \beta^*\|_\infty \leq 4\varphi\lambda) \geq 1 - \sum_{l=0}^1 2s \exp(-n_l \varepsilon^2 c_2) - \sum_{l=0}^1 2s^2 \exp\left(-\frac{c_1 \varepsilon^2 n^2}{4s^2 n_l}\right).$$

■

Proof [Proof of Proposition 23] For simplicity, we will derive the lower bound for one of the two probabilities in the definition:

$$P_0\{C_\alpha^{**} \leq s^*(X) \leq C_\alpha^{**} + \delta | X \in \mathcal{C}\} \geq (1 - \delta_3)M_1\delta, \text{ for } \delta \in (0, \delta^*). \quad (38)$$

The lower bound for the other probability can be derived similarly.

Recall that $\mathcal{C}^0 = \{X \in \mathbb{R}^d : \|X_A - \mu_A^0\| \leq c_1 s^{1/2}(n_0 \wedge n_1)^{1/4} \doteq L\}$ (in the proof of Lemma 10). Let $V^0 = \Sigma_{AA}^{-1/2}(X_A - \mu_A^0) = \tilde{X}_A - \Sigma_{AA}^{-1/2}\mu_A^0$, where $\tilde{X}_A = \Sigma_{AA}^{-1/2}X_A$, then $V^0 \sim \mathcal{N}(0, I_s)$ under P_0 . Define an event

$$\tilde{\mathcal{C}}^0 = \left\{X \in \mathbb{R}^d : \|V^0\| \leq \lambda_m L \doteq \tilde{L}\right\}, \quad (39)$$

where $\lambda_m = \lambda_{\min}(\Sigma_{AA}^{-1/2})$ and $\lambda_{\min}(\cdot)$ denotes the minimum eigenvalue of a matrix. Since $\|V^0\| \geq \lambda_{\min}(\Sigma_{AA}^{-1/2})\|X_A - \mu_A^0\| = \lambda_m\|X_A - \mu_A^0\|$, we have $\tilde{\mathcal{C}}^0 \subset \mathcal{C}^0$. Then inequality (38) holds by invoking Lemma 21 and Lemma 22. ■

Remark 24 *Proof of Proposition 23 indicates that the same conclusion would hold for a general $s \in \mathbb{N}$, if the density of $(\mu_A^1 - \mu_A^0)^\top \Sigma_{AA}^{-1/2} \tilde{X}_A | \tilde{\mathcal{C}}^0$ is bounded below on $(C^1 - \delta^*, C^2 + \delta^*)$ by some constant.*

Proof [Proof of Theorem 15] The first inequality follows from Proposition 1 and the choice of k^* in (12) of the main paper. In the following, we prove the second inequality.

Let $G^* = \{s^* \leq C_\alpha^{**}\}$ and $\hat{G} = \{\hat{s} \leq \hat{C}_\alpha\}$. The excess type II error can be decomposed as

$$P_1(\hat{G}) - P_1(G^*) = \int_{\hat{G} \setminus G^*} |r - C_\alpha| dP_0 + \int_{G^* \setminus \hat{G}} |r - C_\alpha| dP_0 + C_\alpha \{R_0(\phi_\alpha^*) - R_0(\hat{\phi}_{k^*})\}. \quad (40)$$

In the above decomposition, the third part can be bounded via Lemma 8. For the first two parts, let

$$T = \|\hat{s} - s^*\|_{\infty, \mathcal{C}} := \max_{x \in \mathcal{C}} |\hat{s}(x) - s^*(x)|, \text{ and}$$

$$\Delta R_{0, \mathcal{C}} := |R_0(\phi_\alpha^* | \mathcal{C}) - R_0(\hat{\phi}_{k^*} | \mathcal{C})| = |P_0(s^*(X) > C_\alpha^{**} | X \in \mathcal{C}) - P_0(\hat{s}(X) > \hat{C}_\alpha | X \in \mathcal{C})|,$$

where \mathcal{C} is defined in Lemma 10. A high probability bound for $\Delta R_{0, \mathcal{C}}$ was derived in Lemma 11.

