
Minimax Theory for High-dimensional Gaussian Mixtures with Sparse Mean Separation

Martin Azizyan
Machine Learning Department
Carnegie Mellon University
mazizyan@cs.cmu.edu

Aarti Singh
Machine Learning Department
Carnegie Mellon University
aarti@cs.cmu.edu

Larry Wasserman
Department of Statistics
Carnegie Mellon University
larry@stat.cmu.edu

Abstract

While several papers have investigated computationally and statistically efficient methods for learning Gaussian mixtures, precise minimax bounds for their statistical performance as well as fundamental limits in high-dimensional settings are not well-understood. In this paper, we provide precise information theoretic bounds on the clustering accuracy and sample complexity of learning a mixture of two isotropic Gaussians in high dimensions under small mean separation. If there is a sparse subset of relevant dimensions that determine the mean separation, then the sample complexity only depends on the number of relevant dimensions and mean separation, and can be achieved by a simple computationally efficient procedure. Our results provide the first step of a theoretical basis for recent methods that combine feature selection and clustering.

1 Introduction

Gaussian mixture models provide a simple framework for several machine learning problems including clustering, density estimation and classification. Mixtures are especially appealing in high dimensional problems. Perhaps the most common use of Gaussian mixtures is for clustering. Of course, the statistical (and computational) behavior of these methods can degrade in high dimensions. Inspired by the success of variable selection methods in regression, several authors have considered variable selection for clustering. However, there appears to be no theoretical results justifying the advantage of variable selection in high dimensional setting.

To see why some sort of variable selection might be useful, consider clustering n subjects using a vector of d genes for each subject. Typically d is much larger than n which suggests that statistical clustering methods will perform poorly. However, it may be the case that there are only a small number of relevant genes in which case we might expect better behavior by focusing on this small set of relevant genes.

The purpose of this paper is to provide precise bounds on clustering error with mixtures of Gaussians. We consider both the general case where all features are relevant, and the special case where only a subset of features are relevant. Mathematically, we model an irrelevant feature by requiring the mean of that feature to be the same across clusters, so that the feature does not serve to differentiate the groups. Throughout this paper, we use the probability of misclustering an observation, relative to the optimal clustering if we had known the true distribution, as our loss function. This is akin to using excess risk in classification.

This paper makes the following contributions:

- We provide information theoretic bounds on the sample complexity of learning a mixture of two isotropic Gaussians with equal weight in the small mean separation setting that precisely captures the dimension dependence, and matches known sample complexity requirements for some existing algorithms. This also debunks the myth that there is a gap between

statistical and computational complexity of learning mixture of two isotropic Gaussians for small mean separation. Our bounds require non-standard arguments since our loss function does not satisfy the triangle inequality.

- We consider the high-dimensional setting where only a subset of relevant dimensions determine the mean separation between mixture components and show that learning is substantially easier as the sample complexity only depends on the sparse set of relevant dimensions. This provides some theoretical basis for feature selection approaches to clustering.
- We show that a simple computationally feasible procedure nearly achieves the information theoretic sample complexity even in high-dimensional sparse mean separation settings.

Related Work. There is a long and continuing history of research on mixtures of Gaussians. A complete review is not feasible but we mention some highlights of the work most related to ours.

Perhaps the most popular method for estimating a mixture distribution is maximum likelihood. Unfortunately, maximizing the likelihood is NP-Hard. This has led to a stream of work on alternative methods for estimating mixtures. These new algorithms use pairwise distances, spectral methods or the method of moments.

Pairwise methods are developed in Dasgupta (1999), Schulman and Dasgupta (2000) and Arora and Kannan (2001). These methods require the mean separation to increase with dimension. The first one requires the separation to be $\sqrt{d}$ while the latter two improve it to $d^{1/4}$. To avoid this problem, Vempala and Wang (2004) introduced the idea of using spectral methods for estimating mixtures of spherical Gaussians which makes mean separation independent of dimension. The assumption that the components are spherical was removed in Brubaker and Vempala (2008). Their method only requires the components to be separated by a hyperplane and runs in polynomial time, but requires $n = \Omega(d^4 \log d)$ samples. Other spectral methods include Kannan et al. (2005), Achlioptas and McSherry (2005) and Hsu and Kakade (2013). The latter uses clever spectral decompositions together with the method of moments to derive an effective algorithm.

Kalai et al. (2012) use the method of moments to get estimates without requiring separation between components of the mixture components. A similar approach is given in Belkin and Sinha (2010). Chaudhuri et al. (2009) give a modified k -means algorithm for estimating a mixture of two Gaussians. For the large mean separation setting $\mu > 1$, Chaudhuri et al. (2009) show that $n = \tilde{\Omega}(d/\mu^2)$ samples are needed. They also provide an information theoretic bound on the necessary sample complexity of any algorithm which matches the sample complexity of their method (up to log factors) in d and μ . If the mean separation is small $\mu < 1$, they show that $n = \tilde{\Omega}(d/\mu^4)$ samples are sufficient. Our results for the small mean separation setting give a matching necessary condition. Assuming the separation between the component means is *not* too sparse, Chaudhuri and Rao (2008) provide an algorithm for learning the mixture that has polynomial computational and sample complexity.

Most of these papers are concerned with computational efficiency and do not give precise, statistical minimax upper and lower bounds. None of them deal with the case we are interested in, namely, a high dimensional mixture with sparse mean separation.

We should also point out that the results in different papers are not necessarily comparable since different authors use different loss functions. In this paper we use the probability of misclassifying a future observation, relative to how the correct distribution clusters the observation, as our loss function. This should not be confused with the probability of attributing a new observation to the wrong component of the mixture. The latter loss does not typically tend to zero as the sample size increases. Our loss is similar to the excess risk used in classification where we compare the misclassification rate of a classifier to the misclassification rate of the Bayes optimal classifier.

Finally, we remind the reader that our motivation for studying sparsely separated mixtures is that this provides a model for variable selection in clustering problems. There are some relevant recent papers on this problem in the high-dimensional setting. Pan and Shen (2007) use penalized mixture models to do variable selection and clustering simultaneously. Witten and Tibshirani (2010) develop a penalized version of k -means clustering. Related methods include Raftery and Dean (2006); Sun et al. (2012) and Guo et al. (2010). The applied bioinformatics literature also contains a huge number of heuristic methods for this problem. None of these papers provide minimax bounds for the clustering error or provide theoretical evidence of the benefit of using variable selection in unsupervised problems such as clustering.

2 Problem Setup

In this paper, we consider the simple setting of learning a mixture of two isotropic Gaussians with equal mixing weights,¹ given n data points $X_1, \dots, X_n \in \mathbb{R}^d$ drawn i.i.d. from a d -dimensional mixture density function

$$p_\theta(x) = \frac{1}{2}f(x; \mu_1, \sigma^2 I) + \frac{1}{2}f(x; \mu_2, \sigma^2 I),$$

where $f(\cdot; \mu, \Sigma)$ is the density of $\mathcal{N}(\mu, \Sigma)$, $\sigma > 0$ is a fixed constant, and $\theta := (\mu_1, \mu_2) \in \Theta$. We consider two classes Θ of parameters:

$$\begin{aligned}\Theta_\lambda &= \{(\mu_1, \mu_2) : \|\mu_1 - \mu_2\| \geq \lambda\} \\ \Theta_{\lambda,s} &= \{(\mu_1, \mu_2) : \|\mu_1 - \mu_2\| \geq \lambda, \|\mu_1 - \mu_2\|_0 \leq s\} \subseteq \Theta_\lambda.\end{aligned}$$

The first class defines mixtures where the components have a mean separation of at least $\lambda > 0$. The second class defines mixtures with mean separation $\lambda > 0$ along a sparse set of $s \in \{1, \dots, d\}$ dimensions. Also, let P_θ denote the probability measure corresponding to p_θ .

For a mixture with parameter θ , the Bayes optimal classification, that is, assignment of a point $x \in \mathbb{R}^d$ to the correct mixture component, is given by the function

$$F_\theta(x) = \operatorname{argmax}_{i \in \{1,2\}} f(x; \mu_i, \sigma^2 I).$$

Given any other candidate assignment function $F : \mathbb{R}^d \rightarrow \{1, 2\}$, we define the loss incurred by F as

$$L_\theta(F) = \min_{\pi} P_\theta(\{x : F_\theta(x) \neq \pi(F(x))\})$$

where the minimum is over all permutations $\pi : \{1, 2\} \rightarrow \{1, 2\}$. This is the probability of misclustering relative to an oracle that uses the true distribution to do optimal clustering.

We denote by $\hat{F}_n$ any assignment function learned from the data $X_1, \dots, X_n$, also referred to as estimator. The goal of this paper is to quantify how the minimax expected loss (worst case expected loss for the best estimator)

$$R_n \equiv \inf_{\hat{F}_n} \sup_{\theta \in \Theta} \mathbb{E}_\theta L_\theta(\hat{F}_n)$$

scales with number of samples n , the dimension of the feature space d , the number of relevant dimensions s , and the signal-to-noise ratio defined as the ratio of mean separation to standard deviation λ/σ . We will also demonstrate a specific estimator that achieves the minimax scaling.

For the purposes of this paper, we say that feature j is irrelevant if $\mu_1(j) = \mu_2(j)$. Otherwise we say that feature j is relevant.

3 Minimax Bounds

3.1 Small mean separation setting without sparsity

We begin without assuming any sparsity, that is, all features are relevant. In this case, comparing the projections of the data to the projection of the sample mean onto the first principal component suffices to achieve both minimax optimal sample complexity and clustering loss.

Theorem 1 (Upper bound). *Define*

$$\hat{F}_n(x) = \begin{cases} 1 & \text{if } x^T v_1(\hat{\Sigma}_n) \geq \hat{\mu}_n^T v_1(\hat{\Sigma}_n) \\ 2 & \text{otherwise.} \end{cases}$$

where $\hat{\mu}_n = n^{-1} \sum_{i=1}^n X_i$ is the sample mean, $\hat{\Sigma}_n = n^{-1} \sum_{i=1}^n (X_i - \hat{\mu}_n)(X_i - \hat{\mu}_n)^T$ is the sample covariance and $v_1(\hat{\Sigma}_n)$ denotes the eigenvector corresponding to the largest eigenvalue of $\hat{\Sigma}_n$. If $n \geq \max(68, 4d)$, then

$$\sup_{\theta \in \Theta_\lambda} \mathbb{E}_\theta L_\theta(\hat{F}) \leq 600 \max\left(\frac{4\sigma^2}{\lambda^2}, 1\right) \sqrt{\frac{d \log(nd)}{n}}.$$

¹We believe our results should also hold in the unequal mixture weight setting without major modifications.

Furthermore, if $\frac{\lambda}{\sigma} \geq 2 \max(80, 14\sqrt{5d})$, then

$$\sup_{\theta \in \Theta_\lambda} \mathbb{E}_\theta L_\theta(\hat{F}) \leq 17 \exp\left(-\frac{n}{32}\right) + 9 \exp\left(-\frac{\lambda^2}{80\sigma^2}\right).$$

We note that the estimator in Theorem 1 (and that in Theorem 3) does not use knowledge of σ^2 .

Theorem 2 (Lower bound). *Assume that $d \geq 9$ and $\frac{\lambda}{\sigma} \leq 0.2$. Then*

$$\inf_{\hat{F}_n} \sup_{\theta \in \Theta_\lambda} \mathbb{E}_\theta L_\theta(\hat{F}_n) \geq \frac{1}{500} \min \left\{ \frac{\sqrt{\log 2}}{3} \frac{\sigma^2}{\lambda^2} \sqrt{\frac{d-1}{n}}, \frac{1}{4} \right\}.$$

We believe that some of the constants (including lower bound on d and exact upper bound on λ/σ) can be tightened, but the results demonstrate matching scaling behavior of clustering error with d , n and λ/σ . Thus, we see (ignoring constants and log terms) that

$$R_n \approx \frac{\sigma^2}{\lambda^2} \sqrt{\frac{d}{n}}, \quad \text{or equivalently} \quad n \approx \frac{d}{\lambda^4/\sigma^4} \text{ for a constant target value of } R_n.$$

The result is quite intuitive: the dependence on dimension d is as expected. Also we see that the rate depends in a precise way on the signal-to-noise ratio λ/σ . In particular, the results imply that we need $d \leq n$.

In modern high-dimensional datasets, we often have $d > n$ i.e. large number of features and not enough samples. However, inference is usually tractable since not all features are relevant to the learning task at hand. This sparsity of relevant feature set has been successfully exploited in supervised learning problems such as regression and classification. We show next that the same is true for clustering under the Gaussian mixture model.

3.2 Sparse and small mean separation setting

Now we consider the case where there are $s < d$ relevant features. Let S denote the set of relevant features. We begin by constructing an estimator $\hat{S}_n$ of S as follows. Define

$$\hat{\tau}_n = \frac{1 + \alpha}{1 - \alpha} \min_{i \in \{1, \dots, d\}} \hat{\Sigma}_n(i, i), \text{ where}$$

where

$$\alpha = \sqrt{\frac{6 \log(nd)}{n}} + \frac{2 \log(nd)}{n}.$$

Now let

$$\hat{S}_n = \{i \in \{1, \dots, d\} : \hat{\Sigma}_n(i, i) > \hat{\tau}_n\}.$$

Now we use the same method as before, but using only the features in $\hat{S}_n$ identified as relevant.

Theorem 3 (Upper bound). *Define*

$$\hat{F}_n(x) = \begin{cases} 1 & \text{if } x_{\hat{S}_n}^T v_1(\hat{\Sigma}_{\hat{S}_n}) \geq \hat{\mu}_{\hat{S}_n}^T v_1(\hat{\Sigma}_{\hat{S}_n}) \\ 2 & \text{otherwise} \end{cases}$$

where $x_{\hat{S}_n}$ are the coordinates of x restricted to $\hat{S}_n$, and $\hat{\mu}_{\hat{S}_n}$ and $\hat{\Sigma}_{\hat{S}_n}$ are the sample mean and covariance of the data restricted to $\hat{S}_n$. If $n \geq \max(68, 4s)$, $d \geq 2$ and $\alpha \leq \frac{1}{4}$, then

$$\sup_{\theta \in \Theta_{\lambda, s}} \mathbb{E}_\theta L_\theta(\hat{F}) \leq 603 \max\left(\frac{16\sigma^2}{\lambda^2}, 1\right) \sqrt{\frac{s \log(ns)}{n}} + 220 \frac{\sigma \sqrt{s}}{\lambda} \left(\frac{\log(nd)}{n}\right)^{\frac{1}{4}}.$$

Next we find the lower bound.

Theorem 4 (Lower bound). *Assume that $\frac{\lambda}{\sigma} \leq 0.2$, $d \geq 17$, and that $5 \leq s \leq \frac{d+3}{4}$. Then*

$$\inf_{\hat{F}_n} \sup_{\theta \in \Theta_{\lambda, s}} \mathbb{E}_\theta L_\theta(\hat{F}_n) \geq \frac{1}{600} \min \left\{ \sqrt{\frac{8}{45}} \frac{\sigma^2}{\lambda^2} \sqrt{\frac{s-1}{n}} \log\left(\frac{d-1}{s-1}\right), \frac{1}{2} \right\}.$$

We remark again that the constants in our bounds can be tightened, but the results suggest that

$$\frac{\sigma}{\lambda} \left(\frac{s^2 \log d}{n} \right)^{1/4} \succ R_n \succ \frac{\sigma^2}{\lambda^2} \sqrt{\frac{s \log d}{n}},$$

or $n = \Omega \left(\frac{s^2 \log d}{\lambda^4 / \sigma^4} \right)$ for a constant target value of R_n .

In this case, we have a gap between the upper and lower bounds for the clustering loss. Also, the sample complexity can possibly be improved to scale as s (instead of s^2) using a different method. However, notice that the dimension only enters logarithmically. If the number of relevant dimensions is small then we can expect good rates. This provides some justification for feature selection. We conjecture that the lower bound is tight and that the gap could be closed by using a sparse principal component method as in Vu and Lei (2012) to find the relevant features. However, that method is combinatorial and so far there is no known computationally efficient method for implementing it with similar guarantees.

We note that the upper bound is achieved by a two-stage method that first finds the relevant dimensions and then estimates the clusters. This is in contrast to the methods described in the introduction which do clustering and variable selection simultaneously. This raises an interesting question: is it always possible to achieve the minimax rate with a two-stage procedure or are there cases where a simultaneous method outperforms a two-stage procedure? Indeed, it is possible that in the case of general covariance matrices (non-spherical) two-stage methods might fail. We hope to address this question in future work.

4 Proofs of the Lower Bounds

The lower bounds for estimation problems rely on a standard reduction from expected error to hypothesis testing that assumes the loss function is a semi-distance, which the clustering loss isn't. However, a local triangle inequality-type bound can be shown (Proposition 2). This weaker condition can then be used to lower-bound the expected loss, as stated in Proposition 1 (which follows easily from Fano's inequality).

