

The EM Algorithm gives Sample-Optimality for Learning Mixtures of Well-Separated Gaussians

Jeongyeol Kwon

The University of Texas at Austin

KWONCHUNGLI@UTEXAS.EDU

Constantine Caramanis

The University of Texas at Austin

CONSTANTINE@UTEXAS.EDU

Editors: Jacob Abernethy and Shivani Agarwal

Abstract

We consider the problem of spherical Gaussian Mixture models with $k \geq 3$ components when the components are well separated. A fundamental previous result established that separation of $\Omega(\sqrt{\log k})$ is necessary and sufficient for identifiability of the parameters with *polynomial* sample complexity [Regev and Vijayaraghavan \(2017\)](#). In the same context, we show that $\tilde{O}(kd/\epsilon^2)$ samples suffice for any $\epsilon \lesssim 1/k$, closing the gap from polynomial to linear, and thus giving the first optimal sample upper bound for the parameter estimation of well-separated Gaussian mixtures. We accomplish this by proving a new result for the Expectation-Maximization (EM) algorithm: we show that EM converges locally, under separation $\Omega(\sqrt{\log k})$. The previous best-known guarantee required $\Omega(\sqrt{k})$ separation [Yan et al. \(2017\)](#). Unlike prior work, our results do not assume or use prior knowledge of the (potentially different) mixing weights or variances of the Gaussian components. Furthermore, our results show that the finite-sample error of EM does not depend on non-universal quantities such as pairwise distances between means of Gaussian components.

Keywords: Gaussian Mixture Model, EM algorithm, optimal sample complexity

1. Introduction

Learning parameters of a mixture of Gaussian is a fundamental problem in machine learning. In this model, we are given random samples from $k \geq 2$ Gaussian components without observing the label, *i.e.*, the indicator of which component each sample comes from. In this paper, we focus on an important special case of this model where the covariance of each Gaussian component is a multiple of the identity matrix. Formally, we consider a Gaussian mixture model \mathcal{G}^* whose probability density function (p.d.f.) can be represented as $\sum_{j=1}^k \pi_j^* \mathcal{N}(\mu_j^*, \sigma_j^{*2} I_d)$, where d is the dimension, I_d is $d \times d$ identity matrix, and $\mathcal{N}(\mu, \Sigma)$ denotes the p.d.f. of a single Gaussian distribution with mean $\mu \in \mathbb{R}^d$ and covariance $\Sigma \in \mathbb{R}^{d \times d}$. Here π_j^* are mixing weights, $\mu_j^* \in \mathbb{R}^d$ are means, and σ_j^* are scale factors for (identity) covariances of each Gaussian component. This special case is often called the *spherical* Gaussian mixture model. Our goal is to estimate all parameters $\{(\pi_j^*, \mu_j^*, \sigma_j^*), \forall j \in [k]\}$ up to accuracy ϵ .

Learning a mixture of Gaussians has a very long and rich history (see [Section 1.1](#) for an overview of previous works). A variety of algorithms have been proposed for parameter learning. These either require separation assumptions on the means, or structural assumptions on the mean placement, requiring control of the tensor singular values (essentially requiring affine independence). Most tensor-based analysis has therefore been done in a smoothed setting (e.g., [Bhaskara et al. \(2014\)](#);

Ge et al. (2015)), however in the absence of such structural assumptions, mean separation is what controls the hardness of the Gaussian Mixtures problem. Without any separation assumptions, even in one dimension, worst case instances require $\Omega(e^k)$ samples Moitra and Valiant (2010); Hardt and Price (2015). At the other extreme, Dasgupta (1999) demonstrated that under $\Omega(\sqrt{d})$ separation, sample-efficient (polynomial number of samples) identifiability is possible, thus providing the first upper bound on required separation for efficient identifiability. An important recent work by Regev and Vijayaraghavan (2017) characterized the exact threshold for sample-efficient identifiability, establishing that with separation $\Omega(\sqrt{\log k})$ sample-efficient identifiability is possible, where as below that threshold, a super-polynomial number of samples are required.

Perhaps the most widely used algorithm for mixture models is the Expectation-Maximization (EM) algorithm Wu et al. (1983). Recently, Yan et al. (2017); Zhao et al. (2018) established the local convergence of the EM algorithm (*i.e.*, EM converges if initialized from a neighborhood of the ground truth) for mixtures of k spherical Gaussians. These results require $\tilde{\Omega}(\sqrt{k})$ separation between means, and assume all components have identity covariances (*i.e.*, $\sigma_j^* = 1, \forall j \in [k]$). Thus this leaves open the key question as to the (local) behaviour of EM with $\Omega(\sqrt{\log k})$ separation.

The only known (local) algorithm that is guaranteed to converge in the $\Omega(\sqrt{\log k})$ separation regime is the EM-like algorithm proposed in Regev and Vijayaraghavan (2017). However, the sample complexity of their analysis also has a high dependence on k and instance-specific parameters as we explain in more detail below. They also require the initialization to be $O(1/k^2)$ -close to the true parameters. By obtaining guarantees that depend only on $O(1)$ -close initialization, we are able to give an optimal upper bound on sample complexity for learning the parameters of a mixture of spherical Gaussians.

Main Contributions. In this work, we return to the classical EM algorithm in the same $\Omega(\sqrt{\log k})$ separation regime, which thanks to the lower bound of Regev and Vijayaraghavan (2017) we know is optimal. We obtain improved convergence guarantees in this regime, and thereby close some of the existing gaps in the literature. Specifically, our main contributions are as follows:

- We show that with separation $\Omega(\sqrt{\log k})$, population EM converges given a good initialization. This improves the previous results of Yan et al. (2017) that required $\Omega(\sqrt{k})$ -separation. For the initialization, our result only requires $O(1)$ -closeness to the ground truth, hence improving the $O(1/k^2)$ initialization requirement in Regev and Vijayaraghavan (2017). Finally, unlike all prior work we are aware of, our result does not assume prior knowledge of mixing weights or variance parameters, and these need not be the same; we show we can adaptively estimate these quantities along with the means. These improvements enable our last result below on the optimal sample complexity of learning Gaussian mixture model parameters.
- We show that (sample-splitting) finite-sample EM converges to the ground truth given an $O(1)$ -close initialization. Our result has sample complexity $n = \tilde{O}(d\pi_{min}^{-1}/\epsilon^2)$ (where $\pi_{min} = \min_i \pi_i^*$) to recover all parameters up to ϵ accuracy:

$$\forall i \in [k], \|\mu_i - \mu_i^*\| \leq \sigma_i^* \epsilon, |\pi_i - \pi_i^*| \leq \pi_i^* \epsilon, |\sigma_i - \sigma_i^*| \leq \sigma_i^* \epsilon / \sqrt{d}.$$

Note that a natural sample complexity lower bound for the Gaussian mixture model is $\Omega(d\pi_{min}^{-1}/\epsilon^2)$, since we need to collect at least $\Omega(d/\epsilon^2)$ samples from each component. We give the best possible sample complexity in terms of all parameters $\pi_{min}, k, d, \epsilon$. This significantly improves over previous results Yan et al. (2017); Regev and Vijayaraghavan (2017) where even in the

balanced setting ($\pi_{min} = O(1/k)$), the sample complexities are at least worse by $\Omega(k^4 \rho^6)$ with an instance-dependent parameter ρ ¹.

- We show the sample-complexity $\tilde{O}(\max\{kd/\epsilon^2, k^3 d\})$ for learning the parameters of spherical Gaussian mixtures with $\Omega(\sqrt{\log k})$ separation. For $\epsilon \leq 1/k$ this gives $\tilde{O}(kd/\epsilon^2)$ and hence is optimal. The sample complexity guarantee here does not require any initialization, or boundedness of the parameters. The breakthrough in [Regev and Vijayaraghavan \(2017\)](#) was the first to establish polynomial sample complexity learning, but only gave sample complexity in the form $poly(k, d, \rho, 1/\epsilon)$, which is at least $\Omega(k^9 d \rho^4 \epsilon^{-2})$. Our result closes the gap and shows that the information-theoretically necessary sample complexity is also sufficient: as long as the separation is $\Omega(\sqrt{\log k})$, then $\tilde{O}(kd/\epsilon^2)$ samples are sufficient; this matches a lower bound $\Omega(kd/\epsilon^2)$ up to logarithmic factors.

1.1. Prior Art

Separation-Based Algorithms Learning mixtures of Gaussians has a long and rich history since [Dasgupta \(1999\)](#) who gave a first polynomial-time algorithm under $\Omega(\sqrt{d})$ separation. This result has been followed by a number of works [Sanjeev and Kannan \(2001\)](#); [Vempala and Wang \(2004\)](#); [Dasgupta and Schulman \(2007\)](#); [Achlioptas and McSherry \(2005\)](#); [Kannan et al. \(2005\)](#) that improve the result in various separation regimes. Currently, the best algorithmic results are the recent works by [Diakonikolas et al. \(2018\)](#); [Hopkins and Li \(2018\)](#); [Kothari et al. \(2018\)](#), where they provide algorithms that estimate parameters assuming $\Omega(k^{1/\gamma})$ separation, using $\tilde{O}(poly(k^\gamma, d, 1/\epsilon))$ samples (and running time) for arbitrary $\gamma > 0$. In particular, their can recover parameters of spherical Gaussian mixtures under $\Omega(\sqrt{\log k})$ separation using $\Omega(poly(k^{\log k}, d, 1/\epsilon))$ samples. While all works mentioned here aim to get (nearly) polynomial-time algorithms, our work is more in line with [Regev and Vijayaraghavan \(2017\)](#): we focus on upper bounding the sample complexity.

Iterative Algorithms EM is one of the most popular algorithms for mixture problems. The first results on convergence were infinitesimally local, and asymptotic [Redner and Walker \(1984\)](#); [Xu and Jordan \(1996\)](#); [Ma et al. \(2000\)](#). Recently, the work in [Balakrishnan et al. \(2017\)](#) builds off the idea of coupling finite sample and population EM, characterizing the non-asymptotic convergence of EM within a fixed (non-infinitesimal) basin of attraction. A flurry of work has followed in recent years, making substantial progress in the theory of EM. For instance, [Xu et al. \(2016\)](#); [Daskalakis et al. \(2017\)](#) established the global convergence of the EM algorithm from a random initialization for Gaussian mixture models with two components, and [Kwon et al. \(2019\)](#) established the same for two component mixed linear regression. For more than two components, we cannot hope for such a global convergence guarantee, as shown by [Jin et al. \(2016\)](#). On the positive side, some recent works have established local convergence results, showing that EM converges from a well-initialized point, under some minimum separation conditions [Yan et al. \(2017\)](#); [Kwon and Caramanis \(2019\)](#). However, the best known guarantees for the EM algorithm for a mixture of Gaussians with k components require separation of order $\tilde{\Omega}(\sqrt{k})$, and, moreover, are restricted to the equal identity covariance setting across all components [Yan et al. \(2017\)](#). Another popular iterative heuristic is the k -means algorithm (also called Lloyd’s algorithm) [Kumar and Kannan \(2010\)](#); [Awasthi and Sheffet \(2012\)](#); [Lu and Zhou \(2016\)](#). The best known convergence result for this algorithm requires at least $\Omega(\sqrt{k})$ separation.

1. ρ depends on instance-specific parameters such as $\max_{i \neq j} \|\mu_i^* - \mu_j^*\|$ or $\max_{i \neq j} \sigma_i^* / \sigma_j^*$. Since our sample complexity result does not depend on ρ , we do not require boundedness on parameters.

Thus the state of the art analyses of EM and Lloyd’s both leave a significant gap to the limit of $\Omega(\sqrt{\log k})$ mean separation. A variant of EM proposed in [Regev and Vijayaraghavan \(2017\)](#) takes a big step forward: it is shown to converge locally for spherical Gaussian mixtures with $\Omega(\sqrt{\log k})$ separation. While their algorithm is sample-efficient, the bound on samples is a (large) polynomial in k and instance-specific parameters, where as the information theoretic lower bound is $\Omega(kd/\epsilon^2)$.

Moment-Based Methods The method-of-moments is a powerful general-purpose technique for learning a family of parametric distributions. However, even in one dimension, an information-theoretic argument shows that an exponential (in k) number of samples is required to recover the parameters of a Gaussian mixture model in the absence of a minimum separation condition [Moitra and Valiant \(2010\)](#); [Hardt and Price \(2015\)](#). To circumvent such information-theoretic bottleneck of parameter learning, a vast line of work explores tensor-decomposition methods in a smoothed setting [Hsu and Kakade \(2013\)](#); [Ge et al. \(2015\)](#); [Anandkumar et al. \(2014\)](#); [Kalai et al. \(2010\)](#); [Anderson et al. \(2014\)](#). However, such an approach cannot work when the means of Gaussian components lie in a low dimensional subspace. Furthermore, statistical precision of moment-based methods has poor dependence on the geometric properties of an instance, such as singular values of a tensor matrix or the norm of means. Therefore, they are often used in conjunction with an iterative procedure such as the EM algorithm which produces much more accurate estimators as we show in this work.

Lower Bounds Without any separation, $\Omega(e^k)$ samples are necessary [Moitra and Valiant \(2010\)](#); [Hardt and Price \(2015\)](#). In fact, the separation condition has to be at least $\Omega(\sqrt{\log k})$ to get a fully polynomial sample complexity as shown in [Regev and Vijayaraghavan \(2017\)](#). There are also computational lower bound results due to [Diakonikolas et al. \(2017, 2018\)](#) framed in a statistical query (SQ) model [Feldman et al. \(2017\)](#) for general (non-spherical) Gaussian mixtures. It would be interesting to understand the implications of their results in the setting of spherical Gaussian mixtures (though we do not explore this here).

Distribution Learning Another branch of research in Gaussian mixture models is density estimation. [Feldman et al. \(2006\)](#); [Chan et al. \(2014\)](#); [Suresh et al. \(2014\)](#); [Diakonikolas et al. \(2019\)](#); [Li and Schmidt \(2017\)](#); [Ashtiani et al. \(2018b,a\)](#). In this problem, the goal is to learn a distribution \mathcal{G} that minimizes the total variation distance to \mathcal{G}^* . When this hypothesis \mathcal{G} is also in the class of mixtures of k Gaussian components, it is called *proper* learning. Most known proper-learning algorithms are sample-efficient or even sample-optimal [Ashtiani et al. \(2018a\)](#). Our result connects the result in [Ashtiani et al. \(2018a\)](#) to the parameter learning of a well-separated mixture of Gaussians.

Other Related Work Finally, we mention that there are a few related branches of research to learning a mixture of Gaussians such as graph clustering and community detection [Massoulié \(2014\)](#); [Mixon et al. \(2016\)](#); [Yin et al. \(2019\)](#). In such problems, an analogous notion of separation condition is considered for sample-efficient learning (see the recent survey in [Abbe \(2017\)](#)).

1.2. A Roadmap and Proof Outline

Our starting point is the by-now standard procedure where we analyze EM in the population setting first, and then connect the result to finite-sample EM through the concentration of measures.

Analysis of the Population EM Algorithm. The E-step in the EM algorithm constructs weights (soft-label) for all components to construct the expectations of the log-likelihood on the current estimators. If we have a good enough estimation of parameters, then we anticipate that most samples

should have approximately the right weights according to which components they come from. Given the current estimate of the mean parameters, let \mathcal{E}_{good} denote the set where the E-step assigns approximately the right label (weight). Then, the estimation error in means, for example, after one EM step can be decomposed as

$$\mathbb{E}_{\mathcal{D}}[wX] - \mu^* = \underbrace{(\mathbb{E}_{\mathcal{D}}[wX|\mathcal{E}_{good}] - \mu^*)P(\mathcal{E}_{good})}_{\text{errors from good samples}} + \underbrace{(\mathbb{E}_{\mathcal{D}}[wX|\mathcal{E}_{good}^c] - \mu^*)P(\mathcal{E}_{good}^c)}_{\text{errors from bad samples}},$$

where $\mathbb{E}_{\mathcal{D}}[wX]$ is the M-step operator for means, and w is a weight constructed for a sample X in the E-step. In the well-separated regime, we can show that $\mathbb{E}_{\mathcal{D}}[wX|\mathcal{E}_{good}] \approx \mu^*$ for good samples, while $P(\mathcal{E}_{good}^c) \approx 0$ for bad samples. While previous local analysis has a similar flavor [Yan et al. \(2017\)](#); [Balakrishnan et al. \(2017\)](#), it is significantly more challenging to construct these good events for non-symmetric and non-equal variance Gaussian mixtures, since the effective dimension of the EM operator does not shrink as in the symmetrized [Balakrishnan et al. \(2017\)](#) or equal-variance setting [Yan et al. \(2017\)](#). We show that if all parameters are well-initialized (see details in [Theorem 1](#)), then EM converges locally to the true parameters at a linear rate.

Analysis of the Finite-Sample EM Algorithm. In order to get the right order of statistical error for the EM algorithm, our measure concentration analysis treats good and bad samples separately. We adopt the technique used in [Kwon and Caramanis \(2019\)](#), and we split the concentration of measure argument into two parts: (i) concentration due to the sum of independent random variables and (ii) concentration of the empirical probability of the event (see [Proposition 8](#)). This categorization strategy is critical to get a finite-sample error that is instance-independent and does not depend on, for example, pairwise distances between means. This independence is in contrast with results of prior work, e.g., [Yan et al. \(2017\)](#); [Zhao et al. \(2018\)](#); [Balakrishnan et al. \(2017\)](#); [Regev and Vijayaraghavan \(2017\)](#). While technical, this concentration argument is the key differentiator in our results. We use this same technique again later in the paper, to obtain the optimal statistical error for learning a mixture of well-separated Gaussians.

Minimal Initialization Requirements and Sample-Optimal Learning. We next consider the problem of sample-optimal learning. To use our EM result, we need to show a sample-efficient strategy for obtaining a sufficiently good initialization for EM. In a recent work, [Ashtiani et al. \(2018a\)](#) gives the optimal sample-complexity result for proper learning. When applied to our setting, we can learn a candidate mixture of spherical Gaussians \mathcal{G} that satisfies $\|\mathcal{G} - \mathcal{G}^*\|_{TV} \leq \epsilon$ using $\tilde{O}(kd/\epsilon^2)$ samples, where $\|\cdot\|_{TV}$ is the total variation (TV) distance. We use this result to obtain a guarantee on the mean parameters. Specifically, we show that if a candidate spherical Gaussian mixture \mathcal{G} satisfies $\|\mathcal{G} - \mathcal{G}^*\|_{TV} \leq \pi_{min}/4$, then the mean parameters in \mathcal{G} satisfy

$$\|\mu_i - \mu_i^*\| \leq \frac{1}{4} \min_{i \neq j} \|\mu_i^* - \mu_j^*\|, \quad (1)$$

where μ_i are mean parameters of \mathcal{G} and μ_i^* are means of \mathcal{G}^* . From [\(1\)](#), we get an initialization condition that is not yet good enough to plug in our EM result. We develop a bridge to obtaining a sufficiently good initialization of all parameters with (essentially) the k -means algorithm, and conclude that $\tilde{O}(\pi_{min}^{-1}d/\epsilon^2 \vee k\pi_{min}^{-2}d)$ samples are sufficient to recover all parameters up to arbitrary ϵ accuracy. We note that this final result is possible because the initialization requirement we ultimately need is significantly weaker than what prior art requires.

2. Preliminaries and Notation

We consider the mixture of k spherical Gaussian mixtures with parameters $\{(\pi_j^*, \mu_j^*, \sigma_j^*), \forall j \in [k]\}$. True parameters are denoted by $(\cdot)^*$, while estimators are denoted using the same letters without $*$. Each j^{th} component has mean μ_j^* and covariance $\sigma_j^* I_d$ in \mathbb{R}^d . We use \mathcal{D} to represent the distribution of the mixture of Gaussians \mathcal{G}^* , and \mathcal{D}_j to represent the distribution of the j^{th} Gaussian component.

The EM algorithm is composed of two steps, the E-step that constructs the expectation of the log-likelihood on the current estimators, and the M-step that maximizes this expectation. For a mixture of spherical Gaussian distributions, each step of the EM algorithm is as follows:

$$\begin{aligned} \text{(E-step)} : \quad w_i(X) &= \frac{\pi_i \exp(-\|X - \mu_i\|^2 / (2\sigma_i^2) - d \log(\sigma_i^2) / 2)}{\sum_{j=1}^k \pi_j \exp(-\|X - \mu_j\|^2 / (2\sigma_j^2) - d \log(\sigma_j^2) / 2)}, \\ \text{(M-step)} : \quad \pi_i^+ &= \mathbb{E}_{\mathcal{D}}[w_i], \quad \mu_i^+ = \mathbb{E}_{\mathcal{D}}[w_i X] / \mathbb{E}_{\mathcal{D}}[w_i], \\ \sigma_i^{+2} &= \mathbb{E}_{\mathcal{D}}[w_i \|X - \mu_i^+\|^2] / (d \mathbb{E}_{\mathcal{D}}[w_i]), \end{aligned}$$

for all $i \in [k]$. In the above notation, we use $\mathbb{E}_{\mathcal{D}}[\cdot]$ to denote the expectation over the entire mixture distribution. We use $\mathbb{E}_{\mathcal{D}_j}[\cdot]$ to represent the expectation over the single j^{th} Gaussian component. In the E-step, $w_i := w_i(X)$ represents the probability of the sample X being generated from the i^{th} component as computed using the current estimates of parameters $\{(\pi_j, \mu_j, \sigma_j), \forall j \in [k]\}$. After the M-step, we use $(\cdot)^+$ to denote the corresponding updated estimators. When we consider the entire sequence of estimators, we use $(\cdot)^{(t)}$ to denote estimators in the t^{th} step. Finite-sample EM replaces the expectation with the empirical mean constructed with n i.i.d. samples. The corresponding estimators are denoted by $\{(\tilde{\pi}_j, \tilde{\mu}_j, \tilde{\sigma}_j), \forall j \in [k]\}$. In the E-step in the finite-sample EM, we use $w_{1,i}$ to represent the weight for the 1^{st} component constructed with the i^{th} sample X_i .

Finally, we introduce some conventions we use in this paper. We define $\pi_{\min} = \min_i \pi_i^*$, $\rho_\pi = \max_i \pi_i^* / \min_i \pi_i^*$, and $\rho_\sigma = \max_i \sigma_i^* / \min_i \sigma_i^*$. We use $R_{ij}^* = \|\mu_i^* - \mu_j^*\|_2$ to denote the pairwise distance between components. We often use $\|\cdot\|$ without subscript 2 to denote the l_2 norm of a vector in \mathbb{R}^d . The estimation error in the mean of the i^{th} component is defined as $\Delta_{\mu_i} := \mu_i^* - \mu_i$. We use $\mathbb{1}_{\mathcal{E}}$ for the indicator function for the event \mathcal{E} . We use standard complexity analysis notations $o(\cdot), O(\cdot), \tilde{O}(\cdot), \Omega(\cdot)$. Finally, we use “with high probability” in statements when the success probability of the algorithm is at least $1 - \delta$ where $\delta = n^{-\Omega(1)}$.

3. Convergence Analysis of the EM Algorithm

In this section, we give local convergence guarantees for both population EM and finite-sample EM. We first study population EM, and connect the result to the finite-sample setting.

3.1. Analysis of Population EM

We first state our main result for population EM. Compared to previous works, we also consider the setting of unknown and unequal variances, and hence must estimate these along with the means and mixing weights. We focus on handling the means and mixing weights in the main text, as the analysis for variance estimators is significantly more involved and delicate (see Appendix B.6 for the analysis of variance estimation).

Theorem 1 *There exists a universal constant $C \geq 64$ such that the following holds. Suppose a mixture of k spherical Gaussians has parameters $\{(\pi_j^*, \mu_j^*, \sigma_j^*) : j \in [k]\}$ such that*

$$\forall i \neq j \in [k], \|\mu_i^* - \mu_j^*\| \geq C(\sigma_i^* \vee \sigma_j^*) \cdot \left(\sqrt{\log k + \log(\rho_\sigma \rho_\pi)} \right), \quad (2)$$

and suppose the mean initialization $\mu_1^{(0)}, \dots, \mu_k^{(0)}$ satisfies

$$\forall i \in [k], \|\mu_i^{(0)} - \mu_i^*\| \leq \frac{\sigma_i^*}{16} \min_{i \neq j} \|\mu_i^* - \mu_j^*\| / (\sigma_i^* \vee \sigma_j^*). \quad (3)$$

Also, suppose the mixing weights and variances are initialized such that

$$\forall i \in [k], |\pi_i^{(0)} - \pi_i^*| \leq \pi_i^*/2, |(\sigma_i^{(0)})^2 - \sigma_i^{*2}| \leq 0.5\sigma_i^{*2}/\sqrt{d}. \quad (4)$$

Then, population EM converges in $T = O(\log(1/\epsilon))$ iterations to the true solution such that for all $i \in [k]$, we have $\|\mu_i^{(T)} - \mu_i^\| \leq \sigma_i^* \epsilon$, $|\pi_i^{(T)} - \pi_i^*|/\pi_i^* \leq \epsilon$, and $|(\sigma_i^{(T)})^2 - \sigma_i^{*2}|/\sigma_i^{*2} \leq \epsilon/\sqrt{d}$.*

Note that the convergence rate is linear and does not depend on instance-dependent quantities. Furthermore, for estimating variances, we eventually get $O(\epsilon/\sqrt{d})$ accurate estimates. Hence, the final output parameters are also $O(\epsilon)$ -close to the true mixture distribution in total variation distance.

The proof of Theorem 1 starts with the useful fact about the EM operator.

Fact 2 *True parameters $\{(\pi_i^*, \mu_i^*, \sigma_i^*) : i \in [k]\}$ are the fixed point of the EM operator, i.e.,*

$$\pi_i^* = \mathbb{E}_{\mathcal{D}}[w_i^*], \mu_i^* = \mathbb{E}_{\mathcal{D}}[w_i^* X] / \mathbb{E}_{\mathcal{D}}[w_i^*], \sigma_i^{*2} = \mathbb{E}_{\mathcal{D}}[w_i^* \|X - \mu_i^*\|^2] / (d \mathbb{E}_{\mathcal{D}}[w_i^*]),$$

where w_i^* is the weight constructed with true parameters in the E-step. Then, we can represent the estimation error after one EM iteration with the following lemma.

Lemma 3 *Define $\Delta_{w_i} = w_i - w_i^*$. The estimation errors after one EM iteration can be written as*

$$\begin{aligned} \pi_i^+ - \pi_i^* &= \mathbb{E}_{\mathcal{D}}[w_i] - \mathbb{E}_{\mathcal{D}}[w_i^*] = \mathbb{E}_{\mathcal{D}}[\Delta_{w_i}], \\ \mu_i^+ - \mu_i^* &= \mathbb{E}_{\mathcal{D}}[\Delta_{w_i}(X - \mu_i^*)] / \mathbb{E}_{\mathcal{D}}[w_i]. \end{aligned} \quad (5)$$

From this lemma, we observe that errors are proportional to the difference in weights given by the E-step that are constructed with true parameters and current estimates respectively. We focus on the parameters for the first component $i = 1$ and omit the subscript i when the context is clear. The next lemma is key; it defines what we call good samples.

Lemma 4 *Suppose X comes from the j^{th} component ($j \neq 1$). Let $v = X - \mu_j^*$ (thus, $v \sim \mathcal{N}(0, \sigma_j^{*2} I_d)$) and $\beta = R_{j1}^{*2} / (64(\sigma_1^* \vee \sigma_j^*)^2)$. Consider the following events:*

$$\begin{aligned} \mathcal{E}_{j,1} &:= \{-R_{j1}^{*2}/5 \leq \langle v, \mu_j^* - \mu_1^* \rangle\}, \\ \mathcal{E}_{j,2} &:= \{-R_{j1}^{*2}/64 \leq \langle v, \Delta_{\mu_1} \rangle\} \cap \{\langle v, \Delta_{\mu_j} \rangle \leq (\sigma_j^*/\sigma_1^*)^2 R_{j1}^{*2}/64\}, \\ \mathcal{E}_{j,3} &:= \left\{ d \left(1 - 2\sqrt{\beta/d} \right) \leq \|v\|^2 / \sigma_j^{*2} \leq d \left(1 + 2\sqrt{\beta/d} + 2\beta/d \right) \right\}. \end{aligned} \quad (6)$$

If the above three events occur, then the E-step assigns exponentially small weight to the other components, i.e., $w_1 \leq (\pi_1/\pi_j) \exp(-\beta)$ and $w_1^ \leq (\pi_1^*/\pi_j^*) \exp(-\beta)$. We call such X a good sample. This good event happens with probability at least $1 - 5 \exp(-\beta)$.*

Combined with the expression in (5), this lemma implies the following.

Corollary 5 *In the setting of Lemma 4, for any $j \neq 1$ and $s \in \mathbb{S}^{d-1}$,*

$$\begin{aligned} \mathbb{E}_{\mathcal{D}_j}[\Delta_w] &\leq (3(\pi_1^*/\pi_j^*) + 5) \exp(-\beta), \\ |\mathbb{E}_{\mathcal{D}_j}[\Delta_w \langle v, s \rangle]| &\leq (3(\pi_1^*/\pi_j^*)\sigma_j^* + 5R_{j1}^*) \exp(-\beta). \end{aligned} \quad (7)$$

The corollary bounds the estimation errors coming from the j^{th} component for some $j \neq 1$. We obtain this result by decomposing the error term as $\mathbb{E}_{\mathcal{D}_j}[\Delta_w] = \mathbb{E}_{\mathcal{D}_j}[\Delta_w \mathbb{1}_{\mathcal{E}_{\text{good}}}] + \mathbb{E}_{\mathcal{D}_j}[\Delta_w \mathbb{1}_{\mathcal{E}_{\text{good}}^c}]$, for estimation errors in mixing weights. Then we control the error from good samples by small Δ_w , and control the error from bad samples with the small probability of event $P(\mathcal{E}_{\text{good}}^c)$. Finally we sum up all errors from all components, and bound this sum with the following lemma.

Lemma 6 *For well-separated mixtures of Gaussians that satisfy the separation condition (2),*

$$\sum_{j \neq 1} (\pi_1^* + \pi_j^*) R_{j1}^{*q} \exp\left(-R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2\right) \leq c_q \sigma_1^{*q} \pi_1^*, \quad (8)$$

for $q \in \{0, 1, 2\}$ with sufficiently small absolute constants c_q .

By combining these results, we can guarantee that after one single iteration, we get $\|\mu_i^+ - \mu_i^*\| \leq 0.5\sigma_i^*$, $|\pi_i^+ - \pi_i^*| \leq 0.5\pi_i^*$, and $|\sigma_i^{+2} - \sigma_i^{*2}| \leq 0.5\sigma_i^{*2}/\sqrt{d}$ for all $i \in [k]$. Hence when $\|\mu_i - \mu_i^*\| \geq 0.5\sigma_i^*$, the lemmas stated in the main text suffice to guarantee that we get improved estimators through the EM operation. The full proof in this case is given in Appendix B.

When $\|\mu_i - \mu_i^*\| \leq 0.5\sigma_i^*$, we define

$$D_m = \max_i \left[\max \left(\|\mu_i - \mu_i^*\|/\sigma_i^*, |\pi_i - \pi_i^*|/\pi_i^*, \sqrt{d}|\sigma_i^2 - \sigma_i^{*2}|/\sigma_i^{*2} \right) \right], \quad (9)$$

and show that $D_m^+ \leq \gamma D_m$ for some absolute constant $\gamma < 1$. We first get a more fine-grained expression for (5) that is proportional to D_m by differentiating Δ_w using the mean-value theorem. Then we use the (essentially) same lemmas to show that $D_m^+ \leq \gamma D_m$, which guarantees a linear convergence in D_m to ϵ in $O(\log(1/\epsilon))$ steps. Since the proof involves significant algebraic manipulation, we defer the proof for this case to Appendix E.

3.2. Finite-Sample EM Analysis

Now we move on to the finite-sample EM algorithm. For analysis purposes, we use the common variant of the iterative algorithm which is often referred to as the *sample-splitting* scheme. This scheme divides the n samples into T batches of size n/T , and uses a new batch of samples in each iteration, which removes the probabilistic dependency between iterations. We note that the uniform concentration approach typically used to avoid sample-splitting often results in overly pessimistic statistical error. While the analysis with sample-splitting suffices for the optimal sample-complexity guarantee in this work, it will be interesting to remove this dependency in the future. We now state the main theorem for the finite-sample EM algorithm.

Theorem 7 *There exists a universal constant $C \geq 64$ such that the following holds. Suppose a mixture of k spherical Gaussians has parameters $\{(\pi_j^*, \mu_j^*, \sigma_j^*) : j \in [k]\}$ such that*

$$\forall i \neq j \in [k], \|\mu_i^* - \mu_j^*\| \geq C(\sigma_i^* \vee \sigma_j^*) \cdot c\sqrt{\log(\rho_\sigma/\pi_{\min})}, \quad (10)$$

with some given constant $c > 2$, and suppose the initializers $\{(\tilde{\pi}_j^{(0)}, \tilde{\mu}_j^{(0)}, \tilde{\sigma}_j^{(0)}) : j \in [k]\}$ satisfy

$$\forall i \in [k], \|\tilde{\mu}_i^{(0)} - \mu_i^*\| \leq \frac{\sigma_i^*}{16} \min_{i \neq j} \|\mu_i^* - \mu_j^*\| / (\sigma_i^* \vee \sigma_j^*), \quad (11)$$

$$\forall i \in [k], |\tilde{\pi}_i^{(0)} - \pi_i^*| \leq \pi_i^*/2, \quad |(\tilde{\sigma}_i^{(0)})^2 - \sigma_i^{*2}| \leq 0.5\sigma_i^{*2}/\sqrt{d}. \quad (12)$$

Suppose we use $n \geq C'(d\pi_{\min}^{-1} \log^2(k^2T/\delta)/\epsilon^2)$ samples with sufficiently large universal constant C' for every iteration. Then, sample-splitting finite-sample EM converges in $T = O(\log(1/\epsilon))$ iterations to the true solution such that for all $i \in [k]$, we have $\|\tilde{\mu}_i^{(T)} - \mu_i^*\| \leq \sigma_i^*\epsilon$, $|\tilde{\pi}_i^{(T)} - \pi_i^*| \leq \pi_i^*\epsilon$, and $|(\tilde{\sigma}_i^{(T)})^2 - \sigma_i^{*2}| \leq \sigma_i^{*2}\epsilon/\sqrt{d}$, with probability at least $1 - \delta - T/n^{c-2}k^{30}$.

