

High Dimensional Binary Classification under Label Shift: Phase Transition and Regularization

Jiahui Cheng Minshuo Chen Hao Liu Tuo Zhao Wenjing Liao ^{*}

Abstract

Label Shift has been widely believed to be harmful to the generalization performance of machine learning models. Researchers have proposed many approaches to mitigate the impact of the label shift, e.g., balancing the training data. However, these methods often consider the underparametrized regime, where the sample size is much larger than the data dimension. The research under the overparametrized regime is very limited. To bridge this gap, we propose a new asymptotic analysis of the Fisher Linear Discriminant classifier for binary classification with label shift. Specifically, we prove that there exists a phase transition phenomenon: Under certain overparametrized regime, the classifier trained using imbalanced data outperforms the counterpart with reduced balanced data. Moreover, we investigate the impact of regularization to the label shift: The aforementioned phase transition vanishes as the regularization becomes strong.

Keywords: Linear discriminant analysis, Binary classification, Label shift, Underparametrized and overparametrized regime, double descent phenomenon

1 Introduction

Label shift [58] occurs predominantly in classification tasks, in domains like computer vision [4, 72], medical diagnosis [26, 47], fraud detection [56] and others [28, 59]. Label shift often stems from, for example, a non-stationary environment and a biased way that the training and test data sets are collected.

The basic assumption in label shift [44] is that, while the class prior changes, the conditional distributions of data within a class are maintained in training and testing. A well known special case of label shift is learning with imbalanced data [14, 74] where the training are remarkably imbalanced due to some sampling bias, while the test data have a more balanced prior on the labels, e.g., uniform prior. It is commonly believed that training with imbalanced data can significantly undermine the overall performance of the trained classifiers [36, 51].

Typically, the test label distribution is unknown and many methods have been proposed to estimate the test priors [15, 65]. When the test label distribution has been estimated, learning under label shift [10, 21] reduces to the problem of resampling the training data. Common techniques include oversampling the minority class, downampling the majority class [12, 30, 71], and reweighting [33, 70]. For example, Seiffert et al. [61] integrate downampling with boosting for classification with imbalanced training data. The performance (measured by the area under the ROC curve) improves compared to base line methods (e.g., adaboost). However, these balancing method comes with several drawbacks. When the data imbalancing is extreme, downampling incurs significant loss of information, and oversampling can lead to over-fitting [16, 17]. Reweighting methods tend to make the optimization of deep models difficult [17, 33].

The aforementioned findings primarily focus on the underparametrized regime, where the sample size is much larger than the data dimension. Recently, significant progress have been made in training overparametrized models, such as deep neural networks. Such progress have stimulated empirical and theoretical studies on overparametrized models, whose statistical properties surprisingly challenge the conventional wisdom. For example, the typical U-shaped bias-variance trade-off curve is complemented

^{*}Jiahui Cheng and Minshuo Chen contributed equally to this work. Jiahui Cheng and Wenjing Liao are affiliated with the School of Mathematics at Georgia Institute of Technology. Minshuo Chen is affiliated with the Department of Electrical and Computer Engineering at Princeton University. Hao Liu is affiliated with the Department of Mathematics at Hong Kong Baptist University. Tuo Zhao is affiliated with the School of Industrial and Systems Engineering at Georgia Institute of Technology. Email: {jcheng328,tourzhao,wliao60}@gatech.edu, mc0750@princeton.edu, haoliu@hkbu.edu.hk. This research is partially supported by NSF DMS 2012652 and NSF CAREER 2145167.

by the double descent phenomenon observed in various models (see more in the related work section, Bartlett et al. [5], Belkin et al. [6], Hastie et al. [29], Mei and Montanari [52]).

In this paper, we study binary classification with label shift. The classifier is taken as Fisher Linear Discriminant Analysis (LDA, Bishop [11], Fisher [23]). We consider a Gaussian mixture data model, where $x \in \mathbb{R}^p$ is the feature, $y \in \{0, 1\}$ is the label, and $x | y$ is Gaussian distributed. Suppose n training data are collected under certain prior with n_ℓ samples in class ℓ such that $n_0 + n_1 = n$. Specifically, data imbalance refers to the case where the majority and minority class priors do not match, i.e., $n_1/n_0 \neq 1$. The test prior on each class is denoted as $\pi_\ell = \mathbb{P}(y = \ell)$ for $\ell = 0, 1$. We assume the test priors are known and different from the training priors. When the test priors are not known, we can still estimate them from the empirical label distributions in the test data.

Our contributions. We provide a theoretical analysis on the performance of LDA under label shift, in both the under- and over-parametrized regime. We explicitly quantify the misclassification error in the proportional limit of $n \rightarrow \infty$ and $p/n_\ell \rightarrow \gamma_\ell$ for $\ell = 0, 1$, where $\gamma_\ell > 0$ is a constant. Our theory shows a peaking phenomenon when the sample size is close to the data dimension.

We demonstrate a *phase transition* phenomenon about data imbalance: The misclassification error exhibits different behaviors as the two-class ratio n_1/n_0 varies, depending on the value of γ_0 . In particular, when γ_0 is fixed and the ratio n_1/n_0 increases from 1, we observe the following three phases:

- In the underparametrized regime (e.g., $\gamma_0 = 0.5$), the misclassification error first decreases then increases as γ_1 decays, yet the error decrease is marginal.
- In the lightly overparametrized regime (e.g., $\gamma_0 = 2.5$), the misclassification error first increases then decreases as γ_1 decays.
- In the overparametrized regime (e.g., $\gamma_0 = 5$), the misclassification error first decreases then increases, and finally decreases again as γ_1 decays.

Such a phase transition suggests that LDA trained with imbalanced data can outperform the counterpart trained with reduced balanced data, in certain overparametrized regime.

Moreover, we investigate the impact of the ℓ_2 regularization on the performance of LDA under label shift: The aforementioned phase transition vanishes when the regularization is sufficiently strong. While the phase transition persists when the regularization is weak.

Related work. In literature, many methods have been developed to handle classification under label shift. The sampling techniques include the informed upsampling [40, 45], synthetic oversampling [16], cluster-based oversampling [37]. Cost-sensitive methods [33, 66] use a cost matrix to represent the penalty of classifying examples from one class to another. Examples are cost-sensitive decision trees [49] and cost-sensitive neural networks [41]. Kernel-based methods are developed in [31, 46, 68]. Despite the empirical success, there are limited theories about how the classification results are affected by data imbalance.

Statistical properties of LDA has been well established in existing works (Anderson [1], Fukunaga [25, Section 10.2], Sifaou et al. [63], Velilla and Hernández [67], Zollanvari and Dougherty [75], Zollanvari et al. [76]). LDA with balanced training and test data is studied in Raudys and Duin [60] in the overparametrized case, while the assumption is more restrictive than ours, e.g., the feature vector is Gaussian with the identity covariance matrix in Raudys and Duin [60]. In the asymptotic regime (i.e., $p, n \rightarrow \infty$), Bickel and Levina [9] show that, when $p/n \rightarrow \infty$, LDA tends to random guessing. Later, Wang and Jiang [69] consider the proportional scenario where $p/n \rightarrow \gamma \in (0, 1)$. Our theory is more general, and covers both $0 < \gamma < 1$ (underparametrization) and $\gamma > 1$ (overparametrization). The misclassification error of Regularized LDA is analyzed in Elkhail et al. [22]. Our error analysis on LDA can not be implied from Elkhail et al. [22] by taking the limit of the regularization parameter to 0 since the covariance matrix is not invertible in the overparametrized case. We note a parallel line of work studying LDA in sparsity constrained high-dimensional binary classification problems [13, 48, 62].

Our theory demonstrates a peaking phenomenon of LDA, which has been recognized in history [20, 32, 34, 64] and recently for neural networks [6]. This phenomenon has been justified for linear regression [5, 7, 8, 29, 54], random feature regression [52], logistic regression [18], max-margin linear classifier [53], and others [19, 55, 73]. To our knowledge, we are the first to provide a theoretical justification of the peaking phenomenon for LDA under label shift.

The rest of the paper is organized as follows: Section 2 introduces LDA; Section 3 presents an asymptotic analysis of the misclassification error for LDA, and the phase transition phenomenon under data imbalance; Section 4 presents the impact of regularization; Section 5 presents real-data experiments;

Section 6 presents a proof of our main results; Section 7 discusses binary classification with extremely imbalanced data and in the highly overparametrized regime. We also discuss future directions.

Notation: Given a vector $v \in \mathbb{R}^p$, we denote $\|v\|_{\Sigma} = \sqrt{v^{\top} \Sigma v}$ for a positive definite matrix Σ . Given a matrix M , we denote M^{\dagger} as its pseudo-inverse. Let $\Phi(\cdot)$ be the CDF of the standard normal distribution. For two random variables X and Y , we denote $X \stackrel{d}{=} Y$ as X and Y having the same distribution. For a sequence of random variables $\{X_n\}$, we denote $X_n \xrightarrow{\text{a.s.}} X$ as the almost sure convergence.

2 Binary Classification using LDA

Binary classification aims at classifying an input feature $x \in \mathbb{R}^p$ into two classes labeled by $y \in \{0, 1\}$. A linear classifier achieves this goal by predicting the label based on a linear decision boundary in the form of $\beta^{\top}(x - \alpha) = b$ with $\alpha, \beta \in \mathbb{R}^p$ and $b \in \mathbb{R}$.

To be specific, a linear classifier gives the label y of the feature x by

$$f_{\alpha, \beta}^b(x) = \begin{cases} 0, & \text{if } \beta^{\top}(x - \alpha) > b, \\ 1, & \text{if } \beta^{\top}(x - \alpha) \leq b. \end{cases} \quad (1)$$

Given the class priors $\pi_0 + \pi_1 = 1$, we determine α, β and b by minimizing the misclassification error defined as

$$\begin{aligned} \mathcal{R}(f_{\alpha, \beta}^b) &= \pi_0 \mathbb{P}(f_{\alpha, \beta}^b(x) = 1 \mid y = 0) \\ &\quad + \pi_1 \mathbb{P}(f_{\alpha, \beta}^b(x) = 0 \mid y = 1). \end{aligned} \quad (2)$$

LDA approaches the binary classification problem by assuming that the conditional distribution of x given label $y = 0$ (resp. $y = 1$) is Gaussian $\mathcal{N}(\mu_0, \Sigma)$ (resp. $\mathcal{N}(\mu_1, \Sigma)$). Accordingly, the optimal classifier in LDA (also known as the Bayes rule) takes

$$\alpha^* = \frac{\mu_0 + \mu_1}{2}, \quad \beta^* = \Sigma^{-1}(\mu_0 - \mu_1), \quad b^* = \ln \frac{\pi_1}{\pi_0}. \quad (3)$$

The decision boundary β^* coincides with the Fisher linear discriminant rule [23], which maximizes the ratio of between-class variance and within-class variance:

$$\underset{\beta \in \mathbb{R}^p}{\operatorname{argmax}} \frac{(\beta^{\top} \mu_1 - \beta^{\top} \mu_0)^2}{\beta^{\top} \Sigma \beta}. \quad (4)$$

The optimal ratio in (4) at β^* is defined as the Signal-to-Noise Ratio (SNR), i.e., $\text{SNR} = \|\beta^*\|_{\Sigma}^2 = (\beta^*)^{\top} \Sigma \beta^*$.

In practice, we receive n_0 and n_1 i.i.d training data points from class 0 and 1, respectively. We denote the per-class data points as $\{x_i^{\ell}\}_{i=1}^{n_{\ell}}$ for $\ell = 0, 1$. The total number of samples is $n = n_0 + n_1$. We obtain the empirical Fisher linear discriminant classifier $f_{\hat{\alpha}, \hat{\beta}}^{\hat{b}}$ with

$$\hat{\alpha} = \frac{\hat{\mu}_0 + \hat{\mu}_1}{2}, \quad \hat{\beta} = \hat{\Sigma}^{\dagger}(\hat{\mu}_0 - \hat{\mu}_1), \quad \hat{b} = \ln \frac{n_1}{n_0}, \quad (5)$$

where $\hat{\mu}_0, \hat{\mu}_1$ and $\hat{\Sigma}$ are empirical estimators of μ_0, μ_1, Σ :

$$\hat{\mu}_{\ell} = \frac{1}{n_{\ell}} \sum_{i=1}^{n_{\ell}} x_i^{\ell}, \quad \hat{\Sigma} = \frac{1}{n-2} \sum_{\ell=0,1} \sum_{i=1}^{n_{\ell}} (x_i^{\ell} - \hat{\mu}_{\ell})(x_i^{\ell} - \hat{\mu}_{\ell})^{\top}.$$

3 Phase Transition of LDA under Label Shift

In this section, we present our main results on the misclassification error analysis of LDA, which covers both the under- and over-parametrized regime.

3.1 Error Analysis of LDA

We first introduce a data model for our theoretical analysis.

Assumption 1. For both the training and the test data, the conditional distribution of x given $y = \ell$ is Gaussian, i.e.,

$$x | (y = \ell) \sim \mathcal{N}(\mu_\ell, \Sigma), \quad \text{for } \ell = 0, 1.$$

The training data $\{x_i^\ell\}_{i=1}^{n_\ell}$ are i.i.d. sampled for class $\ell = 0, 1$, respectively. The test data have priors π_0, π_1 such that $\pi_0 + \pi_1 = 1$.

Assumption 1 allows arbitrarily imbalanced training data with $n_0 \neq n_1$ where the test label distribution can be different from the training data. Under Assumption 1, we prove an asymptotic behavior of the misclassification error of LDA in the limit of $n, p \rightarrow \infty$.

