

---

# Minimax Gaussian Classification & Clustering

---

Tianyang Li<sup>1</sup>

Xinyang Yi<sup>1</sup>

Constantine Caramanis<sup>1</sup> Pradeep Ravikumar<sup>2</sup>

<sup>1</sup>University of Texas, Austin

<sup>2</sup>Carnegie Mellon University

## Abstract

We present minimax bounds for classification and clustering error in the setting where covariates are drawn from a mixture of two isotropic Gaussian distributions. Here, we define clustering error in a *discriminative* fashion, demonstrating fundamental connections between classification (supervised) and clustering (unsupervised). For both classification and clustering, our lower bounds show that without enough samples, the best any classifier or clustering rule can do is close to random guessing. For classification, as part of our upper bound analysis, we show that Fisher’s linear discriminant achieves a fast minimax rate  $\Theta(1/n)$  with enough samples  $n$ . For clustering, as part of our upper bound analysis, we show that a clustering rule constructed using principal component analysis achieves the minimax rate with enough samples. We also provide lower and upper bounds for the high-dimensional sparse setting where the dimensionality of the covariates  $p$  is potentially larger than the number of samples  $n$ , but where the difference between the Gaussian means is sparse.

## 1 Introduction

We consider lower bounds and upper bounds for Gaussian classification and clustering. We focus on the setting with two classes/clusters of isotropic (spherical) Gaussian distributions. Our results reveal the dependency of classification and clustering errors on the dimension of the covariates/features  $p$ , number of samples  $n$ , sparsity of the optimal classifier/clustering rule  $s = \|\Sigma^{-1}(\mu_1 - \mu_0)\|_0$ , and separation between two classes/clusters  $\rho = \sqrt{(\mu_1 - \mu_0)^T \Sigma^{-1}(\mu_1 - \mu_0)}$ .

---

Proceedings of the 20<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2017, Fort Lauderdale, Florida, USA. JMLR: W&CP volume 54. Copyright 2017 by the author(s).

Note that  $\rho$  is invariant under linear transforms of the features, and we are interested in the difficult case when  $\rho = O(1)$  is small.

Gaussian classification [11] is the problem of labeling a new sample using its features, given already observed samples with both the label and the features, where we have a prior distribution over the labels, and each class’s features are distributed according to a Gaussian distribution. The classification error of classifier is defined as the probability that the classifier mislabels a sample, which we also term as the risk of the classifier. When analyzing classification algorithms, we are interested in bounding the excess risk, which quantifies how much worse the trained classifier is compared to the optimal classifier – the Bayes classifier. In this paper we present a fast  $\Theta(1/n)$  minimax rate for Gaussian classification, even without the classes being completely separable (the Bayes error is not 0). For lower bounds, we also show that without enough samples, the performance of any classifier is close to random guessing. Random guessing uniformly randomly assigns a sample into one of the two classes, and its classification error is 1/2. For upper bounds, we show that Fisher’s linear discriminant [11] achieves the minimax rate for classification of two balanced isotropic Gaussian distributions. In the sparse setting, we use a  $\ell_1$  regularized optimization problem to estimate the Bayes classifier.

Gaussian clustering [11] is the problem of labeling a new sample using its features, given already observed *unlabeled* samples with only the features. The features are distributed according to a mixture of two Gaussian distributions. We define clustering error in a *discriminative* fashion, and in the spirit of classification error: we define it as the minimum classification error over the two mappings of the two clusters to the two labels. It is easy to see that the Bayes classifier is the optimal clustering rule, and the clustering error is equal to the Bayes classification error. We show that a clustering rule constructed using principal component analysis achieves a fast  $\Theta(1/n)$  minimax rate. In the sparse setting, we first use a thresholded estimate of the optimal clustering rule’s support, and then apply principal component analysis to estimate the optimal clustering rule.

We would like to note that when the separation between the two classes/clusters is constant, our results for classification and clustering are tight up to logarithmic factors in dimension  $p$ , number of samples  $n$ , and sparsity  $s$ . One of our key contributions is the derivation of classification and clustering minimax risk lower bounds. This presents particular technical challenges, since unlike metrics, risk lower bounds do not trivially satisfy a triangle inequality (which is critically used in typical lower bound derivations).

**Related Work** Until recently, results giving a  $\Theta(1/n)$  minimax rate relied critically on a strict separation assumption (the Bayes classifier’s classification error is 0) [15]. Without this assumption, the best known results were  $O(1/\sqrt{n})$  classification error rates [21, 6], or  $o(1/\sqrt{n})$  minimax rates [17, 14]. For certain nonparametric problems, matching upper and lower bounds [17] for classification excess risk are known. However, matching upper and lower bounds in parametric problems are not known, although Theorem 13.21 of [5] gives a  $o(1/\sqrt{n})$  fast rate upper bound for parametric classification problems under a VC dimension condition and Tsybakov’s low noise condition. In a recent result, [16] showed that it is possible to achieve a  $O(1/n)$  rate for excess risk in Gaussian classification without such a separation assumption. The question of a lower bound for this setting remained open. Resolving this and proving a matching lower (minimax) bound is one of the contributions of this paper.

For clustering, most previous works have focused on recovering each cluster [19, 10]. Despite other discriminative clustering methods without theoretical guarantees [22, 13], only [1] has formally defined clustering error. Our definition of clustering error draws fundamental connections between classification (supervised) and clustering (unsupervised), and in particular, is different from that of [1]. Using the notations of Section 2.2, for a clustering rule  $C$ , [1] defines the clustering error as  $\min\{\Pr[C \neq C^*], 1 - \Pr[C \neq C^*]\}$ , where  $C^*$  is the clustering rule using the Bayes classifier. In this paper, we define the clustering error as  $\min\{\Pr[C \neq Y], 1 - \Pr[C \neq Y]\}$ , where  $Y$  is the latent indicator in the Gaussian mixture model. In the former definition, the risk of an optimal classification rule is always zero, and hence does not reveal the inherent “hardness” of the problem, which (as is intuitive) is characterized by the distance between the cluster centers – a quantity we denote by  $\rho$  in the sequel. Conversely, defining the clustering error as we do, we obtain results that explicitly depend on  $\rho$ . Section 4.3 gives a more detailed comparison.

## 2 Problem Setup

### 2.1 Classification

The Gaussian classification problem is specified as follows. Suppose we use  $\mathbf{X} \in \mathbb{R}^p$  to denote the features, and use  $Y \in \{0, 1\}$  to denote the label. Then joint distribution over the features and the label is then specified by a simple Bernoulli distribution for the prior distribution over  $Y$ , and that the features conditioned on each of the two labels are distributed according to Gaussian distributions, with the same covariance matrix but different means. We thus have:

$$\begin{aligned} \mathbf{X} | Y &\sim \mathcal{N}(\boldsymbol{\mu}_Y, \boldsymbol{\Sigma}), \\ \Pr[Y = 1] &= \pi_1, \Pr[Y = 0] = \pi_0. \end{aligned} \quad (1)$$

A classifier is a function  $C : \mathbb{R}^p \rightarrow \{0, 1\}$ . A linear classifier with parameter  $\mathbf{w} \in \mathbb{R}^p$ ,  $b \in \mathbb{R}$  is defined as  $\mathbf{1}(\mathbf{w}^T \mathbf{x} + b > 0)$ .

