

Optimal estimation of Gaussian mixtures via denoised method of moments

Yihong Wu and Pengkun Yang*

Working paper: July 15, 2018

Abstract

The Method of Moments is one of the most widely used methods in statistics for parameter estimation, obtained by solving the system of equations that match the population and estimated moments. However, in practice and especially for the important case of mixture models, one frequently needs to contend with the difficulties of non-existence or non-uniqueness of statistically meaningful solutions, as well as the high computational cost of solving large polynomial systems. Moreover, theoretical analysis of method of moments are mainly confined to asymptotic normality style of results established under strong assumptions.

This paper considers estimating a k -component Gaussian location mixture with a common (possibly unknown) variance parameter. To overcome the aforementioned theoretic and algorithmic hurdles, a crucial step is to denoise the moment estimates by projecting to the truncated moment space before executing the method of moments. Not only does this regularization ensures existence and uniqueness of solutions, it also yields fast solvers by means of Gauss quadrature. Furthermore, by proving new moment comparison theorems in Wasserstein distance via polynomial interpolation and majorization, we establish the statistical guarantees and optimality of the proposed procedure. These results can also be viewed as provable algorithms for Generalized Method of Moments which involves non-convex optimization and lacks theoretical guarantees.

Contents

1	Introduction	2
1.1	Gaussian mixture model	2
1.2	Failure of the usual method of moments	3
1.3	Main results	4
1.4	Why Wasserstein distance?	7
1.5	Related work	8
1.6	Notations	9
1.7	Organization	9
2	Preliminaries	9
2.1	Moment space, SDP characterization, and Gauss quadrature	9
2.2	Polynomial interpolation, majorization, and the Neville diagram	11
2.3	Wasserstein distance	12

*Y. Wu is with the Department of Statistics and Data Science, Yale University, New Haven, CT, yihong.wu@yale.edu. P. Yang is with the Department of Electrical and Computer Engineering and the Coordinated Science Lab, University of Illinois at Urbana-Champaign, Urbana, IL, pyang14@illinois.edu.

3	Optimal transport and moment comparison theorems	14
4	Estimators and statistical guarantees	15
4.1	Known variance	15
4.2	Unknown variance	17
4.3	Adaptive results	20
4.4	Unbounded means	21
5	Lower bounds	23
6	Numerical experiments	25
7	Extensions and discussions	28
7.1	Gaussian location-scale mixtures	28
7.2	Multiple dimensions	29
7.3	General finite mixtures	31
8	Proofs	32
8.1	Proof of moments comparison theorems	32
8.2	Proofs of density estimation	37
8.3	Proofs for Section 4.1	39
8.4	Proofs for Section 4.2	40
8.5	Proofs for Section 4.3	43
8.6	Proofs for Section 4.4	44
8.7	Proofs for Section 5	45
A	Standard form of the semidefinite programming (24)	47
B	Gauss quadrature algorithm	47
C	Auxiliary lemmas	47

1 Introduction

1.1 Gaussian mixture model

Consider a k -component Gaussian location mixture model, where each observation is distributed as

$$X \sim \sum_{i=1}^k w_i N(\mu_i, \sigma^2). \tag{1}$$

Here w_i is the mixing weight such that $w_i \geq 0$ and $\sum_i w_i = 1$, μ_i is the mean (center) of the i^{th} component, and σ is the common standard deviation. Equivalently, we can write the distribution of an observation X as a convolution

$$X \sim \nu * N(0, \sigma^2), \tag{2}$$

where $\nu = \sum_{i=1}^k w_i \delta_{\mu_i}$ denotes the *mixing distribution*. Thus, we can write $X = U + \sigma Z$, where $U \sim \nu$ is referred to as the latent variable, and Z is standard normal independent of U .

Generally speaking, there are three formulations of learning mixture models:

- *Parameter estimation*: estimate the means μ_i 's and the weights w_i 's up to a global permutation, and possibly also σ^2 .
- *Density estimation*: estimate the probability density function of the Gaussian mixture under certain loss such as L_2 or Hellinger distance. This task is further divided into the cases of *proper* and *improper* learning, depending on whether the estimator is required to be a k -Gaussian mixture or not; in the latter case, there are more flexibility in designing the estimator but less interpretability.
- *Clustering*: estimate the latent variable of each sample (i.e. U_i , if the i th sample is represented as $X_i = U_i + \sigma Z_i$) with a small misclassification rate.

It is clear that to ensure the possibility of clustering it is necessary to impose certain separation conditions between the clusters; however, as far as estimation is concerned, both parametric and non-parametric, no separation condition should be needed and one can obtain accurate estimates of the parameters even when clustering is impossible. Furthermore, one should be able to learn from the data the order of the mixture model, that is, the number of components. However, in the present literature, most of the estimation procedures are either classification-based, or rely on separation conditions in the analysis. Bridging this conceptual divide is one of the main motivations of the present paper.

The methodologies are largely divided into likelihood-based and moment-based methods. A detailed review of them is given in Section 1.5. Among likelihood-based methods, the *Maximum Likelihood Estimate* (MLE) is not efficiently computable. The most popular heuristic procedure to approximate the MLE is the *Expectation-Maximization* (EM) algorithm [DLR77]. However, no theoretical guarantee is known in general. The moment-based methods include the classical *method of moments* (MM) [Pea94] and many extensions [Han82, AGH⁺14]. The usual method of moments suffers from many issues that are elaborated in the next subsection.

1.2 Failure of the usual method of moments

The method of moments, commonly attributed to Pearson [Pea94], produces an estimator by equating the population moments to the sample moments and solving the system of moment equations. This method is conceptually simple but suffers from the following problems, especially in the context of mixture models:

- *Solubility*: the method of moments entails solving a multivariate polynomial system, in which one frequently encounters non-existence or non-uniqueness of statistically meaningful solutions.
- *Computation*: solving moment equations can be computationally intensive. For instance, for k -component Gaussian mixture models, the system of moment equations consist of $2k - 1$ polynomial equations with $2k - 1$ variables.
- *Accuracy*: existing statistical literature on the method of moments [VdV00, Han82] either shows only consistency under weak assumptions, or proves asymptotic normality assuming very strong regularity conditions (so that delta method works), which generally do not hold in mixture models since the convergence rates can be slower than parametric. Some results on nonparametric rates are known (cf. [VdV00, Theorem 5.52] and [Kos07, Theorem 14.4]) but the conditions are extremely hard to verify.

To explain the failure of the vanilla method of moments in Gaussian mixture models, consider the following simple example of two components:

Example 1. Consider a Gaussian mixture model with two unit variance components. Since there are three parameters μ_1, μ_2 and $w_1 = 1 - w_2$, we use the first three moments and solve the following system of equations:

$$\begin{aligned}\mathbb{E}_n[X] &= \mathbb{E}[X] = w_1\mu_1 + w_2\mu_2, \\ \mathbb{E}_n[X^2] &= \mathbb{E}[X^2] = w_1\mu_1^2 + w_2\mu_2^2 + 1, \\ \mathbb{E}_n[X^3] &= \mathbb{E}[X^3] = w_1\mu_1^3 + w_2\mu_2^3 + 3(w_1\mu_1 + w_2\mu_2),\end{aligned}\tag{3}$$

where $\mathbb{E}_n[X^i] \triangleq \frac{1}{n} \sum_{j=1}^n X_j^i$ denotes the i^{th} moment of the empirical distribution from n i.i.d. samples. The right-hand sides of (3) are related to the moments of the mixing distribution by a linear system, and we can equivalently rewrite the moment equations (3) as:

$$\begin{aligned}\mathbb{E}_n[X] &= \mathbb{E}[U] = w_1\mu_1 + w_2\mu_2, \\ \mathbb{E}_n[X^2 - 1] &= \mathbb{E}[U^2] = w_1\mu_1^2 + w_2\mu_2^2, \\ \mathbb{E}_n[X^3 - 3X] &= \mathbb{E}[U^3] = w_1\mu_1^3 + w_2\mu_2^3.\end{aligned}\tag{4}$$

where $U \sim w_1\delta_{\mu_1} + w_2\delta_{\mu_2}$. It turns out that with finitely many samples, there is always a non-zero chance that (4) has no solution; even with infinite samples, it is possible that the solution does not exist with constant probability. To see this, note that, from the first two equations of (4), the solution does not exist whenever

$$\mathbb{E}_n[X^2] - 1 < \mathbb{E}_n^2[X],\tag{5}$$

i.e., the Cauchy-Schwarz inequality fails. Consider the case $\mu_1 = \mu_2 = 0$, i.e., $X \sim N(0, 1)$. Then (5) is equivalent to

$$n(\mathbb{E}_n[X^2] - \mathbb{E}_n^2[X]) \leq n,$$

where the left-hand side follows the χ^2 -distribution with $n - 1$ degrees of freedom. Thus, (5) occurs with probability approaching $\frac{1}{2}$ as n diverges, according to the central limit theorem.

In view of the above example, we note that the main issue with the usual method of moments is the following: although individually each moment estimate is accurate (\sqrt{n} -consistent), jointly they do not correspond to the moments of any distribution. Moment vectors satisfy many geometric constraints, e.g., the Cauchy-Schwarz and Hölder inequalities, and lie in a convex set known as the *moment space*. The failure occurs with nonzero probability when the population moments are estimated with finite samples; even with infinite samples, it also provably happens with constant probability when the order of the mixture model is strictly less than k , or equivalently, the population moments are on the boundary of moment space (the equivalence is established in Lemma 35).

1.3 Main results

In this paper, we propose the *denoised method of moments* (DMM), where the moment estimates are jointly denoised through a projection step. This method consists of three main steps: compute noisy estimates of moments; project them onto the moment space; execute the usual method of moments. This procedure resolves the three issues of the vanilla method of moments identified in Section 1.2 simultaneously:

- *Solubility*: a unique statistically meaningful solution is guaranteed to exist;
- *Computation*: the solution can be found through an efficient algorithm (Gauss quadrature) instead of invoking a generic solver of polynomial systems;
- *Accuracy*: the solution provably achieves the optimal rate of convergence, and automatically adaptive to the clustering structure of the population.

We emphasize that the denoising (projection) step is explicitly carried out via a convex optimization in Section 4.1, and implicitly used in analyzing Lindsay’s algorithm [Lin89] in Section 4.2, when the variance parameter is known and unknown, respectively. After denoising, the solubility is guaranteed by the classical theory of moments, which is summarized in Section 2.1. The guarantee and optimality of our algorithm are introduced next.

The framework in this paper is different from most previous parameter estimation in mixture models: rather than estimate parameters of each component, our goal is to recover the mixing distribution consisting of those parameters. The main benefits of this framework include:

- *Assumption-free*: the task of recovering individual components requires assumptions to ensure identifiability, such as strictly positive mixing weight and distinct components. In our framework, those unidentifiable instances have the same mixing distribution to be estimated. Furthermore, under the usual assumption such as separation conditions, statistical guarantees in terms of the W_1 distance can be naturally translated to the estimation of the individual parameters.
- *Inference on the number of components*: this framework allows us to deal with misspecified model and estimate the order of finite mixture.

This framework can be equivalently viewed as deconvolution from (2): the goal is to recover the distribution ν using observations following the convolved distribution.

In this framework, it turns out a meaningful and flexible loss function is the 1-Wasserstein distance (see Section 1.4 for a justification in the context of mixture models), defined by

$$W_1(\nu, \nu') \triangleq \inf\{\mathbb{E}[\|X - Y\|] : X \sim \nu, Y \sim \nu'\}, \quad (6)$$

which, in one dimension, coincides with the L_1 -distance between the cumulative distribution functions (CDFs) [Vil03]. Throughout the paper, we assume that the number of components satisfy

$$k = O\left(\frac{\log n}{\log \log n}\right) \quad (7)$$

Denote the underlying model as a convolution of $\nu = \sum_i w_i \delta_{\mu_i}$ and $N(0, \sigma^2)$. Our main result is the following theorem:

Theorem 1 (Optimal rates). *Suppose that $|\mu_i| \leq M$ for $M \geq 1$ and σ is bounded by a constant, and both k and M are given.*

- *If σ is known, then there exists an estimator $\hat{\nu}$ computable in $O(kn)$ time such that, with probability at least $1 - \delta$,*

$$W_1(\nu, \hat{\nu}) \leq O\left(Mk^{1.5} \left(\frac{n}{\log(1/\delta)}\right)^{-\frac{1}{4k-2}}\right); \quad (8)$$

- If σ is unknown, then there exists an estimator $(\hat{\nu}, \hat{\sigma})$ computable in $O(kn)$ time such that, with probability at least $1 - \delta$,

$$W_1(\nu, \hat{\nu}) \leq O \left(Mk^2 \left(\frac{n}{\log(1/\delta)} \right)^{-\frac{1}{4k}} \right), \quad (9)$$

and

$$|\sigma^2 - \hat{\sigma}^2| \leq O \left(M^2 k \left(\frac{n}{\log(1/\delta)} \right)^{-\frac{1}{2k}} \right). \quad (10)$$

The above results are minimax rate-optimal for constant k as shown in Section 5. Note that those results are proved under the worst-case scenario where the centers can be arbitrarily close, e.g., components completely overlap. It is reasonable to expect, for instance, a parametric rate when components are well-separated and generated sufficient samples. Formally, we have the following definition:

Definition 1. The Gaussian mixture (1) has k_0 (γ, ω) -separated clusters if there exists a partition S_1, \dots, S_{k_0} of $[k]$ such that

- $|\mu_i - \mu_{i'}| \geq \gamma$ for any $i \in S_\ell$ and $i' \in S_{\ell'}$ such that $\ell \neq \ell'$;
- $\sum_{i \in S_\ell} w_i \geq \omega$ for each ℓ .

In the absence of the minimal weight condition (i.e. $\omega = 0$), we say the Gaussian mixture has k_0 γ -separated clusters.

In this case our estimates attain adaptive rates:

Theorem 2 (Adaptive rate). *Under the conditions of Theorem 1, suppose there are k_0 (γ, ω) -separated clusters such that $\gamma\omega \geq C\epsilon$ for some absolute constant $C > 2$, where ϵ denotes the right-hand side of (8) and (9) when σ is known and unknown, respectively.*

- If σ is known, then, with probability at least $1 - \delta$,

$$W_1(\nu, \hat{\nu}) \leq O_k \left(M\gamma^{-\frac{2k_0-2}{2(k-k_0)+1}} \left(\frac{n}{\log(k/\delta)} \right)^{-\frac{1}{4(k-k_0)+2}} \right); \quad (11)$$

- if σ is unknown, then, with probability at least $1 - \delta$,¹

$$\sqrt{|\sigma^2 - \hat{\sigma}^2|}, W_1(\nu, \hat{\nu}) \leq O_k \left(M\gamma^{-\frac{k_0-1}{k-k_0+1}} \left(\frac{n}{\log(k/\delta)} \right)^{-\frac{1}{4(k-k_0+1)}} \right). \quad (12)$$

The result (11) is also minimax rate-optimal when k, k_0 and γ are constants, in view of the lower bounds in [HK15]. We also provide a simple proof in Remark 4 by extending the lower bound argument recipe in Section 5. We do not have a matching lower bound for (12). In the fully-separated case ($k_0 = k$), (12) reduces to $n^{-\frac{1}{4}}$ while a parametric rate is achievable.

¹ Note that the estimation rate for the mean part ν is the square root of the rate for estimating the variance parameter σ^2 . Intuitively, this phenomenon is due to the infinite divisibility of the Gaussian distribution: note that for the location mixture model $\nu * N(0, \sigma^2)$ with $\nu \sim N(0, \epsilon^2)$, $\sigma^2 = 1$ has the same distribution as one with $\nu \sim \delta_0$, and $\sigma^2 = 1 + \epsilon^2$.

Finally, we briefly discuss density estimation, where the goal is to recover the density function of the unknown mixture model. We accurately recovered the mixing distribution by $\hat{\nu}$, so the density function \hat{f} of the convolution $\hat{\nu} * N(0, \sigma^2)$ serves as a natural estimate. Theorem 3 below shows that the density estimate \hat{f} is $O(\frac{1}{n})$ -close to the true density f in χ^2 -divergence, defined by $\chi^2(f\|g) \triangleq \int \frac{(f-g)^2}{g}$, which immediately implies other common distance measures such as Kullback-Leibler divergence, total variation, and Hellinger distance.

Theorem 3 (Density estimation). *Under the conditions of Theorem 1, denote the density of the underlying model by $f = \nu * N(0, \sigma^2)$. If σ is given, then there exists an estimate \hat{f} such that*

$$\chi^2(\hat{f}\|f) + \chi^2(f\|\hat{f}) \leq O_k(\log(1/\delta)/n),$$

with probability $1 - \delta$.

The previous results establish properties of the estimated k Gaussian mixture in correctly specified models. In practice, however, it is never known whether the data collected corresponds to the model assumptions. It is important to consider the impact of model misspecification, i.e., the data is generated by some distribution other than a k Gaussian mixture. In this case, our procedure still reports a meaningful estimate that is close to the best k Gaussian fit of the unknown distribution:

Theorem 4 (Misspecified model). *Assume the distribution of the underlying model, denoted by f , is 1-subgaussian. Suppose there exists a k Gaussian mixture g with a given σ such that $\text{TV}(f, g) \leq \epsilon$. Then, there exists an estimate \hat{f} such that*

$$\text{TV}(\hat{f}, f) \leq O_k\left(\epsilon\sqrt{\log(1/\epsilon)} + \sqrt{\log(1/\delta)/n}\right),$$

with probability $1 - \delta$.

1.4 Why Wasserstein distance?

Throughout the paper we estimate the mixing distribution ν under the Wasserstein distance. This is a natural criterion, which is not too stringent to give trivial result (e.g. the Kolmogorov-Smirnov (KS) distance²) and, at the same time, strong enough to yield meaningful guarantees on the means and weights. In fact, the commonly used criterion $\min_{\Pi} \sum_i |\theta_i - \hat{\theta}_{\Pi(i)}|$ over all permutations Π is precisely k times the Wasserstein distance between two equally weighted distributions [Vil03].

Furthermore, we can harvest statistical guarantees on the support sets and weights of the estimated mixing distribution under the usual assumptions in literatures [Das99, KMV10, HP15] that include separation between μ_i 's and lower bound on the w_i 's. See Section 2.3 for a detailed discussion. We highlight the following result, phrased in terms of the parameter estimation error up to a permutation:

Lemma 1. *Suppose*

$$\nu = \sum_{i=1}^k w_i \delta_{\mu_i}, \quad \hat{\nu} = \sum_{i=1}^k \hat{w}_i \delta_{\hat{\mu}_i},$$

such that $W_1(\nu, \hat{\nu}) < \epsilon$. Let

$$\epsilon_1 = \min\{|\mu_i - \mu_j|, |\hat{\mu}_i - \hat{\mu}_j| : 1 \leq i < j \leq k\}, \quad \epsilon_2 = \min\{w_i, \hat{w}_i : i \in [k]\}.$$

² Consider two mixing distributions δ_0 and δ_ϵ with arbitrarily small ϵ , whose KS distance is always one.