Theorem 2. Let γ_0, γ_1 and Δ^2 be positive constants and set $\gamma = \frac{\gamma_0 \gamma_1}{\gamma_0 + \gamma_1}$. Under Assumption 1, we let $n, p \rightarrow \infty$ with

$$p/n_0 \rightarrow \gamma_0, \quad p/n_1 \rightarrow \gamma_1, \quad p/n \rightarrow \gamma, \quad \text{and} \quad \|\beta^*\|_\Sigma^2 \rightarrow \Delta^2.$$

Then the misclassification error $\mathcal{R}(f_{\hat{\alpha}, \hat{\beta}}^b)$ converges to a limit almost surely when $\gamma \in (0, 1) \cup (1, \infty)$, i.e.,

$$\mathcal{R}(f_{\hat{\alpha}, \hat{\beta}}^b) \xrightarrow{\text{a.s.}} \sum_{\ell=0,1} \pi_\ell \Phi \left(\frac{\frac{g(\gamma_0, \gamma_1, \ell)}{1-\gamma} + (-1)^\ell \ln \frac{\gamma_0}{\gamma_1}}{k(\gamma_0, \gamma_1) \frac{1}{(1-\gamma)^{3/2}}} \right) \quad (6)$$

for $0 < \gamma < 1$, and

$$\mathcal{R}(f_{\hat{\alpha}, \hat{\beta}}^b) \xrightarrow{\text{a.s.}} \sum_{\ell=0,1} \pi_\ell \Phi \left(\frac{\frac{g(\gamma_0, \gamma_1, \ell)}{\gamma(\gamma-1)} + (-1)^\ell \ln \frac{\gamma_0}{\gamma_1}}{k(\gamma_0, \gamma_1) \frac{1}{(\gamma-1)^{3/2}}} \right) \quad (7)$$

for $\gamma > 1$, where

$$\begin{aligned} g(\gamma_0, \gamma_1, \ell) &= -\frac{1}{2}(\Delta^2 + (-1)^\ell(\gamma_0 - \gamma_1)) \\ k(\gamma_0, \gamma_1) &= \sqrt{\Delta^2 + \gamma_0 + \gamma_1}. \end{aligned}$$

Theorem 2 is proved in Section 6.2. Theorem 2 demonstrates a peaking phenomenon of LDA. For simplicity, we illustrate this peaking phenomenon for balanced training ($\gamma_0 = \gamma_1$) and test data ($\pi_0 = \pi_1 = 0.5$). In this case, the limit of $\mathcal{R}(f_{\hat{\alpha}, \hat{\beta}}^b)$ can be simplified to

$$\mathcal{R}(f_{\hat{\alpha}, \hat{\beta}}^b) \xrightarrow{\text{a.s.}} \begin{cases} \Phi \left(-\frac{\Delta^2 \sqrt{1-\gamma}}{2\sqrt{\Delta^2 + 4\gamma}} \right), & 0 < \gamma < 1, \\ \Phi \left(-\frac{\Delta^2 \sqrt{\gamma-1}}{2\gamma\sqrt{\Delta^2 + 4\gamma}} \right), & \gamma > 1. \end{cases}$$

Figure 1 (left panel) shows the asymptotic misclassification error as a function of γ with various SNR.

A peak occurs as γ approaches 1, when the training sample size is approximately equal to the data dimension. In the *underparametrized* regime ($0 < \gamma < 1$), the misclassification error increases with respect to γ . In the *overparametrized* regime ($\gamma > 1$), the misclassification error has a local minimum such that the error first decreases and then increases. This peaking phenomenon exists for all levels of SNR.

We interpret the peaking phenomenon as an interplay between the conditioning of the covariance matrix and the variance in the statistical estimation of the means and the covariance matrix. With balanced training data, the training model and the test model are the same so there is no model mismatch. In the underparametrized regime ($0 < \gamma < 1$), the covariance matrix is full rank. In this case, variance dominates the estimation error, and variance decreases as the sample size increases, since a larger number of samples yield better estimations of μ_0, μ_1 , and Σ . Therefore, the misclassification error decreases as γ decreases. In the overparametrized regime ($\gamma > 1$), the covariance matrix is rank-deficient. According to Bai-Yin theorem [3], the condition number of Σ^\dagger is proportional to $(1 - \sqrt{1/\gamma})^{-2}$, which decreases as γ increases from 1. As a result, the misclassification error decreases as γ increases from 1. When γ further increases, variance dominates the error due to limited number of samples, so the error increases again. This explains the local minimum when $\gamma > 1$.

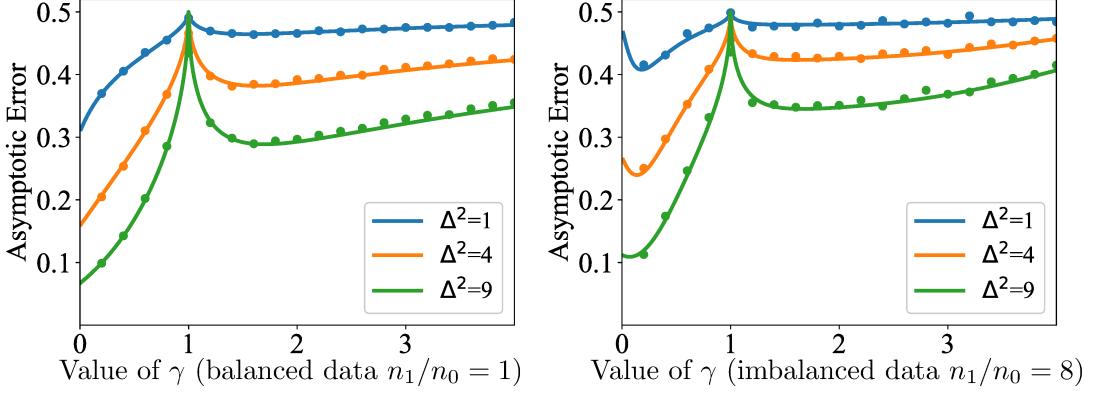


Figure 1: Misclassification error of LDA as a function of γ when $n_1/n_0 = 1$ (left panel) and $n_1/n_0 = 8$ (right panel), with balanced test data ($\pi_0 = \pi_1 = 0.5$), for various levels of asymptotic SNR. The solid curve represents the theoretical error given by Theorem 2, and the scatters denote synthetic data errors, with $n = 200, p = \lceil \gamma n \rceil$.

More interestingly, under the label shift with imbalanced training data ($\gamma_0 \neq \gamma_1$) and balanced test data ($\pi_0 = \pi_1 = 0.5$), there are trade-offs among three factors: (1) the conditioning of the pseudo-inverse of the covariance matrix; (2) the variance in the statistical estimation of the means and the covariance matrix; (3) an additional model mismatch. The misclassification error exhibits intricate behaviors with multiple local minima in the error curve. See Figure 1 (right panel) for an example with $n_1/n_0 = 8$.

3.2 Phase Transition under Label Shift

A critical question for classification under label shift is: When the class priors vary between the training data and the test data, is it beneficial to correct the training data distribution?

In order to evaluate the overall performance of classifier trained with imbalanced data, we usually consider a special case of label shift with balanced test data ($\pi_0 = \pi_1 = 0.5$). The question above is reduced to whether it is beneficial to downsample the majority class. Theorem 2 suggests an interesting dichotomy to the question above. Specifically, we fix γ_0 and investigate how the two-class ratio $n_1/n_0 \in [1, 10]$ affects the performance of LDA. We identify three distinct behaviors of LDA depending on the value of γ_0 . We demonstrate the three behaviors in Figure 2:

- *Behavior I*: When $\gamma_0 = 0.5$, i.e., the class of $\ell = 0$ is underparametrized, the misclassification error first decreases and then increases as a function of n_1/n_0 ;
- *Behavior II*: When $\gamma_0 = 2.5$, i.e., the class of $\ell = 0$ is slightly overparametrized, the misclassification error first increases and then decreases as a function of n_1/n_0 ;
- *Behavior III*: When $\gamma_0 = 5$, i.e., the class of $\ell = 0$ is overparametrized, the misclassification error first decreases and then increases, and finally decreases again, as a function of n_1/n_0 .

We obtain rich insights on training with imbalanced data from Figure 2. When the data imbalance is moderate, for example $n_1/n_0 \in [1, 3]$, training with imbalanced data can outperform the counterpart of using reduced balanced data as in *Behavior I* and *Behavior III*. Nonetheless, the improvement in *Behavior I* is only marginal. On the contrary, *Behavior II* indicates that downsampling the majority class improves the performance of LDA.

As the data imbalance becomes more significant, for example, $n_1/n_0 \in [8, 10]$, *Behavior II* and *III* both indicate that downsampling the majority class hurts the performance. Such behavior is expected, since the downsampling incurs severe information loss.

Theorem 2 also characterizes the misclassification error when n_1/n_0 is extremely large. We can check that the misclassification error converges to 0.5 as $n_1/n_0 \rightarrow \infty$, regardless of the value of γ_0 (see Appendix A). This indicates that extreme label shift renders the trained classifier suffering from the model mismatch. Nonetheless, in such an extreme imbalanced case, the minority group is prone to be outliers, and detection of outliers is also of great interest.

In the sequel, we formally characterize three phases corresponding to the aforementioned different behaviors. We explicitly identify two phase transition knots γ_a and γ_b (derived in Appendix A):

$$\gamma_a = 2 \quad \text{and} \quad \gamma_b = \frac{1}{8}(12 - \Delta^2 + \sqrt{\Delta^4 + 40\Delta^2 + 144}).$$

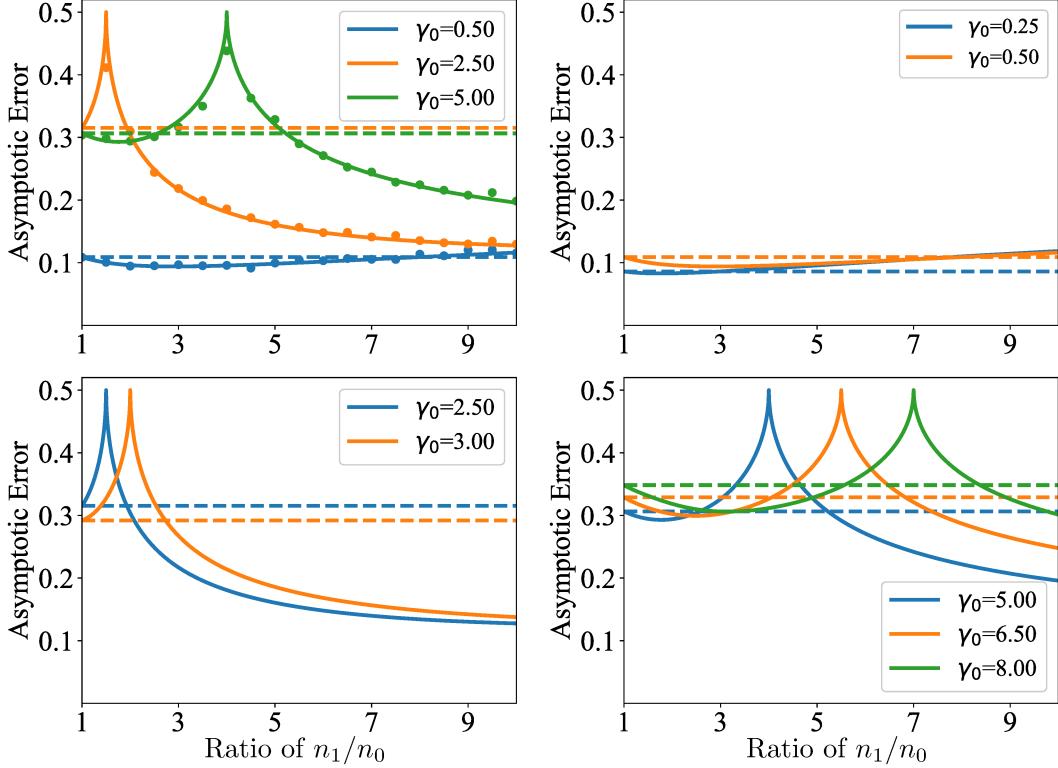


Figure 2: We demonstrate three behaviors of the misclassification error of LDA as a function of n_1/n_0 for various γ_0 and a fixed asymptotic SNR $\Delta^2 = 9$. The solid curve represents the error based on our theoretical analysis in Theorem 2. The scatters denote synthetic data errors, with $n_0 = 40, p = \lceil \gamma_0 n_0 \rceil$. Upper left panel: Plots of an underparametrized ($\gamma_0 = 0.5$), a slightly overparametrized ($\gamma = 2.5$) and an overparametrized ($\gamma = 5$) case. Upper right panel: Behavior I in the underparametrized cases; Bottom Left panel: Behavior II in the slightly overparametrized cases. Bottom right panel: Behavior III in the overparametrized cases.

We claim three phases depending on the value of γ_0 .

- **Phase I:** If $\gamma_0 \in (0, \gamma_a)$, the misclassification error has *Behavior I*;
- **Phase II:** If $\gamma_0 \in (\gamma_a, \gamma_b)$, the misclassification error has *Behavior II*;
- **Phase III:** When $\gamma_0 \in (\gamma_b, \gamma_c)$ for some $\gamma_c > 0$, the misclassification error has *Behavior III*.

We observe that the first transition γ_a appears at the exact parametrized case, i.e., $p = n$. The second transition γ_b depends on the SNR and is always larger than γ_a .

We remark that in **Phase III**, we cut off γ_0 at some threshold γ_c . If γ_0 is extremely large, i.e., the problem is highly overparametrized, we can observe a fourth behavior on the misclassification error. In fact, the misclassification error has multiple local maxima, reflecting a complex interaction between the limited information in the training data and the mismatch of the training and test model. We discuss the highly overparametrized regime in Section 7.

4 Regularization Impact on LDA

In machine learning, regularization is commonly used to stabilize the computation and improve the generalization performance. In this section, we study regularized LDA [24, 27] and analyze its asymptotic misclassification error.