The classification error of a classifier  $C$  is defined as the probability that the classifier mislabels a sample

$$\mathcal{R}(C) = \Pr[C(\mathbf{X}) \neq Y]. \quad (3)$$

For a linear classifier  $C(\mathbf{x}) = \mathbf{1}(\mathbf{w}^T \mathbf{x} + b > 0)$ , the classification error can be written as

$$1 - \pi_1 \Phi\left(\frac{\mathbf{w}^T \boldsymbol{\mu}_1 + b}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}}\right) - \pi_0 \Phi\left(-\frac{\mathbf{w}^T \boldsymbol{\mu}_0 + b}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}}\right). \quad (4)$$

The optimal classifier is the Bayes classifier  $C^*(\mathbf{x}) = \mathbf{1}\left(\frac{\pi_1 \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma})}{\pi_0 \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma})} > 1\right)$ , which is a linear classifier

$$\begin{aligned} C^*(\mathbf{x}) &= \mathbf{1}\left((\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}^{-1} \mathbf{x} \right. \\ &\quad \left. + m + \frac{1}{2}(-\boldsymbol{\mu}_1^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_1 + \boldsymbol{\mu}_0^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_0) + \log \frac{\pi_1}{\pi_0} > 0\right). \end{aligned} \quad (5)$$

A classifier is trained using  $n$  i.i.d. samples  $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)$ , where we have access to both the features and the label.

### 2.2 Clustering

The Gaussian clustering problem is specified as follows. Suppose we use  $\mathbf{X} \in \mathbb{R}^p$  to denote the features. No label is observed. The samples are distributed according to a Gaussian mixture model with two mixture components, each with the same covariance matrix but different means:

$$\mathbf{X} \sim \pi_1 \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}) + \pi_0 \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}). \quad (6)$$

A clustering rule is a function  $C : \mathbb{R}^p \rightarrow \{0, 1\}$ . A linear clustering rule with parameter  $\mathbf{w} \in \mathbb{R}^p$ ,  $b \in \mathbb{R}$  is defined as  $\mathbf{1}(\mathbf{w}^T \mathbf{x} + b > 0)$ . Note that, in the clustering

problem, we can view the classification label  $Y$  as a latent variable, and in particular, it is natural to use an error metric similar to classification error as clustering error. Here, we define the clustering error of a clustering rule  $C$  as

$$\mathcal{R}(C) = \min\{\Pr[C(\mathbf{X}) \neq Y], \Pr[C(\mathbf{X}) \neq 1 - Y]\}, \quad (7)$$

where a minimum is taken over two possible labelings because there is no access to the true label.

For a linear clustering rule  $C(\mathbf{x}) = \mathbf{1}(\mathbf{w}^T \mathbf{x} + b > 0)$ , the clustering error can be written as

$$\min\left\{\pi_1 \Phi\left(\frac{\mathbf{w}^T \boldsymbol{\mu}_1 + b}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}}\right) + \pi_0 \Phi\left(-\frac{\mathbf{w}^T \boldsymbol{\mu}_0 + b}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}}\right), \pi_1 \Phi\left(-\frac{\mathbf{w}^T \boldsymbol{\mu}_1 + b}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}}\right) + \pi_0 \Phi\left(\frac{\mathbf{w}^T \boldsymbol{\mu}_0 + b}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}}\right)\right\}. \quad (8)$$

It is easy to see that using the Bayes classifier as a linear clustering rule minimizes the clustering error, and this error is equal to the Bayes classification error.

A clustering rule is trained using  $n$  i.i.d. samples  $\mathbf{x}_1, \dots, \mathbf{x}_n$ . Here we only have access to the features, but not the latent indicator (the label).

### 2.3 Excess risk

Note that in both classification and clustering, the Bayes classifier minimizes the error, and the minimum error is the Bayes classification error. Taking this into consideration, for a given Gaussian classification or clustering problem, we define the excess risk of a classifier or clustering rule  $C$  as

$$\mathcal{E}(C) = \mathcal{R}(C) - \mathcal{R}^*, \quad (9)$$

where  $\mathcal{R}^* = 1/2 - 1/2 \int |\pi_1 \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}) - \pi_0 \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma})| d\mathbf{x} > 0$  is the Bayes classification error.

### 2.4 Assumptions

In this paper, we consider classification and clustering of two balanced isotropic Gaussian distributions. In classification, we consider distributions of the type  $\Pr[Y = 1] = \Pr[Y = 0] = 1/2$  and  $\mathbf{X} \sim \mathcal{N}(\boldsymbol{\mu}_Y, \sigma^2 \mathbf{I})$ . In clustering, we consider distributions of the type  $\frac{1}{2} \mathcal{N}(\boldsymbol{\mu}_1, \sigma^2 \mathbf{I}) + \frac{1}{2} \mathcal{N}(\boldsymbol{\mu}_0, \sigma^2 \mathbf{I})$ . Our setting is similar to those presented in [4, 8, 1]. We will present results in the general setting without any sparsity, and then results with the following sparsity assumption  $s = \|\boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)\|_0 = \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0\|_0$ .

## 3 Bounds for Gaussian Classification

In this section, we present lower and upper bounds for Gaussian classification, both in the general setting with-

out sparsity, and in the high-dimensional sparse setting where the difference between the means is sparse.

Our lower bounds show that when there are not enough samples, the best any classifier can do is close to random guessing. For the general setting, we show that Fisher's linear discriminant method achieves the minimax rate up to constant factors. Note that the constants in the bounds can be improved.

### 3.1 General Setting without Sparsity

#### 3.1.1 Lower Bound

For the lower bound, we consider Gaussian classification problems where  $\Pr[Y = 1] = \Pr[Y = 0] = 1/2$ ,  $\mathbf{X} | Y \sim \mathcal{N}(\boldsymbol{\mu}_Y, \mathbf{I})$  with  $\boldsymbol{\mu}_1 = -\boldsymbol{\mu}_0 = \boldsymbol{\mu}$ . In this case, the Bayes classification error is  $1 - \Phi(\|\boldsymbol{\mu}\|_2)$ . We will index each classification problem with  $\boldsymbol{\mu}$ .

The following theorem provides a lower bound for Gaussian classification when enough samples are available.

**Theorem 1** (Classification excess risk lower bound). *Let  $\rho > 0$  be a fixed number. For sufficiently large  $p$  and  $n$ , and any classifier  $C$  trained using  $n$  samples, we have*

$$\max_{\|\boldsymbol{\mu}\|_2 = \rho/2} \mathbb{E}_{\boldsymbol{\mu}}[\mathcal{E}_{\boldsymbol{\mu}}(C)] \gtrsim e^{-\rho^2/8} \min\left\{\frac{1}{\rho} \frac{p}{n}, \rho\right\}. \quad (10)$$

The proof of Theorem 1 is based on the following theorem for general classification problems, which provides a ‘‘triangle inequality’’ in Fano's method [25] for a fixed classifier's excess risk in different classification problems.