The proof techniques of the sparse and non-sparse lower bounds are almost identical. The main difference is that in the non-sparse case, we use the Varshamov–Gilbert bound (Lemma 1) to construct a set of sufficiently dissimilar hypotheses, whereas in the sparse case we use an analogous result for sparse hypercubes (Lemma 2). See the supplementary material for complete proofs of all results.

In this section and the next, ϕ and Φ denote the univariate standard normal PDF and CDF.

Lemma 1 (Varshamov–Gilbert bound). *Let $\Omega = \{0, 1\}^m$ for $m \geq 8$. There exists a subset $\{\omega_0, \dots, \omega_M\} \subseteq \Omega$ such that $\omega_0 = (0, \dots, 0)$, $\rho(\omega_i, \omega_j) \geq \frac{m}{8}$ for all $0 \leq i < j \leq M$, and $M \geq 2^{m/8}$, where ρ denotes the Hamming distance between two vectors (Tsybakov (2009)).*

Lemma 2. *Let $\Omega = \{\omega \in \{0, 1\}^m : \|\omega\|_0 = s\}$ for integers $m > s \geq 1$ such that $s \leq m/4$. There exist $\omega_0, \dots, \omega_M \in \Omega$ such that $\rho(\omega_i, \omega_j) > s/2$ for all $0 \leq i < j \leq M$, and $M \geq \left(\frac{m}{s}\right)^{s/5} - 1$ (Massart (2007), Lemma 4.10).*

Proposition 1. *Let $\theta_0, \dots, \theta_M \in \Theta_\lambda$ (or $\Theta_{\lambda, s}$), $M \geq 2$, $0 < \alpha < 1/8$, and $\gamma > 0$. If for all $1 \leq i \leq M$, $\text{KL}(P_{\theta_i}, P_{\theta_0}) \leq \frac{\alpha \log M}{n}$, and if $L_{\theta_i}(\hat{F}) < \gamma$ implies $L_{\theta_j}(\hat{F}) \geq \gamma$ for all $0 \leq i \neq j \leq M$ and clusterings $\hat{F}$, then $\inf_{\hat{F}_n} \max_{i \in [0..M]} \mathbb{E}_{\theta_i} L_{\theta_i}(\hat{F}_n) \geq 0.07\gamma$.*

Proposition 2. *For any $\theta, \theta' \in \Theta_\lambda$, and any clustering $\hat{F}$, let $\tau = L_\theta(\hat{F}) + \sqrt{\text{KL}(P_\theta, P_{\theta'})/2}$. If $L_\theta(F_{\theta'}) + \tau \leq 1/2$, then $L_\theta(F_{\theta'}) - \tau \leq L_{\theta'}(\hat{F}) \leq L_\theta(F_{\theta'}) + \tau$.*

We will also need the following two results. Let $\theta = (\mu_0 - \mu/2, \mu_0 + \mu/2)$ and $\theta' = (\mu_0 - \mu'/2, \mu_0 + \mu'/2)$ for $\mu_0, \mu, \mu' \in \mathbb{R}^d$ such that $\|\mu\| = \|\mu'\|$, and let $\cos \beta = \frac{|\mu^T \mu'|}{\|\mu\|^2}$.

Proposition 3. *Let $g(x) = \phi(x)(\phi(x) - x\Phi(-x))$. Then $2g\left(\frac{\|\mu\|}{2\sigma}\right) \sin \beta \cos \beta \leq L_\theta(F_{\theta'}) \leq \frac{\tan \beta}{\pi}$.*

Proposition 4. *Let $\xi = \frac{\|\mu\|}{2\sigma}$. Then $\text{KL}(P_\theta, P_{\theta'}) \leq \xi^4(1 - \cos \beta)$.*

Proof of Theorem 2. Let $\xi = \frac{\lambda}{2\sigma}$, and define $\epsilon = \min \left\{ \frac{\sqrt{\log 2}}{3} \frac{\sigma^2}{\lambda} \frac{1}{\sqrt{n}}, \frac{\lambda}{4\sqrt{d-1}} \right\}$. Define $\lambda_0^2 = \lambda^2 - (d-1)\epsilon^2$. Let $\Omega = \{0, 1\}^{d-1}$. For $\omega = (\omega(1), \dots, \omega(d-1)) \in \Omega$, let $\mu_\omega = \lambda_0 e_d + \sum_{i=1}^{d-1} (2\omega(i) - 1)\epsilon e_i$ (where $\{e_i\}_{i=1}^d$ is the standard basis for $\mathbb{R}^d$). Let $\theta_\omega = (-\frac{\mu_\omega}{2}, \frac{\mu_\omega}{2}) \in \Theta_\lambda$.

By Proposition 4, $\text{KL}(P_{\theta_\omega}, P_{\theta_\nu}) \leq \xi^4 (1 - \cos \beta_{\omega, \nu})$ where $\cos \beta_{\omega, \nu} = 1 - \frac{2\rho(\omega, \nu)\epsilon^2}{\lambda^2}$, $\omega, \nu \in \Omega$, and ρ is the Hamming distance, so $\text{KL}(P_{\theta_\omega}, P_{\theta_\nu}) \leq \xi^4 \frac{2(d-1)\epsilon^2}{\lambda^2}$. By Proposition 3, since $\cos \beta_{\omega, \nu} \geq \frac{1}{2}$,

$$L_{\theta_\omega}(F_{\theta_\nu}) \leq \frac{1}{\pi} \tan \beta_{\omega, \nu} \leq \frac{1}{\pi} \frac{\sqrt{1 + \cos \beta_{\omega, \nu}}}{\cos \beta_{\omega, \nu}} \sqrt{1 - \cos \beta_{\omega, \nu}} \leq \frac{4}{\pi} \frac{\sqrt{d-1}\epsilon}{\lambda}, \text{ and}$$

$$L_{\theta_\omega}(F_{\theta_\nu}) \geq 2g(\xi) \sin \beta_{\omega, \nu} \cos \beta_{\omega, \nu} \geq g(\xi) \sqrt{1 + \cos \beta_{\omega, \nu}} \sqrt{1 - \cos \beta_{\omega, \nu}} \geq \sqrt{2}g(\xi) \frac{\sqrt{\rho(\omega, \nu)}\epsilon}{\lambda}$$

where $g(x) = \phi(x)(\phi(x) - x\Phi(-x))$. By Lemma 1, there exist $\omega_0, \dots, \omega_M \in \Omega$ such that $M \geq 2^{(d-1)/8}$ and $\rho(\omega_i, \omega_j) \geq \frac{d-1}{8}$ for all $0 \leq i < j \leq M$. For simplicity of notation, let $\theta_i = \theta_{\omega_i}$ for all $i \in [0..M]$. Then, for $i \neq j \in [0..M]$,

$$\text{KL}(P_{\theta_i}, P_{\theta_j}) \leq \xi^4 \frac{2(d-1)\epsilon^2}{\lambda^2}, \quad L_{\theta_i}(F_{\theta_j}) \leq \frac{4}{\pi} \frac{\sqrt{d-1}\epsilon}{\lambda} \quad \text{and} \quad L_{\theta_i}(F_{\theta_j}) \geq \frac{1}{2}g(\xi) \frac{\sqrt{d-1}\epsilon}{\lambda}.$$

Define $\gamma = \frac{1}{4}(g(\xi) - 2\xi^2) \frac{\sqrt{d-1}\epsilon}{\lambda}$. Then for any $i \neq j \in [0..M]$, and any $\hat{F}$ such that $L_{\theta_i}(\hat{F}) < \gamma$,

$$L_{\theta_i}(F_{\theta_j}) + L_{\theta_i}(\hat{F}) + \sqrt{\frac{\text{KL}(P_{\theta_i}, P_{\theta_j})}{2}} < \left(\frac{4}{\pi} + \frac{1}{4}(g(\xi) - 2\xi^2) + \xi^2 \right) \frac{\sqrt{d-1}\epsilon}{\lambda} \leq \frac{1}{2}$$

because, for $\xi \leq 0.1$, by definition of ϵ ,

$$\left(\frac{4}{\pi} + \frac{1}{4}(g(\xi) - 2\xi^2) + \xi^2 \right) \frac{\sqrt{d-1}\epsilon}{\lambda} \leq 2 \frac{\sqrt{d-1}\epsilon}{\lambda} \leq \frac{1}{2}.$$

So, by Proposition 2, $L_{\theta_j}(\hat{F}) \geq \gamma$. Also, $\text{KL}(P_{\theta_i}, P_{\theta_0}) \leq (d-1)\xi^4 \frac{2\epsilon^2}{\lambda^2} \leq \frac{\log M}{9n}$ for all $1 \leq i \leq M$, because, by definition of ϵ , $\xi^4 \frac{2\epsilon^2}{\lambda^2} \leq \frac{\log 2}{72n}$. So by Proposition 1 and the fact that $\xi \leq 0.1$,

$$\inf_{\hat{F}_n} \max_{i \in [0..M]} \mathbb{E}_{\theta_i} L_{\theta_i}(\hat{F}_n) \geq 0.07\gamma \geq \frac{1}{500} \min \left\{ \frac{\sqrt{\log 2}}{3} \frac{\sigma^2}{\lambda^2} \sqrt{\frac{d-1}{n}}, \frac{1}{4} \right\}$$

and to complete the proof we use $\sup_{\theta \in \Theta_\lambda} \mathbb{E}_\theta L_\theta(\hat{F}_n) \geq \max_{i \in [0..M]} \mathbb{E}_{\theta_i} L_{\theta_i}(\hat{F}_n)$ for any $\hat{F}_n$. $\square$

Proof of Theorem 4. For simplicity, we state this construction for $\Theta_{\lambda, s+1}$, assuming $4 \leq s \leq \frac{d-1}{4}$.

Let $\xi = \frac{\lambda}{2\sigma}$, and define $\epsilon = \min \left\{ \sqrt{\frac{8}{45}} \frac{\sigma^2}{\lambda} \sqrt{\frac{1}{n} \log \left(\frac{d-1}{s} \right)}, \frac{1}{2} \frac{\lambda}{\sqrt{s}} \right\}$. Define $\lambda_0^2 = \lambda^2 - s\epsilon^2$. Let $\Omega = \{\omega \in \{0, 1\}^{d-1} : \|\omega\|_0 = s\}$. For $\omega = (\omega(1), \dots, \omega(d-1)) \in \Omega$, let $\mu_\omega = \lambda_0 e_d + \sum_{i=1}^{d-1} \omega(i)\epsilon e_i$ (where $\{e_i\}_{i=1}^d$ is the standard basis for $\mathbb{R}^d$). Let $\theta_\omega = (-\frac{\mu_\omega}{2}, \frac{\mu_\omega}{2}) \in \Theta_{\lambda, s}$. By Lemma 2, there exist $\omega_0, \dots, \omega_M \in \Omega$ such that $M \geq \left(\frac{d-1}{s} \right)^{s/5} - 1$ and $\rho(\omega_i, \omega_j) \geq \frac{s}{2}$ for all $0 \leq i < j \leq M$. The remainder of the proof is analogous to that of Theorem 2 with $\gamma = \frac{1}{4}(g(\xi) - \sqrt{2}\xi^2) \frac{\sqrt{s}\epsilon}{\lambda}$. $\square$

5 Proofs of the Upper Bounds

Propositions 5 and 6 below bound the error in estimating the mean and principal direction, and can be obtained using standard concentration bounds and a variant of the Davis-Kahan theorem. Proposition 7 relates these errors to the clustering loss. For the sparse case, Propositions 8 and 9 bound the added error induced by the support estimation procedure. See supplementary material for proof details.

Proposition 5. Let $\theta = (\mu_0 - \mu, \mu_0 + \mu)$ for some $\mu_0, \mu \in \mathbb{R}^d$ and $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P_\theta$. For any $\delta > 0$, we have $\|\mu_0 - \hat{\mu}_n\| \geq \sigma \sqrt{\frac{2 \max(d, 8 \log \frac{1}{\delta})}{n}} + \|\mu\| \sqrt{\frac{2 \log \frac{1}{\delta}}{n}}$ with probability at least $1 - 3\delta$.

Proposition 6. Let $\theta = (\mu_0 - \mu, \mu_0 + \mu)$ for some $\mu_0, \mu \in \mathbb{R}^d$ and $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P_\theta$ with $d > 1$ and $n \geq 4d$. Define $\cos \beta = |v_1(\sigma^2 I + \mu \mu^T)^T v_1(\widehat{\Sigma}_n)|$. For any $0 < \delta < \frac{d-1}{\sqrt{e}}$, if $\max\left(\frac{\sigma^2}{\|\mu\|^2}, \frac{\sigma}{\|\mu\|}\right) \sqrt{\frac{\max(d, 8 \log \frac{1}{\delta})}{n}} \leq \frac{1}{160}$, then with probability at least $1 - 12\delta - 2 \exp(-\frac{n}{20})$,

$$\sin \beta \leq 14 \max\left(\frac{\sigma^2}{\|\mu\|^2}, \frac{\sigma}{\|\mu\|}\right) \sqrt{d} \sqrt{\frac{10}{n} \log \frac{d}{\delta} \max\left(1, \frac{10}{n} \log \frac{d}{\delta}\right)}.$$

Proposition 7. Let $\theta = (\mu_0 - \mu, \mu_0 + \mu)$, and for some $x_0, v \in \mathbb{R}^d$ with $\|v\| = 1$, let $\widehat{F}(x) = 1$ if $x^T v \geq x_0^T v$, and 2 otherwise. Define $\cos \beta = |v^T \mu|/\|\mu\|$. If $|(x_0 - \mu_0)^T v| \leq \sigma \epsilon_1 + \|\mu\| \epsilon_2$ for some $\epsilon_1 \geq 0$ and $0 \leq \epsilon_2 \leq \frac{1}{4}$, and if $\sin \beta \leq \frac{1}{\sqrt{5}}$, then

$$L_\theta(\widehat{F}) \leq \exp\left\{-\frac{1}{2} \max\left(0, \frac{\|\mu\|}{2\sigma} - 2\epsilon_1\right)^2\right\} \left[2\epsilon_1 + \epsilon_2 \frac{\|\mu\|}{\sigma} + 2 \sin \beta \left(2 \sin \beta \frac{\|\mu\|}{\sigma} + 1\right)\right].$$

Proof. Let $r = \left|\frac{(x_0 - \mu_0)^T v}{\cos \beta}\right|$. Since the clustering loss is invariant to rotation and translation,

$$\begin{aligned} L_\theta(\widehat{F}) &\leq \frac{1}{2} \int_{-\infty}^{\infty} \frac{1}{\sigma} \phi\left(\frac{x}{\sigma}\right) \left[\Phi\left(\frac{\|\mu\| + |x| \tan \beta + r}{\sigma}\right) - \Phi\left(\frac{\|\mu\| - |x| \tan \beta - r}{\sigma}\right)\right] dx \\ &\leq \int_{-\infty}^{\infty} \phi(x) \left[\Phi\left(\frac{\|\mu\|}{\sigma}\right) - \Phi\left(\frac{\|\mu\| - r}{\sigma} - |x| \tan \beta\right)\right] dx. \end{aligned}$$

Since $\tan \beta \leq \frac{1}{2}$ and $\epsilon_2 \leq \frac{1}{4}$, we have $r \leq 2\sigma \epsilon_1 + 2\|\mu\| \epsilon_2$, and $\Phi\left(\frac{\|\mu\|}{\sigma}\right) - \Phi\left(\frac{\|\mu\| - r}{\sigma}\right) \leq 2\left(\epsilon_1 + \epsilon_2 \frac{\|\mu\|}{\sigma}\right) \phi\left(\max\left(0, \frac{\|\mu\|}{2\sigma} - 2\epsilon_1\right)\right)$. Defining $A = \left|\frac{\|\mu\| - r}{\sigma}\right|$,

$$\begin{aligned} \int_{-\infty}^{\infty} \phi(x) \left[\Phi\left(\frac{\|\mu\|}{\sigma}\right) - \Phi\left(\frac{\|\mu\| - r}{\sigma} - |x| \tan \beta\right)\right] dx &\leq 2 \int_0^{\infty} \int_{A - x \tan \beta}^A \phi(x) \phi(y) dy dx \\ &= 2 \int_{-A \sin \beta}^{\infty} \int_{A \cos \beta}^{A \cos \beta + (u + A \sin \beta) \tan \beta} \phi(u) \phi(v) du dv \leq 2\phi(A) \tan \beta (A \sin \beta + 1) \\ &\leq 2\phi\left(\max\left(0, \frac{\|\mu\|}{2\sigma} - 2\epsilon_1\right)\right) \tan \beta \left(\left(2 \frac{\|\mu\|}{\sigma} + 2\epsilon_1\right) \sin \beta + 1\right) \end{aligned}$$

where we used $u = x \cos \beta - y \sin \beta$ and $v = x \sin \beta + y \cos \beta$ in the second step. The bound now follows easily. $\square$