4.1 Error Analysis of Regularized LDA

When an ℓ_2 regularization term is added on β , we consider the following optimization problem based on (4):

$$\operatorname{argmin}_{\beta \in \mathbb{R}^p} -(\beta^\top \mu_1 - \beta^\top \mu_0)^2 + \lambda \|\beta\|_2^2, \quad \text{s.t. } \beta^\top \Sigma \beta = 1.$$

This gives rise to an optimal solution β_λ^* . For simplicity, we formalize it in an equivalent form as

$$\beta_\lambda^* = (\Sigma + \lambda I)^{-1}(\mu_0 - \mu_1). \quad (8)$$

We denote the regularized Fisher linear discriminant classifier by $f_{\alpha^*, \beta_\lambda^*}^{b^*}$ where α^* and b^* are given in (3).

The empirical counterpart of the regularized LDA is given by $f_{\hat{\alpha}, \hat{\beta}_\lambda}^{\hat{b}}$, where the empirical parameters $\hat{\alpha}$, $\hat{\beta}_\lambda$ and \hat{b} are computed through $\hat{\mu}_0, \hat{\mu}_1$ and $\hat{\Sigma}$ according to (5) and

$$\hat{\beta}_\lambda = (\hat{\Sigma} + \lambda I)^{-1}(\hat{\mu}_0 - \hat{\mu}_1).$$

Similar to Theorem 2, we prove an asymptotic behavior of the misclassification error of the regularized LDA.

Theorem 3. *Let γ_0, γ_1 and Δ^2 be positive constants and set $\gamma = \frac{\gamma_0 \gamma_1}{\gamma_0 + \gamma_1}$. Under Assumption 1, we let $n, p \rightarrow \infty$ with*

$$p/n_0 \rightarrow \gamma_0, \quad p/n_1 \rightarrow \gamma_1, \quad p/n \rightarrow \gamma, \quad \text{and} \quad \|\beta^*\|_\Sigma^2 \rightarrow \Delta^2.$$

The misclassification error $\mathcal{R}(f_{\hat{\alpha}, \hat{\beta}_\lambda}^{\hat{b}})$ converges to a limit almost surely when $\gamma \in (0, 1) \cup (1, \infty)$, i.e.,

$$\mathcal{R}(f_{\hat{\alpha}, \hat{\beta}_\lambda}^{\hat{b}}) \xrightarrow{\text{a.s.}} \sum_{\ell=0,1} \pi_\ell \Phi \left(\frac{g(\gamma_0, \gamma_1, \ell) m(-\lambda) + (-1)^\ell \ln \frac{\gamma_0}{\gamma_1}}{k(\gamma_0, \gamma_1) \sqrt{m'(-\lambda)}} \right),$$

where $g(\gamma_0, \gamma_1, \ell)$ and $k(\gamma_0, \gamma_1)$ are defined in Theorem 2, and $m(\lambda) = \int 1/(s-\lambda) dF_\lambda(s)$ with F_λ denoting the Marchenko-Pastur law.

Theorem 3 is proved in Appendix C.1. Theorem 3 implies that regularization has a smoothing effect on the peaking phenomenon. For simplicity, we show this smoothing effect for balanced training ($\gamma_0 = \gamma_1$) and test data ($\pi_0 = \pi_1 = 0.5$).

Figure 3 shows the asymptotic misclassification error of regularized LDA as a function of γ . When the regularization is weak, e.g., $\lambda = 10^{-4}$, we observe a similar peaking phenomenon as in Figure 1, while the peak is lower than that in Figure 1. Compared to the unregularized classifier $f_{\hat{\alpha}, \hat{\beta}}^{\hat{b}}$, regularization improves the conditioning of the estimated covariance matrix, i.e., $\hat{\Sigma} + \lambda I$ is never singular, which in turn mitigates the performance degradation when $\gamma \approx 1$.

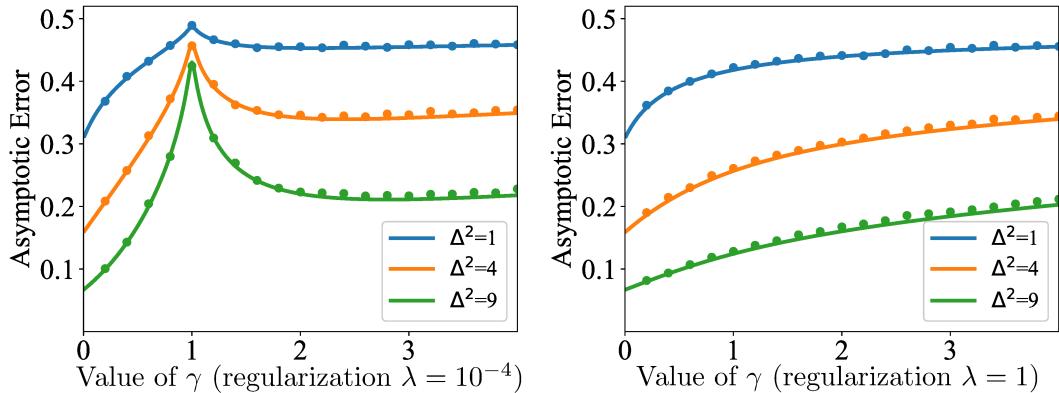


Figure 3: Misclassification error of regularized LDA as a function of γ for various levels of asymptotic SNR. The left panel shows the impact of weak regularization $\lambda = 10^{-4}$, and the right panel uses a strong regularization $\lambda = 1$. The solid line represents the theoretical error given by Theorem 3, and the scatters denote synthetic data errors, with $n = 200, p = \lceil \gamma n \rceil$.

When the regularization is strong, e.g., $\lambda = 1$, the peaking phenomenon disappears in Figure 3 (right panel). In this case, the matrix $\hat{\Sigma} + \lambda I$ is always well-conditioned. The error is dominated by the variance in the statistical estimation of the means and the covariance matrix. As a result, the error increases as γ increases. We remark that proper regularization greatly reduces the misclassification error when the problem is approximately exactly parametrized ($\gamma \approx 1$).

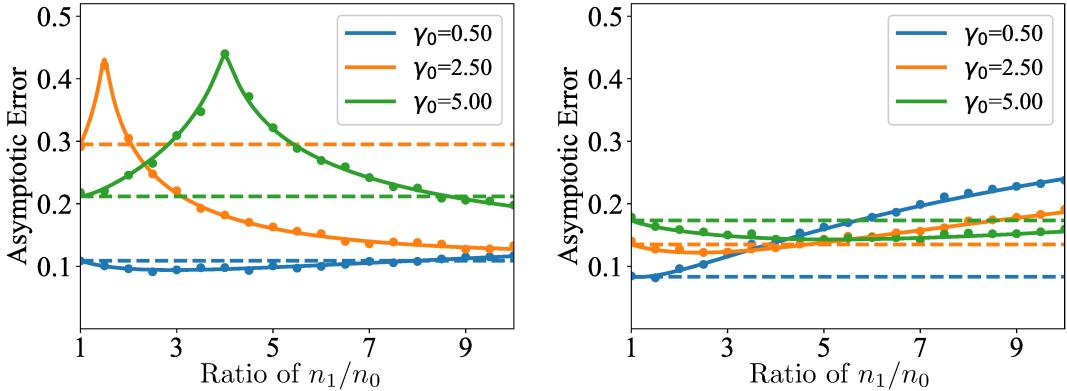


Figure 4: We demonstrate three behaviors of the misclassification error of regularized LDA as a function of n_1/n_0 for various γ_0 when we use a weak regularization with $\lambda = 10^{-4}$ (left panel). They disappears with a strong regularization $\lambda = 1$ (right panel). We fix the asymptotic SNR $\Delta^2 = 9$. The scatters denote synthetic data errors, with $n_0 = 40, p = \lceil \gamma_0 n_0 \rceil$.

4.2 Phase Transition of Regularized LDA

In this section, we study the impact of the regularization on the phase transition phenomenon discussed in Section 3.2. We discuss the impact of weak and strong regularization separately, as they lead to very different behaviors.

When the regularization is weak, e.g., $\lambda = 10^{-4}$, we observe a similar phase transition phenomenon as in Section 3.2. We depict the misclassification error curves as a function of n_1/n_0 in Figure 4.

When the regularization is strong, e.g., $\lambda = 1$, the phase transition phenomenon disappears (See a formal justification in Appendix C.2). Figure 4 shows that the asymptotic misclassification error of regularized LDA as a function of n_1/n_0 for various γ_0 . We observe that the error curve consistently first decreases and then increases. In this case, the matrix $\widehat{\Sigma} + \lambda I$ is always well-conditioned. Therefore the misclassification error is the consequence of the trade-off between two factors: 1) the model mismatch and 2) the variance in the statistical estimation of the means and the covariance matrix.

5 Real-Data Binary Classification

We connect our theoretical findings to real-data binary classification tasks. We consider Neural Network (NN) classifiers for the MNIST (CC-BY 3.0) and CIFAR-10 (MIT) datasets [39, 42]. We focus on the overparametrized regime, which is the working regime for neural networks.

MNIST dataset and NN classifier. The MNIST dataset consists of handwritten digits of resolution 28×28 . We train a neural network classifier with one hidden layer to distinguish digits 3 and 8. We vary the number of hidden units in $\{32, 64, 128\}$. The activation function is ReLU, i.e., $\sigma(\cdot) = \max\{\cdot, 0\}$. We use Adam [38] for training, with default hyperparameters in Pytorch.

CIFAR-10 dataset and NN classifier. The CIFAR-10 dataset consists of RGB images of resolution 32×32 from 10 categories. We pick two similar categories, e.g., horse v.s. deer, for binary classification. We downsample the data to a resolution of 16×16 . We also train an NN classifier with one hidden layer. The number of neurons in the hidden layer is 10, and the activation function is ReLU. We use momentum SGD for training, with momentum coefficient 0.9 and learning rate 0.001.

In both tasks, during training, we fix the number of samples from one category (500 in MNIST and 1000 in CIFAR-10), and vary the samples in the other category. The total number of training epochs is 50 for MNIST and 20 for CIFAR-10. After training, the classifier is tested on a balanced test set, which consists of 900 samples from each class in MNIST, and 1000 samples per-class in CIFAR-10. The test error is averaged over 10 independent runs for MNIST and 5 independent runs for CIFAR-10 with random seeds.

Result. The misclassification error in MNIST and CIFAR-10 as a function of n_1/n_0 is plotted in Figure 5. In both the MNIST and CIFAR-10 experiments, the neural network is overparametrized. Therefore, we expect the misclassification error exhibits *Behavior III* in Figure 2. This is corroborated in Figure 5 as the misclassification error first decreases and then increases as n_1/n_0 grows.

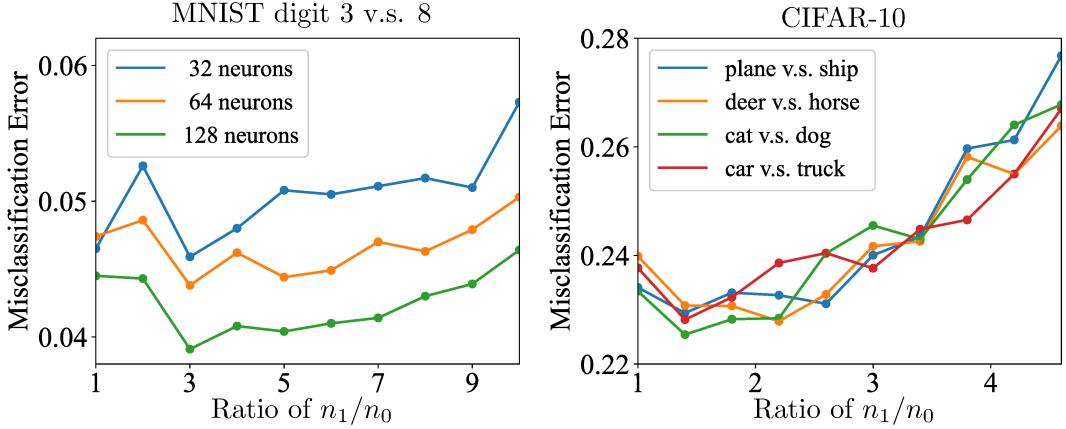


Figure 5: Misclassification error of neural network classifiers on MNIST and CIFAR-10 for binary classification.

More importantly, these experiments consistently indicate that in the overparametrized regime, down-sampling the majority class may hurt the model performance (cf. $n_1/n_0 \approx 3$ in MNIST and $n_1/n_0 \approx 1.5$ in CIFAR-10). Meanwhile, when the data imbalance is relatively severe, down-sampling the majority class can be beneficial, due to its mitigation on label shift between training and testing.

6 Proof of Theorem 2

In this Section, we prove our main Theorem 2. The proof for Theorem 3 about the regularized LDA is similar to the proof of Theorem 2. The proof of Theorem 3 is given in Appendix C.1.

6.1 Lemmas to be used for the proof of Theorem 2

Our analysis relies on change of variables to exploit the independence between the sample mean estimator and the covariance estimator.

Lemma 4 (Independence between sample mean and sample covariance). *Let $x_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mu, \Sigma)$ for $i = 1, \dots, n$, be samples from a Gaussian distribution. We denote $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n x_i$ as the estimator of the sample mean, and $\hat{\Sigma} = \frac{1}{n-1} \sum_{i=1}^n (x_i - \hat{\mu})(x_i - \hat{\mu})^\top$ as the estimator of the sample covariance. Then $\hat{\mu}$ and $\hat{\Sigma}$ are independent.*

We also utilize the asymptotic characterization of the spectrum of Wishart matrix.

Lemma 5 (Isotropicity of wishart matrix). *Assume $\frac{Z^\top Z}{n-2} \sim \mathcal{W}(I_p, n-2)$. For any vector $z \in \mathbb{R}^p$ independent of Z , we have*

$$z^\top \left(\frac{Z^\top Z}{n-2} \right)^\dagger z \stackrel{d}{=} \frac{1}{p} \|z\|_2^2 \text{tr} \left(\left(\frac{Z^\top Z}{n-2} \right)^\dagger \right). \quad (9)$$

Lemma 4 and 5 are proved in Appendix B.1. We next present some helper lemmas.