The separation condition (10) in the statement differs by the constant $c > 2$ from condition (2), which we need to bound the failure probability by $1/n^c$. This inverse-polynomial failure probability arises from the concentration of the empirical probability of rare events.

As mentioned earlier, our finite-sample analysis develops a concentration bound that handles good and bad samples separately. Note that without splitting the analysis for different events, our statistical error may be unnecessarily large. For instance, if we simply apply standard sub-Gaussian tail bounds to $\frac{1}{n} \sum_{i=1}^n w_{1,i}(X_i - \mu_1^*)$ for mean updates, we end up having dependency on the norm of $X_i - \mu_1^*$, which could be as large as $O(\max_{j \neq 1} \|\mu_j^* - \mu_1^*\|)$. This dependence on the pairwise distances of the true means is present in much of the prior art. However, this approach overlooks the fact that the statistical properties of good samples and bad samples are very different. To overcome this issue, we adopt the main statistical tool introduced in [Kwon and Caramanis \(2019\)](#).

Proposition 8 (Proposition 5.3 in [Kwon and Caramanis \(2019\)](#)) *Let X be a random d -dimensional vector, and A be an event in the same probability space with $p = P(A) > 0$. Define random variable $Y = X|A$, i.e., X conditioned on event A , and $Z = \mathbb{1}_{X \in A}$. Let X_i, Y_i, Z_i be the i.i.d. samples from corresponding distributions. Then, the following holds,*

$$\begin{aligned} P\left(\left\|\frac{1}{n} \sum_{i=1}^n X_i \mathbb{1}_{X_i \in A} - \mathbb{E}[X \mathbb{1}_{X \in A}]\right\| \geq t\right) &\leq \max_{m \leq n_e} P\left(\left\|\frac{1}{n} \sum_{i=1}^m (Y_i - \mathbb{E}[Y])\right\| \geq t_1\right) \\ &+ P\left(\left\|\mathbb{E}[Y]\right\| \left|\frac{1}{n} \sum_{i=1}^n Z_i - p\right| \geq t_2\right) + P\left(\left|\sum_{i=1}^n Z_i\right| \geq n_e + 1\right), \end{aligned} \quad (13)$$

for any $0 \leq n_e \leq n$ and $t_1 + t_2 = t$.

We give a short overview of how we use this proposition in the proof for mixing weights. The estimation error in mixing weight after one EM iteration is $\tilde{\pi}_1^+ - \pi_1^* = (1/n \sum_{i=1}^n w_{1,i} - \mathbb{E}_{\mathcal{D}}[w_1]) + (\mathbb{E}_{\mathcal{D}}[w_1] - \mathbb{E}_{\mathcal{D}}[w_1^*])$. The second term is $(\mathbb{E}_{\mathcal{D}}[w_1] - \mathbb{E}_{\mathcal{D}}[w_1^*]) = \pi_1^+ - \pi_1^* \leq \gamma \pi_1^* D_m$, for γ and D_m as in (9). This is the error after one population EM iteration. Hence, if we can show that the first term is less than $\epsilon \pi_1^*$, linear convergence is guaranteed with convergence rate $\gamma < 1$ until D_m reaches the target statistical error ϵ .

Now we decompose a single random variable $w_{1,i}$ into several parts using disjoint indicator functions. Let us define $\mathcal{E}_{j,\text{good}} := \mathcal{E}_{j,1} \cap \mathcal{E}_{j,2} \cap \mathcal{E}_{j,3}$ where good events are as defined in (6), and let

$$w_{1,i} = w_{1,i} \mathbb{1}_{\mathcal{E}_1} + \sum_{j \neq 1} w_{1,i} \mathbb{1}_{\mathcal{E}_j \cap \mathcal{E}_{j,\text{good}}} + w_{1,i} \mathbb{1}_{\mathcal{E}_j \cap \mathcal{E}_{j,\text{good}}^c}.$$

Then, we treat each decomposed $O(k)$ terms as distinct quantities, and find a statistical fluctuation of the sum of each term over samples using Proposition 8.

For each term, the first key step is to find a sub-Gaussian or sub-exponential norm of random variables conditioned on each event. We use this to control the statistical fluctuation of the conditioned random variable on each event, which appears as the first term in (13). This mainly controls the concentration of good samples. For bad events $\mathcal{E}_{j,good}^c$, it is either that there are no bad sample with probability $O(1/poly(n))$, or $O(1)$ bad samples with high probability if $1/n^c < P(\mathcal{E}_{good}^c)$. This controls the concentration of bad samples. The full proof of Theorem 7 is given in Appendix C.

4. Sample Optimal Learning with the EM Algorithm

We now show we can learn well-separated Gaussian mixtures with nearly optimal sample complexity.

Theorem 9 *Suppose a mixture of k spherical Gaussians has parameters $\{(\pi_j^*, \mu_j^*, \sigma_j^*) : j \in [k]\}$ such that*

$$\forall i \neq j \in [k], \|\mu_i^* - \mu_j^*\| \geq C(\sigma_i^* \vee \sigma_j^*) \cdot c\sqrt{\log(\rho\sigma/\pi_{min})}, \quad (14)$$

with a universal constant $C \geq 64$ and some given constant $c > 2$. Then there exists a (possibly inefficient) algorithm that for any $\epsilon > 0$, returns parameters $\{(\pi_j, \mu_j, \sigma_j) : j \in [k]\}$ (up to permutation) such that

$$|\pi_i^* - \pi_i|/\pi_i^* \leq \epsilon, \|\mu_i^* - \mu_i\|/\sigma_i^* \leq \epsilon, |\sigma_i^* - \sigma_i|/\sigma_i^* \leq \epsilon/\sqrt{d}, \forall i \in [k],$$

using $n = \tilde{O}(\pi_{min}^{-1}d/\epsilon^2 \vee k\pi_{min}^{-2}d)$ samples with high probability.

While a polynomial sample upper bound with $\Omega(\sqrt{\log k})$ separation previously has been established in Regev and Vijayaraghavan (2017), our result guarantees the tightest sample complexity for $\epsilon = O(\pi_{min})$. In particular, our result implies that a trivial lower bound $\Omega(d\pi_{min}^{-1}/\epsilon^2)$ is indeed a tight upper bound in the well-separated regime. Furthermore, we do not impose any constraints on the norms of means, or require prior knowledge on mixing weights or variances. Hence Theorem 9 allows any possible realization of Gaussian mixture models that satisfies (14).

We first show that we can relax the initialization conditions (11), (12) such that it is sufficient to have a good initialization *only* for mean parameters. It makes the connection to proper-learning significantly easier since then we do not need any requirement on \mathcal{G} other than being close to \mathcal{G}^* in TV distance. In contrast, Regev and Vijayaraghavan (2017) requires \mathcal{G} to be close in TV distance as well as to have *all* parameters close to \mathcal{G}^* , which raises a technical challenge in connecting the proper-learning and parameter initialization. This initialization requirement results in a much higher sample complexity at least $\tilde{O}(\pi_{min}^{-2}k^3d^3\rho^{16})$ with an instance-specific parameter ρ .

4.1. Better Initialization with the k -Means Algorithm

Algorithm 1 is essentially the k -means algorithm except for Step 3 which estimates the variances. The next lemma is critical: it says Algorithm 1 can help initialize EM.

Lemma 10 *Suppose we are given $\mu_1^{(0)}, \dots, \mu_k^{(0)}$ such that*

$$\|\mu_i^{(0)} - \mu_i^*\| \leq \frac{1}{4} \min_{i \neq j} \|\mu_i^* - \mu_j^*\|, \quad (15)$$

Algorithm 1 One-Step k -means with Good Mean Initializers

Input: n i.i.d. samples from a mixture of well-separated Gaussians \mathcal{G}^* with parameters $\{(\pi_j^*, \mu_j^*, \sigma_j^*), j \in [k]\}$, and initial estimate of means $\{\mu_i^{(0)}, \forall i \in [k]\}$, satisfying (15).

Output: Good initializers for the EM algorithm satisfying (11) and (12).

1. Using n samples from mixtures of well-separated Gaussians, cluster points according to the rule: $C_i = \{X : \|X - \mu_i\| \leq \|X - \mu_j\|, \forall j \neq i\}$.
2. For each cluster C_i , let $\pi_i = |C_i|/n$ and $\mu_i = \text{mean}(C_i)$, the average over all elements in C_i .
3. For each cluster C_i , let samples in C_i stand in any pre-defined order. Compute pairwise distances between all adjacent samples. Let $F(x)$ be a cumulative distribution function of chi-square with d degrees of freedom. Collect all $|C_i| - 1$ computed values, and find $(\alpha_d |C_i|)^{\text{th}}$ quantity among them where $\alpha_d = F(d)$. Set σ_i^2 as the quantity divided by $2d$.

where $\{(\pi_j^*, \mu_j^*, \sigma_j^*) : j \in [k]\}$ are the parameters of well-separated mixtures of Gaussians. Then there exists a universal constant $C' > 0$ such that given with $n \geq C'(d\pi_{\min}^{-1} \log^2(k/\delta))$ samples, Algorithm 1 returns estimators satisfying $\|\mu_i - \mu_i^*\|/\sigma_i^* \leq 4$, $|\pi_i - \pi_i^*|/\pi_i^* \leq 0.5$, and $|\sigma_i^2 - \sigma_i^{*2}|/\sigma_i^{*2} \leq 0.5/\sqrt{d}$ with high probability.

Thus Algorithm 1 succeeds, as long as the initialization satisfies $1/4$ -closeness to the true means, which significantly relaxes the condition required for the EM algorithm. The key elements of the proof of the lemma for mixing weights and means are reminiscent of the ideas we exploit in population EM, as the k -means algorithm can be viewed as a variant of EM with hard-label assignment in the E-step. Estimating variances is trickier since we need to get estimators as good as $O(1/\sqrt{d})$. Controlling the noise to get a $O(1/\sqrt{d})$ estimate requires more than simply computing the average over each cluster (which is what we do in Step 2 to compute the means, and the mixing weights π_i). The key is contained in Step 3 of Algorithm 1, which is essentially computing the median of pairwise distances. The full proof of Lemma 10 is given in Appendix D.1.

4.2. Sample-Optimal Algorithm for Theorem 9

In this section, we first show that if we find a proper hypothesis that is $\pi_{\min}/4$ -close in total variation distance, then all means in the hypothesis are $1/4$ -close to true means of well-separated Gaussian mixtures. The is the content of the following lemma.

Lemma 11 Suppose a mixture of well-separated Gaussians \mathcal{G}^* with parameters $\{(\pi_i^*, \mu_i^*, \sigma_i^*), \forall i \in [k]\}$, and a candidate mixture of (any spherical) Gaussians \mathcal{G} with parameters $\{(\pi_i, \mu_i, \sigma_i), \forall i \in [k]\}$, satisfy $\|\mathcal{G} - \mathcal{G}^*\|_{TV} \leq \pi_{\min}/4$. Then

$$\max_i (\min_j \|\mu_i^* - \mu_j\|/\sigma_i^*) \leq 16\sqrt{\log(1/\pi_{\min})}. \quad (16)$$

The proof of Lemma 11 is given in Appendix D.2. While similar connection between TV distance and parameter distance are stated in Diakonikolas et al. (2017); Regev and Vijayaraghavan (2017), they consider distances smaller than $o(1/k^2)$ either in parameter space or total variation. Matching the sample complexity lower bounds requires a simpler connection to $O(1/k)$ -TV distance. As

mentioned earlier, \mathcal{G} may completely ignore some components if $\|\mathcal{G} - \mathcal{G}^*\|_{TV} > \pi_{min}$, hence $\pi_{min}/4$ (which is $O(1/k)$) is order-wise the minimum possible requirement for this approach.

Now we recall that [Ashtiani et al. \(2018a\)](#) provide a sample-optimal guarantee $\tilde{O}(kd/\epsilon^2)$ for the proper-learning of a mixture of *axis-aligned* Gaussians, *i.e.*, Gaussians with diagonal covariance matrices. From their result, it is straightforward to get the same sample-optimal guarantee for spherical Gaussian mixtures. We then combine it with Lemma 11, which gives a candidate distribution \mathcal{G} that satisfies $\|\mathcal{G} - \mathcal{G}^*\|_{TV} \leq \pi_{min}/4$. Lemma 10 provides the final bridge: executing Algorithm 1 produces an initialization good enough for the EM algorithm, as guaranteed by Theorem 7. The full proof of 9 is given in Appendix D.3.

4.3. Discussion on Initialization

We have shown that $\tilde{O}(d\pi_{min}^{-1}/\epsilon^2)$ samples are sufficient for learning a mixture of well-separated spherical Gaussians. While the combination of proper learning and the EM algorithm gives a sample-optimal algorithm, we note, however, that the algorithm given in [Ashtiani et al. \(2018a\)](#) is not computationally efficient. In fact, as far as we know, no (orderwise) sample-optimal polynomial time algorithm is known for proper-learning even in the $\Omega(\sqrt{\log k})$ -separated regime. Moreover, for general (non-spherical) Gaussian mixtures, work in [Diakonikolas et al. \(2017, 2018\)](#) gives statistical-query based lower bounds, though it is an interesting question to explore what these results imply in the spherical setting. Hence, while our work resolves the sample-complexity question, the computational complexity remains open.

The state-of-the-art algorithms in the $\Omega(\sqrt{\log k})$ separation regime appear in [Diakonikolas et al. \(2018\)](#); [Hopkins and Li \(2018\)](#); [Kothari et al. \(2018\)](#). They run in quasi-polynomial time with the required sample complexity $\tilde{O}(\text{poly}(k^{\log k}, d, 1/\epsilon))$. However, it is still unknown whether it is possible to learn a mixture of well-separated Gaussians in polynomial time with polynomial sample complexity. Achieving polynomial running time with polynomial sample complexity would be a great result that resolves a long lasting open problem in literature.

5. Conclusion

We provide local convergence guarantees for the EM algorithm for a mixture of well-separated spherical Gaussians. We show that EM enjoys desirable local convergence properties in several respects: minimal requirements on the separation condition, optimal sample complexity, and large initialization region. Consequently, our results provide the optimal sample upper bound for learning the parameters of well-separated Gaussian mixture models. Even under structural assumptions or larger sample complexity regimes when other methods apply, EM may still be an appealing local algorithm to amplify the estimation accuracy, as these global algorithms tend to incur large and instance-dependent statistical error. While our analysis is restricted to the well-separated regime, we conjecture that EM locally converges to the true parameters even with smaller separation. It will be an interesting future challenge to establish a similar result in a weaker separation regime.

Acknowledgement

This work was partially funded by NSF grants 1704778, 1646522, 1934932 and the Army Futures Command.

References

- Emmanuel Abbe. Community detection and stochastic block models: recent developments. *The Journal of Machine Learning Research*, 18(1):6446–6531, 2017.
- Dimitris Achlioptas and Frank McSherry. On spectral learning of mixtures of distributions. In *International Conference on Computational Learning Theory*, pages 458–469. Springer, 2005.
- Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *The Journal of Machine Learning Research*, 15(1):2773–2832, 2014.
- Joseph Anderson, Mikhail Belkin, Navin Goyal, Luis Rademacher, and James Voss. The more, the merrier: the blessing of dimensionality for learning large gaussian mixtures. In *Conference on Learning Theory*, pages 1135–1164, 2014.
- Hassan Ashtiani, Shai Ben-David, Nicholas Harvey, Christopher Liaw, Abbas Mehrabian, and Yaniv Plan. Nearly tight sample complexity bounds for learning mixtures of gaussians via sample compression schemes. In *Advances in Neural Information Processing Systems*, pages 3412–3421, 2018a.
- Hassan Ashtiani, Shai Ben-David, and Abbas Mehrabian. Sample-efficient learning of mixtures. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018b.
- Pranjal Awasthi and Or Sheffet. Improved spectral-norm bounds for clustering. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 37–49. Springer, 2012.
- Sivaraman Balakrishnan, Martin J Wainwright, Bin Yu, et al. Statistical guarantees for the EM algorithm: From population to sample-based analysis. *The Annals of Statistics*, 45(1):77–120, 2017.
- Aditya Bhaskara, Moses Charikar, Ankur Moitra, and Aravindan Vijayaraghavan. Smoothed analysis of tensor decompositions. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 594–603. ACM, 2014.
- Siu-On Chan, Ilias Diakonikolas, Rocco A Servedio, and Xiaorui Sun. Efficient density estimation via piecewise polynomial approximation. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 604–613. ACM, 2014.
- Sanjoy Dasgupta. Learning mixtures of gaussians. In *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pages 634–644. IEEE, 1999.
- Sanjoy Dasgupta and Leonard Schulman. A probabilistic analysis of EM for mixtures of separated, spherical gaussians. *Journal of Machine Learning Research*, 8(Feb):203–226, 2007.
- Constantinos Daskalakis, Christos Tzamos, and Manolis Zampetakis. Ten steps of EM suffice for mixtures of two gaussians. In *30th Annual Conference on Learning Theory*, 2017.

- Ilias Diakonikolas, Daniel M Kane, and Alistair Stewart. Statistical query lower bounds for robust estimation of high-dimensional gaussians and gaussian mixtures. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 73–84. IEEE, 2017.
- Ilias Diakonikolas, Daniel M Kane, and Alistair Stewart. List-decodable robust mean estimation and learning mixtures of spherical gaussians. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1047–1060. ACM, 2018.
- Ilias Diakonikolas, Gautam Kamath, Daniel Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high-dimensions without the computational intractability. *SIAM Journal on Computing*, 48(2):742–864, 2019.
- Jon Feldman, Rocco A Servedio, and Ryan O’Donnell. Pac learning axis-aligned mixtures of gaussians with no separation assumption. In *International Conference on Computational Learning Theory*, pages 20–34. Springer, 2006.
- Vitaly Feldman, Elena Grigorescu, Lev Reyzin, Santosh S Vempala, and Ying Xiao. Statistical algorithms and a lower bound for detecting planted cliques. *Journal of the ACM (JACM)*, 64(2):8, 2017.
- Rong Ge, Qingqing Huang, and Sham M Kakade. Learning mixtures of gaussians in high dimensions. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 761–770. ACM, 2015.
- Moritz Hardt and Eric Price. Tight bounds for learning a mixture of two gaussians. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 753–760. ACM, 2015.
- Samuel B Hopkins and Jerry Li. Mixture models, robustness, and sum of squares proofs. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1021–1034, 2018.
- 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.
- Tadeusz Inglot. Inequalities for quantiles of the chi-square distribution. *Probability and Mathematical Statistics*, 30(2):339–351, 2010.
- Chi Jin, Yuchen Zhang, Sivaraman Balakrishnan, Martin J Wainwright, and Michael I Jordan. Local maxima in the likelihood of gaussian mixture models: Structural results and algorithmic consequences. In *Advances in neural information processing systems*, pages 4116–4124, 2016.
- Adam Tauman Kalai, Ankur Moitra, and Gregory Valiant. Efficiently learning mixtures of two gaussians. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 553–562. ACM, 2010.
- Ravindran Kannan, Hadi Salmasian, and Santosh Vempala. The spectral method for general mixture models. In *International Conference on Computational Learning Theory*, pages 444–457. Springer, 2005.

- Pravesh K Kothari, Jacob Steinhardt, and David Steurer. Robust moment estimation and improved clustering via sum of squares. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1035–1046. ACM, 2018.
- Amit Kumar and Ravindran Kannan. Clustering with spectral norm and the k-means algorithm. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 299–308. IEEE, 2010.
- Jeongyeol Kwon and Constantine Caramanis. Em converges for a mixture of many linear regressions. *arXiv preprint arXiv:1905.12106*, 2019.
- Jeongyeol Kwon, Wei Qian, Constantine Caramanis, Yudong Chen, and Damek Davis. Global convergence of the EM algorithm for mixtures of two component linear regression. In *Conference on Learning Theory*, pages 2055–2110, 2019.
- Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.
- Jerry Li and Ludwig Schmidt. Robust and proper learning for mixtures of gaussians via systems of polynomial inequalities. In *Conference on Learning Theory*, pages 1302–1382, 2017.
- Yu Lu and Harrison H Zhou. Statistical and computational guarantees of lloyd’s algorithm and its variants. *arXiv preprint arXiv:1612.02099*, 2016.
- Jinwen Ma, Lei Xu, and Michael I Jordan. Asymptotic convergence rate of the em algorithm for gaussian mixtures. *Neural Computation*, 12(12):2881–2907, 2000.
- Laurent Massoulié. Community detection thresholds and the weak ramanujan property. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 694–703, 2014.
- Dustin G Mixon, Soledad Villar, and Rachel Ward. Clustering subgaussian mixtures by semidefinite programming. *arXiv preprint arXiv:1602.06612*, 2016.
- Ankur Moitra and Gregory Valiant. Settling the polynomial learnability of mixtures of gaussians. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 93–102. IEEE, 2010.
- Richard A Redner and Homer F Walker. Mixture densities, maximum likelihood and the em algorithm. *SIAM review*, 26(2):195–239, 1984.
- Oded Regev and Aravindan Vijayaraghavan. On learning mixtures of well-separated gaussians. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 85–96. IEEE, 2017.
- Arora Sanjeev 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.
- Ananda Theertha Suresh, Alon Orlitsky, Jayadev Acharya, and Ashkan Jafarpour. Near-optimal-sample estimators for spherical gaussian mixtures. In *Advances in Neural Information Processing Systems*, pages 1395–1403, 2014.

- Santosh Vempala and Grant Wang. A spectral algorithm for learning mixture models. *Journal of Computer and System Sciences*, 68(4):841–860, 2004.
- Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- JG Wendel. Note on the gamma function. *The American Mathematical Monthly*, 55(9):563–564, 1948.
- CF Jeff Wu et al. On the convergence properties of the EM algorithm. *The Annals of statistics*, 11(1): 95–103, 1983.
- Ji Xu, Daniel J Hsu, and Arian Maleki. Global analysis of expectation maximization for mixtures of two gaussians. In *Advances in Neural Information Processing Systems*, pages 2676–2684, 2016.
- Lei Xu and Michael I Jordan. On convergence properties of the em algorithm for gaussian mixtures. *Neural computation*, 8(1):129–151, 1996.
- Bowei Yan, Mingzhang Yin, and Purnamrita Sarkar. Convergence of gradient EM on multi-component mixture of gaussians. In *Advances in Neural Information Processing Systems*, pages 6956–6966, 2017.
- Mingzhang Yin, YX Wang, and Purnamrita Sarkar. A theoretical case study of structured variational inference for community detection. *arXiv preprint arXiv:1907.12203*, 2019.
- Ruofei Zhao, Yuanzhi Li, and Yuekai Sun. Statistical convergence of the EM algorithm on gaussian mixture models. *arXiv preprint arXiv:1810.04090*, 2018.

Appendix A. Useful Technical Lemmas

We state some useful lemmas:

Lemma 12 (Sub-Gaussian tails Vershynin (2010)) *Let v be a sub-Gaussian random vector with parameter σ^2 in d -dimensional space. Then for any unit vector $s \in \mathbb{S}^{d-1}$ and $\alpha > 0$,*

$$P(\langle v, s \rangle \geq \alpha) \leq \exp(-\alpha^2/2\sigma^2).$$

Lemma 13 (chi-Square tails Laurent and Massart (2000)) *Let v be a chi-square random variable with d degrees of freedom. Then for any $\alpha > 0$,*

$$\begin{aligned} P(v \geq d + 2\sqrt{d\alpha} + 2\alpha) &\leq \exp(-\alpha), \\ P(v \leq d - 2\sqrt{d\alpha}) &\leq \exp(-\alpha). \end{aligned}$$

Lemma 14 (Lower bound for chi-Square tails Inglot (2010)) *Let v be a chi-square random variable with degree of freedom $d \geq 2$. Then for any $u \geq d - 1$,*

$$\frac{1 - e^{-2}}{2} \frac{u}{u - d + 2\sqrt{d}} \mathcal{E}_d(u) \leq P(v \geq u) \leq \frac{1}{\sqrt{\pi}} \frac{u}{u - d + 2} \mathcal{E}_d(u),$$

where $\mathcal{E}_d(u) = \exp(-(1/2)(u - d - (d - 2) \log u + (d - 1) \log d))$.

Lemma 15 (Sub-Gaussian norm Vershynin (2010)) *v is a sub-Gaussian random vector in d -dimensional space if and only if for any unit vector $s \in \mathbb{S}^{d-1}$, there exists a finite value $K > 0$ such that*

$$\sup_{p \geq 1} p^{-1/2} \mathbb{E}[|\langle v, s \rangle|^p]^{1/p} \leq K.$$

We denote the sub-Gaussian norm of v as $\|v\|_{\psi_2} \leq K$. Furthermore, the tail probability is bounded by

$$P(\langle v, s \rangle \geq t) \leq \exp(-ct^2/K^2),$$

for some universal constant $c > 0$.

Lemma 16 (Sub-exponential norm Vershynin (2010)) *X is a sub-exponential random variable if and only if there exists a finite value $K > 0$ such that*

$$\sup_{p \geq 1} p^{-1} \mathbb{E}[|X|^p]^{1/p} \leq K.$$

We denote the sub-exponential norm of v as $\|v\|_{\psi_1} \leq K$. Furthermore, the tail probability is bounded by

$$P(X \geq t) \leq \exp(-c \min(t/K, t^2/K^2)),$$

for some universal constant $c > 0$.

Lemma 17 Suppose $v \sim \mathcal{N}(0, I_d)$, $u, s \in \mathbb{S}^{d-1}$ are any fixed unit vectors, and $\alpha > 0$ is some constant. Then the following holds:

$$\mathbb{E}_{v \sim \mathcal{N}(0, I_d)}[|\langle v, s \rangle|^p | |\langle v, u \rangle| \geq \alpha] \leq (2\alpha)^p + 4\alpha \exp(-\alpha^2/2)(2p)^{p/2}.$$

Proof Let us first decompose $s = s_u + s_p$, where s_u is parallel to u and s_p is orthogonal to u . We can rewrite the target quantity as

$$\begin{aligned} \left(\frac{\mathbb{E}[|\langle v, s_u \rangle + \langle v, s_p \rangle|^p \mathbf{1}_{\langle v, u \rangle \geq \alpha}]}{P(\langle v, u \rangle \geq \alpha)} \right)^{1/p} &\leq \frac{\mathbb{E}[|\langle v, s_u \rangle|^p \mathbf{1}_{\langle v, u \rangle \geq \alpha}]^{1/p} + \mathbb{E}[|\langle v, s_p \rangle|^p \mathbf{1}_{\langle v, u \rangle \geq \alpha}]^{1/p}}{P(\langle v, u \rangle \geq \alpha)^{1/p}} \\ &= \|s_u\| (\mathbb{E}[|\langle v, u \rangle|^p | \langle v, u \rangle \geq \alpha])^{1/p} + \mathbb{E}[|\langle v, s_p \rangle|^p]^{1/p}. \end{aligned}$$

Thus, it boils down to upper-bound $\mathbb{E}_{Z \sim \mathcal{N}(0,1)}[Z^p | Z \geq \alpha]$. This can be bounded by

$$\begin{aligned} \mathbb{E}_Z[Z^p \mathbf{1}_{Z \geq \alpha}] / P(Z \geq \alpha) &= \mathbb{E}_Z[Z^p \mathbf{1}_{2\alpha \geq Z \geq \alpha}] / P(Z \geq \alpha) + \mathbb{E}_Z[Z^p \mathbf{1}_{Z \geq 2\alpha}] / P(Z \geq \alpha) \\ &\leq (2\alpha)^p + \sqrt{\mathbb{E}_Z[Z^{2p}]} \sqrt{P(Z \geq 2\alpha)} / P(Z \geq \alpha) \\ &\leq (2\alpha)^p + 2^{p/2} \Gamma(p+1/2)^{1/2} / \sqrt[4]{\pi} \exp(-\alpha^2) / (\exp(-\alpha^2/2) / (\sqrt{2\pi} 2\alpha)) \\ &\leq (2\alpha)^p + 4(2^{p/2}) p^{p/2} \alpha \exp(-\alpha^2/2). \end{aligned}$$

■

Lemma 18 Let $v \sim \mathcal{N}(0, I_d)$. Then for any $p \geq 1$,

$$\mathbb{E}[||v||^p] = 2^{p/2} \Gamma((p+d)/2) / \Gamma(d/2) \leq (d+p)^{p/2}.$$

Proof The first equality is standard and given in [Yan et al. \(2017\)](#). From here, we can proceed as

$$\begin{aligned} \Gamma((p+d)/2) &= (d/2 + p/2)(d/2 + p/2 - 1) \dots (d/2 + \alpha + 1) \Gamma(d/2 + \alpha) \\ &\leq ((d+p)/2)^{\lfloor p/2 \rfloor} \Gamma(d/2 + \alpha), \end{aligned}$$

where $\alpha = p/2 - \lfloor p/2 \rfloor \in [0, 1)$ and $\lfloor p/2 \rfloor$ is the largest integer that does not exceed $p/2$. We then use Gautschi's inequality for the ratio of Gamma functions [Wendel \(1948\)](#) which states

$$\Gamma(x+1) / \Gamma(x+s) \leq (x+1)^{1-s}.$$

Applying this with $x+s = d/2$ and $x+1 = d/2 + \alpha$, we get

$$((d+p)/2)^{\lfloor p/2 \rfloor} \Gamma(d/2 + \alpha) / \Gamma(d/2) \leq ((d+p)/2)^{\lfloor p/2 \rfloor} (d+\alpha)^\alpha \leq ((d+p)/2)^{p/2}.$$

Thus, we can conclude

$$\mathbb{E}[||v||^p] = 2^{p/2} \Gamma((p+d)/2) / \Gamma(d/2) \leq (d+p)^{p/2}.$$

■

The following lemmas are the upper bound on the L_p norm of random variables conditioned on some events. We use them in an important way, as they help us bound expected errors from bad events.