It follows from Lemma 7 that if $n'_0 \geq \max\{4/(\alpha\alpha_0), \delta_0^{-2}, (\delta'_0)^{-2}, (\frac{1}{10}M_1\delta^{*\gamma})^{-4}\}$,

$$\xi_{\alpha, \delta_0, n'_0}(\delta'_0) \leq \frac{5}{2}(n'_0)^{-1/4} \leq \frac{1}{4}M_1(\delta^*)^\gamma.$$

Because the lower bound in the detection condition should be smaller than 1 to make sense, $M_1\delta^{*\gamma} < 1$. This together with $n_0 \wedge n_1 \geq [-\log(M_1\delta^{*\gamma}/4)]^2$ implies that $\exp\{-(n_0 \wedge n_1)^{1/2}\} \leq M_1\delta^{*\gamma}/4$.

Let $\mathcal{E}_2 = \{R_{0, \mathcal{C}} \leq 2[\xi_{\alpha, \delta_0, n'_0}(\delta'_0) + \exp\{-(n_0 \wedge n_1)^{1/2}\}]\}$. On the event \mathcal{E}_2 we have

$$\left\{ \frac{R_{0, \mathcal{C}}}{M_1} \right\}^{1/\gamma} \leq \left\{ \frac{2[\xi_{\alpha, \delta_0, n'_0}(\delta'_0) + \exp\{-(n_0 \wedge n_1)^{1/2}\}]}{M_1} \right\}^{1/\gamma} \leq \delta^*.$$

To find the relation between C_α^{**} and \hat{C}_α , we invoke the detection condition as follows:

$$\begin{aligned} & P_0 \left(s^*(X) \geq C_\alpha^{**} + (\Delta R_{0, \mathcal{C}}/M_1)^{1/\gamma} | X \in \mathcal{C} \right) \\ &= R_0(\phi_\alpha^* | \mathcal{C}) - P_0(C_\alpha^{**} < s^*(X) < C_\alpha^{**} + (\Delta R_{0, \mathcal{C}}/M_1)^{1/\gamma} | X \in \mathcal{C}) \\ &\leq R_0(\phi_\alpha^* | \mathcal{C}) - \Delta R_{0, \mathcal{C}} \quad (\text{by detection condition}) \\ &\leq R_0(\hat{\phi}_{k^*} | \mathcal{C}) = P_0(\hat{s}(X) > \hat{C}_\alpha | X \in \mathcal{C}) \\ &\leq P_0(s^*(X) > \hat{C}_\alpha - T | X \in \mathcal{C}). \end{aligned}$$

This implies that $C_\alpha^{**} + (\Delta R_{0, \mathcal{C}}/M_1)^{1/\gamma} \geq \hat{C}_\alpha - T$, which further implies that

$$\hat{C}_\alpha \leq C_\alpha^{**} + (\Delta R_{0, \mathcal{C}}/M_1)^{1/\gamma} + T.$$

Note that

$$\begin{aligned}
 & \mathcal{C} \cap (\widehat{G} \setminus G^*) \\
 &= \mathcal{C} \cap \{s^* > C_\alpha^{**}, \hat{s} \leq \widehat{C}_\alpha\} \\
 &= \mathcal{C} \cap \{s^* > C_\alpha^{**}, \hat{s} \leq C_\alpha^{**} + (\Delta R_{0,\mathcal{C}}/M_1)^{1/\gamma} + T\} \cap \{\hat{s} \leq \widehat{C}_\alpha\} \\
 &\subset \mathcal{C} \cap \{C_\alpha^{**} + (\Delta R_{0,\mathcal{C}}/M_1)^{1/\gamma} + 2T \geq s^* \geq C_\alpha^{**}, \hat{s} \leq C_\alpha^{**} + (\Delta R_{0,\mathcal{C}}/M_1)^{1/\gamma} + T\} \cap \{\hat{s} \leq \widehat{C}_\alpha\} \\
 &\subset \mathcal{C} \cap \{C_\alpha^{**} + (\Delta R_{0,\mathcal{C}}/M_1)^{1/\gamma} + 2T \geq s^* \geq C_\alpha^{**}\}.
 \end{aligned}$$