Lemma 6. *Given a matrix Z with i.i.d. standard normal distributed entries, we have*

$$\text{tr} \left((Z^\top Z)^\dagger \right) = \text{tr} \left((ZZ^\top)^\dagger \right). \quad (10)$$

Lemma 7 (Strong law of large numbers). *Assume $z \sim \mathcal{N}(0, I_p)$ and μ_d is a non-random p -dimensional vector such that $\|\mu_d\|_2^2 \xrightarrow{\text{a.s.}} \Delta^2$ and n_ℓ satisfies Binomial distribution $B(n_0 + n_1, \pi_\ell)$ then we have*

$$\frac{1}{\sqrt{p}} \mu_d^\top z \xrightarrow{\text{a.s.}} 0, \quad \frac{1}{p} z^\top z \xrightarrow{\text{a.s.}} 1, \quad \frac{n_\ell}{n_0 + n_1} \xrightarrow{\text{a.s.}} \pi_\ell.$$

Lemma 8 (Marchenko-Pastur law). *Let F_γ be the Marchenko-Pastur (MP) law of $0 < \gamma < 1$. Then for any real number $\zeta \leq 0$, the Stieltjes transform of MP law and its derivative at ζ are given as*

$$m(\zeta) = \int \frac{1}{s - \zeta} dF_\gamma(s) = \frac{1 - \gamma - \zeta - \sqrt{(\zeta - \gamma - 1)^2 - 4\gamma}}{2\gamma\zeta} \quad (11)$$

and

$$\frac{d}{d\zeta} m(\zeta) = \int \frac{1}{(s - \zeta)^2} dF_\gamma(s) = \frac{2\gamma(\gamma - 1) - \frac{2\gamma(\zeta(\gamma+1) - (\gamma-1)^2)}{\sqrt{(\zeta - \gamma - 1)^2 - 4\gamma}}}{4\gamma^2\zeta^2}. \quad (12)$$

In particular, we have

$$m(0) = \frac{1}{1 - \gamma} \quad \text{and} \quad \frac{d}{d\zeta} m(0) = \frac{1}{(1 - \gamma)^3}.$$

Lemma 6, 7 and 8 are proved in Appendix B.2.

6.2 Proof of Theorem 2

Proof of Theorem 2. To begin with, we recall the misclassification error of binary classification in (2). Substituting the Fisher linear discriminant classifier $f_{\hat{\alpha}, \hat{\beta}}^b$ (with $\hat{\alpha}, \hat{\beta}, \hat{b}$ given in (5)) and prior $\pi_0 = \pi_1 = 1/2$ into (2), we derive

$$\begin{aligned} \mathcal{R}(f_{\hat{\alpha}, \hat{\beta}}^b) &= \frac{1}{2} \mathbb{P} \left(\hat{\beta}^\top (x - \hat{\alpha}) \leq \ln \frac{n_1}{n_0} \mid x \sim \mathcal{N}(\mu_0, \Sigma) \right) \\ &\quad + \frac{1}{2} \mathbb{P} \left(\hat{\beta}^\top (x - \hat{\alpha}) > \ln \frac{n_1}{n_0} \mid x \sim \mathcal{N}(\mu_1, \Sigma) \right) \\ &= \frac{1}{2} \mathbb{P} \left(\frac{\hat{\beta}^\top (x - \mu_0)}{\|\hat{\beta}\|_\Sigma} \leq \frac{\hat{\beta}^\top (\hat{\alpha} - \mu_0) + \frac{n_1}{n_0}}{\|\hat{\beta}\|_\Sigma} \mid x \sim \mathcal{N}(\mu_0, \Sigma) \right) \\ &\quad + \frac{1}{2} \mathbb{P} \left(\frac{\hat{\beta}^\top (x - \mu_1)}{\|\hat{\beta}\|_\Sigma} > \frac{\hat{\beta}^\top (\hat{\alpha} - \mu_1) + \ln \frac{n_1}{n_0}}{\|\hat{\beta}\|_\Sigma} \mid x \sim \mathcal{N}(\mu_1, \Sigma) \right) \\ &= \frac{1}{2} \Phi \left(\left[\hat{\beta}^\top (\hat{\alpha} - \mu_0) + \ln \frac{n_1}{n_0} \right] / \|\hat{\beta}\|_\Sigma \right) + \frac{1}{2} \Phi \left(\left[\hat{\beta}^\top (\mu_1 - \hat{\alpha}) + \ln \frac{n_0}{n_1} \right] / \|\hat{\beta}\|_\Sigma \right). \end{aligned} \quad (13)$$

Therefore, it suffices to find the limits of

$$q_0 = \left(\hat{\beta}^\top (\hat{\alpha} - \mu_0) + \ln \frac{n_1}{n_0} \right) / \|\hat{\beta}\|_\Sigma \quad \text{and} \quad q_1 = \left(\hat{\beta}^\top (\mu_1 - \hat{\alpha}) + \ln \frac{n_0}{n_1} \right) / \|\hat{\beta}\|_\Sigma,$$

since the Gaussian CDF Φ is continuous.

To further aid our analysis, we characterize the distributions of $\hat{\alpha}$, $\hat{\beta}$, and $\hat{\Sigma}$. Specifically, we make the following change of variables:

$$\begin{aligned} \hat{\mu}_0 &\stackrel{d}{=} \frac{1}{\sqrt{n_0}} \Sigma^{\frac{1}{2}} z_0 + \mu_0, \\ \hat{\mu}_1 &\stackrel{d}{=} \frac{1}{\sqrt{n_1}} \Sigma^{\frac{1}{2}} z_1 + \mu_1, \\ \hat{\Sigma} &\stackrel{d}{=} \frac{1}{n-2} \Sigma^{\frac{1}{2}} Z^\top Z \Sigma^{\frac{1}{2}}, \end{aligned} \quad (14)$$

where $z_0, z_1 \sim \mathcal{N}(0, I_p)$ and $Z \in \mathbb{R}^{(n-2) \times p}$ with each element $Z_{i,j} \sim \mathcal{N}(0, 1)$. Note that z_0, z_1 and Z are independent with each other and $Z^\top Z$ is a Wishart matrix by Lemma 4.

• **Case 1.** $0 < \gamma < 1$. We present in detail how to characterize the asymptotic limit of q_0 , and q_1 follows a similar argument. We tackle the numerator and denominator of q_0 separately. Using Lemma 7, we check

$$\ln \frac{n_1}{n_0} \xrightarrow{\text{a.s.}} \ln \frac{\pi_1}{\pi_0} = \ln \frac{\gamma_0}{\gamma_1}.$$

Therefore, we temporarily omit the threshold term $\ln(n_1/n_0)$ in both q_0, q_1 to ease the presentation. By the change of variables formula in (14) and some manipulation, we deduce

$$\begin{aligned}\widehat{\beta}^\top(\widehat{\alpha} - \mu_0) &= -\frac{1}{2}(\mu_d + \frac{1}{\sqrt{n_0}}z_0 - \frac{1}{\sqrt{n_1}}z_1)^\top \left(\frac{Z^\top Z}{n-2}\right)^\dagger (\mu_d - \frac{1}{\sqrt{n_0}}z_0 - \frac{1}{\sqrt{n_1}}z_1) \\ &= \frac{1}{2n_0}z_0^\top \left(\frac{Z^\top Z}{n-2}\right)^\dagger z_0 - \frac{1}{2}(\mu_d - z_1/\sqrt{n_1})^\top \left(\frac{Z^\top Z}{n-2}\right)^\dagger (\mu_d - z_1/\sqrt{n_1})\end{aligned}\quad (15)$$

$$\stackrel{d}{=} \underbrace{\left[-\frac{1}{2}\|\mu_d - \frac{1}{\sqrt{n_1}}z_1\|_2^2 + \frac{1}{2n_0}\|z_0\|_2^2\right]}_A \times \underbrace{\frac{1}{p}\text{tr}\left(\frac{Z^\top Z}{n-2}\right)^\dagger}_B. \quad (16)$$

where $\mu_d = \Sigma^{-1/2}(\mu_0 - \mu_1)$ and the last equality follows from the isotropicity of the Wishart matrix $Z^\top Z$ in lemma 5, In the sequel, we establish the limits of terms A and B , since they are independent.

Asymptotic convergence of A . We show

$$A \xrightarrow{\text{a.s.}} -\frac{1}{2}(\Delta^2 + \gamma_1 - \gamma_0). \quad (17)$$

To see the result above, we expand term A as

$$A = -\frac{1}{2}\|\mu_d\|^2 + \frac{1}{\sqrt{n_1}}\mu_d^\top z_1 - \frac{1}{2n_1}\|z_1\|^2 + \frac{1}{2n_0}\|z_0\|^2.$$

Recall that in Lemma 7, we verify that $\frac{1}{n_1}\|z_1\|_2^2 = \frac{p}{n_1}\frac{1}{p}\|z_1\|_2^2 \xrightarrow{\text{a.s.}} \gamma_1$ and $\frac{1}{n_0}\|z_0\|_2^2 \xrightarrow{\text{a.s.}} \gamma_0$ by the strong law of large numbers again. Using the Borel-Cantelli lemma, we claim $\frac{1}{\sqrt{n_1}}\mu_d^\top z_1 \xrightarrow{\text{a.s.}} 0$. Invoking the assertion in Theorem 2 that $\|\beta^*\|_\Sigma^2 = \|\mu_d\|_2^2 = \Delta^2$, we deduce the desired convergence of A .

Asymptotic convergence of B . We next show

$$B \xrightarrow{\text{a.s.}} \frac{1}{1-\gamma}. \quad (18)$$

The convergence result above utilizes the Stieljes transformation of the Marchenko-Pastur law. Specifically, by the Bai-Yin theorem [3], for $\gamma < 1$,

$$\zeta_{\min}\left(\frac{Z^\top Z}{n-2}\right) \geq \frac{1}{2}(1-\sqrt{\gamma})^2,$$

which implies that $Z^\top Z$ is almost surely invertible. Conditioned on $Z^\top Z$ being invertible, we rewrite B as

$$B = \frac{1}{p}\text{tr}\left(\frac{Z^\top Z}{n-2}\right)^{-1} = \frac{1}{p}\sum_{i=1}^p \frac{1}{s_i} = \int \frac{1}{s}dF_{\frac{Z^\top Z}{n-2}}(s),$$

where s_i 's denote the eigenvalues of $Z^\top Z/(n-2)$ and $F_M(a) = \frac{1}{p}\sum_{i=1}^p \mathbb{1}\{\lambda_i(M) \leq a\}$ is the empirical measure of eigenvalues of M .

Now apply the Marchenko-Pastur theorem [50], which says that $F_{Z^\top Z/(n-2)}$ converges weakly, almost surely, to the Marchenko-Pastur law F_γ (depending only on γ). Invoking the Portmanteau theorem [50], weak convergence is equivalent to the convergence in expectation of all bounded functions h , that are continuous except on a set of zero probability under the limiting measure. Defining $h(s) = 1/s \cdot \mathbb{1}\{s \geq a/2\}$, where we abbreviate $a = (1-\sqrt{\gamma})^2$, it follows that as $n, p \rightarrow \infty$, almost surely,

$$\int_{a/2}^\infty \frac{1}{s}dF_{\frac{Z^\top Z}{n-2}}(s) \rightarrow \int_{a/2}^\infty \frac{1}{s}dF_\gamma(s)$$

We can remove the lower limit of integration on both sides above; for the right-hand side, this follows since support of the Marchenko-Pastur law F_γ is $[a, b]$, where $b = (1+\sqrt{\gamma})$; for the left-hand side, this follows again by the Bai-Yin theorem [3] (which as already stated, implies the smallest eigenvalues of $Z^\top Z/(n-2)$) is almost surely greater than $a/2$ for large enough n). Thus the last display implies that as $n, p \rightarrow \infty$, almost surely,

$$B \rightarrow \int \frac{1}{s}F_\gamma(s)$$

Thanks to the Stieljes transformation of the Marchenko-Pastur law, we can explicitly compute $\int \frac{1}{s} dF_\gamma(s)$. In particular, the Stieljes transformation of F_γ is defined as

$$m(\zeta) = \int \frac{1}{s - \zeta} dF_\gamma(s) \quad \text{with } \zeta \in \mathbb{C}.$$

Taking $\zeta \rightarrow 0$, we obtain $m(0) = \lim_{\zeta \rightarrow 0} m(\zeta) = \frac{1}{1-\gamma}$ by Lemma 8. Consequently, we establish $B \xrightarrow{\text{a.s.}} \frac{1}{1-\gamma}$.