**Theorem 2** (A ‘‘triangle inequality’’ for classification excess risk). *Let  $i$  correspond to a classification problem where  $\Pr[Y = y] = \pi_y^{(i)}$  and  $\mathbf{X} | Y \sim p_Y^{(i)}(\mathbf{X})$ . Denote the Bayes classifier as  $C^{*(i)}(\mathbf{x}) = \mathbf{1}((\pi_1^{(i)} p_1^{(i)}(\mathbf{x})) / (\pi_0^{(i)} p_0^{(i)}(\mathbf{x})) > 1)$ . Suppose we have another classification problem  $j$ . Then for any fixed classifier  $C$ ,*

$$\mathcal{E}_i(C) + \mathcal{E}_j(C) \geq \int_{C^{*(i)} \neq C^{*(j)}} \min\{|\pi_1^{(i)} p_1^{(i)} - \pi_0^{(i)} p_0^{(i)}|, |\pi_1^{(j)} p_1^{(j)} - \pi_0^{(j)} p_0^{(j)}|\} d\mathbf{x}. \quad (11)$$

Indeed, for a classification problem where  $\Pr[Y = y] = \pi_y$  and  $\mathbf{X} | Y \sim p_Y(\mathbf{X})$ , a classifier  $C$ 's excess risk  $\mathcal{E}(C) = \int_{C \neq C^*} |\pi_1 p_1 - \pi_0 p_0| d\mathbf{x}$  where  $C^*$  is the Bayes classifier (Theorem 2.2, [7]).

To construct a packing for applying Fano's method, we use the following lemma on existence of sparse sets (Lemma 4.10, [18]). This lemma is also known as the Gilbert-Varshamov bound (Theorem 17.2, [12]). It is stated here for completeness.

**Lemma 1** (Existence of sparse sets (Lemma 4.10, [18])). *Let  $\Psi = \{\psi \in \{-1, +1\}^p : \|\psi\|_0 = s\}$  for positive integers  $p$  and  $1 \leq s < p/4$ . Then there exists  $\psi_1, \dots, \psi_N$  such that the Hamming distance  $\delta(\psi_i, \psi_j) > s/2$  for all  $1 \leq i < j \leq N$ , and  $\log N \geq \frac{s}{5} \log \frac{p}{s}$ .*

We briefly sketch Theorem 1's proof. Let us consider the set  $M = \{\boldsymbol{\mu} \mid (\sqrt{\rho^2/4 - s\alpha^2}, \pm\alpha\psi_1, \dots, \pm\alpha\psi_{p-1}) \in \mathbb{R}^p\}$  where  $s = (p-1)/6 > p/8$ , and  $\psi_1, \dots, \psi_{p-1}$  are as given in Lemma 1. And we set  $\alpha = 0.001 \min\{\sqrt{1/n}, \rho/(2\sqrt{s})\}$ . Using Theorem 2 we can show that for any fixed  $\boldsymbol{\mu}$ , if  $\mathcal{E}_\mu > 4 \times 10^{-9} \frac{e^{-\rho^2/8}}{\rho} s\alpha^2$  then we must have  $\mathcal{E}_{\boldsymbol{\mu}'} < 4 \times 10^{-9} \frac{e^{-\rho^2/8}}{\rho} s\alpha^2$  for all other  $\boldsymbol{\mu}' \neq \boldsymbol{\mu}$ . This observation reduces the problem to a testing problem. So using Fano's method, we can show that

$$\begin{aligned} & \Pr[\mathcal{E}_\mu(\mathcal{C}) \geq 4 \times 10^{-9} \frac{e^{-\rho^2/8}}{\rho} s\alpha^2] \\ & \geq 1 - \frac{n \max_{\boldsymbol{\mu} \neq \boldsymbol{\mu}'} \text{KL}(\mathbb{P}_\mu \parallel \mathbb{P}_{\boldsymbol{\mu}'}) + \log 2}{\log |M|} \\ & \geq m1 - \frac{ns\alpha^2 + \log 2}{\log |M|} = \Omega(1), \end{aligned} \quad (12)$$

where  $\mathbb{P}_\mu$  is the joint distribution of  $\mathbf{X}, Y$  with parameter  $\boldsymbol{\mu}$ , and the probability is over the uniform distribution of  $\boldsymbol{\mu}$  on  $M$ . Here, we used  $\text{KL}(\mathbb{P}_\mu \parallel \mathbb{P}_{\boldsymbol{\mu}'}) = \frac{1}{2} \|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_2^2$ . Finally, we have

$$\begin{aligned} & \max_{\boldsymbol{\mu}} \mathbb{E}_\mu[\mathcal{E}_\mu(\mathcal{C})] \geq \frac{1}{|M|} \sum_{\boldsymbol{\mu} \in M} \mathbb{E}_\mu[\mathcal{E}_\mu(\mathcal{C})] \\ & \gtrsim \frac{e^{-\rho^2/8}}{\rho} s\alpha^2 \Pr[\mathcal{E}_\mu(\mathcal{C}) \geq 4 \times 10^{-9} \frac{e^{-\rho^2/8}}{\rho} s\alpha^2] \\ & \gtrsim \frac{e^{-\rho^2/8}}{\rho} s\alpha^2 \gtrsim e^{-\rho^2/8} \min\{\frac{1}{\rho n}, \rho\}. \end{aligned} \quad (13)$$

The next theorem states that when there are not enough samples, the best that any classifier can do is close to random guessing. Random guessing uniformly randomly assigns a sample into one of the two classes, and its classification error is  $1/2$ .

**Theorem 3** (Impossibility of classification). *For sufficiently large  $p$  and  $n$ , let  $\rho > 0$ . When  $n \max\{\rho^2, \rho^4\}/p \rightarrow 0$  and  $p, n \rightarrow +\infty$ , then for any classifier  $\mathcal{C}$  trained using  $n$  samples, we have*

$$\inf_{\mathcal{C}} \max_{\|\boldsymbol{\mu}\|_2 = \rho/2} \mathbb{E}_\mu[\mathcal{R}_\mu(\mathcal{C})] \rightarrow \frac{1}{2}. \quad (14)$$

Indeed, notice that  $\max_{\boldsymbol{\mu}} \mathbb{E}_\mu[R_\mu(\mathcal{C})] \geq \mathbb{E}[\mathbb{E}_\mu[R_\mu(\mathcal{C})]]$ , where the outer expectation is taken over the uniform distribution of  $\boldsymbol{\mu}$  on  $\|\boldsymbol{\mu}\|_2 = \rho/2$ , and RHS is equivalent to a Bayesian problem with a uniform prior over  $\boldsymbol{\mu}$ . In the Bayesian problem, it is easy to see that any classifier's classification error is lower bounded by that

of the MAP classifier. The MAP classifier  $\mathcal{C}_{\text{MAP}}$  is a linear classifier, and it can be written as

$$\begin{aligned} \mathcal{C}_{\text{MAP}}(\mathbf{x}) &= \arg \max_y \Pr[Y = y \mid \mathbf{X} = \mathbf{x}, (\mathbf{x}_1, y_1), \dots] \\ &= \mathbf{1}((\sum (2y_i - 1)\mathbf{x}_i^T)\mathbf{x} > 0). \end{aligned} \quad (15)$$

Thus, we can use (4) to compute its classification error. Let  $\hat{\boldsymbol{w}} = \frac{1}{n} \sum (2y_i - 1)\mathbf{x}_i$ . When conditioned on  $\boldsymbol{\mu}$ ,  $\mathbb{E}_\mu[\mathcal{R}(\mathcal{C}_{\text{MAP}})] \rightarrow 1/2$  uniformly for all  $\boldsymbol{\mu}$ , because  $|\boldsymbol{\mu}^T \hat{\boldsymbol{w}} / \|\hat{\boldsymbol{w}}\|_2| \rightarrow 0$  uniformly. This establishes that the MAP classifier's classification error approaches  $1/2$  as  $n \max\{\rho^2, \rho^4\}/p \rightarrow 0$  and  $p, n \rightarrow +\infty$ .