If $\epsilon < \epsilon_1 \epsilon_2 / 4$, then, there exists a permutation Π such that

$$\|\mu - \Pi \hat{\mu}\|_\infty < \epsilon / \epsilon_2, \quad \|w - \Pi \hat{w}\|_\infty < 2\epsilon / \epsilon_1,$$

where $\mu = (\mu_1, \dots, \mu_k)$, $w = (w_1, \dots, w_k)$ denote the atoms and weights of ν , respectively, and $\hat{\mu}, \hat{w}$ denote those of $\hat{\nu}$,

1.5 Related work

There exists a vast literature on mixture models, in particular Gaussian mixtures, and the method of moments. For a comprehensive review see [Lin95, FS06]. Below we highlight a few existing results that are related to the present paper.

Likelihood-based methods. Maximum likelihood estimation (MLE) is one of the most useful method in parameter estimation, and applicable in Gaussian mixtures. Under strong assumptions, MLE is consistent and asymptotically normal [RW84], and the density function of the estimated model is studied by [GW00, GVDV01]. However, those assumptions are difficult to verify, and it is computationally hard to obtain the global maximizer.

Expectation-Maximization (EM) algorithm [DLR77] is an iterative method to approximate the MLE. It has been widely applied Gaussian mixture models [RW84, XJ96] and more recently in high-dimensional settings [BWY17]. This method converges to a local maximizer of the likelihood function in general, but possibly suffers poor quality, and no guarantee exists that it converges to the global maximizer. In practice we need to employ heuristic selections of initial guess [KX03] and stopping criterion [SMA00], and possibly data augmentation techniques [PL01]; its slow convergence rate is also widely observed through empirical experiments [RW84, KX03]. Additionally, this method accesses all raw data in each iteration, which is particularly expensive in the presence of large data sets in modern applications.

Moment-based methods. The simplest moment-based method is the method of moments (MM) introduced by Pearson [Pea94]. The failure of the vanilla MM described in Section 1.2 has motivated various modifications including the minimum chi-square estimation [Ber80] and the *generalized method of moments* [Han82].

Generalized method of moments (GMM) is proposed in [Han82] and is widely applied to analyze economic and financial data. See a list in [Hal05, Table 1.1]. Instead of to exactly solve moment equations, GMM aims to minimize the difference between the sample moments and the fitted moments. This method enjoys various nice asymptotic properties [Han82]. However, the optimization problem is not efficiently solvable in Gaussian mixture models, where the moment conditions are non-convex. In this problem, our procedure provably solves GMM exactly and efficiently, which is another algorithmic contribution of the present paper. We also note that minimizing a quadratic form in GMM is not crucial, and minimizing any distance gives the same guarantee. See a discussion in Section 4.1.

There are many recent works studying provable results [MV10, KMV10, BS10] and the optimal rate of accuracy [HP15] using moment-based methods. A key observation is that a finite order Gaussian mixture is identifiable (up to a permutation of components) from a finite number of moments. However, [MV10, KMV10, BS10] rely on the exhaustive search over a grid of parameters and the comparison with estimated moments, which is too expensive in practice; [HP15] follows a similar approach as [Pea94] by carefully analyzing a polynomial equation obtained from the usual method of moments, and shows the optimal rate $\Theta(n^{-1/12})$ in the case of two components, but this approach is difficult to generalize to more components.

Other methods. From a frequentist perspective, the minimum distance estimation is studied by [Che95, HK15], where they consider the Gaussian mixture model that is closest to the empirical distribution under the KS distance. The optimal statistical guarantees of the estimated mixing distribution under the Wasserstein distance is obtained in [HK15] (which corrects the previous bound in [Che95]). However, their estimator is in general computationally expensive. In contrast, denoised method of moments is efficiently computable and adaptively achieves the optimal rate of accuracy as given in Theorem 2.

1.6 Notations

A discrete distribution supported on k atoms is called a k -atomic distribution. The expectation of a given function f under a distribution μ is denoted by $\mathbb{E}_\mu f = \mathbb{E}_\mu[f(X)] = \int f(x)\mu(dx)$, and the subscript μ may be omitted if it is specified from the context. The empirical mean of f from n samples is denoted as $\mathbb{E}_n[f(X)] = \frac{1}{n} \sum_{i=1}^n f(X_i)$, where X_1, \dots, X_n are i.i.d. copies of X . The r^{th} moment of a distribution μ is denoted by $m_r(\mu) \triangleq \mathbb{E}_\mu X^r$. The moment matrix associated with m_0, m_1, \dots, m_{2r} is a Hankel matrix of order $r + 1$:

$$\mathbf{M}_r = \begin{bmatrix} m_0 & m_1 & \cdots & m_r \\ m_1 & m_2 & \cdots & m_{r+1} \\ \vdots & \vdots & \ddots & \vdots \\ m_r & m_{r+1} & \cdots & m_{2r} \end{bmatrix}. \quad (13)$$

The χ^2 -divergence between two distributions μ and ν is denoted by $\chi^2(\mu\|\nu) = \int \frac{(d\mu - d\nu)^2}{d\nu}$. For matrices $A \succeq B$ stands for $A - B$ being positive semidefinite. The interval $[x - a, x + a]$ is abbreviated as $[x \pm a]$. For two vectors $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$, let $\langle x, y \rangle \triangleq \sum_i x_i y_i$. A distribution π is called σ -subgaussian if $\mathbb{E}_\pi[e^{tX}] \leq \exp(t^2\sigma^2/2)$ for all $t \in \mathbb{R}$. We use standard big- O notations, e.g., for two positive sequence $\{a_n\}$ and $\{b_n\}$, $a_n = O(b_n)$ if $a_n \leq Cb_n$ for some constant $C > 0$; $a_n = \Omega(b_n)$ if $b_n = O(a_n)$; $a_n = \Theta(b_n)$ if $a_n = O(b_n)$ and $a_n = \Omega(b_n)$. We write $a_n = O_\beta(b_n)$ if C depends on another parameter β .

1.7 Organization

The paper is organized as follows. In Section 2 we provide some basic results of the theory of moments, polynomial interpolation, and the Wasserstein distance. In Section 3 we introduce the moment comparison theorems, which bound the Wasserstein distance between two discrete distributions in terms of the discrepancy of their moments. These are key results to prove our main theorems. In Section 4, we propose estimation algorithms and provide their statistical guarantees. Matching minimax lower bounds are given in Section 5. Section 6 contains numerical experiments and comparison with other methods such as the EM algorithm. Section 7 discusses extensions and open problems. Proofs are given in Section 8 and auxiliary results are deferred to Appendix C.

2 Preliminaries

2.1 Moment space, SDP characterization, and Gauss quadrature

The theory of moments plays a key role in the developments of analysis, probability, statistics, and optimization. See the classics [ST43, KS53] and the recent monographs [Las09, Sch17] for a detailed treatment. Below, we briefly review a few that are related to this paper.

The r^{th} moment vector of a distribution π is an r -tuple $\mathbf{m}_r(\pi) = (m_1(\pi), \dots, m_r(\pi))$. The r^{th} moment space on $K \subseteq \mathbb{R}$ is defined as

$$\mathcal{M}_r(K) = \{\mathbf{m}_r(\pi) : \pi \text{ is supported on } K\},$$

which is the convex hull of $\{(x, x^2, \dots, x^r) : x \in K\}$. This convex set satisfies many geometric constraints such as the Cauchy-Schwarz and Hölder inequalities. When $K = [a, b]$ is a compact interval, $\mathcal{M}_r([a, b])$ is completely described by (see [ST43, Theorem 3.1], and also [KS53, Las09])

$$\begin{cases} \mathbf{M}_{0,r} \succeq 0, & (a+b)\mathbf{M}_{1,r-1} \succeq ab\mathbf{M}_{0,r-2} + \mathbf{M}_{2,r}, & r \text{ even,} \\ b\mathbf{M}_{0,r-1} \succeq \mathbf{M}_{1,r} \succeq a\mathbf{M}_{0,r-1}, & & r \text{ odd,} \end{cases} \quad (14)$$

where $\mathbf{M}_{i,j}$ denotes the Hankel matrix with entries m_i, m_{i+1}, \dots, m_j :

$$\mathbf{M}_{i,j} = \begin{bmatrix} m_i & m_{i+1} & \cdots & m_{\frac{i+j}{2}} \\ m_{i+1} & m_{i+2} & \cdots & m_{\frac{i+j}{2}+1} \\ \vdots & \vdots & \ddots & \vdots \\ m_{\frac{i+j}{2}} & m_{\frac{i+j}{2}+1} & \cdots & m_j \end{bmatrix}.$$

Example 2 (Moment spaces on $[0, 1]$). $\mathcal{M}_2([0, 1])$ is simply described by $m_1 \geq m_2 \geq 0$ and $m_2 \geq m_1^2$. $\mathcal{M}_3([0, 1])$ is described by

$$\begin{bmatrix} 1 & m_1 \\ m_1 & m_2 \end{bmatrix} \succeq \begin{bmatrix} m_1 & m_2 \\ m_2 & m_3 \end{bmatrix} \succeq 0.$$

Using Sylvester's criterion (see [HJ12, Theorem 7.2.5]), they are equivalent to

$$\begin{aligned} 0 \leq m_1 \leq 1, \quad m_2 \geq m_3 \geq 0, \\ m_1 m_3 \geq m_2^2, \quad (1 - m_1)(m_2 - m_3) \geq (m_1 - m_2)^2. \end{aligned}$$

The necessity of the above inequalities are apparent: the first two follow from the range $[0, 1]$, and the last two follow from the Cauchy-Schwarz inequality. It turns out that they are also sufficient.

Moment matrices of discrete distributions satisfy more structural properties. For instances, the moment matrix of a k -atomic distribution of any order is of rank at most k , and is a deterministic function of \mathbf{m}_{2k-1} ; the number of atoms can be characterized using the determinants of moment matrices (see [Usp37, p. 362] or [Lin89, Theorem 2A]) by:

Theorem 5. *A sequence m_1, \dots, m_{2r} is the moments of a distribution with exactly r points of support if and only if $\det(\mathbf{M}_{r-1}) > 0$ and $\det(\mathbf{M}_r) = 0$.*

Gauss quadrature is a discrete approximation for a given distribution. Given π supported on $K \subseteq \mathbb{R}$, a k -point Gauss quadrature is a k -atomic distribution $\pi_k = \sum_{i=1}^k w_i \delta_{x_i}$, also supported on K , such that, for any polynomial P of degree at most $2k - 1$,

$$\mathbb{E}_\pi P = \mathbb{E}_{\pi_k} P = \sum_{i=1}^k w_i P(x_i). \quad (15)$$

It is known that, due to Christoffel [Chr77], Gauss quadrature always exists and is uniquely determined by $\mathbf{m}_{2k-1}(\pi)$. A basic algorithm to compute Gauss quadrature is Algorithm 5 provided in Appendix B and many algorithms with improved computational efficiency and numerical stability have been proposed; cf. [Gau04, Chapter 3].

2.2 Polynomial interpolation, majorization, and the Neville diagram

Given a function f and a set of distinct points (commonly referred to as nodes) $\{x_0, \dots, x_k\}$, there exists a unique polynomial P of degree k that coincides with f on every node. P is referred to as the interpolating polynomial and can be expressed in the *Lagrange form* as

$$P(x) = \sum_{i=0}^k f(x_i) \frac{\prod_{j \neq i} (x - x_j)}{\prod_{j \neq i} (x_i - x_j)}, \quad (16)$$


and alternatively, in the *Newton form* as

$$P(x) = a_0 + a_1(x - x_0) + \dots + a_k(x - x_0) \cdots (x - x_{k-1}). \quad (17)$$

The coefficients in (17) can be successively calculated using $a_0 = f(x_0)$, $a_0 + a_1(x_1 - x_0) = f(x_1)$, etc. In general they are divided differences $a_r = f[x_0, \dots, x_r]$ that can be recursively evaluated by

$$f[x_i] = f(x_i) \quad f[x_i, \dots, x_{i+r}] = \frac{f[x_{i+1}, \dots, x_{i+r}] - f[x_i, \dots, x_{i+r-1}]}{x_{i+r} - x_i}. \quad (18)$$

The above recursion can be calculated on Nevilles diagram:



In Neville's diagram, the r^{th} order divided differences are computed in the r^{th} column, and are determined by the previous column and the nodes. The coefficients in (17) are found in the top diagonal of the Neville's diagram. In the present paper Neville's diagram will be used to bound the coefficients in Newton formula (17).

Interpolating polynomials are the main tool to prove moment comparison theorems in Section 3. Specifically, we will interpolate step functions by polynomials in order to bound the difference of two CDFs via their moment difference. Therefore, it is critical to have a good control over the coefficients of the interpolating polynomial. To this end, it turns out the Newton form is more convenient to use than the Lagrange form because the former takes into account the cancellation between each terms in the polynomial. Indeed, in the Lagrange form (16), if two nodes are very close, then each term can be arbitrarily large, even if f itself is a smooth function. In contrast, each term of (17) is stable when f is smooth since divided differences are closely related to derivatives. The following example illustrates this point:

Example 3 (Lagrange versus Newton form). Given three points $x_1 = 0, x_2 = \epsilon, x_3 = 1$ with $f(x_1) = 1, f(x_2) = 1 + \epsilon, f(x_3) = 2$, the interpolating polynomial is $P(x) = x + 1$. The next equation gives the interpolating polynomial in Lagrange's and Newton's form respectively.

$$\text{Lagrange: } P(x) = \frac{(x - \epsilon)(x - 1)}{\epsilon} + (1 + \epsilon) \frac{x(x - 1)}{\epsilon(\epsilon - 1)} + 2 \frac{x(x - \epsilon)}{1 - \epsilon};$$

$$\text{Newton: } P(x) = 1 + x + 0.$$

Polynomial interpolation can be generalized to incorporate information on derivatives, known as *Hermite interpolation*. Formally, given a function f and distinct nodes $x_0 < x_1 < \dots < x_m$, there exists a unique polynomial P of degree k satisfying $P^{(j)}(x_i) = f^{(j)}(x_i)$ for $i = 0, \dots, m$ and $j = 0, \dots, k_i - 1$, where $k + 1 = \sum_i k_i$. Analogous to the Lagrange formula (16), P can be explicitly constructed with the help of the generalized Lagrange polynomials, and an explicit formula is given in [SB02, pp. 52–53]. Newton formula (17) can also be extended by using generalized divided differences with

$$f[x_i, \dots, x_{i+r}] = \frac{f^{(r)}(x_i)}{r!}, \quad x_i = x_{i+1} = \dots = x_{i+r}, \quad (19)$$

To this end, we define an expanded sequence by replacing each x_i by k_i copies of itself:

$$\underbrace{x_0 = \dots = x_0}_{k_0} < \underbrace{x_1 = \dots = x_1}_{k_1} < \dots < \underbrace{x_m = \dots = x_m}_{k_m}. \quad (20)$$

The Hermite interpolating polynomial is obtained by (17) using this new sequence and generalized divided differences, which can also be calculated from the Nevilles diagram by replacing differences by derivatives whenever we have repeated nodes. Below we give an example using Hermite interpolation to construct polynomial majorization, which will be used to prove moment comparison theorems in Section 3.

Example 4 (Hermite interpolation as polynomial majorization). Let $f(x) = \mathbf{1}_{\{x \leq 0\}}$. We want to find a polynomial majorization $P \geq f$ such that $P(x) = f(x)$ on $x = \pm 1$. To this end we interpolate f on $\{-1, 0, 1\}$ with the following:

x	-1	0	1
$P(x)$	1	1	0
$P'(x)$	0	any	0

The Hermite interpolation P is of degree four and majorizes f [Akh65, p. 65]. To see this, we note that $P'(\xi) = 0$ for some $\xi \in (-1, 0)$ by Rolle's theorem. Since $P'(-1) = P'(1) = 0$, P has no other stationary point than $-1, \xi, 1$, and thus decreases monotonically in $(\xi, 1)$. Hence, $-1, 1$ are the only local minimum points of P , and thus $P \geq f$ everywhere. The polynomial P is shown in Fig. 1(b).

To explicitly construct the polynomial, we have an expanded sequence $-1, -1, 0, 1, 1$ by (20). Applying Newton formula (17) with generalized divided differences from the Neville's diagram Fig. 1(a), we obtain that $P(x) = 1 - \frac{1}{4}x(x + 1)^2 + \frac{1}{2}x(x + 1)^2(x - 1)$.

2.3 Wasserstein distance

Wasserstein distance measures the cost of optimal transportation, whose development has broad impact in economics and motivates the study of linear programming. In this paper, we mainly

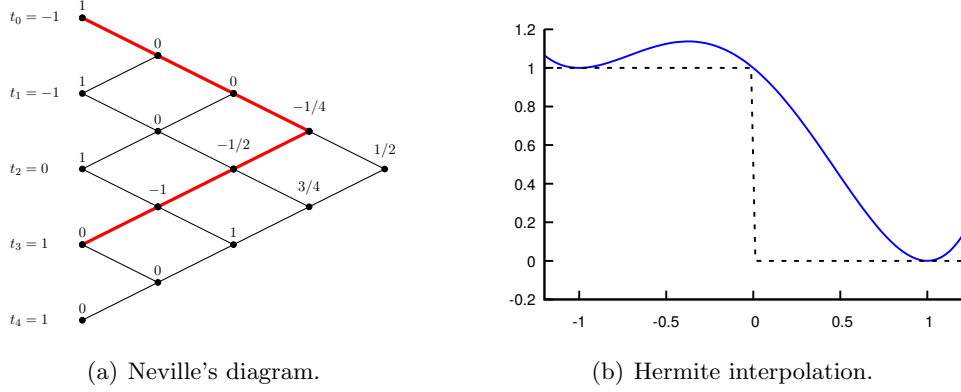


Figure 1: Neville's diagram and Hermite interpolation. In (a), values are recursively calculated from left to right. For example, the red thick line shows that $f[-1, -1, 0, 1]$ is calculated by $\frac{-1/2-0}{1-(-1)} = -1/4$.

focus on the 1-Wasserstein distance defined in (6), which can be equivalently expressed, through the Kantorovich duality (see [Vil03]), as

$$W_1(\nu, \nu') = \sup\{\mathbb{E}_\nu[\varphi] - \mathbb{E}_{\nu'}[\varphi] : \varphi \text{ is 1-Lipschitz}\}. \quad (21)$$

The optimal coupling in the definition (6) has many equivalent characterization [Vil08] but is often difficult to compute analytically in general. Nevertheless, the situation is especially simple for distributions on the real line, where the quantile coupling is optimal, and we have

$$W_1(\nu, \nu') = \int |F_\nu(t) - F_{\nu'}(t)| dt, \quad (22)$$

where F_ν and $F_{\nu'}$ denote the CDFs of ν and ν' , respectively. Both (21) and (22) provide convenient characterizations to bound the Wasserstein distance in Section 3.