Substituting (17) and (18) into (16) yields

$$\widehat{\beta}^\top (\widehat{\alpha} - \mu_0) \xrightarrow{\text{a.s.}} -\frac{1}{2}(\Delta^2 + \gamma_1 - \gamma_0) \frac{1}{1-\gamma}. \quad (19)$$

Next we consider the denominator $\|\widehat{\beta}\|_\Sigma$ in q_0 . Applying the change of variables in (14) and using Lemma 5, analogous to (16), we derive

$$\begin{aligned} \|\widehat{\beta}\|_\Sigma^2 &= (\widehat{\mu}_0 - \widehat{\mu}_1)^\top \widehat{\Sigma}^\dagger \Sigma \widehat{\Sigma}^\dagger (\widehat{\mu}_0 - \widehat{\mu}_1) \\ &= (\mu_d + \frac{1}{\sqrt{n_0}} z_0 - \frac{1}{\sqrt{n_1}} z_1)^\top \left(\left(\frac{Z^\top Z}{n-2} \right)^\dagger \right)^2 (\mu_d + \frac{1}{\sqrt{n_0}} z_0 - \frac{1}{\sqrt{n_1}} z_1) \\ &\stackrel{d}{=} \underbrace{\left\| \mu_d + \frac{1}{\sqrt{n_0}} z_0 - \frac{1}{\sqrt{n_1}} z_1 \right\|_2^2}_{A'} \times \underbrace{\frac{1}{p} \text{tr} \left(\left(\frac{Z^\top Z}{n-2} \right)^\dagger \right)^2}_{B'}. \end{aligned} \quad (20)$$

The convergence of term A' follows the same argument of term A in the numerator, and we have

$$A' \xrightarrow{\text{a.s.}} \Delta^2 + \gamma_1 + \gamma_0. \quad (21)$$

Conditioned on $Z^\top Z$ being invertible, term B' can be written as

$$B' = \int \frac{1}{s^2} dF_{\frac{Z^\top Z}{n-2}}(s) \longrightarrow \int \frac{1}{s^2} dF_\gamma(s).$$

To compute the limiting integral above, we differentiate the Stieljes transformation $m(\zeta)$. By sending $\zeta \rightarrow 0$ again, we can derive

$$B' \xrightarrow{\text{a.s.}} \lim_{\zeta \rightarrow 0} \frac{d}{d\zeta} m(\zeta) = \frac{1}{(1-\zeta)^3}. \quad (22)$$

Substituting (21) and (22) into (20) yields

$$\|\widehat{\beta}\|_\Sigma \xrightarrow{\text{a.s.}} \sqrt{\Delta^2 + \gamma_1 + \gamma_0} \frac{1}{(1-\gamma)^{3/2}}. \quad (23)$$

Combining (19) and (23), as well as putting the threshold term $\ln(n_1/n_0)$ back, we obtain

$$q_0 \xrightarrow{\text{a.s.}} \frac{-\frac{\Delta^2 + \gamma_1 - \gamma_0}{2(1-\gamma)} + \ln \frac{\gamma_0}{\gamma_1}}{\sqrt{\Delta^2 + \gamma_1 + \gamma_0} \frac{1}{(1-\gamma)^{3/2}}}. \quad (24)$$

The same argument of analyzing q_0 applies to q_1 , and therefore, we have

$$q_1 \xrightarrow{\text{a.s.}} \frac{-\frac{\Delta^2 + \gamma_0 - \gamma_1}{2(1-\gamma)} + \ln \frac{\gamma_1}{\gamma_0}}{\sqrt{\Delta^2 + \gamma_1 + \gamma_0} \frac{1}{(1-\gamma)^{3/2}}}. \quad (25)$$

To complete the proof in the case of $0 < \gamma < 1$, we plugging (24) and (25) into (13).

• **Case 2.** $\gamma > 1$. The goal is still to find the limits of q_0 and q_1 . Consider q_0 first. We observe that both (16) and (20) are valid for $\gamma > 1$. However, a key difference is that $Z^\top Z$ is rank deficient in the limit considering $\gamma > 1$. To resolve this issue, we observe that $Z^\top Z$ and ZZ^\top share all nonzero eigenvalues.

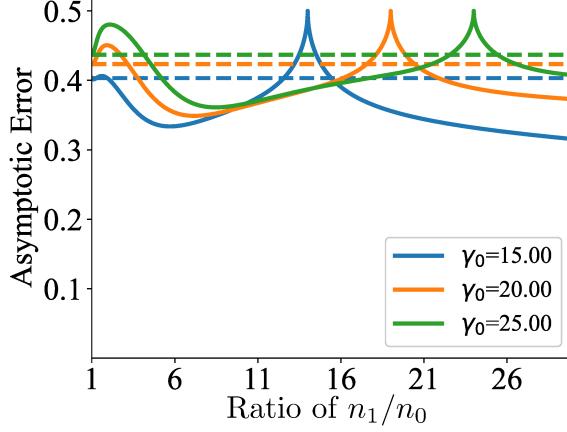


Figure 6: Misclassification error in the highly overparametrized regime as a function of n_1/n_0 , with $\Delta^2 = 9$.

Therefore, we can replace $Z^\top Z$ by ZZ^\top in (16) and (20) without changing their values. Such a reasoning is justified by Lemma 6. We can now rewrite (16) as

$$\begin{aligned} \widehat{\beta}^\top (\widehat{\alpha} - \mu_0) &\stackrel{d}{=} \left[-\frac{1}{2} \|\mu_d - \frac{1}{\sqrt{n_1}} z_1\|_2^2 + \frac{1}{2n_0} \|z_0\|_2^2 \right] \times \frac{(n-2)^2}{p^2} \frac{1}{n-2} \text{tr} \left(\frac{ZZ^\top}{p} \right)^\dagger. \\ &\xrightarrow{\text{a.s.}} -\frac{1}{2} [\Delta^2 + \gamma_1 - \gamma_0] \times \frac{1}{\gamma(\gamma-1)}. \end{aligned} \quad (26)$$

Similarly we can rewrite (20) as

$$\begin{aligned} \|\widehat{\beta}\|_\Sigma^2 &\stackrel{d}{=} \left\| \mu_d + \frac{1}{\sqrt{n_0}} z_0 - \frac{1}{\sqrt{n_1}} z_1 \right\|_2^2 \times \frac{(n-2)^3}{p^3} \frac{1}{n-2} \text{tr} \left(\left(\frac{Z^\top Z}{n-2} \right)^\dagger \right)^2 \\ &\xrightarrow{\text{a.s.}} (\Delta^2 + \gamma_0 + \gamma_1) \times \frac{1}{(\gamma-1)^3}. \end{aligned} \quad (27)$$

(28)

Combining (26) and (27), as well as putting the threshold term $\ln(n_1/n_0)$ back, we obtain

$$q_0 \xrightarrow{\text{a.s.}} \frac{-\frac{\Delta^2 + \gamma_1 - \gamma_0}{2\gamma(\gamma-1)} + \ln \frac{\gamma_0}{\gamma_1}}{\sqrt{\Delta^2 + \gamma_1 + \gamma_0} \frac{1}{(\gamma-1)^{3/2}}},$$

The same argument of analyzing q_0 applies to q_1 , and therefore, we have

$$q_1 \xrightarrow{\text{a.s.}} \frac{-\frac{\Delta^2 + \gamma_0 - \gamma_1}{2\gamma(\gamma-1)} + \ln \frac{\gamma_1}{\gamma_0}}{\sqrt{\Delta^2 + \gamma_1 + \gamma_0} \frac{1}{(\gamma-1)^{3/2}}}.$$

The misclassification error in the case of $\gamma > 1$ follows by substituting q_0, q_1 above into (13). The proof is complete. \square

7 Conclusion and Discussion

This paper provides a theoretical analysis on the performance of LDA under label shift, in both the under- and over-parametrized regime. We explicitly quantify the misclassification error in the proportional limit of $n \rightarrow \infty$ and $p/n_\ell \rightarrow \gamma_\ell$ for $\ell = 0, 1$, where $\gamma_\ell > 0$ is a constant. Our theory shows a peaking phenomenon when the sample size is close to the data dimension. We demonstrate a *phase transition* phenomenon about data imbalance: The misclassification error exhibits different behaviors as the two-class ratio n_1/n_0 varies, depending on the value of γ_0 . We clearly characterize the three behaviors of the misclassification error in the underparametrized, lightly overparametrized, and overparametrized

regions depending on γ_0 . We also investigate the regularized LDA, and show that the peaking and phase transition phenomenons disappear when the regularization becomes strong.

Additionally, in the highly overparametrized regime, the misclassification error under label shift has multiple maxima, as shown in Figure 6. We believe that the multi-peaking behavior of the misclassification error reflects the intertwined interaction between the lack of information in data and the model mismatch under label shift.

A Proof of Phase Transition in Section 3.2

Misclassification error as $n_1/n_0 \rightarrow \infty$ In this section, we will give a proof to show the misclassification error tends to 0.5 when $n_1/n_0 \rightarrow \infty$. When the training data set is extremely imbalanced with $n_1/n_0 \rightarrow \infty$, i.e., $\gamma_0/\gamma_1 \rightarrow \infty$, the Bayes classifier tends to classify all data points to class 1. This leads to the following limits,

$$\frac{\Delta^2 + \gamma_0 - \gamma_1}{\sqrt{\Delta^2 + \gamma_0 + \gamma_1}} \rightarrow +\infty, \quad \frac{\Delta^2 + \gamma_1 - \gamma_0}{\sqrt{\Delta^2 + \gamma_0 + \gamma_1}} \rightarrow -\infty.$$

Then the limit of misclassification error is given by

$$\mathcal{R}(f_{\hat{\alpha}, \hat{\beta}}^b) \xrightarrow{\text{a.s.}} \frac{1}{2}\Phi(-\infty) + \frac{1}{2}\Phi(+\infty) = 0.5.$$

The proof is complete.

Phase transition knots In this section, we provide theoretical justifications of the phase transition knots given in section 3.2. Denote the asymptotic misclassification error in (6) and (7) as $\mathcal{R}(\gamma_0, \gamma_1)$. We fix γ_0 and let γ_1 vary starting from the balanced case with $\gamma_1 = \gamma_0$.

The transition knots are obtained by a local analysis about the instantaneous change of misclassification error as n_1/n_0 slightly increases from 1. We observe that, as n_1/n_0 slightly increases from 1, the misclassification error decreases in Phase I and III, and increases in Phase II. Notice that γ_1 slightly decreases from γ_0 as n_1/n_0 slightly increases from 1.

The instantaneous change of $\mathcal{R}(\gamma_0, \gamma_1)$ with respect to γ_1 can be characterized by the following partial derivative, $\frac{\partial}{\partial \gamma_1} \mathcal{R}(\gamma_0, \gamma_1) \big|_{\gamma_1=\gamma_0}$:

$$\frac{\partial}{\partial \gamma_1} \mathcal{R} \big|_{\gamma_1=\gamma_0} = \frac{\Delta^2 \phi \left(\frac{-\Delta^2(1-\frac{1}{2}\gamma_0)^{\frac{1}{2}}}{2\sqrt{\Delta^2+2\gamma_0}} \right)}{16(\Delta^2+2\gamma_0)^{\frac{3}{2}}(1-\frac{1}{2}\gamma_0)^{\frac{1}{2}}} \times [4 + \Delta^2],$$

for $\gamma_0 < 2$, and

$$\frac{\partial}{\partial \gamma_1} \mathcal{R} \big|_{\gamma_1=\gamma_0} = \frac{\Delta^2 \phi \left(\frac{-\Delta^2(\frac{1}{2}\gamma_0-1)^{\frac{1}{2}}}{\gamma_0\sqrt{\Delta^2+2\gamma_0}} \right)}{4\gamma_0^2(\Delta^2+2\gamma_0)^{\frac{3}{2}}(\frac{1}{2}\gamma_0-1)^{\frac{1}{2}}} \times \underbrace{[4\gamma_0^2 - (12 - \Delta^2)\gamma_0 - 4\Delta^2]}_{Q(\gamma_0, \Delta)},$$

for $\gamma_0 > 2$, where ϕ is the probability density function of the standard normal distribution. The $Q(\gamma_0, \Delta)$ term is a quadratic function with two roots of opposite signs. The positive root is $\gamma_b = \frac{1}{8}(12 - \Delta^2 + \sqrt{\Delta^4 + 40\Delta^2 + 144})$. The sign of the above partial derivative has the following cases:

- When $\gamma_0 \in (0, 2)$, $\frac{\partial}{\partial \gamma_1} \mathcal{R}(\gamma_0, \gamma_1) \big|_{\gamma_1=\gamma_0}$ is always positive. As a result, $\mathcal{R}(\gamma_0, \gamma_1)$ decreases as γ_1 decreases from γ_0 , which corresponds to Phase I.
- When $\gamma_0 \in (2, \gamma_b)$, $\frac{\partial}{\partial \gamma_1} \mathcal{R}(\gamma_0, \gamma_1) \big|_{\gamma_1=\gamma_0}$ is negative. In this case $\mathcal{R}(\gamma_0, \gamma_1)$ increases as γ_1 decreases from γ_0 , which corresponds to Phase II.
- When $\gamma_0 \in (\gamma_b, +\infty)$, $\frac{\partial}{\partial \gamma_1} \mathcal{R}(\gamma_0, \gamma_1) \big|_{\gamma_1=\gamma_0}$ is positive. In this case $\mathcal{R}(\gamma_0, \gamma_1)$ decreases as γ_1 decreases from γ_0 , which includes Phase III.

B Proofs of Lemmas in Section 6.1

B.1 Proofs of Lemma 4 and Lemma 5

Proof of lemma 4. We will prove this result by Basu theorem, i.e., we will show that $(\hat{\mu} - \mu)$ is a complete and sufficient statistic and $\hat{\Sigma}$ is an auxiliary statistic w.r.t. μ .

First, we show $\hat{\mu}$ is a complete statistic. We need to check that for any μ and measurable function g , $\mathbb{E}[g(\hat{\mu})] = 0$ for any μ implies $\mathbb{P}(g(\hat{\mu}) = 0) = 1$ for any μ . Indeed, for any measurable function g such that the expectation of $g(\hat{\mu})$ over sample space (x_1, x_2, \dots, x_n) is zero, i.e.,

$$\mathbb{E}[g(\hat{\mu})] = 0 \quad \text{for any } \mu, \quad (29)$$

we can derive $\mathbb{P}(g(\hat{\mu} - \mu) = 0) = 1$ by taking derivatives of Equation (29) w.r.t. μ recursively,

$$\mathbb{E}\left[h(\hat{\mu})g(\hat{\mu})\right] = 0, \quad \text{for any polynomial } h,$$

and therefore $\hat{\mu}$ is a complete statistic w.r.t. parameter μ .