Proof of Theorem 1. Using Propositions 5 and 6 with $\delta = \frac{1}{\sqrt{n}}$, Proposition 7, and the fact that $(C + x) \exp(-\max(0, x - 4)^2/8) \leq (C + 6) \exp(-\max(0, x - 4)^2/10)$ for all $C, x > 0$,

$$\mathbb{E}_\theta L_\theta(\widehat{F}) \leq 600 \max\left(\frac{4\sigma^2}{\lambda^2}, 1\right) \sqrt{\frac{d \log(nd)}{n}}$$

(it is easy to verify that the bounds are decreasing with $\|\mu\|$, so we use $\|\mu\| = \frac{\lambda}{2}$ to bound the supremum). In the $d = 1$ case Proposition 6 need not be applied, since the principal directions agree trivially. The bound for $\frac{\lambda}{\sigma} \geq 2 \max(80, 14\sqrt{5d})$ can be shown similarly, using $\delta = \exp(-\frac{n}{32})$. $\square$

Proposition 8. Let $\theta = (\mu_0 - \mu, \mu_0 + \mu)$ for some $\mu_0, \mu \in \mathbb{R}^d$ and $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P_\theta$. For any $0 < \delta < \frac{1}{\sqrt{e}}$ such that $\sqrt{\frac{6 \log \frac{1}{\delta}}{n}} \leq \frac{1}{2}$, with probability at least $1 - 6d\delta$, for all $i \in [d]$,

$$|\widehat{\Sigma}_n(i, i) - (\sigma^2 + \mu(i)^2)| \leq \sigma^2 \sqrt{\frac{6 \log \frac{1}{\delta}}{n}} + 2\sigma |\mu(i)| \sqrt{\frac{2 \log \frac{1}{\delta}}{n}} + (\sigma + |\mu(i)|)^2 \frac{2 \log \frac{1}{\delta}}{n}.$$

Proposition 9. Let $\theta = (\mu_0 - \mu, \mu_0 + \mu)$ for some $\mu_0, \mu \in \mathbb{R}^d$ and $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P_\theta$. Define

$$S(\theta) = \{i \in [d] : \mu(i) \neq 0\} \quad \text{and} \quad \widetilde{S}(\theta) = \{i \in [d] : |\mu(i)| \geq 4\sigma\sqrt{\alpha}\}.$$

Assume that $n \geq 1$, $d \geq 2$, and $\alpha \leq \frac{1}{4}$. Then $\widetilde{S}(\theta) \subseteq \widehat{S}_n \subseteq S(\theta)$ with probability at least $1 - \frac{6}{n}$.

Proof. By Proposition 8, with probability at least $1 - \frac{6}{n}$,

$$|\widehat{\Sigma}_n(i, i) - (\sigma^2 + \mu(i)^2)| \leq \sigma^2 \sqrt{\frac{6 \log(nd)}{n}} + 2\sigma |\mu(i)| \sqrt{\frac{2 \log(nd)}{n}} + (\sigma + |\mu(i)|)^2 \frac{2 \log(nd)}{n}$$

for all $i \in [d]$. Assume the above event holds. If $S(\theta) = [d]$, then of course $\widehat{S}_n \subseteq S(\theta)$. Otherwise, for $i \notin S(\theta)$, we have $(1 - \alpha)\sigma^2 \leq \widehat{\Sigma}_n(i, i) \leq (1 + \alpha)\sigma^2$, so it is clear that $\widehat{S}_n \subseteq S(\theta)$. The remainder of the proof is trivial if $\widetilde{S}(\theta) = \emptyset$ or $S(\theta) = \emptyset$. Assume otherwise. For any $i \in S(\theta)$,

$$\widehat{\Sigma}_n(i, i) \geq (1 - \alpha)\sigma^2 + \left(1 - \frac{2 \log(nd)}{n}\right) \mu(i)^2 - 2\alpha\sigma |\mu(i)|.$$

By definition, $|\mu(i)| \geq 4\sigma\sqrt{\alpha}$ for all $i \in \widetilde{S}(\theta)$, so $\frac{(1+\alpha)^2}{1-\alpha}\sigma^2 \leq \widehat{\Sigma}_n(i, i)$ and $i \in \widehat{S}_n$ (we ignore strict equality above as a measure 0 event), i.e. $\widetilde{S}(\theta) \subseteq \widehat{S}_n$, which concludes the proof. $\square$

Proof of Theorem 3. Define $S(\theta) = \{i \in [d] : \mu(i) \neq 0\}$ and $\widetilde{S}(\theta) = \{i \in [d] : |\mu(i)| \geq 4\sigma\sqrt{\alpha}\}$. Assume $\widetilde{S}(\theta) \subseteq \widehat{S}_n \subseteq S(\theta)$ (by Proposition 9, this holds with probability at least $1 - \frac{6}{n}$). If $\widetilde{S}(\theta) = \emptyset$, then we simply have $\mathbb{E}_\theta L_\theta(\widehat{F}_n) \leq \frac{1}{2}$.

Assume $\widetilde{S}(\theta) \neq \emptyset$. Let $\cos \widehat{\beta} = |v_1(\widehat{\Sigma}_{\widehat{S}_n})^T v_1(\Sigma)|$, $\cos \widetilde{\beta} = |v_1(\Sigma_{\widehat{S}_n})^T v_1(\Sigma)|$, and $\cos \beta = |v_1(\widehat{\Sigma}_{\widehat{S}_n})^T v_1(\Sigma_{\widehat{S}_n})|$ where $\Sigma = \sigma^2 I + \mu\mu^T$, and for simplicity we define $\widehat{\Sigma}_{\widehat{S}_n}$ and $\Sigma_{\widehat{S}_n}$ to be the same as $\widehat{\Sigma}_n$ and Σ in $\widehat{S}_n$, respectively, and 0 elsewhere. Then $\sin \widehat{\beta} \leq \sin \widetilde{\beta} + \sin \beta$, and

$$\sin \widetilde{\beta} = \frac{\|\mu - \mu_{\widehat{S}(\theta)}\|}{\|\mu\|} \leq \frac{\|\mu - \mu_{\widetilde{S}(\theta)}\|}{\|\mu\|} \leq \frac{4\sigma\sqrt{\alpha}\sqrt{|S(\theta)| - |\widetilde{S}(\theta)|}}{\|\mu\|} \leq 8\frac{\sigma\sqrt{s\alpha}}{\lambda}.$$

Using the same argument as the proof of Theorem 1, as long as the above bound is smaller than $\frac{1}{2\sqrt{3}}$,

$$\mathbb{E}_\theta L_\theta(\widehat{F}) \leq 600 \max \left(\frac{\sigma^2}{\left(\frac{\lambda}{2} - 4\sigma\sqrt{s\alpha}\right)^2}, 1 \right) \sqrt{\frac{s \log(ns)}{n}} + 104 \frac{\sigma\sqrt{s\alpha}}{\lambda} + \frac{3}{n}.$$

Using the fact $L_\theta(\widehat{F}) \leq \frac{1}{2}$ always, and that $\alpha \leq \frac{1}{4}$ implies $\frac{\log(nd)}{n} \leq 1$, the bound follows. $\square$

6 Conclusion

We have provided minimax lower and upper bounds for estimating high dimensional mixtures. The bounds show explicitly how the statistical difficulty of the problem depends on dimension d , sample size n , separation λ and sparsity level s .

For clarity, we focused on the special case where there are two spherical components with equal mixture weights. In future work, we plan to extend the results to general mixtures of k Gaussians.

One of our motivations for this work is the recent interest in variable selection methods to facilitate clustering in high dimensional problems. Existing methods such as Pan and Shen (2007); Witten and Tibshirani (2010); Raftery and Dean (2006); Sun et al. (2012) and Guo et al. (2010) provide promising numerical evidence that variable selection does improve high dimensional clustering. Our results provide some theoretical basis for this idea.

However, there is a gap between the results in this paper and the above methodology papers. Indeed, as of now, there is no rigorous proof that the methods in those papers outperform a two stage approach where the first stage screens for relevant features and the second stage applies standard clustering methods on the features found in the first stage. We conjecture that there are conditions under which simultaneous feature selection and clustering outperforms a two stage method. Settling this question will require the aforementioned extension of our results to the general mixture case.

Acknowledgements

This research is supported in part by NSF grants IIS-1116458 and CAREER award IIS-1252412, as well as NSF Grant DMS-0806009 and Air Force Grant FA95500910373.

References

- Dimitris Achlioptas and Frank McSherry. On spectral learning of mixtures of distributions. In *Learning Theory*, pages 458–469. Springer, 2005.
- Sanjeev Arora and Ravi Kannan. Learning mixtures of arbitrary gaussians. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 247–257. ACM, 2001.
- Mikhail Belkin and Kaushik Sinha. Polynomial learning of distribution families. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 103–112. IEEE, 2010.
- S Charles Brubaker and Santosh S Vempala. Isotropic pca and affine-invariant clustering. In *Building Bridges*, pages 241–281. Springer, 2008.
- Kamalika Chaudhuri and Satish Rao. Learning mixtures of product distributions using correlations and independence. In *COLT*, pages 9–20, 2008.
- Kamalika Chaudhuri, Sanjoy Dasgupta, and Andrea Vattani. Learning mixtures of gaussians using the k-means algorithm. *arXiv preprint arXiv:0912.0086*, 2009.
- Sanjoy Dasgupta. Learning mixtures of gaussians. In *Foundations of Computer Science, 1999. 40th Annual Symposium on*, pages 634–644. IEEE, 1999.
- Jian Guo, Elizaveta Levina, George Michailidis, and Ji Zhu. Pairwise variable selection for high-dimensional model-based clustering. *Biometrics*, 66(3):793–804, 2010.
- Daniel Hsu and Sham M Kakade. Learning mixtures of spherical gaussians: moment methods and spectral decompositions. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 11–20. ACM, 2013.
- Adam Tauman Kalai, Ankur Moitra, and Gregory Valiant. Disentangling gaussians. *Communications of the ACM*, 55(2):113–120, 2012.
- Ravindran Kannan, Hadi Salmasian, and Santosh Vempala. The spectral method for general mixture models. In *Learning Theory*, pages 444–457. Springer, 2005.
- Pascal Massart. Concentration inequalities and model selection. 2007.
- Wei Pan and Xiaotong Shen. Penalized model-based clustering with application to variable selection. *The Journal of Machine Learning Research*, 8:1145–1164, 2007.
- Adrian E Raftery and Nema Dean. Variable selection for model-based clustering. *Journal of the American Statistical Association*, 101(473):168–178, 2006.
- Leonard J. Schulman and Sanjoy Dasgupta. A two-round variant of em for gaussian mixtures. In *Proc. 16th UAI (Conference on Uncertainty in Artificial Intelligence)*, pages 152–159, 2000.
- Wei Sun, Junhui Wang, and Yixin Fang. Regularized k-means clustering of high-dimensional data and its asymptotic consistency. *Electronic Journal of Statistics*, 6:148–167, 2012.
- Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer, 2009.
- Santosh Vempala and Grant Wang. A spectral algorithm for learning mixture models. *Journal of Computer and System Sciences*, 68(4):841–860, 2004.
- Vincent Q Vu and Jing Lei. Minimax sparse principal subspace estimation in high dimensions. *arXiv preprint arXiv:1211.0373*, 2012.
- Daniela M Witten and Robert Tibshirani. A framework for feature selection in clustering. *Journal of the American Statistical Association*, 105(490), 2010.

Supplementary file for NIPS submission: Minimax Theory for High-dimensional Gaussian Mixtures with Sparse Mean Separation

Martin Azizyan
Machine Learning Department
Carnegie Mellon University
Pittsburgh, PA 15213
mazizyan@cs.cmu.edu

Aarti Singh
Machine Learning Department
Carnegie Mellon University
Pittsburgh, PA 15213
aarti@cs.cmu.edu

Larry Wasserman
Department of Statistics
Carnegie Mellon University
Pittsburgh, PA 15213
larry@stat.cmu.edu

1 Notation

For $\theta = (\mu_1, \mu_2) \in \mathbb{R}^{2 \times d}$, define

$$p_\theta(x) = \frac{1}{2}f(x; \mu_1, \sigma^2 I) + \frac{1}{2}f(x; \mu_2, \sigma^2 I),$$

where $f(\cdot; \mu, \Sigma)$ is the density of $\mathcal{N}(\mu, \Sigma)$, $\sigma > 0$ is a fixed constant. Let P_θ denote the probability measure corresponding to p_θ . We consider two classes Θ of parameters:

$$\begin{aligned}\Theta_\lambda &= \{(\mu_1, \mu_2) : \|\mu_1 - \mu_2\| \geq \lambda\} \\ \Theta_{\lambda,s} &= \{(\mu_1, \mu_2) : \|\mu_1 - \mu_2\| \geq \lambda, \|\mu_1 - \mu_2\|_0 \leq s\} \subseteq \Theta_\lambda.\end{aligned}$$

Throughout this document, ϕ and Φ denote the standard normal density and distribution functions.

For a mixture with parameter θ , the Bayes optimal classification, that is, assignment of a point $x \in \mathbb{R}^d$ to the correct mixture component, is given by the function

$$F_\theta(x) = \operatorname{argmax}_{i \in \{1,2\}} f(x; \mu_i, \sigma^2 I).$$

Given any other candidate assignment function $F : \mathbb{R}^d \rightarrow \{1, 2\}$, we define the loss incurred by F as

$$L_\theta(F) = \min_{\pi} P_\theta(\{x : F_\theta(x) \neq \pi(F(x))\})$$

where the minimum is over all permutations $\pi : \{1, 2\} \rightarrow \{1, 2\}$.

For $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} P_\theta$, let $\hat{\mu}_n$ and $\hat{\Sigma}_n$ be the mean and covariance of the corresponding empirical distribution.

Also, for a matrix B , $v_i(B)$ and $\lambda_i(B)$ are the i 'th eigenvector and eigenvalue of B (assuming B is symmetric), arranged so that $\lambda_i(B) \geq \lambda_{i+1}(B)$, and $\|B\|_2$ is the spectral norm.

2 Upper bounds

2.1 Standard concentration bounds

2.1.1 Concentration bounds for estimating the mean

Proposition 1. *Let $X \sim \chi_d^2$. Then for any $\epsilon > 0$,*

$$\mathbb{P}(X > (1 + \epsilon)d) \leq \exp \left\{ -\frac{d}{2} (\epsilon - \log(1 + \epsilon)) \right\}.$$

If $\epsilon < 1$, then

$$\mathbb{P}(X < (1 - \epsilon)d) \leq \exp \left\{ \frac{d}{2}(\epsilon + \log(1 - \epsilon)) \right\}.$$

Proof. Since $\mathbb{E}e^{tX} = (1 - 2t)^{-\frac{d}{2}}$ for $0 < t < \frac{1}{2}$,

$$\begin{aligned} \mathbb{P}(X > (1 + \epsilon)d) &= \mathbb{P}(e^{tX} > e^{t(1+\epsilon)d}) \\ &\leq e^{-t(1+\epsilon)d} (1 - 2t)^{-\frac{d}{2}} \\ &= \exp \left[-t(1 + \epsilon)d + \frac{d}{2} \log \frac{1}{1 - 2t} \right]. \end{aligned}$$

To minimize the right hand side, we differentiate the exponent with respect to t to obtain the equation

$$-(1 + \epsilon)d + \frac{d}{1 - 2t} = 0$$

which can be satisfied by setting $t = \frac{1}{2} \left(1 - \frac{1}{1+\epsilon} \right) < \frac{1}{2}$ (it is easy to verify that this is a global minimum). Using this value for t , the first bound follows.