We decompose as follows

$$\int_{(\widehat{G} \setminus G^*)} |r - C_\alpha| dP_0 = \int_{(\widehat{G} \setminus G^*) \cap \mathcal{C}} |r - C_\alpha| dP_0 + \int_{(\widehat{G} \setminus G^*) \cap \mathcal{C}^c} |r - C_\alpha| dP_0 =: \text{(I)} + \text{(II)}.$$

To bound (I), recall that

$$r(x) = \frac{f_1(x)}{f_0(x)} = \exp\left(s^*(x) - \mu_a^\top \Sigma^{-1} \mu_d\right),$$

and that, $r(x) > C_\alpha$ is equivalent to $s^*(x) > C_\alpha^{**} = \log C_\alpha + \mu_a^\top \Sigma^{-1} \mu_d$. By the mean value theorem, we have

$$|r(x) - C_\alpha| = e^{-\mu_a^\top \Sigma^{-1} \mu_d} |e^{s^*(x)} - e^{C_\alpha^{**}}| = e^{-\mu_a^\top \Sigma^{-1} \mu_d} \cdot e^{z'} |s^*(x) - C_\alpha^{**}|,$$

where z' is some quantity between $s^*(x)$ and C_α^{**} . Denote by $\mathcal{C}_1 = \{x : C_\alpha^{**} + (\Delta R_{0,\mathcal{C}}/M_1)^{1/\gamma} + 2T \geq s^*(x) \geq C_\alpha^{**}\}$. Restricting to $\mathcal{C} \cap \mathcal{C}_1$, we have

$$z' \leq C_\alpha^{**} + (\Delta R_{0,\mathcal{C}}/M_1)^{1/\gamma} + 2T.$$

This together with $\mathcal{C} \cap (\widehat{G} \setminus G^*) \subset \mathcal{C} \cap \mathcal{C}_1$ implies that

$$\begin{aligned}
 \text{(I)} &\leq \int_{\mathcal{C} \cap \mathcal{C}_1} |r - C_\alpha| dP_0 \\
 &= \int_{\mathcal{C} \cap \mathcal{C}_1} \exp\{z' - \mu_a^\top \Sigma^{-1} \mu_d\} |s^*(x) - C_\alpha^{**}| dP_0 \\
 &\leq \int_{\mathcal{C} \cap \mathcal{C}_1} \exp\left\{C_\alpha^{**} + (\Delta R_{0,\mathcal{C}}/M_1)^{1/\gamma} + 2T - \mu_a^\top \Sigma^{-1} \mu_d\right\} |s^*(x) - C_\alpha^{**}| dP_0.
 \end{aligned}$$

Since C_α and $\mu_a^\top \Sigma^{-1} \mu_d$ are assumed to be bounded, $C_\alpha^{**} = \log C_\alpha + \mu_a^\top \Sigma^{-1} \mu_d$ is also bounded. Let $\mathcal{E}_2 = \{R_{0,\mathcal{C}} \leq 2[\xi_{\alpha,\delta_0,n'_0}(\delta'_0) + \exp\{-(n_0 \wedge n_1)^{1/2}\}]\}$. By Lemma 11, $\mathbb{P}(\mathcal{E}_2) \geq 1 - \delta_0 - \delta'_0$. Let $\mathcal{E}_3 = \{T \leq 4c_1 \varphi \lambda s(n_0 \wedge n_1)^{1/4}\}$. By Lemma 10, $\mathbb{P}(\mathcal{E}_3) \geq 1 - \delta_1 - \delta_2$. Restricting to the event $\mathcal{E}_2 \cap \mathcal{E}_3$, $R_{0,\mathcal{C}}$ and T are bounded. Therefore on the event $\mathcal{E}_2 \cap \mathcal{E}_3$, there exists a positive constant c' such that

$$\begin{aligned}
 \text{(I)} &\leq c' \int_{\mathcal{C} \cap \mathcal{C}_1} |s^*(x) - C_\alpha^{**}| dP_0 \\
 &\leq c' \left((\Delta R_{0,\mathcal{C}}/M_1)^{1/\gamma} + 2T \right) P_0(\mathcal{C} \cap \mathcal{C}_1).
 \end{aligned}$$