To prove $\hat{\mu}$ is also a sufficient statistic for μ , we need to show that given the statistic $\hat{\mu}$ the conditional distribution of x_1, \dots, x_n does not depend on μ . Note that $\hat{\mu}$ has a multivariate normal distribution, i.e., $\hat{\mu} \sim \mathcal{N}(\mu, \frac{1}{n}\Sigma)$, since $\hat{\mu} = \frac{1}{n}\sum_{i=1}^n x_i$ is a linear combination of i.i.d. multivariate normal vectors x_1, x_2, \dots, x_n . The pdf of $\hat{\mu}$ and the joint distribution of x_1, x_2, \dots, x_n are given by

$$\begin{aligned} f(\hat{\mu}) &= \frac{1}{(2\pi)^{\frac{p}{2}} |\frac{1}{n}\Sigma|^{\frac{1}{2}}} \exp\left(-\frac{n}{2}(\hat{\mu} - \mu)^\top \Sigma^{-1}(\hat{\mu} - \mu)\right), \\ f(x_1, \dots, x_n) &= \frac{1}{(2\pi)^{\frac{np}{2}} |\Sigma|^{\frac{n}{2}}} \exp\left(-\sum_{i=1}^n \frac{1}{2}(x_i - \mu)^\top \Sigma^{-1}(x_i - \mu)\right). \end{aligned} \quad (30)$$

The joint density function of x_1, \dots, x_n and $\hat{\mu}$ is given by

$$f(x_1, \dots, x_n, \hat{\mu}) = f(x_1, \dots, x_n) \mathbb{1}\left(\hat{\mu} = \frac{1}{n}(x_1 + x_2 + \dots + x_n)\right). \quad (31)$$

By taking the fraction of (30) and (31), the conditional density of x_1, \dots, x_n given $\hat{\mu}$ is

$$f(x_1, \dots, x_n | \hat{\mu}) = C \exp\left(-\frac{1}{2}(x - \hat{\mu})^\top \Sigma^{-1}(x - \hat{\mu})\right), \quad (32)$$

where C is a constant. By Fisher-Neyman factorization theorem [43], given the statistic $\hat{\mu}$ the conditional distribution of x_1, \dots, x_n does not depend on μ and therefore $\hat{\mu}$ is a sufficient statistic for μ .

Sample covariance has a distribution which doesn't depend on the parameter μ .

$$\hat{\Sigma} = \sum_{i=1}^n (x_i - \hat{\mu})^\top (x_i - \hat{\mu}) = \sum_{i=1}^n z_i^\top z_i, \quad (33)$$

and therefore it is a auxiliary statistic.

Combining $\hat{\mu}$ being a complete and sufficient statistic and $\hat{\Sigma}$ being an auxiliary statistic, we obtain that $\hat{\mu}$ and $\hat{\Sigma}$ are independent, by Basu Theorem. \square

Proof of lemma 5. The following isotropic property of Wishart distribution has been given by Wang and Jiang [69]. For any orthogonal matrix $U \in \mathbb{R}^{p \times p}$, we have

$$U^\top \left(\frac{Z^\top Z}{n-2} \right) U \sim \mathcal{W}(I_p, n-2). \quad (34)$$

We next apply this property to the left-hand side of equation (9),

$$\begin{aligned} z^\top \left(\frac{Z^\top Z}{n-2} \right)^\dagger z &= z^\top U_i U_i^\top \left(\frac{Z^\top Z}{n-2} \right)^\dagger U_i U_i^\top z \\ &= \|z\|^2 e_i^\top \left(U_i^\top \left(\frac{Z^\top Z}{n-2} \right) U_i \right)^\dagger e_i \\ &\stackrel{d}{=} \|z\|^2 e_i^\top \left(\frac{Z^\top Z}{n-2} \right)^\dagger e_i, \end{aligned} \quad (35)$$

where U_i is a orthogonal matrix that transforms the vector z to canonical basis vector e_i , i.e.,

$$U_i^\top z = \|z\|_2 e_i, \quad i = 1, 2, \dots, p. \quad (36)$$

We can further simplify the product $z^\top (\frac{Z^\top Z}{n-2})^\dagger z$ by taking an average over the index i ,

$$\begin{aligned} z^\top \left(\frac{Z^\top Z}{n-2} \right)^\dagger z &\stackrel{d}{=} \frac{1}{p} \|z\|_2^2 \sum_{i=1, \dots, p} e_i^\top \left(\frac{Z^\top Z}{n-2} \right)^\dagger e_i \\ &= \frac{1}{p} \|z\|_2^2 \text{tr} \left(\left(\frac{Z^\top Z}{n-2} \right)^\dagger \right), \end{aligned} \quad (37)$$

where we get the isotropicity of Wishart distribution in Equation (9). \square

B.2 Proofs of Lemma 6, Lemma 7 and Lemma 8

Proof of Lemma 6. We use the eigenvalue decomposition of $Z^\top Z = U^\top D U$ to simplify the left-hand side of Equation (10),

$$\begin{aligned} \text{tr}[(Z^\top Z)^\dagger] &= \text{tr}(UD^\dagger U^\top) \\ &= \text{tr}(D^\dagger) \\ &= \sum_{s \in \lambda(Z^\top Z), s \neq 0} \frac{1}{s}. \end{aligned}$$

The result above implies that the trace of the pseudo-inverse of $Z^\top Z$ is equal to the sum of the reciprocal of its eigenvalues. By the same arguments on ZZ^\top , we can show that

$$\text{tr}[(ZZ^\top)^\dagger] = \sum_{s \in \lambda(ZZ^\top), s \neq 0} \frac{1}{s}.$$

Then we deduce the desired result by the fact that the set of non-zero eigenvalues of ZZ^\top matches that of $Z^\top Z$. \square

Proof of Lemma 7. We first compute the limit of $\frac{1}{\sqrt{p}} \mu_d^\top z$. The linear combination of multivariate normal random vector $z \sim \mathcal{N}(0, I_p)$ is a normal random variable, namely, $\frac{1}{\sqrt{p}} \mu_d^\top z \sim \mathcal{N}(0, \frac{1}{p} \mu_d^\top \mu_d)$. From the concentration inequality of the normal random variable [57], we have

$$\mathbb{P} \left(\left| \frac{1}{\sqrt{p}} \mu_d^\top z \right| \geq \frac{\epsilon}{\sqrt{p}} \|\mu_d\|_2 \right) \leq 2 \frac{1}{\sqrt{2\pi}\epsilon} e^{-\epsilon^2/2} \quad \text{for all } x \geq 0. \quad (38)$$

Combining Equation (38), $e^{-x} < \frac{1}{x}$ for $x > 0$ and $\|\mu_d\|_2 < 2\Delta$ for a sufficiently large p , the sum of the probabilities of $\frac{1}{\sqrt{p}} |\mu_d^\top z| > \epsilon$ is finite, for any positive $\epsilon > 0$, i.e.,

$$\begin{aligned} \sum_{p=1}^{\infty} \mathbb{P} \left(\left| \frac{1}{\sqrt{p}} \mu_d^\top z \right| > \epsilon \right) &\leq \sum_{p=1}^{\infty} \frac{\|\mu_d\|_2}{\sqrt{2\pi p}\epsilon} e^{-(\epsilon^2 p)/(2\|\mu_d\|_2^2)} \\ &< \sum_{p=1}^{\infty} \frac{2\|\mu_d\|_2^3}{\sqrt{2\pi}\epsilon^3} \frac{1}{p^{3/2}} \\ &< \infty. \end{aligned}$$

By the Borel-Cantelli lemma, we have

$$\frac{1}{\sqrt{p}} \mu_d^\top z \xrightarrow{\text{a.s.}} 0.$$

We next consider the limit of $\frac{1}{p} z_\ell^\top z_\ell$. Since z_ℓ satisfies the chi-squared distribution independently with expectation $\mathbb{E}[z_{\ell,i}^2] = 1$ and finite variance $\text{Var}(z_{\ell,i}^2) = 2$, we know the average of the squared elements in z_ℓ converges to the expectation almost surely by the strong law of the large numbers, namely,

$$\frac{1}{p} z_\ell^\top z_\ell \xrightarrow{\text{a.s.}} 1.$$

By the same arguments given above, n_ℓ satisfies the binomial distribution $B(n_0 + n_1, \pi_\ell)$ which is composed of $n_0 + n_1$ independent Bernoulli distribution with expectation π_0 and Variance $\pi_0\pi_1$. From the strong law of large numbers, we have

$$\frac{n_\ell}{n_0 + n_1} \xrightarrow{\text{a.s.}} \pi_\ell.$$

□

Proof of Lemma 8. We first derive the expression of $m(\zeta)$. The Marchenko-Pastur law is supported on a compact subset of \mathcal{R}^+ , i.e., $\text{supp}(F_\gamma) \subset [a, b]$ where

$$a = (1 - \sqrt{\gamma})^2, \quad \text{and} \quad b = (1 + \sqrt{\gamma})^2.$$

Let $\{z_k\}$ be a sequence of complex numbers such that $\text{Im}(z_k) > 0, \text{Re}(z_k) = \zeta$ for any k and $\lim_{k \rightarrow \infty} z_k = \zeta$. Consider the sequence of integral

$$\int_a^b \frac{1}{s - z_k} dF_\gamma(s).$$

For any k , $0 < \gamma < 1$ and $s > a$, we have

$$\left| \frac{1}{s - z_k} \right| \leq \left| \frac{1}{s} \right| < \frac{1}{(1 - \sqrt{\gamma})^2} < \infty.$$

By the dominated convergence theorem, we have

$$\begin{aligned} \int \frac{1}{s - \zeta} dF_\gamma(s) &= \int_a^b \lim_{k \rightarrow \infty} \frac{1}{s - z_k} dF_\gamma(s) = \lim_{k \rightarrow \infty} \int_a^b \frac{1}{s - z_k} dF_\gamma(s) \\ &= \lim_{k \rightarrow \infty} \int \frac{1}{s - z_k} dF_\gamma(s). \end{aligned} \tag{39}$$

To compute $\int \frac{1}{s - z_k} dF_\gamma(s)$, Bai and Silverstein [2, Lemma 3.11] gives

$$\int \frac{1}{s - z_k} dF_\gamma(s) = \frac{1 - \gamma - z_k + \sqrt{(z_k - \gamma - 1)^2 - 4\gamma}}{2\gamma z_k}. \tag{40}$$

According to the definition of the square root of complex numbers in Bai and Silverstein [2, Equation (2.3.2)], the real part of $\sqrt{(z_k - \gamma - 1)^2 - 4\gamma}$ has the same sign as that of $z_k - \gamma - 1$. Since $\text{Re}(z_k) = \zeta \leq 0, \gamma > 0$, the real part of $\sqrt{(z_k - \gamma - 1)^2 - 4\gamma}$ is negative and gives

$$\lim_{k \rightarrow \infty} \sqrt{(z_k - \gamma - 1)^2 - 4\gamma} = -\sqrt{(\zeta - \gamma - 1)^2 - 4\gamma}. \tag{41}$$

Substituting (41) and (40) into (39) gives rise to (11).

We then compute $m(0)$. When substituting $\zeta = 0$ into (11), both the numerator and the denominator are 0. Here we apply L'Hospital's rule:

$$\begin{aligned} m(0) &= \lim_{\zeta \rightarrow 0} \frac{1 - \gamma - \zeta - \sqrt{(\zeta - \gamma - 1)^2 - 4\gamma}}{2\gamma\zeta} \\ &= \lim_{\zeta \rightarrow 0} \frac{1}{2\gamma} \left(-1 - \frac{\zeta - \gamma - 1}{\sqrt{(\zeta - \gamma - 1)^2 - 4\gamma}} \right) \\ &= \frac{1}{2\gamma} \left(-1 - \frac{-\gamma - 1}{\sqrt{(-\gamma - 1)^2 - 4\gamma}} \right) \\ &= \frac{1}{2\gamma} \left(-1 + \frac{1 + \gamma}{1 - \gamma} \right) \\ &= \frac{1}{1 - \gamma}. \end{aligned}$$

We next derive the expression of $\frac{d}{d\zeta}m(\zeta)$. To derive the expression, we first show that

$$\int \frac{1}{(s-\zeta)^2} dF_\gamma(s) = \lim_{z \rightarrow \zeta} \frac{d}{dz} \int \frac{1}{s-z} dF_\gamma(s) \quad \text{for } z \in \mathbb{C} \text{ with } \operatorname{Re}(z) = \zeta, \operatorname{Im}(z) > 0. \quad (42)$$

Let $\{h_k\}$ be a set of complex numbers such that $|\operatorname{Re}(h_k)| \leq |\zeta|/2$ for any k and $\lim_{k \rightarrow \infty} h_k = 0$. For any k and $s \geq a$, we have

$$\left| \frac{1}{(s-z-h_k)(s-z)} \right| \leq \frac{1}{(1-\sqrt{\gamma})^4} < \infty.$$

By the dominated convergence theorem, we have

$$\begin{aligned} \frac{\partial}{\partial z} \int \frac{1}{s-z} dF_\gamma(s) &= \lim_{k \rightarrow \infty} \frac{1}{h_k} \left[\int \frac{1}{s-(z+h_k)} dF_\gamma(s) - \int \frac{1}{s-z} dF_\gamma(s) \right] \\ &= \lim_{k \rightarrow \infty} \frac{1}{h_k} \int_a^b \frac{1}{s-z-h_k} - \frac{1}{s-z} dF_\gamma(s) \\ &= \lim_{k \rightarrow \infty} \int_a^b \frac{1}{(s-z-h_k)(s-z)} dF_\gamma(s) \\ &= \int_a^b \frac{1}{(s-z)^2} dF_\gamma(s), \end{aligned}$$

where the second equality holds since $F_\gamma(s)$ is supported on $[a, b]$.