Also, for $t > 0$ and $\epsilon < 1$,

$$\begin{aligned} \mathbb{P}(X < (1 - \epsilon)d) &= \mathbb{P}(e^{-tX} > e^{-t(1-\epsilon)d}) \\ &\leq e^{t(1-\epsilon)d} (1 + 2t)^{-\frac{d}{2}} \\ &= \exp \left[t(1 - \epsilon)d - \frac{d}{2} \log(1 + 2t) \right] \end{aligned}$$

and setting $t = \frac{1}{2} \left(\frac{1}{1-\epsilon} - 1 \right)$,

$$\mathbb{P}(X < (1 - \epsilon)d) \leq \exp \left[\frac{d}{2}(\epsilon + \log(1 - \epsilon)) \right].$$

□

Proposition 2. Let $Z_1, \dots, Z_n \stackrel{i.i.d.}{\sim} \mathcal{N}(0, I_d)$. Then for any $\epsilon > 0$,

$$\mathbb{P} \left(\left\| \frac{1}{n} \sum_{i=1}^n Z_i \right\| \geq \sqrt{\frac{(1 + \epsilon)d}{n}} \right) \leq \exp \left\{ -\frac{d}{2} (\epsilon - \log(1 + \epsilon)) \right\}.$$

Proof. Using Proposition 1,

$$\begin{aligned} \mathbb{P} \left(\left\| \frac{1}{n} \sum_{i=1}^n Z_i \right\| \geq \sqrt{\frac{(1 + \epsilon)d}{n}} \right) &= \mathbb{P} \left(\left\| \frac{1}{\sqrt{n}} \sum_{i=1}^n Z_i \right\|^2 \geq (1 + \epsilon)d \right) \\ &= \mathbb{P}(X \geq (1 + \epsilon)d) \\ &\leq \exp \left\{ -\frac{d}{2} (\epsilon - \log(1 + \epsilon)) \right\} \end{aligned}$$

where $X \sim \chi_d^2$.

□

2.1.2 Concentration bounds for estimating principal direction

Proposition 3. Let $Z_1, \dots, Z_n \stackrel{i.i.d.}{\sim} \mathcal{N}(0, I_d)$ and $\delta > 0$. If $n \geq d$ then with probability at least $1 - 3\delta$,

$$\begin{aligned} \|\widehat{\Sigma}_n - I_d\|_2 &\leq 3 \left(1 + \sqrt{\frac{2 \log \frac{1}{\delta}}{d}} \right) \sqrt{\frac{d}{n}} \max \left(1, \left(1 + \sqrt{\frac{2 \log \frac{1}{\delta}}{d}} \right) \sqrt{\frac{d}{n}} \right) \\ &\quad + \left(1 + \sqrt{\frac{8 \log \frac{1}{\delta}}{d}} \max \left(1, \frac{8 \log \frac{1}{\delta}}{d} \right) \right) \frac{d}{n} \end{aligned}$$

where $\widehat{\Sigma}_n$ is the empirical covariance of Z_i .

Proof. Let $\bar{Z}_n = \frac{1}{n} \sum_{i=1}^n Z_i$. Then

$$\begin{aligned} \|\hat{\Sigma}_n - I_d\|_2 &= \left\| \frac{1}{n} \sum_{i=1}^n Z_i Z_i^T - I_d - \bar{Z}_n \bar{Z}_n^T \right\|_2 \\ &\leq \left\| \frac{1}{n} \sum_{i=1}^n Z_i Z_i^T - I_d \right\|_2 + \|\bar{Z}_n\|^2. \end{aligned}$$

It is well known that for any $\epsilon_1 > 0$,

$$\mathbb{P} \left(\left\| \frac{1}{n} \sum_{i=1}^n Z_i Z_i^T - I_d \right\|_2 \geq 3(1 + \epsilon_1) \sqrt{\frac{d}{n}} \max \left(1, (1 + \epsilon_1) \sqrt{\frac{d}{n}} \right) \right) \leq 2 \exp \left\{ -\frac{d\epsilon_1^2}{2} \right\}.$$

Using this along with Proposition 2, we have for any $\epsilon_2 > 0$,

$$\begin{aligned} \mathbb{P} \left(\|\hat{\Sigma}_n - I_d\|_2 \geq 3(1 + \epsilon_1) \sqrt{\frac{d}{n}} \max \left(1, (1 + \epsilon_1) \sqrt{\frac{d}{n}} \right) + \frac{(1 + \epsilon_2)d}{n} \right) \\ \leq 2 \exp \left\{ -\frac{d\epsilon_1^2}{2} \right\} + \exp \left\{ -\frac{d}{2} (\epsilon_2 - \log(1 + \epsilon_2)) \right\}. \end{aligned}$$

Setting $\epsilon_1 = \sqrt{\frac{2 \log \frac{1}{\delta}}{d}}$,

$$\begin{aligned} \mathbb{P} \left(\|\hat{\Sigma}_n - I_d\|_2 \geq 3 \left(1 + \sqrt{\frac{2 \log \frac{1}{\delta}}{d}} \right) \sqrt{\frac{d}{n}} \max \left(1, \left(1 + \sqrt{\frac{2 \log \frac{1}{\delta}}{d}} \right) \sqrt{\frac{d}{n}} \right) + \frac{(1 + \epsilon_2)d}{n} \right) \\ \leq 2\delta + \exp \left\{ -\frac{d}{2} (\epsilon_2 - \log(1 + \epsilon_2)) \right\} \\ \leq 2\delta + \exp \left\{ -\frac{d}{8} \epsilon_2 \min(1, \epsilon_2) \right\} \end{aligned}$$

and, setting $\epsilon_2 = \sqrt{\frac{8 \log \frac{1}{\delta}}{d} \max \left(1, \frac{8 \log \frac{1}{\delta}}{d} \right)}$, with probability at least $1 - 3\delta$,

$$\begin{aligned} \|\hat{\Sigma}_n - I_d\|_2 &\leq 3 \left(1 + \sqrt{\frac{2 \log \frac{1}{\delta}}{d}} \right) \sqrt{\frac{d}{n}} \max \left(1, \left(1 + \sqrt{\frac{2 \log \frac{1}{\delta}}{d}} \right) \sqrt{\frac{d}{n}} \right) \\ &\quad + \left(1 + \sqrt{\frac{8 \log \frac{1}{\delta}}{d} \max \left(1, \frac{8 \log \frac{1}{\delta}}{d} \right)} \right) \frac{d}{n}. \end{aligned}$$

□

Proposition 4. Let $X_1, Y_1, \dots, X_n, Y_n \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$. Then for any $\epsilon > 0$,

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n X_i Y_i \right| > \frac{\epsilon}{2} \right) \leq 2 \exp \left\{ -\frac{n\epsilon \min(1, \epsilon)}{10} \right\}.$$

Proof. Let $Z = XY$ where $X, Y \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$. Then for any t such that $|t| < 1$,

$$\mathbb{E} e^{tZ} = \frac{1}{\sqrt{1 - t^2}}.$$

So for $0 < t < 1$,

$$\begin{aligned}
\mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n X_i Y_i > \epsilon\right) &= \mathbb{P}\left(\exp\left\{\sum_{i=1}^n t X_i Y_i\right\} > \exp(n\epsilon t)\right) \\
&\leq \mathbb{E}\left(\exp\left\{\sum_{i=1}^n t X_i Y_i\right\}\right) \exp(-n\epsilon t) \\
&= (\mathbb{E} \exp(t X_i Y_i))^n \exp(-n\epsilon t) \\
&= (1 - t^2)^{-\frac{n}{2}} \exp(-n\epsilon t) \\
&= \exp\left\{-\frac{n}{2} (2\epsilon t + \log(1 - t^2))\right\}.
\end{aligned}$$

The bound is minimized by $t = \frac{1}{2\epsilon} (\sqrt{1 + 4\epsilon^2} - 1) < 1$, so

$$\mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n X_i Y_i > \epsilon\right) \leq \exp\left\{-\frac{n}{2} h(2\epsilon)\right\}$$

where

$$h(u) = \left(\sqrt{1 + u^2} - 1\right) + \log\left(1 - \frac{1}{u^2} \left(\sqrt{1 + u^2} - 1\right)^2\right).$$

Since $h(u) \geq \frac{u}{5} \min(1, u)$,

$$\mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n X_i Y_i > \epsilon\right) \leq \exp\left\{-\frac{n}{2} \frac{2\epsilon}{5} \min(1, 2\epsilon)\right\}$$

and the proof is complete by noting that the distribution of $X_i Y_i$ is symmetric. $\square$

2.2 Davis–Kahan

Lemma 1. Let $A, E \in \mathbb{R}^{d \times d}$ be symmetric matrices, and $u \in \mathbb{R}^{d-1}$ such that

$$u_i = v_{i+1}(A)^T E v_1(A).$$

If $\lambda_1(A) - \lambda_2(A) > 0$ and

$$\|E\|_2 \leq \frac{\lambda_1(A) - \lambda_2(A)}{5}$$

then

$$\sqrt{1 - (v_1(A)^T v_1(A + E))^2} \leq \frac{4\|u\|}{\lambda_1(A) - \lambda_2(A)}$$

(Corollary 8.1.11 of Golub and Van Loan (1996)).

2.3 Bounding error in estimating the mean

Proposition 5. Let $\theta = (\mu_0 - \mu, \mu_0 + \mu)$ for some $\mu_0, \mu \in \mathbb{R}^d$ and $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P_\theta$. For any $\delta > 0$,

$$\mathbb{P}\left(\|\mu_0 - \hat{\mu}_n\| \geq \sigma \sqrt{\frac{2 \max(d, 8 \log \frac{1}{\delta})}{n}} + \|\mu\| \sqrt{\frac{2 \log \frac{1}{\delta}}{n}}\right) \leq 3\delta.$$

Proof. Let $Z_1, \dots, Z_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I)$ and $Y_1, \dots, Y_n$ i.i.d. such that $\mathbb{P}(Y_i = -1) = \mathbb{P}(Y_i = 1) = \frac{1}{2}$. Then for any $\epsilon_1, \epsilon_2 > 0$,

$$\begin{aligned}
& \mathbb{P} \left(\|\mu_0 - \hat{\mu}_n\| \geq \sigma \sqrt{\frac{(1 + \epsilon_1)d}{n}} + \|\mu\| \epsilon_2 \right) \\
&= \mathbb{P} \left(\left\| \mu_0 - \frac{1}{n} \sum_{i=1}^d (\sigma Z_i + \mu_0 + \mu Y_i) \right\| \geq \sigma \sqrt{\frac{(1 + \epsilon_1)d}{n}} + \|\mu\| \epsilon_2 \right) \\
&= \mathbb{P} \left(\left\| \sigma \frac{1}{n} \sum_{i=1}^d Z_i + \mu \frac{1}{n} \sum_{i=1}^d Y_i \right\| \geq \sigma \sqrt{\frac{(1 + \epsilon_1)d}{n}} + \|\mu\| \epsilon_2 \right) \\
&\leq \mathbb{P} \left(\sigma \left\| \frac{1}{n} \sum_{i=1}^d Z_i \right\| + \|\mu\| \left| \frac{1}{n} \sum_{i=1}^d Y_i \right| \geq \sigma \sqrt{\frac{(1 + \epsilon_1)d}{n}} + \|\mu\| \epsilon_2 \right) \\
&\leq \mathbb{P} \left(\left\| \frac{1}{n} \sum_{i=1}^d Z_i \right\| \geq \sqrt{\frac{(1 + \epsilon_1)d}{n}} \right) + \mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^d Y_i \right| \geq \epsilon_2 \right) \\
&\leq \exp \left\{ -\frac{d}{2} (\epsilon_1 - \log(1 + \epsilon_1)) \right\} + 2 \exp \left\{ -\frac{n\epsilon_2^2}{2} \right\}
\end{aligned}$$

where the last step is using Hoeffding's inequality and Proposition 2. Setting $\epsilon_2 = \sqrt{\frac{2 \log \frac{1}{\delta}}{n}}$,

$$\begin{aligned}
& \mathbb{P} \left(\|\mu_0 - \hat{\mu}_n\| \geq \sigma \sqrt{\frac{(1 + \epsilon_1)d}{n}} + \|\mu\| \sqrt{\frac{2 \log \frac{1}{\delta}}{n}} \right) \\
&\leq \exp \left\{ -\frac{d}{2} (\epsilon_1 - \log(1 + \epsilon_1)) \right\} + 2\delta.
\end{aligned}$$

Since $\epsilon_1 - \log(1 + \epsilon_1) \geq \frac{\epsilon_1}{4} \min(1, \epsilon_1)$,

$$\exp \left\{ -\frac{d}{2} (\epsilon_1 - \log(1 + \epsilon_1)) \right\} \leq \exp \left\{ -\frac{d}{8} \epsilon_1 \min(1, \epsilon_1) \right\}.$$

Setting

$$\epsilon_1 = \sqrt{\frac{8 \log \frac{1}{\delta}}{d} \max \left(1, \frac{8 \log \frac{1}{\delta}}{d} \right)},$$

we have

$$\mathbb{P} \left(\|\mu_0 - \hat{\mu}_n\| \geq \sigma \sqrt{\frac{d}{n} \left(1 + \sqrt{\frac{8 \log \frac{1}{\delta}}{d} \max \left(1, \frac{8 \log \frac{1}{\delta}}{d} \right)} \right)} + \|\mu\| \sqrt{\frac{2 \log \frac{1}{\delta}}{n}} \right) \leq 3\delta$$

and the bound follows. $\square$

2.4 Bounding error in estimating principal direction

Proposition 6. Let $\theta = (\mu_0 - \mu, \mu_0 + \mu)$ for some $\mu_0, \mu \in \mathbb{R}^d$ and $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P_\theta$. If $n \geq d$ then for any $\delta, \delta_1 > 0$, with probability at least $1 - 5\delta - 2\delta_1$,

$$\begin{aligned} & \|\widehat{\Sigma}_n - (\sigma^2 I_d + \mu\mu^T)\|_2 \\ & \leq 3\sigma^2 \left(1 + \sqrt{\frac{2 \log \frac{1}{\delta}}{d}}\right) \sqrt{\frac{d}{n}} \max \left(1, \left(1 + \sqrt{\frac{2 \log \frac{1}{\delta}}{d}}\right) \sqrt{\frac{d}{n}}\right) \\ & \quad + \sigma^2 \left(1 + \sqrt{\frac{8 \log \frac{1}{\delta}}{d} \max \left(1, \frac{8 \log \frac{1}{\delta}}{d}\right)}\right) \frac{d}{n} \\ & \quad + 4\sigma\|\mu\| \sqrt{\left(1 + \sqrt{\frac{8 \log \frac{1}{\delta}}{d} \max \left(1, \frac{8 \log \frac{1}{\delta}}{d}\right)}\right) \frac{d}{n} + \frac{2\|\mu\|^2 \log \frac{1}{\delta_1}}{n}} \end{aligned}$$

where $\widehat{\Sigma}_n$ is the empirical covariance of X_i .

Proof. We can express X_i as $X_i = \sigma Z_i + \mu Y_i + \mu_0$ where $Z_1, \dots, Z_n \stackrel{i.i.d.}{\sim} \mathcal{N}(0, I_d)$ and $Y_1, \dots, Y_n$ i.i.d. such that $\mathbb{P}(Y_i = -1) = \mathbb{P}(Y_i = 1) = \frac{1}{2}$. Then

$$\begin{aligned} \widehat{\Sigma}_n - (\sigma^2 I_d + \mu\mu^T) &= \sigma^2 (\widehat{\Sigma}_n^Z - I_d) - \mu\mu^T \bar{Y}^2 \\ & \quad + \sigma \left(\frac{1}{n} \sum_{i=1}^n Y_i Z_i - \bar{Y} \bar{Z} \right) \mu^T \\ & \quad + \sigma \mu \left(\frac{1}{n} \sum_{i=1}^n Y_i Z_i - \bar{Y} \bar{Z} \right)^T \end{aligned}$$

where $\widehat{\Sigma}_n^Z$ is the empirical covariance of Z_i and $\bar{Y}$ and $\bar{Z}$ are the empirical means of Y_i and Z_i . So

$$\begin{aligned} \|\widehat{\Sigma}_n - (\sigma^2 I_d + \mu\mu^T)\|_2 &\leq \sigma^2 \|\widehat{\Sigma}_n^Z - I_d\|_2 + \|\mu\|^2 \bar{Y}^2 \\ & \quad + 2\sigma\|\mu\| \left(\left\| \frac{1}{n} \sum_{i=1}^n Y_i Z_i \right\| + |\bar{Y}| \|\bar{Z}\| \right). \end{aligned}$$

By Hoeffding's inequality,

$$\mathbb{P} \left(\|\mu\|^2 \bar{Y}^2 \geq \frac{2\|\mu\|^2 \log \frac{1}{\delta_1}}{n} \right) \leq 2\delta_1.$$

Since $|\bar{Y}| \leq 1$ and since $Y_i Z_i$ has the same distribution as Z_i , by Proposition 2, for any $\epsilon > 0$,

$$\begin{aligned} & \mathbb{P} \left(2\sigma\|\mu\| \left(\left\| \frac{1}{n} \sum_{i=1}^n Y_i Z_i \right\| + |\bar{Y}| \|\bar{Z}\| \right) \geq 4\sigma\|\mu\| \sqrt{\frac{(1+\epsilon)d}{n}} \right) \\ & \leq 2 \exp \left\{ -\frac{d}{2} (\epsilon - \log(1+\epsilon)) \right\} \leq 2 \exp \left\{ -\frac{d}{8} \epsilon \min(1, \epsilon) \right\}. \end{aligned}$$