Lemma 19 *Let $v \sim \mathcal{N}(0, I_d)$. Let $p \geq 1$ and $r^2 \geq d$. Then,*

$$\mathbb{E}[\|v\|^p | \|v\|^2 \geq r^2] \leq (2r)^p + 4(d + 2p)^{p/2} \exp(-r^2/8).$$

Proof Similarly to the one-dimensional case, we start with

$$\begin{aligned} \mathbb{E}[\|v\|^p | \|v\|^2 \geq r^2] &\leq (2r)^p + \sqrt{E[\|v\|^{2p}]} \frac{\sqrt{P(\|v\|^2 \geq 4r^2)}}{P(\|v\|^2 \geq r^2)} \\ &\leq (2r)^p + (d + 2p)^{p/2} \frac{\sqrt{P(\|v\|^2 \geq 4r^2)}}{P(\|v\|^2 \geq r^2)}. \end{aligned}$$

We can use the inequalities for lower and upper bounds for the tail probability of a chi-square distribution from Proposition 3.1 in [Inglot \(2010\)](#). From the inequality,

$$\begin{aligned} P(\|v\|^2 \geq 4r^2) &\leq 2 \exp\left(-\frac{1}{2}(4r^2 - d - (d-2)\log(4r^2) + (d-1)\log d)\right), \\ P(\|v\|^2 \geq r^2) &\geq \frac{1}{2} \exp\left(-\frac{1}{2}(r^2 - d - (d-2)\log(r^2) + (d-1)\log d)\right). \end{aligned}$$

Using this, we can bound

$$\begin{aligned} \frac{\sqrt{P(\|v\|^2 \geq 4r^2)}}{P(\|v\|^2 \geq r^2)} &\leq 2\sqrt{2} \exp\left(-\frac{1}{2}r^2 - \frac{1}{4}(d + (d-2)\log(r^2)) + \frac{1}{4}((d-2)\log(4d) + \log d)\right), \\ &\leq 2\sqrt{2} \exp(-r^2/8). \end{aligned}$$

Plugging in this relation, the lemma follows. ■

Lemma 20 *Let $v \sim \mathcal{N}(0, I_d)$. Let $p \geq 1$ and $r^2 = d + 2\sqrt{\alpha d} + 2\alpha$ with $\alpha > 8$. Then for any fixed unit vector $s \in \mathbb{S}^{d-1}$,*

$$\mathbb{E}[|\langle v, s \rangle|^p | \|v\|^2 \geq r^2] \leq (64\alpha)^{p/2} + 4(8\alpha + 2p)^{p/2}.$$

Proof Due to the rotational invariance of standard Gaussian distribution, without loss of generality, we can investigate $\mathbb{E}[|v_1|^p | \|v\|^2 \geq r^2]$. Let $a = r^2 - d$ and $b = a - (d-2)\log(1 + a/d)$. We first look at the case when $\alpha < d/8$ (thus $d > 64$ if $\alpha > 8$). In this case, first observe that $a/d = (2\sqrt{\alpha d} + 2\alpha)/d < 1$ and,

$$\begin{aligned} a^2/d &= 4(\alpha d + \alpha^2 + 2\alpha\sqrt{\alpha d})/d = 4(\alpha + \alpha/8 + \alpha/\sqrt{2}) \leq 8\alpha, \\ b &= a - (d-2)\log(1 + a/d) \leq a - (d-2)(a/d - (a/d)^2/6) \\ &\leq 2a/d + (a^2/6d)(d-2)/d \leq 2\alpha. \end{aligned}$$

Then we change the quantity as

$$\frac{\mathbb{E}[|v_1|^p \mathbb{E}[\mathbf{1}_{\|v\|^2 \geq r^2 - v_1^2} | v_1]]}{P(\|v\|^2 \geq r^2)} = \frac{\mathbb{E}[|v_1|^p P(\|u\|^2 \geq r^2 - v_1^2)]}{P(\|v\|^2 \geq r^2)},$$

where $u \sim \mathcal{N}(0, I_{d-1})$. Continuing the process,

$$\begin{aligned} \frac{\mathbb{E}[|v_1|^p P(\|u\|^2 \geq r^2 - v_1^2)]}{P(\|v\|^2 \geq r^2)} &= \frac{\mathbb{E}[|v_1|^p \mathbf{1}_{v_1^2 \geq b}] + \mathbb{E}[|v_1|^p \mathbf{1}_{v_1^2 \leq b} P(\|u\|^2 \geq r^2 - v_1^2)]}{P(\|v\|^2 \geq r^2)} \\ &\leq \frac{\mathbb{E}[|v_1|^p |v_1^2 \geq b] P(v_1^2 \geq b) + \mathbb{E}[|v_1|^p \mathbf{1}_{v_1^2 \leq b} P(\|u\|^2 \geq r^2 - v_1^2)]}{P(\|v\|^2 \geq r^2)}. \end{aligned}$$

Now using Lemma 14, we have

$$\begin{aligned} P(\|v\|^2 \geq r^2) &\geq \frac{1 - e^{-2}}{2} \frac{r^2}{r^2 - d + 2\sqrt{d}} \exp\left(-\frac{1}{2}(r^2 - d - (d-2)\log(r^2) + (d-1)\log d)\right), \\ &\geq \frac{1 - e^{-2}}{4} \exp\left(-\frac{1}{2}(a - (d-2)\log(r^2/d))\right), \end{aligned}$$

and

$$\begin{aligned} P(\|u\|^2 \geq r^2 - v_1^2) &\leq \frac{1}{\sqrt{\pi}} \frac{r^2 - v_1^2}{r^2 - v_1^2 - d + 1} \exp\left(-\frac{1}{2}(r^2 - v_1^2 - (d-1) - (d-3)\log(r^2 - v_1^2) + (d-2)\log d)\right) \\ &= \frac{1}{\sqrt{\pi}} \frac{1}{a - v_1^2 + 1} \exp\left(-\frac{1}{2}(1 + a - v_1^2 - (d-2)\log((r^2 - v_1^2)/d))\right). \end{aligned}$$

Therefore, we have that

$$\begin{aligned} \frac{\mathbb{E}[|v_1|^p \mathbf{1}_{v_1^2 \leq b} P(\|u\|^2 \geq r^2 - v_1^2)]}{P(\|v\|^2 \geq r^2)} &\leq \frac{4}{\sqrt{\pi}(1 - e^{-2})\sqrt{2\pi}} \int_{-\sqrt{b}}^{\sqrt{b}} \frac{|v_1|^{p-1} |v_1|}{1 + a - v_1^2} \exp\left(-\frac{1}{2}(1 - (d-2)\log((r^2 - v_1^2)/r^2))\right) \\ &\leq \frac{8b^{(p-1)/2}}{\pi} \int_0^{\sqrt{b}} \frac{v_1}{1 + a - v_1^2} \exp\left(-\frac{1}{2} + \frac{d-2}{2}\log(1 - v_1^2/r^2)\right) \\ &\leq \frac{8b^{(p-1)/2}}{\pi} \int_0^{\sqrt{b}} \frac{v_1}{1 + a - v_1^2} \leq \frac{-4b^{(p-1)/2}}{\pi} \ln(1 + a - v_1^2)|_0^{\sqrt{b}} \\ &\leq \frac{4b^{(p-1)/2}}{\pi} \frac{1 + a}{1 + (d-2)\log(1 + a/d)}. \end{aligned}$$

Recall $a < d$, $\log(1 + a/d) \geq a/d - (a/d)^2/6$. Plugging this in, we have

$$\begin{aligned} \frac{4b^{(p-1)/2}}{\pi} \frac{1 + a}{1 + (d-2)\log(1 + a/d)} &\leq \frac{4b^{(p-1)/2}}{\pi} \frac{1 + a}{1 + a(d-2)/d - (a/d)^2/6} \\ &\leq 4b^{(p-1)/2} \leq 4(2\alpha)^{(p-1)/2} \leq (2\alpha)^{p/2}. \end{aligned}$$

On the other hand, we know from above that,

$$\mathbb{E}[|v_1|^p |v_1^2 \geq b] \leq (2\sqrt{b})^p + (2p)^{p/2} \leq (8\alpha)^{p/2} + 4\sqrt{b} \exp(-b/2) (2p)^{p/2}.$$

Furthermore, we know that

$$\frac{P(v_1^2 \geq b)}{P(\|v\|^2 \geq r^2)} \leq \frac{2 \exp(-b/2)}{(1 - e^{-2})/2(\sqrt{d}/2) \exp(-b/2)} \leq 1/(1 - e^{-2}),$$

where we used that $d > 8\alpha > 64$. Thus, when $\alpha < d/8$, we can conclude that

$$\mathbb{E}[|v_1|^p \|v\|^2 \geq r^2] \leq 2(8\alpha)^{p/2} + 4(2p)^{p/2}.$$

Now let us consider the other side, when $\alpha > d/8$. In this case, we simply use Lemma 3.8. Observe that

$$\mathbb{E}[|v_1|^p \|v\|^2 \geq r^2] \leq \mathbb{E}[\|v\|^p \|v\|^2 \geq r^2] \leq (2r)^p + 4(d + 2p)^{p/2} \exp(-r^2/8).$$

Also, we can observe that when $\alpha > d/8$,

$$\begin{aligned} b &\leq a = 2\sqrt{\alpha d} + 2\alpha \leq (4\sqrt{2} + 2)\alpha \leq 8\alpha, \\ 2\alpha &\leq r^2 = d + a \leq 16\alpha. \end{aligned}$$

Therefore, we can apply these as is to obtain

$$\mathbb{E}[|v_1|^p \|v\|^2 \geq r^2] \leq (64\alpha)^{p/2} + 4(8\alpha + 2p)^{p/2} \exp(-\alpha/4).$$

Now combining all inequalities, we can conclude that

$$\mathbb{E}[|\langle v, s \rangle|^p \|v\|^2 \geq r^2] \leq (64\alpha)^{p/2} + 4(8\alpha + 2p)^{p/2}.$$

■

Lemma 21 *Let $v \sim \mathcal{N}(0, I_d)$. Then for any $p \geq 1$ and for any fixed unit vector $s \in \mathbb{S}^{d-1}$,*

$$\mathbb{E}[|v|^p |\langle v, s \rangle| \geq \alpha]^{1/p} \leq (2\alpha) + (2p)^{1/2} + (d + p - 1)^{1/2}.$$

Proof It can be easily shown that

$$\begin{aligned} \mathbb{E}[|v|^p |\langle v, s \rangle| \geq \alpha]^{1/p} &= \mathbb{E}[(v_1^2 + \|v_{2:d}\|^2)^{p/2} |v_1| \geq \alpha]^{1/p} \\ &\leq \mathbb{E}[(|v_1| + \|v_{2:d}\|)^p |v_1| \geq \alpha]^{1/p} \\ &\leq \mathbb{E}[|v_1|^p |v_1| \geq \alpha]^{1/p} + \mathbb{E}[\|v_{2:d}\|^p |v_1| \geq \alpha]^{1/p} \\ &\leq (2\alpha) + (2p)^{1/2} + \mathbb{E}[\|v_{2:d}\|^p]^{1/p} \\ &\leq (2\alpha) + (2p)^{1/2} + (d + p - 1)^{1/2}. \end{aligned}$$

Here we used the fact that $\|X\|_2 \leq \|X\|_1$ for any vector $X \in \mathbb{R}^2$ in the first step, Minkowski's inequality in the second step, and applied the previous lemmas. ■

Appendix B. Proof for the Convergence of the Population EM when $D_m \geq 1/2$.

B.1. Proofs for Fact 2 and Lemma 3

Proof for Fact 2. Note that for each sample X , $w_i^*(X) = \pi_i^* f_{\mathcal{D}_i}(X) / f_{\mathcal{D}}(X)$ where $f_{\mathcal{D}}$ is a p.d.f. of the mixture distribution, and $f_{\mathcal{D}_i}(X)$ is a p.d.f. of the i^{th} component. Then,

$$\begin{aligned}\mathbb{E}_{\mathcal{D}}[w_i^*] &= \int w_i^*(X) f_{\mathcal{D}}(X) = \int \pi_i^* f_{\mathcal{D}_i}(X) = \pi_i^*, \\ \mathbb{E}_{\mathcal{D}}[w_i^* X] &= \int w_i^*(X) X f_{\mathcal{D}}(X) = \int \pi_i^* X f_{\mathcal{D}_i}(X) = \pi_i^* \mu_i^*.\end{aligned}$$

Dividing by π_i^* recovers μ_i^* . For variances,

$$\mathbb{E}_{\mathcal{D}}[w_i^* \|X - \mu_i^*\|^2] = \int \pi_i^* \|X - \mu_i^*\|^2 f_{\mathcal{D}_i}(X) = \pi_i^* d\sigma_i^{*2}.$$

Dividing by $d\mathbb{E}_{\mathcal{D}}[w_i^*] = d\pi_i^*$ gives σ_i^{*2} . ■

Proof for Lemma 3. For mixing weights, $\pi_i^+ = \mathbb{E}_{\mathcal{D}}[w_i]$ by construction, $\pi_i^* = \mathbb{E}_{\mathcal{D}}[w_i^*]$ by Fact 2. Hence it is obvious that $\pi_i^+ - \pi_i^* = \mathbb{E}_{\mathcal{D}}[\Delta_{w_i}]$.

For means, note that $\mu_i^+ = \mathbb{E}_{\mathcal{D}}[w_i X] / \mathbb{E}_{\mathcal{D}}[w_i]$. Then,

$$\begin{aligned}\mu_i^+ - \mu_i^* &= \mathbb{E}_{\mathcal{D}}[w_i X] / \mathbb{E}_{\mathcal{D}}[w_i] - \mu_i^* = (\mathbb{E}_{\mathcal{D}}[w_i X] - \mathbb{E}_{\mathcal{D}}[w_i] \mu_i^*) / \mathbb{E}_{\mathcal{D}}[w_i] \\ &= (\mathbb{E}_{\mathcal{D}}[w_i X] - \mathbb{E}_{\mathcal{D}}[w_i \mu_i^*]) / \mathbb{E}_{\mathcal{D}}[w_i] \\ &= \mathbb{E}_{\mathcal{D}}[w_i (X - \mu_i^*)] / \mathbb{E}_{\mathcal{D}}[w_i].\end{aligned}$$

Now note that

$$\mathbb{E}_{\mathcal{D}}[w_i^* (X - \mu_i^*)] = \mathbb{E}_{\mathcal{D}}[w_i^* X] - \mathbb{E}[w_i^*] \mu_i^* = \mu_i^* \pi_i^* - \pi_i^* \mu_i^* = 0.$$

Hence we prove that $\mu_i^+ - \mu_i^* = \mathbb{E}_{\mathcal{D}}[\Delta_{w_i} (X - \mu_i^*)] / \mathbb{E}_{\mathcal{D}}[w_i]$. ■

B.2. Proof of Lemma 4

Proof Let us examine the exponents in w_1 . By definition of the weight constructed in the E-step, we can see that

$$w_1 \leq \frac{\pi_1}{\pi_j} \exp \left(-\frac{\|v + \mu_j^* - \mu_1\|^2}{2\sigma_1^2} + \frac{\|v + \mu_j^* - \mu_j\|^2}{2\sigma_j^2} - \frac{d}{2} \log(\sigma_1^2 / \sigma_j^2) \right).$$

Our goal is to find conditions for good event where $w_1 \leq \exp(\cdot)$. The *sufficient* condition for this is when the sum of these three terms is smaller than $-C \log(R_{j_1}^* k \rho_{\pi} / \sigma_1^*)$. If this is the case, then the weight given to this sample is less than $O(\sigma_1^* / (k R_{j_1}^*))$ which can cancel out errors from the mismatch in labels. The cases are divided based on whether $\sigma_1 \geq \sigma_j$ or $\sigma_1 \leq \sigma_j$. We first rearrange the inside of the exponent,

$$-\frac{\|v + \mu_j^* - \mu_1\|^2}{2\sigma_1^2} + \frac{\|v + \Delta_j\|^2}{2\sigma_j^2} - \frac{d}{2} \log(\sigma_1^2 / \sigma_j^2)$$

$$\begin{aligned}
&= -\frac{\|v + \mu_j^* - \mu_1\|^2}{2\sigma_1^2} + \frac{\|v + \Delta_j\|^2}{2\sigma_j^2} - \frac{d}{2} \log(\sigma_1^2/\sigma_j^2) \\
&= -\frac{\|v\|^2 + \|\mu_j^* - \mu_1\|^2 + 2\langle v, \mu_j^* - \mu_1 \rangle}{2\sigma_1^2} + \frac{\|v\|^2 + 2\langle v, \Delta_j \rangle + \|\Delta_j\|^2}{2\sigma_j^2} - \frac{d}{2} \log(\sigma_1^2/\sigma_j^2) \\
&\leq \underbrace{-\frac{7R_{j1}^{*2}}{16\sigma_1^2}}_I + \underbrace{\frac{\|\Delta_j\|^2}{2\sigma_j^2}}_{II} + \underbrace{(-\langle v, \mu_j^* - \mu_1 \rangle/\sigma_1^2)}_{III} + \underbrace{\langle v, \Delta_j \rangle/\sigma_j^2}_{IV} + \underbrace{\left(-\frac{\|v\|^2}{2} \left(\frac{1}{\sigma_1^2} - \frac{1}{\sigma_j^2}\right) - \frac{d}{2} \log(\sigma_1^2/\sigma_j^2)\right)}_V.
\end{aligned} \tag{17}$$

II , III and IV can be controlled in a fairly straight-forward manner. Check that,

$$II \leq \frac{\|\Delta_j\|^2}{2\sigma_j^2} \leq \frac{R_{j1}^{*2}}{512(\sigma_1^* \vee \sigma_j^*)^2} \frac{\sigma_j^{*2}}{\sigma_j^2} \leq \frac{3R_{j1}^{*2}}{512\sigma_1^2},$$

$$\begin{aligned}
P\left(III \geq 7R_{j1}^{*2}/32\sigma_1^2\right) &= P\left(\langle v, \mu_j^* - \mu_1 \rangle \leq -7R_{j1}^{*2}/32\right) \\
&\leq P\left(\langle v, \mu_j^* - \mu_1^* \rangle \leq -R_{j1}^{*2}/5\right) + P\left(\langle v, \Delta_1 \rangle \leq -R_{j1}^{*2}/64\right) \\
&\leq 2 \exp\left(-R_{j1}^{*2}/(64\sigma_j^{*2})\right),
\end{aligned}$$

$$\begin{aligned}
P\left(IV \geq R_{j1}^{*2}/20\sigma_1^2\right) &\leq P\left(\langle v, \Delta_j \rangle \geq R_{j1}^{*2}\sigma_j^2/20\sigma_1^2\right) \\
&\leq P\left(\langle v, \Delta_j \rangle \geq R_{j1}^{*2}(\sigma_j^*/\sigma_1^*)^2/64\right) \leq \exp\left(-R_{j1}^{*2}/(64\sigma_1^{*2})\right),
\end{aligned}$$

Thus, $I + II + III + IV$ is smaller than $(R_{j1}^{*2}/\sigma_1^2)(-7/16 + 7/32 + 1/20 + 3/512) \leq -5/32$ with high probability. Now the remaining main challenge is to bound V . We should consider cases separately when $\sigma_1 \geq \sigma_j$ and $\sigma_1 \leq \sigma_j$.

Let us first consider $\sigma_1 \geq \sigma_j$. Overall, we want $V \leq R_{j1}^{*2}/(8\sigma_1^2)$ so that the entire sum inside exponent is small enough to kill this sample. That is, we want that

$$-\frac{\|v\|^2}{2} \left(\frac{1}{\sigma_1^2} - \frac{1}{\sigma_j^2}\right) - \frac{d}{2} \log(\sigma_1^2/\sigma_j^2) \leq R_{j1}^{*2}/(8\sigma_1^2). \tag{18}$$

Let us introduce some auxiliary variables to simplify the expression. Let $x = (\sigma_1^2 - \sigma_j^2)/\sigma_1^2$ and $a = R_{j1}^{*2}/(8\sigma_1^2)$. Then the above equation can be written as

$$\begin{aligned}
\frac{\|v\|^2}{2}(x/\sigma_j^2) + \frac{d}{2} \log(1-x) \leq a &\iff \frac{\|v\|^2}{\sigma_j^2} \leq \frac{1}{x}(-d \log(1-x) + 2a) \\
&\iff \frac{1}{d} \frac{\|v\|^2}{\sigma_j^{*2}} \leq \frac{\sigma_j^2}{\sigma_j^{*2}} \frac{1}{x}(-\log(1-x) + 2a/d).
\end{aligned}$$

Note that $\|v\|^2/\sigma_j^{*2}$ is a degree- d chi-square random variable and we can apply standard tail bound for CDF of χ_d distribution (which is essentially upper tail bound for sub-exponential random variable). The following useful inequality will help us to bound it in more convenient form:

$$-\log(1-x) \geq x + x^2/2,$$

so that the sufficient condition for this is,

$$\frac{1}{d} \frac{\|v\|^2}{\sigma_j^{*2}} \leq \left(1 - 1/(2\sqrt{d})\right) \left(1 + \sqrt{2a/d} + a/d\right) \leq \frac{\sigma_j^2}{\sigma_j^{*2}} \left(1 + \frac{x}{2} + \frac{2a}{dx}\right).$$

More sufficient condition is

$$\frac{1}{d} \frac{\|v\|^2}{\sigma_j^{*2}} \leq 1 + \left(\sqrt{a/d} + a/2d\right),$$

where the initialization condition for $|\sigma_j^2 - \sigma_j^{*2}|/\sigma_j^{*2} \leq 1/(4\sqrt{d}) \ll \sqrt{a/d}$ is used for the first inequality. From the Lemma 13, this is true with probability at least $1 - \exp(-a/4) \geq 1 - \exp(-R_{j_1}^{*2}/(64\sigma_1^{*2}))$. In summary, when $\sigma_1 \geq \sigma_j$, we have $w_1 \leq (\pi_1/\pi_j) \exp(-R_{j_1}^{*2}/(64\sigma_1^{*2}))$ with probability $1 - 3 \exp(-R_{j_1}^{*2}/(64\sigma_1^{*2}))$. Then, $\mathbb{E}_{\mathcal{D}_j}[w_1] \leq 3(1 + \pi_1^*/\pi_j^*) \exp(-R_{j_1}^{*2}/(64\sigma_1^{*2}))$.

Now we consider the second case when $\sigma_1 \leq \sigma_j$. We again use the same formulation as in (18). This time, let $x = (\sigma_j^2 - \sigma_1^2)/\sigma_1^2$ and find the probability for

$$\begin{aligned} -\frac{\|v\|^2}{2\sigma_j^2}x + \frac{d}{2}\log(1+x) \leq a &\iff \frac{\|v\|^2}{\sigma_j^2} \geq \frac{1}{x}(d\log(1+x) - 2a) \\ &\iff \frac{\|v\|^2}{d\sigma_j^{*2}} \geq \frac{\sigma_j^2}{\sigma_j^{*2}} \frac{1}{x}(d\log(1+x) - 2a/d), \end{aligned}$$

where x ranges from 0 to infinity. To control this, divide cases when $0 \leq x \leq 3/4$ and $x \geq 3/4$. If $0 \leq x \leq 3/4$, then

$$\log(1+x) \leq x - x^2/2 + x^3/3 \leq x - x^2/4.$$

Using this, it is enough to give a probability bound for

$$\frac{\|v\|^2}{d\sigma_j^{*2}} \geq (1 + 1/(2\sqrt{d}))(1 - 2\sqrt{a/(2d)}) \geq \frac{\sigma_j^2}{\sigma_j^{*2}}(1 - x/4 - 2a/(xd)).$$

Therefore, the sufficient condition is

$$\frac{\|v\|^2}{d\sigma_j^{*2}} \geq (1 - \sqrt{a/2d}).$$

When $3/4 \leq x$, note that $\log(1+x)/x \leq 3/4$. Therefore,

$$\frac{\|v\|^2}{d\sigma_j^{*2}} \geq (1 - \sqrt{a/(xd)}) \geq \frac{\sigma_j^2}{\sigma_j^{*2}}(1 - 1/4 - 2a/(xd)).$$

Note that $a/x = R_{j_1}^{*2}/(8(\sigma_j^2 - \sigma_1^2)) \geq R_{j_1}^{*2}/(16\sigma_j^{*2})$. The sufficient condition for $\|v\|^2$ is thus,

$$\frac{\|v\|^2}{d\sigma_j^{*2}} \geq \max \left(1 - \sqrt{R_{j_1}^{*2}/(16d\sigma_1^{*2})}, 1 - \sqrt{R_{j_1}^{*2}/(16d\sigma_j^{*2})} \right),$$

which will hold with probability at least $1 - \exp(-R_{j_1}^{*2}/64\sigma_j^{*2})$. Note that these are all sufficient conditions to ensure $w_1 \leq 3(\pi_1^*/\pi_j^*) \exp(-R_{j_1}^{*2}/(64\sigma_j^{*2}))$.

Combining all cases, when events defined in (6) happen, then w_1 is small enough. As we have seen already, this is true with probability at least $1 - 5 \exp(-R_{j_1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2)$. \blacksquare

B.3. Proof of Lemma 5

Proof By the Lemma 4, the first equation is easy to show. Define $\mathcal{E}_{j,\text{good}} = \mathcal{E}_{j,1} \cap \mathcal{E}_{j,2} \cap \mathcal{E}_{j,3}$. For the ease of notation, let $\beta = R_{j_1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2$.

$$\begin{aligned} \mathbb{E}_{\mathcal{D}_j}[w_1] &= \mathbb{E}_{\mathcal{D}_j} [w_1 \mathbf{1}_{\mathcal{E}_{j,\text{good}}}] + \mathbb{E}_{\mathcal{D}_j} [w_1 \mathbf{1}_{\mathcal{E}_{j,\text{good}}^c}] \\ &\leq 3(\pi_1^*/\pi_j^*) \exp(-\beta) + P(\mathbf{1}_{\mathcal{E}_{j,\text{good}}^c}) \leq 3(\pi_1^*/\pi_j^*) \exp(-\beta) + 5 \exp(-\beta). \end{aligned}$$

For the second equation,

$$\begin{aligned} |\mathbb{E}_{\mathcal{D}_j}[w_1 \langle v, s \rangle]| &= |\mathbb{E}_{\mathcal{D}_j} [w_1 \langle v, s \rangle \mathbf{1}_{\mathcal{E}_{j,\text{good}}}]| + |\mathbb{E}_{\mathcal{D}_j} [w_1 \langle v, s \rangle \mathbf{1}_{\mathcal{E}_{j,\text{good}}^c}]| \\ &\leq 3(\pi_1^*/\pi_j^*) \exp(-\beta) \mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle|] + \mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle| | \mathcal{E}_{j,1}^c] P(\mathcal{E}_{j,1}^c) \\ &\quad + \mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle| | \mathcal{E}_{j,2}^c] P(\mathcal{E}_{j,2}^c) + \mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle| | \mathcal{E}_{j,3}^c] P(\mathcal{E}_{j,3}^c). \end{aligned}$$

$\mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle| | \mathcal{E}_{j,1}^c]$ can be bounded with Lemma 17, with $p = 1$ and $\alpha = R_{j_1}^*/5\sigma_j^*$.

$$\mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle| | \langle v, R_{j_1}^* \rangle \geq R_{j_1}^{*2}/5] \leq \sigma_j^* \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [|\langle v, s \rangle| | |\langle v, u \rangle| \geq \alpha] \leq \sigma_j^* (2\alpha + \sqrt{2}) \leq R_{j_1}^*.$$

Similarly, we can bound $\mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle| | \mathcal{E}_{j,2}^c] P(\mathcal{E}_{j,2}^c) \leq 2R_{j_1}^*$ using the same Lemma 17 with $p = 1$ and $\alpha = R_{j_1}^*/4\sigma_j^*$. For the third term, we use Lemma 20, with $p = 1$ and $\alpha = R_{j_1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2 = \beta$. Then,

$$\begin{aligned} \sigma_j^* \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [|\langle v, s \rangle| | \|v\|^2 \geq d + 2\sqrt{\alpha d} + 2\alpha] &\leq \sigma_j^* \left((64\alpha)^{1/2} + 4 \exp(-\alpha/2) (8\alpha + 2)^{1/2} \right) \leq 2R_{j_1}^*, \\ \sigma_j^* \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [|\langle v, s \rangle| | \|v\|^2 \leq d - 2\sqrt{\alpha d}] &\leq \sigma_j^* \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [|\langle v, s \rangle|] \leq \sigma_j^*, \end{aligned}$$

Collecting these three components, we can conclude that

$$|\mathbb{E}_{\mathcal{D}_j}[w_1 \langle v, s \rangle]| \leq (3(\pi_1^*/\pi_j^*)\sigma_j^* + 5R_{j_1}^*) \exp(-\beta).$$

The same argument holds for $\Delta_w = w_1 - w_1^*$ since $0 \leq w_1, w_1^* \leq 3(\pi_1^*/\pi_j^*) \exp(-\beta)$ ensures $|\Delta_w| \leq 3(\pi_1^*/\pi_j^*) \exp(-\beta)$. \blacksquare

B.4. Proof of Lemma 6

We will focus on $q = 2$ case. Due to the separation condition (2), we have $R_{j_1}^{*2}/(\sigma_1^* \vee \sigma_j^*)^2 \geq C^2 \log(\rho_\sigma/\pi_{min}) \geq 4096 \log(\rho_\sigma/\pi_{min})$. Let $x := R_{j_1}^{*2}/(\sigma_1^* \vee \sigma_j^*)^2$. One useful fact is, if $x \geq 4096$, then $x \geq 128 \ln x$. Hence,

$$\begin{aligned} \sum_{j \neq 1} R_{j_1}^{*2} \exp(-R_{j_1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2) &= \sum_{j \neq 1} (\sigma_j^* \vee \sigma_1^*)^2 x \exp(-x/64) \leq \sum_{j \neq 1} (\sigma_j^* \vee \sigma_1^*)^2 \exp(-x/128) \\ &\leq \sum_{j \neq 1} (\sigma_j^* \vee \sigma_1^*)^2 (\rho_\sigma/\pi_{min})^{-32} \ll c\sigma_1^{*2}, \end{aligned} \quad (19)$$

for some small constant c . The similar result holds for $q = 0$ and $q = 1$. Similarly,

$$\begin{aligned} \sum_{j \neq 1} \pi_j^* R_{j_1}^{*2} \exp(-R_{j_1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2) &= \sum_{j \neq 1} \pi_j^* (\sigma_j^* \vee \sigma_1^*)^2 x \exp(-x/64) \leq \sum_{j \neq 1} \pi_j^* (\sigma_j^* \vee \sigma_1^*)^2 \exp(-x/128) \\ &\leq \max_{j \neq 1} (\sigma_j^* \vee \sigma_1^*)^2 (\rho_\sigma/\pi_{min})^{-32} \ll c\sigma_1^{*2} \pi_1^*. \end{aligned} \quad (20)$$

Summing up the result of $\pi_1^*(19) + (20) \leq c'\pi_1^*\sigma_1^{*2}$ gives the Lemma for $q = 2$. The cases for $q = 0, 1$ can be shown similarly.

B.5. Convergence of Means and Mixing Weights

Proof First we consider the error that comes from other components.

When $j \neq 1$: We will primarily focus on bounding this quantity by analyzing the errors from each components separately. Then we will give a bound to estimators after one population EM iteration. Note that the Corollary 5 also holds for $\mathbb{E}_{\mathcal{D}_j}[\Delta_w]$ as in the corollary. With the Lemma 6, we can bound the errors for mixing weights from other components. For mixing weights,

$$\sum_{j \neq 1} \pi_j^* \mathbb{E}_{\mathcal{D}_j}[\|w_1 - w_1^*\|] \leq 5 \sum_{j \neq 1} \pi_j^* \exp(-R_{j_1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2) + 3\pi_1^* \sum_{j \neq 1} \exp(-R_{j_1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2) \leq c\pi_1^*,$$

for some small constant c .

Similarly, we can bound the errors to the mean estimator from other components. First observe that

$$\begin{aligned} \|\mu_1^+ - \mu_1^*\| &= \|\mathbb{E}_{\mathcal{D}}[\Delta_w(X - \mu_1^*)]\|/\mathbb{E}_{\mathcal{D}}[w_1] \\ &\leq \sum_j \pi_j^* \sup_{s \in \mathbb{S}^{d-1}} |\mathbb{E}_{\mathcal{D}_j}[\Delta_w \langle X - \mu_1^*, s \rangle]| \leq \sum_j \pi_j^* \sup_{s \in \mathbb{S}^{d-1}} |\mathbb{E}_{\mathcal{D}}[\Delta_w \langle v + \mu_j^* - \mu_1^*, s \rangle]|. \end{aligned}$$

The errors from other components are thus,

$$\begin{aligned} \sum_{j \neq 1} \pi_j^* \left(\sup_{s \in \mathbb{S}^{d-1}} |\mathbb{E}_{\mathcal{D}_j}[\Delta_w \langle v, s \rangle]| + R_{j_1}^* |\mathbb{E}_{\mathcal{D}_j}[\Delta_w]| \right) &\leq 10 \sum_{j \neq 1} (\pi_1^* + \pi_j^*) R_{j_1}^* \exp(-R_{j_1}^*/64(\sigma_1^* \vee \sigma_j^*)^2) \\ &\leq c'\sigma_1^* \pi_1^*, \end{aligned}$$

for some small constant c' .

When $j = 1$: From the correct component, we expect the weight is mostly close to 1, and 0 only rarely to bad samples. For this case, we can consider the weights given to other components. That is,

$$\begin{aligned}\mathbb{E}_{\mathcal{D}_1}[1 - w_1] &= \sum_{l \neq 1} \mathbb{E}_{\mathcal{D}_1}[w_l] \leq \sum_{l \neq 1} (3\pi_l^*/\pi_1^* + 5) \exp(-R_{l1}^{*2}/64(\sigma_l^* \vee \sigma_1^*)^2) \\ &\leq (1/\pi_1^*) \left(3 \sum_{l \neq 1} \pi_l^* \exp(-R_{l1}^{*2}/64(\sigma_l^* \vee \sigma_1^*)^2) + 5\pi_1^* \sum_{l \neq 1} \exp(-R_{l1}^{*2}/64(\sigma_l^* \vee \sigma_1^*)^2) \right) \\ &\leq (1/\pi_1^*)(c\pi_1^*) \leq c,\end{aligned}$$

for some small constant c . The same result holds for $\mathbb{E}_{\mathcal{D}_1}[|\Delta_w|] = \mathbb{E}_{\mathcal{D}_1}[|(1 - w_1) - (1 - w_1^*)|]$.

Similarly, for the means,

$$\begin{aligned}|\mathbb{E}_{\mathcal{D}_1}[(1 - w_1)\langle v, s \rangle]| &\leq \sum_{l \neq 1} \mathbb{E}_{\mathcal{D}_1}[|w_l \langle v, s \rangle|] \leq (1/\pi_1^*) \sum_{l \neq 1} (3\sigma_l^* \pi_l^* + 5R_{j1}^* \pi_1^*) \exp(-R_{j1}^{*2}/64(\sigma_l^* \vee \sigma_1^*)^2) \\ &\leq (1/\pi_1^*)(c' \pi_1^* \sigma_1^*) \leq c' \sigma_1^*.\end{aligned}$$

The same result also holds for $|\mathbb{E}_{\mathcal{D}_1}[\Delta_w \langle v, s \rangle]| = |\mathbb{E}_{\mathcal{D}_1}[((1 - w_1) - (1 - w_1^*)) \langle v, s \rangle]|$.

Errors from all components: Now we can give a bound for the estimation errors after one population EM operation. For mixing weights,

$$|\pi_1^+ - \pi_1^*| \leq \pi_1^* |\mathbb{E}_{\mathcal{D}_j}[\Delta_w]| + \sum_{j \neq 1} \pi_j^* |\mathbb{E}_{\mathcal{D}_j}[\Delta_w]| \leq c_\pi \pi_1^*,$$

and

$$\pi_1^+ \|\mu_1^+ - \mu_1^*\| \leq \pi_1^* \|\mathbb{E}_{\mathcal{D}_1}[\Delta_w(X - \mu_1^*)]\| + \sum_{j \neq 1} \pi_j^* \|\mathbb{E}_{\mathcal{D}_j}[\Delta_w(X - \mu_1^*)]\| \leq c\pi_1^* \sigma_1^*.$$

Thus, $|\pi_1^+ - \pi_1^*| \leq c_\pi \pi_1^*$, $\|\mu_1^+ - \mu_1^*\| \leq c_\mu \sigma_1^*$ for some small constants $c_\pi, c_\mu \leq 0.5$. \blacksquare

B.6. Convergence of Variances

Proof We need some sharper bound on weights and probability of bad events, we need to go through another case study if we also have to estimate σ_1^2 . We need to show that σ_1^{+2} will be very close to σ_1^{*2} , i.e., $|\sigma_1^{+2} - \sigma_1^{*2}|/\sigma_1^{*2} \leq 0.5/\sqrt{d}$. First, let us arrange the EM operator for σ_1^+ .

$$\begin{aligned}\sigma_1^{+2} &= \mathbb{E}_{\mathcal{D}}[w_1 \|X - \mu_1^+\|^2] / (d\pi_1^+) \\ &= \mathbb{E}_{\mathcal{D}}[w_1 (\|X - \mu_1^* + \mu_1^* - \mu_1^+\|^2)] / (d\pi_1^+) \\ &= (\mathbb{E}_{\mathcal{D}}[w_1 (\|X - \mu_1^*\|^2 + \|\mu_1^* - \mu_1^+\|^2)] + 2\langle \mathbb{E}_{\mathcal{D}}[w_1(X - \mu_1^*)], \mu_1^* - \mu_1^+ \rangle) / (d\pi_1^+) \\ &= (\mathbb{E}_{\mathcal{D}}[w_1 \|X - \mu_1^*\|^2] - \pi_1^+ \|\mu_1^* - \mu_1^+\|^2) / (d\pi_1^+) \\ &= \mathbb{E}_{\mathcal{D}}[w_1 \|X - \mu_1^*\|^2] / (d\pi_1^+) - \|\mu_1^* - \mu_1^+\|^2 / d.\end{aligned}$$

We need to further change the expression to get a tight bound for the error. The difference from the ground truth is,

$$\sigma_1^{+2} - \sigma_1^{*2} = \mathbb{E}_{\mathcal{D}}[w_1 \|X - \mu_1^*\|^2] / (d\pi_1^+) - \sigma_1^{*2} - \|\mu_1^* - \mu_1^+\|^2 / d$$

$$\begin{aligned}
 &= \mathbb{E}_{\mathcal{D}}[w_1(\|X - \mu_1^*\|^2 - d\sigma_1^{*2})]/(d\pi_1^+) - \|\mu_1^* - \mu_1^+\|^2/d \\
 &= \mathbb{E}_{\mathcal{D}}[\Delta_w(\|X - \mu_1^*\|^2 - d\sigma_1^{*2})]/(d\pi_1^+) - \|\mu_1^* - \mu_1^+\|^2/d \\
 &= \frac{\sum_j \pi_j^* \mathbb{E}_{\mathcal{D}_j}[\Delta_w(\|\mu_j^* - \mu_1^*\|^2 + 2\langle v, \mu_j^* - \mu_1^* \rangle + \|v\|^2 - d\sigma_1^{*2})]}{d\pi_1^+} - \frac{\|\mu_1^* - \mu_1^+\|^2}{d}, \\
 |\sigma_1^{+2} - \sigma_1^{*2}| &\leq \frac{\sum_j \pi_j^* |\mathbb{E}_{\mathcal{D}_j}[\Delta_w]| R_{j1}^{*2}}{d\pi_1^+} + \frac{\sum_j 2\pi_j^* |\mathbb{E}_{\mathcal{D}_j}[\Delta_w \langle v, \mu_j^* - \mu_1^* \rangle]|}{d\pi_1^+} \\
 &\quad + \frac{\sum_j \pi_j^* |\mathbb{E}_{\mathcal{D}_j}[\Delta_w(\|v\|^2 - d\sigma_j^{*2})]|}{d\pi_1^+} + \frac{\sum_j \pi_j^* |\mathbb{E}_{\mathcal{D}_j}[\Delta_w] d(\sigma_j^{*2} - \sigma_1^{*2})|}{d\pi_1^+} + \frac{\|\mu_1^* - \mu_1^+\|^2}{d}.
 \end{aligned}$$

Let us consider the terms one by one. We can use Corollary 5 for the first and second terms. The third term can be bounded with Cauchy-Schwartz inequality.

$$\begin{aligned}
 |\mathbb{E}_{\mathcal{D}_j}[\Delta_w(\|v\|^2 - d\sigma_j^{*2})]| &\leq \sqrt{\mathbb{E}_{\mathcal{D}_j}[\Delta_w]} \sqrt{\mathbb{E}_{\mathcal{D}_j}[(\|v\|^2 - d\sigma_j^{*2})^2]} \\
 &\leq \sqrt{(3\pi_1^*/\pi_j^* + 5) \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2)} \sqrt{2d\sigma_j^{*4}},
 \end{aligned}$$

since the variance of chi-Square distribution with d degrees of freedom is $2d$.