Since

$$\left| \frac{1}{(s-z)^2} \right| \leq \frac{1}{(1-\sqrt{\gamma})^4} < \infty,$$

for any $s \geq a$, we have

$$\begin{aligned} \int \frac{1}{(s-\zeta)^2} dF_\gamma(s) &= \int_a^b \lim_{z \rightarrow \zeta} \frac{1}{(s-z)^2} dF_\gamma(s) = \lim_{z \rightarrow \zeta} \int_a^b \frac{1}{(s-z)^2} dF_\gamma(s) \\ &= \lim_{z \rightarrow \zeta} \frac{\partial}{\partial z} \int \frac{1}{s-z} dF_\gamma(s). \end{aligned}$$

Using (40), we have

$$\begin{aligned} &\frac{\partial}{\partial z} \int \frac{1}{s-z} dF_\gamma(s) \\ &= \frac{(2\gamma z) \left[-1 + \frac{z-\gamma-1}{\sqrt{(z-\gamma-1)^2-4\gamma}} \right] - (2\gamma) \left(1 - \gamma - z + \sqrt{(z-\gamma-1)^2 - 4\gamma} \right)}{4\gamma^2 z^2} \\ &= \frac{(2\gamma)(\gamma-1) + \frac{(2\gamma z)(z-\gamma-1)}{\sqrt{(z-\gamma-1)^2-4\gamma}} - 2\gamma\sqrt{(z-\gamma-1)^2 - 4\gamma}}{4\gamma^2 z^2}. \end{aligned} \quad (43)$$

Letting $z \rightarrow \zeta$ in (43) and recall that the real part of $\sqrt{(z-\gamma-1)^2 - 4\gamma}$ is negative, one gets (12).

To compute $\frac{d}{d\zeta}m(0)$, by L'Hopital's rule, we deduce

$$\frac{d}{d\zeta}m(0) = \lim_{\zeta \rightarrow 0} \frac{(2\gamma)(\gamma-1) - \frac{(2\gamma z)(\zeta-\gamma-1)}{\sqrt{(\zeta-\gamma-1)^2-4\gamma}} + 2\gamma\sqrt{(\zeta-\gamma-1)^2 - 4\gamma}}{4\gamma^2 \zeta^2} = \frac{1}{(1-\gamma)^3}. \quad (44)$$

□

C Proofs in Section 4

C.1 Proof of Theorem 3

Proof of Theorem 3. The proof uses the same technique as in the Theorem 2, the misclassification error is the same as (13), and we only need to show the limits of q_0 and q_1 . By the change of variables formula

in (14) and Lemma 8, we deduce

$$\begin{aligned}
& \widehat{\beta}^\top (\widehat{\alpha} - \mu_0) \\
&= -\frac{1}{2} (\mu_d + \frac{1}{\sqrt{n_0}} z_0 - \frac{1}{\sqrt{n_1}} z_1)^\top \left(\frac{Z^\top Z}{n-2} + \lambda I_p \right)^\dagger (\mu_d - \frac{1}{\sqrt{n_0}} z_0 - \frac{1}{\sqrt{n_1}} z_1) \\
&= \frac{1}{2n_0} z_0^\top \left(\frac{Z^\top Z}{n-2} + \lambda I_p \right)^\dagger z_0 - \frac{1}{2} (\mu_d - z_1/\sqrt{n_1})^\top \left(\frac{Z^\top Z}{n-2} + \lambda I_p \right)^\dagger (\mu_d - z_1/\sqrt{n_1}) \\
&\stackrel{d}{=} \left[-\frac{1}{2} \|\mu_d - \frac{1}{\sqrt{n_1}} z_1\|_2^2 + \frac{1}{2n_0} \|z_0\|_2^2 \right] \times \frac{1}{p} \text{tr} \left(\frac{Z^\top Z}{n-2} + \lambda I_p \right)^\dagger \\
&\xrightarrow{\text{a.s.}} -\frac{1}{2} (\Delta^2 + \gamma_0 - \gamma_1) m(-\lambda).
\end{aligned} \tag{45}$$

Similarly we can derive

$$\begin{aligned}
\|\widehat{\beta}\|_\Sigma^2 &= (\widehat{\mu}_0 - \widehat{\mu}_1)^\top (\widehat{\Sigma} + \lambda I_p)^\dagger \Sigma (\widehat{\Sigma} + \lambda I_p)^\dagger (\widehat{\mu}_0 - \widehat{\mu}_1) \\
&= (\mu_d + \frac{1}{\sqrt{n_0}} z_0 - \frac{1}{\sqrt{n_1}} z_1)^\top \left(\left(\frac{Z^\top Z}{n-2} + \lambda I_p \right)^\dagger \right)^2 (\mu_d + \frac{1}{\sqrt{n_0}} z_0 - \frac{1}{\sqrt{n_1}} z_1) \\
&\stackrel{d}{=} \left\| \mu_d + \frac{1}{\sqrt{n_0}} z_0 - \frac{1}{\sqrt{n_1}} z_1 \right\|_2^2 \times \frac{1}{p} \text{tr} \left(\left(\frac{Z^\top Z}{n-2} + \lambda I_p \right)^\dagger \right)^2 \\
&\xrightarrow{\text{a.s.}} (\Delta^2 + \gamma_0 + \gamma_1) m'(-\lambda).
\end{aligned} \tag{46}$$

Combining (45) and (46), as well as putting the threshold term $\ln(n_1/n_0)$ back, we obtain

$$q_0 \xrightarrow{\text{a.s.}} \frac{-\frac{1}{2} (\Delta^2 - \gamma_0 + \gamma_1) m(-\lambda) + \ln \frac{\gamma_0}{\gamma_1}}{\sqrt{(\Delta^2 + \gamma_1 + \gamma_0) m'(-\lambda)}}.$$

The same argument of analyzing q_0 applies to q_1 and therefore, we have

$$q_1 \xrightarrow{\text{a.s.}} \frac{-\frac{1}{2} (\Delta^2 + \gamma_0 - \gamma_1) m(-\lambda) + \ln \frac{\gamma_1}{\gamma_0}}{\sqrt{(\Delta^2 + \gamma_1 + \gamma_0) m'(-\lambda)}}.$$

We complete the proof by substituting q_0, q_1 above into (13). \square

C.2 Proof of Regularized Phase Transition in Section 4.2

In this section, we show with a strong regularization, the phase transition phenomenon will vanish. Denote the asymptotic misclassification error in Theorem 3 as

$$\mathcal{R}_\lambda(\gamma_0, \gamma_1) = \sum_{\ell=0,1} \Phi \left(\frac{g(\gamma_0, \gamma_1, \ell) m(-\lambda) + (-1)^\ell \ln \frac{\gamma_0}{\gamma_1}}{k(\gamma_0, \gamma_1) \sqrt{m'(-\lambda)}} \right),$$

and we use the shorthand $\mathcal{R}_\lambda(\gamma)$ to denote $\mathcal{R}_\lambda(\gamma_0, \gamma_1)$ with the balanced data, i.e., $\gamma_0 = \gamma_1 = 2\gamma$,

$$\mathcal{R}_\lambda(\gamma) := \mathcal{R}_\lambda(2\gamma, 2\gamma) = \Phi \left(\frac{-\Delta^2 m(-\lambda)}{2\sqrt{(\Delta^2 + 4\gamma) m'(-\lambda)}} \right).$$

We show the phase transition phenomenon vanishes with a strong regularization, namely,

$$\frac{\partial}{\partial \gamma_1} \mathcal{R}_\lambda(\gamma_0, \gamma_1) \Big|_{\gamma_0=\gamma_1=2\gamma} > 0 \quad \text{for a strong } \lambda > 0.$$

To see the result above, we need to show that $\frac{\partial \mathcal{R}_\lambda(\gamma)}{\partial \gamma} > 0$ with a strong regularization. Specifically, invoking Chain rule and by some manipulation, we have

$$\frac{\partial}{\partial \gamma_1} \mathcal{R}_\lambda(\gamma_0, \gamma_1) \Big|_{\gamma_0=\gamma_1=2\gamma} = \frac{\partial \mathcal{R}_\lambda(\gamma)}{\partial \gamma} \frac{\partial \gamma}{\partial \gamma_1} \Big|_{\gamma_0=\gamma_1=2\gamma}.$$

By Mathematica Software [35], we check

$$\frac{\partial \mathcal{R}_\lambda(\gamma)}{\partial \gamma} = \phi \left(\frac{-\Delta^2 m(-\lambda)}{2\sqrt{(\Delta^2 + 4\gamma)m'(-\lambda)}} \right) \underbrace{\frac{\partial}{\partial \gamma} \left(\frac{-\Delta^2 m(-\lambda)}{2\sqrt{(\Delta^2 + 4\gamma)m'(-\lambda)}} \right)}_{(*)},$$

where ϕ is the pdf of the standard normal distribution. By Mathematica Software [35], the denominator of $(*)$ is also always positive, and is given by

$$\begin{aligned} & 2\sqrt{2}\gamma^3\lambda^3 \frac{[(\gamma + \lambda + 1)^2 - 4\gamma]^{3/2} (4\gamma + \Delta^2)^{3/2}}{[\gamma\lambda^2((\gamma + \lambda + 1)^2 - 4\gamma)]^{3/2}} \\ & \times \left[\gamma^3 + \gamma^2(\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} + 2\lambda - 3) \right. \\ & + \gamma \left(-2\sqrt{\gamma^2 + 2\gamma(\lambda - 1) + (\lambda + 1)^2} + \lambda\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} + \lambda^2 + 3 \right) \\ & \left. + (\lambda + 1)(\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} - \lambda - 1) \right]^{3/2} \end{aligned}$$

As a result, the sign of $\frac{\partial \mathcal{R}_\lambda(\gamma)}{\partial \gamma}$ is determined by the numerator of $(*)$ given as

$$\begin{aligned} & \Delta^2 \left\{ \gamma(\lambda + 1) \left[\Delta^2 \left(-3\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} - 2\lambda\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} \right. \right. \right. \\ & + \lambda^2 + 5\lambda + 4 \left. \right) + 8(\lambda + 1)^2 \left(\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} - \lambda - 1 \right) \left. \right] \\ & + \gamma^2 \left[\Delta^2 \left(2\lambda\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} + 3\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} \right. \right. \\ & - 6\lambda - 6 \left. \right) - 4(\lambda + 1) \left(2\lambda\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} - 9\lambda - 9 \right. \\ & \left. \left. + 7\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} \right) \right] + \gamma^3 \left[4\lambda \left(\sqrt{\gamma^2 + 2\gamma(\lambda - 1) + (\lambda + 1)^2} \right. \right. \\ & + \lambda\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} + \lambda^2 - 9 \left. \right) + 36\sqrt{\gamma^2 + 2\gamma(\lambda - 1) + (\lambda + 1)^2} \\ & + \Delta^2 \left(-\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} + \lambda + 4 \right) - 64 \right] + \gamma^4 \left[-\Delta^2 + 56 \right. \\ & - 20\sqrt{\gamma^2 + 2\gamma(\lambda - 1) + (\lambda + 1)^2} + 4\lambda \left(2\sqrt{\gamma^2 + 2\gamma(\lambda - 1) + (\lambda + 1)^2} + 3\lambda - 4 \right) \left. \right] \\ & + 4\gamma^5 \left[\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} + 3\lambda - 6 \right] + 4\gamma^6 \\ & \left. + \Delta^2(\lambda + 1)^3 \left(\sqrt{2(\gamma + 1)\lambda + (\gamma - 1)^2 + \lambda^2} - \lambda - 1 \right) \right]. \end{aligned}$$

Combining the denominator and numerator, $\frac{\partial}{\partial \gamma} \mathcal{R}_\lambda(\gamma)$ is positive only when one of the following case happens,

1. $\Delta > 0$ and $0 < \gamma \leq 1$ and $\lambda > 0$.
2. $\Delta > 0$ and $\gamma \geq \frac{\sqrt{\Delta^2 + 4}}{2} + 1$ and $\lambda > 0$.
3. $\Delta > 0$ and $1 < \gamma < \frac{\sqrt{\Delta^2 + 4}}{2} + 1$ and $\lambda \geq$ the smallest real root of

$$[\#1^4(32\gamma + 4\Delta^2) + \#1^3(96\gamma^2 + 8\gamma\Delta^2 + 128\gamma + 16\Delta^2) + \#1^2(96\gamma^3 + 112\gamma^2 + 16\gamma\Delta^2 + 192\gamma + \Delta^4 + 24\Delta^2) + \#1(-8\gamma^3\Delta^2 - 16\gamma^2\Delta^2 + 32\gamma^4 - 96\gamma^3 - 64\gamma^2 - 2\gamma\Delta^4 + 8\gamma\Delta^2 + 128\gamma + 2\Delta^4 + 16\Delta^2) - 4\gamma^4\Delta^2 + 16\gamma^3\Delta^2 + \gamma^2\Delta^4 - 16\gamma^2\Delta^2 - 16\gamma^4 + 64\gamma^3 - 80\gamma^2 - 2\gamma\Delta^4 + 32\gamma + \Delta^4 + 4\Delta^2].$$

Consequently, We deduce that the misclassification error increases when γ grows in the interval $(0, 1)$ or $(\frac{\sqrt{\Delta^2 + 4}}{2}, \infty)$; when γ grows in $(1, \frac{\sqrt{\Delta^2 + 4}}{2})$, the misclassification error decreases when λ is small, yet increases when λ is large. For example, when $\Delta^2 = 9$ and $\lambda = 1$, $\mathcal{R}_\lambda(\gamma)$ increases monotonically with respect to γ , and the peaking phenomenon disappears. Meanwhile we have the instantaneous derivative $\frac{\partial \mathcal{R}_\lambda(\gamma_0, \gamma_1)}{\partial \gamma_1} |_{\gamma_0=\gamma_1=2\gamma} > 0$ for any γ_0 , which implies that the phase transition phenomenon vanishes.