Note that by the margin assumption (we know $\bar{\gamma} = 1$, but we choose to reserve the explicit dependency of $\bar{\gamma}$ by not substituting the numerical value),

$$\begin{aligned} P_0(\mathcal{C} \cap \mathcal{C}_1) &= P_0(C_\alpha^{**} + (\Delta R_{0,\mathcal{C}}/M_1)^{1/\gamma} + 2T \geq s^* \geq C_\alpha^{**}, \mathcal{C}) \\ &\leq P_0(C_\alpha^{**} + (\Delta R_{0,\mathcal{C}}/M_1)^{1/\gamma} + 2T \geq s^* \geq C_\alpha^{**} | \mathcal{C}) \\ &\leq M_0 \left((\Delta R_{0,\mathcal{C}}/M_1)^{1/\gamma} + 2T \right)^{\bar{\gamma}}. \end{aligned}$$

Therefore,

$$(I) \leq c' M_0 \left((\Delta R_{0,\mathcal{C}}/M_1)^{1/\gamma} + 2T \right)^{1+\bar{\gamma}}.$$

Regarding (II), by Lemma 10 we have

$$\begin{aligned} (II) &\leq \int_{\mathcal{C}^c} |r - C_\alpha| dP_0 \leq \int_{\mathcal{C}^c} r dP_0 + C_\alpha \int_{\mathcal{C}^c} dP_0 = P_1(\mathcal{C}^c) + C_\alpha P_0(\mathcal{C}^c) \\ &\leq (1 + C_\alpha) \exp\{-(n_0 \wedge n_1)^{1/2}\}. \end{aligned}$$

Therefore,

$$\int_{(\widehat{G} \setminus G^*)} |r - C_\alpha| dP_0 \leq c' M_0 \left((\Delta R_{0,\mathcal{C}}/M_1)^{1/\gamma} + 2T \right)^{1+\bar{\gamma}} + (1 + C_\alpha) \exp\{-(n_0 \wedge n_1)^{1/2}\}.$$

To bound $\int_{(G^* \setminus \widehat{G})} |r - C_\alpha| dP_0$, we decompose

$$\int_{(G^* \setminus \widehat{G})} |r - C_\alpha| dP_0 = \int_{(G^* \setminus \widehat{G}) \cap \mathcal{C}} |r - C_\alpha| dP_0 + \int_{(G^* \setminus \widehat{G}) \cap \mathcal{C}^c} |r - C_\alpha| dP_0 =: (I') + (II'').$$

To bound (I'), we invoke both the margin assumption and the detection condition, and we need to define a new quantity $\bar{\Delta} R_{0,\mathcal{C}} := P_0(s^*(X) > C_\alpha^{**} | X \in \mathcal{C}) - P_0(\hat{s}(X) > \widehat{C}_\alpha | X \in \mathcal{C})$. When $\bar{\Delta} R_{0,\mathcal{C}} \geq 0$, we have

$$\begin{aligned} &P_0 \left(s^*(X) \geq C_\alpha^{**} + (\bar{\Delta} R_{0,\mathcal{C}}/M_0)^{1/\bar{\gamma}} | X \in \mathcal{C} \right) \\ &= P_0(s^*(X) > C_\alpha^{**} | X \in \mathcal{C}) - P_0(C_\alpha^{**} < s^*(X) < C_\alpha^{**} + (\bar{\Delta} R_{0,\mathcal{C}}/M_0)^{1/\bar{\gamma}} | X \in \mathcal{C}) \\ &\geq P_0(s^*(X) > C_\alpha^{**} | X \in \mathcal{C}) - \bar{\Delta} R_{0,\mathcal{C}} \quad (\text{by margin assumption}) \\ &= P_0(\hat{s}(X) > \widehat{C}_\alpha | X \in \mathcal{C}) \\ &\geq P_0(s^*(X) > \widehat{C}_\alpha + T | X \in \mathcal{C}). \end{aligned}$$