The fourth term requires redefinition of good events to get an error that only scales with \sqrt{d} when d grows large. In order to bound this term, we need to give a sharper bound on $\mathbb{E}_{\mathcal{D}_j}[\Delta_w]$. That is, we need the following lemma:

Lemma 22 *For $j \neq 1$, if $|\sigma_j^{*2} - \sigma_1^{*2}| \geq 10R_{j1}^*(\sigma_j^* \vee \sigma_1^*)/\sqrt{d}$, then*

$$|\mathbb{E}_{\mathcal{D}_j}[\Delta_w]|, |\mathbb{E}_{\mathcal{D}_j}[w_1]| \leq O\left(\exp(-d \min(1, t^2)/256) \exp(-R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2)\right),$$

where $t = |\sigma_j^{*2} - \sigma_1^{*2}|/\sigma_1^{*2}$. This implies,

$$d|\sigma_j^{*2} - \sigma_1^{*2}| |\mathbb{E}_{\mathcal{D}_j}[\Delta_w]| \leq O\left(\sqrt{d}(1 + \pi_1^*/\pi_j^*) R_{j1}^*(\sigma_j^* \vee \sigma_1^*) \exp(-R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2)\right).$$

The intuition is as follows: what if $(\sigma_j^{*2} - \sigma_1^{*2}) < 10(R_{j1}^*\sigma_1^*)/\sqrt{d}$? Then we have nothing to worry about, since \sqrt{d} is already canceled out and $\mathbb{E}_{\mathcal{D}_j}[\Delta_w]$ will be small enough to cancel the rest which is $R_{j1}^*\sigma_1^*$. The problem is when $|\sigma_j^{*2} - \sigma_1^{*2}| > 10(R_{j1}^*\sigma_1^*)/\sqrt{d}$. Here, we expect that the error bound $\mathbb{E}_{\mathcal{D}_j}[\Delta_w]$ would crucially depend on the quantity $t = |\sigma_j^{*2} - \sigma_1^{*2}|/\sigma_1^{*2} > 10R_{j1}^*/(\sigma_1^*\sqrt{d})$. Recall that there are three main terms in the exponent of weights as in (17):

$$\begin{aligned}
 III &= -\langle v, \mu_j^* - \mu_1 \rangle / \sigma_1^2, \\
 IV &= \langle v, \Delta_j \rangle / \sigma_j^2, \\
 V &= -\frac{\|v\|^2}{2} \left(\frac{1}{\sigma_1^2} - \frac{1}{\sigma_j^2} \right) - \frac{d}{2} \log(\sigma_1^2 / \sigma_j^2).
 \end{aligned}$$

Proof

When $\sigma_1^* \geq \sigma_j^*$ and $|\sigma_j^{*2} - \sigma_1^{*2}| > 10R_{j1}^*\sigma_1^*/\sqrt{d}$: Let us refine good events so that each events can take this case into account. Note that this can only happen for very large $d > 64^2 \cdot 100$ given our separation condition. We first show that if $\sigma_1^{*2} - \sigma_j^{*2} \geq 10R_{j1}^*\sigma_1^*/\sqrt{d}$, then $x > 4t/5$. To see this,

$$\begin{aligned} \sigma_1^2 - \sigma_j^2 &\geq (1 - 1/(2\sqrt{d}))\sigma_1^{*2} - (1 + 1/(2\sqrt{d}))\sigma_j^{*2} \\ &\geq (\sigma_1^{*2} - \sigma_j^{*2}) - 1/(2\sqrt{d})(\sigma_1^{*2} + \sigma_j^{*2}) \geq 9R_{j1}^*\sigma_1^*/\sqrt{d} \end{aligned}$$

which then is connected to

$$x \geq (\sigma_1^{*2} - \sigma_j^{*2})/\sigma_1^2 - 1/(2\sqrt{d})(\sigma_1^{*2} + \sigma_j^{*2})/\sigma_1^2 \geq 4t/5,$$

given good initialization of σ_1 and σ_j .

Instead of just requiring $V \leq R_{j1}^{*2}/(8\sigma_1^2)$ as in mixing weight case, let us require $V \leq -dx^2/16 + a$ where $x = (\sigma_1^2 - \sigma_j^2)/\sigma_1^2$ and $a = R_{j1}^{*2}/(8\sigma_1^2)$. Then,

$$\frac{1}{d} \frac{\|v\|^2}{\sigma_j^{*2}} \leq \frac{\sigma_j^2}{\sigma_j^{*2}} \left(1 + \frac{x}{2} - \frac{x}{8} + \frac{2a}{dx} \right) = \frac{\sigma_j^2}{\sigma_j^{*2}} \left(1 + \frac{3x}{8} + \frac{2a}{dx} \right),$$

is a sufficient condition to guarantee $V \leq -dx^2/16 + a$. The probability of this event is

$$\begin{aligned} P \left(\frac{\|v\|^2}{d\sigma_j^{*2}} \geq \frac{\sigma_j^2}{\sigma_j^{*2}} \left(1 + \frac{3x}{8} + \frac{2a}{dx} \right) \right) &\leq P \left(\frac{\|v\|^2}{d\sigma_j^{*2}} \geq \left(1 - \frac{1}{2\sqrt{d}} \right) \left(1 + \frac{3x}{8} + \frac{2a}{dx} \right) \right) \\ &\leq \exp \left(-dx/32 - R_{j1}^*/(64\sigma_1^{*2}) \right). \end{aligned}$$

For quantities *III* and *IV*, we will require

$$\begin{aligned} P \left(\text{III} \geq dx^2/32 + 7R_{j1}^{*2}/(32\sigma_1^2) \right) &\leq P \left(\langle v, \mu_j^* - \mu_1 \rangle \geq \sqrt{dx}\sqrt{d}(\sigma_1^2 - \sigma_j^2)/32 + 7R_{j1}^{*2}/32 \right) \\ &\leq P \left(\langle v, \mu_j^* - \mu_1^* \rangle \geq \sqrt{dx}\sqrt{d}(\sigma_1^2 - \sigma_j^2)/40 + R_{j1}^{*2}/5 \right) + P \left(\langle v, \Delta_1 \rangle \geq \sqrt{dx}\sqrt{d}(\sigma_1^2 - \sigma_j^2)/160 + R_{j1}^{*2}/32 \right) \\ &\leq \exp \left(-dx^2/2 \underbrace{(\sqrt{d}(\sigma_1^2 - \sigma_j^2)/40\sigma_j^*R_{j1}^*)^2}_{(i)} - R_{j1}^{*2}/(64\sigma_j^{*2}) \right) \\ &\quad + \exp \left(-dx^2/2 \underbrace{(\sqrt{d}(\sigma_1^2 - \sigma_j^2)/10\sigma_j^*R_{j1}^*)^2}_{(i')} - R_{j1}^{*2}/(16\sigma_j^{*2}) \right). \end{aligned}$$

For *IV*,

$$\begin{aligned} P \left(\langle v, \Delta_j \rangle \geq dx^2\sigma_j^2/128 + R_{j1}^{*2}\sigma_j^2/(20\sigma_1^2) \right) \\ &\leq P \left(Z \geq dx^2(\sigma_1^* \vee \sigma_j^*)(\sigma_j^2/\sigma_j^{*2})/(8R_{j1}^*) + 4R_{j1}^*(\sigma_j^2/\sigma_j^{*2})/(5\sigma_1^2) \right) \\ &\leq \exp \left(-dx^2/2 \underbrace{(\sqrt{d}(\sigma_1^2 - \sigma_j^2)/16\sigma_1^*R_{j1}^*)^2}_{(i'')} - R_{j1}^{*2}/(64\sigma_1^{*2}) \right), \end{aligned}$$

When three events happens at the same time, we are guaranteed that $w_1 \leq \exp(-dx^2/64 - R_{j1}^{*2}/(64\sigma_1^{*2}))$. Meanwhile, by the assumption on this case study we have $(i), (i'), (i'') \geq 1/32$.

That is, we get $w_1 \leq \exp(-dt^2/128 - R_{j_1}^{*2}/(64\sigma_1^{*2}))$ with probability $1 - 4\exp(-dt^2/64 - R_{j_1}^{*2}/(64\sigma_1^{*2}))$. Under the same events, we can show the same for w_1^* . Therefore,

$$\mathbb{E}_{\mathcal{D}_j}[\Delta_w] \leq 10(\pi_1^*/\pi_j^* + 1) \exp(-dt^2/64 - R_{j_1}^{*2}/(64\sigma_1^{*2})).$$

This is enough to bound the fourth term (note that if $\sqrt{dt} > 8R_{j_1}^*/\sigma_1^* \gg 512$, then $dt^2 \geq 64 \log(\sqrt{dt})$ will be guaranteed). The consequence of this relation is that

$$|\mathbb{E}_{\mathcal{D}_j}[\Delta_w]d(\sigma_j^{*2} - \sigma_1^{*2})| \leq \sigma_1^{*2}\sqrt{d} \exp(-R_{j_1}^{*2}/(64\sigma_1^{*2})).$$

When $\sigma_1^* \leq \sigma_j^*$ and $|\sigma_j^{*2} - \sigma_1^{*2}| > 10R_{j_1}^*\sigma_j^*/\sqrt{d}$: We can go through the other side (when $\sigma_1 \leq \sigma_j$) similarly. For this case, let $x = (\sigma_j^2 - \sigma_1^2)/\sigma_1^2$. In this case, we consider the case when $t = (\sigma_j^{*2} - \sigma_1^{*2})/\sigma_1^{*2} > 10R_{j_1}^*\sigma_j^*/(\sigma_1^{*2}\sqrt{d})$, and we show that $x \geq 4t/5$ as before. Note that still the case can only happen when $d > 640^2$.

$$\begin{aligned} \sigma_j^2 - \sigma_1^2 &\geq (1 - 1/2\sqrt{d})\sigma_j^{*2} - (1 + 1/2\sqrt{d})\sigma_1^{*2} \\ &= (\sigma_j^{*2} - \sigma_1^{*2}) - 1/(2\sqrt{d})(\sigma_j^{*2} + \sigma_1^{*2}) \\ &\geq 9R_{j_1}^*\sigma_j^*/\sqrt{d}, \\ x &\geq 4t/5 \geq 8R_{j_1}^*\sigma_j^*/(\sigma_1^{*2}\sqrt{d}). \end{aligned}$$

We define the condition for V as

$$\begin{aligned} P(V \geq -dx^2/32 + R_{j_1}^{*2}/(8\sigma_1^2)) &= P(-\|v\|^2x/(2\sigma_j^2) + \frac{d}{2} \log(1+x) \geq -dx^2/32 + a) \\ &\leq P(-\|v\|^2x/\sigma_j^2 + d(x - x^2/4) \geq -dx^2/16 + 2a) \\ &= P(\|v\|^2/\sigma_j^2 \leq d(1 - 3x/16 - 2a/(dx))) \\ &\leq \exp(-dx/128 - R_{j_1}^{*2}/(64\sigma_1^{*2})), \quad \text{for } 0 \leq x \leq 3/4, \end{aligned}$$

We also similarly compute the bad probabilities for other quantities *III* and *IV*.

$$\begin{aligned} P(\text{III} \geq dx^2/64 + 7R_{j_1}^{*2}/(32\sigma_1^2)) &\leq P(\langle v, \mu_j^* - \mu_1^* \rangle \geq \sqrt{dx}\sqrt{d}(\sigma_j^2 - \sigma_1^2)/80 + R_{j_1}^{*2}/5) + P(\langle v, \Delta_1 \rangle \geq \sqrt{dx}\sqrt{d}(\sigma_j^2 - \sigma_1^2)/320 + R_{j_1}^{*2}/64) \\ &\leq \exp(-dx^2/2 \underbrace{(\sqrt{d}(\sigma_j^2 - \sigma_1^2)/(80R_{j_1}^*\sigma_j^*)^2 - R_{j_1}^{*2}/(32\sigma_j^{*2}))}_{(i)}) \\ &\quad + \exp(-dx^2/2 \underbrace{(\sqrt{d}(\sigma_j^2 - \sigma_1^2)/(20R_{j_1}^*\sigma_j^*)^2 - R_{j_1}^{*2}/(64\sigma_j^{*2}))}_{(i')})), \end{aligned}$$

$$\begin{aligned} P(\text{IV} \geq dx^2/128 + R_{j_1}^{*2}/(20\sigma_1^2)) &\leq P(\langle x, \Delta_j \rangle \sigma_1^2/\sigma_j^2 \geq dx^2\sigma_1^2/16 + 4R_{j_1}^{*2}/5) \\ &\leq P(Z \geq \sqrt{dx}(\sqrt{d}(\sigma_j^2 - \sigma_1^2)/(16R_{j_1}^*\sigma_1^*)) + R_{j_1}^{*2}/(4\sigma_1^{*2})) \\ &\leq \exp(-dx^2/2 \underbrace{(\sqrt{d}(\sigma_j^2 - \sigma_1^2)/(16R_{j_1}^*\sigma_1^*)^2 - R_{j_1}^{*2}/(64\sigma_1^{*2}))}_{(i'')}) \end{aligned}$$

Since we assumed $|\sigma_j^2 - \sigma_1^2| \geq 9R_{j_1}^* \sigma_j^* / \sqrt{d}$, we can see that $(i), (i'), (ii') \geq 1/100$. Thus, similarly we can get $w_1 \leq 3(\pi_1^*/\pi_j^*) \exp(-dt^2/256 - R_{j_1}^{*2}/(64\sigma_1^{*2}))$ with probability $1 - \exp(-dt^2/256 - R_{j_1}^{*2}/(64\sigma_1^{*2}))$. Using the same argument, $\mathbb{E}_{\mathcal{D}_j}[\Delta_w]$ can be bounded by $\exp(-dt^2/256 - R_{j_1}^{*2}/(64\sigma_1^{*2}))$.

When $3/4 \leq x$, we target:

$$\begin{aligned} P(V \geq -dx/16 + R_{j_1}^{*2}/(8\sigma_1^2)) &= P(-\|v\|^2 x / (2\sigma_j^2) + (d/2) \log(1+x) \geq -dx/8 + R_{j_1}^{*2}/(8\sigma_1^2)) \\ &\leq P(\|v\|^2 / \sigma_j^2 \leq d(\log(1+x)/x + 1/8) - R_{j_1}^{*2}/(4\sigma_1^2)/x) \\ &\leq P(\|v\|^2 / \sigma_j^2 \leq d(1 - 1/8 - R_{j_1}^{*2}/(4d(\sigma_j^2 - \sigma_1^2)))) \\ &\leq \exp(-d/256 - R_{j_1}^{*2}/(64\sigma_j^{*2})), \end{aligned}$$

Note that when $x = (\sigma_j^2 - \sigma_1^2)/\sigma_1^2 > 6(R_{j_1}^* \sigma_j^*) / (\sigma_1^{*2} \sqrt{d})$, it is true that $d \gg 512 \log \sqrt{d}$. For *III* and *IV*, when $x \geq 3/4$, we find a probability for

$$\begin{aligned} P(III \geq dx/32 + 7R_{j_1}^{*2}/(32\sigma_1^2)) &\leq P(\langle v, \mu_j^* - \mu_1^* \rangle \geq \sqrt{d}\sqrt{d}(\sigma_j^2 - \sigma_1^2)/64 + R_{j_1}^{*2}/5) \\ &\quad + P(\langle v, \Delta_1 \rangle \geq \sqrt{d}\sqrt{d}(\sigma_j^2 - \sigma_1^2)/64 + R_{j_1}^{*2}/64) \\ &\leq \exp(-d/2 \underbrace{(\sqrt{d}(\sigma_j^2 - \sigma_1^2)/(64R_{j_1}^* \sigma_j^*)^2 - R_{j_1}^{*2}/(64\sigma_j^{*2}))}_{(i)}) \\ &\quad + \exp(-d/2 \underbrace{(\sqrt{d}(\sigma_j^2 - \sigma_1^2)/(4R_{j_1}^* \sigma_j^*)^2 - R_{j_1}^{*2}/(64\sigma_j^{*2}))}_{(i)}), \end{aligned}$$

and,

$$\begin{aligned} P(IV \geq dx/64 + R_{j_1}^{*2}/(20\sigma_1^2)) &\leq P(ZR_{j_1}^* \sigma_j^* \geq dx\sigma_1^2/4 + 4R_{j_1}^{*2}/5) \\ &\leq \exp(-d/2 \underbrace{(\sqrt{d}(\sigma_j^2 - \sigma_1^2)/(4R_{j_1}^* \sigma_j^*)^2 - R_{j_1}^{*2}/(64\sigma_1^{*2}))}_{(i')}), \end{aligned}$$

Again, the similar result holds for $\mathbb{E}_{\mathcal{D}_j}[\Delta_w]$. Therefore,

$$\mathbb{E}_{\mathcal{D}_j}[\Delta_w] \leq 3(1 + \pi_1^*/\pi_j^*) \exp(-d/256 - R_{j_1}^{*2}/(64\sigma_j^{*2})),$$

Collecting all cases yields the Lemma. ■

Errors from own component $j = 1$: Note that if $j = 1$, first, second, and fourth terms are gone automatically. For the third term, $\mathbb{E}_{\mathcal{D}_1}[\Delta_w] \leq c$ for some small c as we have seen several times, and $\mathbb{E}_{\mathcal{D}_1}[(\|v\|^2 - d\sigma_1^{*2})^2] \leq 2d\sigma_1^{*4}$. The fifth term is less than $O(\sigma_1^{*2}/d)$ as we have already seen that next estimates for means are already within $c_\mu \sigma_1^*$. Hence the error from own component is less than c_1 for some small constant c_1 .

Errors from all components: Now we can collect every terms to give a bound the error of σ_1^+ ,

$$\sigma_1^{+2} - \sigma_1^{*2} \leq c_1 \sigma_1^{*2} + \frac{\sum_{j \neq 1} \pi_j^* R_{j_1}^{*2} \mathbb{E}_{\mathcal{D}_j}[\Delta_w]}{d\pi_1^+} + \frac{2 \sum_{j \neq 1} \pi_j^* R_{j_1}^* \mathbb{E}_{\mathcal{D}_j}[\Delta_w]}{d\pi_1^+}$$

$$\begin{aligned}
& + \frac{\sum_{j \neq 1} 5\sigma_j^{*2}(\pi_1^* + \pi_j^*)\sqrt{d} \exp(-R_{j1}^{*2}/128(\sigma_1^* \vee \sigma_j^*)^2)}{d\pi_1^+} \\
& + \frac{\sum_{j \neq 1} 50(\pi_1^* + \pi_j^*)\sqrt{d}R_{j1}^*(\sigma_j^* \vee \sigma_1^*) \exp(-R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2)}{d\pi_1^+} + \frac{c_\mu^2 \sigma_1^{*2}}{d},
\end{aligned}$$

which gives $|\sigma_1^{+2} - \sigma_1^{*2}| \leq c_\sigma \sigma_1^{*2}/\sqrt{d}$, for some small constant $c_\sigma < 0.5$ given good enough SNR condition in (2). \blacksquare

Appendix C. Proof for Finite-Sample EM

We define some additional notations. We use \mathcal{E}_j to denote the event that the i^{th} sample comes from j^{th} component. Define $\mathcal{E}_{j,\text{good}} := \mathcal{E}_{j,1} \cap \mathcal{E}_{j,2} \cap \mathcal{E}_{j,3}$ where $\mathcal{E}_{j,\cdot}$ are as defined in (6). For the simplicity in notation, we now use looser upper bound ρ_π for $(1 \vee \pi_1^*/\pi_j^*)$. We use \lesssim when the inequality holds up to some universal constants. Finally, under the modified condition in (10), we will use slightly modified version of Lemma 6:

Lemma 23 *For well-separated mixture of Gaussians, for $q \in \{0, 1, 2\}$,*

$$\rho_\pi \sum_{j \neq 1} R_{j1}^{*q} \exp\left(-R_{j1}^{*2}/(128c)(\sigma_1^* \vee \sigma_j^*)^2\right) \leq c_q \sigma_1^{*q} \pi_{\min}, \quad (21)$$

for some small constants c_q given separation condition as in (10).

Proof The proof is similar to that of Lemma 6. Note that $R_{j1}^{*2}/(\sigma_1^* \vee \sigma_j^*)^2 \geq C^2 c^2 \log(\rho_\sigma/\pi_{\min})$ where the universal constant C is such that $C^2 \geq 4096$. Let $x := R_{j1}^{*2}/(\sigma_1^* \vee \sigma_j^*)^2$. Then, since $\log(x)/x$ is decreasing in x whenever $x \geq e$,

$$\frac{\log(x)}{x} \leq \frac{\log(C^2 c^2 \log(\rho_\sigma/\pi_{\min}))}{C^2 c^2 \log(\rho_\sigma/\pi_{\min})} \leq \frac{1}{256c}.$$

Applying this to the equation (21) with $q = 2$,

$$\begin{aligned}
\rho_\pi \sum_{j \neq 1} R_{j1}^{*2} \exp\left(-R_{j1}^{*2}/(128c)(\sigma_1^* \vee \sigma_j^*)^2\right) & \leq \rho_\pi \sum_{j \neq 1} (\sigma_1^* \vee \sigma_j^*)^2 \exp\left(-R_{j1}^{*2}/(256c)(\sigma_1^* \vee \sigma_j^*)^2\right) \\
& \leq \rho_\pi \sum_{j \neq 1} (\sigma_1^* \vee \sigma_j^*)^2 (\rho_\sigma/\pi_{\min})^{-32} \ll c_2 \sigma_1^{*2} \pi_{\min},
\end{aligned}$$

which gives the lemma with some small constant c_2 . Similar claims hold for $q = 0, 1$. \blacksquare

We also restate here the Proposition 8.

Proposition 24 (Restatement of Proposition 8) *Let X be a random d -dimensional vector, and A be an event in the same probability space with $p = P(A) > 0$. Define random variable $Y = X|A$, i.e., X conditioned on event A , and $Z = \mathbf{1}_{X \in A}$. Let X_i, Y_i, Z_i be the i.i.d samples from corresponding distributions. Then, the following holds,*

$$P\left(\left\|\frac{1}{n} \sum_{i=1}^n X_i \mathbf{1}_{X_i \in A} - \mathbb{E}[X \mathbf{1}_{X \in A}]\right\| \geq t\right) \leq \max_{m \leq n_e} P\left(\frac{1}{n} \left\|\sum_{i=1}^m (Y_i - \mathbb{E}[Y])\right\| \geq t_1\right)$$

$$+ P \left(\left\| \mathbb{E}[Y] \right\| \left| \frac{1}{n} \sum_{i=1}^n Z_i - p \right| \geq t_2 \right) + P \left(\left| \sum_{i=1}^n Z_i \right| \geq n_e + 1 \right). \quad (22)$$

for any $0 \leq n_e \leq n$ and $t_1 + t_2 = t$.

C.1. Concentration in Mixing Weights

Proof We give a concentration result for mixing weights first. We can first check that

$$\tilde{\pi}_1^+ - \pi_1^+ = \frac{1}{n} \sum_{i=1}^n w_{1,i} - \mathbb{E}_{\mathcal{D}}[w_1].$$

Then, single $w_{1,i}$ can be decomposed using indicators. Then,

$$w_{1,i} = w_{1,i} \mathbb{1}_{\mathcal{E}_1} + \sum_{j \neq 1} w_{1,i} \mathbb{1}_{\mathcal{E}_j \cap \mathcal{E}_{j,good}} + w_{1,i} \mathbb{1}_{\mathcal{E}_j \cap \mathcal{E}_{j,good}^c}.$$

Now let us apply proposition 8 step by step.

With \mathcal{E}_1 : Note that $P(\mathcal{E}_1) = \pi_1^*$. We can pick $n_e = 2n\pi_1^*$. By multiplicative version of concentration inequality for Bernoulli random variable, the second and third terms will be safely killed. Also, we note that $w_{1,i}$ is bounded random variable. Therefore,

$$P(1/n \left| \sum_{i=1}^{n_e} (w_{1,i} - \mathbb{E}_{\mathcal{D}_1}[w_{1,i}]) \right| \geq t) \leq \exp(-2n_e(nt/n_e)^2) = \exp(-2n^2/n_e t^2).$$

Thus, $t = O\left(\sqrt{n_e/n} \sqrt{\ln(k^2 T/\delta)/n}\right) = O\left(\sqrt{\pi_1^* \ln(k^2 T/\delta)/n}\right)$ gives $\delta/(k^2 T)$ error bound.

With $\mathcal{E}_j \cap \mathcal{E}_{j,good}$: When a good sample comes from $j^{th} \neq 1$ component, the weight given to first component is very small, *i.e.*, $w_{1,i} \mathbb{1}_{\mathcal{E}_j \cap \mathcal{E}_{j,good}} \leq 5\rho_\pi \exp(-R_{j1}^{*2}/(64(\sigma_j^* \vee \sigma_1^*)^2))$. Thus, it is a bounded random variable, therefore its statistical error can be bounded by

$$t = O\left(\rho_\pi \sqrt{\pi_j^*} \sqrt{\ln(k^2 T/\delta)/n}\right) \exp\left(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2\right),$$

with probability at least $1 - (k^2 T/\delta)$.

With $\mathcal{E}_{j,good}^c$: This is a very special case, since the chance of this event to happen is $p := 5\pi_j^* \exp(-R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2)$, *i.e.*, exponentially small. We first need to bound the number of samples that have fallen into this bad event with high probability. We divide the case as when $n \geq p^{1/c}$ and $n \leq p^{1/c}$ for some constant $c > 2$.

Let us first consider when $n \geq p^{1/c}$. Recall the Bernstein's inequality, which states for Bernoulli random variable that

$$P\left(\left|\frac{1}{n} \sum_{i=1}^n Z_i - p\right| \geq t\right) \leq \exp(-nt^2/(2p + 2/3t)).$$

Solving the right hand side to get a $\delta/(k^2 T)$ probability bound, we get

$$n_e = O(c \ln(k^2 T/\delta) + c\sqrt{np \ln(k^2 T/\delta)}).$$

Then using the proposition, we can decide how large the sum of $1/n \sum_{i=1}^n w_{1,i} \mathbb{1}_{\mathcal{E}_{j,\text{good}}^c}$, which will be

$$t = O\left(\sqrt{p} \vee 1/n \sqrt{\ln^2(k^2 T/\delta)/n}\right) = O\left(p^{1/2c} \sqrt{\ln^2(k^2 T/\delta)/n}\right) = \tilde{O}\left(\pi_{\min}^{32} \sqrt{1/n}\right),$$

with probability $1 - n^c$.

On the other side, if $n \leq p^{1/c}$, then we will have $n_e = 0$ with probability at least $1 - np \geq 1 - p^{1-1/c}$. Note that $p^{1-1/c} \leq O(\pi_{\min}^{32}/n^{c-2})$ given SNR condition as in the theorem. Thus, in this case, with probability at least $1 - 1/(n^{c-2}k^{32})$, we have

$$\left| \frac{1}{n} \sum_{i=1}^n w_{1,i} \mathbb{1}_{\mathcal{E}_{j,\text{good}}^c} - \mathbb{E}_{\mathcal{D}}[w_{1,i} \mathbb{1}_{\mathcal{E}_{j,\text{good}}^c}] \right| \leq \mathbb{E}_{\mathcal{D}_j}[w_{1,i} \mathbb{1}_{\mathcal{E}_{j,\text{good}}^c}] \leq \exp(-R_{j1}^{*2}/(64(2c)(\sigma_1^* \vee \sigma_j^*)^2)).$$

Collect all errors: Now we can collect all items we found for each cases. Taking union bound over all $O(k)$ items, with probability $1 - O(\delta/kT) - O(1/(n^{c-2}k^{31}))$,

$$\begin{aligned} \left| \frac{1}{n} \sum_{i=1}^n w_{1,i} - \mathbb{E}_{\mathcal{D}}[w_{1,i}] \right| &\leq O\left(\sqrt{\pi_1^* \ln(k^2 T/\delta)/n}\right) + O\left(\rho_\pi \sqrt{\ln^2(k^2 T/\delta)/n}\right) \sum_{j \neq 1} \exp(-R_{j1}^{*2}/(128c(\sigma_j^* \vee \sigma_1^*)^2)) \\ &\leq O(\pi_1^* \epsilon), \end{aligned}$$

where we used Lemma 23. Thus, $|\tilde{\pi}_1^+ - \pi_1^*|/\pi_1^* \leq |\tilde{\pi}_1^+ - \pi_1^+|/\pi_1^* + |\pi_1^+ - \pi_1^*|/\pi_1^* \leq \epsilon + \gamma D_m$. Thus, after $T = O(\log(1/\epsilon))$ iteration, we get $|\tilde{\pi}_1^{(T)} - \pi_1^*| \leq \pi_1^* \epsilon$ with probability $1 - O(\delta/k) - O(T/(n^{c-2}k^{31}))$. We can take union bound over all $O(k)$ components to get the result for all components with probability $1 - O(\delta) - O(\log(1/\epsilon)/n^{c-2}k^{30})$. \blacksquare

C.2. Concentration in Means

Proof Now let us look at the iteration for means. First, we should observe that

$$\begin{aligned} \tilde{\mu}_1^+ - \mu_1^* &= \left(\frac{1}{n} \sum_{i=1}^n w_{1,i} (X - \mu_1^*) \right) / \left(\frac{1}{n} \sum_{i=1}^n w_{1,i} \right) \\ &= \left(\frac{1}{n} \sum_{i=1}^n w_{1,i} (X_i - \mu_1^*) - \mathbb{E}_{\mathcal{D}}[w_1 (X - \mu_1^*)] + \mathbb{E}_{\mathcal{D}}[w_1 (X - \mu_1^*)] - \mathbb{E}_{\mathcal{D}}[w_1^* (X - \mu_1^*)] \right) / \tilde{\pi}_1^+ \\ &= \left(\underbrace{\frac{1}{n} \sum_{i=1}^n w_{1,i} (X_i - \mu_1^*) - \mathbb{E}_{\mathcal{D}}[w_1 (X - \mu_1^*)]}_{e_\mu} + \underbrace{\mathbb{E}_{\mathcal{D}}[\Delta_w (X - \mu_1^*)]}_{B_\mu} \right) / \tilde{\pi}_1^+. \end{aligned}$$

B_μ is decreasing as we have seen for population EM, we focus on the fluctuation of the sum of random variables $W = w_1(X - \mu_1^*)$. We further decompose this random variable using disjoint events. That is,

$$W_i = W_i \mathbb{1}_{\mathcal{E}_1} + \sum_{j \neq 1}^n \left(W_i \mathbb{1}_{\mathcal{E}_{j,\text{good}}} + W_i \mathbb{1}_{\mathcal{E}_{j,1}^c} + W_i \mathbb{1}_{\mathcal{E}_{j,1} \cap \mathcal{E}_{j,2}^c} + W_i \mathbb{1}_{\mathcal{E}_{j,1} \cap \mathcal{E}_{j,2} \cap \mathcal{E}_{j,3}^c} \right).$$

Now for each decomposed sample, we compute ψ_2 or ψ_1 norm conditioned on each event, and sum everything at the end.