Setting

$$\epsilon = \sqrt{\frac{8 \log \frac{1}{\delta}}{d} \max \left(1, \frac{8 \log \frac{1}{\delta}}{d}\right)}$$

we have

$$\begin{aligned} & \mathbb{P} \left(2\sigma\|\mu\| \left(\left\| \frac{1}{n} \sum_{i=1}^n Y_i Z_i \right\| + |\bar{Y}| \|\bar{Z}\| \right) \geq 4\sigma\|\mu\| \sqrt{\left(1 + \sqrt{\frac{8 \log \frac{1}{\delta}}{d} \max \left(1, \frac{8 \log \frac{1}{\delta}}{d}\right)}\right) \frac{d}{n}} \right) \\ & \leq 2\delta. \end{aligned}$$

Finally, by Proposition 3, with probability at least $1 - 3\delta$,

$$\begin{aligned} \sigma^2 \|\widehat{\Sigma}_n^Z - I_d\|_2 &\leq 3\sigma^2 \left(1 + \sqrt{\frac{2 \log \frac{1}{\delta}}{d}} \right) \sqrt{\frac{d}{n}} \max \left(1, \left(1 + \sqrt{\frac{2 \log \frac{1}{\delta}}{d}} \right) \sqrt{\frac{d}{n}} \right) \\ &\quad + \sigma^2 \left(1 + \sqrt{\frac{8 \log \frac{1}{\delta}}{d} \max \left(1, \frac{8 \log \frac{1}{\delta}}{d} \right)} \right) \frac{d}{n} \end{aligned}$$

and we complete the proof by combining the three bounds. $\square$

Proposition 7. Let $\theta = (\mu_0 - \mu, \mu_0 + \mu)$ for some $\mu_0, \mu \in \mathbb{R}^d$ and $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P_\theta$. If $n \geq d > 1$ then for any $0 < \delta \leq \frac{1}{\sqrt{e}}$ and $i \in [2..d]$, with probability at least $1 - 7\delta$,

$$\begin{aligned} &\left| v_i(\sigma^2 I + \mu\mu^T)^T (\widehat{\Sigma}_n - (\sigma^2 I + \mu\mu^T)) v_i(\sigma^2 I + \mu\mu^T) \right| \\ &\leq \sigma^2 \frac{1}{2} \sqrt{\frac{10 \log \frac{1}{\delta}}{n} \max \left(1, \frac{10 \log \frac{1}{\delta}}{n} \right)} + \sigma \|\mu\| \sqrt{\frac{2 \log \frac{1}{\delta}}{n}} + (\sigma^2 + \sigma \|\mu\|) \frac{2 \log \frac{1}{\delta}}{n}. \end{aligned}$$

Proof. Let $Z_1, W_1, \dots, Z_n, W_n \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$ and $Y_1, \dots, Y_n$ i.i.d. such that $\mathbb{P}(Y_i = -1) = \mathbb{P}(Y_i = 1) = \frac{1}{2}$. It is easy to see that the quantity of interest is equal in distribution to

$$\left| \frac{1}{n} \sum_{j=1}^n (\sigma Z_i - \sigma \bar{Z})(\sigma W_i - \sigma \bar{W} + \|\mu\| Y_i - \|\mu\| \bar{Y}) \right|$$

where $\bar{Z}, \bar{W}, \bar{Y}$ are the respective empirical means. Moreover,

$$\begin{aligned} &\left| \frac{1}{n} \sum_{j=1}^n (\sigma Z_i - \sigma \bar{Z})(\sigma W_i - \sigma \bar{W} + \|\mu\| Y_i - \|\mu\| \bar{Y}) \right| \\ &\leq \sigma^2 \left| \frac{1}{n} \sum_{j=1}^n Z_i W_i \right| + \sigma^2 |\bar{Z}| |\bar{W}| + \sigma \|\mu\| \left| \frac{1}{n} \sum_{j=1}^n Z_i Y_i \right| + \sigma \|\mu\| |\bar{Z}| |\bar{Y}|. \end{aligned}$$

From Proposition 4, we have

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n Z_i W_i \right| > \frac{1}{2} \sqrt{\frac{10 \log \frac{1}{\delta}}{n} \max \left(1, \frac{10 \log \frac{1}{\delta}}{n} \right)} \right) \leq 2\delta;$$

using Hoeffding's inequality,

$$\mathbb{P} \left(|\bar{Y}| \geq \sqrt{\frac{2 \log \frac{1}{\delta}}{n}} \right) \leq 2\delta;$$

and using the Gaussian tail bound, for $\delta \leq \frac{1}{\sqrt{e}}$,

$$\mathbb{P} \left(|\bar{Z}| \geq \sqrt{\frac{2 \log \frac{1}{\delta}}{n}} \right) \leq \delta$$

and the final result follows easily. $\square$

Proposition 8. Let $\theta = (\mu_0 - \mu, \mu_0 + \mu)$ for some $\mu_0, \mu \in \mathbb{R}^d$ and $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P_\theta$ with $d > 1$ and $n \geq 4d$. For any $0 < \delta < \frac{d-1}{\sqrt{e}}$, if

$$\max \left(\frac{\sigma^2}{\|\mu\|^2}, \frac{\sigma}{\|\mu\|} \right) \sqrt{\frac{\max(d, 8 \log \frac{1}{\delta})}{n}} \leq \frac{1}{160}$$

then with probability at least $1 - 12\delta - 2 \exp(-\frac{n}{20})$,

$$\sqrt{1 - (v_1(\sigma^2 I + \mu\mu^T)^T v_1(\widehat{\Sigma}_n))^2} \leq 14 \max\left(\frac{\sigma^2}{\|\mu\|^2}, \frac{\sigma}{\|\mu\|}\right) \sqrt{d} \sqrt{\frac{10 \log \frac{d}{\delta}}{n} \max\left(1, \frac{10 \log \frac{d}{\delta}}{n}\right)}.$$

Proof. By Proposition 6 (with $\delta_1 = \exp(-\frac{n}{20})$), Proposition 7 (with $\delta_2 = \frac{\delta}{d-1}$), and Lemma 1, with probability at least $1 - 12\delta - 2 \exp(-\frac{n}{20})$,

$$\begin{aligned} & \sqrt{1 - (v_1(\sigma^2 I + \mu\mu^T)^T v_1(\widehat{\Sigma}_n))^2} \\ & \leq \frac{4\sqrt{d-1}}{\|\mu\|^2} \left[\sigma^2 \frac{1}{2} \sqrt{\frac{10 \log \frac{d-1}{\delta}}{n} \max\left(1, \frac{10 \log \frac{d-1}{\delta}}{n}\right)} + \sigma \|\mu\| \sqrt{\frac{2 \log \frac{d-1}{\delta}}{n}} + (\sigma^2 + \sigma \|\mu\|) \frac{2 \log \frac{d-1}{\delta}}{n} \right] \end{aligned}$$

and the result follows after some simplifications. $\square$

2.5 General result relating error in estimating mean and principal direction to clustering loss

Proposition 9. Let $\theta = (\mu_0 - \mu, \mu_0 + \mu)$ and let

$$\widehat{F}(x) = \begin{cases} 1 & \text{if } x^T v \geq x_0^T v \\ 2 & \text{otherwise} \end{cases}$$

for some $x_0, v \in \mathbb{R}^d$, with $\|v\| = 1$. Define $\cos \beta = |v^T \mu| / \|\mu\|$. If $|(x_0 - \mu_0)^T v| \leq \sigma \epsilon_1 + \|\mu\| \epsilon_2$ for some $\epsilon_1 \geq 0$ and $0 \leq \epsilon_2 \leq \frac{1}{4}$, and if $\sin \beta \leq \frac{1}{\sqrt{5}}$, then

$$L_\theta(\widehat{F}) \leq \exp\left\{-\frac{1}{2} \max\left(0, \frac{\|\mu\|}{2\sigma} - 2\epsilon_1\right)^2\right\} \left[2\epsilon_1 + \epsilon_2 \frac{\|\mu\|}{\sigma} + 2 \sin \beta \left(2 \sin \beta \frac{\|\mu\|}{\sigma} + 1\right)\right].$$

Proof.

$$\begin{aligned} L_\theta(\widehat{F}) &= \min_{\pi} P_\theta(\{x : F_\theta(x) \neq \pi(\widehat{F}(x))\}) \\ &= \min \{P_\theta[\{x : ((x - \mu_0)^T \mu)((x - x_0)^T v) \geq 0\}], \\ &\quad P_\theta[\{x : ((x - \mu_0)^T \mu)((x - x_0)^T v) \leq 0\}]\}. \end{aligned}$$

WLOG assume $v^T \mu \geq 0$ (otherwise we can simply replace v with $-v$, which does not affect the bound). Then

$$\begin{aligned} L_\theta(\widehat{F}) &= P_\theta[\{x : ((x - \mu_0)^T \mu)((x - x_0)^T v) \leq 0\}] \\ &= P_\theta[\{x : ((x - \mu_0)^T \mu)((x - \mu_0)^T v - (x_0 - \mu_0)^T v) \leq 0\}] \\ &= P_\theta\left[\left\{x : \left((x - \mu_0)^T \frac{\mu}{\|\mu\|}\right)((x - \mu_0)^T v - (x_0 - \mu_0)^T v) \leq 0\right\}\right]. \end{aligned}$$

Define

$$\begin{aligned} \check{\mu} &= \frac{\mu}{\|\mu\|}, \\ \check{x} &= (x - \mu_0)^T \check{\mu}, \end{aligned}$$

and

$$\check{y} = (x - \mu_0)^T \frac{v - \check{\mu}\check{\mu}^T v}{\|v - \check{\mu}\check{\mu}^T v\|} \equiv (x - \mu_0)^T \frac{v - \check{\mu}\check{\mu}^T v}{\sin \beta}$$

so that

$$\begin{aligned} L_\theta(\widehat{F}) &= P_\theta[\{x : \check{x}(\check{y} \sin \beta + \check{x} \cos \beta - (x_0 - \mu_0)^T v) \leq 0\}] \\ &= P_\theta[\{x : \min(0, B(\check{y})) \leq \check{x} \leq \max(0, B(\check{y}))\}] \end{aligned}$$

where

$$B(\check{y}) = \frac{(x_0 - \mu_0)^T v}{\cos \beta} - \check{y} \tan \beta.$$

Since $\check{x}$ and $\check{y}$ are projections of $x - \mu_0$ onto orthogonal unit vectors, and since $\check{x}$ is exactly the component of $x - \mu_0$ that lies in the direction of μ , we can integrate out all other directions and obtain

$$L_\theta(\widehat{F}) = \int_{-\infty}^{\infty} \phi_\sigma(\check{y}) \int_{\min(0, B(\check{y}))}^{\max(0, B(\check{y}))} \left(\frac{1}{2} \phi_\sigma(\check{x} + \|\mu\|) + \frac{1}{2} \phi_\sigma(\check{x} - \|\mu\|) \right) d\check{x} d\check{y}$$

where ϕ_σ is the density of $\mathcal{N}(0, \sigma^2)$. But,

$$\begin{aligned} & \int_{\min(0, B(\check{y}))}^{\max(0, B(\check{y}))} \left(\frac{1}{2} \phi_\sigma(\check{x} + \|\mu\|) + \frac{1}{2} \phi_\sigma(\check{x} - \|\mu\|) \right) d\check{x} \\ &= \frac{1}{2} \int_{\min(0, B(\check{y}))}^{\max(0, B(\check{y}))} \phi_\sigma(\check{x} + \|\mu\|) d\check{x} + \frac{1}{2} \int_{\min(0, B(\check{y}))}^{\max(0, B(\check{y}))} \phi_\sigma(\check{x} - \|\mu\|) d\check{x} \\ &= \frac{1}{2} \left(\Phi \left(\frac{\max(0, B(\check{y})) + \|\mu\|}{\sigma} \right) - \Phi \left(\frac{\min(0, B(\check{y})) + \|\mu\|}{\sigma} \right) \right) \\ &+ \frac{1}{2} \left(-\Phi \left(\frac{-\max(0, B(\check{y})) + \|\mu\|}{\sigma} \right) + \Phi \left(\frac{-\min(0, B(\check{y})) + \|\mu\|}{\sigma} \right) \right) \\ &= \frac{1}{2} \left(\Phi \left(\frac{\|\mu\| + |B(\check{y})|}{\sigma} \right) - \Phi \left(\frac{\|\mu\| - |B(\check{y})|}{\sigma} \right) \right). \end{aligned}$$

Since the above quantity is increasing in $|B(\check{y})|$, and since $|B(\check{y})| \leq |\check{y}| \tan \beta + r$ where

$$r = \left| \frac{(x_0 - \mu_0)^T v}{\cos \beta} \right|,$$

we have that, replacing $\check{y}$ by x ,

$$\begin{aligned} L_\theta(\widehat{F}) &\leq \frac{1}{2} \int_{-\infty}^{\infty} \frac{1}{\sigma} \phi \left(\frac{x}{\sigma} \right) \left[\Phi \left(\frac{\|\mu\| + |x| \tan \beta + r}{\sigma} \right) - \Phi \left(\frac{\|\mu\| - |x| \tan \beta - r}{\sigma} \right) \right] dx \\ &\leq \int_{-\infty}^{\infty} \frac{1}{\sigma} \phi \left(\frac{x}{\sigma} \right) \left[\Phi \left(\frac{\|\mu\|}{\sigma} \right) - \Phi \left(\frac{\|\mu\| - |x| \tan \beta - r}{\sigma} \right) \right] dx \\ &= \int_{-\infty}^{\infty} \phi(x) \left[\Phi \left(\frac{\|\mu\|}{\sigma} \right) - \Phi \left(\frac{\|\mu\| - r}{\sigma} - |x| \tan \beta \right) \right] dx \\ &+ \left[\Phi \left(\frac{\|\mu\|}{\sigma} \right) - \Phi \left(\frac{\|\mu\| - r}{\sigma} \right) \right]. \end{aligned}$$

Since $\tan \beta \leq \frac{1}{2}$, we have that $r \leq 2|(x_0 - \mu_0)^T v| \leq 2\sigma\epsilon_1 + 2\|\mu\|\epsilon_2$ and

$$\begin{aligned} \Phi \left(\frac{\|\mu\|}{\sigma} \right) - \Phi \left(\frac{\|\mu\| - r}{\sigma} \right) &\leq \frac{r}{\sigma} \phi \left(\max \left(0, \frac{\|\mu\| - r}{\sigma} \right) \right) \\ &\leq \left(2\epsilon_1 + 2\epsilon_2 \frac{\|\mu\|}{\sigma} \right) \phi \left(\max \left(0, (1 - 2\epsilon_2) \frac{\|\mu\|}{\sigma} - 2\epsilon_1 \right) \right), \end{aligned}$$

and since $\epsilon_2 \leq \frac{1}{4}$,

$$\Phi \left(\frac{\|\mu\|}{\sigma} \right) - \Phi \left(\frac{\|\mu\| - r}{\sigma} \right) \leq 2 \left(\epsilon_1 + \epsilon_2 \frac{\|\mu\|}{\sigma} \right) \phi \left(\max \left(0, \frac{\|\mu\|}{2\sigma} - 2\epsilon_1 \right) \right).$$

Defining $A = \left\lfloor \frac{\|\mu\| - r}{\sigma} \right\rfloor$,

$$\begin{aligned}
& \int_{-\infty}^{\infty} \phi(x) \left[\Phi\left(\frac{\|\mu\| - r}{\sigma}\right) - \Phi\left(\frac{\|\mu\| - r}{\sigma} - |x| \tan \beta\right) \right] dx \\
& \leq 2 \int_0^{\infty} \int_{A - x \tan \beta}^A \phi(x) \phi(y) dy dx = 2 \int_{-A \sin \beta}^{\infty} \int_{A \cos \beta}^{A \cos \beta + (x + A \sin \beta) \tan \beta} \phi(x) \phi(y) dy dx \\
& \leq 2\phi(A \cos \beta) \tan \beta \int_{-A \sin \beta}^{\infty} (x + A \sin \beta) \phi(x) dx \\
& = 2\phi(A \cos \beta) \tan \beta (A \sin \beta \Phi(A \sin \beta) + \phi(A \sin \beta)) \\
& \leq 2\phi(A) \tan \beta (A \sin \beta + 1) \\
& \leq 2\phi\left(\max\left(0, \frac{\|\mu\| - r}{\sigma}\right)\right) \tan \beta \left(\left(\frac{\|\mu\| + r}{\sigma}\right) \sin \beta + 1\right)
\end{aligned}$$

and

$$\begin{aligned}
& \int_{-\infty}^{\infty} \phi(x) \left[\Phi\left(\frac{\|\mu\| - r}{\sigma}\right) - \Phi\left(\frac{\|\mu\| - r}{\sigma} - |x| \tan \beta\right) \right] dx \\
& \leq 2\phi\left(\max\left(0, \frac{\|\mu\|}{2\sigma} - 2\epsilon_1\right)\right) \tan \beta \left(\left(2\frac{\|\mu\|}{\sigma} + 2\epsilon_1\right) \sin \beta + 1\right).
\end{aligned}$$

So we have that

$$\begin{aligned}
L_{\theta}(\widehat{F}) & \leq 2\left(\epsilon_1 + \epsilon_2 \frac{\|\mu\|}{\sigma}\right) \phi\left(\max\left(0, \frac{\|\mu\|}{2\sigma} - 2\epsilon_1\right)\right) \\
& + 2\phi\left(\max\left(0, \frac{\|\mu\|}{2\sigma} - 2\epsilon_1\right)\right) \tan \beta \left(\left(2\frac{\|\mu\|}{\sigma} + 2\epsilon_1\right) \sin \beta + 1\right) \\
& \leq \phi\left(\max\left(0, \frac{\|\mu\|}{2\sigma} - 2\epsilon_1\right)\right) \times \\
& \times \left[2\epsilon_1 + 2\epsilon_2 \frac{\|\mu\|}{\sigma} + 4 \sin \beta \tan \beta \frac{\|\mu\|}{\sigma} + 4\epsilon_1 \sin \beta \tan \beta + 2 \tan \beta\right] \\
& \leq \exp\left\{-\frac{1}{2} \max\left(0, \frac{\|\mu\|}{2\sigma} - 2\epsilon_1\right)^2\right\} \left[2\epsilon_1 + \epsilon_2 \frac{\|\mu\|}{\sigma} + \tan \beta \left(2 \sin \beta \frac{\|\mu\|}{\sigma} + 1\right)\right].
\end{aligned}$$

□

2.6 Non-sparse upper bound

Theorem 1. For any $\theta \in \Theta_{\lambda}$ and $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P_{\theta}$, let

$$\widehat{F}(x) = \begin{cases} 1 & \text{if } x^T v_1(\widehat{\Sigma}_n) \geq \widehat{\mu}_n^T v_1(\widehat{\Sigma}_n) \\ 2 & \text{otherwise,} \end{cases}$$

and let $n \geq \max(68, 4d)$, $d \geq 1$.