So when $\bar{\Delta} R_{0,\mathcal{C}} \geq 0$, $\widehat{C}_\alpha \geq C_\alpha^{**} - (\bar{\Delta} R_{0,\mathcal{C}}/M_0)^{1/\bar{\gamma}} - T$. On the other hand, when $\bar{\Delta} R_{0,\mathcal{C}} < 0$,

$$\begin{aligned} &P_0 \left(s^*(X) \geq C_\alpha^{**} - (-\bar{\Delta} R_{0,\mathcal{C}}/M_1)^{1/\gamma} | X \in \mathcal{C} \right) \\ &= P_0(s^*(X) > C_\alpha^{**} | X \in \mathcal{C}) + P_0(C_\alpha^{**} \geq s^*(X) \geq C_\alpha^{**} - (-\bar{\Delta} R_{0,\mathcal{C}}/M_1)^{1/\gamma} | X \in \mathcal{C}) \\ &\geq P_0(s^*(X) > C_\alpha^{**} | X \in \mathcal{C}) + |\bar{\Delta} R_{0,\mathcal{C}}| \quad (\text{by detection condition}) \\ &= P_0(\hat{s}(X) > \widehat{C}_\alpha | X \in \mathcal{C}) \\ &\geq P_0(s^*(X) > \widehat{C}_\alpha + T | X \in \mathcal{C}). \end{aligned}$$

So when $\bar{\Delta}R_{0,c} < 0$, $\widehat{C}_\alpha \geq C_\alpha^{**} - (-\bar{\Delta}R_{0,c}/M_1)^{1/\gamma} - T$. Note that $\Delta R_{0,c} = |\bar{\Delta}R_{0,c}|$. Therefore we have in both cases,

$$\widehat{C}_\alpha \geq C_\alpha^{**} - (\Delta R_{0,c}/M_0)^{1/\bar{\gamma}} \wedge (\Delta R_{0,c}/M_1)^{1/\gamma} - T.$$

Using the above inequality, we have

$$\begin{aligned} & \mathcal{C} \cap (G^* \setminus \widehat{G}) \\ &= \mathcal{C} \cap \{s^* \leq C_\alpha^{**}, \hat{s} > \widehat{C}_\alpha\} \\ &= \mathcal{C} \cap \{s^* \leq C_\alpha^{**}, \hat{s} \geq C_\alpha^{**} - (\Delta R_{0,c}/M_0)^{1/\bar{\gamma}} \wedge (\Delta R_{0,c}/M_1)^{1/\gamma} - T\} \cap \{\hat{s} > \widehat{C}_\alpha\} \\ &\subset \mathcal{C} \cap \{C_\alpha^{**} - (\Delta R_{0,c}/M_0)^{1/\bar{\gamma}} \wedge (\Delta R_{0,c}/M_1)^{1/\gamma} - 2T \leq s^* \leq C_\alpha^{**}\} \cap \{\hat{s} \geq \widehat{C}_\alpha\} \\ &\subset \mathcal{C} \cap \{C_\alpha^{**} - (\Delta R_{0,c}/M_0)^{1/\bar{\gamma}} \wedge (\Delta R_{0,c}/M_1)^{1/\gamma} - 2T \leq s^* \leq C_\alpha^{**}\}. \end{aligned}$$

Denote by $\mathcal{C}_2 = \{x : C_\alpha^{**} - (\Delta R_{0,c}/M_0)^{1/\bar{\gamma}} \wedge (\Delta R_{0,c}/M_1)^{1/\gamma} - 2T \leq s^*(x) \leq C_\alpha^{**}\}$. Then we just showed that $\mathcal{C} \cap (G^* \setminus \widehat{G}) \subset \mathcal{C} \cap \mathcal{C}_2$. Recall that

$$|r(x) - C_\alpha| = e^{-\mu_a^\top \Sigma^{-1} \mu_d} |e^{s^*(x)} - e^{C_\alpha^{**}}| = e^{-\mu_a^\top \Sigma^{-1} \mu_d} \cdot e^{z'} |s^*(x) - C_\alpha^{**}|,$$