The commands of Mathematica are provided as follows.

```

In[1]: m[ $\gamma$ _,  $\lambda$ _]:=  $\frac{\sqrt{(\gamma+\lambda+1)^2-4\gamma}+\gamma-\lambda-1}{2\gamma\lambda}$ 
In[2]: R[ $\gamma$ _,  $\lambda$ _,  $\Delta$ _]:= -  $\frac{\Delta^2 m(\gamma,\lambda)}{2\sqrt{-(4\gamma+\Delta^2)} \frac{\partial m(\gamma,\lambda)}{\partial \lambda}}$ 
In[3]: de[ $\gamma$ _,  $\lambda$ _,  $\Delta$ _]:= Evaluate[Denominator[FullSimplify[Together[ $\frac{\partial R(\gamma,\lambda,\Delta)}{\partial \gamma}$ ]]]]
In[4]: nu[ $\gamma$ _,  $\lambda$ _,  $\Delta$ _]:= Evaluate[Numerator[FullSimplify[Together[ $\frac{\partial R(\gamma,\lambda,\Delta)}{\partial \gamma}$ ]]]]
In[5]: Reduce[de( $\gamma$ ,  $\lambda$ ,  $\Delta$ )  $\geq$  0  $\wedge$   $\lambda$   $>$  0  $\wedge$   $\gamma$   $>$  0  $\wedge$   $\Delta$   $>$  0, { $\gamma$ ,  $\lambda$ }]
In[6]: Reduce[nu( $\gamma$ ,  $\lambda$ ,  $\Delta$ )  $\geq$  0  $\wedge$   $\lambda$   $>$  0  $\wedge$   $\gamma$   $>$  0  $\wedge$   $\Delta$   $>$  0, { $\gamma$ ,  $\lambda$ }]
In[7]: Reduce[nu( $\gamma$ , 1, 3)  $\geq$  0  $\wedge$   $\gamma$   $>$  0, { $\gamma$ }]

```

References

- [1] Theodore Wilbur Anderson. An introduction to multivariate statistical analysis. Technical report, Wiley New York, 1962.
- [2] Z. Bai and Jack Silverstein. Spectral analysis of large dimensional random matrices, 01 2010.
- [3] Z. D. Bai and Y. Q. Yin. Limit of the smallest eigenvalue of a large dimensional sample covariance matrix. *Ann. Probab.*, 21(3):1275–1294, 07 1993. doi:10.1214/aop/1176989118.
- [4] Ricardo Barandela, E Rangel, José Salvador Sánchez, and Francesc J Ferri. Restricted decontamination for the imbalanced training sample problem. In *Iberoamerican congress on pattern recognition*, pages 424–431. Springer, 2003.
- [5] Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- [6] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- [7] Mikhail Belkin, Daniel Hsu, and Ji Xu. Two models of double descent for weak features. *SIAM Journal on Mathematics of Data Science*, 2(4):1167–1180, 2020.
- [8] Koby Bibas, Yaniv Fogel, and Meir Feder. A new look at an old problem: A universal learning approach to linear regression. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pages 2304–2308. IEEE, 2019.
- [9] Peter J Bickel and Elizaveta Levina. Some theory for fisher’s linear discriminant function, naive bayes’, and some alternatives when there are many more variables than observations. *Bernoulli*, 10(6):989–1010, 2004.
- [10] Christopher M. Bishop. Neural networks for pattern recognition. 1995.
- [11] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [12] Paula Branco, Luis Torgo, and Rita Ribeiro. A survey of predictive modelling under imbalanced distributions, 2015.
- [13] Tony Cai and Weidong Liu. A direct estimation approach to sparse linear discriminant analysis. *Journal of the American statistical association*, 106(496):1566–1577, 2011.
- [14] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. 2019. arXiv:1906.07413 [cs.LG].
- [15] Yee Seng Chan and Hwee Tou Ng. Word sense disambiguation with distribution estimation. page 1010–1015, 2005.
- [16] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.

[17] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9268–9277, 2019.

[18] Zeyu Deng, Abla Kammoun, and Christos Thrampoulidis. A model of double descent for high-dimensional binary linear classification. 11 2019.

[19] Michał Dereziński, Feynman Liang, and Michael W Mahoney. Exact expressions for double descent and implicit regularization via surrogate random design. *arXiv preprint arXiv:1912.04533*, 2019.

[20] Robert PW Duin. Small sample size generalization. In *Proceedings of the Scandinavian Conference on Image Analysis*, volume 2, pages 957–964. PROCEEDINGS PUBLISHED BY VARIOUS PUBLISHERS, 1995.

[21] Charles Elkan. The foundations of cost-sensitive learning. page 973–978, 2001.

[22] Khalil Elkhailil, Abla Kammoun, Romain Couillet, Tareq Y Al-Naffouri, and Mohamed-Slim Alouini. A large dimensional study of regularized discriminant analysis. *IEEE Transactions on Signal Processing*, 68:2464–2479, 2020.

[23] Ronald A Fisher. The use of multiple measurements in taxonomic problems. *Annals of eugenics*, 7 (2):179–188, 1936.

[24] Jerome H Friedman. Regularized discriminant analysis. *Journal of the American statistical association*, 84(405):165–175, 1989.

[25] Keinosuke Fukunaga. *Introduction to statistical pattern recognition*. Elsevier, 2013.

[26] Jerzy W Grzymala-Busse, Linda K Goodwin, Witold J Grzymala-Busse, and Xinqun Zheng. An approach to imbalanced data sets based on changing rule strength. In *Rough-neural computing*, pages 543–553. Springer, 2004.

[27] Yaqian Guo, Trevor Hastie, and Robert Tibshirani. Regularized linear discriminant analysis and its application in microarrays. *Biostatistics*, 8(1):86–100, 2007.

[28] Guo Haixiang, Li Yijing, Jennifer Shang, Gu Mingyun, Huang Yuanyue, and Gong Bing. Learning from class-imbalanced data: Review of methods and applications. *Expert Systems with Applications*, 73:220–239, 2017.

[29] Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560*, 2019.

[30] Haibo He and Edwardo A Garcia. Learning from imbalanced data. *IEEE Transactions on knowledge and data engineering*, 21(9):1263–1284, 2009.

[31] Xia Hong, Sheng Chen, and Chris J Harris. A kernel-based two-class classifier for imbalanced data sets. *IEEE Transactions on neural networks*, 18(1):28–41, 2007.

[32] Jianping Hua, Zixiang Xiong, James Lowey, Edward Suh, and Edward R Dougherty. Optimal number of features as a function of sample size for various classification rules. *Bioinformatics*, 21 (8):1509–1515, 2005.

[33] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5375–5384, 2016.

[34] Gordon Hughes. On the mean accuracy of statistical pattern recognizers. *IEEE transactions on information theory*, 14(1):55–63, 1968.

[35] Wolfram Research, Inc. Mathematica, Version 12.2. URL <https://www.wolfram.com/mathematica>. Champaign, IL, 2020.

[36] Nathalie Japkowicz and Shaju Stephen. The class imbalance problem: A systematic study. *Intelligent data analysis*, 6(5):429–449, 2002.

[37] Taeho Jo and Nathalie Japkowicz. Class imbalances versus small disjuncts. *ACM Sigkdd Explorations Newsletter*, 6(1):40–49, 2004.

[38] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[39] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. 2009.

[40] Miroslav Kubat, Stan Matwin, et al. Addressing the curse of imbalanced training sets: one-sided selection. In *International Conference on Machine Learning*, volume 97, pages 179–186. Citeseer, 1997.

[41] Matjaz Kukar, Igor Kononenko, et al. Cost-sensitive learning with neural networks. In *ECAI*, volume 15, pages 88–94. Citeseer, 1998.

[42] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[43] Erich L Lehmann. *Fisher, Neyman, and the creation of classical statistics*. Springer Science & Business Media, 2011.

[44] Zachary Lipton, Yu-Xiang Wang, and Alexander Smola. Detecting and correcting for label shift with black box predictors. In *International Conference on Machine Learning*, pages 3122–3130. PMLR, 2018.

[45] Xu-Ying Liu, Jianxin Wu, and Zhi-Hua Zhou. Exploratory undersampling for class-imbalance learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(2):539–550, 2008.

[46] Yang Liu, Aijun An, and Xiangji Huang. Boosting prediction accuracy on imbalanced datasets with svm ensembles. In *Pacific-Asia conference on knowledge discovery and data mining*, pages 107–118. Springer, 2006.

[47] Brian Mac Namee, Padraig Cunningham, Stephen Byrne, and Owen I Corrigan. The problem of bias in training data in regression problems in medical decision support. *Artificial intelligence in medicine*, 24(1):51–70, 2002.

[48] Qing Mai, Hui Zou, and Ming Yuan. A direct approach to sparse discriminant analysis in ultra-high dimensions. *Biometrika*, 99(1):29–42, 2012.

[49] Marcus A Maloof. Learning when data sets are imbalanced and when costs are unequal and unknown. In *ICML-2003 workshop on learning from imbalanced data sets II*, volume 2, pages 2–1, 2003.

[50] V.A. Marčenko and Leonid Pastur. Distribution of eigenvalues for some sets of random matrices. *Math USSR Sb*, 1:457–483, 01 1967.

[51] Maciej A Mazurowski, Piotr A Habas, Jacek M Zurada, Joseph Y Lo, Jay A Baker, and Georgia D Tourassi. Training neural network classifiers for medical decision making: The effects of imbalanced datasets on classification performance. *Neural networks*, 21(2-3):427–436, 2008.

[52] Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and double descent curve. *arXiv preprint arXiv:1908.05355*, 2019.

[53] Andrea Montanari, Feng Ruan, Youngtak Sohn, and Jun Yan. The generalization error of max-margin linear classifiers: High-dimensional asymptotics in the overparametrized regime. 11 2019.

[54] Vidya Muthukumar, Kailas Vodrahalli, Vignesh Subramanian, and Anant Sahai. Harmless interpolation of noisy data in regression. *IEEE Journal on Selected Areas in Information Theory*, 1(1):67–83, 2020.

[55] Preetum Nakkiran. More data can hurt for linear regression: Sample-wise double descent. *arXiv preprint arXiv:1912.07242*, 2019.

[56] K Philip and SJS Chan. Toward scalable learning with non-uniform class and cost distributions: A case study in credit card fraud detection. In *Proceeding of the Fourth International Conference on Knowledge Discovery and Data Mining*, pages 164–168, 1998.

[57] Mark A. Pinsky. An introduction to probability theory and its applications, vol. 2 (william feller). *SIAM Rev.*, 14(4):662–663, October 1972. ISSN 0036-1445. doi:10.1137/1014119.

[58] Joaquin Quiñonero-Candela, Masashi Sugiyama, Anton Schwaighofer, and N. Lawrence. When training and test sets are different: Characterizing learning transfer. 2009.

[59] Predrag Radivojac, Nitesh V Chawla, A Keith Dunker, and Zoran Obradovic. Classification and knowledge discovery in protein databases. *Journal of Biomedical Informatics*, 37(4):224–239, 2004.

[60] Sarunas Raudys and Robert PW Duin. Expected classification error of the fisher linear classifier with pseudo-inverse covariance matrix. *Pattern recognition letters*, 19(5-6):385–392, 1998.

[61] Chris Seiffert, Taghi M Khoshgoftaar, Jason Van Hulse, and Amri Napolitano. Rusboost: A hybrid approach to alleviating class imbalance. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 40(1):185–197, 2009.

[62] Jun Shao, Yazhen Wang, Xinwei Deng, Sijian Wang, et al. Sparse linear discriminant analysis by thresholding for high dimensional data. *Annals of Statistics*, 39(2):1241–1265, 2011.

[63] Houssem Sifaou, Abla Kammoun, and Mohamed-Slim Alouini. High-dimensional linear discriminant analysis classifier for spiked covariance model. *Journal of Machine Learning Research*, 21, 2020.

[64] Chao Sima and Edward R Dougherty. The peaking phenomenon in the presence of feature-selection. *Pattern Recognition Letters*, 29(11):1667–1674, 2008.

[65] Amos J Storkey. When training and test sets are different: characterising learning transfer. pages 3–28, 2009.

[66] Yanmin Sun, Mohamed S Kamel, Andrew KC Wong, and Yang Wang. Cost-sensitive boosting for classification of imbalanced data. *Pattern Recognition*, 40(12):3358–3378, 2007.

[67] Santiago Velilla and Adolfo Hernández. On the consistency properties of linear and quadratic discriminant analyses. *Journal of multivariate analysis*, 96(2):219–236, 2005.

[68] Benjamin X Wang and Nathalie Japkowicz. Boosting support vector machines for imbalanced data sets. *Knowledge and information systems*, 25(1):1–20, 2010.

[69] Cheng Wang and Binyan Jiang. On the dimension effect of regularized linear discriminant analysis. *Electron. J. Statist.*, 12(2):2709–2742, 2018. doi:10.1214/18-EJS1469.

[70] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 7032–7042, 2017.

[71] Gary M Weiss. Mining with rarity: a unifying framework. *ACM Sigkdd Explorations Newsletter*, 6(1):7–19, 2004.

[72] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3485–3492. IEEE, 2010.

[73] Ji Xu and Daniel Hsu. On the number of variables to use in principal component regression. *arXiv preprint arXiv:1906.01139*, 2019.

[74] Eric Zhao, Anqi Liu, Animashree Anandkumar, and Yisong Yue. Active learning under label shift. 2021. [arXiv:2007.08479 \[cs.LG\]](https://arxiv.org/abs/2007.08479).

[75] Amin Zollanvari and Edward R Dougherty. Generalized consistent error estimator of linear discriminant analysis. *IEEE transactions on signal processing*, 63(11):2804–2814, 2015.

[76] Amin Zollanvari, Ulisses M Braga-Neto, and Edward R Dougherty. Analytic study of performance of error estimators for linear discriminant analysis. *IEEE Transactions on Signal Processing*, 59(9):4238–4255, 2011.