As previously mentioned in Section 1.4, two discrete distributions close in the Wasserstein distance have similar support sets and weights. This is made precise in Lemma 2 and 3 next:

Lemma 2. *Suppose ν and ν' are discrete distributions supported on S and S' , respectively. Let $\epsilon = \min\{\nu(x) : x \in S\} \wedge \min\{\nu'(x) : x \in S'\}$. Then,*

$$d_H(S, S') \leq W_1(\nu, \nu')/\epsilon,$$

where d_H denotes the Hausdorff distance defined as

$$d_H(S, S') = \max\left\{\sup_{x \in S} \inf_{x' \in S'} |x - x'|, \sup_{x' \in S'} \inf_{x \in S} |x - x'|\right\}. \quad (23)$$

Proof. For any coupling P_{XY} such that $X \sim \nu$ and $Y \sim \nu'$,

$$\mathbb{E}|X - Y| = \sum_x \mathbb{P}[X = x] \mathbb{E}[|X - Y| | X = x] \geq \sum_x \epsilon \cdot \inf_{x' \in S'} |x - x'| \geq \epsilon \cdot \sup_{x \in S} \inf_{x' \in S'} d(x, x').$$

Interchanging X and Y completes the proof. \square

Lemma 3. *For any $\delta > 0$,*

$$\nu(x) - \nu'([x \pm \delta]) \leq W_1(\nu, \nu')/\delta, \quad \nu'(x) - \nu([x \pm \delta]) \leq W_1(\nu, \nu')/\delta.$$

Proof. Using the optimal coupling P_{XY}^* such that $X \sim \nu$ and $Y \sim \nu'$, applying Markov inequality yields that

$$\mathbb{P}[|X - Y| > \delta] \leq \mathbb{E}|X - Y|/\delta = W_1(\nu, \nu')/\delta.$$

By Strassen's theorem (see [Vil03, Corollary 1.28]), for any Borel set B , we have $\nu(B) \leq \nu'(B^\delta) + W_1(\nu, \nu')/\delta$ and $\nu'(B) \leq \nu(B^\delta) + W_1(\nu, \nu')/\delta$, where $B^\delta \triangleq \{x : \inf_{y \in B} |x - y| \leq \delta\}$ denotes the δ -fattening of B . The conclusion follows by considering a singleton $B = \{x\}$. \square

Lemma 2 and 3 together yield a bound on the parameter estimation error (up to a permutation) in terms of the Wasserstein distance, which was previously given in Lemma 1:

Proof. Denote the support sets of ν and ν' by $S = \{\mu_1, \dots, \mu_k\}$ and $S' = \{\hat{\mu}_1, \dots, \hat{\mu}_k\}$, respectively. Applying Lemma 2 yields that $d_H(S, S') < \epsilon/\epsilon_2$, which is less than $\epsilon_1/4$ by the assumption $\epsilon < \epsilon_1\epsilon_2/4$. Since $|\mu_i - \mu_j| \geq \epsilon$ for every $i \neq j$, then there exists a permutation Π such that

$$\|\mu - \Pi\hat{\mu}\|_\infty < \epsilon/\epsilon_2.$$

Applying Lemma 3 twice with $\delta = \epsilon/2$, $x = \mu_i$ and $x = (\Pi\hat{\mu})_i$, respectively, we obtain the desired

$$w_i - (\Pi\hat{w})_i \leq 2\epsilon/\epsilon_1, \quad (\Pi\hat{w})_i - w_i \leq 2\epsilon/\epsilon_1. \quad \square$$

3 Optimal transport and moment comparison theorems

A discrete distribution with k atoms has $2k - 1$ free parameters. Therefore it is reasonable to expect that it can be uniquely determined by its first $2k - 1$ moments. Indeed, we have the following simple identifiability results for discrete distributions:

Lemma 4. *Let ν and ν' be distributions on the real line.*

1. *If ν and ν' are both k -atomic, then $\nu = \nu'$ if and only if $\mathbf{m}_{2k-1}(\nu) = \mathbf{m}_{2k-1}(\nu')$.*
2. *If ν is k -atomic, then $\nu = \nu'$ if and only if $\mathbf{m}_{2k}(\nu) = \mathbf{m}_{2k}(\nu')$.*

In the context of statistical estimation, we only have access to samples and noisy estimates of moments. To solve the inverse problems from moments to distributions, our theory relies on the following stable version of the identifiability in Lemma 4, which show that closeness of moments implies closeness of distributions in Wasserstein distance. In the sequel we refer to Propositions 1 and 2 as moment comparison theorems.

Proposition 1. *Let ν and ν' be k -atomic distributions supported on $[-1, 1]$. If $|m_i(\nu) - m_i(\nu')| \leq \delta$ for $i = 1, \dots, 2k - 1$, then*

$$W_1(\nu, \nu') \leq O\left(k\delta^{\frac{1}{2k-1}}\right).$$

Proposition 2. *Let ν be a k -atomic distribution supported on $[-1, 1]$. If $|m_i(\nu) - m_i(\nu')| \leq \delta$ for $i = 1, \dots, 2k$, then*

$$W_1(\nu, \nu') \leq O\left(k\delta^{\frac{1}{2k}}\right).$$

Remark 1. The exponents in Proposition 1 and 2 are optimal. To see this, we first note that the number of moments needed for identifiability in Lemma 4 cannot be reduced:

1. Given any $2k$ distinct points, there exist two k -atomic distributions with disjoint support sets but identical first $2k - 2$ moments (see Lemma 29).

2. Given any continuous distribution, its k -point Gauss quadrature is k -atomic and have identical first $2k - 1$ moments (see Section 2.1).

By the first observation, there exists two k -atomic distributions ν and ν' such that

$$m_i(\nu) = m_i(\nu'), \quad i = 1, \dots, 2k - 2, \quad |m_{2k-1}(\nu) - m_{2k-1}(\nu')| = c_k, \quad W_1(\nu, \nu') = d_k,$$

where c_k and d_k are strictly positive constants that depend on k . Let $\tilde{\nu}$ and $\tilde{\nu}'$ denote the distributions of ϵX and $\epsilon X'$ such that $X \sim \nu$ and $X' \sim \nu'$, respectively. Then, we have

$$\max_{i \in [2k-1]} |m_i(\tilde{\nu}) - m_i(\tilde{\nu}')| = \epsilon^{2k-1} c_k, \quad W_1(\tilde{\nu}, \tilde{\nu}') = \epsilon d_k.$$

This concludes the tightness of the exponent in Proposition 1. Similarly, the exponent in Proposition 2 is also tight using the second observation.

Remark 2. Classical moments comparison theorems aim to show convergence of distributions by comparing a *growing* number of moments. For example, Chebyshev's theorem states if $\mathbf{m}_r(\pi) = \mathbf{m}_r(N(0, 1))$, then (see [Dia87, Theorem 2])

$$\sup_{x \in \mathbb{R}} |F_\pi(x) - \Phi(x)| \leq \sqrt{\frac{\pi}{2r}},$$

where F_π and Φ denote the CDFs of π and $N(0, 1)$, respectively. For two compactly supported distributions, the above estimate can be sharpened to $O(\frac{\log r}{r})$ [Kra32]. In contrast, in the context of estimating finite mixtures we are dealing with discrete mixing distributions, which can be identified by a *fixed* number of moments. However, with finitely many samples, it is impossible to exactly determine the moments, and measuring the error in the KS distance is too much to ask (see Section 1.4). It turns out that W_1 -distance is a suitable metric for this purpose, and the closeness of moments does imply the closeness of distribution in the W_1 distance, which is the integrated difference (L_1 -distance) between two CDFs as opposed the uniform error (L_∞ -distance).

4 Estimators and statistical guarantees

4.1 Known variance

The denoised method of moments for estimating Gaussian location mixture models (2) with known variance parameter σ^2 consists of three main steps:

1. estimate $\mathbf{m}_{2k-1}(\nu)$ by $\tilde{\mathbf{m}} = (\tilde{m}_1, \dots, \tilde{m}_{2k-1})$ (using Hermite polynomials);
2. denoise $\tilde{\mathbf{m}}$ by its projection $\hat{\mathbf{m}}$ to the moment space (semidefinite programming);
3. find a k -atomic distribution $\hat{\nu}$ such that $\mathbf{m}_{2k-1}(\hat{\nu}) = \hat{\mathbf{m}}$ (Gauss quadrature).

The complete algorithm is summarized in Algorithm 1.

We estimate the moments of the mixing distribution in lines 1 to 4. The unique unbiased estimators for the polynomials of the mean parameter in a Gaussian location model are Hermite polynomials

$$H_r(x) = r! \sum_{j=0}^{\lfloor r/2 \rfloor} \frac{(-1/2)^j}{j!(r-2j)!} x^{r-2j}, \quad (25)$$

Algorithm 1 Denoised method of moments with known variance.

Input: n independent samples X_1, \dots, X_n , order k , variance σ^2 , interval $I = [a, b]$.

Output: estimated mixing distribution.

- 1: **for** $r = 1$ **to** $2k - 1$ **do**
- 2: $\hat{\gamma}_r = \frac{1}{n} \sum_i X_i^r$
- 3: $\tilde{m}_r = r! \sum_{i=0}^{\lfloor r/2 \rfloor} \frac{(-1/2)^i}{i!(r-2i)!} \hat{\gamma}_{r-2i} \sigma^{2i}$
- 4: **end for**
- 5: Let \hat{m} be the optimal solution of the following:

$$\min\{\|\tilde{m} - \hat{m}\| : \hat{m} \text{ satisfies (14)}\}, \quad (24)$$

where $\tilde{m} = (\tilde{m}_1, \dots, \tilde{m}_{2k-1})$.

- 6: Report the outcome of Gauss quadrature (Algorithm 5) with input \hat{m} .
-

such that $\mathbb{E}H_r(X) = \mu^r$ when $X \sim N(\mu, 1)$. Thus, if we define

$$\gamma_r(x, \sigma) = \sigma^r H_r(x/\sigma) = r! \sum_{j=0}^{\lfloor r/2 \rfloor} \frac{(-1/2)^j}{j!(r-2j)!} \sigma^{2j} x^{r-2j}, \quad (26)$$

then $\mathbb{E}\gamma_r(X, \sigma) = \mu^r$ when $X \sim N(\mu, \sigma^2)$. Hence, \tilde{m}_r is an unbiased estimate of $m_r(\nu)$. The variance of \tilde{m}_r is analyzed in the following lemma:

Lemma 5. *If $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} \nu * N(0, \sigma^2)$ and ν is supported on $[-M, M]$, then*

$$\text{var}[\tilde{m}_r] \leq \frac{1}{n} (O(M + \sigma\sqrt{r}))^{2r}.$$

As observed in Section 1.2, the major reason for the failure of the usual method of moments is that the unbiased estimate \tilde{m} needs not constitute a legitimate moment sequence, despite the consistency of each \tilde{m}_i . To resolve this issue, we project \tilde{m} to the moment space using (24). As explained in Section 2.1, (14) consists of positive semidefinite constraints, and thus the optimal solution of (24) can be obtained by semidefinite programming (SDP).³ In fact, it suffices to solve a *feasibility* program and find any valid moment vector \hat{m} that is within the desired $\frac{1}{\sqrt{n}}$ statistical accuracy.

Now that \hat{m} is indeed a valid moment sequence, we use the Gauss quadrature introduced in Section 2.1 (see Algorithm 5 in Appendix B) to find the unique k -atomic distribution $\hat{\nu}$ such that $\mathbf{m}_{2k-1}(\hat{\nu}) = \hat{m}$. Using Algorithm 1, \tilde{m} is computed in $O(kn)$ time, the semidefinite programming is solvable in $O(k^{6.5})$ time using the interior-point method (see [WSV12]), and the Gauss quadrature can be evaluated in $O(k^3)$ time [GW69]. In view of the global assumption (7), Algorithm 1 can be executed in $O(kn)$ time.

We now prove the statistical guarantee (8) for the DMM estimator previously announced in Theorem 1:

Proof. By scaling it suffices consider $M = 1$. We use Algorithm 1 with Euclidean norm in (24). Using the variance of \tilde{m} in Lemma 5 and Chebyshev inequality yield that, for each $r = 1, \dots, 2k-1$, with probability $1 - \frac{1}{8k}$,

$$|\tilde{m}_r - m_r(\nu)| \leq \sqrt{k/n} (c\sqrt{r})^r, \quad (27)$$

³ The formulation (24) with Euclidean norm can already be implemented in popular modeling languages for convex optimization problem such as CVXPY [DB16]. A standard form of SDP is given in Appendix A.

for some absolute constant c . By the union bound, with probability $3/4$, (27) holds simultaneously for every $r = 1, \dots, 2k - 1$, and thus

$$\|\tilde{m} - \mathbf{m}_{2k-1}(\nu)\|_2 \leq \epsilon, \quad \epsilon \triangleq \frac{(\sqrt{ck})^{2k+1}}{\sqrt{n}}.$$

Since $\mathbf{m}_{2k-1}(\nu)$ satisfies (14) and thus is one feasible solution for (24), we have $\|\tilde{m} - \hat{m}\|_2 \leq \epsilon$. Note that $\hat{m} = \mathbf{m}_{2k-1}(\hat{\nu})$. Hence, by triangle inequality, we obtain the following statistical accuracy:

$$\|\mathbf{m}_{2k-1}(\hat{\nu}) - \mathbf{m}_{2k-1}(\nu)\|_2 \leq \epsilon, \quad (28)$$

Applying Proposition 1 yields that, with probability $3/4$,

$$W_1(\hat{\nu}, \nu) \leq O\left(k^{1.5}n^{-\frac{1}{4k-2}}\right).$$

The confidence $1 - \delta$ in (8) can be obtained if we estimate \tilde{m}_r by taking the median of $T = \log \frac{2k}{\delta}$ estimates, where each uses a batch of n/T samples. Then Hoeffding's inequality and the union bound imply that, with probability $1 - \delta$,

$$|\tilde{m}_r - m_r(\nu)| \leq \sqrt{\frac{\log(2k/\delta)}{n}} (c\sqrt{r})^r, \quad \forall r = 1, \dots, 2k - 1, \quad (29)$$

and the conclusion follows. \square

To conclude this subsection, we discuss the connection to the Generalized Method of Moment (GMM), introduced by Hansen [Han82, Hal05]. Instead of solving the moment equations, GMM aims to minimize the difference between estimated and fitted moments:

$$Q(\theta) = (\hat{m} - m(\theta))^\top W(\hat{m} - m(\theta)), \quad (30)$$

where \hat{m} is the estimated moment, θ is the fitted model parameter, and W is a positive semidefinite weighting matrix. The minimizer of $Q(\theta)$ serves as the GMM estimate for the unknown model parameter θ_0 . The objective function Q in θ is rightly nonconvex in general, notably under the Gaussian mixture model with θ corresponding to the unknown means and weights, which is hard to optimize. Note that (24) with Euclidean norm is equivalent to GMM with the identity weighting matrix. The optimal weighting matrix W^* is the inverse of $\lim_{n \rightarrow \infty} \text{cov}[\sqrt{n}(\hat{m} - m(\theta_0))]$ but involving the unknown model parameters θ_0 . A popular approach is a two-step estimator [Hal05]:

1. a suboptimal weighting matrix, e.g., identity matrix, is used to obtain a consistent estimate of θ_0 , and then a consistent estimate \hat{W} for W^* ;
2. θ_0 is re-estimated using weighting matrix \hat{W} .

The above two-step approach can be similarly implemented in our denoised method of moments.

4.2 Unknown variance

When the variance parameter σ^2 is unknown, unbiased estimator for the moments of the mixing distribution no longer exists (see Lemma 30). It is not difficult to consistently estimate the variance⁴, then plug into the DMM estimator in Section 4.1 to obtain a consistent estimate of the

⁴ For instance, the simple estimator $\hat{\sigma} = \frac{\max_i X_i}{\sqrt{2 \log n}}$ satisfies $|\sigma - \hat{\sigma}| = O_P(\log n)^{-\frac{1}{2}}$.

mixing distribution ν ; however, the convergence rate is far from optimal. In fact, to achieve the optimal rate it is crucial to simultaneously estimate both the means and the variance parameters. To this end, again we take a moment-based approach. The following result provides a guarantee for any joint estimate of both the mixing distribution and the variance parameter in terms of its moments accuracy.

Proposition 3. *Let*

$$\pi = \nu * N(0, \sigma^2), \quad \hat{\pi} = \hat{\nu} * N(0, \hat{\sigma}^2),$$

where $\nu, \hat{\nu}$ are k -atomic distributions supported on $[-M, M]$, and $\sigma, \hat{\sigma}$ are bounded. If $|m_r(\pi) - m_r(\hat{\pi})| \leq \epsilon$ for $r = 1, \dots, 2k$, then

$$|\sigma^2 - \hat{\sigma}^2| \leq O(M^2 \epsilon^{\frac{1}{k}}), \quad W_1(\nu, \hat{\nu}) \leq O(Mk^{1.5} \epsilon^{\frac{1}{2k}}).$$

To apply Proposition 3, we can solve the method of moments equations, namely, find a k -atomic distribution $\hat{\nu}$ and $\hat{\sigma}^2$ such that

$$\mathbb{E}_n[X^r] = \mathbb{E}_{\hat{\pi}}[X^r], \quad r = 1, \dots, 2k \quad (31)$$

where $\hat{\pi} = \hat{\mu} * N(0, \hat{\sigma}^2)$ is the fitted Gaussian mixture. Here both the number of equations and the number of variables are equal to $2k$. Suppose (31) has a solution $(\hat{\mu}, \hat{\sigma})$. Then applying Proposition 3 with $\delta = O_k(\frac{1}{\sqrt{n}})$ yields the result in Theorem 1, which is minimax optimal (see Section 5). In stark contrast to the known σ case, where we have shown in Section 1.2 that the vanilla method of moments equation can have no solution unless we denoise by projection to the moment space, here with one extra scale parameter σ , one can show that (31) has a solution with probability one!⁵ Furthermore, an efficient method of finding a solution to (31) is due to Lindsay [Lin89] and summarized in Algorithm 2. Indeed, the sample moments are computable in $O(kn)$ time, and the smallest non-negative root of the polynomial of degree $k(k+1)$ can be found in $O(k^2)$ time using Newton's method (see [Atk08]).

In [Lin89] the consistency of this estimator was proved under extra assumptions. In fact we will that it unconditionally achieves the minimax optimal rate (9) and (10) previously accounced in Theorem 1. In this section we show that Lindsay's algorithm chooses a $\hat{\sigma}$ so that the corresponding the moment estimates lie in the moment space with probability one. In this sense, although no explicit projection is involved, the noisy estimates are *implicitly* denoised.