where z' is some quantity between $s^*(x)$ and C_α^{**} . Restricting to $\mathcal{C} \cap \mathcal{C}_2$, we have

$$z' \leq C_\alpha^{**}.$$

This together with $\mathcal{C} \cap (G^* \setminus \widehat{G}) \subset \mathcal{C} \cap \mathcal{C}_2$ implies that

$$\begin{aligned} (I') &\leq \int_{\mathcal{C} \cap \mathcal{C}_2} |r - C_\alpha| dP_0 \\ &= \int_{\mathcal{C} \cap \mathcal{C}_2} \exp\{z' - \mu_a \Sigma^{-1} \mu_d\} |s^*(x) - C_\alpha^{**}| dP_0 \\ &\leq \int_{\mathcal{C} \cap \mathcal{C}_2} c'' |s^*(x) - C_\alpha^{**}| dP_0 \\ &\leq c'' \left((\Delta R_{0,c}/M_0)^{1/\bar{\gamma}} \wedge (\Delta R_{0,c}/M_1)^{1/\gamma} + 2T \right) P_0(\mathcal{C} \cap \mathcal{C}_2). \end{aligned}$$

Note that by the margin assumption,

$$\begin{aligned} & P_0(\mathcal{C} \cap \mathcal{C}_2) \\ &\leq P_0(C_\alpha^{**} - (\Delta R_{0,c}/M_0)^{1/\bar{\gamma}} \wedge (\Delta R_{0,c}/M_1)^{1/\gamma} - 2T \leq s^*(X) \leq C_\alpha^{**}) \\ &\leq M_0 \left((\Delta R_{0,c}/M_0)^{1/\bar{\gamma}} \wedge (\Delta R_{0,c}/M_1)^{1/\gamma} + 2T \right)^{\bar{\gamma}}. \end{aligned}$$

Therefore,

$$(I') \leq c'' M_0 \left((\Delta R_{0,c}/M_0)^{1/\bar{\gamma}} \wedge (\Delta R_{0,c}/M_1)^{1/\gamma} + 2T \right)^{1+\bar{\gamma}}.$$

Regarding (II'), by Lemma 10 we have

$$(II') \leq \int_{\mathcal{C}^c} rdP_0 + C_\alpha \int_{\mathcal{C}^c} dP_0 = P_1(\mathcal{C}^c) + C_\alpha P_0(\mathcal{C}^c) \leq (1 + C_\alpha) \exp\{-(n_0 \wedge n_1)^{1/2}\}.$$

Therefore, by the excess type II error decomposition equation (40),

$$P_1(\widehat{G}) - P_1(G^*) = (I) + (II) + (I') + (II') + C_\alpha \{R_0(\phi_\alpha^*) - R_0(\widehat{\phi}_{k^*})\}.$$

Using the upper bounds for (I), (II), (I') and (II') and Lemma 8, With probability at least $1 - \delta_0 - \delta'_0 - \delta_1 - \delta_2$, we have

$$\begin{aligned} P_1(\widehat{G}) - P_1(G^*) &\leq c' M_0 \left((\Delta R_{0,\mathcal{C}}/M_1)^{1/\underline{\gamma}} + 2T \right)^{1+\bar{\gamma}} \\ &\quad + c'' M_0 \left((\Delta R_{0,\mathcal{C}}/M_0)^{1/\bar{\gamma}} \wedge (\Delta R_{0,\mathcal{C}}/M_1)^{1/\underline{\gamma}} + 2T \right)^{1+\bar{\gamma}} \\ &\quad + 2(1 + C_\alpha) \exp\{-(n_0 \wedge n_1)^{1/2}\} + C_\alpha \cdot \xi_{\alpha,\delta_0,n'_0}(\delta'_0) \\ &\leq \bar{c}'_1 \Delta R_{0,\mathcal{C}}^{(1+\bar{\gamma})/\underline{\gamma}} + \bar{c}'_2 T^{1+\bar{\gamma}} + \bar{c}'_3 \exp\{-(n_0 \wedge n_1)^{1/2}\} + C_\alpha \cdot \xi_{\alpha,\delta_0,n'_0}(\delta'_0), \end{aligned}$$