### 3.1.2 Upper Bound

We show that Fisher's linear discriminant achieves the minimax rate. Here we consider isotropic Gaussian classification problems with  $\Pr[Y = 1] = \Pr[Y = 0] = 1/2$  and  $\mathbf{X} \mid Y \sim \mathcal{N}(\boldsymbol{\mu}_Y, \sigma^2 \mathbf{I})$ , where  $\boldsymbol{\mu}_1, \boldsymbol{\mu}_0$ , and  $\sigma$  are unknown. We assume that  $\|\boldsymbol{\mu}_1\|_2 = O(1)$ ,  $\|\boldsymbol{\mu}_0\|_2 = O(1)$ , and  $\sigma = \Theta(1)$ . Thus  $\rho = \sqrt{(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)} = O(1)$  and  $e^{-\rho^2/8} = \Omega(1)$ .

For  $y = 1, 0$ , let  $n_y = \sum \mathbf{1}(y_i = y)$ . We estimate each class's mean  $\hat{\boldsymbol{\mu}}_y = \frac{1}{n_y} \sum_{y_i=y} \mathbf{x}_i$ . And the trained classifier is  $\hat{\mathcal{C}}(\mathbf{x}) = \mathbf{1}(\hat{\boldsymbol{w}}^T \mathbf{x} + \hat{b} > 0)$  where  $\hat{\boldsymbol{w}} = \hat{\boldsymbol{\mu}}_1 - \hat{\boldsymbol{\mu}}_0$ , and  $\hat{b} = \frac{1}{2}(-\hat{\boldsymbol{\mu}}_1^T \hat{\boldsymbol{\mu}}_1 + \hat{\boldsymbol{\mu}}_0^T \hat{\boldsymbol{\mu}}_0)$ .

**Theorem 4** (Fisher's linear discriminant upper bound). *Let  $\rho = \sqrt{(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)} = O(1)$ . When  $n \gtrsim p \log \delta^{-1} / \rho^4$ , with probability at least  $1 - \delta$ , for Fisher's linear discriminant, we have*

$$\mathcal{E}(\hat{\mathcal{C}}) \lesssim \frac{1}{\rho n} \log \frac{1}{\delta}. \quad (16)$$

The proof of Theorem 4 is based on Theorem 1 of [16], which connects parameter estimation error with classification excess risk. We state this theorem here for completeness.

**Theorem 5** (Classification excess risk upper bound (Theorem 1, [16])). *In a Gaussian classification problem with  $\Pr[Y = 1] = \Pr[Y = 0] = 1/2$  and  $\mathbf{X} \mid Y \sim \mathcal{N}(\boldsymbol{\mu}_Y, \boldsymbol{\Sigma})$ , the excess risk of a linear classifier  $\mathcal{C}(\mathbf{x}) = \mathbf{1}(\mathbf{w}^T \mathbf{x} + b > 0)$  is bounded by*

$$\mathcal{E}(\mathcal{C}) \lesssim \rho e^{-\rho^2/8} (e_1^2 + e_0^2) + e_1^3 + e_0^3, \quad (17)$$

where  $e_1 = |\frac{\mathbf{w}^T \boldsymbol{\mu}_1 + b}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}} - \rho/2|$  and  $e_0 = |\frac{\mathbf{w}^T \boldsymbol{\mu}_0 + b}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}} + \rho/2|$  with  $\rho = \sqrt{(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)}$ .

### 3.2 Sparse Setting

In the sparse setting, we consider Gaussian classification problems with  $\Pr[Y = 1] = \Pr[Y = 0] = 1/2$  and  $\mathbf{X} \mid Y \sim \mathcal{N}(\boldsymbol{\mu}_Y, \sigma^2 \mathbf{I})$ , with the additional requirement that  $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0$  is  $s$ -sparse ( $\|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0\|_0 = s$ ).

### 3.2.1 Lower Bound

Similar to Section 3.1.1, for the lower bound we consider the Gaussian classification problems where  $\Pr[Y = 1] = \Pr[Y = 0] = 1/2$ ,  $\mathbf{X} | Y \sim \mathcal{N}(\boldsymbol{\mu}_Y, \mathbf{I})$  with  $\boldsymbol{\mu}_1 = -\boldsymbol{\mu}_0 = \boldsymbol{\mu}$ . The additional requirement here is that  $\boldsymbol{\mu}$  is  $s$ -sparse ( $\|\boldsymbol{\mu}\|_0 = s$ ).

The following theorem, in parallel with Theorem 1, establishes a lower bound for the excess risk of sparse classification.

**Theorem 6** (Sparse classification excess risk lower bound). *Let  $\rho > 0$  be a fixed number. For sufficiently large  $p$  and  $n$ , let  $1 \leq s < p$ , then for any classifier  $\mathcal{C}$  trained using  $n$  samples, we have*

$$\max_{\substack{\|\boldsymbol{\mu}\|_2 = \rho/2 \\ \|\boldsymbol{\mu}\|_0 = s}} \mathbb{E}_{\boldsymbol{\mu}}[\mathcal{E}_{\boldsymbol{\mu}}(\mathcal{C})] \gtrsim e^{-\rho^2/8} \min\left\{\frac{1}{\rho} \frac{s \log \frac{p}{s}}{n}, \rho\right\}. \quad (18)$$

The proof of Theorem 6 is similar to that of Theorem 1, and it is also based on Lemma 1 and Theorem 2. Theorem 6's proof is omitted for brevity.

### 3.2.2 Upper Bound

We show that a modified version Fisher's linear discriminant achieves the minimax rate in sparse classification. Here we consider isotropic Gaussian classification problems with  $\Pr[Y = 1] = \Pr[Y = 0] = 1/2$  and  $\mathbf{X}|Y \sim \mathcal{N}(\boldsymbol{\mu}_Y, \sigma^2 \mathbf{I})$ , where  $\boldsymbol{\mu}_1$ ,  $\boldsymbol{\mu}_0$ , and  $\sigma$  are unknown. In the sparse setting,  $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0$  is  $s$ -sparse ( $\|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0\|_0 = s$ ). We assume that  $\|\boldsymbol{\mu}_1\|_2 = O(1)$ ,  $\|\boldsymbol{\mu}_0\|_2 = O(1)$ , and  $\sigma = \Theta(1)$ . Thus  $\rho = \sqrt{(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)} = O(1)$  and  $e^{-\rho^2/8} = \Omega(1)$ .