We first describe the intuition of choosing $\hat{\sigma}$ in Lindsay's algorithm, i.e., line 6 of Algorithm 2. Let $X \sim \nu * N(0, \sigma^2)$. For any $\sigma' \leq \sigma$, we have

$$\mathbb{E}[\gamma_j(X, \sigma')] = m_j(\nu * N(0, \sigma^2 - \sigma'^2)).$$

Let $d_k(\sigma')$ denote determinant of moment matrix $\{\mathbb{E}[\gamma_{i+j}(X, \sigma')]\}_{i,j=0}^k$, an even polynomial in σ' of degree $k(k+1)$. According to Theorem 5, $d_k(\sigma') > 0$ when $0 \leq \sigma' < \sigma$ and becomes zero at $\sigma' = \sigma$, and thus σ is characterized by the smallest positive zero of d_k . In lines 5 – 6, d_k is estimated by \hat{d}_k using the empirical moments, and σ is estimated by the smallest positive zero of \hat{d}_k . We first note that \hat{d}_k indeed has a positive zero:

⁵ It is possible that the equation (31) has no solution, for instance, when $k = 2, n = 7$ and the empirical distribution is $\pi_7 = \frac{1}{7}\delta_{-\sqrt{7}} + \frac{1}{7}\delta_{\sqrt{7}} + \frac{5}{7}\delta_0$. The first four empirical moments are $\mathbf{m}_4(\pi_7) = (0, 2, 0, 14)$, which cannot be realized by any two-component Gaussian mixture (1). Indeed, suppose $\hat{\pi} = w_1 N(\mu_1, \sigma^2) + (1 - w_1)N(\mu_2, \sigma^2)$ is a solution to (31). Eliminating variables leads to the contradiction that $2\mu_1^4 + 2 = 0$. Assuringly, as we will show later in Lemma 7, such cases occur with probability zero.

Algorithm 2 Lindsay’s estimator for normal mixtures with an unknown common variance

Input: n samples X_1, \dots, X_n .

Output: estimated mixing distribution $\hat{\nu}$, and estimated variance $\hat{\sigma}^2$.

- 1: **for** $r = 1$ **to** $2k$ **do**
 - 2: $\hat{\gamma}_r = \frac{1}{n} \sum_i X_i^r$
 - 3: $\hat{m}_r(\sigma) = r! \sum_{i=0}^{\lfloor r/2 \rfloor} \frac{(-1/2)^i}{i!(r-2i)!} \hat{\gamma}_{r-2i} \sigma^{2i}$
 - 4: **end for**
 - 5: Let $\hat{d}_k(\sigma)$ be the determinant of the matrix $\{\hat{m}_{i+j}(\sigma)\}_{i,j=0}^k$.
 - 6: Let $\hat{\sigma}$ be the smallest positive root of $\hat{d}_k(\sigma) = 0$.
 - 7: **for** $r = 1$ **to** $2k$ **do**
 - 8: $\hat{m}_r = \hat{m}_r(\hat{\sigma})$
 - 9: **end for**
 - 10: Let $\hat{\nu}$ be the outcome of the Gauss quadrature (Algorithm 5) with input $\hat{m}_1, \dots, \hat{m}_{2k-1}$
 - 11: Report $\hat{\nu}$ and $\hat{\sigma}^2$.
-

Lemma 6. *Assume $n > k$ and the mixture distribution has a density. Then, almost surely, \hat{d}_k has a positive root within $(0, s]$, where $s^2 \triangleq \frac{1}{n} \sum_{i=1}^n (X_i - \mathbb{E}_n[X])^2$ denotes the sample variance.*

The next result shows that, with the above choice of $\hat{\sigma}$, the moment estimates $\hat{m}_j = \mathbb{E}_n[\gamma_j(X, \hat{\sigma})]$ for $j = 1, \dots, 2k$ given in line 8 are implicitly denoised and lie in the moment space with probability one. Thus (31) has a solution, and the estimated mixing distribution $\hat{\nu}$ can be found by the Gauss quadrature. This result was previously shown in [Lin89] but under extra conditions.

Lemma 7. *Assume $n \geq 2k - 1$ and the mixture distribution has a density. Then, almost surely, there exists a k -atomic distribution $\hat{\nu}$ such that $m_j(\hat{\nu}) = \hat{m}_j$ for $j \leq 2k$, where \hat{m}_j is from Algorithm 2.*

With the above analysis, we now prove the statistical guarantee (9) and (10) for Lindsay’s algorithm announced in Theorem 1:

Proof. It suffices to consider $M = 1$. Let $\hat{\pi} = \hat{\nu} * N(0, \hat{\sigma}^2)$ and $\pi = \nu * N(0, \sigma^2)$ denote the estimated mixture distribution and the ground truth, respectively. Let $\hat{m}_r = \mathbb{E}_n[X^r]$ and $m_r = m_r(\pi)$. The variance of \hat{m}_r is upper bounded by

$$\text{var}[\hat{m}_r] = \frac{1}{n} \text{var}[X_1^r] \leq \frac{1}{n} \mathbb{E}[X^{2r}] \leq \frac{(\sqrt{cr})^{2r}}{n},$$

for some absolute constant c . Using Chebyshev inequality, for each $r = 1, \dots, 2k$, with probability $1 - \frac{1}{8k}$, we have,

$$|\hat{m}_r - m_r| \leq (\sqrt{cr})^r \sqrt{k/n}. \quad (32)$$

By the union bound, with probability $3/4$, the above holds simultaneously for every $r = 1, \dots, 2k$. It follows from Lemma 6 and 7 that (31) holds with probability one. Therefore,

$$|m_r(\hat{\pi}) - m_r(\pi)| \leq (\sqrt{cr})^r \sqrt{k/n}, \quad r = 1, \dots, 2k.$$

for some absolute constant c . In the following, the error of variance estimate is denoted by $\tau^2 = |\sigma^2 - \hat{\sigma}^2|$.

- If $\sigma \leq \hat{\sigma}$, let $\nu' = \hat{\nu} * N(0, \tau^2)$. Using $\mathbb{E}_\pi[\gamma_r(X, \sigma)] = m_r(\nu)$ and $\mathbb{E}_{\hat{\pi}}[\gamma_r(X, \sigma)] = m_r(\nu')$, where γ_r is the Hermite polynomial (26), we obtain that (see Lemma 26)

$$|m_r(\nu') - m_r(\nu)| \leq (\sqrt{c'k})^{2k} \sqrt{k/n}, \quad r = 1, \dots, 2k, \quad (33)$$

for an absolute constant c' . Applying Proposition 3 yields that,

$$|\sigma^2 - \hat{\sigma}^2| \leq O(kn^{-\frac{1}{2k}}), \quad W_1(\nu, \hat{\nu}) \leq O(k^2 n^{-\frac{1}{4k}}).$$

- If $\sigma \geq \hat{\sigma}$, let $\nu' = \nu * N(0, \tau^2)$. Similar to (33), we have

$$|m_r(\hat{\nu}) - m_r(\nu')| \leq (\sqrt{c'k})^{2k} \sqrt{k/n} \triangleq \epsilon, \quad r = 1, \dots, 2k.$$

To apply Proposition 3, we also need to ensure that $\hat{\nu}$ has a bounded support, which is not obvious. To circumvent this issue, we apply a truncation argument thanks to the following tail probability bound for $\hat{\nu}$ (see Lemma 16):

$$\mathbb{P}[|\hat{U}| \geq \sqrt{c_0 k}] \leq \epsilon(\sqrt{c_1 k}/t)^{2k}, \quad \hat{U} \sim \hat{\nu}, \quad (34)$$

for absolute constants c and c' . To this end, consider $\tilde{U} = \hat{U} \mathbf{1}_{\{|\hat{U}| \leq \sqrt{c_0 k}\}} \sim \tilde{\nu}$. Note that \tilde{U} is k -atomic supported on $[-\sqrt{c_0 k}, \sqrt{c_0 k}]$, we have $W_1(\nu, \tilde{\nu}) \leq \epsilon e^{O(k)}$ and $|m_r(\tilde{\nu}) - m_r(\hat{\nu})| \leq k\epsilon(c_1 k)^k$ for $r = 1, \dots, 2k$. Using the triangle inequality yields that

$$|m_r(\tilde{\nu}) - m_r(\nu')| \leq \epsilon + k\epsilon(c_1 k)^k.$$

Now we apply Proposition 3 with $\tilde{\nu}$ and $\nu * N(0, \tau^2)$ where both $\tilde{\nu}$ and ν are k -atomic supported on $[-\sqrt{c_0 k}, \sqrt{c_0 k}]$. In the case $\tilde{\nu}$ is discrete, the dependence on k in Proposition 3 can be improved (by improving (65) in the proof) and we obtain that

$$|\sigma^2 - \hat{\sigma}^2| \leq O(kn^{-\frac{1}{2k}}), \quad W_1(\nu, \tilde{\nu}) \leq O(k^2 n^{-\frac{1}{4k}}).$$

Using $k \leq O(\frac{\log n}{\log \log n})$, we also obtain $W_1(\nu, \hat{\nu}) \leq O(k^2 n^{-\frac{1}{2k}})$ by the triangle inequality.

To obtain a confidence $1 - \delta$ in (9) and (10), we can replace the empirical moments \hat{m}_r by the median of $T = \log \frac{1}{\delta}$ independent estimates similar to (29). \square

4.3 Adaptive results

In sections 4.1 and 4.2, we proved the statistical guarantees of our estimators under the worst-case scenario where the means can be arbitrarily close. Under separation conditions on the means (see Definition 1), our estimators automatically achieve a strictly better accuracy than the one claimed in Theorem 1. The goal in this subsection is to show those adaptive results. The key is the following adaptive version of the moment comparison theorems (cf. Propositions 1 and 2):

Proposition 4. *Suppose both ν and ν' are supported on a set of ℓ atoms in $[-1, 1]$, and each atom is at least γ away from all but at most ℓ' other atoms. Let $\delta = \max_{i \in [\ell-1]} |m_i(\nu) - m_i(\nu')|$. Then,*

$$W_1(\nu, \nu') \leq \ell \left(\frac{\ell 4^{\ell-1} \delta}{\gamma^{\ell-\ell'-1}} \right)^{\frac{1}{\ell'}}.$$

Proposition 5. *Suppose ν is supported on k atoms in $[-1, 1]$ and any $t \in \mathbb{R}$ is at least γ away from all but k' atoms. Let $\delta = \max_{i \in [2k]} |m_i(\nu) - m_i(\nu')|$. Then,*

$$W_1(\nu, \nu') \leq 8k \left(\frac{k4^{2k}\delta}{\gamma^{2(k-k')}} \right)^{\frac{1}{2k'}}.$$

The adaptive result (11) in the known variance parameter case is obtained using Proposition 4 in place of Proposition 1. To deal with unknown variance parameter case, using Proposition 5, we first show the following adaptive version of Proposition 3:

Proposition 6. *Under the conditions of Proposition 3, if both Gaussian mixtures both have k_0 γ -separated clusters in the sense of Definition 1, then,*

$$\sqrt{|\sigma^2 - \hat{\sigma}^2|}, W_1(\nu, \hat{\nu}) \leq O_k \left(\left(\frac{\epsilon}{\gamma^{2(k_0-1)}} \right)^{\frac{1}{2(k-k_0+1)}} \right).$$

Using these propositions, we now prove the adaptive rate of the denoised method of moments previously announced in Theorem 2:

Proof of Theorem 2. By scaling it suffices to consider $M = 1$. Recall that the Gaussian mixture is assumed to have k_0 (γ, ω) -separated clusters in the sense of Definition 1, that is, there exists a partition S_1, \dots, S_{k_0} of $[k]$ such that $|\mu_i - \mu_{i'}| \geq \gamma$ for any $i \in S_\ell$ and $i' \in S_{\ell'}$ such that $\ell \neq \ell'$, and $\sum_{i \in S_\ell} w_i \geq \omega$ for each ℓ .

Let $\hat{\nu}$ be the estimated mixing distribution which satisfies $W_1(\nu, \hat{\nu}) \leq \epsilon$ by Theorem 1. Since $\gamma\omega \geq C\epsilon$ by assumption, for each S_ℓ , there exists $i \in S_\ell$ such that μ_i is within distance $c\gamma$, where $c = 1/C$, to some atom of $\hat{\nu}$. Therefore, the estimated mixing distribution $\hat{\nu}$ has k_0 $(1 - 2c)\gamma$ -separated clusters. Denote the union of the support sets of ν and $\hat{\nu}$ by \mathcal{S} .

- When σ is known, each atom in \mathcal{S} is $\Omega(\gamma)$ away from at least $2(k_0 - 1)$ other atoms. Then (11) follows from Proposition 4 with $\ell = 2k$ and $\ell' = (2k - 1) - 2(k_0 - 1)$.
- When σ is unknown, (12) follows from a similar proof of (9) and (10) with Proposition 3 replaced by Proposition 6. \square

The rate in (11) as well as its optimality is previously obtained in [HK15], but their minimum-distance estimator is computationally expensive. Finally, we note that if one only assumes the separation condition but not the lower bound on the weights, we can obtain an intermediate result that is stronger than (8) but weaker than (11).

Theorem 6. *Under the conditions of Theorem 1, suppose σ is known and the Gaussian mixture has k_0 γ -separated clusters. Then, with probability at least $1 - \delta$,*

$$W_1(\nu, \hat{\nu}) \leq O_k \left(M\gamma^{-\frac{k_0-1}{2k-k_0}} \left(\frac{n}{\log(k/\delta)} \right)^{-\frac{1}{4k-2k_0}} \right). \quad (35)$$

4.4 Unbounded means

In the previous subsections, we assume that the means lie in a bounded interval. In the unbounded case, it is in fact impossible to estimate the mixing distribution under the Wasserstein distance⁶.

⁶ Let $\pi_\epsilon = \frac{1+\epsilon}{2}\delta_0 + \frac{1-\epsilon}{2}\delta_M$. Then $W_1(\pi_0, \pi_\epsilon) = M\epsilon$, but $D(\pi_0 \| \pi_\epsilon) \leq O(\epsilon^2)$. Choosing $\epsilon = o(1/\sqrt{n})$ and $M \gg 1/\epsilon$ leads to arbitrarily large estimation error.

Nevertheless, provided that the weights are bounded away from zero, it is possible to estimate the support set of the mixing distribution with respect to the Hausdorff distance (cf. (23)). This is the goal of this subsection.

In the unbounded case, blindly applying the previous moment-based methods does not work, because the estimated moments suffer from large variance due to the wide range of values of the means (cf. Lemma 5). To resolve this issue, we shall apply the *divide and conquer* technique to reduce the range in each subprogram. Specifically, we will divide the real line into small intervals, estimate means in each interval separately, and report the union. The complete algorithm is given in Algorithm 3.

Algorithm 3 Estimate means of a Gaussian mixture model in the unbounded case.

Input: n samples X_1, \dots, X_n , variance parameter σ^2 (optional), cluster parameter L , and weights threshold τ , test sample size n' .

Output: a set of estimated means \hat{S}

- 1: Merge overlapping intervals $[X_i \pm L]$ for $i \leq n'$ into disjoint ones I_1, \dots, I_s .
 - 2: **for** $j = 1$ **to** s **do**
 - 3: Let c_j, ℓ_j be such that $I_j = [c_j \pm \ell_j]$.
 - 4: Let $C_j = \{X_i - c_j : X_i \in I_j, i > n'\}$.
 - 5: **if** σ^2 is specified **then**
 - 6: Let $(\hat{w}, \hat{\mu})$ be the outcome of Algorithm 1 with input C_j, σ^2 , and $[-\ell_j, \ell_j]$.
 - 7: **else**
 - 8: Let $(\hat{w}, \hat{\mu})$ be the outcome of Algorithm 2 with input C_j .
 - 9: **end if**
 - 10: Let $\hat{S}_j = \{\hat{x}_i + c_j : \hat{w}_i \geq \tau\}$.
 - 11: **end for**
 - 12: Report $\hat{S} = \cup_j \hat{S}_j$.
-

The first step is to apply a clustering method that partitions the samples into a small number of groups. There are many clustering algorithms in practice such as the popular Lloyd's k -means clustering [Llo82]. In lines 1 – 4, we present a conservative yet simple clustering with the following guarantees (see Lemma 18):

- each interval is of length at most $O(kL)$;
- a sample $X_i = U_i + \sigma Z_i$ is always in the same interval as the latent variable U_i .

In the present clustering method, each cluster C_j only contains samples that are not used line 1 so that the intervals are independent of each C_j . This is a commonly used *sample splitting* technique in statistics to simplify the analysis. Note that only a small number of samples are needed to determine the intervals (see Theorem 7). In the second step, we estimate means in each I_j using samples C_j and report the union of all means.

The statistical guarantee of Algorithm 3 is analyzed in Theorem 7. Note that Theorem 7 holds in the worst-case, and can be improved in many situations: the number of samples in each C_j increases proportionally to the total weights; the adaptive rate in Theorem 2 is applicable when separation is present within one interval; we can postulate fewer components in one interval based on information from other intervals.

Theorem 7. *Assume in the Gaussian mixture (1) $w_i \geq \epsilon$, σ is bounded. Let $S = \text{supp}(\nu)$ be the set of means of the Gaussian mixture, and \hat{S} be the output of Algorithm 3 with $L = \Theta(\sqrt{\log n})$ and*

$\tau = \epsilon/(2k)$. If $n \geq 2n' \geq \Omega(\frac{\log(k/\delta)}{\epsilon})$, then, with probability $1 - \delta - n^{-\Omega(1)}$, we have

$$d_H(\hat{S}, S) \leq \begin{cases} O\left(Lk^{3.5}\left(\frac{\epsilon n}{\log(1/\delta)}\right)^{-\frac{1}{4k-2}}/\epsilon\right), & \sigma \text{ is known,} \\ O\left(Lk^4\left(\frac{\epsilon n}{\log(1/\delta)}\right)^{-\frac{1}{4k}}/\epsilon\right), & \sigma \text{ is unknown,} \end{cases}$$

where d_H denotes the Hausdorff distance (see (23)).

5 Lower bounds

This section introduces minimax lower bounds for estimating Gaussian location mixture models which certify the optimality of our estimators. We will apply Le Cam's two-point method, namely, find two Gaussian mixtures that are statistically close but have different parameters. Then any estimate suffers a loss at least proportional to their parameter difference.

To show a vanishing statistical distance between two mixture models, one commonly used proxy is *moment matching*, i.e., $\nu * N(0, 1)$ and $\nu' * N(0, 1)$ are statistically close if $\mathbf{m}_\ell(\nu) = \mathbf{m}_\ell(\nu')$ for some large ℓ . This is demonstrated in Fig. 2, and is made precise in Lemma 8. Statistical closeness via moment matching has been established, for instance, by orthogonal expansion [WV10, CL11], by Taylor expansion [HP15, WY16], and by the best polynomial approximation [WY15]. Similar results to this lemma were previously obtained in [WV10, CL11, HP15].




Figure 2: Statistical closeness via moment matching. In (a), two different mixing distributions coincide on their first six moments; in (b), the mixing distributions are convolved with the standard normal distribution (the black dashed line), and the Gaussian mixtures are statistically close.

Lemma 8. Suppose ν and ν' are centered distributions such that $\mathbf{m}_\ell(\nu) = \mathbf{m}_\ell(\nu')$.