for some positive constants \bar{c}'_1 , \bar{c}'_2 and \bar{c}'_3 . In the last inequality of the above chain, we used $\underline{\gamma} \geq \bar{\gamma}$. Note that on the event $\mathcal{E}_2 \cap \mathcal{E}_3$, Lemma 11 guarantees $\Delta R_{0,\mathcal{C}} \leq 2[\xi_{\alpha,\delta_0,n'_0}(\delta'_0) + \exp\{-(n_0 \wedge n_1)^{1/2}\}]$. Lemma 10 guarantees that $T \leq 4c'_1 \varphi \lambda s (n_0 \wedge n_1)^{1/4}$. Therefore,

$$\begin{aligned} P_1(\widehat{G}) - P_0(G^*) &\leq \bar{c}''_1 \xi_{\alpha,\delta_0,n'_0}(\delta'_0)^{(1+\bar{\gamma})/\underline{\gamma} \wedge 1} + \bar{c}''_2 [4c'_1 \varphi \lambda s (n_0 \wedge n_1)^{1/4}]^{1+\bar{\gamma}} \\ &\quad + \bar{c}''_3 [\exp\{-(n_0 \wedge n_1)^{1/2}\}]^{(1+\bar{\gamma})/\underline{\gamma}}. \end{aligned}$$

Lemma 7 guarantees that $\xi_{\alpha,\delta_0,n'_0}(\delta'_0) \leq (5/2)(n'_0)^{-1/4}$. Then the excess type II error is bounded by

$$\begin{aligned} P_1(\widehat{G}) - P_0(G^*) &\leq \bar{c}_1 (n'_0)^{-(\frac{1}{4} \wedge \frac{1+\bar{\gamma}}{4\underline{\gamma}})} + \bar{c}_2 (\lambda s)^{1+\bar{\gamma}} (n_0 \wedge n_1)^{\frac{1+\bar{\gamma}}{4}} \\ &\quad + \bar{c}_3 \exp \left\{ -(n_0 \wedge n_1)^{\frac{1}{2}} \left(\frac{1+\bar{\gamma}}{\underline{\gamma}} \wedge 1 \right) \right\}. \end{aligned}$$

■

References

A. Bloemendal, L. Erdos, A. Knowles, H.-T. Yau, and J. Yin. Isotropic local laws for sample covariance and generalized wigner matrices. *Electronic Journal of Probability*, 19:1–53, 2015.

- Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- Tony Cai and Weidong Liu. A direct estimation approach to sparse linear discriminant analysis. *J. Amer. Statist. Assoc.*, 106:1566–1577, 2011.
- A. Cannon, J. Howse, D. Hush, and C. Scovel. Learning with the neyman-pearson and min-max criteria. *Technical Report LA-UR-02-2951*, 2002.
- D. Casasent and X. Chen. Radial basis function neural networks for nonlinear fisher discrimination and neyman-pearson classification. *Neural Networks*, 16(5-6):529 – 535, 2003.
- Koei Chin, Sandy DeVries, Jane Fridlyand, Paul T Spellman, Ritu Roydasgupta, Wen-Lin Kuo, Anna Lapuk, Richard M Neve, Zuwei Qian, Tom Ryder, Fanqing Chen, Heidi Feiler, Taku Tokuyasu, Chris Kingsley, Shanaz Dairkee, Zhenhang Meng, Karen Chew, Daniel Pinkel, Ajay Jain, Britt Marie Ljung, Laura Esserman, Donna G Albertson, Frederic M Waldman, and Joe W Gray. Genomic and transcriptional aberrations linked to breast cancer pathophysiology. *Cancer Cell*, 10(6):529–541, December 2006.
- C. Elkan. The foundations of cost-sensitive learning. *In Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, pages 973–978, 2001.
- J. Fan, Y. Feng, and X. Tong. 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, 2012.
- Gavin J GJ Gordon, Roderick V RV Jensen, Li-Li LL Hsiao, Steven R SR Gullans, Joshua E JE Blumenstock, Sridhar S Ramaswamy, William G WG Richards, David J DJ Sugarbaker, and Raphael R Bueno. Translation of Microarray Data into Clinically Relevant Cancer Diagnostic Tests Using Gene Expression Ratios in Lung Cancer and Mesothelioma. *Cancer Research*, 62(17):4963–4967, September 2002.
- Yaqian Guo, Trevor Hastie, and Robert Tibshirani. Regularized discriminant analysis and its application in microarrays. *Biostatistics*, 1:1–18, 2005.
- Min Han, Dirong Chen, and Zhaoxu Sun. Analysis to Neyman-Pearson classification with convex loss function. *Analysis in Theory and Applications*, 24(1):18–28, 2008. ISSN 1672-4070. doi: 10.1007/s10496-008-0018-3.
- Daniel Hsu, Sham Kakade, and Tong Zhang. A tail inequality for quadratic forms of subgaussian random vectors. *Electronic Communications in Probability*, 17(52):1–6, 2012.
- Jingyi Jessica Li and Xin Tong. Genomic applications of the neyman-pearson classification paradigm. In *Big Data Analytics in Genomics*, pages 145–167. Springer, 2016.
- Q. Mai, H. Zou, and M. Yuan. A direct approach to sparse discriminant analysis in ultra-high dimensions. *Biometrika*, 99:29–42, 2012.
- E. Mammen and A.B. Tsybakov. Smooth discrimination analysis. *Annals of Statistics*, 27: 1808–1829, 1999.