With $j \neq 1, \mathcal{E}_{j,good}$: Let $Y_i = W_i \mathbb{1}_{\mathcal{E}_{j,good}} | \mathcal{E}_j$ and $Z_i = \mathbb{1}_j$. Then using the proposition,

$$\begin{aligned} P\left(\left\|\frac{1}{n} \sum_i^n W_i \mathbb{1}_{j,good} - \mathbb{E}_{\mathcal{D}}[W \mathbb{1}_{j,good}]\right\| \geq t\right) &\leq P\left(\frac{1}{n} \left\|\sum_i^{n_e} Y_i - \mathbb{E}_{\mathcal{D}_j}[Y]\right\| \geq t_1\right) \\ &+ P\left(\mathbb{E}_{\mathcal{D}_j}[Y] \left|\frac{1}{n} \sum_i^n Z_i - \pi_j^*\right| \geq t_2\right) + P\left(\sum_{i=1}^n Z_i \geq n_e + 1\right). \end{aligned}$$

Now we find a sub-Gaussian norm of Y_i , which can be computed as

$$\begin{aligned} \|Y\|_{\psi_2} &= \sup_{p \geq 1} p^{-1/2} \mathbb{E}_{\mathcal{D}_j} [|w_1 \mathbb{1}_{j,good} \langle v + \mu_j^* - \mu_1^*, s \rangle|^p]^{1/p} \\ &\leq 10\rho_\pi \exp(-R_{j1}^{*2}/64\sigma_1^{*2}) \sup_{p \geq 1} p^{-1/2} (\mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle|^p]^{1/p} + R_{j1}^*) \\ &\leq 20\rho_\pi (R_{j1}^* + \sigma_1^*) \exp(-R_{j1}^{*2}/64\sigma_1^{*2}). \end{aligned}$$

Meanwhile, we can set $n_e = n\pi_j^* + O(\sqrt{np \ln(k^2 T/\delta)}) \leq 2n\pi_j^*$ as previously to get a high probability bound for the number of samples from j^{th} component. Using standard 1/2-covering argument for d -dimensional sub-Gaussian vector, we have

$$P\left(\left\|\sum_{i=1}^{n_e} Y_i - \mathbb{E}_{\mathcal{D}_j}[Y]\right\| \geq nt_1\right) \leq \exp\left(-\frac{(nt_1)^2}{n_e \|Y\|_{\psi_2}^2} + Cd\right),$$

for some universal constant C . That is, $t_1 = O\left(\|Y\|_{\psi_2} \sqrt{\frac{n_e}{n}} \sqrt{\frac{d + \log(k^2 T/\delta)}{n}}\right)$ with probability at least $1 - \delta/(k^2 T)$.

With $j \neq 1, \mathcal{E}_{j,1}^c$: We can use the same trick with $Y_i = W_i | \mathcal{E}_j \cap \mathcal{E}_{j,1}^c$. Note that $p := P(\mathcal{E}_j \cap \mathcal{E}_{j,1}^c) \leq \pi_j^* \exp(-R_{j1}^{*2}/64\sigma_1^{*2})$. Sub-Gaussian norm of Y_i in this case can be bounded with using one of the lemmas.

$$\begin{aligned} \|Y\|_{\psi_2} &= \sup_{p \geq 1} p^{-1/2} \mathbb{E}_{\mathcal{D}_j} [|w_1 \langle v + \mu_j^* - \mu_1^*, s \rangle|^p | \langle v, u \rangle \geq R_{j1}^*/5]^{1/p} \\ &\leq \sup_{p \geq 1} p^{-1/2} (R_{j1}^* + \mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle|^p | \langle v, u \rangle \geq R_{j1}^*/5]^{1/p}) \\ &\leq R_{j1}^* + \sup_{p \geq 1} p^{-1/2} \sigma_j^* (2R_{j1}^*/5\sigma_j^* + (2p)^{1/2}) \\ &\leq 2R_{j1}^*, \end{aligned}$$

where u is a unit vector in direction $\mu_j^* - \mu_1^*$. Since the probability of the bad event is very small, we divide the cases into $n \geq p^{-1/c}$ and $n \leq p^{-1/c}$ for some $c > 2$, as we have done for handling bad events for mixing weights. From Bernstein's inequality, if $n \geq p^{-1/c}$, with probability at least $1 - \delta/k^2 T$,

$$n_e \leq O\left(\ln(k^2 T/\delta) + \sqrt{np \ln(k^2 T/\delta)}\right).$$

Using this, we can give a good bound for t_1 with high probability,

$$\begin{aligned} t_1 &= O\left(\|Y\|_{\psi_2} \sqrt{n_e/n} \sqrt{(d + \ln(k^2 T/\delta))/n}\right) \\ &\leq O\left(R_{j_1}^* \sqrt{1/n} \vee \sqrt{p/n} \sqrt{\ln(k^2 T/\delta)} \sqrt{(d + \ln(k^2 T/\delta))/n}\right) \\ &\leq \tilde{O}\left(R_{j_1}^* p^{1/2c} \sqrt{d/n}\right), \end{aligned}$$

and

$$t_2 = O(R_{j_1}^* \ln(k^2 T/\delta)/n) = O\left(R_{j_1}^* p^{1/2c} \sqrt{\ln^2(k^2 T/\delta)/n}\right),$$

getting a similar scale of fluctuation.

For the other case when $n \leq p^{-1/c}$, we can again get $n_e = 0$ with probability at least $1 - p^{1-1/c}$, which is again greater than $1 - (1/(n^{c-2} k^{32}))$. In this case, $t_1 = 0$ and $t_2 = \mathbb{E}[Y]p \leq 2R_{j_1}^* p^{1-1/c}/n \lesssim \sigma_1^* \pi_{\min}/n$, which is again sufficiently small.

In all cases, we have that the fluctuation conditioned on this bad event is $\tilde{O}(\sigma_1^* \pi_{\min} \sqrt{d/n})$ with probability at least $1 - \delta/(k^2 T) - 1/(n^{c-2} k^{32})$.

With $j \neq 1$, $\mathcal{E}_{j,1} \cap \mathcal{E}_{j,2}^c$: Let $Y_i = W_i \mathbb{1}_{\mathcal{E}_{j,1}} |\mathcal{E}_{j,2}^c$. Then $p := P(\mathcal{E}_{j,2}^c) \leq 2\pi_j^* \exp(-R_{j_1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2)$. We can again follow the same path as we have done for other bad events. The key step is to get a sub-Gaussian norm.

$$\begin{aligned} \|Y\|_{\psi_2} &= \sup_{p \geq 1} p^{-1/2} \mathbb{E}_{\mathcal{D}_j} [|w_1 \mathbb{1}_{\mathcal{E}_{j,1}} \langle v + \mu_j^* - \mu_1^*, s \rangle|^p | \mathcal{E}_{j,2}^c]^{1/p} \\ &\leq R_{j_1}^* + \sup_{p \geq 1} p^{-1/2} \mathbb{E}_{\mathcal{D}_j} [|w_1 \langle v, s \rangle|^p | \langle v, \Delta_1 \rangle \geq R_{j_1}^{*2}/64]^{1/p} \\ &\quad + \sup_{p \geq 1} p^{-1/2} \mathbb{E}_{\mathcal{D}_j} [|w_1 \langle v, s \rangle|^p | \langle v, \Delta_j \rangle \geq (\sigma_j^*/\sigma_1^*)^2 R_{j_1}^{*2}/64]^{1/p} \\ &\leq R_{j_1}^* + 2\sigma_j^* \sup_{p \geq 1} p^{-1/2} \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [|\langle v, s \rangle|^p | \langle v, u \rangle \geq R_{j_1}^*/4(\sigma_j^* \vee \sigma_1^*)]^{1/p} \\ &\leq R_{j_1}^* + 2 \sup_{p \geq 1} p^{-1/2} (\sigma_j^* (2R_{j_1}^*/4(\sigma_j^* \vee \sigma_1^*) + \sqrt{2p})) \\ &\leq 3R_{j_1}^*. \end{aligned}$$

The rest of the step is similar to the previous case. We can thus again get a results that the deviation in this case is also $\tilde{O}\left(R_{j_1}^* \exp(-R_{j_1}^{*2}/(64 \cdot (2c)(\sigma_1^* \vee \sigma_j^*)^2)) \sqrt{d/n}\right)$ with probability at least $1 - \delta/(k^2 T) - 1/(n^{c-2} k^{32})$.

With $j \neq 1$, $\mathcal{E}_{j,1} \cap \mathcal{E}_{j,2} \cap \mathcal{E}_{j,3}^c$: Let $Y_i = W_i \mathbb{1}_{\mathcal{E}_{j,1} \cap \mathcal{E}_{j,2}} |\mathcal{E}_{j,3}^c$. Again, $p := P(\mathcal{E}_{j,3}^c) \leq 2\pi_j^* \exp(-R_{j_1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2)$ again. This time, we can invoke lemma 3.9 with $\alpha = R_{j_1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2$, to get

$$\begin{aligned} \|Y\|_{\psi_2} &= \sup_{p \geq 1} p^{-1/2} \mathbb{E}_{\mathcal{D}_j} [|w_1 \mathbb{1}_{\mathcal{E}_{j,1}} \langle v + \mu_j^* - \mu_1^*, s \rangle|^p | \mathcal{E}_{j,3}^c]^{1/p} \\ &\leq R_{j_1}^* + \sup_{p \geq 1} p^{-1/2} \mathbb{E}_{\mathcal{D}_j} [|w_1 \langle v, s \rangle|^p | \|v\|^2/\sigma_j^{*2} \geq d + 2\sqrt{d\alpha} + 2\alpha]^{1/p} \end{aligned}$$

$$\begin{aligned}
 & + \sup_{p \geq 1} p^{-1/2} \mathbb{E}_{\mathcal{D}_j} [|w_1 \langle v, s \rangle|^p \|v\|^2 / \sigma_j^{*2}] \leq d - 2\sqrt{d\alpha}]^{1/p} \\
 & \leq R_{j1}^* + c_1 \sigma_j^* + \sigma_j^* \sup_{p \geq 1} p^{-1/2} \mathbb{E}_{\mathcal{D}_j} [| \langle v, s \rangle|^p \|v\|^2] \geq d + 2\sqrt{d\alpha} + 2\alpha]^{1/p} \\
 & \leq R_{j1}^* + c_1 \sigma_j^* + \sigma_j^* \sup_{p \geq 1} p^{-1/2} ((64\alpha)^{1/2} + 4^{1/p} (8\alpha + p)^{1/2}) \\
 & \leq 3R_{j1}^*.
 \end{aligned}$$

The rest of the step is similar to the previous case. We can thus again get a results that the deviation in this case is also $\tilde{O} \left(R_{j1}^* \exp(-R_{j1}^{*2} / (64 \cdot 2c(\sigma_1^* \vee \sigma_j^*)^2)) \sqrt{d/n} \right)$ with probability at least $1 - \delta / (k^2 T) - 1 / (n^{c-2} k^{32})$.

With \mathcal{E}_1 : Let $Y_i = W_i |_{\mathcal{E}_1}$, $p := P(\mathcal{E}_1) = \pi_1^*$. $\|Y\|_{\psi_2}$ can be easily verified such that

$$\|Y\|_{\psi_2} = \sup_{p \geq 1} p^{-1/2} \mathbb{E}_{\mathcal{D}_1} [|w_1 \langle v, s \rangle|^p]^{1/p} \leq K \sigma_1^*,$$

for some constant K . We can set $n_e = 2n\pi_1^*$ as usual, to get the statistical error by those samples are

$$\tilde{O} \left(\sigma_1^* \sqrt{\pi_1^*} \sqrt{d/n} \right),$$

with probability at least $1 - \delta / (k^2 T)$.

Summing up every terms, the entire error is

$$\begin{aligned}
 e_\mu & \lesssim \sigma_1^* \left(\sqrt{\pi_1^*} \sqrt{d/n} + \sqrt{d/n} \rho_\pi \sum_{j \neq 1} R_{j1}^* \exp \left(-R_{j1}^{*2} / (128c(\sigma_j^* \vee \sigma_1^*)^2) \right) \right) \\
 & \lesssim \sigma_1^* \pi_{\min} \epsilon,
 \end{aligned}$$

with probability $1 - O(\delta/kT + 1/n^{c-2} \cdot 1/k^{31})$. In consequence,

$$\|\tilde{\mu}_1^+ - \mu_1^*\| \leq \pi_{\min} \sigma_1^* (\epsilon + \gamma \|\tilde{\mu}_1 - \mu_1^*\| / \sigma_1^*) / \tilde{\pi}_1^+ \leq \sigma_1^* (\epsilon + \gamma D_m).$$

Similarly to mixing weights, after taking union bound over all k components and $T = O(\log(1/\epsilon))$ iterations, we get $\|\tilde{\mu}_1^{(T)} - \mu_1^*\| \leq \sigma_1^* \epsilon$ with probability $1 - O(\delta + T/n^{c-2} \cdot 1/k^{30})$. \blacksquare

C.3. Concentration in Variances

Proof

With finite samples, the finite-sample EM iteration for variance is

$$\begin{aligned}
 \tilde{\sigma}_1^{+2} - \sigma_1^{*2} & = \left(\sum_{i=1}^n w_{1,i} \|X_i - \tilde{\mu}_1^+\|^2 \right) / \left(d \sum_{i=1}^n w_{1,i} \right) - \sigma_1^{*2} \\
 & = \left(1/n \sum_{i=1}^n w_{1,i} \|X_i - \mu_1^*\|^2 \right) / (d\tilde{\pi}_1^+) - \sigma_1^{*2} - \|\tilde{\mu}_1^+ - \mu_1^*\|^2 / d
 \end{aligned}$$

$$= \underbrace{\left(\frac{1}{n} \sum_{i=1}^n w_{1,i} (\|X_i - \mu_1^*\|^2 - d\sigma_1^{*2}) - \mathbb{E}[w_1 (\|X - \mu_1^*\|^2 - d\sigma_1^{*2})] \right)}_{\epsilon_\sigma} / (d\tilde{\pi}_1^+) + \frac{\mathbb{E}[\Delta_w (\|X - \mu_1^*\|^2 - d\sigma_1^{*2})]}{d\tilde{\pi}_1^+} - \frac{\|\tilde{\mu}_1^+ - \mu_1^*\|^2}{d}.$$

In order to be more precise, we may need to target $\tilde{O}(\sqrt{1/nd})$ for statistical precision. But it is enough to proceed more roughly, since $\tilde{O}(\sqrt{1/n})$ is enough to guarantee ϵ/\sqrt{d} statistical error with $n = \Omega(d)$ samples (we need this for estimating means). Let us define $W_i = w_{1,i} (\|X_i - \mu_1^*\|^2 - d\sigma_1^{*2})$ and use the decomposition strategy as we have done for μ .

With $j \neq 1, \mathcal{E}_{j,\text{good}}$: Let $Y_i = W_i \mathbb{1}_{j,\text{good}} | \mathcal{E}_j$, $p := \pi_j^*$. First task is, similarly, to find a sub-exponential norm (since now Y are sum of squared variables). We first compute it,

$$\begin{aligned} \|Y\|_{\psi_1} &= \sup_{p \geq 1} p^{-1} \mathbb{E}_{\mathcal{D}_j} [|w_1 \mathbb{1}_{j,\text{good}} (\|v + \mu_j^* - \mu_1^*\|^2 - d\sigma_1^{*2})|^p]^{1/p} \\ &\leq |w_1 \mathbb{1}_{j,\text{good}}| \sup_{p \geq 1} p^{-1} \left(R_{j1}^{*2} + \mathbb{E}_{\mathcal{D}_j} [\|v\|^{2p}]^{1/p} + 2\mathbb{E}_{\mathcal{D}_j} [|\langle v, \mu_j^* - \mu_1^* \rangle|^p]^{1/p} + \mathbb{E}_{\mathcal{D}_j} [d\sigma_1^{*2}]^{1/p} \right) \\ &\leq 10\rho_\pi \exp\left(-\frac{R_{j1}^{*2}}{64(\sigma_1^* \vee \sigma_j^*)^2}\right) \sup_{p \geq 1} p^{-1} \left((R_{j1}^{*2} + d\sigma_1^{*2}) + KR_{j1}^* \sigma_j^* \sqrt{p} + K'd\sigma_j^{*2} p \right) \\ &\leq C\rho_\pi \exp\left(-R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2\right) \left(R_{j1}^{*2} + d\sigma_1^{*2} + d\sigma_j^{*2} + R_{j1}^* \sigma_j^* \right) \\ &\leq C\rho_\pi \left(R_{j1}^{*2} + d(\sigma_j^* \vee \sigma_1^*)^2 \right) \exp\left(-R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2\right). \end{aligned}$$

Here we bound this term with the tail bound for sub-exponential random variable with ψ_1 -norm $\|Y\|_{\psi_1}$. Note that for sub-exponential random variable, from [Vershynin \(2010\)](#),

$$\left| \sum_{i=1}^{n_e} (Y - \mathbb{E}[Y]) \right| \leq \|Y\|_{\psi_1} O\left(\sqrt{n_e \log(1/\delta')} + \log(1/\delta')\right),$$

with probability $1 - \delta'$. In order to decide the statistical fluctuation, we just need to pick the maximum among $\sqrt{p/n}$ and $1/n$, which is in effect same to the case when Y is sub-Gaussian. For this event, we can set $n_e = 2n\pi_j^*$ as before to bound the number of samples. Thus, we get bound the statistical error as

$$\tilde{O}\left(\rho_\pi \sqrt{\pi_j^*} \sqrt{1/n}\right) \left(R_{j1}^{*2} + d(\sigma_j^* \vee \sigma_1^*)^2 \right) \exp\left(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2\right),$$

with probability at least $1 - \delta/(k^2T)$. Note that d term will be canceled out with division by d at the end.

With $j \neq 1, \mathcal{E}_{j,1}^c$: Similarly, we find sub-exponential norm of $Y = W | \mathcal{E}_j \cap \mathcal{E}_{j,1}^c$.

$$\|Y\|_{\psi_1} = \sup_{p \geq 1} p^{-1} \mathbb{E}_{\mathcal{D}_j} [|w_1 (\|v + \mu_j^* - \mu_1^*\|^2 - d\sigma_1^{*2})|^p | \langle v, \mu_j^* - \mu_1^* \rangle \geq R_{j1}^*/5]^{1/p}$$

$$\begin{aligned}
 &\leq \sup_{p \geq 1} p^{-1} \left((R_{j_1}^{*2} + d\sigma_1^{*2}) + \mathbb{E}_{\mathcal{D}_j} [\|v\|^{2p} |\langle v, \mu_j^* - \mu_1^* \rangle \geq R_{j_1}^*/5]^{1/p} \right. \\
 &\quad \left. + 2\mathbb{E}_{\mathcal{D}_j} [|\langle v, \mu_j^* - \mu_1^* \rangle|^p |\langle v, \mu_j^* - \mu_1^* \rangle \geq R_{j_1}^*/5]^{1/p} \right) \\
 &\leq R_{j_1}^{*2} + d\sigma_1^{*2} + \sup_{p \geq 1} p^{-1} \left(2R_{j_1}^* \sigma_j^* (2R_{j_1}^*/(5\sigma_j^*) + \sqrt{2p}) + (8R_{j_1}^{*2}/25 + 4p + 2\pi^{1/p}(d+p-1)) \right) \\
 &\leq R_{j_1}^{*2} + d\sigma_1^{*2} + \left(R_{j_1}^{*2} + 5\sigma_j^{*2} + 3d\sigma_j^{*2} \right) \leq 3R_{j_1}^{*2} + d\sigma_1^{*2} + 3d\sigma_j^{*2}.
 \end{aligned}$$

Rest of the procedure follows similarly to the cases handled bad cases on means. We will get the statistical error of

$$\tilde{O} \left(\left(R_{j_1}^{*2} + d(\sigma_1^* \vee \sigma_j^*)^2 \right) p^{1/2c} \sqrt{1/n} \right) = \tilde{O} \left(\sqrt{1/n} \right) \left(R_{j_1}^{*2} + d(\sigma_1^* \vee \sigma_j^*)^2 \right) \exp \left(-R_{j_1}^{*2}/(128c(\sigma_j^* \vee \sigma_1^*)^2) \right),$$

with probability $1 - 1/(n^{c-2}k^{32})$.

With $j \neq 1, \mathcal{E}_{j,2}^c$: For this, we can follow exactly same procedure to $\mathcal{E}_{j,1}^c$ case to get the same result.

With $j \neq 1, \mathcal{E}_{j,3}^c$: We need to bound the p^{th} norm conditioned on $\|v\|^2/\sigma_j^{*2} \geq r^2$ where $r^2 := d + 2\sqrt{d\alpha} + 2\alpha$, where $\alpha = R_{j_1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2$.

$$\begin{aligned}
 \|Y\|_{\psi_1} &= \sup_{p \geq 1} p^{-1} \left(R_{j_1}^{*2} + d\sigma_1^{*2} + \mathbb{E}_{\mathcal{D}_j} [\|v\|^{2p} \|v\|^2/\sigma_j^{*2} \geq r^2]^{1/p} + 2R_{j_1}^* \mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle|^p \|v\|^2/\sigma_j^{*2} \geq r^2]^{1/p} \right) \\
 &\leq R_{j_1}^{*2} + d\sigma_1^{*2} \\
 &\quad + \sup_{p \geq 1} p^{-1} \left(\sigma_j^{*2} \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [\|v\|^{2p} \|v\|^2 \geq r^2]^{1/p} + 2R_{j_1}^* \sigma_j^* \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [|\langle v, s \rangle|^p \|v\|^2 \geq r^2]^{1/p} \right),
 \end{aligned}$$

where s is unit vector in direction $\mu_j^* - \mu_1^*$. We can invoke Lemma 19 to get

$$\begin{aligned}
 \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [\|v\|^{2p} \|v\|^2 \geq r^2]^{1/p} &\leq (4r^2) + 4^{1/p}(d+4p) \exp(-r^2/8p) \\
 &\leq 4d + (R_{j_1}^*/(\sigma_1^* \vee \sigma_j^*))\sqrt{d} + (R_{j_1}^{*2}/8(\sigma_1^* \vee \sigma_j^*)^2) + 4^{1/p}(d+4p) \exp(-d^2/8p) \\
 &\leq 5d + 5p + (R_{j_1}^*/\sigma_j^*)\sqrt{d} + R_{j_1}^{*2}/(8\sigma_j^{*2}).
 \end{aligned}$$

We can also invoke Lemma 20 to get

$$\begin{aligned}
 \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [|\langle v, s \rangle|^p \|v\|^2 \geq r^2]^{1/p} &\leq (64\alpha)^{1/2} + 4^{1/p}(8\alpha + 2p)^{1/2} \exp(-\alpha/2p) \\
 &\leq 2R_{j_1}^*/(\sigma_1^* \vee \sigma_j^*) + 4p^{1/2}.
 \end{aligned}$$

Now we can further continuing to bound sub-exponential norm as

$$\begin{aligned}
 \|Y\|_{\psi_1} &\leq R_{j_1}^{*2} + d\sigma_1^{*2} + \sup_{p \geq 1} p^{-1} \left(R_{j_1}^{*2}/8 + R_{j_1}^* \sigma_j^* \sqrt{d} + 5d\sigma_j^{*2} + 5p\sigma_j^{*2} + 2R_{j_1}^* \sigma_j^* + 4p^{1/2}\sigma_j^{*2} \right) \\
 &\leq 4R_{j_1}^{*2} + d\sigma_1^{*2} + 15d\sigma_j^{*2} + 3R_{j_1}^* \sigma_j^* \sqrt{d}.
 \end{aligned}$$

Hence following the same procedure for bad events, statistical fluctuation will be again smaller than

$$\tilde{O} \left(\sqrt{1/n} \left(R_{j_1}^{*2} + d(\sigma_1^* \vee \sigma_j^*)^2 \right) \exp \left(-R_{j_1}^{*2}/(128c(\sigma_j^* \vee \sigma_1^*)^2) \right) \right).$$

With $j = 1$: Finally, we need to handle this case. We recall Lemma 18, sub-exponential norm will be less than

$$\begin{aligned} \|Y\|_{\psi_1} &= \sup_{p \geq 1} p^{-1} \mathbb{E}_{\mathcal{D}_1} [|w_1(\|v\|^2 - d\sigma_1^{*2})|^p]^{1/p} \\ &\leq 3d\sigma_1^{*2}. \end{aligned}$$

We can use Proposition 8 with $n_e = 2n\pi_1^*$. Similar to $\mathcal{E}_{j,good}$ case, the statistical fluctuation is $\tilde{O}(d\sigma_1^{*2} \sqrt{\pi_1^*} \sqrt{1/n})$ with probability at least $1 - \delta/(k^2T)$.

Now collecting all $O(k)$ error terms,

$$e_\sigma \lesssim \sqrt{1/n} \left(d\sigma_1^{*2} \sqrt{\pi_1^*} + \rho_\pi \sum_{j \neq 1} (R_{j1}^{*2} + d(\sigma_j^* \vee \sigma_1^*)^2) \exp \left(-R_{j1}^{*2} / (128c(\sigma_j^* \vee \sigma_1^*)^2) \right) \right) \leq \sqrt{d}\sigma_1^{*2} \epsilon \pi_1^*,$$

with probability at least $1 - \delta/(kT) - 1/(n^{c-2} \cdot k^{31})$. Now we can conclude that,

$$\begin{aligned} |\tilde{\sigma}_1^{+2} - \sigma_1^{*2}| &\leq e_\sigma / (d\pi_1^*) + (\pi_1^+ / \tilde{\pi}_1^+) \left(\sigma_1^{+2} - \sigma_1^{*2} + \|\mu_1^+ - \mu_1^*\|^2 / d \right) - \|\tilde{\mu}_1^+ - \mu_1^*\|^2 / d \\ &\leq \sigma_1^{*2} (\epsilon / \sqrt{d} + (1 + \epsilon) \gamma D_m / \sqrt{d} + 3\epsilon D_m / d) \\ &\leq \sigma_1^{*2} (\epsilon' + \gamma D_m) / \sqrt{d}, \end{aligned}$$

with some constant rescaling of ϵ to ϵ' . ■

Appendix D. Proofs for Section 4

D.1. Proof for Lemma 10

Proof Let us first check the correctness of π_i and μ_i . This proof is reminiscent of the analysis on population EM when $D_m \geq 1/2$. The step 1 and 2, which are essentially the steps of the k-mean algorithm, can be also considered as a variant of E-step and M-step, with a rule (for 1st component):

$$\begin{aligned} \text{(E'-step)} : w_1 &= \mathbb{1}_{\|X - \mu_1\|^2 \leq \|X - \mu_j\|^2, \forall j \neq 1}, \\ \text{(M'-step)} : \pi_1^+ &= \mathbb{E}_{\mathcal{D}}[w_1], \\ \mu_1^+ &= \mathbb{E}_{\mathcal{D}}[w_1 X] / \mathbb{E}_{\mathcal{D}}[w_1]. \end{aligned} \tag{23}$$

Let us follow the proof strategy of population EM. As before, note that $\pi_1^* = \mathbb{E}_{\mathcal{D}}[w_1^*]$, where w_1^* is a weight constructed at E-step with the standard EM algorithm. Regardless of different weight assignment rules, the estimation error after one step can be represented as

$$\begin{aligned} \pi_1^+ - \pi_1^* &= \mathbb{E}_{\mathcal{D}}[w_1] - \mathbb{E}_{\mathcal{D}}[w_1^*] = \mathbb{E}_{\mathcal{D}}[\Delta_w], \\ \mu_1^+ - \mu_1^* &= \mathbb{E}_{\mathcal{D}}[w_1(X - \mu_1^*)] / \mathbb{E}_{\mathcal{D}}[w_1] = \mathbb{E}_{\mathcal{D}}[\Delta_w(X - \mu_1^*)] / \mathbb{E}_{\mathcal{D}}[w_1], \end{aligned}$$

which is exactly the same as in standard population EM. We similarly split the errors from other components and own component.

When $j \neq 1$: Let $v = X - \mu_j^*$ and define good event as

$$\mathcal{E}'_j = \{\langle v, \mu_j - \mu_1 \rangle \geq -R_{j1}^{*2}/4\}.$$

Since $\|\mu_j - \mu_1\| \leq \|\mu_j^* - \mu_1^*\| + \|\mu_j^* - \mu_j\| + \|\mu_1^* - \mu_1\| \leq 3R_{j1}^*/2$, we have $P(\mathcal{E}'_j^c) \leq \exp(-R_{j1}^{*2}/(72\sigma_j^{*2}))$. Observe that,

$$\begin{aligned} \|X - \mu_1\|^2 \leq \|X - \mu_j\|^2 &\iff 2\langle v, \mu_j^* - \mu_1 \rangle + \|\mu_j^* - \mu_1\|^2 \leq 2\langle v, \mu_j^* - \mu_j \rangle + \|\mu_j^* - \mu_j\|^2 \\ &\iff 2\langle v, \mu_j - \mu_1 \rangle \leq \|\mu_j^* - \mu_j\|^2 - \|\mu_j^* - \mu_1\|^2 \\ &\implies \langle v, \mu_j - \mu_1 \rangle \leq -R_{j1}^{*2}/4. \end{aligned}$$

Note that by the initialization condition, $\|\mu_j^* - \mu_j\| \leq R_{j1}^*/4$ and $\|\mu_j^* - \mu_1\| \geq 3R_{j1}^*/4$. That is, if $\mathbb{1}_{\mathcal{E}'_j} = 1$, then $w_1 = 0$. We can conclude that $\mathbb{E}_{\mathcal{D}_j}[w_1] \leq \exp(-R_{j1}^{*2}/(72\sigma_j^{*2}))$ for all $j \neq 1$. Now using Lemma 4, we can also see that

$$|\mathbb{E}_{\mathcal{D}_j}[\Delta_w]| \leq \mathbb{E}_{\mathcal{D}_j}[|w_1 - w_1^*|] \leq 5(1 + \pi_1^*/\pi_j^*) \exp(-R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2) + \exp(-R_{j1}^{*2}/(72\sigma_j^{*2})).$$

Summing up all errors from $j \neq 1$,

$$\sum_{j \neq 1} \pi_j^* |\mathbb{E}_{\mathcal{D}_j}[\Delta_w]| \leq 6 \sum_{j \neq 1} (\pi_1^* + \pi_j^*) \exp(-R_{j1}^{*2}/128(\sigma_1^* \vee \sigma_j^*)^2),$$

which can be bounded by some small constants $c_1 < 0.01$ with Lemma 23, given good separation condition.

Similarly, the errors to means are also small: for any unit vector $s \in \mathbb{S}^{d-1}$,

$$\begin{aligned} |\mathbb{E}_{\mathcal{D}_j}[\Delta_w \langle v, \mu_j^* - \mu_1^* + s \rangle]| &\leq R_{j1}^* |\mathbb{E}_{\mathcal{D}_j}[\Delta_w]| + |\mathbb{E}_{\mathcal{D}_j}[\Delta_w \langle v, s \rangle]| \\ &\lesssim R_{j1}^* (1 + \pi_1^*/\pi_j^*) \exp(-R_{j1}^{*2}/128(\sigma_1^* \vee \sigma_j^*)^2), \end{aligned}$$

where we applied the same technique to bound as in Corollary 5. Summing up over $j \neq 1$ and applying Lemma 23 gives the similar result, $\sum_{j \neq 1} \pi_j^* \mathbb{E}_{\mathcal{D}_j}[\Delta_w \langle X - \mu_1^*, s \rangle] \leq c_2 \sigma_1^* \pi_1^*$ for some small constant c_2 .

When $j = 1$: Recall that when we compute errors from its own component, we bounded $\mathbb{E}_{\mathcal{D}_1}[1 - w_1]$.

$$\mathbb{E}_{\mathcal{D}_1}[1 - w_1] = \sum_{l \neq 1} \mathbb{E}_{\mathcal{D}_1}[w_l] \leq \sum_{l \neq 1} \exp(-R_{l1}^{*2}/72\sigma_1^{*2}) \leq c_1,$$

for some small constant $c_1 < 0.01$. Meanwhile, in the population EM, we have shown that $\mathbb{E}_{\mathcal{D}_1}[1 - w_1^*] \leq c_2$ for small constant c_2 . Hence, $\mathbb{E}_{\mathcal{D}_1}[\Delta_w] = \mathbb{E}_{\mathcal{D}_1}[(1 - w_1) - (1 - w_1^*)] \leq c_3$ for small constant c_3 . Similarly, we can bound the errors for means,

$$|\mathbb{E}_{\mathcal{D}_1}[\Delta_w \langle v, s \rangle]| \leq c_4 \sigma_1^*,$$

for some small constant c_4 .

Collecting errors from all components gives that $|\pi_1^+ - \pi_1^*| \leq c_\pi \pi_1^*$ and $\|\mu_1^+ - \mu_1^*\| \leq c_\mu \sigma_1^*$ for some small constants $c_\pi < 0.5, c_\mu < 4$.

Population to Finite-Sample: We can reproduce the proof for finite-sample EM with modified rule (23). To see this, observe that

$$\begin{aligned}\tilde{\pi}_1^+ - \pi_1^* &= \frac{1}{n} \sum_{i=1}^n w_{1,i} - \mathbb{E}_{\mathcal{D}}[w_1], \\ \tilde{\mu}_1^+ - \mu_1^* &= \left(\frac{1}{n} \sum_{i=1}^n w_{1,i}(X - \mu_1^*) - \mathbb{E}_{\mathcal{D}}[w_1(X - \mu_1^*)] + \mathbb{E}_{\mathcal{D}}[\Delta_w(X - \mu_1^*)] \right) / \left(\frac{1}{n} \sum_{i=1}^n w_{1,i} \right),\end{aligned}$$

which is exactly in the same format as when we used standard EM iteration. Note that the proof of concentration in finite-sample EM holds for any different rule of assigning weights in E-step, as long as the probability of bad events is exponentially small. In this case it is as small as $\exp(-R_{j_1}^{*2}/72\sigma_j^{*2})$. Hence the same procedure in Appendix C can give a desired finite-sample error bound with high probability.

Estimating σ_i^{*2} : From the analysis of estimating the mixing weights, we can conclude that the elements in each cluster C_i are mostly from i^{th} component and only a few fraction of elements are from the other components (say, less than 1%). Furthermore, only less than 1% of samples from i^{th} component are missing. Thus, each cluster C_i can be considered as 2%-corrupted data from the i^{th} component. In order to retrieve σ_i^2 such that $|\sigma_i^2 - \sigma_i^{*2}| \leq 0.5\sigma_i^{*2}/\sqrt{d}$, we can consider taking a median-like quantity among pairwise distances of samples.