- If ν and ν' are ϵ -subgaussian for $\epsilon < 1$, then

$$\chi^2(\nu * N(0, 1) \| \nu' * N(0, 1)) \leq O\left(\frac{1}{\sqrt{\ell}} \frac{\epsilon^{2\ell+2}}{1 - \epsilon^2}\right). \quad (36)$$

- If ν and ν' are supported on $[-\epsilon, \epsilon]$ for $\epsilon < 1$, then

$$\chi^2(\nu * N(0, 1) \| \nu' * N(0, 1)) \leq O\left(\left(\frac{\epsilon\epsilon^2}{\ell + 1}\right)^{\ell+1}\right). \quad (37)$$

Remark 3 (Tightness of Lemma 8). When ℓ is odd, there exists a pair of ϵ -subgaussian distributions ν and ν' such that $\mathbf{m}_\ell(\nu) = \mathbf{m}_\ell(\nu')$, while $\chi^2(\nu * N(0, 1) \| \nu' * N(0, 1)) \geq \Omega_\ell(\epsilon^{2\ell+2})$. They can be constructed using Gauss quadrature introduced in Section 2.1. To this end, let $\ell = 2k - 1$ and we set $\nu = N(0, \epsilon^2)$ and $\nu' = g_k$ to be its k -point Gauss quadrature. Then $\mathbf{m}_{2k-1}(\nu) = \mathbf{m}_{2k-1}(g_k)$, and g_k is also ϵ -subgaussian (see Lemma 21). It is obtained in [WV10, (54)] that

$$\chi^2(g_k * N(0, 1) \| \nu * N(0, 1)) = \sum_{j \geq 2k} \frac{1}{j!} \left(\frac{\epsilon^2}{1 + \epsilon^2} \right)^j |\mathbb{E}_{\tilde{g}_k}[H_j]|^2,$$

where \tilde{g}_k is the k -point Gauss quadrature of the standard normal distribution, and H_k is the Hermite polynomial of degree k (see (25)). Since $\mathbb{E}_{\tilde{g}_k}[H_{2k}] = -k!$ (see Lemma 23), we have

$$\chi^2(g_k * N(0, 1) \| \nu * N(0, 1)) \geq \frac{(k!)^2}{(2k)!} \left(\frac{\epsilon^2}{1 + \epsilon^2} \right)^{2k} \geq (\Omega(\epsilon))^{4k},$$

when $\epsilon < 1$.

In view of Lemma 8, the best lower bound follows from two different mixing distributions ν and ν' such that $\mathbf{m}_\ell(\nu) = \mathbf{m}_\ell(\nu')$ with the largest degree ℓ . See Lemma 4 and Remark 1 for a discussion on the largest possible ℓ . Next we provide the precise minimax lower bounds for known and unknown variance cases separately.

Known variance. We shall assume a unit variance. First, we define the space of all k Gaussian location mixtures as

$$\mathcal{P}_k = \{\nu * N(0, 1) : \nu \text{ is } k\text{-atomic supported on } [-1, 1]\},$$

and we consider the worst-case risk over all mixture models in \mathcal{P}_k . From the identifiability of discrete distributions in Lemma 4, two different k -atomic distributions can match up to $2k - 2$ moments. Therefore, using Lemma 8, the best minimax lower bound using Le Cam's method is obtained from the optimal pair of distributions for the following:

$$\begin{aligned} & \max W_1(\nu, \nu') \\ & \text{s.t. } \mathbf{m}_{2k-2}(\nu) = \mathbf{m}_{2k-2}(\nu'), \\ & \quad \nu, \nu' \text{ are } k\text{-atomic on } [-\epsilon, \epsilon]. \end{aligned} \tag{38}$$

The value of the above optimization problem is $\Omega(\epsilon/k)$ (see Lemma 20). Using $\epsilon = \sqrt{kn}^{-\frac{1}{4k-2}}$, we obtain the following minimax lower bound:

Proposition 7.

$$\inf_{\hat{\nu}} \sup_{P \in \mathcal{P}_k} \mathbb{E}_P W_1(\nu, \hat{\nu}) \geq \Omega \left(\frac{1}{\sqrt{k}} n^{-\frac{1}{4k-2}} \right)$$

where $\hat{\nu}$ is an estimator measurable with respect to $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P = \nu * N(0, 1)$.

Remark 4. The above lower bound argument can be easily extended to prove the optimality of (11) in Theorem 2, where the mixture satisfies further separation conditions in the sense of Definition 1. In this case, the main difficulty is to estimate parameters in the biggest cluster. When there are k_0

γ -separated clusters, the biggest cluster is of order at most $k' = k - k_0 + 1$. Similar to (38), let $\tilde{\nu}$ and $\tilde{\nu}'$ be two k' -atomic distributions on $[-\epsilon, \epsilon]$. Consider the following mixing distributions

$$\nu = \frac{k_0 - 1}{k_0} \nu_0 + \frac{1}{k_0} \tilde{\nu}, \quad \nu' = \frac{k_0 - 1}{k_0} \nu_0 + \frac{1}{k_0} \tilde{\nu}',$$

where ν_0 is the uniform distribution over $\{\pm 2\gamma, \pm 3\gamma, \dots\}$ of cardinality $k_0 - 1$. Then both mixture models have k_0 $(\gamma, \frac{1}{k_0})$ -separated clusters. Thus the minimax lower bound $\Omega(\frac{1}{\sqrt{k'}} n^{-\frac{1}{4k'-2}})$ analogously follows from Le Cam's method.

Unknown variance. In this case the collection of mixture models is defined as

$$\mathcal{P}'_k = \{\nu * N(0, \sigma^2) : \nu \text{ is } k\text{-atomic supported on } [-1, 1], \sigma \leq 1\}.$$

In Lemma 8, mixing distributions are not restricted to be k -atomic but can be Gaussian location mixtures themselves, thanks to the infinite divisibility of the Gaussian distributions, e.g., $N(0, \epsilon^2) * N(0, 1) = N(0, 1 + \epsilon^2)$. Let g_k be the k -point Gauss quadrature of $N(0, \epsilon^2)$. Then g_k has the same first $2k - 1$ moments as $N(0, \epsilon^2)$, and $g_k * N(0, 1)$ is a k Gaussian mixture with a unit variance. Applying (36) yields that

$$\chi^2(g_k * N(0, 1) \| N(0, 1 + \epsilon^2)) \leq O(\epsilon^{4k}).$$

Using $W_1(g_k, \delta_0) \geq \Omega(\epsilon/\sqrt{k})$ (see Lemma 22), and choosing $\epsilon = n^{-\frac{1}{4k}}$, we obtain the following minimax lower bound:

Proposition 8. For $k \geq 2$,

$$\inf_{\hat{\nu}} \sup_{P \in \mathcal{P}_k} \mathbb{E}_P W_1(\nu, \hat{\nu}) \geq \Omega\left(\frac{1}{\sqrt{k}} n^{-\frac{1}{4k}}\right),$$

$$\inf_{\hat{\nu}} \sup_{P \in \mathcal{P}_k} \mathbb{E}_P |\sigma^2 - \hat{\sigma}^2| \geq \Omega\left(n^{-\frac{1}{2k}}\right),$$

where the infimum is taken over estimators $\hat{\nu}, \hat{\sigma}^2$ measurable with respect to $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} P = \nu * N(0, \sigma^2)$.

6 Numerical experiments

The algorithms of this paper are implemented in Python.⁷ In Algorithm 1, the explicit denoising via semidefinite programming uses CVXPY [DB16] and CVXOPT [ADV13], and the Gauss quadrature is calculated based on [GW69]. In this section, we compare the performance of our algorithms with the EM algorithm, also implemented in Python, and the GMM algorithm using the popular package `gmm` [Cha10] implemented in R. We omit the comparison with the vanilla method of moments which constantly fails to output a meaningful solution (see Section 1.2). In all figures presented in this section, we omit the running time of `gmm`, which is on the order of hours as compared to seconds using our algorithms; the slowness of `gmm` is mainly due to the heuristic solver of the non-convex optimization (30).

We first clarify the parameters used in the experiments. EM and the iterative solver for (30) in `gmm` both require an initialization and a stop criterion. We use the best over five random

⁷ The implementations are available at <https://github.com/Albuso0/mixture>.

initializations: the means are drawn independently from a uniform distribution, and the weights are from a Dirichlet distribution; then we pick the estimate that maximizes the likelihood and the minimal moment discrepancy (30) in EM and GMM, respectively. The EM algorithm terminates when log-likelihood increases less than 10^{-3} or 5,000 iterations are reached; we use the default stop criterion in `gmm` [Cha10].

Known variance. We generated a random instance of Gaussian mixture model with five components and an unit variance. The means are drawn uniformly from $[-1, 1]$; the weights are drawn from the Dirichlet distribution with parameters $(1, 1, 1, 1, 1)$, i.e., uniform over the probability simplex. It has the following parameters:

Weights	0.123	0.552	0.010	0.080	0.235
Centers	-0.236	-0.168	-0.987	0.299	0.150

We repeat the experiments 20 times and plot the average and the standard deviation of the errors in the Wasserstein distance. We also plot the running time at each sample size. The results are shown in Fig. 3. These three algorithms have comparable accuracies, but EM is significantly




Figure 3: Comparison of different methods under a randomly generated five-component Gaussian mixture model.

slower than DMM: it is 15 times slower with 5,000 samples and is increasing with the number of samples. This is because EM accesses all samples in each iteration, instead of first summarizing data into a few moments.

EM becomes slower when samples from different components have more overlaps since the maximizer of the likelihood function lies in a flat area [RW84, KX03]. In this case, a loose stop criterion will terminate the algorithm early before convergence, while a stringent one incurs substantially longer running time. To show this, we do an experiment in which a two-component Gaussian mixture is to be estimated. However, the two components completely overlap, i.e., samples are drawn from $N(0, 1)$. To see the effect of the stop criterion, we additionally run the EM algorithm that terminates when the log-likelihood increases less than 10^{-4} instead of 10^{-3} , shown as EM^+ in the figures. The setup is the same as before and the results are shown in Fig. 4. Again the accuracies are similar, but EM^+ is much slower than EM without substantial gain in the accuracy. Specifically, at 5,000 samples, EM is still 15 times slower than DMM, but EM^+ is 60 times slower.

Lastly, we demonstrate a faster rate in the well-separated case as shown in Theorem 2. In this experiment, the samples are drawn from $\frac{1}{2}N(1, 1) + \frac{1}{2}N(1, -1)$. The results are shown in Fig. 5.

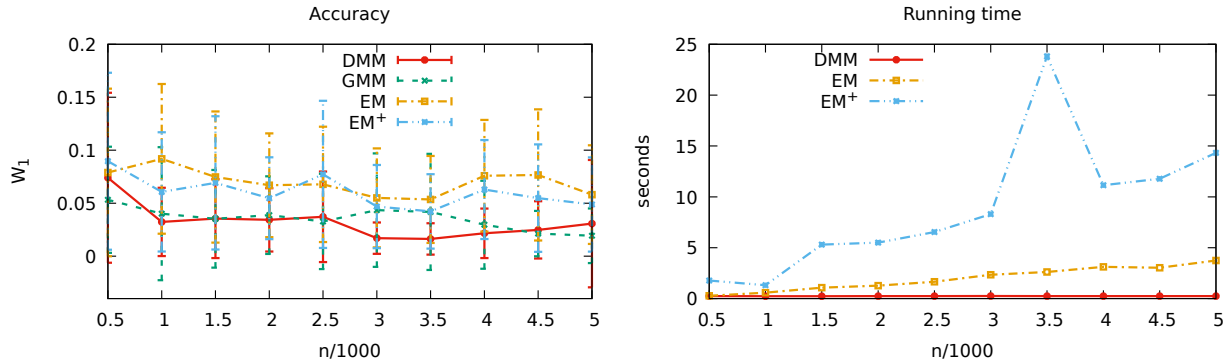


Figure 4: Comparison of different methods when components completely overlap.




Figure 5: Comparison of different methods when components are separated.

In this case, the estimation error decays faster than the one shown in Fig. 4. The larger absolute values of the Wasserstein distance is an artifact of the range of the means.

Unknown variance. We conduct an experiment under the same five-component Gaussian mixture as before, but now the estimators no longer have access to the true variance parameter. In this case, Lindsay’s algorithm (see Algorithm 2) involves the empirical moments of degrees up to 10, among which higher order moments are hard to estimate with limited samples. Indeed, the standard deviation of $\mathbb{E}_n[X^{10}]$ is $\frac{1}{\sqrt{n}}\sqrt{\text{var}[X^{10}]} \approx 473$ under this specified model with $n = 5000$ samples. To resolve this issue, we introduce an extra screening threshold τ to determine whether an empirical moment is too noisy and accept the empirical moment of order j only when its empirical variance satisfies

$$\frac{\mathbb{E}_n[X^{2j}] - (\mathbb{E}_n[X^j])^2}{n} \leq \tau, \quad (39)$$

where the left-hand side of (39) is an estimate of the variance of $\mathbb{E}_n[X^j]$. The estimated mixture model consists of \tilde{k} components for the largest \tilde{k} such that the first $2\tilde{k}$ empirical moments are all accepted. In the experiment, we choose $\tau = 0.5$. The results are shown in Fig. 6. The performance of the Lindsay and EM estimators are similar and better than GMM, which is possibly due to the large variance of higher order empirical moments. The running time comparison are similar




Figure 6: Comparison of different methods with unknown variance.

to before and thus are omitted. The experiments under the models of Fig. 4 and Fig. 5 also yield similar results.

7 Extensions and discussions

7.1 Gaussian location-scale mixtures

In this paper we focus on the Gaussian location mixture model (1), where all components share the same (possibly unknown) variance. One immediate extension is the Gaussian location-scale mixture model with heteroscedastic components:

$$\sum_{i=1}^k w_i N(\mu_i, \sigma_i^2) \quad (40)$$

Parameter estimation for this model turns out to be significantly more difficult than the location mixture model, in particular

- The likelihood function is unbounded. In fact, it is well-known that the maximum likelihood estimator is ill-defined [KW56, p. 905]. For instance, consider $k = 2$, for any sample size n , we have

$$\sup_{p_1, p_2, \theta_1, \theta_2, \sigma} \prod_{i=1}^n \left[\frac{p_1}{\sigma_1} \varphi\left(\frac{X_i - \theta_1}{\sigma_1}\right) + \frac{p_2}{\sigma_2} \varphi\left(\frac{X_i - \theta_2}{\sigma_2}\right) \right] = \infty,$$

achieved by, e.g., $\theta_1 = X_1, p_1 = 1/2, \sigma_2 = 1$, and $\sigma_1 \rightarrow 0$.

- In this model, the identifiability result based on moments is not completely settled and we do not have a counterpart of Lemma 4. Note that the model (40) comprises $3k - 1$ free parameters (k means, k variances, and k weights normalized to one), so it is expected to be identified through its first $3k - 1$ moments. However, the intuition of equating the number of parameters and the number of equations is already known to be wrong as pointed out by Pearson [Pea94], who showed that for $k = 2$, five moments are insufficient and six moments are enough. The recent result [ARS16] showed that, if the parameters are in general positions, then $3k - 1$ moments can identify the Gaussian mixture distribution up to finitely many solutions (known as algebraic identifiability). Whether $3k$ moments can identify the model uniquely (known as rational identifiability) remains open, except for $k = 2$.

Besides the issue of identifiability, the optimal estimation rate under the Gaussian location-scale mixture model is resolved only in special cases. The sharp rate is only known in the case of two components to be $\Theta(n^{-1/12})$ for estimating means and $\Theta(n^{-1/6})$ for estimating variances [HP15], achieved by a robust variation of Pearson’s method of moment equations [Pea94]. For k components, the optimal rate is known to be $n^{-\Theta(1/k)}$ [MV10, KMV10], achieved by an exhaustive grid search on the parameter space. In addition, the above results all aim to recover parameters of all components (up to a global permutation), which necessarily requires many assumptions including lower bounds on mixing weights and separation between components; recovering the mixing distribution with respect to, say, Wasserstein distance, remains open.

7.2 Multiple dimensions

So far we have focused on Gaussian mixtures in one dimension. The multivariate version of this problem has been studied in the context of clustering, or classification, which typically usually requires nonoverlapping components [Das99, VW04]. One commonly used approach is dimensionality reduction: projecting data onto some lower dimensional subspace, clustering samples in that subspace, and mapping back to the original space. Common choices of the subspace include random subspaces and subspaces obtained from the singular value decomposition. The approach using random subspace is analyzed in [Das99, AK01], and requires a pairwise separation polynomial in the dimensions; the subspace from singular value decomposition is analyzed in [VW04, KSV05, AM05, BV08], and requires a pairwise separation that grows polynomially in the number of components. Tensor decomposition for spherical Gaussian mixtures has been studied in [HK13], which requires the stronger assumption that that means are linear independent and is inapplicable in lower dimensions, say, two or three dimensions.

When components are allowed to overlap significantly, the random projection approach is also adopted by [MV10, KMV10, HP15], where the estimation problem in high dimensions is reduced to that in one dimension, so that univariate methodologies can be invoked as a primitive. We provide an algorithm (Algorithm 4) using similar random projection ideas to estimate the parameters of a Gaussian mixture model in d dimensions for known covariance matrices, using the univariate algorithm in Section 4.1 as a subroutine, and obtain the estimation guarantee in Theorem 8; the unknown covariance case can be handled analogously using the algorithm in Section 4.2 instead. However, the dependency of the performance guarantee on the dimension is highly suboptimal,⁸ which stems from the fact that the method based on random projections estimates each coordinate independently. Moreover, this method needs to match the Gaussian components of the estimated model in each random direction, which necessarily requires lower bounds on the mixing weights and separation between the means.

Theorem 8. *Suppose in a d -dimensional Gaussian mixture $\sum_{j=1}^k w_j N(\mu_j, \Sigma)$,*

$$\|\mu_j\|_2 \leq M, \quad \|\mu_i - \mu_j\|_2 \leq \epsilon, \quad w_j \geq \epsilon', \quad \forall i \neq j.$$

Then Algorithm 4 with $n > (\Omega_k(\frac{M}{\tilde{\epsilon}\epsilon'}))^{4k-2} \log \frac{d}{\delta}$ samples, $\tau = \frac{\tilde{\epsilon}}{2M}$, and $\rho = M$, where $\tilde{\epsilon} = \frac{\delta\epsilon}{k^2\sqrt{d}}$, yields $\hat{\pi}$ such that, with probability $1 - 2\delta$,

$$W_1(\pi, \hat{\pi}) < O_k \left(\sqrt{d} \frac{M\epsilon_n}{\tau\epsilon'} \right),$$

where $\pi = \sum_j w_j \delta_{\mu_j}$ and $\epsilon_n = \min\left\{ \left(\frac{n}{\log(d/\delta)}\right)^{-\frac{1}{4k-2}}, \tilde{\epsilon}^{2-2k} \sqrt{\frac{\log(d/\delta)}{n}} \right\}$.

⁸ Specifically, in d dimensions, estimating each coordinate independently suffers an ℓ_2 -loss proportional to \sqrt{d} ; however, it is possible to achieve $d^{1/4}$. See in Lemma 34 for an example.

Algorithm 4 Learning a k -component Gaussian mixture in d dimensions.

Input: n samples $X_1, X_2, \dots, X_n \in \mathbb{R}^d$, common covariance matrix Σ , and separation parameter τ , radius parameter ρ .