To train the classifier  $\mathcal{C}(\mathbf{x}) = \hat{\mathbf{w}}^T \mathbf{x} + \hat{b}$ , we split the i.i.d. samples into two parts of size  $n^{(1)} \approx n^{(2)} \approx n/2$ . On the first part of the data, we estimate each class's mean  $\hat{\boldsymbol{\mu}}_y = \frac{1}{n_y} \sum_{y_i^{(1)}=y} \mathbf{x}_i^{(1)}$  where  $n_y = \sum \mathbf{1}(y_i^{(1)} = y)$  for  $y = 1, 0$ . We estimate  $\hat{\mathbf{w}}$  by solving the following  $\ell_1$  regularized optimization problem

$$\min_{\mathbf{w}} \|\mathbf{w} - (\hat{\boldsymbol{\mu}}_1 - \hat{\boldsymbol{\mu}}_0)\|_2^2 + \lambda \|\mathbf{w}\|_1, \quad (19)$$

with  $\lambda = \Theta(\sigma \sqrt{\log p/n})$ . When  $\sigma$  is unknown, empirically cross validation is known to be effective in selecting a suitable regularization parameter  $\lambda$  [9]. To estimate  $\hat{b}$ , we use the second part of the data and  $\hat{\mathbf{w}}$

$$\hat{b} = \frac{1}{2} \hat{\mathbf{w}}^T \left( \frac{1}{n^{(2)}} \sum \mathbf{x}_i^{(2)} \right). \quad (20)$$

The following theorem, which is a special case of the result for parameter recovery using  $\ell_1$  regularized optimization [20], shows that we can successfully recover  $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0$  using  $\hat{\mathbf{w}}$ .

**Theorem 7** (Sparse mean estimation using  $\ell_1$  regularized least squares). *Let  $\boldsymbol{\theta}^* \in \mathbb{R}^p$  be a  $s$ -sparse vector ( $\|\boldsymbol{\theta}^*\|_0 = s < p/2$ ). Suppose we observe  $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}^* + \boldsymbol{\epsilon}$ . Set  $\lambda = 4\|\boldsymbol{\epsilon}\|_{\infty}$ . If we estimate  $\boldsymbol{\theta}^*$  using*

$$\hat{\mathbf{w}} = \arg \min_{\mathbf{w}} \|\mathbf{w} - \hat{\boldsymbol{\theta}}\|_2^2 + \lambda \|\mathbf{w}\|_1, \quad (21)$$

then we have

$$\begin{aligned} \|\hat{\mathbf{w}} - \boldsymbol{\theta}^*\|_1 &\lesssim s \|\boldsymbol{\epsilon}\|_{\infty}, \\ \|\hat{\mathbf{w}} - \boldsymbol{\theta}^*\|_2 &\lesssim \sqrt{s} \|\boldsymbol{\epsilon}\|_{\infty}. \end{aligned} \quad (22)$$

A key component in the proof of Theorem 7 is that  $\mathbf{I}$  satisfies the restricted eigenvalue condition, which follows from the fact that  $\mathbf{I}$  satisfies the Restricted Isometry Property [23].

The next theorem, in parallel with Theorem 4, shows that sparse Fisher's linear discriminant achieves the minimax rate up to logarithmic factors.

**Theorem 8** (Sparse Fisher's linear discriminant upper bound). *Let  $\rho = \sqrt{(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)} = O(1)$ . When  $n \gtrsim s \log p \log \delta^{-1}/\rho^4$  and  $1 \leq s < p/2$ , then with probability at least  $1 - \delta$ , sparse Fisher's linear discriminant satisfies*

$$\mathcal{E}(\hat{\mathcal{C}}) \lesssim \frac{1}{\rho} \frac{s \log p}{n} \log \frac{1}{\delta}. \quad (23)$$

## 4 Bounds for Gaussian clustering

In this section, we present lower and upper bounds for Gaussian clustering, both in the general setting without sparsity, and in the high-dimensional sparse setting where the difference between the means is sparse.

### 4.1 General Setting without Sparsity

#### 4.1.1 Lower Bound

For the lower bound, we consider Gaussian clustering problems where  $\mathbf{X} \sim \frac{1}{2} \mathcal{N}(\boldsymbol{\mu}_1, \mathbf{I}) + \frac{1}{2} \mathcal{N}(\boldsymbol{\mu}_0, \mathbf{I})$  with  $\boldsymbol{\mu}_1 = -\boldsymbol{\mu}_0 = \boldsymbol{\mu}$ . In this case, the optimal clustering rule's error is equal to the Bayes classification error  $1 - \Phi(\|\boldsymbol{\mu}\|_2)$ . We will index each clustering problem with  $\boldsymbol{\mu}$ .

Similar to Theorem 1, the following theorem provides a lower bound for Gaussian clustering, but with a different dependency on the separation between the two clusters ( $\rho = \sqrt{(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)}$ ).

**Theorem 9** (Clustering excess risk lower bound). *Let  $\rho > 0$  be a fixed number. For sufficiently large  $p$  and  $n$ , and any clustering rule  $\mathcal{C}$  trained using  $n$  samples, we have*

$$\max_{\|\boldsymbol{\mu}\|_2 = \rho/2} \mathbb{E}_{\boldsymbol{\mu}}[\mathcal{E}_{\boldsymbol{\mu}}(\mathcal{C})] \gtrsim e^{-\rho^2/8} \min\left\{\frac{1}{\rho^3} \frac{p}{n}, \rho\right\}. \quad (24)$$

Similar to the proof of Theorem 1, the proof of Theorem 9 is based on the following theorem for general clustering problems, which provides a “triangle inequality” in Fano’s method [25] for a fixed clustering rule’s excess risk in different clustering problems.

**Theorem 10** (A “triangle inequality” for clustering excess risk). *Let  $i$  correspond to a clustering problem where  $\mathbf{X} \sim \pi_1 p_1^{(i)}(\mathbf{X}) + \pi_0 p_0^{(i)}(\mathbf{X})$ . Denote the optimal clustering rule as  $\mathcal{C}^{*(i)}(\mathbf{x}) = \mathbf{1}((\pi_1^{(i)} p_1^{(i)}(\mathbf{x})) / (\pi_0^{(i)} p_0^{(i)}(\mathbf{x})) > 1)$ . Suppose we have another clustering problem  $j$ . Then for any fixed clustering rule  $\mathcal{C}$ ,*

$$\begin{aligned} & \mathcal{E}_i(\mathcal{C}) + \mathcal{E}_j(\mathcal{C}) \\ & \geq \min\left\{ \int_{\mathcal{C}^{*(i)} \neq \mathcal{C}^{*(j)}} \min\{|\pi_1^{(i)} p_1^{(i)} - \pi_0^{(i)} p_0^{(i)}|, \right. \\ & \quad \left. |\pi_1^{(j)} p_1^{(j)} - \pi_0^{(j)} p_0^{(j)}|\} d\mathbf{x}, \right. \\ & \quad \left. \int_{\mathcal{C}^{*(i)} = \mathcal{C}^{*(j)}} \min\{|\pi_1^{(i)} p_1^{(i)} - \pi_0^{(i)} p_0^{(i)}|, \right. \\ & \quad \left. |\pi_1^{(j)} p_1^{(j)} - \pi_0^{(j)} p_0^{(j)}|\} d\mathbf{x} \right\}. \end{aligned} \quad (25)$$

Indeed, for a clustering problem where  $\mathbf{X} \sim \pi_1 p_1(\mathbf{X}) + \pi_0 p_0(\mathbf{X})$ , a clustering rule  $\mathcal{C}$ ’s excess risk  $\mathcal{E}(\mathcal{C}) = \min\{\int_{\mathcal{C} \neq \mathcal{C}^*} |\pi_1 p_1 - \pi_0 p_0| d\mathbf{x}, \int_{\mathcal{C} = \mathcal{C}^*} |\pi_1 p_1 - \pi_0 p_0| d\mathbf{x}\}$  where  $\mathcal{C}^*$  is the optimal clustering rule. This follows from the definition of clustering error (7), and the Bayes classifier’s classification error (Theorem 2.2, [7]).