- Rahul Mukerjee and SH Ong. Variance and covariance inequalities for truncated joint normal distribution via monotone likelihood ratio and log-concavity. *Journal of Multivariate Analysis*, 139:1–6, 2015.
- W. Polonik. Measuring mass concentrations and estimating density contour clusters-an excess mass approach. *Annals of Statistics*, 23:855–881, 1995.
- P. Rigollet and X. Tong. Neyman-pearson classification, convexity and stochastic constraints. *Journal of Machine Learning Research*, 12:2831–2855, 2011.
- C. Scott. Comparison and design of neyman-pearson classifiers. Unpublished, 2005.
- C. Scott and R. Nowak. A neyman-pearson approach to statistical learning. *IEEE Transactions on Information Theory*, 51(11):3806–3819, 2005.
- Jun Shao, Yazhen Wang, Xinwei Deng, and Sijian Wang. Sparse linear discriminant analysis by thresholding for high dimensional data. *Annals of Statistics*, 39:1241–1265, 2011.
- Xin Tong. A plug-in approach to neyman-pearson classification. *Journal of Machine Learning Research*, 14:3011–3040, 2013.
- Xin Tong, Yang Feng, and Anqi Zhao. A survey on neyman-pearson classification and suggestions for future research. *Wiley Interdisciplinary Reviews: Computational Statistics*, 8(2):64–81, 2016.
- Xin Tong, Yang Feng, and Jingyi Li. Neyman-Pearson (NP) Classification algorithms and NP receiver operating characteristics (NP-ROC). *Science Advances*, page eaao1659, 2018.
- Vladimir Vapnik. *The nature of statistical learning theory*. Springer, 1999.
- Charles Wang, Binsheng Gong, Pierre R Bushel, Jean Thierry-Mieg, Danielle Thierry-Mieg, Joshua Xu, Hong Fang, Huixiao Hong, Jie Shen, Zhenqiang Su, et al. The concordance between rna-seq and microarray data depends on chemical treatment and transcript abundance. *Nature biotechnology*, 32(9):926, 2014.
- D. Witten and R. Tibshirani. Penalized classification using fisher’s linear discriminant. *Journal of the Royal Statistical Society Series B*, 73:753–772, 2012.
- B. Zadrozny, J. Langford, and N. Abe. Cost-sensitive learning by cost-proportionate example weighting. *IEEE International Conference on Data Mining*, page 435, 2003.
- Anqi Zhao, Yang Feng, Lie Wang, and Xin Tong. Neyman-Pearson classification under high dimensional settings. *Journal of Machine Learning Research*, 17(213):1–39, 2016.