In each cluster, let the elements be in some fixed order which is pre-defined before we see the entire dataset. First let us consider the case when there is no corruption in each cluster. That is, each cluster has the true samples from its own component. Without loss of generality, let us focus on the first cluster C_1 . Let the elements in C_1 as X_1, X_2, \dots, X_m where $m = |C_1|$. Since all $X_i \sim \mathcal{N}(\mu_1^*, \sigma_1^{*2} I_d)$, distribution of $X_i - X_{i+1}$ follows $\mathcal{N}(0, 2\sigma_1^{*2} I_d)$. Hence, $\|X_i - X_{i+1}\|^2 / (2\sigma_1^{*2})$ is a chi-square random variable with d degrees of freedom.

Let $F(x)$ be the cdf function of a chi-square distribution with d degrees of freedom. Consider cdf value $x_l := F(d - \sqrt{d}/2)$ and $x_r := F(d + \sqrt{d}/2)$. We can numerically check that $x_r - \alpha_d \geq 0.1$ and $\alpha_d - x_l \geq 0.1$ where $\alpha_d = F(d)$ as defined in the Algorithm 1 (for large d , the pdf of chi-square distribution is very well-approximated by normal distribution).

Now let us define $r_i = \|X_{2i} - X_{2i-1}\|^2 / (2\sigma_1^{*2})$ and $r'_i = \|X_{2i+1} - X_{2i}\|^2 / (2\sigma_1^{*2})$ for $i = 1, 2, \dots, m/2 - 1$. Then let Z, Z' be the portion of r_i s such that $r_i \leq d - \sqrt{d}/2$ and $r'_i \leq d - \sqrt{d}/2$ respectively. By standard concentration of Bernoulli random variable, both Z and Z' are well concentrated around x_l with probability at least $1 - \delta/k$, given $m = \Omega(\log(k/\delta))$ samples (this holds since we generate $n = \Omega(\pi_{\min}^{-1} \log(k/\delta))$ samples from mixture distribution). Note that the key point here is, there is no probabilistic dependency between r_i s for all i , and similarly between r'_i s for all i .

Finally, we return to the 2% corrupted data from the first component. In this set, we see all adjacent pairs $\|X_{i+1} - X_i\|^2 / (2\sigma_1^{*2})$ for all $i = 1, 2, \dots, |C_i| - 1$. This is because due to the insertion of wrong samples and deletion of authentic samples, the parity of original index might have changed. By looking at all adjacent pairs, we can look at both r_i and r'_i . Note that 2% corruption can at most corrupt 4% of original r_i s and r'_i s respectively. Fortunately, we have 10% margin from α_d . That is, in the corrupted set, it is still guaranteed that α_d^{th} value is greater than $d - \sqrt{d}/2$, which is a standard argument for median-type estimators. The similar argument holds for the other direction.

In conclusion, if we take α_d^{th} value among distances of all adjacent pairs in cluster C_1 , that value is within $\sigma_1^{*2}[2d - \sqrt{d}, 2d + \sqrt{d}]$ range with high probability. We get a desired initialization parameter for variances by dividing the quantity by $2d$. \blacksquare

D.2. Proof for Lemma 11

We define some additional notations that will be used in this section. We use $P_{\mathcal{G}}(\cdot)$ to denote the probability of some event when underlying distribution is the candidate \mathcal{G} . Similarly, $P_{\mathcal{G}^*}(\cdot)$ denotes the probability when underlying distribution is true mixture \mathcal{G}^* . We use $P_{\mathcal{G}}(\cdot|X \sim j^{th})$ to denote the probability of event when X comes from j^{th} component in candidate distribution \mathcal{G} . $P_{\mathcal{G}^*}(\cdot|X \sim j^{th})$ is defined in a similar way. We use R_{j1} to denote $\|\mu_j - \mu_1^*\|$ for $j \neq 1$.

Proof Suppose the conclusion is not true, i.e., $\exists i \in [k]$ s.t. $\|\mu_i^* - \mu_j\|/\sigma_i^* \geq 16\sqrt{\log(1/\pi_{min})}$, $\forall j \in [k]$. Without loss of generality, let μ_1^* is far from all μ_j by at least $16\sqrt{\log(1/\pi_{min})}$. We consider the cases when $d \geq 128 \log(1/\pi_{min})$ and $d \leq 128 \log(1/\pi_{min})$.

Case I. $d \geq 128 \log(1/\pi_{min})$: We define an event,

$$\mathcal{E} = \{\|X - \mu_1^*\|^2/\sigma_1^{*2} + d \log \sigma_1^{*2} \leq \|X - \mu_j\|^2/\sigma_j^2 + d \log \sigma_j^2, \quad \forall j \in [k]\}.$$

Our goal is to show that $P_{\mathcal{G}^*}(\mathcal{E}) \geq 3\pi_{min}/4$ and $P_{\mathcal{G}}(\mathcal{E}) \leq \pi_{min}/2$. Then, by the definition of total variation distance, $\|\mathcal{G} - \mathcal{G}^*\|_{TV} \geq |P_{\mathcal{G}^*}(\mathcal{E}) - P_{\mathcal{G}}(\mathcal{E})| \geq \pi_{min}/4$.

Probability from true distribution: Let us first show $P_{\mathcal{G}^*}(\mathcal{E}) \geq 3\pi_{min}/4$. It suffices to show that $P_{\mathcal{G}^*}(\mathcal{E}|X \sim 1^{st}) \geq 3/4$. Thus, we are considering the event

$$\|v\|^2/\sigma_1^{*2} + d \log(\sigma_1^{*2}) \leq \|v + \mu_1^* - \mu_j\|^2/\sigma_j^2 + d \log(\sigma_j^2), \quad (24)$$

where $v \sim \mathcal{N}(0, \sigma_1^{*2}I)$. Similarly to we have seen in previous proofs for EM, we divide the cases into when $\sigma_1^* \geq \sigma_j$ and $\sigma_1^* \leq \sigma_j$.

When $\sigma_1^* \geq \sigma_j$, let $x = (\sigma_1^{*2} - \sigma_j^2)/\sigma_j^2$. After rearranging (24), we get

$$-\frac{\|v\|^2}{\sigma_1^{*2}}x + d \log(1+x) \leq \|\mu_1^* - \mu_j\|^2/\sigma_j^2 + 2\langle v, \mu_1^* - \mu_j \rangle/\sigma_j^2,$$

Then the probability of $P_{\mathcal{G}^*}(\mathcal{E}^c|X \sim 1^{st})$ is less than

$$\begin{aligned} P\left(\langle v, \mu_1^* - \mu_j \rangle \leq -\frac{\|\mu_1^* - \mu_j\|^2}{4}\right) + P\left(-\frac{\|v\|^2}{\sigma_1^{*2}}x + d \log(1+x) \geq \frac{\|\mu_1^* - \mu_j\|^2}{2\sigma_j^2}\right) \\ \leq \exp(-R_{j1}^2/32\sigma_1^{*2}) + P\left(\frac{\|v\|^2}{d\sigma_1^{*2}} \leq \frac{1}{x} \left(\log(1+x) - \frac{R_{j1}^2}{2d\sigma_j^2}\right)\right). \end{aligned}$$

Using the similar trick as before, first consider when $0 \leq x \leq 3/4$. Then,

$$P\left(\frac{\|v\|^2}{d\sigma_1^{*2}} \leq \frac{\log(1+x)}{x} - \frac{R_{j1}^2}{2d\sigma_j^2x}\right) \leq P\left(\frac{\|v\|^2}{d\sigma_1^{*2}} \leq 1 - \frac{x}{4} - \frac{R_{j1}^2}{2d\sigma_j^2x}\right)$$

$$\leq P\left(\frac{\|v\|^2}{d\sigma_1^{*2}} \leq 1 - 2\sqrt{\frac{R_{j1}^2}{2d\sigma_j^2}}\right) \leq \exp(-R_{j1}^2/2\sigma_j^2).$$

When $x \geq 3/4$,

$$\begin{aligned} P\left(\frac{\|v\|^2}{d\sigma_1^{*2}} \leq \frac{\log(1+x)}{x} - \frac{R_{j1}^2}{2d\sigma_j^2 x}\right) &\leq P\left(\frac{\|v\|^2}{d\sigma_1^{*2}} \leq 1 - \frac{1}{4} - \frac{R_{j1}^2}{2d(\sigma_1^{*2} - \sigma_j^2)}\right) \\ &\leq P\left(\frac{\|v\|^2}{d\sigma_1^{*2}} \leq 1 - 2\sqrt{\frac{R_{j1}^2}{8d(\sigma_1^{*2} - \sigma_j^2)}}\right) \leq \exp(-R_{j1}^2/8\sigma_1^{*2}). \end{aligned}$$

In either case, $P_{G^*}(\mathcal{E}^c | X \sim 1^{st}) \leq 2\exp(-R_{j1}^2/32\sigma_1^{*2}) \leq 2\pi_{min}^8 \leq 1/4$.

When $\sigma_1^* \leq \sigma_j$, let $x = (\sigma_j^2 - \sigma_1^{*2})/\sigma_j^2$ and we similarly rearrange (24).

$$\frac{\|v\|^2}{\sigma_1^{*2}} x + d \log(1-x) \leq \|\mu_1^* - \mu_j\|^2/\sigma_j^2 + 2\langle v, \mu_1^* - \mu_j \rangle/\sigma_j^2.$$

We need to divide the cases when $\sigma_j^2 \leq 8\sigma_1^{*2}$ and $\sigma_j^2 \geq 8\sigma_1^{*2}$. When $\sigma_j^2 \leq 8\sigma_1^{*2}$, we proceed similarly to previous cases,

$$\begin{aligned} P(\mathcal{E}^c | X \sim 1^{st}) &\leq P\left(\langle v, \mu_1^* - \mu_j \rangle \leq -\frac{\|\mu_1^* - \mu_j\|^2}{4}\right) + P\left(\frac{\|v\|^2}{\sigma_1^{*2}} x + d \log(1-x) \geq \frac{\|\mu_1^* - \mu_j\|^2}{2\sigma_j^2}\right) \\ &\leq \exp(-R_{j1}^2/32\sigma_1^{*2}) + P\left(\frac{\|v\|^2}{d\sigma_1^{*2}} \geq \frac{1}{x} \left(-\log(1-x) + \frac{R_{j1}^2}{2d\sigma_j^2}\right)\right) \\ &\leq \exp(-R_{j1}^2/32\sigma_1^{*2}) + P\left(\frac{\|v\|^2}{d\sigma_1^{*2}} \geq 1 + \frac{x}{2} + \frac{R_{j1}^2}{2dx\sigma_j^2}\right) \\ &\leq \exp(-R_{j1}^2/32\sigma_1^{*2}) + P\left(\frac{\|v\|^2}{d\sigma_1^{*2}} \geq 1 + 2\sqrt{\frac{R_{j1}^2}{8d\sigma_j^2}} + 2\frac{R_{j1}^2}{8dx\sigma_j^2}\right) \\ &\leq \exp(-R_{j1}^2/32\sigma_1^{*2}) + \exp(-R_{j1}^2/8\sigma_j^2) \leq \exp(-R_{j1}^2/32\sigma_1^{*2}) + \exp(-R_{j1}^2/64\sigma_1^{*2}) \\ &\leq \pi_{min}^8 + \pi_{min}^4 \leq 1/4. \end{aligned}$$

When $\sigma_j^2 \geq 8\sigma_1^{*2}$, we first note that $x \geq 7/8$. Thus, $-\log(1-x)/x \geq 2.376$. We can then bound the term simply as

$$\begin{aligned} P(\mathcal{E}^c | X \sim 1^{st}) &\leq P\left(\langle v, \mu_1^* - \mu_j \rangle \leq -\frac{\|\mu_1^* - \mu_j\|^2}{4}\right) + P\left(\frac{\|v\|^2}{\sigma_1^{*2}} x + d \log(1-x) \geq \frac{\|\mu_1^* - \mu_j\|^2}{2\sigma_j^2}\right) \\ &\leq \pi_{min}^8 + P\left(\frac{\|v\|^2}{d\sigma_1^{*2}} \geq 1 + 1.376\right). \end{aligned}$$

Using standard tail probability for chi-Square distribution, we get

$$P\left(\frac{\|v\|^2}{d\sigma_1^{*2}} \geq 1 + 1.376\right) \leq \exp(-1.376d/8) \leq \pi_{min}^{20},$$

since we are considering the case when $d \geq 128 \log(1/\pi_{min})$. Therefore, we get $P(\mathcal{E}^c | X \sim 1^{st}) \leq 1/4$ in all cases. It concludes that $P_{G^*}(\mathcal{E}) \geq 3\pi_{min}/4$.

Probability from candidates: Now we show that $P_{\mathcal{G}}(\mathcal{E}) \leq \pi_{\min}/2$. Toward this goal, we need to show that for each $j \in [k]$, $P_{\mathcal{G}}(\mathcal{E}|X \sim j^{th}) \leq \pi_{\min}/2$. The corresponding event becomes

$$\|v + \mu_j - \mu_1^*\|^2/\sigma_1^{*2} + d \log(\sigma_1^{*2}) \leq \|v\|^2/\sigma_j^2 + d \log(\sigma_j^2). \quad (25)$$

Let us first consider the case when $\sigma_j \leq \sigma_1^*$. Now we set $x = (\sigma_1^{*2} - \sigma_j^2)/\sigma_1^{*2}$. Then the rearrangement of (25) gives

$$\|\mu_1^* - \mu_j\|^2/\sigma_1^{*2} + 2\langle v, \mu_j - \mu_1^* \rangle/\sigma_1^{*2} \leq \frac{\|v\|^2}{\sigma_j^2} x + d \log(1 - x).$$

We bound the probability of this event similarly to previous cases.

$$\begin{aligned} P(\mathcal{E}|X \sim j^{th}) &\leq P\left(\langle v, \mu_j - \mu_1^* \rangle \leq -\frac{\|\mu_1^* - \mu_j\|^2}{4}\right) + P\left(\frac{\|v\|^2}{\sigma_j^2} x + d \log(1 - x) \geq \frac{\|\mu_1^* - \mu_j\|^2}{2\sigma_1^{*2}}\right) \\ &\leq \exp(-R_{j1}^2/32\sigma_j^2) + P\left(\frac{\|v\|^2}{d\sigma_j^2} \geq \frac{1}{x} \left(-\log(1 - x) + \frac{R_{j1}^2}{2d\sigma_1^{*2}}\right)\right) \\ &\leq \exp(-R_{j1}^2/32\sigma_1^{*2}) + P\left(\frac{\|v\|^2}{d\sigma_j^2} \geq 1 + \frac{R_{j1}^2}{2dx\sigma_j^2}\right) \\ &\leq \exp(-R_{j1}^2/32\sigma_1^{*2}) + \exp(-R_{j1}^2/(16(\sigma_1^{*2} - \sigma_j^2))) \\ &\leq \pi_{\min}^8 + \pi_{\min}^{16} \leq \pi_{\min}/2. \end{aligned}$$

When $\sigma_j \geq \sigma_1^*$, the proof should be more delicate. First we rearrange (25) to see that

$$\|\mu_1^* - \mu_j\|^2/\sigma_1^{*2} + 2\langle v, \mu_j - \mu_1^* \rangle/\sigma_1^{*2} \leq -\frac{\|v\|^2}{\sigma_j^2} x + d \log(1 + x),$$

is the event to bound, where $x = (\sigma_j^2 - \sigma_1^{*2})/\sigma_1^{*2}$. Here, we will consider three cases, $0 \leq x \leq 3/4$, $3/4 \leq x \cap R_{j1}^2/\sigma_j^2 \geq 32 \log(1/\pi_{\min})$, and $R_{j1}^2/\sigma_j^2 \leq 32 \log(1/\pi_{\min})$. First, if $0 \leq x \leq 3/4$, then $\log(1 + x) \leq x - x^2/4$ and $\sigma_j^2 \leq 2\sigma_1^{*2}$, thus

$$\begin{aligned} P(\mathcal{E}|X \sim j^{th}) &\leq P\left(\langle v, \mu_j - \mu_1^* \rangle \leq -\frac{7}{16}\|\mu_1^* - \mu_j\|^2\right) + P\left(\frac{-\|v\|^2}{\sigma_j^2} x + d \log(1 + x) \geq \frac{\|\mu_1^* - \mu_j\|^2}{8\sigma_1^{*2}}\right) \\ &\leq \exp(-49R_{j1}^2/512\sigma_j^2) + P\left(\frac{\|v\|^2}{d\sigma_j^2} \leq \frac{1}{x} \left(\log(1 + x) - \frac{R_{j1}^2}{8d\sigma_1^{*2}}\right)\right) \\ &\leq \exp(-R_{j1}^2/32\sigma_1^{*2}) + P\left(\frac{\|v\|^2}{d\sigma_j^2} \leq 1 - \frac{x}{4} - \frac{R_{j1}^2}{8dx\sigma_1^{*2}}\right) \\ &\leq \exp(-R_{j1}^2/32\sigma_1^{*2}) + \exp(-R_{j1}^2/(32\sigma_1^{*2})) \leq \pi_{\min}^8 + \pi_{\min}^8 \leq \pi_{\min}/2. \end{aligned}$$

If $3/4 \leq x$ and $R_{j1}^2/\sigma_j^2 \geq 32 \log(1/\pi_{\min})$, then $\log(1 + x)/x \leq 3/4$, thus

$$P(\mathcal{E}|X \sim j^{th}) \leq P\left(\langle v, \mu_j - \mu_1^* \rangle \leq -\frac{7}{16}\|\mu_1^* - \mu_j\|^2\right) + P\left(\frac{-\|v\|^2}{\sigma_j^2} x + d \log(1 + x) \geq \frac{\|\mu_1^* - \mu_j\|^2}{8\sigma_1^{*2}}\right)$$

$$\begin{aligned}
 &\leq \exp(-49R_{j1}^2/512\sigma_j^2) + P\left(\frac{\|v\|^2}{d\sigma_j^2} \leq \frac{1}{x} \left(\log(1+x) - \frac{R_{j1}^2}{8d\sigma_1^{*2}}\right)\right) \\
 &\leq \exp(-6\log(1/\pi_{min})) + P\left(\frac{\|v\|^2}{d\sigma_j^2} \leq 1 - \frac{1}{4} - \frac{R_{j1}^2}{8dx\sigma_1^{*2}}\right) \\
 &\leq \exp(-6\log(1/\pi_{min})) + P\left(\frac{\|v\|^2}{d\sigma_j^2} \leq 1 - 2\sqrt{\frac{1}{64} + \frac{R_{j1}^2}{64d(\sigma_j^2 - \sigma_1^{*2})}}\right) \\
 &\leq \exp(-6\log(1/\pi_{min})) + \exp(-d/64 - R_{j1}^2/64\sigma_j^2) \leq \pi_{min}^6 + \pi_{min}^{2.5} \leq \pi_{min}/2.
 \end{aligned}$$

Finally, if $R_{j1}^2/\sigma_j^2 \leq 32\log(1/\pi_{min})$, we take a different path. First of all, this can only happen when $\sigma_j^2 \geq 8\sigma_1^{*2}$ and $x \geq 7$, and $R_{j1}^2/\sigma_j^2 \leq d/4$. Using rotational invariance property of Gaussian with (scale of) identity covariance, without loss of generality, we can set $v = v_1\hat{e}_{j1} + v_{2:d}$ where \hat{e}_{j1} is a unit vector in direction $\mu_j - \mu_1^*$, and $v_{2:d}$ is the rest $d-1$ dimensional orthogonal component. Then, we can rearrange the event as

$$\begin{aligned}
 &(v_1 + R_{j1})^2/\sigma_1^{*2} - v_1^2/\sigma_j^2 \leq -\frac{\|v_{2:d}\|^2}{\sigma_j^2}x + d\log(1+x), \\
 \implies &\left(\frac{1}{\sigma_1^{*2}} - \frac{1}{\sigma_j^2}\right) \left(v_1 - \left(\frac{1}{\sigma_1^{*2}} - \frac{1}{\sigma_j^2}\right)^{-1} \frac{R_{j1}}{\sigma_1^{*2}}\right)^2 - \left(\frac{1}{\sigma_1^{*2}} - \frac{1}{\sigma_j^2}\right)^{-1} \frac{R_{j1}^2}{\sigma_1^{*4}} + \frac{R_{j1}^2}{\sigma_1^{*2}} \\
 &= \frac{x}{\sigma_j^2} \left(v_1 - \frac{\sigma_j^2 R_{j1}}{x \sigma_1^{*2}}\right)^2 - \frac{R_{j1}^2}{\sigma_j^2 - \sigma_1^{*2}} \leq -\frac{\|v_{2:d}\|^2}{\sigma_j^2}x + d\log(1+x) \\
 \implies &\frac{\|v_{2:d}\|^2}{\sigma_j^2} \leq d\frac{\log(1+x)}{x} + \frac{1}{x} \frac{R_{j1}^2}{\sigma_j^2 - \sigma_1^{*2}} \\
 \implies &\frac{\|v_{2:d}\|^2}{(d-1)\sigma_j^2} \leq \left(1 + \frac{1}{d-1}\right) \left(\frac{\log 8}{7} + \frac{2}{49}\right) \leq \frac{1}{2}.
 \end{aligned}$$

And, the probability of $P(\|v_{2:d}\|^2/\sigma_j^2 \leq (d-1)/2) \leq \exp(-(d-1)/16) \leq \pi_{min}^4$ upper-bounds the probability $P_{\mathcal{G}}(\mathcal{E}|X \sim j^{th})$.

Combining all cases, we can conclude that $P_{\mathcal{G}}(\mathcal{E}) \leq \pi_{min}/2$. Hence, we get $\|\mathcal{G}^* - \mathcal{G}\|_{TV} \geq \pi_{min}/4$ as desired when $d \geq 128\log(1/\pi_{min})$.

Case II. $d \leq 128\log(1/\pi_{min})$: In this case, we consider the following two events.

$$\begin{aligned}
 \mathcal{E}_1 &= \{\|X - \mu_1^*\| \leq \sigma_1^* \sqrt{2d}\}, \\
 \mathcal{E}_2 &= \{\|X - \mu_1^*\| \leq 3\sigma_1^* \sqrt{2d}\}.
 \end{aligned}$$

We will show that either $P_{\mathcal{G}^*}(\mathcal{E}_1) - P_{\mathcal{G}}(\mathcal{E}_1) \geq \pi_{min}/4$ or $P_{\mathcal{G}}(\mathcal{E}_1^c \cap \mathcal{E}_2) - P_{\mathcal{G}^*}(\mathcal{E}_1^c \cap \mathcal{E}_2) \geq \pi_{min}/4$. As a first step, we show that $P_{\mathcal{G}^*}(\mathcal{E}_1|X \sim 1^{st}) > 3/4$, and $P_{\mathcal{G}^*}(\mathcal{E}_2|X \sim j^{th}) \ll \pi_{min}/100$ for $j = 2, \dots, k$.

It is easy to see that

$$P_{\mathcal{G}^*}(\mathcal{E}_1|X \sim 1^{st}) = P_{v \sim \mathcal{N}(0, I)}(\|v\|^2 \leq 2d) \geq 3/4,$$

using the pre-computed cdf value of chi-square distribution with degree d (we can numerically check it for small d , and we can approximate it with approximation to normal distribution for large d). For other components,

$$\begin{aligned}
P_{\mathcal{G}^*}(\mathcal{E}_2 | X \sim j^{th}) &= P_{v \sim \mathcal{N}(0, \sigma_j^{*2} I)}(\|v + \mu_j^* - \mu_1^*\|^2 \leq 18\sigma_1^* d) \\
&\leq P_{v \sim \mathcal{N}(0, \sigma_j^{*2} I)}((\langle v, \hat{e}_{j1} \rangle + R_{j1}^*)^2 \leq 18\sigma_1^* d) \\
&\leq P_{v \sim \mathcal{N}(0, \sigma_j^{*2} I)}(\langle v, \hat{e}_{j1} \rangle \leq 3\sigma_1^* \sqrt{2d} - R_{j1}^*) \\
&\leq P_{v \sim \mathcal{N}(0, \sigma_j^{*2} I)}(\langle v, \hat{e}_{j1} \rangle \leq -R_{j1}^*/4) \\
&\leq \exp(-R_{j1}^{*2}/32\sigma_j^{*2}) \leq \pi_{min}^{64},
\end{aligned}$$

where \hat{e}_{j1} is a unit vector in direction $\mu_j^* - \mu_1^*$. Note that $3\sigma_1^* \sqrt{2d} \leq 48\sigma_1^* \sqrt{\log(1/\pi_{min})} \leq 3R_{j1}^*/4$. Combining two facts, it is easy to see that $P_{\mathcal{G}^*}(\mathcal{E}_1) \geq 3\pi_{min}/4$ and $P_{\mathcal{G}^*}(\mathcal{E}_1^c \cap \mathcal{E}_2) \leq \pi_{min}/4$.

Now we show that either $P_{\mathcal{G}}(\mathcal{E}_1) \leq \pi_{min}/2$ or $P_{\mathcal{G}}(\mathcal{E}_1^c \cap \mathcal{E}_2) \geq \pi_{min}/2$ is true. Suppose $P_{\mathcal{G}}(\mathcal{E}_1) \geq \pi_{min}/2$. Observe that $\|\mu_j - \mu_1^*\|/\sigma_1^* \geq 16\sqrt{\log(1/\pi_{min})} \geq \sqrt{2d}$ for all j . That is, all μ_j are outside of the sphere that \mathcal{E}_1 considers. Therefore, if we imagine a bigger ball of radius $3\sqrt{2d}$, for any j , there exists a ball of radius $\sqrt{2d}$ in $\mathcal{E}_2 \cap \mathcal{E}_1^c$ where the contribution from j^{th} component is larger than the ball considered in \mathcal{E}_1 . Since this is true for all j , we can conclude that $P_{\mathcal{G}}(\mathcal{E}_1^c \cap \mathcal{E}_2) \geq \pi_{min}/2$.

In conclusion, $\|\mathcal{G} - \mathcal{G}^*\|_{TV} \geq \pi_{min}/4$ if any one of μ_i^* cannot find a good initializer in candidate parameters. Note that this result does not assume any separation condition in candidate distributions. Neither, this lower bound for TV distance does not depend on any other parameters but π_{min} . ■

D.3. Proof of Theorem 9

Proof Lemma 11 indicates that if $\|\mathcal{G} - \mathcal{G}^*\|_{TV} \leq \pi_{min}/4$, then we have initializers that satisfy the requirement (15) up to some permutation in \mathcal{G} . Note that since the true mixture distribution \mathcal{G}^* satisfies the separation condition (14), when (16) holds true, each i^{th} component has its unique j for initial mean μ_j : if $\|\mu_j - \mu_i^*\|$ is less than $16\sigma_i^* \sqrt{\log(1/\pi_{min})} < \frac{1}{4}\|\mu_i^* - \mu_{i'}^*\|$ for $i' \neq i$, then μ_j is at least $\frac{3}{4}\|\mu_{i'}^* - \mu_i^*\| > 48\sigma_{i'}^* \sqrt{\log(1/\pi_{min})}$ far apart from other $\mu_{i'}^*$. Hence, one μ_j can only be associated with only one μ_i^* and vice versa.

We first show how to get the sample-optimal guarantee for the proper-learning of spherical Gaussian mixtures using the sample-compression scheme introduced in Ashtiani et al. (2018a). The compressibility of a distribution is (informally) defined as follows:

Definition 25 (Informal Definition of Compressibility in Ashtiani et al. (2018a)) For any $\epsilon > 0$, a distribution \mathcal{F} is called $(\tau(\epsilon), t(\epsilon), m(\epsilon))$ -compressible if the following holds: if $m(\epsilon)$ samples are drawn from \mathcal{F} , we can carefully select $\tau(\epsilon)$ samples (among $m(\epsilon)$ samples) and additional $t(\epsilon)$ -bits such that a pre-defined systematic procedure (decoder) takes them as inputs, and returns a distribution $\hat{\mathcal{F}}$ that satisfies $\|\hat{\mathcal{F}} - \mathcal{F}\|_{TV} \leq \epsilon$ with high probability.

See Definition 3.1 and 3.2 in their work Ashtiani et al. (2018a) for more details. Their key result (see Theorem 3.5 in Ashtiani et al. (2018a)) states that if a distribution is $(\tau(\epsilon), t(\epsilon), m(\epsilon))$ -compressible, then $\tilde{O}(m(\epsilon) + (\tau(\epsilon) + t(\epsilon))/\epsilon^2)$ samples suffice to learn a ϵ -close distribution in TV distance. The optimal sample upper bound $\tilde{O}(kd/\epsilon^2)$ for learning a mixture of axis-aligned Gaussians then follows

by (i) showing that a single axis-aligned Gaussian is $(O(d), O(d \log(d/\epsilon)), O(d))$ -compressible, and (ii) using their Lemma 3.7 in [Ashtiani et al. \(2018a\)](#) to conclude that a mixture of k axis-aligned Gaussians is $(O(kd), O(kd \log(d/\epsilon)), \tilde{O}(kd/\epsilon))$ -compressible.

Given their argument for a mixture of axis-aligned Gaussians, it is straight-forward to get the same result for spherical Gaussians. We only need to show that a single spherical Gaussian is also $(O(d), O(d \log(d/\epsilon)), O(d))$ -compressible. Then we can use the same argument using their Lemma 3.7 to conclude that a mixture of k spherical Gaussians is also $(O(kd), O(kd \log(d/\epsilon)), \tilde{O}(kd/\epsilon))$ -compressible, hence $\tilde{O}(kd/\epsilon^2)$ samples suffice to learn a ϵ -close distribution in TV distance.

Suppose a single axis-aligned Gaussian with mean $\mu = (\mu_1, \mu_2, \dots, \mu_d) \in \mathbb{R}^d$ and covariance $\Sigma = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_d)$. The decoder they construct for an axis-aligned Gaussian outputs $\hat{\mu}$ and $\hat{\Sigma} = \text{diag}(\hat{\sigma}_1, \hat{\sigma}_2, \dots, \hat{\sigma}_d)$ such that

$$|\mu_i - \hat{\mu}_i| \leq \sigma_i \epsilon / d, \quad |\sigma_i - \hat{\sigma}_i| \leq \sigma_i \epsilon / d, \quad \forall i \in [k],$$

which hence guarantees that $\|\mathcal{N}(\mu, \Sigma) - \mathcal{N}(\hat{\mu}, \hat{\Sigma})\|_{TV} \leq \epsilon$. For a spherical Gaussian, we can use the same decoder by considering it as an axis-aligned Gaussian, and simply pick σ_1 as a common scale factor of an identity matrix for a spherical Gaussian. Therefore, a spherical Gaussian is also compressible with the same parameters $(O(d), O(d \log(d/\epsilon)), O(d))$. Combining with Lemma 3.7 and Theorem 3.5 in [Ashtiani et al. \(2018a\)](#), we obtain a $\tilde{O}(kd/\epsilon^2)$ sample-complexity guarantee for the proper-learning of spherical Gaussian mixtures.

Now we can first get the candidate distribution \mathcal{G} using the algorithm given in [Ashtiani et al. \(2018a\)](#) with $\tilde{O}(kd\pi_{min}^{-2})$ samples to get $\|\mathcal{G} - \mathcal{G}^*\|_{TV} \leq \pi_{min}/4$. Then we can run Algorithm 1 and then the EM algorithm using $\tilde{O}(d\pi_{min}^{-1}/\epsilon^2)$ samples. This gives the algorithm for Theorem 9. \blacksquare

Appendix E. Deferred Proof: Convergence of Population EM when $D_m \leq 1/2$.

We define a target error $D_m = \max_j \left(\|\mu_j - \mu_j^*\|/\sigma_j^*, |\pi_j - \pi_j^*|/\pi_j^*, \sqrt{d}|\sigma_j^2 - \sigma_j^{*2}|/\sigma_j^{*2} \right) \leq 1/2$.

Proof First of all, we differentiate our EM operator with respect to all variables being estimated (σ_i^2 are considered as a single variable). For instance,

$$\begin{aligned} \pi_1^+ - \pi_1^* &= \mathbb{E}_{\mathcal{D}}[w_1] - \mathbb{E}_{\mathcal{D}}[w_1^*] = \mathbb{E}_{\mathcal{D}}[\Delta_w^u], \\ \mu_1^+ - \mu_1^* &= \mathbb{E}_{\mathcal{D}}[\Delta_w^u(X - \mu_1^*)]/\mathbb{E}_{\mathcal{D}}[w_1], \\ \sigma_1^{+2} - \sigma_1^{*2} &= \mathbb{E}_{\mathcal{D}}[\Delta_w^u(\|X - \mu_1^*\|^2 - d\sigma_1^{*2})]/(d\pi_1^+) - \|\mu_1^+ - \mu_1^*\|^2/d, \end{aligned}$$

where

$$\begin{aligned} \Delta_w^u &= -w_1^u(1 - w_1^u)(X - \mu_1^u)^T(\mu_1 - \mu_1^*)/\sigma_1^{u2} + \sum_{l \neq 1} w_1^u w_l^u (X - \mu_l^u)^T(\mu_l - \mu_l^*)/\sigma_l^{u2} \\ &\quad - w_1^u(1 - w_1^u)(\pi_1 - \pi_1^*)/\pi_1^u + \sum_{l \neq 1} w_1^u w_l^u (\pi_l - \pi_l^*)/\pi_l^u \\ &\quad - w_1^u(1 - w_1^u)(+\|X - \mu_1^u\|^2/(2\sigma_1^{u4}) - d/(2\sigma_1^{u2}))(\sigma_1^2 - \sigma_1^{*2}) \\ &\quad + \sum_{l \neq 1} w_1^u w_l^u (+\|X - \mu_l^u\|^2/(2\sigma_l^{u4}) - d/(2\sigma_l^{u2}))(\sigma_l^2 - \sigma_l^{*2}), \end{aligned}$$

where w_1^u is a weight constructed with $\mu_1^u := \mu_1^* + u(\mu_1 - \mu_1^*)$ for some $u \in [0, 1]$, and other u scripted variables are defined similarly. In addition to previous technical lemmas, we state a few more helper lemmas. The same result holds for $\mathbb{E}_{\mathcal{D}_j}[\Delta_w]$ as in this corollary. With the above lemma, we can bound the errors for mixing weights from other components. We need one more lemma for bounding the sum of errors.