Then

$$\sup_{\theta \in \Theta_{\lambda}} \mathbb{E} L_{\theta}(\widehat{F}) \leq 600 \max\left(\frac{4\sigma^2}{\lambda^2}, 1\right) \sqrt{\frac{d \log(nd)}{n}}.$$

Furthermore, if $\frac{\lambda}{\sigma} \geq 2 \max(80, 14\sqrt{5d})$, then

$$\sup_{\theta \in \Theta_{\lambda}} \mathbb{E} L_{\theta}(\widehat{F}) \leq 17 \exp\left(-\frac{n}{32}\right) + 9 \exp\left(-\frac{\lambda^2}{80\sigma^2}\right).$$

Proof. Using Propositions 5 and 8 with $\delta = \frac{1}{\sqrt{n}}$, Proposition 9, and the fact that $(C + x) \exp(-\max(0, x - 4)^2/8) \leq (C + 6) \exp(-\max(0, x - 4)^2/10)$ for all $C, x > 0$,

$$\mathbb{E}L_\theta(\hat{F}) \leq 600 \max\left(\frac{4\sigma^2}{\lambda^2}, 1\right) \sqrt{\frac{d \log(nd)}{n}}$$

(it is easy to verify that the bounds are decreasing with $\|\mu\|$, so we use $\|\mu\| = \frac{\lambda}{2}$ to bound the supremum). Note that the $d = 1$ case must be handled separately, but results in a bound that agrees with the above.

Also, when $\frac{\lambda}{\sigma} \geq 2 \max(80, 14\sqrt{5d})$, using $\delta = \exp(-\frac{n}{32})$,

$$\mathbb{E}L_\theta(\hat{F}) \leq 17 \exp\left(-\frac{n}{32}\right) + 9 \exp\left(-\frac{\lambda^2}{80\sigma^2}\right).$$

□

2.7 Estimating the support in the sparse case

Proposition 10. Let $\theta = (\mu_0 - \mu, \mu_0 + \mu)$ for some $\mu_0, \mu \in \mathbb{R}^d$ and $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P_\theta$. For any $0 < \delta < \frac{1}{\sqrt{e}}$ such that $\sqrt{\frac{6 \log \frac{1}{\delta}}{n}} \leq \frac{1}{2}$, with probability at least $1 - 6d\delta$,

$$|\hat{\Sigma}_n(i, i) - (\sigma^2 + \mu(i)^2)| \leq \sigma^2 \sqrt{\frac{6 \log \frac{1}{\delta}}{n}} + 2\sigma|\mu(i)| \sqrt{\frac{2 \log \frac{1}{\delta}}{n}} + (\sigma + |\mu(i)|)^2 \frac{2 \log \frac{1}{\delta}}{n}$$

for all $i \in [d]$.

Proof. Consider any $i \in [d]$. Let $Z_1, \dots, Z_n \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$ and $Y_1, \dots, Y_n$ i.i.d. such that $\mathbb{P}(Y_j = -1) = \mathbb{P}(Y_j = 1) = \frac{1}{2}$. Then $\hat{\Sigma}_n(i, i)$ is equal in distribution to

$$\frac{1}{n} \sum_{j=1}^n (\sigma Z_j + \mu(i) Y_j - \sigma \bar{Z} - \mu(i) \bar{Y})^2$$

where $\bar{Z}$ and $\bar{Y}$ are the respective empirical means, and

$$\begin{aligned} \frac{1}{n} \sum_{j=1}^n (\sigma Z_j + \mu(i) Y_j - \sigma \bar{Z} - \mu(i) \bar{Y})^2 &= \frac{1}{n} \sum_{j=1}^n (\sigma Z_j + \mu(i) Y_j)^2 - (\sigma \bar{Z} + \mu(i) \bar{Y})^2 \\ &= \sigma^2 \frac{1}{n} \sum_{j=1}^n Z_j^2 + \mu(i)^2 + 2\sigma\mu(i) \frac{1}{n} \sum_{j=1}^n Z_j Y_j \\ &\quad - \sigma^2 \bar{Z}^2 - \mu(i)^2 \bar{Y}^2 - 2\sigma\mu(i) \bar{Z} \bar{Y} \end{aligned}$$

So, by Hoeffding's inequality, a Gaussian tail bound, and Proposition 1, we have that for any $0 < \delta < \frac{1}{\sqrt{e}}$, with probability at least $1 - 6\delta$,

$$|\hat{\Sigma}_n(i, i) - (\sigma^2 + \mu(i)^2)| \leq \sigma^2 \sqrt{\frac{6 \log \frac{1}{\delta}}{n}} + 2\sigma|\mu(i)| \sqrt{\frac{2 \log \frac{1}{\delta}}{n}} + (\sigma + |\mu(i)|)^2 \frac{2 \log \frac{1}{\delta}}{n}$$

where we have used the fact that for $\epsilon \in (0, 0.5]$,

$$\max\{-\epsilon + \log(1 + \epsilon), \epsilon + \log(1 - \epsilon)\} \leq -\frac{\epsilon^2}{3}$$

and the result follows easily. □

Proposition 11. Let $\theta = (\mu_0 - \mu, \mu_0 + \mu)$ for some $\mu_0, \mu \in \mathbb{R}^d$ and $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P_\theta$. Define

$$\begin{aligned} S(\theta) &= \{i \in [d] : \mu(i) \neq 0\}, \\ \alpha &= \sqrt{\frac{6 \log(nd)}{n}} + \frac{2 \log(nd)}{n}, \\ \tilde{S}(\theta) &= \{i \in [d] : |\mu(i)| \geq 4\sigma\sqrt{\alpha}\}, \\ \hat{\tau}_n &= \frac{1+\alpha}{1-\alpha} \min_{i \in [d]} \hat{\Sigma}_n(i, i), \end{aligned}$$

and

$$\hat{S}_n = \{i \in [d] : \hat{\Sigma}_n(i, i) > \hat{\tau}_n\}.$$

Assume that $n \geq 1$, $d \geq 2$, and $\alpha \leq \frac{1}{4}$. Then $\tilde{S}(\theta) \subseteq \hat{S}_n \subseteq S(\theta)$ with probability at least $1 - \frac{6}{n}$.

Proof. By Proposition 10, with probability at least $1 - \frac{6}{n}$,

$$|\hat{\Sigma}_n(i, i) - (\sigma^2 + \mu(i)^2)| \leq \sigma^2 \sqrt{\frac{6 \log(nd)}{n}} + 2\sigma|\mu(i)| \sqrt{\frac{2 \log(nd)}{n}} + (\sigma + |\mu(i)|)^2 \frac{2 \log(nd)}{n}$$

for all $i \in [d]$. Assume the above event holds. If $S(\theta) = [d]$, then of course $\hat{S}_n \subseteq S(\theta)$. Otherwise, for $i \notin S(\theta)$,

$$(1 - \alpha)\sigma^2 \leq \hat{\Sigma}_n(i, i) \leq (1 + \alpha)\sigma^2$$

so it is clear that $\hat{S}_n \subseteq S(\theta)$.

The remainder of the proof is trivial if $\tilde{S}(\theta) = \emptyset$ or $S(\theta) = \emptyset$. Assume otherwise. For any $i \in S(\theta)$,

$$\begin{aligned} \hat{\Sigma}_n(i, i) &\geq (1 - \alpha)\sigma^2 + \mu(i)^2 - 2\sigma|\mu(i)| \sqrt{\frac{2 \log(nd)}{n}} - 2\sigma|\mu(i)| \frac{2 \log(nd)}{n} - \mu(i)^2 \frac{2 \log(nd)}{n} \\ &\geq (1 - \alpha)\sigma^2 + \left(1 - \frac{2 \log(nd)}{n}\right) \mu(i)^2 - 2\alpha\sigma|\mu(i)|. \end{aligned}$$

By definition, $|\mu(i)| \geq 4\sigma\sqrt{\alpha}$ for all $i \in \tilde{S}(\theta)$, so

$$\frac{(1 + \alpha)^2}{1 - \alpha} \sigma^2 \leq (1 - \alpha)\sigma^2 + \left(1 - \frac{2 \log(nd)}{n}\right) \mu(i)^2 - 2\alpha\sigma|\mu(i)| \leq \hat{\Sigma}_n(i, i)$$

and $i \in \hat{S}_n$ (we ignore strict equality above as a measure 0 event), i.e. $\tilde{S}(\theta) \subseteq \hat{S}_n$, which concludes the proof. $\square$

2.8 Sparse upper bound

Theorem 2. For any $\theta = (\mu_0 - \mu, \mu_0 + \mu) \in \Theta_{\lambda, s}$ and $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P_\theta$ with $n \geq \max(68, 4s)$ and $s \geq 1$, define

$$\begin{aligned} \alpha &= \sqrt{\frac{6 \log(nd)}{n}} + \frac{2 \log(nd)}{n}, \\ \hat{\tau}_n &= \frac{1+\alpha}{1-\alpha} \min_{i \in [d]} \hat{\Sigma}_n(i, i), \end{aligned}$$

and

$$\hat{S}_n = \{i \in [d] : \hat{\Sigma}_n(i, i) > \hat{\tau}_n\}.$$

Assume that $d \geq 2$, and $\alpha \leq \frac{1}{4}$. Let

$$\hat{F}_n(x) = \begin{cases} 1 & \text{if } x_{\hat{S}_n}^T v_1(\hat{\Sigma}_{\hat{S}_n}) \geq \hat{\mu}_{\hat{S}_n}^T v_1(\hat{\Sigma}_{\hat{S}_n}) \\ 2 & \text{otherwise} \end{cases}$$

where $\hat{\mu}_{\hat{S}_n}$ and $\hat{\Sigma}_{\hat{S}_n}$ are the empirical mean and covariance of X_i for the dimensions in $\hat{S}_n$, and 0 elsewhere. Then

$$\sup_{\theta \in \Theta_{\lambda, s}} \mathbb{E} L_\theta(\hat{F}) \leq 603 \max\left(\frac{16\sigma^2}{\lambda^2}, 1\right) \sqrt{\frac{s \log(ns)}{n}} + 220 \frac{\sigma\sqrt{s}}{\lambda} \left(\frac{\log(nd)}{n}\right)^{\frac{1}{4}}.$$

Proof. Define

$$S(\theta) = \{i \in [d] : \mu(i) \neq 0\}$$

and

$$\tilde{S}(\theta) = \{i \in [d] : |\mu(i)| \geq 4\sigma\sqrt{\alpha}\},$$

Assume $\tilde{S}(\theta) \subseteq \hat{S}_n \subseteq S(\theta)$ (by Proposition 11, this holds with probability at least $1 - \frac{6}{n}$). If $\tilde{S}(\theta) = \emptyset$, then we simply have $\mathbb{E}L_\theta(\hat{F}_n) \leq \frac{1}{2}$.

Assume $\tilde{S}(\theta) \neq \emptyset$. Let

$$\cos \hat{\beta} = |v_1(\hat{\Sigma}_{\hat{S}_n})^T v_1(\Sigma)|,$$

$$\cos \tilde{\beta} = |v_1(\Sigma_{\tilde{S}_n})^T v_1(\Sigma)|,$$

and

$$\cos \beta = |v_1(\hat{\Sigma}_{\hat{S}_n})^T v_1(\Sigma_{\hat{S}_n})|$$

where $\Sigma = \sigma^2 I + \mu\mu^T$, and $\Sigma_{\hat{S}_n}$ is the same as Σ in $\hat{S}_n$, and 0 elsewhere. Then

$$\sin \hat{\beta} \leq \sin \tilde{\beta} + \sin \beta.$$

Also

$$\begin{aligned} \sin \tilde{\beta} &= \frac{\|\mu - \mu_{\tilde{S}(\theta)}\|}{\|\mu\|} \\ &\leq \frac{\|\mu - \mu_{\tilde{S}(\theta)}\|}{\|\mu\|} \\ &\leq \frac{4\sigma\sqrt{\alpha}\sqrt{|S(\theta)| - |\tilde{S}(\theta)|}}{\|\mu\|} \\ &\leq 8\frac{\sigma\sqrt{s\alpha}}{\lambda}. \end{aligned}$$

Using the same argument as the proof of Theorem 1, we have that as long as the above bound is smaller than $\frac{1}{2\sqrt{5}}$,

$$\begin{aligned} \mathbb{E}L_\theta(\hat{F}) &\leq 600 \max\left(\frac{\sigma^2}{\left(\frac{\lambda}{2} - 4\sigma\sqrt{s\alpha}\right)^2}, 1\right) \sqrt{\frac{s \log(ns)}{n}} + 104\frac{\sigma\sqrt{s\alpha}}{\lambda} + \frac{3}{n} \\ &\leq 603 \max\left(16\frac{\sigma^2}{\lambda^2}, 1\right) \sqrt{\frac{s \log(ns)}{n}} + 104\frac{\sigma\sqrt{s\alpha}}{\lambda}. \end{aligned}$$

However, when $8\frac{\sigma\sqrt{s\alpha}}{\lambda} > \frac{1}{2\sqrt{5}}$, the above bound is bigger than $\frac{1}{2}$, which is a trivial upper bound on the clustering error, hence the bound can be stated without further conditions. Finally, since $\alpha \leq \frac{1}{4}$, we must have $\frac{\log(nd)}{n} \leq 1$, so $\alpha \leq (\sqrt{6} + 2)\sqrt{\frac{\log(nd)}{n}}$, which completes the proof. $\square$

3 Lower bounds

3.1 Standard tools

Lemma 2. Let $P_0, P_1, \dots, P_M$ be probability measures satisfying

$$\frac{1}{M} \sum_{i=1}^M \text{KL}(P_i, P_0) \leq \alpha \log M$$

where $0 < \alpha < 1/8$ and $M \geq 2$. Then

$$\inf_{\psi} \max_{i \in [0..M]} P_i(\psi \neq i) \geq 0.07$$

(Tsybakov (2009)).

Lemma 3. (Varshamov–Gilbert bound) Let $\Omega = \{0, 1\}^m$ for $m \geq 8$. Then there exists a subset $\{\omega_0, \dots, \omega_M\} \subseteq \Omega$ such that $\omega_0 = (0, \dots, 0)$,

$$\rho(\omega_i, \omega_j) \geq \frac{m}{8}, \quad \forall 0 \leq i < j \leq M,$$

and

$$M \geq 2^{m/8},$$

where ρ denotes the Hamming distance between two vectors (Tsybakov (2009)).