Output: estimated mixing distribution $\hat{\pi}$ with weights and means $(\hat{w}_j, \hat{\mu}_j)$ for $j = 1, \dots, k$.

- 1: Let (b_1, \dots, b_d) be a set of random orthonormal basis in \mathbb{R}^d , and $r = b_1$.
 - 2: Let $\{(w_j, \mu_j)\}$ be the outcome of Algorithm 1 using n projected samples $\langle X_1, r \rangle, \dots, \langle X_n, r \rangle$, variance $r^\top \Sigma r$, and interval $[-\rho, \rho]$.
 - 3: Reordering the indices such that $\mu_1 < \mu_2 < \dots < \mu_k$.
 - 4: Initialize k weights $\hat{w}_j = w_j$ and means $\hat{\mu}_j = (0, \dots, 0)$.
 - 5: **for** $i = 1$ **to** d **do**
 - 6: Let $r' = r + \tau b_i$.
 - 7: Let $\{\mu'_j\}$ be the estimated means (weights are ignored) from Algorithm 1 using n projected samples $\langle X_1, r' \rangle, \dots, \langle X_n, r' \rangle$, variance $r'^\top \Sigma r'$, and interval $[-\rho - \tau, \rho + \tau]$.
 - 8: Reordering the indices such that $\mu'_1 < \mu'_2 < \dots < \mu'_k$.
 - 9: Let $\hat{\mu}_j := \hat{\mu}_j + b_i \frac{\mu'_j - \mu_j}{\tau}$ for $j = 1, \dots, k$.
 - 10: **end for**
-

Proof. By the distribution of random direction r on the unit sphere (see Lemma 33) and the union bound, we obtain that, with probability $1 - \delta$,

$$|\langle \mu_i - \mu_j, r \rangle| > 2\tilde{\epsilon}, \quad \forall i \neq j.$$

Without loss of generality, assume $\langle \mu_1, r \rangle < \dots < \langle \mu_k, r \rangle$. Applying Theorem 1 yields that, with probability $1 - \frac{\delta}{d+1}$,

$$W_1(\pi_r, \hat{\pi}_r) \leq O_k \left(M \left(\frac{n}{\log(d/\delta)} \right)^{-\frac{1}{4k-2}} \right),$$

where π_r denotes the Gaussian mixture projected on r and $\hat{\pi}_r$ is its estimate. The right-hand side of the above inequality is less than $c\epsilon_r\epsilon'$ for some constant $c < 0.5$ when $n > (\Omega_k(\frac{M}{\tilde{\epsilon}\epsilon'}))^{4k-2} \log \frac{d}{\delta}$. Applying Theorem 2 yields that

$$W_1(\pi_r, \hat{\pi}_r) \leq O_k \left(M\tilde{\epsilon}^{2-2k} \sqrt{\frac{\log(d/\delta)}{n}} \right).$$

Hence, we obtained $W_1(\pi_r, \hat{\pi}_r) \leq O_k(M\epsilon_n)$. It follows from Lemma 1 that, after reordering indices,

$$|\langle \mu_j, r \rangle - \tilde{\mu}_j| < O_k(M\epsilon_n/\epsilon'), \quad |w_j - \hat{w}_j| < O_k(M\epsilon_n/\tilde{\epsilon}). \quad (41)$$

On each direction $r_\ell = r + \tau b_\ell$, the means are separated by $|\langle \mu_i - \mu_j, r_\ell \rangle| > 2\tilde{\epsilon} - 2M\tau > \tilde{\epsilon}$ and the ordering of the means remains the same as on direction r . Therefore the accuracy similar to (41) continues to hold for the estimated means $\tilde{\mu}_{\ell,j}$ (μ'_j in lines 7 and 8). Note that $\mu_j = \sum_\ell b_\ell \frac{\langle \mu_j, r_\ell \rangle - \langle \mu_j, r \rangle}{\tau}$ and $\hat{\mu}_j = \sum_\ell b_\ell \frac{\tilde{\mu}_{\ell,j} - \hat{\mu}_j}{\tau}$. Therefore,

$$\|\hat{\mu}_j - \mu_j\|_2^2 \leq \sum_{\ell=1}^d \left(\frac{O_k(M\epsilon_n/\epsilon')}{\tau} \right)^2.$$

Applying the triangle equality yields that

$$W_1(\pi, \hat{\pi}) < \sqrt{d} O_k(M\epsilon_n/\epsilon')/\tau + M O_k(M\epsilon_n/\tilde{\epsilon}) < O_k \left(\sqrt{d} \frac{M\epsilon_n}{\tau\epsilon'} \right). \quad \square$$

It is interesting to directly extend the DMM methodology to multiple dimensions, which is challenging both theoretically and algorithmically:

- To apply our method in multiple dimensions, the challenge is to obtain a multi-dimensional moment comparison theorem analogous to Proposition 1 or 2, the key step leading to the optimal rate. These results are proved in Section 8.1 by the primal formulation of the Wasserstein distance and its simple formula (22) in one dimension [Vil03]. Alternatively, they can be proved via the dual formula (21) which holds in any dimension; however, the proof relies on the Newton's interpolation formula, which is again difficult to generalize or analyze in multiple dimensions.
- To obtain a computationally efficient algorithm, we rely on the semidefinite characterization of the moment space in one dimension to denoise the noisy estimates of moments. In multiple dimensions, however, it remains open how to efficiently describe the moment space [Las09] as well as how to extend the Gauss quadrature rule to multivariate distributions.

7.3 General finite mixtures

Though this paper focuses on Gaussian location mixture models, the moments comparison theorems in Section 3 are independent of properties of Gaussian. As long as moments of the mixing distribution are estimated accurately, similar theory and algorithms can be obtained. Unbiased estimate of moments exists in many useful mixture models, including exponential mixtures [Jew82], Poisson mixtures [KX05], and more generally the quadratic variance exponential family (QVEF) whose variance is at most a quadratic function of the mean [Mor82, (8.8)].

As a closely related topic of this paper, we discuss the Gaussian scale mixture model in details, which has been extensively studied in the statistics literature [AM74] and is widely used in image and video processing [WS00, PSWS03]. In a Gaussian scale mixture, a sample is distributed as

$$X \sim \sum_{i=1}^k w_i N(0, \sigma_i^2) = \int N(0, \sigma^2) d\nu(\sigma^2),$$

where $\nu = \sum_{i=1}^k w_i \delta_{\sigma_i^2}$ is a k -atomic mixing distribution. Equivalently, a sample can be represented as $X = \sqrt{V}Z$, where $V \sim \nu$ and Z is standard normal independent of V . In this model, samples from different components significantly overlap, so clustering-based algorithms will fail. Nevertheless, moments of ν can be easily estimated, for instance, using $\mathbb{E}_n[X^{2r}]/\mathbb{E}[Z^{2r}]$ for $m_r(\nu)$ with accuracy $O_r(1/\sqrt{n})$. Applying a similar algorithm to DMM in Section 4.1, we obtain an estimate $\hat{\nu}$ such that

$$W_1(\nu, \hat{\nu}) \leq O_k(n^{-\frac{1}{4k-2}}),$$

with high probability.

Moreover, using similar recipe to Section 5, a minimax lower bound can be established. Analogous to (38), let ν and ν' be a pair of k -atomic distributions supported on $[0, \epsilon]$ such that they match the first $2k - 2$ moments, and let

$$\pi = \int N(0, \sigma^2) d\nu(\sigma^2), \quad \pi' = \int N(0, \sigma^2) d\nu'(\sigma^2),$$

which match their first $4k - 3$ moments and are $\sqrt{\epsilon}$ -subgaussian. Applying Lemma 8 with $\pi * N(0, 0.5)$, $\pi' * N(0, 0.5)$, and $\epsilon = O_k(n^{-\frac{1}{4k-2}})$ yields a minimax lower bound

$$\inf_{\hat{\nu}} \sup_{P \in \mathcal{G}_k} \mathbb{E}_P W_1(\nu, \hat{\nu}) \geq \Omega_k \left(n^{-\frac{1}{4k-2}} \right),$$

where the estimator $\hat{\nu}$ is measurable with respect to $X_1, \dots, X_n \sim P$, and the space of k Gaussian scale mixtures is defined as

$$\mathcal{G}_k = \left\{ \int N(0, \sigma^2) d\nu(\sigma^2) : \nu \text{ is } k\text{-atomic supported on } [0, 1] \right\}.$$

8 Proofs

8.1 Proof of moments comparison theorems

In this subsection we prove Propositions 1 and 2. As a warm-up, we start by proving Lemma 4, with the purpose of introducing the apparatus of interpolating polynomials. Throughout this section, we use

$$F_\pi(x) = \int_{-\infty}^x \pi(dt).$$

to denote the CDF of a distribution π .

Proof of Lemma 4. We only need to prove the “if” part.

1. Denote the union of the support sets of ν and ν' by S . Here S is of size at most $2k$. For any $t \in \mathbb{R}$, there exists a polynomial P of degree at most $2k - 1$ to interpolate $x \mapsto \mathbf{1}_{\{x \leq t\}}$ on S . Since $m_i(\nu) = m_i(\nu')$ for $i = 1, \dots, 2k - 1$, we have

$$F_\nu(t) = \mathbb{E}_\nu[\mathbf{1}_{\{X \leq t\}}] = \mathbb{E}_\nu[P(X)] = \mathbb{E}_{\nu'}[P(X)] = \mathbb{E}_{\nu'}[\mathbf{1}_{\{X \leq t\}}] = F_{\nu'}(t).$$

2. Denote the support set of ν by $S' = \{x_1, \dots, x_k\}$. Let $Q(x) = \prod_i (x - x_i)^2$, a non-negative polynomial of degree $2k$. Since $m_i(\nu) = m_i(\nu')$ for $i = 1, \dots, 2k$, we have

$$\mathbb{E}_{\nu'}[Q(X)] = \mathbb{E}_\nu[Q(X)] = 0.$$

Therefore, ν' is also supported on S' and thus is k -atomic. The conclusion follows from the first case of Lemma 4. \square

Next we prove Proposition 9, which is slightly stronger than Proposition 1. We provide three proofs: the first two are based on the primal (coupling) formulation of W_1 distance (22), and the third proof uses the dual formulation (21). Specifically,

- The first proof uses polynomials to interpolate step functions, whose expected values are the CDFs. The closeness of moments imply the closeness of distribution functions and thus, by (22), a small Wasserstein distance. Similar idea applies to the proof of Proposition 2 later.
- The second proof finds a polynomial that preserves the sign of the difference between two CDFs, and then relate the Wasserstein distance to the integral of that polynomial. Related idea has been used in [MV10, Lemma 20] which finds a polynomial that preserves the sign of the difference between two Gaussian mixture densities.
- The third proof uses polynomials to approximate 1-Lipschitz functions, whose expected values are related to the Wasserstein distance via the dual formulation (21).

Proposition 9. *Let ν and ν' be discrete distributions supported on a total of ℓ atoms in $[-1, 1]$. If*

$$|m_i(\nu) - m_i(\nu')| \leq \delta, \quad i = 1, \dots, \ell - 1, \quad (42)$$

then

$$W_1(\nu, \nu') \leq O\left(\ell \delta^{\frac{1}{\ell-1}}\right).$$

First proof of Proposition 9. Suppose ν and ν' are supported on

$$S = \{t_1, \dots, t_\ell\}, \quad t_1 < t_2 < \dots < t_\ell. \quad (43)$$

Then, using the integral representation (22), the W_1 distance reduces to

$$W_1(\nu, \nu') = \sum_{r=1}^{\ell-1} |F_\nu(t_r) - F_{\nu'}(t_r)| \cdot |t_{r+1} - t_r|. \quad (44)$$

For each r , let $f_r(x) = \mathbf{1}_{\{x \leq t_r\}}$, and P_r be the unique polynomial of degree $\ell - 1$ to interpolate f_r on S . In this way we have $f_r = P_r$ almost surely under both ν and ν' , and thus

$$|F_\nu(t_r) - F_{\nu'}(t_r)| = |\mathbb{E}_\nu f_r - \mathbb{E}_{\nu'} f_r| = |\mathbb{E}_\nu P_r - \mathbb{E}_{\nu'} P_r|. \quad (45)$$

P_r can be expressed using Newton formula (17) as

$$P_r(x) = 1 + \sum_{i=r+1}^{\ell} f_r[t_1, \dots, t_i] g_{i-1}(x), \quad (46)$$

where $g_r(x) = \prod_{j=1}^r (x - t_j)$ and we used $f_r[t_1, \dots, t_i] = 0$ for $i = 1, \dots, r$. In (46), the absolute values of divided differences are obtained in Lemma 24:

$$|f_r[t_1, \dots, t_i]| \leq \frac{\binom{i-2}{r-1}}{(t_{r+1} - t_r)^{i-1}}. \quad (47)$$

In the summation of (46), let $g_{i-1}(x) = \sum_{j=0}^{i-1} a_j x^j$. Since $|t_j| \leq 1$ for every j , we have $\sum_{j=0}^{i-1} |a_j| \leq 2^{i-1}$ (see Lemma 25). Applying (42) yields that

$$|\mathbb{E}_\nu[g_{i-1}] - \mathbb{E}_{\nu'}[g_{i-1}]| \leq \sum_{j=1}^{i-1} |a_j| \delta \leq 2^{i-1} \delta. \quad (48)$$

Then we obtain from (45) and (46) that

$$|F_\nu(t_r) - F_{\nu'}(t_r)| \leq \sum_{i=r+1}^{\ell} \frac{\binom{i-2}{r-1} 2^{i-1} \delta}{(t_{r+1} - t_r)^{i-1}} \leq \frac{\ell 4^{\ell-1} \delta}{(t_{r+1} - t_r)^{\ell-1}}. \quad (49)$$

Also, $|F_\nu(t_r) - F_{\nu'}(t_r)| \leq 1$ trivially. Therefore,

$$W_1(\nu, \nu') \leq \sum_{r=1}^{\ell-1} \left(\frac{\ell 4^{\ell-1} \delta}{(t_{r+1} - t_r)^{\ell-1}} \wedge 1 \right) \cdot |t_{r+1} - t_r| \leq 4e(\ell - 1) \delta^{\frac{1}{\ell-1}}, \quad (50)$$

where we used $\max\{\frac{\alpha}{x^{\ell-2}} \wedge x : x > 0\} = \alpha^{\frac{1}{\ell-1}}$ and $x^{\frac{1}{x-1}} \leq e$ for $x \geq 1$. \square

Second proof of Proposition 9. Suppose on the contrary that

$$W_1(\nu, \nu') \geq C \ell \delta^{\frac{1}{\ell-1}}, \quad (51)$$

for some absolute constant C . We will show that $\max_{i \in [\ell-1]} |m_i(\nu) - m_i(\nu')| \geq \delta$. We continue to use S in (43) to denote the support of ν and ν' . Let $\Delta F(t) = F_\nu(t) - F_{\nu'}(t)$ denote the difference between two CDFs. Using (44), there exists $r \in [\ell - 1]$ such that

$$|\Delta F(t_r)| \cdot |t_{r+1} - t_r| \geq C \delta^{\frac{1}{\ell-1}}. \quad (52)$$

We first construct a polynomial L that preserves the sign of ΔF . To this end, let $S' = \{s_1, \dots, s_m\} \subseteq S$ such that $t_1 = s_1 < s_2 < \dots < s_m = t_\ell$ be the set of points where ΔF changes sign, i.e., $\Delta F(x)\Delta F(y) \leq 0$ for every $x \in (s_i, s_{i+1})$, $y \in (s_{i+1}, s_{i+2})$, for every i . Let $L(x) \in \pm \prod_{i=2}^{m-1} (x - s_i)$ be a polynomial of degree at most $\ell - 2$ that also changes sign on S' such that

$$\Delta F(x)L(x) \geq 0, \quad t_1 \leq x \leq t_\ell.$$

Consider the integral of the above positive function. Applying integral by parts, and using $\Delta F(t_\ell) = \Delta F(t_1) = 0$ yields that

$$\int_{t_1}^{t_\ell} \Delta F(x)L(x)dx = - \int_{t_1}^{t_\ell} P(x)d\Delta F(x) = \mathbb{E}_{\nu'}[P(X)] - \mathbb{E}_\nu[P(X)], \quad (53)$$

where $P(x)$ is a polynomial of degree at most $\ell - 1$ such that $P'(x) = L(x)$. If we write $L(x) = \sum_{j=0}^{\ell-2} a_j x^j$, then $P(x) = \sum_{j=0}^{\ell-2} \frac{a_j}{j+1} x^{j+1}$. Since $|s_j| \leq 1$ for every j , we have $\sum_{j=0}^{\ell-2} |a_j| \leq 2^{\ell-2}$ (see Lemma 25), and thus $\sum_{j=0}^{\ell-2} \frac{|a_j|}{j+1} \leq 2^{\ell-2}$. Hence,

$$|\mathbb{E}_{\nu'}[P(X)] - \mathbb{E}_\nu[P(X)]| \leq 2^{\ell-2} \max_{i \in [\ell-1]} |m_i(\nu) - m_i(\nu')|. \quad (54)$$

Since $\Delta F(x)L(x)$ is always non-negative, applying (52) to (53) yields that

$$|\mathbb{E}_{\nu'}[P(X)] - \mathbb{E}_\nu[P(X)]| \geq \int_{t_r}^{t_{r+1}} |\Delta F(x)L(x)|dx \geq \frac{C\delta^{\frac{1}{\ell-1}}}{|t_{r+1} - t_r|} \int_{t_r}^{t_{r+1}} |L(x)|dx. \quad (55)$$

Recall that $|L(x)| = \prod_{i=2}^{m-1} |x - s_i|$. Then for $x \in (t_r, t_{r+1})$, we have $|x - s_i| \geq x - t_r$ if $s_i \leq t_r$, and $|x - s_i| \geq t_{r+1} - x$ if $s_i \geq t_{r+1}$. Hence,

$$|L(x)| \geq (t_{r+1} - x)^a (x - t_r)^b,$$

for some $a, b \in \mathbb{N}$ such that $a, b \geq 1$ and $a + b \leq \ell - 2$. The integral of the right-hand side of the above inequality can be expressed as (see [AS64, 6.2.1])

$$\int_{t_r}^{t_{r+1}} (t_{r+1} - x)^a (x - t_r)^b dx = \frac{(t_{r+1} - t_r)^{a+b+1}}{(a+1)\binom{a+b+1}{b}}.$$

Since $|t_{r+1} - t_r| \geq |\Delta F(t_r)| \cdot |t_{r+1} - t_r| \geq C\delta^{\frac{1}{\ell-1}}$ and $\binom{a+b+1}{b} \leq 2^{a+b+1}$, and $a + b + 1 \leq \ell - 1$, we obtain from (55) that

$$|\mathbb{E}_{\nu'}[P(X)] - \mathbb{E}_\nu[P(X)]| \geq \delta \frac{(C/2)^{\ell-1}}{\ell}. \quad (56)$$

We obtain from (54) and (56) that

$$\max_{i \in [\ell-1]} |m_i(\nu) - m_i(\nu')| \geq \delta \frac{(C/4)^{\ell-1}}{\ell}. \quad \square$$