Lemma 26 *For values $q \in \{0, 1, 2, 3, 4\}$, the following summations are bounded by*

$$\sum_{j \neq 1} (\pi_1^* + \pi_j^*) (R_{j1}^*/\sigma_1^*)^q \exp(-R_{j1}^{*2}/128(\sigma_1^* \vee \sigma_j^*)^2) \leq c_q \pi_1^*, \quad (26)$$

$$\sum_{j \neq 1} (\pi_1^* + \pi_j^*) (R_{j1}^*/\sigma_j^*)^q \exp(-R_{j1}^{*2}/128(\sigma_1^* \vee \sigma_j^*)^2) \leq c_q \pi_1^*, \quad (27)$$

with some small constant c_q given large enough universal constant C .

The proof is similar to Lemma 23, and we will allow much larger universal constant $C \geq 128$ in the separation condition (2) to lighten the algebraic burden. The proof will be given in Appendix E.4.

From this point, due to the heavy calculation and algebra, we give up tracking most constants in error bounds. Bounds will be often given in $O(\cdot)$ notation, but we note that the hidden constants will not be too large.

E.1. Convergence of Mixing Weights

Let us start with the simplest targets. Similarly to $D_m \geq 1/2$, we consider errors from other components first.

Errors from other components $j \neq 1$: We first bound

$$\begin{aligned} e_{j1} &= |\mathbb{E}_{\mathcal{D}_j}[w_1^u(1-w_1^u)(X-\mu_1^u)^T(\mu_1-\mu_1^*)]/\sigma_1^{u2}| \\ &\leq 2\|\mathbb{E}_{\mathcal{D}_j}[w_1^u(1-w_1^u)(X-\mu_1^u)]/\sigma_1^*\| \|(\mu_1-\mu_1^*)/\sigma_1^*\| \\ &\leq 2D_m \sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_j}[w_1^u \langle X-\mu_1^u, s \rangle] / \sigma_1^*. \end{aligned}$$

Using the Lemma 6,

$$\mathbb{E}_{\mathcal{D}_j}[|w_1^u \langle X-\mu_1^u, s \rangle|] \leq O\left(\left(1 + \pi_1^*/\pi_j^*\right) R_{j1}^* \exp(-R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2)\right),$$

which yields $e_{j1} \leq D_m O\left(\left(1 + \pi_1^*/\pi_j^*\right) (R_{j1}^*/\sigma_1^*) \exp(-R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2)\right)$.

The second target is

$$\begin{aligned} e_{j2} &= |\mathbb{E}_{\mathcal{D}_j}[\sum_{l \neq 1} w_1^u w_l^u (X-\mu_l^u)^T(\mu_l-\mu_l^*)]/\sigma_l^{u2}| \\ &\leq 2 \sum_{l \neq 1} \|\mathbb{E}_{\mathcal{D}_j}[w_1^u w_l^u (X-\mu_l^u)]\| \|\mu_l-\mu_l^*\| / \sigma_l^{*2} \\ &\leq 2D_m \sum_{l \neq 1} \|\mathbb{E}_{\mathcal{D}_j}[w_1^u w_l^u (X-\mu_l^u)]/\sigma_l^*\| \end{aligned}$$

$$\leq 2D_m \sqrt{\mathbb{E}_{\mathcal{D}_j}[w_1^u]} \sum_{l \neq 1} \underbrace{\sqrt{\sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_j}[w_l^u \langle X - \mu_l^u, s \rangle^2 / \sigma_l^{*2}]}_I.$$

Main challenge is to show that $\sum_{l \neq 1} \sqrt{\sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_j}[w_l^u \langle X - \mu_l^u, s \rangle^2 / \sigma_l^{*2}] = O(1)}$. As this will appear several times, we state a helping lemma.

Lemma 27 *The summation of term I over $l \neq j$ is bounded by*

$$\sum_{l \neq j} \sqrt{\sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_j}[w_l^u \langle X - \mu_l^u, s \rangle^2 / \sigma_l^{*2}] \leq c,$$

for some small constant c . When $l = j$, $I = O(1)$.

The proof of lemma is given in Section E.4. By the lemma, the summation over the entire term is $O(1)$. The term $\sqrt{\mathbb{E}_{\mathcal{D}_j}[w_1^u]}$ are less than

$$O\left(\sqrt{(1 + \pi_1^*/\pi_j^*) \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2)}\right).$$

Therefore, we can conclude that

$$e_{j2} \leq D_m O\left((1 + \pi_1^*/\pi_j^*) \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2)\right).$$

The third target is

$$\begin{aligned} e_{j3} &= |\mathbb{E}_{\mathcal{D}_j}[w_1^u(1 - w_1^u)(\pi_1 - \pi_1^*)/\pi_1^u]| \\ &\leq 2D_m |\mathbb{E}_{\mathcal{D}_j}[w_1^u]| \leq D_m O\left((1 + \pi_1^*/\pi_j^*) \exp(-R_{j1}^{*2}/64\sigma_1^{*2})\right). \end{aligned}$$

The fourth target is

$$\begin{aligned} e_{j4} &= |\mathbb{E}_{\mathcal{D}_j}[\sum_{l \neq 1} w_1^u w_l^u (\pi_1 - \pi_1^*)/\pi_1^u]| \\ &\leq 2D_m |\mathbb{E}_{\mathcal{D}_j}[\sum_{l \neq 1} w_1^u w_l^u]| \leq 2D_m \mathbb{E}_{\mathcal{D}_j}[w_1^u] \leq D_m O\left((1 + \pi_1^*/\pi_j^*) \exp(-R_{j1}^{*2}/64\sigma_1^{*2})\right), \end{aligned}$$

since $\sum_{l \neq 1} w_l^u \leq 1$.

The fifth target is

$$\begin{aligned} e_{j5} &= \left| \mathbb{E}_{\mathcal{D}_j} [w_1^u(1 - w_1^u)(\|X - \mu_1^u\|^2 - d\sigma_1^{u2})/(2\sigma_1^{u2})] (\sigma_1^2 - \sigma_1^{*2})/\sigma_1^{u2} \right| \\ &\leq \left(2D_m/\sqrt{d} \right) |\mathbb{E}_{\mathcal{D}_j} [w_1^u(1 - w_1^u)(\|X - \mu_1^u\|^2 - d\sigma_1^{u2})/(2\sigma_1^{u2})]| \\ &\leq \left(2D_m/\sigma_1^{*2}\sqrt{d} \right) |\mathbb{E}_{\mathcal{D}_j} [w_1^u(1 - w_1^u)(\|v\|^2 - d\sigma_j^{*2} + 2\langle v, \mu_j^* - \mu_1^u \rangle \\ &\quad + \|\mu_j^* - \mu_1^u\|^2 + d(\sigma_j^{*2} - \sigma_1^{*2}) + d(\sigma_1^{*2} - \sigma_1^{u2}))]| \\ &\leq \left(2D_m/\sigma_1^{*2}\sqrt{d} \right) \mathbb{E}_{\mathcal{D}_j} \left[\left| w_1^u(\|v\|^2 - d\sigma_j^{*2}) \right| + 2|w_1^u \langle v, \mu_j^* - \mu_1^u \rangle| \right] \end{aligned}$$

$$+ 2 \left[w_1^u R_{j1}^{*2} \right] + \left| w_1^u d(\sigma_j^{*2} - \sigma_1^{*2}) \right| + w_1^u d \left(\sigma_1^{*2} D_m / \sqrt{d} \right) \Big].$$

We have seen similar terms in $D_m \geq 1/2$ case. Each term we can bound as

$$\begin{aligned} \mathbb{E}_{\mathcal{D}_j} [|w_1^u (\|v\|^2 - d\sigma_j^{*2})|] &\leq \sqrt{\mathbb{E}_{\mathcal{D}_j} [|w_1^u|]} \sqrt{\mathbb{E}_{\mathcal{D}_j} [(\|v\|^2 - d\sigma_j^{*2})^2]} \\ &\leq 4\sigma_j^{*2} \sqrt{d(1 + \pi_1^*/\pi_j^*)} \exp(-R_{j1}^{*2}/128(\sigma_1^* \vee \sigma_j^*)^2), \end{aligned}$$

$$\mathbb{E}_{\mathcal{D}_j} [|w_1^u \langle v, \mu_j^* - \mu_1^u \rangle|] \leq O \left((1 + \pi_1^*/\pi_j^*) R_{j1}^* \sigma_j^* \exp(-R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2) \right),$$

$$d(\sigma_j^{*2} - \sigma_1^{*2}) \mathbb{E}_{\mathcal{D}_j} [|w_1^u|] \leq O \left(R_{j1}^* (\sigma_j^* \vee \sigma_1^*) (1 + \pi_1^*/\pi_j^*) \sqrt{d} \exp(-R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2) \right),$$

and the rest terms can also be easily bounded. Thus, we have shown that

$$e_{j5} \leq D_m O \left(R_{j1}^* (\sigma_j^* \vee \sigma_1^*) / \sigma_1^{*2} \cdot (1 + \pi_1^*/\pi_j^*) \exp(-R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2) \right).$$

Finally, we control the last term.

$$\begin{aligned} e_{j6} &= \left| \mathbb{E}_{\mathcal{D}_j} \left[\sum_{l \neq 1} w_1^u w_l^u (\|X - \mu_l^u\|^2 - d\sigma_l^{u2}) / (2\sigma_l^{u2}) \cdot (\sigma_l^2 - \sigma_l^{*2}) / \sigma_l^{u2} \right] \right| \\ &\leq \sum_{l \neq 1} \left(2D_m / (\sigma_l^{*2} \sqrt{d}) \right) |\mathbb{E}_{\mathcal{D}_j} [w_1^u w_l^u (\|X - \mu_l^u\|^2 - d\sigma_j^{*2} + d(\sigma_j^{*2} - \sigma_l^{u2}))]| \\ &\leq (2D_m / \sqrt{d}) \sum_{l \neq 1} \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_1^u]} \\ &\times \underbrace{\left(\sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u (\|v\|^2 - d\sigma_j^{*2})^2]} + 2\sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u \langle v, \mu_j^* - \mu_l^u \rangle^2]} + 2R_{jl}^{*2} \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u]} + d(\sigma_j^{*2} - \sigma_l^{u2}) \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u]} \right)}_{II} / \sigma_l^{*2}. \end{aligned}$$

As II will appear frequently, we state a helping lemma.

Lemma 28 *The summation of II over all $l \neq j$ is much less than \sqrt{d} . That is,*

$$\begin{aligned} \sum_{l \neq j} \left(1 / \sigma_l^{*2} \right) &\left(\sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u (\|v\|^2 - d\sigma_j^{*2})^2]} + 2\sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u \langle v, \mu_j^* - \mu_l^u \rangle^2]} \right. \\ &\left. + 2R_{jl}^{*2} \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u]} + d(\sigma_j^{*2} - \sigma_l^{u2}) \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u]} \right) \leq c\sqrt{d}, \end{aligned}$$

for some small constant c . When $l = j$, $II = O(\sqrt{d})$.

We will also prove this lemma in the last Section E.4. Thus, the entire summation is $O(\sqrt{d})$. Now we can conclude that

$$e_{j6} \leq (2D_m/\sqrt{d})\sqrt{\mathbb{E}_{\mathcal{D}_j}[w_1^u]}O(\sqrt{d}) \leq O\left((1 + \pi_1^*/\pi_j^*)D_m \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2)\right).$$

Collecting all components, we can conclude that

$$\mathbb{E}_{\mathcal{D}_j}[w_1^u] \leq D_m O\left((R_{j1}^{*2}/\sigma_1^{*2})(1 + \pi_1^*/\pi_j^*) \exp\left(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2\right)\right).$$

The summation over all other components $j \neq 1$ can be bounded thus with Lemma 26.

Computing errors from $j = 1$: Reproducing the equation (), we start from

$$\begin{aligned} e_{11} &\leq 2D_m \sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_1}[w_1^u(1 - w_1^u)\langle v + \mu_1^* - \mu_1^u, s \rangle]/\sigma_1^* \\ &\leq 2D_m \sup_{s \in \mathbb{S}^{d-1}} (\mathbb{E}_{\mathcal{D}_1}[(1 - w_1^u)\langle v, s \rangle]/\sigma_1^* + \mathbb{E}_{\mathcal{D}_1}[1 - w_1^u]D_m). \end{aligned}$$

Observe that $\mathbb{E}_{\mathcal{D}_1}[1 - w_1^u] = \sum_{l \neq 1} \mathbb{E}_{\mathcal{D}_1}[w_l^u] \leq \sum_{l \neq 1} 8(\pi_l^*/\pi_1^*) \exp(-R_{l1}^{*2}/64(\sigma_1^* \vee \sigma_l^*)^2)$. This is smaller than 1 by Lemma 26. Similarly, we can see that

$$\mathbb{E}_{\mathcal{D}_1}[(1 - w_1^u)\langle v, s \rangle] = \sum_{l \neq 1} \mathbb{E}_{\mathcal{D}_1}[w_l^u \langle v, s \rangle] \leq O\left(\sum_{l \neq 1} \sigma_1^* \sqrt{1 + \pi_l^*/\pi_1^*} \exp(-R_{l1}^{*2}/128(\sigma_1^* \vee \sigma_l^*)^2)\right).$$

This gives that $e_{11} \leq D_m \sum_{l \neq 1} O\left((1 + \pi_l^*/\pi_1^*) \exp(-R_{l1}^{*2}/128(\sigma_1^* \vee \sigma_l^*)^2)\right) \ll D_m$.

The second error term is

$$\begin{aligned} e_{12} &\leq 2D_m \sum_{l \neq 1} \sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_1}[w_l^u \langle v + \mu_1^* - \mu_l^u, s \rangle]/\sigma_l^* \\ &\leq D_m \sum_{l \neq 1} O\left((R_{l1}^*/\sigma_l^*)(1 + \pi_l^*/\pi_1^*) \exp(-R_{l1}^{*2}/64(\sigma_l^* \vee \sigma_1^*)^2)\right), \end{aligned}$$

which is guaranteed to be $e_{12} \ll D_m$.

The third and fourth terms are smaller than $2D_m \sum_{l \neq 1} \mathbb{E}_{\mathcal{D}_1}[w_l^u] \ll D_m$. Now we need to deal with fifth and sixth terms, which again require some algebraic manipulation. We can start from ()...

$$e_{15} \leq (2D_m/\sigma_1^{*2}\sqrt{d})\mathbb{E}_{\mathcal{D}_1}\left[(1 - w_1^u)\left(\|v\|^2 - d\sigma_1^{*2}\right) + 2|\langle v, \mu_1^* - \mu_1^u \rangle| + D_m^2 + d(\sigma_1^{*2} - \sigma_1^{u2})\right].$$

For each item, we can say that

$$\begin{aligned} \mathbb{E}_{\mathcal{D}_1}[(1 - w_1^u)(\|v\|^2 - d\sigma_1^{*2})] &\leq \sqrt{\mathbb{E}_{\mathcal{D}_1}[1 - w_1^u]} \sqrt{\mathbb{E}_{\mathcal{D}_1}[(\|v\|^2 - d\sigma_1^{*2})^2]} \\ &\leq O\left(\sqrt{\sum_{l \neq 1} (1 + \pi_l^*/\pi_1^*) \exp(-R_{l1}^{*2}/64(\sigma_1^* \vee \sigma_l^*)^2)} \sqrt{2d\sigma_1^{*4}}\right), \end{aligned}$$

and for two other terms,

$$\mathbb{E}_{\mathcal{D}_1}[(1 - w_1^u)\langle v, \mu_1^* - \mu_1^u \rangle] \leq D_m \sigma_1^* \sqrt{\mathbb{E}_{\mathcal{D}_1}[1 - w_1^u]} \sqrt{\mathbb{E}_{\mathcal{D}_1}[\langle v, s \rangle^2]}$$

$$\leq D_m \sigma_1^{*2} O \left(\sqrt{\sum_{l \neq 1} (1 + \pi_l^*/\pi_1^*) \exp(-R_{l1}^{*2}/64(\sigma_1^* \vee \sigma_l^*)^2)} \right),$$

and

$$\mathbb{E}_{\mathcal{D}_1} [(1 - w_1^u) d(\sigma_1^{*2} - \sigma_1^{u2})] \leq (\sigma_1^{*2} D_m \sqrt{d}) O \left(\sum_{l \neq 1} (1 + \pi_l^*/\pi_1^*) \exp(-R_{l1}^{*2}/64(\sigma_1^* \vee \sigma_l^*)^2) \right).$$

Thus $e_{15} \leq 2D_m/(\sigma_1^{*2} \sqrt{d}) \cdot c\sqrt{d}\sigma_1^{*2} \leq c'D_m$ for small constants c, c' .

Finally, the sixth term can be similarly bounded as

$$\begin{aligned} e_{16} &\leq \sum_{l \neq 1} (2D_m/\sigma_l^{*2} \sqrt{d}) |\mathbb{E}_{\mathcal{D}_1} [w_1^u w_l^u (\|X - \mu_l^u\|^2 - d\sigma_l^{*2})]| \\ &\leq \sum_{l \neq 1} (2D_m/\sigma_l^{*2} \sqrt{d}) \left| \mathbb{E}_{\mathcal{D}_1} \left[w_l^u \left((\|v\|^2 - d\sigma_1^{*2}) + 2\langle v, \mu_1^* - \mu_l^u \rangle + \|\mu_1^* - \mu_l^u\|^2 + d(\sigma_1^{*2} - \sigma_l^{*2}) \right) \right] \right| \\ &\leq D_m \sum_{l \neq 1} (2/\sigma_l^{*2} \sqrt{d}) \left(|\mathbb{E}_{\mathcal{D}_1} [w_l^u (\|v\|^2 - d\sigma_1^{*2})]| \right. \\ &\quad \left. + 2D_m \sigma_1^* |\mathbb{E}_{\mathcal{D}_1} [w_l^u \langle v, s \rangle]| + \mathbb{E}_{\mathcal{D}_1} [w_l^u] D_m^2 \sigma_1^{*2} + \sqrt{d} D_m \sigma_1^{*2} \mathbb{E}_{\mathcal{D}_1} [w_l^u] \right) \\ &\leq D_m \sum_{l \neq 1} (2/\sigma_l^{*2} \sqrt{d}) (5 + 3\pi_l^*/\pi_1^*) \exp(-R_{l1}^{*2}/128(\sigma_l^* \vee \sigma_1^*)^2) \left(\sqrt{2d}\sigma_1^{*2} + 2D_m \sigma_1^{*2} + D_m^2 \sigma_1^{*2} + \sqrt{d} D_m \sigma_1^{*2} \right) \\ &\leq 20D_m/\pi_1^* \sum_{l \neq 1} (\sigma_1^{*2}/\sigma_l^{*2}) (\pi_1^* + \pi_l^*) \exp(-R_{l1}^{*2}/128(\sigma_l^* \vee \sigma_1^*)^2) \leq cD_m \end{aligned}$$

for small constant c . Collecting all components, we can conclude that $\mathbb{E}_{\mathcal{D}_1} [\Delta_w^u] \leq c'D_m$.

Errors from all components: Collecting the errors from other components and own components, now we can conclude that $\mathbb{E}_{\mathcal{D}} [\Delta_w^u] = \sum_j \pi_j^* \mathbb{E}_{\mathcal{D}_j} [\Delta_w^u] \leq c_\mu \pi_1^* D_m$ for some small $c_\mu < 1$.

E.2. Convergence of Means

Computing errors from $j \neq 1$: Let us proceed in a very similar way we did for mixing weights. Let us first handle $-w_1^u(1 - w_1^u)(X - \mu_1^u)^T(\mu_1 - \mu_1^*)/\sigma_1^2$.

$$\begin{aligned} e_{j1} &= \|\mathbb{E}_{\mathcal{D}_j} [w_1^u(1 - w_1^u)(X - \mu_1^u)^T(\mu_1 - \mu_1^*)(X - \mu_1^*)]/\sigma_1^{u2}\| \\ &\leq 2\|\mathbb{E}_{\mathcal{D}_j} [w_1^u(1 - w_1^u)(X - \mu_1^*)(X - \mu_1^u)^T]/\sigma_1^* \|_{op} \|(\mu_1 - \mu_1^*)/\sigma_1^*\| \\ &\leq 2D_m \sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_j} [w_1^u(1 - w_1^u) \langle X - \mu_1^*, s \rangle \langle X - \mu_1^u, s \rangle] / \sigma_1^*. \end{aligned}$$

Therefore, it is enough to show that for any fixed unit vector s ,

$$|\mathbb{E}_{\mathcal{D}_j} [w_1^u(1 - w_1^u) \langle X - \mu_1^*, s \rangle \langle X - \mu_1^u, s \rangle]|,$$

is exponentially small. We state one more helper lemma that bounds

Lemma 29 For $j \neq 1$,

$$\mathbb{E}_{\mathcal{D}_j} [w_1^u \langle v, s \rangle^2] \leq 5(1 + \pi_1^*/\pi_j^*) R_{j1}^{*2} \exp\left(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2\right). \quad (28)$$

Proof We can reproduce the proof of Corollary 5. Let $\beta = R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2$. Then,

$$\begin{aligned} |\mathbb{E}_{\mathcal{D}_j} [w_1 \langle v, s \rangle^2]| &= |\mathbb{E}_{\mathcal{D}_j} [w_1 \langle v, s \rangle^2 \mathbb{1}_{\mathcal{E}_{j,good}}]| + \left| \mathbb{E}_{\mathcal{D}_j} [w_1 \langle v, s \rangle^2 \mathbb{1}_{\mathcal{E}_{j,good}^c}] \right| \\ &\leq 3(\pi_1^*/\pi_j^*) \exp(-\beta) \mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle|^2] + \mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle|^2 | \mathcal{E}_{j,1}^c] P(\mathcal{E}_{j,1}^c) \\ &\quad + \mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle|^2 | \mathcal{E}_{j,2}^c] P(\mathcal{E}_{j,2}^c) + \mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle|^2 | \mathcal{E}_{j,3}^c] P(\mathcal{E}_{j,3}^c). \end{aligned}$$

$\mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle|^2 | \mathcal{E}_{j,1}^c]$ can be bounded with Lemma 17, with $p = 2$ and $\alpha = R_{j1}^*/5\sigma_j^*$.

$$\mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle|^2 | \langle v, R_{j1}^* \rangle \geq R_{j1}^{*2}/5] \leq \sigma_j^* \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [|\langle v, s \rangle|^2 | |\langle v, u \rangle| \geq \alpha] \leq \sigma_j^{*2} (4\alpha^2 + 4) \leq R_{j1}^{*2}.$$

Similarly, we can bound $\mathbb{E}_{\mathcal{D}_j} [|\langle v, s \rangle|^2 | \mathcal{E}_{j,2}^c] P(\mathcal{E}_{j,2}^c) \leq 2R_{j1}^*$ using the same Lemma 17 with $p = 2$ and $\alpha = R_{j1}^*/4\sigma_j^*$. For the third term, we use Lemma 20, with $p = 2$ and $\alpha = R_{j1}^{*2}/64(\sigma_1^* \vee \sigma_j^*)^2 = \beta$. Then,

$$\begin{aligned} \sigma_j^* \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [|\langle v, s \rangle|^2 \|v\|^2 \geq d + 2\sqrt{\alpha d} + 2\alpha] &\leq \sigma_j^{*2} ((64\alpha) + 4 \exp(-\alpha/2)(8\alpha + 2)) \leq 2R_{j1}^{*2}, \\ \sigma_j^{*2} \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [|\langle v, s \rangle|^2 \|v\|^2 \leq d - 2\sqrt{\alpha d}] &\leq \sigma_j^{*2} \mathbb{E}_{v \sim \mathcal{N}(0, I_d)} [|\langle v, s \rangle|^2] \leq \sigma_j^{*2}, \end{aligned}$$

Collecting these three components, we can conclude that

$$|\mathbb{E}_{\mathcal{D}_j} [w_1 \langle v, s \rangle^2]| \leq (3(\pi_1^*/\pi_j^*) \sigma_j^{*2} + 5R_{j1}^{*2}) \exp(-\beta).$$

This yields the equation (28). ■

Then we can proceed as

$$\begin{aligned} |\mathbb{E}_{\mathcal{D}_j} [w_1^u (1 - w_1^u) \langle X - \mu_1^*, s \rangle \langle X - \mu_1^u, s \rangle]| &\leq \mathbb{E}_{\mathcal{D}_j} \left[w_1^u \left(\langle v, s \rangle^2 + 2|\langle v, s \rangle| R_{j1}^* + 2R_{j1}^{*2} \right) \right] \\ &\leq O \left((1 + \pi_1^*/\pi_j^*) R_{j1}^{*2} \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2) \right), \end{aligned}$$

which yields $e_{j1} \leq D_m \sigma_1^* O \left((1 + \pi_1^*/\pi_j^*) (R_{j1}^*/\sigma_1^*)^2 \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2) \right)$.

Similarly, we bound the second term.

$$\begin{aligned} e_{j2} &\leq 2D_m \left| \mathbb{E}_{\mathcal{D}_j} \left[\sum_{l \neq 1} w_1^u w_l^u \langle X - \mu_1^*, s \rangle \langle X - \mu_l^u, s \rangle / \sigma_l^* \right] \right| \\ &\leq 2D_m \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_1^u \langle X - \mu_1^*, s \rangle^2]} \sum_{l \neq 1} \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u \langle X - \mu_l^u, s \rangle^2] / \sigma_l^{*2}}. \end{aligned}$$

Using the Lemma 27 as in mixing weights, we get

$$e_{j2} \leq D_m \sigma_1^* O \left((1 + \pi_1^*/\pi_j^*) (R_{j1}^*/\sigma_1^*) \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2) \right).$$

Third term and fourth term are straight-forward to bound.

$$e_{j3} \leq 2D_m |\mathbb{E}_{\mathcal{D}_j} [w_1^u \langle X - \mu_1^*, s \rangle]| \leq D_m O \left(R_{j1}^* \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2) \right),$$

$$e_{j4} \leq 2D_m \left| \mathbb{E}_{\mathcal{D}_j} \left[\sum_{l \neq 1} w_1^u w_l^u \langle X - \mu_1^*, s \rangle \right] \right| \leq 2D_m |\mathbb{E}_{\mathcal{D}_j} [w_1^u \langle X - \mu_1^*, s \rangle]|,$$

which is again smaller than $D_m O \left(R_{j1}^* \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2) \right)$.

The challenging fifth and sixth term is also bounded using similar algebra.

$$\begin{aligned} e_{j5} &= \left| \mathbb{E}_{\mathcal{D}_j} [w_1^u (1 - w_1^u) \langle X - \mu_1^*, s \rangle (\|X - \mu_1^u\|^2 - d\sigma_1^{u2}) / (2\sigma_1^{u2})] (\sigma_1^2 - \sigma_1^{*2}) / \sigma_1^{u2} \right| \\ &\leq \left(2D_m / \sqrt{d} \right) |\mathbb{E}_{\mathcal{D}_j} [w_1^u \langle X - \mu_1^*, s \rangle (\|X - \mu_1^u\|^2 - d\sigma_1^{u2}) / (2\sigma_1^{u2})]| \\ &\leq \left(2D_m / \sigma_1^{*2} \sqrt{d} \right) \left| \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_1^u \langle X - \mu_1^*, s \rangle^2]} \left(\sqrt{\mathbb{E}_{\mathcal{D}_j} [\|v\|^2 - d\sigma_j^{*2}]} + 2\sqrt{\mathbb{E}_{\mathcal{D}_j} [\langle v, \mu_j^* - \mu_1^u \rangle^2]} \right) \right. \\ &\quad \left. + 2R_{j1}^{*2} |\mathbb{E}_{\mathcal{D}_j} [w_1^u \langle X - \mu_1^*, s \rangle]| + d(\sigma_j^{*2} - \sigma_1^{u2}) \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_1^u]} \sqrt{\mathbb{E}_{\mathcal{D}_j} [\langle X - \mu_1^*, s \rangle^2]} \right| \\ &\leq \left(2D_m / \sigma_1^{*2} \sqrt{d} \right) (1 + \pi_1^* / \pi_j^*) O \left(\sqrt{R_{j1}^{*2} \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2)} (\sqrt{2d} + 4R_{j1}^* \sigma_j^*) \right. \\ &\quad \left. + 2R_{j1}^{*3} \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2) + \sqrt{d} R_{j1}^{*2} (\sigma_j^* \vee \sigma_1^*) \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2) \right) \\ &\leq D_m \sigma_1^* O \left((1 + \pi_1^* / \pi_j^*) (R_{j1}^* / \sigma_1^*)^3 \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2) \right). \end{aligned}$$

Finally,

$$\begin{aligned} e_{j6} &= \left| \mathbb{E}_{\mathcal{D}_j} \left[\sum_{l \neq 1} w_1^u w_l^u \langle X - \mu_1^*, s \rangle (\|X - \mu_l^u\|^2 - d\sigma_l^{u2}) / (2\sigma_l^{u2}) \right] (\sigma_l^2 - \sigma_l^{*2}) / \sigma_l^{u2} \right| \\ &\leq \left(2D_m / \sqrt{d} \right) \sum_{l \neq 1} (1 / \sigma_l^{*2}) \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_1^u \langle X - \mu_1^*, s \rangle^2]} \\ &\quad \times \left(\sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u (\|v\|^2 - d\sigma_j^{*2})]} + 2\sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u \langle v, \mu_j^* - \mu_l^u \rangle^2]} + 2R_{jl}^{*2} \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u]} + d(\sigma_j^{*2} - \sigma_l^{u2}) \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u]} \right). \end{aligned}$$

Using Lemma 28, the summation is less than $c\sqrt{d}$ and hence

$$e_{j6} \leq D_m \sigma_1^* O \left((1 + \pi_1^* / \pi_j^*) (R_{j1}^* / \sigma_1^*) \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2) \right).$$

So all guarantee that

$$\sum_{j \neq 1} \pi_j^* \|\mathbb{E}_{\mathcal{D}_j} [\Delta_w^u (X - \mu_1^*)]\| \leq D_m \sigma_1^* \sum_{j \neq 1} O \left((\pi_1^* + \pi_j^*) (R_{j1}^* / \sigma_1^*)^3 \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2) \right). \quad (29)$$

Errors from $j = 1$: We repeat the process of bounding six terms as always.

$$\begin{aligned} e_{11} &\leq 2D_m \sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_1} [w_1^u (1 - w_1^u) \langle X - \mu_1^*, s \rangle \langle X - \mu_1^u, s \rangle] / \sigma_1^* \\ &\leq 2D_m \sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_1} [(1 - w_1^u) (\langle v, s \rangle^2 + \langle v, s \rangle D_m \sigma_1^*)] / \sigma_1^*. \end{aligned}$$

Then,

$$\begin{aligned} \mathbb{E}_{\mathcal{D}_1} [(1 - w_1^u) \langle v, s \rangle^2] &= \sum_{l \neq 1} \mathbb{E}_{\mathcal{D}_1} [w_l^u \langle v, s \rangle^2] \leq \sum_{l \neq 1} O \left((1 + \pi_l^* / \pi_1^*) R_{l1}^{*2} \exp(-R_{l1}^{*2} / 64(\sigma_l^* \vee \sigma_1^*)^2) \right) \\ &\leq (1 / \pi_1^*) \sum_{l \neq 1} (\pi_1^* + \pi_l^*) R_{l1}^{*2} \exp(-R_{l1}^{*2} / 64(\sigma_l^* \vee \sigma_1^*)^2) \leq c \sigma_1^{*2}, \end{aligned}$$

for some small constant c with the Lemma 26.

The second term will be similarly,

$$\begin{aligned} e_{12} &\leq 2D_m \mathbb{E}_{\mathcal{D}_1} \left[\sum_{l \neq 1} |w_1^u w_l^u \langle v, s \rangle \langle X - \mu_l^u, s \rangle| / \sigma_l^* \right] \\ &\leq 2D_m \sigma_1^* \sum_{l \neq 1} \sqrt{\mathbb{E}_{\mathcal{D}_1} [w_l^u \langle X - \mu_l^u, s \rangle^2] / \sigma_l^{*2}} \leq c D_m \sigma_1^*, \end{aligned}$$

for some small constant c .

Third and fourth term is easy to handle,

$$e_{13}, e_{14} \leq D_m \sum_{l \neq 1} O \left(R_{l1}^* \exp(-R_{l1}^{*2} / 64(\sigma_l^* \vee \sigma_1^*)^2) \right),$$

which is again much less than D_m .

For the fifth term, we again start from

$$\begin{aligned} e_{15} &= \left| \mathbb{E}_{\mathcal{D}_1} [w_1^u (1 - w_1^u) \langle X - \mu_1^*, s \rangle (\|X - \mu_1^u\|^2 - d\sigma_1^{u2}) / (2\sigma_1^{u2})] (\sigma_1^2 - \sigma_1^{*2}) / \sigma_1^{u2} \right| \\ &\leq \left(2D_m / \sqrt{d} \right) \left| \mathbb{E}_{\mathcal{D}_1} [(1 - w_1^u) \langle v, s \rangle (\|v + \mu_1^* - \mu_1^u\|^2 - d\sigma_1^{u2}) / (2\sigma_1^{u2})] \right| \\ &\leq \left(2D_m / \sigma_1^{*2} \sqrt{d} \right) \left| \sqrt{\mathbb{E}_{\mathcal{D}_1} [(1 - w_1^u) \langle v, s \rangle^2]} \left(\sqrt{\mathbb{E}_{\mathcal{D}_1} [\|v\|^2 - d\sigma_1^{*2}]^2} + 2\sqrt{\mathbb{E}_{\mathcal{D}_1} [\langle v, \mu_1^* - \mu_1^u \rangle^2]} \right) \right. \\ &\quad \left. + 2D_m^2 \sigma_1^{*2} |\mathbb{E}_{\mathcal{D}_1} [(1 - w_1^u) \langle v, s \rangle]| + D_m \sqrt{d} \sigma_1^{*2} |\mathbb{E}_{\mathcal{D}_1} [(1 - w_1^u) \langle v, s \rangle]| \right|, \end{aligned}$$

The first term in the above can be bounded with Lemma 29,

$$\begin{aligned} \mathbb{E}_{\mathcal{D}_1} [(1 - w_1^u) \langle v, s \rangle^2] &= \sum_{l \neq 1} \mathbb{E}_{\mathcal{D}_1} [w_l^u \langle v, s \rangle^2] \leq \sigma_1^{*2} O \left(\sum_{l \neq 1} (1 + \pi_l^* / \pi_1^*) R_{l1}^{*2} \exp(-R_{l1}^{*2} / 64(\sigma_l^* \vee \sigma_1^*)^2) \right) \\ &\leq c \sigma_1^{*2}. \end{aligned}$$

and

$$\mathbb{E}_{\mathcal{D}_1}[(\|v\|^2 - d\sigma_1^{*2})^2] = 2d\sigma_1^{*4}, \quad \mathbb{E}_{\mathcal{D}_1}[\langle v, \mu_1^* - \mu_1^u \rangle^2] \leq D_m^2 \sigma_1^{*4}.$$

Similarly, we have that $\mathbb{E}_{\mathcal{D}_1}[(1 - w_1^u)\langle v, s \rangle] \leq c\sigma_1^*$. Collecting all components, we can bound this fifth term $e_{15} \ll D_m \sigma_1^*$.