The proof of Theorem 9 is similar to that of Theorem 1. However, in clustering we have  $\text{KL}(\mathbb{P}_{\boldsymbol{\mu}}, \mathbb{P}_{\boldsymbol{\mu}'}) \lesssim \rho^4(1 - \boldsymbol{\mu}^T \boldsymbol{\mu}' / (\|\boldsymbol{\mu}\|_2 \|\boldsymbol{\mu}'\|_2))$  (Proposition 24, [1]). Thus, in clustering the dependency on  $\rho$  is different from classification.

Notice that, for a classification problem, we can always remove the labels, and treat it as a clustering problem. Denote the learned clustering rule as  $\hat{\mathcal{C}}$ , then both  $\hat{\mathcal{C}}$  and  $1 - \hat{\mathcal{C}}$  can be used as classifiers for the original classification problem. Thus we have the following corollary to Theorem 3, which states that, without enough samples the best any clustering rule can do is close to random guessing.

**Corollary 1** (Impossibility of clustering). *For sufficiently large  $p$  and  $n$ , let  $\rho > 0$ . When  $n \max\{\rho^2, \rho^4\}/p \rightarrow 0$  and  $p, n \rightarrow +\infty$ , then for any clustering rule  $\mathcal{C}$  trained using  $n$  samples, we have*

$$\inf_{\mathcal{C}} \max_{\|\boldsymbol{\mu}\|_2 = \rho/2} \mathbb{E}_{\boldsymbol{\mu}}[\mathcal{R}_{\boldsymbol{\mu}}(\mathcal{C})] \rightarrow \frac{1}{2}. \quad (26)$$

#### 4.1.2 Upper Bound

For the upper bound, we will consider the setting with  $\mathbf{X} \sim \mathbf{X} \sim \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}_1, \sigma^2\mathbf{I}) + \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}_0, \sigma^2\mathbf{I})$  where  $\boldsymbol{\mu}_1$ ,  $\boldsymbol{\mu}_0$ , and  $\sigma$  are unknown. We assume that

$\|\boldsymbol{\mu}_1\| = O(1)$ ,  $\|\boldsymbol{\mu}_0\| = O(1)$ , and  $\sigma = \Theta(1)$ . Thus  $\rho = \sqrt{(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)} = O(1)$  and  $e^{-\rho^2/8} = \Omega(1)$ .

Our method for clustering is similar to that of [1]. To estimate a clustering rule, we first compute the sample mean  $\hat{\mathbf{m}} = \frac{1}{n} \sum \mathbf{x}_i$ , then we compute the largest eigenvalue’s normalized eigenvector  $\hat{\mathbf{v}}$  of the sample covariance matrix  $\hat{\boldsymbol{\Sigma}} = \frac{1}{n} \sum (\mathbf{x}_i - \hat{\mathbf{m}})(\mathbf{x}_i - \hat{\mathbf{m}})^T$ . Next, we use  $\hat{\mathcal{C}}(\mathbf{x}) = \hat{\mathbf{v}}^T \mathbf{x} - \hat{\mathbf{v}}^T \hat{\mathbf{m}}$  as the estimated clustering rule.

**Theorem 11** (Clustering upper bound). *For sufficiently large  $n$  and  $p$ , let  $\rho = \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0\|_2 / \sigma = O(1)$ . When  $n \gtrsim p \log(pn) / \rho^6$ , the expected excess risk of  $\hat{\mathcal{C}}$  satisfies*

$$\mathbb{E}[\mathcal{E}(\hat{\mathcal{C}})] \lesssim \frac{1}{\rho^3} \frac{p}{n} \log(pn). \quad (27)$$

To prove Theorem 11, we will use the following corollary on the connection between parameter estimation error and clustering error, which follows from Theorem 5 and the definition of clustering error (7).

**Corollary 2** (Clustering excess risk upper bound). *In a Gaussian clustering problem with  $\mathbf{X} \sim \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}) + \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma})$ , the excess risk of a linear clustering rule  $\mathcal{C}(\mathbf{x}) = \mathbf{1}(\mathbf{w}^T \mathbf{x} + b > 0)$  is bounded by*

$$\begin{aligned} \mathcal{E}(\mathcal{C}) \lesssim \min\{ & \rho e^{-\rho^2/8} (e_1^2 + e_0^2) + e_1^3 + e_0^3, \\ & \rho e^{-\rho^2/8} (f_1^2 + f_0^2) + f_1^3 + f_0^3\}, \end{aligned} \quad (28)$$

where  $e_1 = |\frac{\mathbf{w}^T \boldsymbol{\mu}_1 + b}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}} - \rho/2|$ ,  $e_0 = |\frac{\mathbf{w}^T \boldsymbol{\mu}_0 + b}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}} + \rho/2|$ ,  $f_1 = |\frac{\mathbf{w}^T \boldsymbol{\mu}_1 + b}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}} + \rho/2|$ ,  $f_0 = |\frac{\mathbf{w}^T \boldsymbol{\mu}_0 + b}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}} - \rho/2|$  and  $\rho = \sqrt{(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)}$ .

The proof of Theorem 11 uses Proposition 6 of [1], which gives the error of using  $\hat{\mathbf{v}}$  to estimate the direction of  $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0$ . We state it here for completeness.

**Proposition 1** (Proposition 6, [1]). *Suppose  $n > 4p$ , define  $\cos \beta = \boldsymbol{\Delta} \boldsymbol{\mu}^T \hat{\mathbf{v}} / \|\boldsymbol{\Delta} \boldsymbol{\mu}\|_2$ . For any  $0 \leq \delta < (p - 1) / \sqrt{e}$ , if  $\max\{4/\rho^2, 2/\rho\} \sqrt{\frac{\max\{d, 8 \log \delta^{-1}\}}{n}} < 1/180$ , then with probability at least  $1 - 12\delta - 2e^{-n/20}$ , we have*

$$\sin \beta \leq 14 \max\{4/\rho^2, 2/\rho\} \sqrt{p} \sqrt{\frac{10}{n} \log \frac{p}{\delta} \max\{1, \frac{10}{n} \log \frac{p}{\delta}\}}. \quad (29)$$

## 4.2 Sparse Setting

In the sparse setting, we consider Gaussian clustering problems with  $\mathbf{X} \sim \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}_1, \sigma^2\mathbf{I}) + \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}_0, \sigma^2\mathbf{I})$ , with the additional requirement that  $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0$  is  $s$ -sparse ( $\|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0\|_0 = s$ ).

### 4.2.1 Lower Bound

Similar to Section 4.1.1, for the lower bound we consider the Gaussian clustering problems where  $\mathbf{X} \sim \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}_1, \sigma^2\mathbf{I}) + \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}_0, \sigma^2\mathbf{I})$  with  $\boldsymbol{\mu}_1 = -\boldsymbol{\mu}_0 = \boldsymbol{\mu}$ . The additional requirement here is that  $\boldsymbol{\mu}$  is  $s$ -sparse ( $\|\boldsymbol{\mu}\|_0 = s$ ).