Third proof of Proposition 9. We continue to use S in (43) to denote the support of ν and ν' . For any 1-Lipschitz function f , $\mathbb{E}_\nu f$ and $\mathbb{E}_{\nu'} f$ only pertain to function values $f(t_1), \dots, f(t_\ell)$, which can be interpolated by a polynomial of degree $\ell - 1$. However, the coefficients of the interpolating

polynomial can be arbitrarily large.⁹ To fix this issue, we slightly modify the function f on S to \tilde{f} , and then interpolate \tilde{f} with bounded coefficients. In this way we have

$$|\mathbb{E}_\nu f - \mathbb{E}_{\nu'} f| \leq 2 \max_{x \in \{t_1, \dots, t_\ell\}} |\tilde{f}(x) - f(x)| + |\mathbb{E}_\nu P - \mathbb{E}_{\nu'} P|.$$

To this end, we define the values of \tilde{f} recursively by

$$\tilde{f}(t_1) = f(t_1), \quad \tilde{f}(t_i) = \tilde{f}(t_{i-1}) + (f(t_i) - f(t_{i-1})) \mathbf{1}_{\{t_i - t_{i-1} > \tau\}}, \quad (57)$$

where $\tau \leq 2$ is a parameter we will optimize later. From the above definition $|\tilde{f}(x) - f(x)| \leq \tau \ell$ for $x \in S$. The interpolating polynomial P can be expressed using Newton formula (17) as

$$P(x) = \sum_{i=1}^{\ell} \tilde{f}[t_1, \dots, t_i] g_{i-1}(x),$$

where $g_r(x) = \prod_{j=1}^r (x - t_j)$ such that $|\mathbb{E}_\nu[g_r] - \mathbb{E}_{\nu'}[g_r]| \leq 2^r \delta$ by (48) for $r \leq \ell - 1$. Since f is 1-Lipschitz, we have $|\tilde{f}[t_i, t_{i+1}]| \leq 1$ for every i . Higher order divided differences are recursively evaluated by (18). We now prove

$$\tilde{f}[t_i, \dots, t_{i+j}] \leq (2/\tau)^{j-1}, \quad \forall i, j. \quad (58)$$

by induction on j . Assume (58) holds for every i and some fixed j . The recursion (18) gives

$$\tilde{f}[t_i, \dots, t_{i+j+1}] = \frac{\tilde{f}[t_{i+1}, \dots, t_{i+j+1}] - \tilde{f}[t_i, \dots, t_{i+j}]}{t_{i+j+1} - t_i}.$$

If $t_{i+j+1} - t_i < \tau$, then $\tilde{f}[t_i, \dots, t_{i+j+1}] = 0$ by (57); otherwise, $\tilde{f}[t_i, \dots, t_{i+j+1}] \leq (\frac{2}{\tau})^j$ by triangle inequality. Using (58), we obtain that

$$|\mathbb{E}_\nu f - \mathbb{E}_{\nu'} f| \leq 2\tau\ell + \sum_{i=2}^{\ell} \left(\frac{2}{\tau}\right)^{i-2} 2^{i-1} \delta \leq 2\ell \left(\tau + \frac{4^{\ell-2}}{\tau^{\ell-2}} \delta\right).$$

The conclusion follows by letting $\tau = 4\delta^{\frac{1}{\ell-1}}$. □

The proof of Proposition 2 uses a similar idea as the first proof of Proposition 9 to approximate step functions for all values of ν and ν' ; however, this is clearly impossible for non-discrete ν' . For this reason, we turn from interpolation to majorization. A classical method to bound a distribution function by moments is to construct two polynomials that majorizes and minorizes a step function, respectively. Then the expectations of these two polynomials provide a sandwich bound for the distribution function. This idea is used, for example, in the proof of Chebyshev-Markov-Stieltjes inequality (cf. [Akh65, Theorem 2.5.4]).

Proof of Proposition 2. Suppose ν is supported on $x_1 < x_2 < \dots < x_k$. Fix $t \in \mathbb{R}$ and let $f_t(x) = \mathbf{1}_{\{x \leq t\}}$. Suppose $x_m < t < x_{m+1}$. Similar to Example 4, we construct polynomial majorant and minorant using Hermite interpolation. To this end, let P_t and Q_t be the unique degree- $2k$ polynomials to interpolate f_t with the following:

⁹ For example, the polynomial to interpolate $f(-\epsilon) = f(\epsilon) = \epsilon, f(\epsilon) = 0$ is $P(x) = x^2/\epsilon$.

	x_1	\dots	x_m	t	x_{m+1}	\dots	x_k
P	1	\dots	1	1	0	\dots	0
P'	0	\dots	0	any	0	\dots	0
Q	1	\dots	1	0	0	\dots	0
Q'	0	\dots	0	any	0	\dots	0

As a consequence of Rolle's theorem, $P_t \geq f_t \geq Q_t$ (cf. [Akh65, p. 65], and an illustration in Fig. 7): Using Lagrange formula of Hermite interpolation [SB02, pp. 52–53], P_t and Q_t differ by




Figure 7: Polynomial majorant P_t and minorant Q_t that coincide with the step function on 6 red points. The polynomials are of degree 12, obtained by Hermite interpolation in Section 2.2.

$$P_t(x) - Q_t(x) = R_t(x) \triangleq \prod_i \left(\frac{x - x_i}{t - x_i} \right)^2.$$

The sandwich bound for f_t yields a sandwich bound for the CDFs:

$$\begin{aligned} \mathbb{E}_{\nu'}[Q_t] \leq F_{\nu'}(t) \leq \mathbb{E}_{\nu'}[P_t] &= \mathbb{E}_{\nu'}[Q_t] + \mathbb{E}_{\nu'}[R_t], \\ \mathbb{E}_{\nu}[Q_t] \leq F_{\nu}(t) \leq \mathbb{E}_{\nu}[P_t] &= \mathbb{E}_{\nu}[Q_t]. \end{aligned}$$

Then the CDFs differ by

$$\begin{aligned} |F_{\nu'}(t) - F_{\nu}(t)| &\leq (f(t) + g(t)) \wedge 1 \leq f(t) \wedge 1 + g(t) \wedge 1, \\ f(t) &\triangleq |\mathbb{E}_{\nu'}[Q_t] - \mathbb{E}_{\nu}[Q_t]|, \quad g(t) \triangleq \mathbb{E}_{\nu'}[R_t]. \end{aligned} \tag{59}$$

The conclusion will be obtained from the integral of CDF difference using (22). Since R_t is almost surely zero under ν , we also have $g(t) = |\mathbb{E}_{\nu'}[R_t] - \mathbb{E}_{\nu}[R_t]|$. Similar to (48), we obtain that

$$g(t) = |\mathbb{E}_{\nu'}[R_t] - \mathbb{E}_{\nu}[R_t]| \leq \frac{2^{2k} \delta}{\prod_{i=1}^k (t - x_i)^2}.$$

Hence,

$$\int (g(t) \wedge 1) dt \leq \int \left(\frac{2^{2k} \delta}{\prod_{i=1}^k (t - x_i)^2} \wedge 1 \right) dt \leq 16k \delta^{\frac{1}{2k}}, \tag{60}$$

where the last inequality is proved in Lemma 28.

Next we analyze $f(t)$. The polynomial Q_t (and also P_t) can be expressed using Newton formula (17) as

$$Q_t(x) = 1 + \sum_{i=2m+1}^{2k+1} f_t[t_1, \dots, t_i] g_{i-1}(x), \quad (61)$$

where t_1, \dots, t_{2k+1} denotes the expanded sequence

$$x_1, x_1, \dots, x_m, x_m, t, x_{m+1}, x_{m+1}, \dots, x_k, x_k$$

obtained by (20), $g_r(x) = \prod_{j=1}^r (x - t_j)$, and we used $f_t[t_1, \dots, t_i] = 0$ for $i = 1, \dots, 2m$. In (61), the absolute values of divided differences are obtained in Lemma 24:

$$f_t[t_1, \dots, t_i] \leq \frac{\binom{i-2}{2m-1}}{(t - x_m)^{i-1}}.$$

Using (61), and applying the upper bound for $|\mathbb{E}_{\nu'}[g_{i-1}] - \mathbb{E}_{\nu}[g_{i-1}]|$ in (48), we obtain that,

$$f(t) = |\mathbb{E}_{\nu'}[Q_t] - \mathbb{E}_{\nu}[Q_t]| \leq \sum_{i=2m+1}^{2k+1} \frac{\binom{i-2}{2m-1} 2^{i-1} \delta}{(t - x_m)^{i-1}} \leq \frac{k 4^{2k} \delta}{(t - x_m)^{2k}}, \quad x_m < t < x_{m+1}, \quad m \geq 1.$$

If $t < x_1$, then $Q_t = 0$ and thus $f(t) = 0$. Then, analogous to (60), we obtain that

$$\int (f(t) \wedge 1) dt \leq 16k\delta^{\frac{1}{2k}}. \quad (62)$$

Using (60) and (62), the conclusion follows by applying (59) to the integral representation of Wasserstein distance (22). \square

8.2 Proofs of density estimation

Lemma 9 (Bound χ^2 -divergence using moments difference). *Suppose all moments of ν and ν' exist, and ν' is centered with variance σ^2 . Then,*

$$\chi^2(\nu * N(0, 1) \| \nu' * N(0, 1)) \leq e^{\frac{\sigma^2}{2}} \sum_{j \geq 1} \frac{(\Delta m_j)^2}{j!},$$

where $\Delta m_j = m_j(\nu) - m_j(\nu')$ denotes the j^{th} moment difference.

Proof. The densities of two mixture distributions $\nu * N(0, 1)$ and $\nu' * N(0, 1)$ are

$$f(x) = \int \phi(x - u) d\nu(u) = \phi(x) \sum_{j \geq 1} H_j(x) \frac{m_j(\nu)}{j!},$$

$$g(x) = \int \phi(x - u) d\nu'(u) = \phi(x) \sum_{j \geq 1} H_j(x) \frac{m_j(\nu')}{j!},$$

respectively, where ϕ denotes the density of $N(0, 1)$, and we used $\phi(x - u) = \phi(x) \sum_{j \geq 0} H_j(x) \frac{x^j}{j!}$ (see the exponential generating function of Hermite polynomials [AS64, 22.9.17]). Since $x \mapsto e^x$ is convex, applying Jensen's inequality yields that

$$g(x) = \phi(x) \mathbb{E}[\exp(U'x - U'^2/2)] \geq \phi(x) \exp(-\sigma^2/2).$$

Consequently,

$$\begin{aligned} \chi^2(\nu * N(0, 1) \| \nu' * N(0, 1)) &= \int \frac{(f(x) - g(x))^2}{g(x)} dx \\ &\leq e^{\frac{\sigma^2}{2}} \mathbb{E} \left[\left(\sum_{j \geq 1} H_j(Z) \frac{\Delta m_j}{j!} \right)^2 \right] = e^{\frac{\sigma^2}{2}} \sum_{j \geq 1} \frac{(\Delta m_j)^2}{j!}, \end{aligned}$$

where $Z \sim N(0, 1)$ and the last step follows from the orthogonality of Hermite polynomials:

$$\mathbb{E}[H_i(Z)H_j(Z)] = j! \mathbf{1}_{\{i=j\}}. \quad (63)$$

□

Lemma 10. *If U and U' each takes at most k values in $[-1, 1]$, and $|\mathbb{E}[U^j] - \mathbb{E}[U'^j]| \leq \epsilon$ for $j = 1, \dots, 2k - 1$, then, for any $\ell \geq 2k$,*

$$|\mathbb{E}[U^\ell] - \mathbb{E}[U'^\ell]| \leq 3^\ell \epsilon.$$

Proof. Let $f(x) = x^\ell$ and denote the atoms of U and U' by $x_1 < \dots < x_{k'}$ for $k' \leq 2k$. The function f can be interpolated on $x_1, \dots, x_{k'}$ using a polynomial P of degree at most $2k - 1$, which, in the Newton form (17), is

$$P(x) = \sum_{i=1}^{k'} f[x_1, \dots, x_i] g_{i-1}(x) = \sum_{i=1}^{k'} \frac{f^{(i-1)}(\xi_i)}{(i-1)!} g_{i-1}(x),$$

for some $\xi_i \in [x_1, x_i]$, where $g_r(x) = \prod_{j=1}^r (x - x_j)$ and we used the intermediate value theorem for the divided differences (see [SB02, (2.1.4.3)]). Note that for $\xi_i \in [-1, 1]$, $|f^{(i-1)}(\xi_i)| \leq \frac{\ell!}{(\ell-1+i)!}$. Similar to (48), we obtain that

$$|\mathbb{E}[U^\ell] - \mathbb{E}[U'^\ell]| = |\mathbb{E}[P(U)] - \mathbb{E}[P(U')]| \leq \sum_{i=1}^{k'} \binom{\ell}{i-1} 2^{i-1} \epsilon \leq 3^\ell \epsilon. \quad \square$$

Proof of Theorem 3. By scaling it suffices to consider $M = 1$. Similar to (28) and (29), we obtain an estimated mixing distribution $\hat{\nu}$ supported on k atoms in $[-1, 1]$ such that, with probability $1 - \delta$,

$$\|\mathbf{m}_{2k-1}(\hat{\nu}) - \mathbf{m}_{2k-1}(\nu)\|_2 \leq \sqrt{c_k \log(1/\delta)/n},$$

for some constant c_k that depends on k . The conclusion follows from Lemmas 9 and 10. □

Proof of Theorem 4. Recall that f is 1-subgaussian and σ is a fixed constant. Similar to (29), we obtain an estimate \tilde{m}_r for $\mathbb{E}_f[\gamma_r(X, \sigma)]$ (see the definition of $\gamma_r(\cdot, \sigma)$ in (26)) for $r = 1, \dots, 2k - 1$ such that, with probability $1 - \delta$,

$$|\tilde{m}_r - \mathbb{E}_f[\gamma_r(X, \sigma)]| \leq \sqrt{c_k \log(1/\delta)/n},$$

for some constant c_k that depends on k . By assumption, $\text{TV}(f, g) \leq \epsilon$ where both f and g are 1-subgaussian. Let $g = \nu * N(0, \sigma^2)$. Then, using Lemma 11 below and the triangle inequality, we have

$$|\tilde{m}_r - m_r(\nu)| \leq O_k \left(\epsilon \sqrt{\log(1/\epsilon)} + \sqrt{\log(1/\delta)/n} \right), \quad r = 1, \dots, 2k - 1.$$

Using the projection (24), we obtain $\hat{\nu}$ similar to (28) such that

$$\|\mathbf{m}_{2k-1}(\hat{\nu}) - \mathbf{m}_{2k-1}(\nu)\|_2 \leq O_k \left(\epsilon \sqrt{\log(1/\epsilon)} + \sqrt{\log(1/\delta)/n} \right).$$

Let $\hat{f} = \hat{\nu} * N(0, \sigma^2)$. Using the moment comparison in Lemmas 9 and 10, and applying the upper bound $\text{TV}(\hat{f}, g) \leq \sqrt{\chi^2(\hat{f}\|g)/2}$, we obtain that

$$\text{TV}(\hat{f}, g) \leq O_k \left(\epsilon \sqrt{\log(1/\epsilon)} + \sqrt{\log(1/\delta)/n} \right).$$

The conclusion follows from the triangle inequality. \square

Lemma 11. *Let σ be a constant. If f and g are 1-subgaussian, and $\text{TV}(f, g) \leq \epsilon$, then,*

$$|\mathbb{E}_f[\gamma_r(X, \sigma)] - \mathbb{E}_g[\gamma_r(X, \sigma)]| \leq O_r(\epsilon \sqrt{\log(1/\epsilon)}).$$

Proof. The total variation distance has the following variational representation:

$$\text{TV}(f, g) = \frac{1}{2} \sup_{\|h\|_\infty \leq 1} |\mathbb{E}_f h - \mathbb{E}_g h|. \quad (64)$$

Here the function $\gamma_r(\cdot, \sigma)$ is a polynomial and unbounded, so the above representation cannot be directly applied. Instead, we apply a truncation argument, thanks to the subgaussianity of f and g , and obtain that, for both $X \sim f$ and g (see Claims 1 and 3),

$$\mathbb{E}[\gamma_r(X, \sigma) \mathbf{1}_{\{|X| \geq \alpha\}}] \leq (O(\sqrt{r}))^r \mathbb{E}|X^r \mathbf{1}_{\{|X| \geq \alpha\}}| \leq (O(\alpha\sqrt{r}))^r e^{-\alpha^2/2}.$$

Note that by definition (26), $\gamma_r(x, \sigma)$ on $|x| \leq \alpha$ is at most $(O(\alpha\sqrt{r}))^r$. Applying (64) yields that, for $h(x) = \gamma_r(x, \sigma) \mathbf{1}_{\{|x| \leq \alpha\}}$,

$$|\mathbb{E}_f h - \mathbb{E}_g h| \leq \epsilon (O(\alpha\sqrt{r}))^r.$$

The conclusion follows by choosing $\alpha = O_r(\sqrt{\log(1/\epsilon)})$ and using the triangle inequality. \square

8.3 Proofs for Section 4.1

Proof of Lemma 5. Note that $\tilde{m}_r = \frac{1}{n} \sum_{i=1}^n \gamma_r(X_i, \sigma)$. Then we have

$$\text{var}[\tilde{m}_r] = \frac{1}{n} \text{var}[\gamma_r(X, \sigma)],$$

where $X \sim \nu * N(0, \sigma^2)$. Since the standard deviation of a summation is at most the sum of individual standard deviations, using (26), we have

$$\sqrt{\text{var}[\gamma_r(X, \sigma)]} \leq r! \sum_{j=0}^{\lfloor r/2 \rfloor} \frac{(1/2)^j}{j!(r-2j)!} \sigma^{2j} \sqrt{\text{var}[X^{r-2j}]}.$$

X can be viewed as $U + \sigma Z$ where $U \sim \nu$ and $Z \sim N(0, 1)$ independent of U . Since ν is supported on $[-M, M]$, for any $\ell \in \mathbb{N}$, we have

$$\text{var}[X^\ell] \leq \mathbb{E}[X^{2\ell}] \leq 2^{2\ell-1} (M^{2\ell} + \mathbb{E}|\sigma Z|^{2\ell}) \leq ((2M)^\ell + \mathbb{E}|3\sigma Z|^\ell)^2,$$

where in the last step we used the inequality $\mathbb{E}|Z|^{2\ell} \leq 2^\ell (\mathbb{E}|Z|^\ell)^2$ (see Lemma 12 below). Therefore,

$$\begin{aligned} \sqrt{\text{var}[\gamma_r(X, \sigma)]} &\leq r! \sum_{j=0}^{\lfloor r/2 \rfloor} \frac{(1/2)^j}{j!(r-2j)!} \sigma^{2j} ((2M)^{r-2j} + \mathbb{E}|3\sigma Z|^{r-2j}) \\ &= \mathbb{E}(2M + \sigma Z')^r + \mathbb{E}(3\sigma|Z| + \sigma Z')^r, \end{aligned}$$

where $Z' \sim N(0, 1)$ independent of Z . The conclusion follows by the moments of the standard normal distribution (see [BK80]). \square