Lemma 4. Let $\Omega = \{\omega \in \{0, 1\}^m : \|\omega\|_0 = s\}$ for integers $m > s \geq 1$. For any $\alpha, \beta \in (0, 1)$ such that $s \leq \alpha\beta m$, there exists $\omega_0, \dots, \omega_M \in \Omega$ such that for all $0 \leq i < j \leq M$,

$$\rho(\omega_i, \omega_j) > 2(1 - \alpha)s$$

and

$$\log(M + 1) \geq cs \log\left(\frac{m}{s}\right)$$

where

$$c = \frac{\alpha}{-\log(\alpha\beta)}(-\log \beta + \beta - 1).$$

In particular, setting $\alpha = 3/4$ and $\beta = 1/3$, we have that $\rho(\omega_i, \omega_j) > s/2$, $\log(M + 1) \geq \frac{s}{5} \log\left(\frac{m}{s}\right)$, as long as $s \leq m/4$ (Massart (2007), Lemma 4.10).

3.2 A reduction to hypothesis testing without a general triangle inequality

Proposition 12. Let $\theta_0, \dots, \theta_M \in \Theta_\lambda$ (or $\Theta_{\lambda, s}$), $M \geq 2$, $0 < \alpha < 1/8$, and $\gamma > 0$. If

$$\max_{i \in [M]} \text{KL}(P_{\theta_i}, P_{\theta_0}) \leq \frac{\alpha \log M}{n}$$

and for all $0 \leq i \neq j \leq M$ and clusterings $\hat{F}$,

$$L_{\theta_i}(\hat{F}) < \gamma \text{ implies } L_{\theta_j}(\hat{F}) \geq \gamma,$$

then

$$\inf_{\hat{F}_n} \max_{i \in [0..M]} \mathbb{E}_{\theta_i} L_{\theta_i}(\hat{F}_n) \geq 0.07\gamma.$$

Proof. Using Markov's inequality,

$$\inf_{\hat{F}_n} \max_{i \in [0..M]} \mathbb{E}_{\theta_i} L_{\theta_i}(\hat{F}_n) \geq \gamma \inf_{\hat{F}_n} \max_{i \in [0..M]} P_{\theta_i}^n \left(L_{\theta_i}(\hat{F}_n) \geq \gamma \right).$$

Define $\psi^*(\hat{F}_n) = \underset{i \in [0..M]}{\operatorname{argmin}} L_{\theta_i}(\hat{F}_n)$. By assumption, $L_{\theta_i}(\hat{F}_n) < \gamma$ implies $L_{\theta_j}(\hat{F}_n) \geq \gamma$ for any $j \neq i$, so $L_{\theta_i}(\hat{F}_n) < \gamma$ only when $\psi^*(\hat{F}_n) = i$. Hence,

$$P_{\theta_i}^n \left(\psi^*(\hat{F}_n) = i \right) \geq P_{\theta_i}^n \left(L_{\theta_i}(\hat{F}_n) < \gamma \right)$$

and

$$\begin{aligned} \inf_{\hat{F}_n} \max_{i \in [0..M]} P_{\theta_i}^n \left(L_{\theta_i}(\hat{F}_n) \geq \gamma \right) &\geq \max_{i \in [0..M]} P_{\theta_i}^n \left(\psi^*(\hat{F}_n) \neq i \right) \\ &\geq \inf_{\hat{\psi}_n} \max_{i \in [0..M]} P_{\theta_i}^n \left(\hat{\psi}_n \neq i \right) \\ &\geq 0.07 \end{aligned}$$

where the last step is by Lemma 2. □

3.3 Properties of the clustering error

Proposition 13. For any $\theta, \theta' \in \Theta_\lambda$, and any clustering $\hat{F}$, if

$$L_\theta(F_{\theta'}) + L_\theta(\hat{F}) + \sqrt{\frac{\text{KL}(P_\theta, P_{\theta'})}{2}} \leq \frac{1}{2},$$

then

$$L_\theta(F_{\theta'}) - L_\theta(\hat{F}) - \sqrt{\frac{\text{KL}(P_\theta, P_{\theta'})}{2}} \leq L_{\theta'}(\hat{F}) \leq L_\theta(F_{\theta'}) + L_\theta(\hat{F}) + \sqrt{\frac{\text{KL}(P_\theta, P_{\theta'})}{2}}.$$

Proof. WLOG assume $F_\theta, F_{\theta'}$, and $\hat{F}$ are such that, using simplified notation,

$$L_\theta(F_{\theta'}) = P_\theta(F_\theta \neq F_{\theta'})$$

and

$$L_\theta(\hat{F}) = P_\theta(F_\theta \neq \hat{F}).$$

Then

$$\begin{aligned} P_\theta(F_{\theta'} \neq \hat{F}) &= P_\theta\left((F_\theta = F_{\theta'}) \cap (F_\theta \neq \hat{F}) \cup (F_\theta \neq F_{\theta'}) \cap (F_\theta = \hat{F})\right) \\ &= P_\theta\left((F_\theta = F_{\theta'}) \cap (F_\theta \neq \hat{F})\right) + P_\theta\left((F_\theta \neq F_{\theta'}) \cap (F_\theta = \hat{F})\right). \end{aligned}$$

Since

$$0 \leq P_\theta\left((F_\theta = F_{\theta'}) \cap (F_\theta \neq \hat{F})\right) \leq P_\theta(F_\theta \neq \hat{F}) = L_\theta(\hat{F}),$$

$$P_\theta\left((F_\theta \neq F_{\theta'}) \cap (F_\theta = \hat{F})\right) \leq P_\theta(F_\theta \neq F_{\theta'}) = L_\theta(F_{\theta'}),$$

and

$$L_\theta(F_{\theta'}) - L_\theta(\hat{F}) = P_\theta(F_\theta \neq F_{\theta'}) - P_\theta(F_\theta \neq \hat{F}) \leq P_\theta\left((F_\theta \neq F_{\theta'}) \cap (F_\theta = \hat{F})\right),$$

we have that

$$L_\theta(F_{\theta'}) - L_\theta(\hat{F}) \leq P_\theta(F_{\theta'} \neq \hat{F}) \leq L_\theta(F_{\theta'}) + L_\theta(\hat{F})$$

and

$$L_\theta(F_{\theta'}) - L_\theta(\hat{F}) - \text{TV}(P_\theta, P_{\theta'}) \leq P_{\theta'}(F_{\theta'} \neq \hat{F}) \leq L_\theta(F_{\theta'}) + L_\theta(\hat{F}) + \text{TV}(P_\theta, P_{\theta'}).$$

It is easy to see that if $L_\theta(F_{\theta'}) + L_\theta(\hat{F}) + \text{TV}(P_\theta, P_{\theta'}) \leq \frac{1}{2}$, then the above bound implies

$$L_\theta(F_{\theta'}) - L_\theta(\hat{F}) - \text{TV}(P_\theta, P_{\theta'}) \leq L_{\theta'}(\hat{F}) \leq L_\theta(F_{\theta'}) + L_\theta(\hat{F}) + \text{TV}(P_\theta, P_{\theta'}).$$

The final step is to use the fact that $\text{TV}(P_\theta, P_{\theta'}) \leq \sqrt{\frac{\text{KL}(P_\theta, P_{\theta'})}{2}}$. □

Proposition 14. For some $\mu_0, \mu, \mu' \in \mathbb{R}^d$ such that $\|\mu\| = \|\mu'\|$, let

$$\theta = \left(\mu_0 - \frac{\mu}{2}, \mu_0 + \frac{\mu}{2}\right)$$

and

$$\theta' = \left(\mu_0 - \frac{\mu'}{2}, \mu_0 + \frac{\mu'}{2}\right).$$

Then

$$2g\left(\frac{\|\mu\|}{2\sigma}\right) \sin \beta \cos \beta \leq L_\theta(F_{\theta'}) \leq \frac{1}{\pi} \tan \beta$$

where $\cos \beta = \frac{|\mu^T \mu'|}{\|\mu\|^2}$ and $g(x) = \phi(x)(\phi(x) - x\Phi(-x))$.

Proof. It is easy to see that

$$L_\theta(F_{\theta'}) = \frac{1}{2} \int_{\mathbb{R}} \frac{1}{\sigma} \phi\left(\frac{x}{\sigma}\right) \left(\Phi\left(\frac{\|\mu\|}{2\sigma} + \frac{|x| \tan \beta}{\sigma}\right) - \Phi\left(\frac{\|\mu\|}{2\sigma} - \frac{|x| \tan \beta}{\sigma}\right) \right) dx.$$

Define $\xi = \frac{\|\mu\|}{2\sigma}$. With a change of variables, we have

$$\begin{aligned} L_\theta(F_{\theta'}) &= \frac{1}{2} \int_{\mathbb{R}} \phi(x) (\Phi(\xi + |x| \tan \beta) - \Phi(\xi - |x| \tan \beta)) dx \\ &= \int_0^\infty \phi(x) (\Phi(\xi + x \tan \beta) - \Phi(\xi - x \tan \beta)) dx. \end{aligned}$$

For any $a \leq b$, $\Phi(b) - \Phi(a) \leq \frac{b-a}{\sqrt{2\pi}}$, so

$$\begin{aligned} L_\theta(F_{\theta'}) &= \int_0^\infty \phi(x) (\Phi(\xi + x \tan \beta) - \Phi(\xi - x \tan \beta)) dx \\ &\leq \int_0^\infty \phi(x) (\Phi(x \tan \beta) - \Phi(-x \tan \beta)) dx \\ &\leq \tan \beta \sqrt{\frac{2}{\pi}} \int_0^\infty x \phi(x) dx \\ &= \frac{1}{\pi} \tan \beta. \end{aligned}$$

Also,

$$\begin{aligned} L_\theta(F_{\theta'}) &= \int_0^\infty \phi(x) (\Phi(\xi + x \tan \beta) - \Phi(\xi - x \tan \beta)) dx \\ &\geq 2 \tan \beta \int_0^\infty x \phi(x) \phi(\xi + x \tan \beta) dx \\ &= 2 \tan \beta \frac{1}{\sqrt{2\pi}} \int_0^\infty x \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{x^2 + (\xi + x \tan \beta)^2}{2}\right\} dx \\ &= 2 \tan \beta \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{\xi^2}{2}\left(1 - \frac{\tan^2 \beta}{1 + \tan^2 \beta}\right)\right\} \int_0^\infty x \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{\left(x + \frac{\xi \tan \beta}{1 + \tan^2 \beta}\right)^2}{2\left(\frac{1}{\sqrt{1 + \tan^2 \beta}}\right)^2}\right\} dx \\ &\geq 2 \tan \beta \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{\xi^2}{2}\right\} \int_0^\infty x \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{(x + \xi \sin \beta \cos \beta)^2}{2 \cos^2 \beta}\right\} dx \\ &= 2 \tan \beta \phi(\xi) \left[\frac{\cos^2 \beta}{\sqrt{2\pi}} \exp\left\{-\frac{\xi^2 \sin^2 \beta}{2}\right\} - \xi \sin \beta \cos^2 \beta \Phi(-\xi \sin \beta) \right] \\ &= 2 \sin \beta \cos \beta \phi(\xi) [\phi(\xi \sin \beta) - \xi \sin \beta \Phi(-\xi \sin \beta)] \\ &\geq 2 \sin \beta \cos \beta \phi(\xi) [\phi(\xi) - \xi \Phi(-\xi)]. \end{aligned}$$

□

3.4 A KL divergence bound of the necessary order

Proposition 15. For some $\mu_0, \mu, \mu' \in \mathbb{R}^d$ such that $\|\mu\| = \|\mu'\|$, let

$$\theta = \left(\mu_0 - \frac{\mu}{2}, \mu_0 + \frac{\mu}{2} \right)$$

and

$$\theta' = \left(\mu_0 - \frac{\mu'}{2}, \mu_0 + \frac{\mu'}{2} \right).$$

Then

$$\text{KL}(P_\theta, P_{\theta'}) \leq \xi^4 (1 - \cos \beta)$$

where $\xi = \frac{\|\mu\|}{2\sigma}$ and $\cos \beta = \frac{|\mu^T \mu'|}{\|\mu\| \|\mu'\|}$.

Proof. Since the KL divergence is invariant to affine transformations, it is easy to see that

$$\text{KL}(P_\theta, P_{\theta'}) = \int_{\mathbb{R}} \int_{\mathbb{R}} p_1(x, y) \log \frac{p_1(x, y)}{p_2(x, y)} dx dy$$

where

$$p_1(x, y) = \frac{1}{2} \phi(x + \xi_x) \phi(y + \xi_y) + \frac{1}{2} \phi(x - \xi_x) \phi(y - \xi_y),$$

$$p_2(x, y) = \frac{1}{2} \phi(x + \xi_x) \phi(y - \xi_y) + \frac{1}{2} \phi(x - \xi_x) \phi(y + \xi_y),$$

$$\xi_x = \xi \cos \frac{\beta}{2}, \quad \xi_y = \xi \sin \frac{\beta}{2}.$$

Since

$$\begin{aligned} \frac{p_1(x, y)}{p_2(x, y)} &= \frac{\phi(x + \xi_x) \phi(y + \xi_y) + \phi(x - \xi_x) \phi(y - \xi_y)}{\phi(x + \xi_x) \phi(y - \xi_y) + \phi(x - \xi_x) \phi(y + \xi_y)} \\ &= \frac{\exp(-x\xi_x - y\xi_y) + \exp(x\xi_x + y\xi_y)}{\exp(-x\xi_x + y\xi_y) + \exp(x\xi_x - y\xi_y)} \end{aligned}$$

we have

$$\log \frac{p_1(x, y)}{p_2(x, y)} = \log \frac{\cosh(x\xi_x + y\xi_y)}{\cosh(x\xi_x - y\xi_y)}.$$

Furthermore,

$$\begin{aligned} & \int_{\mathbb{R}} \int_{\mathbb{R}} \frac{1}{2} \phi(x + \xi_x) \phi(y + \xi_y) \log \frac{\cosh(x\xi_x + y\xi_y)}{\cosh(x\xi_x - y\xi_y)} dx dy \\ &= \int_{\mathbb{R}} \int_{\mathbb{R}} \frac{1}{2} \phi(-x + \xi_x) \phi(-y + \xi_y) \log \frac{\cosh(-x\xi_x - y\xi_y)}{\cosh(-x\xi_x + y\xi_y)} dx dy \\ &= \int_{\mathbb{R}} \int_{\mathbb{R}} \frac{1}{2} \phi(x - \xi_x) \phi(y - \xi_y) \log \frac{\cosh(x\xi_x + y\xi_y)}{\cosh(x\xi_x - y\xi_y)} dx dy \end{aligned}$$

so

$$\begin{aligned} \text{KL}(P_\theta, P_{\theta'}) &= \int_{\mathbb{R}} \int_{\mathbb{R}} \phi(x - \xi_x) \phi(y - \xi_y) \log \frac{\cosh(x\xi_x + y\xi_y)}{\cosh(x\xi_x - y\xi_y)} dx dy \\ &= \int_{\mathbb{R}} \int_{\mathbb{R}} \phi(x) \phi(y) \log \frac{\cosh(x\xi_x + \xi_x^2 + y\xi_y + \xi_y^2)}{\cosh(x\xi_x + \xi_x^2 - y\xi_y - \xi_y^2)} dx dy. \end{aligned}$$

But for any x

$$\begin{aligned}
& - \int_{\mathbb{R}} \phi(x) \phi(y) \log \cosh(x\xi_x + \xi_x^2 - y\xi_y - \xi_y^2) dy \\
& = - \int_{\mathbb{R}} \phi(x) \phi(-y) \log \cosh(x\xi_x + \xi_x^2 + y\xi_y - \xi_y^2) dy \\
& = - \int_{\mathbb{R}} \phi(x) \phi(y) \log \cosh(x\xi_x + \xi_x^2 + y\xi_y - \xi_y^2) dy,
\end{aligned}$$

thus,

$$\begin{aligned}
\text{KL}(P_\theta, P_{\theta'}) &= \int_{\mathbb{R}} \int_{\mathbb{R}} \phi(x) \phi(y) \log \frac{\cosh(x\xi_x + \xi_x^2 + y\xi_y + \xi_y^2)}{\cosh(x\xi_x + \xi_x^2 + y\xi_y - \xi_y^2)} dx dy \\
&= \int_{\mathbb{R}} \phi(z) \log \frac{\cosh(z\sqrt{\xi_x^2 + \xi_y^2} + \xi_x^2 + \xi_y^2)}{\cosh(z\sqrt{\xi_x^2 + \xi_y^2} + \xi_x^2 - \xi_y^2)} dz \\
&= \int_{\mathbb{R}} \phi(z) \log \frac{\cosh(\xi z + \xi_x^2 + \xi_y^2)}{\cosh(\xi z + \xi_x^2 - \xi_y^2)} dz
\end{aligned}$$

since $\xi_x^2 + \xi_y^2 = \xi^2$. By the mean value theorem and the fact that $\tanh$ is monotonically increasing,

$$\begin{aligned}
\log \frac{\cosh(\xi z + \xi_x^2 + \xi_y^2)}{\cosh(\xi z + \xi_x^2 - \xi_y^2)} &\leq 2\xi_y^2 \tanh(\xi z + \xi_x^2 + \xi_y^2) \\
&= 2\xi_y^2 \tanh(\xi z + \xi^2)
\end{aligned}$$

for all z . Since $\tanh$ is an odd function,

$$\begin{aligned}
\text{KL}(P_\theta, P_{\theta'}) &\leq 2\xi_y^2 \int_{\mathbb{R}} \phi(z) \tanh(\xi z + \xi^2) dz \\
&= 2\xi_y^2 \int_{\mathbb{R}} \phi(z) (\tanh(\xi z + \xi^2) - \tanh(\xi z)) dz.
\end{aligned}$$