The bound for the final term follows similarly.

$$\begin{aligned} e_{16} &= \left| \mathbb{E}_{\mathcal{D}_1} \left[\sum_{l \neq 1} w_1^u w_l^u \langle X - \mu_1^*, s \rangle (\|X - \mu_l^u\|^2 - d\sigma_l^{u2}) / (2\sigma_l^{u2}) \right] (\sigma_l^2 - \sigma_l^{*2}) / \sigma_l^{u2} \right| \\ &\leq \left(2D_m / \sqrt{d} \right) \sum_{l \neq 1} (1/\sigma_l^{*2}) \sqrt{\mathbb{E}_{\mathcal{D}_1}[w_l^u \langle v, s \rangle^2]} \\ &\quad \times \left(\sqrt{\mathbb{E}_{\mathcal{D}_1}[w_l^u (\|v\|^2 - d\sigma_j^{*2})^2]} + 2\sqrt{\mathbb{E}_{\mathcal{D}_1}[w_l^u \langle v, \mu_1^* - \mu_l^u \rangle^2]} + 2R_{jl}^{*2} \sqrt{\mathbb{E}_{\mathcal{D}_1}[w_l^u]} + d(\sigma_1^{*2} - \sigma_l^{u2}) \sqrt{\mathbb{E}_{\mathcal{D}_1}[w_l^u]} \right). \end{aligned}$$

We can again use Lemma 28 and get $e_{j6} \leq cD_m \sigma_1^*$ for some small constant c .

Finally, we collect all error terms to conclude that $\mathbb{E}_{\mathcal{D}_1}[\Delta_w^u(X - \mu_1^*)] \leq cD_m \sigma_1^*$ for some small constant c . Now with the equation (29),

$$\begin{aligned} \|\mu_1^+ - \mu_1^*\| &= \left(\sum_j \pi_j^* \mathbb{E}_{\mathcal{D}_j}[\Delta_w^u(X - \mu_1^*)] \right) / \pi_1^+ \\ &\leq D_m \sigma_1^* O \left(\sum_{j \neq 1} (\pi_1^* + \pi_j^*) (R_{j1}^* / \sigma_1^*)^3 \exp(-R_{j1}^{*2} / 128(\sigma_j^* \vee \sigma_1^*)^2) \right) / \pi_1^+ \\ &\leq cD_m \sigma_1^* \pi_1^* / \pi_1^+ \leq c' D_m \sigma_1^*, \end{aligned}$$

for some small constant c' , where we used Lemma 26 with $q = 3$.

E.3. Convergence of the Variance

The most challenging part is again to show the convergence of variance estimators. We start with each term by term as other quantities.

Errors from other components $j \neq 1$: Let us start from the first error term.

$$\begin{aligned} e_{j1} &= |\mathbb{E}_{\mathcal{D}_j}[w_1^u (1 - w_1^u) (X - \mu_1^u)^T (\mu_1 - \mu_1^*) (\|X - \mu_1^*\|^2 - d\sigma_1^{*2})] / \sigma_1^{u2}| \\ &\leq 2 \left\| \mathbb{E}_{\mathcal{D}_j}[w_1^u (\|X - \mu_1^*\|^2 - d\sigma_1^{*2}) (X - \mu_1^*)] / \sigma_1^* \right\| \|(\mu_1 - \mu_1^*) / \sigma_1^*\| \\ &\leq 2D_m \sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_j}[w_1^u (\|X - \mu_1^*\|^2 - d\sigma_1^{*2}) \langle X - \mu_1^u, s \rangle] / \sigma_1^* \\ &\leq 2D_m \sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_j} \left[w_1^u \left((\|v\|^2 - d\sigma_j^{*2}) + 2\langle v, \mu_j^* - \mu_1^* \rangle + \|\mu_j^* - \mu_1^*\|^2 + d(\sigma_j^{*2} - \sigma_1^{*2}) \right) \langle X - \mu_1^u, s \rangle \right] / \sigma_1^*. \end{aligned}$$

For each item,

$$\mathbb{E}_{\mathcal{D}_j}[w_1^u \langle X - \mu_1^u, s \rangle (\|v\|^2 - d\sigma_j^{*2})] \leq \sqrt{\mathbb{E}_{\mathcal{D}_j}[w_1^u \langle X - \mu_1^u, s \rangle^2]} \sqrt{\mathbb{E}_{\mathcal{D}_j}[(\|v\|^2 - d\sigma_j^{*2})^2]}$$

$$\begin{aligned} &\leq O\left(\sqrt{(1 + \pi_1^*/\pi^*)R_{j1}^{*2} \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2)}\sigma_j^{*2}\sqrt{2d}\right) \\ &\leq O\left(\sqrt{d}\sigma_j^{*2}R_{j1}^*(1 + \pi_1^*/\pi_j^*) \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2)\right). \end{aligned}$$

$$\begin{aligned} 2R_{j1}^*\mathbb{E}_{\mathcal{D}_j}[w_1^u\langle X - \mu_1^u, s \rangle\langle v, s' \rangle] &\leq 2R_{j1}^*\sqrt{\mathbb{E}_{\mathcal{D}_j}[w_1^u\langle X - \mu_1^u, s \rangle^2]}\sqrt{\mathbb{E}_{\mathcal{D}_j}[w_1^u\langle v, s' \rangle^2]} \\ &\leq O\left((1 + \pi_1^*/\pi_j^*)R_{j1}^{*2}\sigma_j^* \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2)\right), \end{aligned}$$

and the third item is

$$R_{j1}^{*2}\mathbb{E}_{\mathcal{D}_j}[w_1^u\langle X - \mu_1^u, s \rangle] \leq O\left((1 + \pi_1^*/\pi^*)R_{j1}^{*3} \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2)\right).$$

The final item, we again need Lemma 22 to get

$$\begin{aligned} d(\sigma_j^{*2} - \sigma_1^{*2})\mathbb{E}_{\mathcal{D}_j}[w_1^u\langle X - \mu_1^u, s \rangle] &\leq d(\sigma_j^{*2} - \sigma_1^{*2})\sqrt{\mathbb{E}_{\mathcal{D}_j}[w_1^u]}\sqrt{\mathbb{E}_{\mathcal{D}_j}[w_1^u\langle X - \mu_1^u, s \rangle^2]} \\ &\leq O\left((1 + \pi_1^*/\pi_j^*)R_{j1}^{*2}(\sigma_j^* \vee \sigma_1^*)\sqrt{d} \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2)\right). \end{aligned}$$

Hence, the first error term is bounded as

$$e_{j1} \leq \sigma_1^{*2}O\left((1 + \pi_1^*/\pi_j^*)\sqrt{d}(R_{j1}^*/\sigma_1^*)^3 \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2)\right).$$

The error bound for second error term starts with arranging equations as usual,

$$\begin{aligned} e_{j2} &= \sum_{l \neq 1} |\mathbb{E}_{\mathcal{D}_j}[w_1^u w_l^u (X - \mu_l^u)^T (\mu_l - \mu_l^*) (\|X - \mu_1^*\|^2 - d\sigma_1^{*2})]| / \sigma_l^{*2} \\ &\leq 2D_m \sum_{l \neq 1} \sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_j} \left[w_1^u w_l^u (\|X - \mu_1^*\|^2 - d\sigma_1^{*2}) \langle X - \mu_l^u, s \rangle \right] / \sigma_l^* \\ &\leq 2D_m \sqrt{\mathbb{E}_{\mathcal{D}_j} \left[w_1^u (\|X - \mu_1^*\|^2 - d\sigma_1^{*2})^2 \right]} \sum_{l \neq 1} \sqrt{\mathbb{E}_{\mathcal{D}_j} \left[w_l^u \langle X - \mu_l^u, s \rangle^2 / \sigma_l^{*2} \right]}. \end{aligned}$$

For the first square root, useful inequality is $(a + b + c + d)^2 \leq 4(a^2 + b^2 + c^2 + d^2)$ for any real a, b, c, d . Using this,

$$\begin{aligned} &\mathbb{E}_{\mathcal{D}_j} \left[w_1^u \left(\|X - \mu_1^*\|^2 - d\sigma_1^{*2} \right)^2 \right] \\ &\leq 4\mathbb{E}_{\mathcal{D}_j} \left[w_1^u \left((\|v\|^2 - d\sigma_j^{*2})^2 + 4\langle v, \mu_j^* - \mu_1^* \rangle^2 + \|\mu_j^* - \mu_1^*\|^4 + d^2(\sigma_j^{*2} - \sigma_1^{*2})^2 \right) \right] \\ &\leq 4\sqrt{\mathbb{E}_{\mathcal{D}_j}[w_1^u]}\sqrt{\mathbb{E}_{\mathcal{D}_j}[(\|v\|^2 - d\sigma_j^{*2})^4]} \\ &\quad + 16R_{j1}^{*2}\mathbb{E}_{\mathcal{D}_j}[w_1^u\langle v, s' \rangle^2] + 4R_{j1}^{*4}\mathbb{E}_{\mathcal{D}_j}[w_1^u] + 4d^2(\sigma_j^{*2} - \sigma_1^{*2})^2\mathbb{E}_{\mathcal{D}_j}[w_1^u] \\ &\leq (1 + \pi_1^*/\pi_j^*) \left(O\left(d\sigma_j^{*4}\right) \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2) + O\left(dR_{j1}^{*4}\right) \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2) \right), \end{aligned} \tag{30}$$

where we used Lemma 22. Meanwhile, the summation over $l \neq 1$ in right hand side is $O(1)$ by Lemma 27. Therefore,

$$e_{j2} \leq \sigma_1^{*2} O\left(\sqrt{d}(1 + \pi_1^*/\pi_j^*)(R_{j1}^*/\sigma_1^*)^2 \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2)\right),$$

As we can imagine, e_{j3} and e_{j4} can be shown to be bounded using the same procedure for e_{j1} and e_{j2} .

Now we jump to e_{j5} and e_{j6} .

$$\begin{aligned} e_{j5} &= \left| \mathbb{E}_{\mathcal{D}_j} \left[w_1^u (1 - w_1^u) (\|X - \mu_1^*\|^2 - d\sigma_1^{*2}) (\|X - \mu_1^u\|^2 - d\sigma_1^{u2}) / (2\sigma_1^{u2}) \right] (\sigma_1^2 - \sigma_1^{*2}) / \sigma_1^{u2} \right| \\ &\leq \left(2D_m / \sigma_1^{*2} \sqrt{d} \right) \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_1^u (\|X - \mu_1^*\|^2 - d\sigma_1^{*2})^2]} \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_1^u (\|X - \mu_1^u\|^2 - d\sigma_1^{u2})^2]}. \end{aligned}$$

The rest of the procedure is repetition of (30) (we get a same order-wise bound for both square roots). Therefore, we get

$$e_{j5} \leq D_m \sigma_1^{*2} O\left(\left(1 + \pi_1^*/\pi_j^*\right) \sqrt{d} (R_{j1}^*/\sigma_1^*)^4 \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2)\right).$$

Finally, with the similar strategy, we have

$$\begin{aligned} e_{j6} &= \left| \sum_{l \neq 1} \mathbb{E}_{\mathcal{D}_j} \left[w_1^u w_l^u (\|X - \mu_1^*\|^2 - d\sigma_1^{*2}) (\|X - \mu_l^u\|^2 - d\sigma_l^{u2}) / (2\sigma_l^{u2}) \right] (\sigma_l^2 - \sigma_l^{*2}) / \sigma_l^{u2} \right| \\ &\leq \left(2D_m / \sqrt{d} \right) \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_1^u (\|X - \mu_1^*\|^2 - d\sigma_1^{*2})^2]} \sum_{l \neq 1} (1/\sigma_l^{*2}) \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u (\|X - \mu_l^u\|^2 - d\sigma_l^{u2})^2]}, \end{aligned}$$

The summation over $l \neq 1$ is $O(1)$ as shown in Lemma 28. Therefore,

$$e_{j6} \leq D_m \sigma_1^{*2} O\left(\left(1 + \pi_1^*/\pi_j^*\right) \sqrt{d} (R_{j1}^*/\sigma_1^*)^2 \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2)\right).$$

Collecting all error terms, every term is less than

$$\sigma_1^{*2} (1 + \pi_1^*/\pi_j^*) \sqrt{d} O\left(\left(R_{j1}^*/\sigma_1^*\right)^3 \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2) + \left(R_{j1}^*/\sigma_1^*\right)^4 \exp(-R_{j1}^{*2}/64(\sigma_j^* \vee \sigma_1^*)^2)\right).$$

By Lemma 26, the entire summation is smaller than $c\sigma_1^{*2}\pi_1^*\sqrt{d}$ for some small constant c . Recall that \sqrt{d} will be divided by d in the end.

Errors from own component $j = 1$: We will walk through the same procedure. The first term is:

$$\begin{aligned} e_{11} &= |\mathbb{E}_{\mathcal{D}_1} [w_1^u (1 - w_1^u) (X - \mu_1^u)^T (\mu_1 - \mu_1^*) (\|v\|^2 - d\sigma_1^{*2})] / \sigma_1^{u2}| \\ &\leq 2D_m \sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_1} [(1 - w_1^u) (\|v\|^2 - d\sigma_1^{*2}) \langle X - \mu_1^u, s \rangle] / \sigma_1^* \\ &\leq 2D_m \sup_{s \in \mathbb{S}^{d-1}} \sum_{l \neq 1} \mathbb{E}_{\mathcal{D}_1} [w_l^u (\|v\|^2 - d\sigma_1^{*2}) \langle X - \mu_1^u, s \rangle] / \sigma_1^* \\ &\leq 2D_m \sqrt{\mathbb{E}_{\mathcal{D}_1} [(\|v\|^2 - d\sigma_1^{*2})^2]} / \sigma_1^* \sup_{s \in \mathbb{S}^{d-1}} \sum_{l \neq 1} \sqrt{\mathbb{E}_{\mathcal{D}_1} [w_l^u \langle X - \mu_1^u, s \rangle^2]} \end{aligned}$$

$$\leq D_m \sqrt{d} \sigma_1^* O \left(\sum_{l \neq 1} (1 + \pi_l^*/\pi_1^*) \sigma_1^* \exp(-R_{l1}^{*2}/128(\sigma_l^* \vee \sigma_1^*)^2) \right),$$

which is smaller than $cD_m \sqrt{d} \sigma_1^*$ for some small c by Lemma 26.

Similarly, the second term can be bounded as

$$\begin{aligned} e_{12} &= \left| \mathbb{E}_{\mathcal{D}_1} \left[\sum_{l \neq 1} w_1^u w_l^u (X - \mu_l^u)^T (\mu_l - \mu_l^*) (\|v\|^2 - d\sigma_1^{*2}) \right] / \sigma_l^{*2} \right| \\ &\leq 2D_m \sum_{l \neq 1} \sup_{s \in \mathbb{S}^{d-1}} \mathbb{E}_{\mathcal{D}_1} [w_l^u (\|v\|^2 - d\sigma_1^{*2}) \langle X - \mu_l^u, s \rangle] / \sigma_l^* \\ &\leq 2D_m \sup_{s \in \mathbb{S}^{d-1}} \sum_{l \neq 1} \mathbb{E}_{\mathcal{D}_1} [w_l^u (\|v\|^2 - d\sigma_1^{*2}) \langle X - \mu_l^u, s \rangle] / \sigma_l^* \\ &\leq 2D_m \sqrt{\mathbb{E}_{\mathcal{D}_1} [(\|v\|^2 - d\sigma_1^{*2})^2]} \sup_{s \in \mathbb{S}^{d-1}} \sum_{l \neq 1} \sqrt{\mathbb{E}_{\mathcal{D}_1} [w_l^u \langle X - \mu_l^u, s \rangle^2] / \sigma_l^{*2}} \\ &\leq cD_m \sqrt{d} \sigma_1^{*2}, \end{aligned}$$

for small constant c , where in the last step we used Lemma 27.

e_{13} and e_{14} can be bounded similarly. Finally, e_{15} can be bounded as

$$\begin{aligned} e_{15} &= \left| \mathbb{E}_{\mathcal{D}_1} \left[w_1^u (1 - w_1^u) (\|v\|^2 - d\sigma_1^{*2}) (\|X - \mu_1^u\|^2 - d\sigma_1^{u2}) / (2\sigma_1^{u2}) \right] (\sigma_1^2 - \sigma_1^{*2}) / \sigma_1^{u2} \right| \\ &\leq \left(2D_m / \sigma_1^{*2} \sqrt{d} \right) \sum_{l \neq 1} \mathbb{E}_{\mathcal{D}_1} \left[(\|v\|^2 - d\sigma_1^{*2}) w_l^u (\|v + (\mu_1^* - \mu_1^u)\|^2 - d\sigma_1^{*2} + d(\sigma_1^{*2} - \sigma_1^{u2}))^2 \right] \\ &\leq \left(2D_m / \sigma_1^{*2} \sqrt{d} \right) \sum_{l \neq 1} \mathbb{E}_{\mathcal{D}_1} \left[w_l^u (\|v\|^2 - d\sigma_1^{*2}) \left((\|v\|^2 - d\sigma_1^{*2}) + 2D_m \sigma_1^* \langle v, s \rangle + D_m^2 \sigma_1^{*2} + D_m \sigma_1^{*2} \sqrt{d} \right) \right] \\ &\leq \left(2D_m / \sigma_1^{*2} \sqrt{d} \right) \sum_{l \neq 1} \sqrt{\mathbb{E}_{\mathcal{D}_1} [w_l^u]} \sqrt{\mathbb{E}_{\mathcal{D}_1} [(\|v\|^2 - d\sigma_1^{*2})^4]} + 2D_m \sigma_1^* \sqrt{\mathbb{E}_{\mathcal{D}_1} [w_l^u \langle v, s \rangle^2]} \sqrt{\mathbb{E}_{\mathcal{D}_1} [(\|v\|^2 - d\sigma_1^{*2})^2]} \\ &\quad + (D_m \sqrt{d} + D_m^2) \sigma_1^{*2} \sqrt{\mathbb{E}_{\mathcal{D}_1} [w_l^u]} \sqrt{\mathbb{E}_{\mathcal{D}_1} [(\|v\|^2 - d\sigma_1^{*2})^2]} \end{aligned}$$

Then, we are summing over $O \left((1 + \pi_l^*/\pi_1^*) (R_{l1}^*/\sigma_1^*) \exp(-R_{l1}^{*2}/128(\sigma_l^* \vee \sigma_1^*)^2) \right)$ for $l \neq 1$, which is bounded again with Lemma 26. Dominating error bound is therefore $e_{15} \leq cD_m \sqrt{d} \sigma_1^{*2}$ again for small constant c . We follow the similar procedure for the final term e_{16} .

$$\begin{aligned} e_{16} &= \left| \mathbb{E}_{\mathcal{D}_1} \left[\sum_{l \neq 1} w_1^u w_l^u (\|X - \mu_1^*\|^2 - d\sigma_1^{*2}) (\|X - \mu_l^u\|^2 - d\sigma_l^{u2}) / (2\sigma_l^{u2}) \right] (\sigma_l^2 - \sigma_l^{*2}) / \sigma_l^{u2} \right| \\ &\leq \left(2D_m / \sqrt{d} \right) \sum_{l \neq 1} (1/\sigma_l^{*2}) \mathbb{E}_{\mathcal{D}_1} \left[w_1^u w_l^u (\|v\|^2 - d\sigma_1^{*2}) ((\|v\|^2 - d\sigma_1^{*2}) + 2R_{l1}^* \langle v, s \rangle + 2R_{l1}^{*2} + d(\sigma_1^{*2} - \sigma_l^{u2})) \right] \\ &\leq \left(2D_m / \sqrt{d} \right) \sum_{l \neq 1} (1/\sigma_l^{*2}) \left(\mathbb{E}_{\mathcal{D}_1} [w_l^u (\|v\|^2 - d\sigma_1^{*2})^2] + 2\mathbb{E}_{\mathcal{D}_1} [R_{l1}^* w_l^u (\langle v, s \rangle + R_{l1}^*) (\|v\|^2 - d\sigma_1^{*2})] \right) \end{aligned}$$

$$+ \mathbb{E}_{\mathcal{D}_1}[w_l^u(\|v\|^2 - d\sigma_1^{*2})d(\sigma_1^{*2} - \sigma_l^{*2})] + \mathbb{E}_{\mathcal{D}_1}[w_l^u(\|v\|^2 - d\sigma_1^{*2})d(\sigma_l^{*2} - \sigma_l^{u2})] \Big).$$

The remaining steps are bounding each four term in the summation.

$$\begin{aligned} \mathbb{E}_{\mathcal{D}_1}[w_l^u(\|v\|^2 - d\sigma_1^{*2})^2] &\leq \sqrt{\mathbb{E}_{\mathcal{D}_1}[w_l^u]^2} \sqrt{\mathbb{E}_{\mathcal{D}_1}[(\|v\|^2 - d\sigma_1^{*2})^4]} \\ &\leq O\left(d\sigma_1^{*4}(1 + \pi_l^*/\pi_1^*) \exp(-R_{l1}^{*2}/128(\sigma_l^* \vee \sigma_1^*)^2)\right), \end{aligned}$$

where we used (31) for bounding fourth order central moment of degree- d chi-square random variable.

The second one is

$$\begin{aligned} 2R_{l1}^* \mathbb{E}_{\mathcal{D}_1}[w_l^u(\langle v, s \rangle + R_{l1}^*)(\|v\|^2 - d\sigma_1^{*2})] &\leq 2R_{l1}^* \sqrt{2\mathbb{E}_{\mathcal{D}_1}[w_l^u(\langle v, s \rangle^2 + R_{l1}^{*2})]} \sqrt{\mathbb{E}_{\mathcal{D}_1}[(\|v\|^2 - d\sigma_1^{*2})^2]} \\ &\leq O\left(\sqrt{d}R_{l1}^{*2}\sigma_1^{*2}(1 + \pi_l^*/\pi_1^*) \exp(-R_{l1}^{*2}/128(\sigma_l^* \vee \sigma_1^*)^2)\right). \end{aligned}$$

Third term we will use Lemma 22 to get,

$$\begin{aligned} d(\sigma_1^{*2} - \sigma_l^{*2})\mathbb{E}_{\mathcal{D}_1}[w_l^u(\|v\|^2 - d\sigma_1^{*2})] &\leq d(\sigma_1^{*2} - \sigma_l^{*2})\sqrt{\mathbb{E}_{\mathcal{D}_1}[w_l^u]^2} \sqrt{\mathbb{E}_{\mathcal{D}_1}[(\|v\|^2 - d\sigma_1^{*2})^2]} \\ &\leq O\left(\sqrt{d}R_{l1}^*(\sigma_l^* \vee \sigma_1^*) \exp(-R_{l1}^{*2}/128(\sigma_l^* \vee \sigma_1^*)^2)\sqrt{2d}\sigma_1^{*2}\right), \end{aligned}$$

and finally we have

$$\begin{aligned} d(\sigma_1^{*2} - \sigma_l^{u2})\mathbb{E}_{\mathcal{D}_1}[w_l^u(\|v\|^2 - d\sigma_1^{*2})] &\leq \sqrt{d}D_m\sigma_l^{*2}\sqrt{\mathbb{E}_{\mathcal{D}_1}[w_l^u]^2} \sqrt{\mathbb{E}_{\mathcal{D}_1}[(\|v\|^2 - d\sigma_1^{*2})^2]} \\ &\leq O\left(\sqrt{d}\sigma_l^{*2} \exp(-R_{l1}^{*2}/128(\sigma_l^* \vee \sigma_1^*)^2)\sqrt{2d}\sigma_1^{*2}\right). \end{aligned}$$

Now, we can collect all terms and bound e_{16} as

$$e_{16} \leq \sqrt{d}D_m\sigma_1^{*2}O\left(\sum_{l \neq 1} (R_{l1}^*/\sigma_l^*)^2(1 + \pi_l^*/\pi_1^*) \exp(-R_{l1}^{*2}/128(\sigma_l^* \vee \sigma_1^*)^2)\right) \leq c\sqrt{d}D_m\sigma_1^{*2},$$

with small constant c .

Errors from all components: Now we have that errors across all components can be bounded as

$$\begin{aligned} \sum_{j=1}^k \pi_j^*(e_{j1} + \dots + e_{j6}) &\leq \pi_1^*c(D_m\sqrt{d}\sigma_1^{*2}) \\ &\quad + c(D_m\sqrt{d}\sigma_1^{*2})O\left(\sum_{j \neq 1} (\pi_1^* + \pi_j^*)(R_{j1}/\sigma_1^*)^4 \exp(-R_{j1}^{*2}/128(\sigma_j^* \vee \sigma_1^*)^2)\right) \\ &\leq c'D_m\sqrt{d}\sigma_1^{*2}\pi_1^*, \end{aligned}$$

using Lemma 26 with $q = 4$. Finally, recalling the errors for variances,

$$|\sigma_1^{+2} - \sigma_1^{*2}| \leq c'D_m\sqrt{d}\sigma_1^{*2}\pi_1^*/d\pi_1^+ + \|\mu_1^+ - \mu_1^*\|^2/d \leq c_\sigma D_m\sigma_1^{*2}/\sqrt{d},$$

where we considered that the mean estimator at the next iteration is also improved. The constant is $c_\sigma < 1$ given large enough separation between centers of Gaussians. \blacksquare

In all cases, we can conclude that $D_m^+ = \max_j (\|\mu_j^+ - \mu_j^*\|/\sigma_j^*, |\pi_j^+ - \pi_j^*|/\pi_j^*, \sqrt{d}|\sigma_j^{+2} - \sigma_j^{*2}|/\sigma_j^{*2}) \leq \gamma D_m$, for some constant $\gamma < 1$.

E.4. Proof of Auxiliary Lemmas

E.4.1. PROOF OF LEMMA 26

Proof We will only show for $q = 4$ and the other cases will follow similarly. Again, let $x := R_{j1}^{*2}/(\sigma_1^* \vee \sigma_j^*)^2$. Then, since $x \geq C^2 \geq 128^2$ by the separation condition, $x \geq 512 \log x$. Then the proof is again trivial:

$$\begin{aligned} \sum_{j \neq 1} (\pi_1^* + \pi_j^*) \rho_\sigma^4 x^2 \exp(-x^2/128) &\leq \sum_{j \neq 1} (\pi_1^* + \pi_j^*) \rho_\sigma^4 \exp(-x^2/256) \\ &\leq \sum_{j \neq 1} (\pi_1^* + \pi_j^*) \rho_\sigma^4 (\rho_\sigma/\pi_{\min})^{-32} \ll c \pi_{\min}, \end{aligned}$$

for small constant c . \blacksquare

E.4.2. PROOF OF LEMMA 27

Proof If $l = j$, upper bound is simply obtained by setting $w_j^u = 1$.

$$\mathbb{E}_{\mathcal{D}_j} [w_j^u \langle v + \mu_j^* - \mu_j^u, s \rangle^2] \leq \mathbb{E}_{\mathcal{D}_j} [2 \langle v, s \rangle^2 + 2\sigma_j^{*2} D_m^2] \leq 4\sigma_j^{*2}.$$

If $l \neq j$, then

$$\mathbb{E}_{\mathcal{D}_j} [w_l^u \langle v + \mu_j^* - \mu_l^u, s \rangle^2] \leq \mathbb{E}_{\mathcal{D}_j} [w_l^u (2 \langle v, s \rangle^2 + 4R_{jl}^{*2})] \leq O \left(R_{jl}^{*2} (1 + \pi_l^*/\pi_j^*) \exp \left(-R_{jl}^{*2}/64(\sigma_j^* \vee \sigma_l^*)^2 \right) \right),$$

where we used Lemma 29. Then, summation over $l \neq j$ yields

$$\sum_{l \neq j} \sqrt{\mathbb{E}_{\mathcal{D}_j} [w_l^u \langle v + \mu_j^* - \mu_l^u, s \rangle^2] / \sigma_l^{*2}} \leq \sum_{l \neq j} O \left((R_{jl}^*/\sigma_l^*) (1 + \pi_l^*/\pi_j^*) \exp \left(-R_{jl}^{*2}/128(\sigma_j^* \vee \sigma_l^*)^2 \right) \right) \leq c,$$

for some small constant c by Lemma 26. \blacksquare

E.4.3. PROOF OF LEMMA 28

Proof If $l = j$, then the upper bound can be found by setting $w_l^u = 1$. That is,

$$\begin{aligned} &\left(1/\sigma_j^{*2} \right) \left(\sqrt{\mathbb{E}_{\mathcal{D}_j} [(\|v\|^2 - d\sigma_j^{*2})^2]} + 2\sqrt{\mathbb{E}_{\mathcal{D}_j} [\langle v, \mu_j^* - \mu_j^u \rangle^2]} + d(\sigma_j^{*2} - \sigma_l^{u2}) \right) \\ &\leq \left(\sqrt{2d}\sigma_j^{*2} + 4\sigma_j^{*2} D_m + D_m \sigma_j^{*2} \sqrt{d} \right) / \sigma_j^{*2} = O(\sqrt{d}). \end{aligned}$$

Now consider when $l \neq j$. In addition to previous lemmas, we need the fact about chi-square distribution with degree d . Its fourth central moment is,

$$\mathbb{E}_{v \sim \mathcal{N}(0, I_d)}[(\|v\|^2 - d)^4] = 12d(d+4). \quad (31)$$

With this fact,

$$\mathbb{E}_{\mathcal{D}_j}[w_l^u (\|v\|^2 - d\sigma_j^{*2})^2] \leq \sqrt{\mathbb{E}_{\mathcal{D}_j}[w_l^u]} \sqrt{\mathbb{E}_{\mathcal{D}_j}[(\|v\|^2 - d\sigma_j^{*2})^4]} \leq 8d\sigma_j^{*4}(1 + \pi_l^*/\pi_j^*) \exp(-R_{jl}^{*2}/128(\sigma_j^* \vee \sigma_l^*)^2).$$

Other terms can also be bounded similarly as

$$\begin{aligned} \mathbb{E}_{\mathcal{D}_j}[w_l^u \langle v, \mu_j^* - \mu_l^u \rangle^2] &\leq 2R_{jl}^{*2} \mathbb{E}_{\mathcal{D}_j}[w_l^u \langle v, s \rangle^2] \leq O\left((1 + \pi_l^*/\pi_j^*)R_{jl}^{*4}\right) \exp(-R_{jl}^{*2}/64(\sigma_j^* \vee \sigma_l^*)^2), \\ \mathbb{E}_{\mathcal{D}_j}[w_l^u] &\leq O\left((1 + \pi_l^*/\pi_j^*)\right) \exp(-R_{jl}^{*2}/64(\sigma_j^* \vee \sigma_l^*)^2). \end{aligned}$$

We use Lemma 22 to bound the value of $d(\sigma_j^{*2} - \sigma_l^{*2})\mathbb{E}_{\mathcal{D}_j}[w_l^u]$. We first remove u superscript as

$$\begin{aligned} d(\sigma_j^{*2} - \sigma_l^{*2})\mathbb{E}_{\mathcal{D}_j}[w_l^u] &= (d(\sigma_j^{*2} - \sigma_l^{*2}) + d(\sigma_l^{*2} - \sigma_l^{*2}))\mathbb{E}_{\mathcal{D}_j}[w_l^u] \\ &\leq d(\sigma_j^{*2} - \sigma_l^{*2})\mathbb{E}_{\mathcal{D}_j}[w_l^u] + \sqrt{d}\sigma_l^{*2}\mathbb{E}_{\mathcal{D}_j}[w_l^u]. \end{aligned}$$

From Lemma 22, either we have $(\sigma_j^{*2} - \sigma_l^{*2}) \leq 10R_{jl}^*(\sigma_j^* \vee \sigma_l^*)/\sqrt{d}$ or

$$\mathbb{E}_{\mathcal{D}_j}[w_l^u] \leq O\left(\exp(-d \min(1, t^2)/256) \exp(-R_{jl}^{*2}/64(\sigma_j^* \vee \sigma_l^*)^2)\right),$$

where $t = (\sigma_j^{*2} - \sigma_l^{*2})/\sigma_l^{*2}$. This gives us

$$d(\sigma_j^{*2} - \sigma_l^{*2})/\sigma_l^{*2}\mathbb{E}_{\mathcal{D}_j}[w_l^u] \leq O\left((1 + \pi_l^*/\pi_j^*)R_{jl}^*(\sigma_j^* \vee \sigma_l^*)\sqrt{d} \exp(-R_{jl}^{*2}/64(\sigma_j^* \vee \sigma_l^*)^2)\right).$$

Now summing up all terms over $l \neq j$, we get

$$\begin{aligned} &\sum_{l \neq j} \sqrt{d} \sqrt{1 + \pi_l^*/\pi_j^*} O\left(R_{jl}^{*2}/\sigma_l^{*2} \exp(-R_{jl}^{*2}/128(\sigma_j^* \vee \sigma_l^*)^2) + \sigma_j^{*2}/\sigma_l^{*2} \exp(-R_{jl}^{*2}/256(\sigma_j^* \vee \sigma_l^*)^2)\right) \\ &\leq \sqrt{d}/\pi_j^* \sum_{l \neq j} (\pi_j^* + \pi_l^*) O\left(R_{jl}^{*2}/\sigma_l^{*2} \exp(-R_{jl}^{*2}/128(\sigma_j^* \vee \sigma_l^*)^2) + \sigma_j^{*2}/\sigma_l^{*2} \exp(-R_{jl}^{*2}/256(\sigma_j^* \vee \sigma_l^*)^2)\right) \\ &\leq c\sqrt{d}, \end{aligned}$$

for some small constant c . ■