The following theorem, in parallel with Theorem 9, establishes a lower bound for the excess risk of sparse clustering.

**Theorem 12** (Sparse clustering excess risk lower bound). *Let  $\rho > 0$  be a fixed number. For sufficiently large  $p$  and  $n$ , let  $1 \leq s < p$ , then for any clustering rule  $\mathbf{C}$  trained using  $n$  samples, we have*

$$\max_{\substack{\|\boldsymbol{\mu}\|_2 = \rho/2 \\ \|\boldsymbol{\mu}\|_0 = s}} \mathbb{E}_{\boldsymbol{\mu}}[\mathcal{E}_{\boldsymbol{\mu}}(\mathbf{C})] \gtrsim e^{-\rho^2/8} \min\left\{\frac{1}{\rho^3} \frac{s \log \frac{p}{s}}{n}, \rho\right\}. \quad (30)$$

The proof of Theorem 12 is similar to that of Theorem 9, and it is omitted for brevity.

### 4.2.2 Upper Bound

Here we consider isotropic Gaussian clustering problems with  $\mathbf{X} \sim \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}_1, \sigma^2\mathbf{I}) + \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}_0, \sigma^2\mathbf{I})$ , where  $\boldsymbol{\mu}_1$ ,  $\boldsymbol{\mu}_0$ , and  $\sigma$  are unknown. In the sparse setting,  $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0$  is  $s$ -sparse ( $\|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0\|_0 = s$ ). We assume that  $\|\boldsymbol{\mu}_1\| = O(1)$ ,  $\|\boldsymbol{\mu}_0\| = O(1)$ , and  $\sigma = \Theta(1)$ . Thus  $\rho = \sqrt{(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)} = O(1)$  and  $e^{-\rho^2/8} = \Omega(1)$ .

Our method for sparse clustering is similar to that of [1]. Let  $S = \{i : \boldsymbol{\mu}_1 - \boldsymbol{\mu}_0 \neq 0\}$  be the set of relevant features. We first construct an estimate  $\hat{S}$  of the set of relevant features. Let  $\hat{\mathbf{m}} = \frac{1}{n} \sum \mathbf{x}_i$  and  $\hat{\boldsymbol{\Sigma}} = \frac{1}{n} \sum (\mathbf{x}_i - \hat{\mathbf{m}})(\mathbf{x}_i - \hat{\mathbf{m}})^T$ . Let  $\hat{\tau} = \frac{1+\alpha}{1-\alpha} \min_{1 \leq i \leq p} \hat{\boldsymbol{\Sigma}}_{ii}$  where  $\alpha = \sqrt{6 \log(np)/n} + 2 \log(np)/n$ . Now let  $\hat{S} = \{i : \hat{\boldsymbol{\Sigma}}_{ii} > \hat{\tau}\}$ . Next, we compute the largest eigenvalue's normalized eigenvector  $\hat{\mathbf{v}}$  of the sample covariance matrix restricted to coordinates in  $\hat{S}$ . Finally, we use  $\hat{\mathbf{C}}(\mathbf{x}) = \hat{\mathbf{v}}^T \mathbf{x}_{\hat{S}} - \hat{\mathbf{v}}^T \hat{\mathbf{m}}_{\hat{S}}$  as the estimated clustering rule, where  $\mathbf{x}_{\hat{S}}$  and  $\hat{\mathbf{m}}_{\hat{S}}$  are  $\mathbf{x}$  and  $\hat{\mathbf{m}}$  restricted to coordinates in  $\hat{S}$ , respectively.

**Theorem 13** (Clustering upper bound). *For sufficiently large  $n$ ,  $p$ , and  $s < n/4$ , let  $\rho = \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0\|_2/\sigma$ . When  $\alpha < 1/4$  and  $n \gtrsim s^2 \log(pn)/\rho^8$ , the expected excess risk of  $\hat{\mathbf{C}}$  satisfies*

$$\mathbb{E}[\mathcal{E}(\hat{\mathbf{C}})] \lesssim \frac{1}{\rho^3} \frac{s \log(ns)}{n} + \frac{s}{\rho^2} \left(\frac{\log(pn)}{n}\right)^{\frac{1}{2}}. \quad (31)$$

Although the above bound is not as tight as other previous bounds, we would like to note that sparse Gaussian mixture problems and sparse principal component analysis problems are both computationally and statistically challenging [3, 2, 19, 24, 10].

The proof of Theorem 13 uses Proposition 9 of [1] regarding performance of support recovery.

**Proposition 2** (Proposition 9, [1]). *Assume that  $n \geq 1$ ,  $p \geq 2$ , and  $\alpha < 1/4$ . Define  $\tilde{S} = \{i : |(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)_i| \geq 4\sigma\sqrt{\alpha}\}$ . Then  $\tilde{S} \subseteq \hat{S} \subseteq S$  with probability at least  $1 - 6/n$ .*

We briefly sketch Theorem 13's proof. Let  $\hat{\beta}$  be the angle between  $\hat{\mathbf{v}}_{\hat{S}}$  and  $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0$ ,  $\bar{\beta}$  the angle between  $(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)_{\hat{S}}$  and  $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0$ , and  $\beta$  the angle between  $\hat{\mathbf{v}}_{\hat{S}}$  and  $(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)_{\hat{S}}$ . Proposition 1 shows  $\beta$  is small. Proposition 2 shows  $\bar{\beta}$  is small. Using the triangle inequality in spherical geometry we have  $\hat{\beta} \leq \bar{\beta} + \beta$ . Thus establishing Theorem 13.

### 4.3 Comparison of Different Clustering Error Definitions

Here we compare our definition of clustering error with that of [1], and show relationships between the two definitions. We define clustering error as

$$\min\{\Pr[\mathbf{C} \neq Y], 1 - \Pr[\mathbf{C} \neq Y]\} \quad (32)$$

with  $Y$  being the latent indicator in the Gaussian mixture model, whereas in [1] it is defined as  $\min\{\Pr[\mathbf{C} \neq \mathbf{C}^*], 1 - \Pr[\mathbf{C} \neq \mathbf{C}^*]\}$ .

One advantage of our definition is that we can empirically evaluate the clustering error if true labels are given. This is because our definition of clustering error is the same as classification error in the corresponding classification problem. However, empirically evaluating [1]'s clustering error is not straightforward even if true labels are available.

**Proposition 3** ([1]'s clustering error upper bound (Proposition 7, [1])). *For the clustering problem considered in 4.1.2, [1]'s clustering error of a linear clustering rule  $\mathbf{C}(\mathbf{x}) = \mathbf{1}(w^T \mathbf{x} + b > 0)$  is bounded by*

$$\min\{\Pr[\mathbf{C} \neq \mathbf{C}^*], 1 - \Pr[\mathbf{C} \neq \mathbf{C}^*]\} \lesssim (\epsilon_1 + \epsilon_2 \rho + |\sin \beta|), \quad (33)$$

if  $|b + w^T \frac{(\boldsymbol{\mu}_1 + \boldsymbol{\mu}_0)}{2}| \leq (\sigma \epsilon_1 + \|\frac{\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0}{2}\|_2 \epsilon_2) \|w\|_2$  for some  $\epsilon_1 \geq 0$  and  $0 \leq \epsilon_2 \leq \frac{1}{4}$ , and  $|\sin \beta| \leq 1/\sqrt{5}$ , where  $\cos \beta = \frac{w^T (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)}{\|w\|_2 \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0\|_2}$ .