Using the mean value theorem again,

$$\begin{aligned}
\tanh(\xi z + \xi^2) - \tanh(\xi z) &\leq \xi^2 \max_{x \in [\xi z, \xi z + \xi^2]} (1 - \tanh^2(x)) \\
&\leq \xi^2
\end{aligned}$$

for all z , so

$$\begin{aligned}
\text{KL}(P_\theta, P_{\theta'}) &\leq 2\xi^2 \xi_y^2 \\
&= 2\xi^4 \sin^2 \frac{\beta}{2} \\
&= \xi^4 (1 - \cos \beta).
\end{aligned}$$

□

3.5 Non-sparse lower bound

Theorem 3. Assume that $d \geq 9$ and $\frac{\lambda}{\sigma} \leq 0.2$. Then

$$\inf_{\hat{F}_n} \sup_{\theta \in \Theta_\lambda} \mathbb{E}_\theta L_\theta(\hat{F}_n) \geq \frac{1}{500} \min \left\{ \frac{\sqrt{\log 2} \sigma^2}{3 \lambda^2} \sqrt{\frac{d-1}{n}}, \frac{1}{4} \right\}.$$

Proof. Let $\xi = \frac{\lambda}{2\sigma}$, and define

$$\epsilon = \min \left\{ \frac{\sqrt{\log 2}}{3} \frac{\sigma^2}{\lambda} \frac{1}{\sqrt{n}}, \frac{\lambda}{4\sqrt{d-1}} \right\}.$$

Define $\lambda_0^2 = \lambda^2 - (d-1)\epsilon^2$. Let $\Omega = \{0, 1\}^{d-1}$. For $\omega = (\omega(1), \dots, \omega(d-1)) \in \Omega$, let $\mu_\omega = \lambda_0 e_d + \sum_{i=1}^{d-1} (2\omega(i) - 1)\epsilon e_i$ (where $\{e_i\}_{i=1}^d$ is the standard basis for $\mathbb{R}^d$). Let $\theta_\omega = \left(-\frac{\mu_\omega}{2}, \frac{\mu_\omega}{2}\right) \in \Theta_\lambda$.

By Proposition 15, for any $\omega, \nu \in \Omega$,

$$\text{KL}(P_{\theta_\omega}, P_{\theta_\nu}) \leq \xi^4 (1 - \cos \beta_{\omega, \nu})$$

where

$$\cos \beta_{\omega, \nu} = \frac{|\mu_\omega^T \mu_\nu|}{\lambda^2} = 1 - \frac{2\rho(\omega, \nu)\epsilon^2}{\lambda^2}$$

and ρ is the Hamming distance, so

$$\begin{aligned} \text{KL}(P_{\theta_\omega}, P_{\theta_\nu}) &\leq \xi^4 \frac{2\rho(\omega, \nu)\epsilon^2}{\lambda^2} \\ &\leq \xi^4 \frac{2(d-1)\epsilon^2}{\lambda^2}. \end{aligned}$$

By Proposition 14, since $\cos \beta_{\omega, \nu} \geq \frac{1}{2}$,

$$\begin{aligned} L_{\theta_\omega}(F_{\theta_\nu}) &\leq \frac{1}{\pi} \tan \beta_{\omega, \nu} \\ &\leq \frac{1}{\pi} \frac{\sqrt{1 + \cos \beta_{\omega, \nu}}}{\cos \beta_{\omega, \nu}} \sqrt{1 - \cos \beta_{\omega, \nu}} \\ &\leq \frac{4}{\pi} \frac{\sqrt{d-1}\epsilon}{\lambda} \end{aligned}$$

and

$$\begin{aligned} L_{\theta_\omega}(F_{\theta_\nu}) &\geq 2g(\xi) \sin \beta_{\omega, \nu} \cos \beta_{\omega, \nu} \\ &\geq g(\xi) \sin \beta_{\omega, \nu} \\ &\geq \sqrt{2}g(\xi) \frac{\sqrt{\rho(\omega, \nu)}\epsilon}{\lambda} \end{aligned}$$

where $g(x) = \phi(x)(\phi(x) - x\Phi(-x))$. By Lemma 3, there exist $\omega_0, \dots, \omega_M \in \Omega$ such that $M \geq 2^{(d-1)/8}$ and

$$\rho(\omega_i, \omega_j) \geq \frac{d-1}{8}, \quad \forall 0 \leq i < j \leq M.$$

For simplicity of notation, let $\theta_i = \theta_{\omega_i}$ for all $i \in [0..M]$. Then, for $i \neq j \in [0..M]$,

$$\text{KL}(P_{\theta_i}, P_{\theta_j}) \leq \xi^4 \frac{2(d-1)\epsilon^2}{\lambda^2},$$

and

$$L_{\theta_i}(F_{\theta_j}) \leq \frac{4}{\pi} \frac{\sqrt{d-1}\epsilon}{\lambda}$$

and

$$L_{\theta_i}(F_{\theta_j}) \geq \frac{1}{2}g(\xi) \frac{\sqrt{d-1}\epsilon}{\lambda}.$$

Define

$$\gamma = \frac{1}{4}(g(\xi) - 2\xi^2) \frac{\sqrt{d-1}\epsilon}{\lambda}.$$

Then for any $i \neq j \in [0..M]$, and any $\hat{F}$ such that $L_{\theta_i}(\hat{F}) < \gamma$,

$$L_{\theta_i}(F_{\theta_j}) + L_{\theta_i}(\hat{F}) + \sqrt{\frac{\text{KL}(P_{\theta_i}, P_{\theta_j})}{2}} < \left(\frac{4}{\pi} + \frac{1}{4}(g(\xi) - 2\xi^2) + \xi^2 \right) \frac{\sqrt{d-1}\epsilon}{\lambda} \leq \frac{1}{2}$$

because, for $\xi \leq 0.1$, by definition of ϵ ,

$$\left(\frac{4}{\pi} + \frac{1}{4}(g(\xi) - 2\xi^2) + \xi^2 \right) \frac{\sqrt{d-1}\epsilon}{\lambda} \leq 2 \frac{\sqrt{d-1}\epsilon}{\lambda} \leq \frac{1}{2}.$$

So, by Proposition 13,

$$L_{\theta_j}(\hat{F}) \geq L_{\theta_i}(F_{\theta_j}) - L_{\theta_i}(\hat{F}) - \sqrt{\frac{\text{KL}(P_{\theta_i}, P_{\theta_j})}{2}} \geq \gamma.$$

Also,

$$\begin{aligned} \max_{i \in [M]} \text{KL}(P_{\theta_i}, P_{\theta_0}) &\leq (d-1)\xi^4 \frac{2\epsilon^2}{\lambda^2} \\ &\leq \frac{\log M}{9n} \end{aligned}$$

because, by definition of ϵ ,

$$\xi^4 \frac{2\epsilon^2}{\lambda^2} \leq \frac{\log 2}{72n}.$$

So by Proposition 12 and the fact that $\xi \leq 0.1$,

$$\begin{aligned} \inf_{\hat{F}_n} \max_{i \in [0..M]} \mathbb{E}_{\theta_i} L_{\theta_i}(\hat{F}_n) &\geq 0.07\gamma \\ &= 0.07 \frac{1}{4}(g(\xi) - 2\xi^2) \frac{\sqrt{d-1}\epsilon}{\lambda} \\ &\geq \frac{1}{500} \min \left\{ \frac{\sqrt{\log 2}}{3} \frac{\sigma^2}{\lambda^2} \sqrt{\frac{d-1}{n}}, \frac{1}{4} \right\} \end{aligned}$$

and to complete the proof we use the fact that

$$\inf_{\hat{F}_n} \sup_{\theta \in \Theta_\lambda} \mathbb{E}_\theta L_\theta(\hat{F}_n) \geq \inf_{\hat{F}_n} \max_{i \in [0..M]} \mathbb{E}_{\theta_i} L_{\theta_i}(\hat{F}_n).$$

□

3.6 Sparse lower bound

Theorem 4. Assume that $\frac{\lambda}{\sigma} \leq 0.2$, $d \geq 17$, and

$$5 \leq s \leq \frac{d-1}{4} + 1.$$

Then

$$\inf_{\hat{F}_n} \sup_{\theta \in \Theta_{\lambda,s}} \mathbb{E}_\theta L_\theta(\hat{F}_n) \geq \frac{1}{600} \min \left\{ \sqrt{\frac{8}{45}} \frac{\sigma^2}{\lambda^2} \sqrt{\frac{s-1}{n} \log \left(\frac{d-1}{s-1} \right)}, \frac{1}{2} \right\}.$$

Proof. For simplicity, we state this proof for $\Theta_{\lambda,s+1}$, assuming $4 \leq s \leq \frac{d-1}{4}$. Let $\xi = \frac{\lambda}{2\sigma}$, and define

$$\epsilon = \min \left\{ \sqrt{\frac{8}{45}} \frac{\sigma^2}{\lambda} \sqrt{\frac{1}{n} \log \left(\frac{d-1}{s} \right)}, \frac{1}{2} \frac{\lambda}{\sqrt{s}} \right\}.$$

Define $\lambda_0^2 = \lambda^2 - s\epsilon^2$. Let $\Omega = \{\omega \in \{0, 1\}^{d-1} : \|\omega\|_0 = s\}$. For $\omega = (\omega(1), \dots, \omega(d-1)) \in \Omega$, let $\mu_\omega = \lambda_0 e_d + \sum_{i=1}^{d-1} \omega(i) \epsilon e_i$ (where $\{e_i\}_{i=1}^d$ is the standard basis for $\mathbb{R}^d$). Let $\theta_\omega = (-\frac{\mu_\omega}{2}, \frac{\mu_\omega}{2}) \in \Theta_{\lambda,s}$.

By Proposition 15,

$$\text{KL}(P_{\theta_\omega}, P_{\theta_\nu}) \leq \xi^4(1 - \cos \beta_{\omega,\nu})$$

where

$$\cos \beta_{\omega,\nu} = \frac{|\mu_\omega^T \mu_\nu|}{\lambda^2} = 1 - \frac{\rho(\omega, \nu)\epsilon^2}{2\lambda^2}$$

and ρ is the Hamming distance, so

$$\begin{aligned} \text{KL}(P_{\theta_\omega}, P_{\theta_\nu}) &\leq \xi^4 \frac{\rho(\omega, \nu)\epsilon^2}{2\lambda^2} \\ &\leq \xi^4 \frac{s\epsilon^2}{\lambda^2}. \end{aligned}$$

By Proposition 14, since $\cos \beta_{\omega,\nu} \geq \frac{1}{2}$,

$$\begin{aligned} L_{\theta_\omega}(F_{\theta_\nu}) &\leq \frac{1}{\pi} \tan \beta_{\omega,\nu} \\ &\leq \frac{2}{\pi} \sin \beta_{\omega,\nu} \\ &\leq \frac{2\sqrt{2}}{\pi} \frac{\sqrt{s\epsilon}}{\lambda} \end{aligned}$$

and

$$\begin{aligned} L_{\theta_\omega}(F_{\theta_\nu}) &\geq 2g(\xi) \sin \beta_{\omega,\nu} \cos \beta_{\omega,\nu} \\ &\geq g(\xi) \sin \beta_{\omega,\nu} \\ &\geq \frac{g(\xi)}{\sqrt{2}} \frac{\sqrt{\rho(\omega, \nu)\epsilon}}{\lambda} \end{aligned}$$

where $g(x) = \phi(x)(\phi(x) - x\Phi(-x))$. By Lemma 4, there exist $\omega_0, \dots, \omega_M \in \Omega$ such that $\log(M+1) \geq \frac{s}{5} \log\left(\frac{d-1}{s}\right)$ and

$$\rho(\omega_i, \omega_j) \geq \frac{s}{2}, \quad \forall 0 \leq i < j \leq M.$$

For simplicity of notation, let $\theta_i = \theta_{\omega_i}$ for all $i \in [0..M]$. Then, for $i \neq j \in [0..M]$,

$$\text{KL}(P_{\theta_i}, P_{\theta_j}) \leq \xi^4 \frac{s\epsilon^2}{\lambda^2},$$

and

$$L_{\theta_i}(F_{\theta_j}) \leq \frac{2\sqrt{2}}{\pi} \frac{\sqrt{s\epsilon}}{\lambda}$$

and

$$L_{\theta_i}(F_{\theta_j}) \geq \frac{g(\xi)}{2} \frac{\sqrt{s\epsilon}}{\lambda}.$$

Define

$$\gamma = \frac{1}{4}(g(\xi) - \sqrt{2}\xi^2) \frac{\sqrt{s\epsilon}}{\lambda}.$$

Then for any $i \neq j \in [0..M]$, and any $\hat{F}$ such that $L_{\theta_i}(\hat{F}) < \gamma$,

$$L_{\theta_i}(F_{\theta_j}) + L_{\theta_i}(\hat{F}) + \sqrt{\frac{\text{KL}(P_{\theta_i}, P_{\theta_j})}{2}} < \left(\frac{2\sqrt{2}}{\pi} + \frac{1}{4}(g(\xi) - \sqrt{2}\xi^2) + \frac{\xi^2}{\sqrt{2}} \right) \frac{\sqrt{s\epsilon}}{\lambda} \leq \frac{1}{2}$$

because, for $\xi \leq 0.1$, by definition of ϵ ,

$$\left(\frac{2\sqrt{2}}{\pi} + \frac{1}{4}(g(\xi) - \sqrt{2}\xi^2) + \frac{\xi^2}{\sqrt{2}} \right) \frac{\sqrt{s\epsilon}}{\lambda} \leq \frac{\sqrt{s\epsilon}}{\lambda} \leq \frac{1}{2}.$$

So, by Proposition 13,

$$L_{\theta_j}(\hat{F}) \geq L_{\theta_i}(F_{\theta_j}) - L_{\theta_i}(\hat{F}) - \sqrt{\frac{\text{KL}(P_{\theta_i}, P_{\theta_j})}{2}} \geq \gamma.$$

Also,

$$\begin{aligned} \max_{i \in [M]} \text{KL}(P_{\theta_i}, P_{\theta_0}) &\leq \xi^4 \frac{s\epsilon^2}{\lambda^2} \\ &\leq \frac{1}{18n} \log \left(\left(\frac{d-1}{s} \right)^{\frac{s}{5}} \right) \\ &\leq \frac{1}{9n} \log \left(\left(\left(\frac{d-1}{s} \right)^{\frac{s}{5}} - 1 \right) \right) \\ &\leq \frac{\log M}{9n} \end{aligned}$$

because, by definition of ϵ ,

$$\xi^4 \frac{s\epsilon^2}{\lambda^2} \leq \frac{s}{90n} \log \left(\frac{d-1}{s} \right).$$

So by Proposition 12 and the fact that $\xi \leq 0.1$,

$$\begin{aligned} \inf_{\hat{F}_n} \max_{i \in [0..M]} \mathbb{E}_{\theta_i} L_{\theta_i}(\hat{F}_n) &\geq 0.07\gamma \\ &\geq 0.07 \frac{0.1}{4} \frac{\sqrt{s}\epsilon}{\lambda} \\ &\geq \frac{1}{600} \min \left\{ \sqrt{\frac{8}{45}} \frac{\sigma^2}{\lambda^2} \sqrt{\frac{s}{n} \log \left(\frac{d-1}{s} \right)}, \frac{1}{2} \right\} \end{aligned}$$

and to complete the proof we use the fact that

$$\inf_{\hat{F}_n} \sup_{\theta \in \Theta_{\lambda, s}} \mathbb{E}_{\theta} L_{\theta}(\hat{F}_n) \geq \inf_{\hat{F}_n} \max_{i \in [0..M]} \mathbb{E}_{\theta_i} L_{\theta_i}(\hat{F}_n).$$

□

References

- G.H. Golub and C.F. Van Loan. *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, 3rd edition, 1996. ISBN 9780801854149.
- Pascal Massart. Concentration inequalities and model selection. 2007.
- Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer, 2009.