Also, our clustering error definition captures the intrinsic hardness of the clustering problem better than [1]'s definition. For [1]'s definition, a linear clustering rule's clustering error can be upper bounded using Proposition 3. Following [1], with high probability we have  $\min\{\Pr[\mathbf{C} \neq \mathbf{C}^*], 1 - \Pr[\mathbf{C} \neq \mathbf{C}^*]\} \lesssim \sin \beta$ . For our definition, from Corollary 2 it is easy to derive that  $\min\{\Pr[\mathbf{C} \neq Y], 1 - \Pr[\mathbf{C} \neq Y]\} \lesssim \rho |\sin \beta|^2 + |\sin \beta|^3$  holds with high probability. Thus when  $|\sin \beta| \lesssim \rho$ , which requires  $n \gg p/\rho^6$ , we can easily "convert" these

two definitions into one another up to a constant. Despite this observation, [1]’s definition does not, in a straightforward fashion, exhibit an impossibility result similar to Corollary 1. When  $n \ll p/\rho^2$ , Corollary 1 shows the clustering error defined by us approaches  $1/2$ , which can be interpreted as random guessing. But it is not straightforward to make such an interpretation using [1]’s definition, and to our knowledge an impossibility result using [1]’s definition of clustering error is not known.

## 5 Conclusion

In this paper, we presented minimax lower and upper bounds for Gaussian classification and clustering. Our results explicitly show how the statistical difficulty depends on dimension  $p$ , number of samples  $n$ , sparsity of the optimal classifier/clustering rule  $s$ , and separation between two classes/clusters  $\rho = \sqrt{(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)}$ . Our results here focus on classification and clustering of two isotropic (spherical) Gaussian distributions with equal proportions.

In future work, we plan to improve the bounds’ dependency on separation  $\rho$ . Here we have shown that, in the general non-sparse setting when  $n \max\{\rho^2, \rho^4\}/p \rightarrow 0$  and  $p, n \rightarrow +\infty$ , both classification and clustering are close to impossible. But efficient procedures for classification (Fisher’s linear discriminant) and clustering (principal component analysis) achieve minimax rates when  $n\rho^4/p \rightarrow +\infty$  and  $n\rho^6/p \rightarrow +\infty$ . In the sparse setting, for classification and clustering we have sample lower bounds  $n = \Omega(s \log \frac{p}{s}/\rho^2)$  for classification and  $n = \Omega(s \log \frac{p}{s}/\rho^4)$  for clustering. But efficient procedures for classification (sparse Fisher’s linear discriminant) and clustering (sparse principal component analysis) requires  $n\rho^4/(s \log p) \rightarrow +\infty$  for classification and  $n\rho^8/(s \log p) \rightarrow +\infty$  for clustering. We conjecture that there is a statistical and computational tradeoff in both classification and clustering problems, similar to the result on sparse principal component detection [2]. Another part of future work is to extend this paper’s results to general Gaussian distributions.

## 6 Acknowledgments

C.C. acknowledges partial support of NSF Grants 1609279, 1056028 and 1302435, and the U.S. DoT through the Data-Supported Transportation Operations and Planning (D-STOP) Tier 1 University Transportation Center. P.R. acknowledges the support of ARO via W911NF-12-1-0390 and NSF via IIS-1149803, IIS-1320894, IIS-1447574, and DMS-1264033, and NIH via R01 GM117594-01 as part of the Joint DMS/NIGMS Initiative to Support Research at the

Interface of the Biological and Mathematical Sciences.

## References

- [1] M. Azizyan, A. Singh, and L. Wasserman. Minimax Theory for High Dimensional Gaussian Mixtures with Sparse Mean Separation. In *Advances in Neural Information Processing Systems*, pages 2139–2147, 2013.
- [2] Q. Berthet and P. Rigollet. Complexity theoretic lower bounds for sparse principal component detection. In *Conference on Learning Theory*, pages 1046–1066, 2013.
- [3] Q. Berthet and P. Rigollet. Optimal detection of sparse principal components in high dimension. *The Annals of Statistics*, 41(4):1780–1815, 2013.
- [4] P. J. Bickel and E. Levina. Some theory for Fisher’s linear discriminant function, ‘naive Bayes’, and some alternatives when there are many more variables than observations. *Bernoulli*, pages 989–1010, 2004.
- [5] S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- [6] T. Cai and W. Liu. A direct estimation approach to sparse linear discriminant analysis. *Journal of the American Statistical Association*, 2012.
- [7] L. Devroye, L. Györfi, and G. Lugosi. *A probabilistic theory of pattern recognition*. Springer Science & Business Media, 2013.
- [8] J. Fan and Y. Fan. High dimensional classification using features annealed independence rules. *Annals of Statistics*, 36(6):2605, 2008.
- [9] C. Giraud. *Introduction to high-dimensional statistics*, volume 138. CRC Press, 2014.
- [10] M. Hardt and E. Price. Tight bounds for learning a mixture of two Gaussians. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 753–760. ACM, 2015.
- [11] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics. Springer New York, 2013.
- [12] S. Jukna. *Extremal combinatorics: with applications in computer science*. Springer Science & Business Media, 2011.



- [13] A. Krause, P. Perona, and R. G. Gomes. Discriminative clustering by regularized information maximization. In *Advances in Neural Information Processing Systems*, pages 775–783, 2010.
- [14] G. Lecué. Optimal rates of aggregation in classification under low noise assumption. *Bernoulli*, 13(4):1000–1022, 2007.
- [15] G. Lecué. Classification with minimax fast rates for classes of Bayes rules with sparse representation. *Electronic Journal of Statistics*, 2:741–773, 2008.
- [16] T. Li, A. Prasad, and P. Ravikumar. Fast Classification Rates for High-dimensional Gaussian Generative Models. In *Advances in Neural Information Processing Systems*, pages 1054–1062, 2015.
- [17] E. Mammen and A. Tsybakov. Smooth discrimination analysis. *The Annals of Statistics*, 27(6):1808–1829, 1999.
- [18] P. Massart. *Concentration inequalities and model selection*. Springer, 2007.
- [19] A. Moitra and G. Valiant. Settling the Polynomial Learnability of Mixtures of Gaussians. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 93–102. IEEE, 2010.
- [20] S. Negahban, B. Yu, M. Wainwright, and P. Ravikumar. A unified framework for high-dimensional analysis of  $M$ -estimators with decomposable regularizers. In *Advances in Neural Information Processing Systems*, pages 1348–1356, 2009.
- [21] J. Shao, Y. Wang, X. Deng, and S. Wang. Sparse linear discriminant analysis by thresholding for high dimensional data. *The Annals of statistics*, 39(2):1241–1265, 2011.
- [22] Y. Shi and F. Sha. Information-Theoretical Learning of Discriminative Clusters for Unsupervised Domain Adaptation. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*, pages 1079–1086, 2012.
- [23] S. A. van de Geer and P. Bühlmann. On the conditions used to prove oracle results for the Lasso. *Electronic Journal of Statistics*, 3:1360–1392, 2009.
- [24] N. Verzelen and E. Arias-Castro. Detection and feature selection in sparse mixture models. *arXiv preprint arXiv:1405.1478*, 2014.
- [25] B. Yu. Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer, 1997.