Lemma 12. *Let $Z \sim N(0, 1)$. For $\ell \in \mathbb{N}$, we have*

$$\sqrt{\frac{\pi}{8}} \leq \frac{\mathbb{E}|Z|^{2\ell}}{2^\ell (\mathbb{E}|Z|^\ell)^2} \leq \sqrt{\frac{2}{\pi}}.$$

Proof. Direct calculations lead to (see [GR07, 3.461.2-3]):

$$\frac{\mathbb{E}|Z|^{2\ell}}{2^\ell (\mathbb{E}|Z|^\ell)^2} = \begin{cases} \frac{\binom{2\ell}{\ell}}{(\ell/2)^{2\ell}}, & \ell \text{ even,} \\ \frac{\pi \ell}{8^\ell} \binom{2\ell}{\ell} \binom{\ell-1}{\frac{\ell-1}{2}}, & \ell \text{ odd.} \end{cases}$$

Using $\frac{2^n}{\sqrt{2n}} \leq \binom{n}{n/2} \leq 2^n \sqrt{\frac{2}{\pi n}}$ [Ash65, Lemma 4.7.1], we obtain that

$$\begin{aligned} \sqrt{\frac{\pi}{8}} &\leq \frac{\binom{2\ell}{\ell}}{(\ell/2)^{2\ell}} \leq \sqrt{\frac{2}{\pi}}, \\ \frac{\pi}{4} \sqrt{\frac{\ell}{2(\ell-1)}} &\leq \frac{\pi \ell}{8^\ell} \binom{2\ell}{\ell} \binom{\ell-1}{\frac{\ell-1}{2}} \leq \sqrt{\frac{\ell}{2(\ell-1)}}, \end{aligned}$$

which prove this lemma for $\ell \geq 5$. For $\ell \leq 4$ the lemma follows from the above equalities. \square

8.4 Proofs for Section 4.2

Proof of Proposition 3. By scaling it suffices to consider $M = 1$. Without loss of generality assume $\sigma \geq \hat{\sigma}$ and otherwise we can interchange π and $\hat{\pi}$. Let $\tau^2 = \sigma^2 - \hat{\sigma}^2$ and $\nu' = \nu * N(0, \tau^2)$. Similar to (33), we obtain that

$$|m_r(\nu') - m_r(\hat{\nu})| \leq (c\sqrt{k})^{2k} \epsilon, \quad r = 1, \dots, 2k, \quad (65)$$

for some absolute constant c . Using Lemma 13 below yields that $\tau \leq O(\epsilon^{\frac{1}{2k}})$. It follows from Proposition 2 that

$$W_1(\nu', \hat{\nu}) \leq O\left(k^{1.5} \epsilon^{\frac{1}{2k}}\right).$$

The conclusion follows from $W_1(\nu', \nu) \leq O(\tau)$ and the triangle inequality. \square

Lemma 13. *Suppose $\pi = \nu * N(0, \tau^2)$ and π' is k -atomic supported on $[-1, 1]$. Let $\epsilon = \max_{i \in [2k]} |m_i(\pi) - m_i(\pi')|$. Then,*

$$\tau \leq 2(\epsilon/k!)^{\frac{1}{2k}}.$$

Proof. Denote the support of π' by $\{x'_1, \dots, x'_k\}$. Consider the polynomial $P(x) = \prod_{i=1}^k (x - x'_i)^2 = \sum_{i=0}^{2k} a_i x^i$ that is almost surely zero under π' . Since every $|x'_i| \leq 1$, similar to (48), we obtain that

$$\mathbb{E}_\pi[P] = |\mathbb{E}_\pi[P] - \mathbb{E}_{\pi'}[P]| \leq 2^{2k}\epsilon.$$

Since $\pi = \nu * N(0, \tau^2)$, we have

$$\mathbb{E}_\pi[P] \geq \min_x \mathbb{E}[P(x + \tau Z)] \geq \tau^{2k} \min_{y_1, \dots, y_k} \mathbb{E} \left[\prod_i (Z + y_i)^2 \right] = k! \tau^{2k},$$

where $Z \sim N(0, 1)$, and in the last step we used Lemma 14 below. \square

Lemma 14. *Let $Z \sim N(0, 1)$. Then,*

$$\min\{\mathbb{E}[p^2(Z)] : \deg(p) \leq k, p \text{ is monic}\} = k!$$

achieved by $p = H_k$.

Proof. Since p is monic, it can be written as $p = H_k + \sum_{j=0}^{k-1} \alpha_j H_j$, where H_j is the Hermite polynomial (25). By the orthogonality (63), we have $\mathbb{E}[p^2(Z)] = k! + \sum_{j=0}^{k-1} \alpha_j^2 j!$ and the conclusion follows. \square

Proof of Lemma 6. The proof is similar to [Lin89, Theorem 5B]. Let $\hat{\mathbf{M}}_r(\sigma)$ denote the moment matrix associated with the empirical moments of $\gamma_i(X, \sigma)$ for $i \leq 2r$, and let

$$\hat{\sigma}_r = \inf\{\sigma > 0 : \det(\hat{\mathbf{M}}_r(\sigma)) = 0\}. \quad (66)$$

The smallest positive zero of \hat{d}_k is given by $\hat{\sigma}_k$. Direct calculation shows that $\hat{\sigma}_1 = s$. Since the mixture distribution has a density, then almost surely, the empirical distribution has n points of support. By Theorem 5, the matrix $\hat{\mathbf{M}}_r(0)$ is positive definite and thus $\hat{\sigma}_r > 0$ for any $r < n$. For any $q < r$, if $\hat{\mathbf{M}}_r(\sigma)$ is positive definite, then $\hat{\mathbf{M}}_q(\sigma)$ as a leading principal submatrix is also positive definite. Since eigenvalues of $\hat{\mathbf{M}}_r(\sigma)$ are continuous functions of σ , we have $\hat{\sigma}_r > \sigma \Rightarrow \hat{\sigma}_q > \sigma$, and thus

$$\hat{\sigma}_q \geq \hat{\sigma}_r, \quad \forall q < r. \quad (67)$$

In particular, $\hat{\sigma}_k \leq \hat{\sigma}_1$. \square

Proof of Lemma 7. We continue to use the notation in (66). Applying (67) and Lemma 6 yields that

$$0 < \hat{\sigma} = \hat{\sigma}_k \leq \hat{\sigma}_{k-1} \leq \dots \leq \hat{\sigma}_1 = s,$$

and for any $\sigma < \hat{\sigma}_j$, the matrix $\hat{\mathbf{M}}_j(\sigma)$ is positive definite. Since $\det(\hat{\mathbf{M}}_k(\hat{\sigma})) = 0$, then, for some $r \in \{1, \dots, k\}$, we have $\det(\hat{\mathbf{M}}_j(\hat{\sigma})) = 0$ for $j = r, \dots, k$, and $\det(\hat{\mathbf{M}}_j(\hat{\sigma})) > 0$ for $j = 0, \dots, r-1$. By Theorem 5, there exist a r -atomic distribution whose j^{th} moment coincides with $\hat{\gamma}_j(\hat{\sigma})$ for $j \leq 2r$. It suffices to show that $r = k$ almost surely.

Since the mixture distribution has a density, in the following we condition on the event that all samples X_1, \dots, X_n are distinct, which happens almost surely, without loss of generality. We first show that the empirical moments $(\hat{\gamma}_1, \dots, \hat{\gamma}_n)$, where $\hat{\gamma}_j = \frac{1}{n} \sum_i X_i^j$, have a joint density in \mathbb{R}^n . The Jacobian matrix of this transformation is

$$\frac{1}{n} \begin{bmatrix} 1 & & & \\ & 2 & & \\ & & \ddots & \\ & & & n \end{bmatrix} \begin{bmatrix} 1 & \dots & 1 \\ X_1 & \dots & X_n \\ \vdots & \ddots & \vdots \\ X_1^{n-1} & \dots & X_n^{n-1} \end{bmatrix},$$

which is invertible. Since those n samples (X_1, \dots, X_n) have a joint density, then the empirical moments $(\hat{\gamma}_1, \dots, \hat{\gamma}_n)$ also have a joint density.

Suppose, for the sake of contradiction, that $r \leq k-1$. Then $\det(\hat{\mathbf{M}}_{r-1}(\hat{\sigma})) > 0$ and $\det(\hat{\mathbf{M}}_r(\hat{\sigma})) = \det(\hat{\mathbf{M}}_{r+1}(\hat{\sigma})) = 0$. In this case, $\hat{m}_{2r+1}(\hat{\sigma})$ is a deterministic function of $\hat{m}_1(\hat{\sigma}), \dots, \hat{m}_{2r}(\hat{\sigma})$ (see Lemma 31). Since $\hat{\sigma}$ is the smallest positive root of $\hat{d}_r(\sigma) = 0$, it is uniquely determined by $(\hat{\gamma}_1, \dots, \hat{\gamma}_{2r})$. Therefore, $\hat{m}_{2r+1}(\hat{\sigma})$, and thus $\hat{\gamma}_{2r+1}$, are both deterministic functions of $(\hat{\gamma}_1, \dots, \hat{\gamma}_{2r})$, which happens with probability zero, since the sequence $(\hat{\gamma}_1, \dots, \hat{\gamma}_{2r+1})$ has a joint density. Consequently, $r \leq k-1$ with probability zero. \square

The proof of (34) relies on the following result, which obtains a tail probability bound by comparing moments.

Lemma 15. *Let $\epsilon = \max_{i \in [2k]} |m_i(\nu) - m_i(\nu')|$. If either ν or ν' is k -atomic, and ν is supported on $[-1, 1]$, then, for any $t > 1$,*

$$\mathbb{P}[|Y| \geq t] \leq 2^{2k+1} \epsilon (t-1)^{-2k}, \quad Y \sim \nu'.$$

Proof. We only show the upper tail bound $\mathbb{P}[Y \geq t]$. The lower tail bound of Y is equal to the upper tail bound of $-Y$.

- Suppose ν is k -atomic supported on $\{x_1, \dots, x_k\}$. Consider a polynomial $P(x) = \prod_i (x - x_i)^2$ of degree $2k$ that is almost surely zero under ν . Since every $|x_i| \leq 1$, similar to (48), we obtain that

$$\mathbb{E}_{\nu'}[P] = |\mathbb{E}_{\nu}[P] - \mathbb{E}_{\nu'}[P]| \leq 2^{2k} \epsilon.$$

Using Markov inequality, for any $t > 1$, we have

$$\mathbb{P}[Y \geq t] \leq \mathbb{P}[P(Y) \geq P(t)] \leq \frac{\mathbb{E}[P(Y)]}{P(t)} \leq \frac{2^{2k} \epsilon}{(t-1)^{2k}}.$$

- Suppose ν' is k -atomic supported on $\{x_1, \dots, x_k\}$. If those values are all within $[-1, 1]$, then we are done. If there are at most $k-1$ values, denoted by $\{x_1, \dots, x_{k-1}\}$, are within $[-1, 1]$, then we consider a polynomial $P(x) = (x^2 - 1) \prod_i (x - x_i)^2$ of degree $2k$ that is almost surely non-positive under ν . Similar to (48), we obtain that

$$\mathbb{E}_{\nu'}[P] \leq \mathbb{E}_{\nu'}[P] - \mathbb{E}_{\nu}[P] \leq 2^{2k} \epsilon.$$

Since $P \geq 0$ almost surely under ν' , the conclusion follows analogously using Markov inequality. \square

Lemma 16. *Let*

$$\pi = \nu * N(0, \tau^2), \quad \hat{\pi} = \hat{\nu},$$

where ν and ν are both k -atomic, ν is supported on $[-1, 1]$, and $\tau \leq 1$. If $|m_i(\pi) - m_i(\hat{\pi})| \leq \epsilon$ for $i \leq 2k$, then, for any $t \geq \sqrt{18k}$,

$$\mathbb{P}[|\hat{U}| \geq t] \leq 2^{2k+1} \epsilon \left(\frac{t}{\sqrt{18k}} - 1 \right)^{-2k}, \quad \hat{U} \sim \hat{\nu}.$$

Proof. Let g be the $(k+1)$ -point Gauss quadrature of the standard normal distribution. Furthermore, g is supported on $[-\sqrt{4k+6}, \sqrt{4k+6}]$ for some absolute constant c (see the bound on the zeros of Hermite polynomials in [Sze75, p. 129]). Let $G \sim g$, $U \sim \nu$, and $\hat{U} \sim \hat{\nu}$. Denote the maximum absolute value of $U + \tau G$ by M which is at most $1 + \sqrt{4k+6} \leq \sqrt{18k}$ for $k \geq 1$. Applying Lemma 15 to the distributions of $\frac{U + \tau G}{\sqrt{18k}}$ and $\frac{\hat{U}}{\sqrt{18k}}$ yields the desired conclusion. \square

8.5 Proofs for Section 4.3

Proof of Theorem 6. Note that U and \hat{U} are both supported on a set of $2k$ atoms, and the largest cluster of U is of size at most $k - k_0 + 1$. Since different clusters of U are separated by γ , then each atom of either U and \hat{U} is at least $\gamma/2$ away from all but $2k - k_0$ atoms. From the proof of (8), we have $|m_r(\hat{U}) - m_r(U)| < (O(\sqrt{k}))^{2k} \sqrt{\frac{\log(k/\delta)}{n}}$. The conclusion follows from Proposition 4. \square

Proof of Proposition 4. The proof is analogous to the first proof of Proposition 9, apart from a more careful analysis of polynomial coefficients. When each atom is at least γ away from all but at most ℓ' other atoms, the left-hand side of (49) is upper bounded by

$$|F_\nu(t_r) - F_{\nu'}(t_r)| \leq \frac{\ell 4^{\ell-1} \delta}{(t_{r+1} - t_r)^{\ell'} \gamma^{\ell-\ell'-1}},$$

The remaining proof is similar. \square

Proof of Proposition 5. Similar to the proof of Proposition 4, this proof is analogous to Proposition 2 apart from a more careful analysis of polynomial coefficients. When every $t \in \mathbb{R}$ is at least γ away from all but k' atoms, the left-hand sides of (60) and (62) are upper bounded by

$$\begin{aligned} \int (g(t) \wedge 1) dt &\leq 4k \left(\frac{2^{2k} \delta}{\gamma^{2(k-k')}} \right)^{1/(2k')}, \\ \int (f(t) \wedge 1) dt &\leq 4k \left(\frac{k 4^{2k} \delta}{\gamma^{2(k-k')}} \right)^{1/(2k')}. \end{aligned}$$

The remaining proof is similar. \square

Proof of Proposition 6. The proof is similar to Proposition 3, except that moment comparison theorem Proposition 2 is replaced by its adaptive version Proposition 5. Recall (65):

$$|m_r(\nu') - m_r(\hat{\nu})| \leq (c\sqrt{k})^{2k} \epsilon, \quad r = 1, \dots, 2k,$$

where $\nu' = \nu * N(0, \tau^2)$ and $\tau^2 = |\sigma^2 - \hat{\sigma}^2|$. Since $\hat{\nu} * N(0, 1)$ has k_0 γ -separated clusters, any $t \in \mathbb{R}$ can be $\gamma/2$ close to at most $k - k_0 + 1$ atoms of $\hat{\nu}$. Applying Proposition 5 yields that

$$W_1(\nu', \hat{\nu}) \leq 8k \left(\frac{k(4c\sqrt{k})^{2k} \epsilon}{(\gamma/2)^{2(k_0-1)}} \right)^{\frac{1}{2(k-k_0+1)}}.$$

Using Lemma 17 below yields that $\tau \leq O_k(W_1(\nu', \hat{\nu}))$. The conclusion follows from $W_1(\nu', \nu) \leq O(\tau)$ and the triangle inequality. \square

Lemma 17. *Suppose $\pi = \nu * N(0, \tau^2)$ and π' is k -atomic. Then*

$$\tau \leq O_k(W_1(\pi, \pi')).$$

Proof. In this proof we write $W_1(X, Y) = W_1(P_X, P_Y)$. Let $Z \sim N(0, 1)$, $U \sim \nu$, and $U' \sim \pi'$. For any $x \in \mathbb{R}$, we have

$$W_1(x + \tau Z, U') = \tau W_1(Z, (U' - x)/\tau) \geq c_k \tau,$$

where $c_k = \inf\{W_1(Z, Y) : Y \text{ is } k\text{-atomic}\}$ ¹⁰. For any coupling between $U + \tau Z$ and U' ,

$$\mathbb{E}|U + \tau Z - U'| = \mathbb{E}[\mathbb{E}[|U + \tau Z - U'| | U]] \geq c_k \tau. \quad \square$$

¹⁰ We can prove that $c_k \geq \Omega(1/k)$ using the dual formula (21).

Proof of Theorem 6. By scaling it suffices to consider $M = 1$. Recall that the Gaussian mixture is assumed to have k_0 γ -separated clusters in the sense of Definition 1, that is, there exists a partition S_1, \dots, S_{k_0} of $[k]$ such that $|\mu_i - \mu_{i'}| \geq \gamma$ for any $i \in S_\ell$ and $i' \in S_{\ell'}$ such that $\ell \neq \ell'$. Denote the union of the support sets of ν and $\hat{\nu}$ by \mathcal{S} . Each atom in \mathcal{S} is at least $\gamma/2$ away from at least $k_0 - 1$ other atoms. Then (35) follows from Proposition 4 with $\ell = 2k$ and $\ell' = (2k - 1) - (k_0 - 1)$. \square

8.6 Proofs for Section 4.4

Lemma 18. *Assume in the Gaussian mixture (1) $w_i \geq \epsilon$, $\sigma = 1$. Suppose $L = \sqrt{c \log n}$ in Algorithm 3. Then, with probability at least $1 - ke^{-n'\epsilon} - n^{-(\frac{c}{8}-1)}$, the following holds:*

- $\ell_j \leq 3kL$ for every j .
- Let $X_i = U_i + Z_i$ for $i \in [n]$, where $U_i \sim \nu$ is the latent variable and $Z_i \sim N(0, 1)$. Then, $|Z_i| \leq 0.5L$ for every $i \in [n]$; $X_i \in I_j$ if and only if $U_i \in I_j$.

Proof. By the union bound, with probability $1 - ke^{-n'\epsilon} - n^{-(\frac{c}{8}-1)}$, the following holds:

- $|Z_i| \leq 0.5L$ for every $i \in [n]$.
- For every $j \in [k]$, there exists $i \leq n'$ such that $U_i = \mu_j$.

Recall the disjoint intervals $I_1 \cup \dots \cup I_s = \cup_{i=1}^{n'} [X_i \pm L]$. Then, we obtain that

$$\bigcup_{j=1}^k [\mu_j \pm 0.5L] \subseteq I_1 \cup \dots \cup I_s \subseteq \bigcup_{j=1}^k [\mu_j \pm 1.5L].$$

The total length of all intervals is at most $3kL$. Since $|Z_i| \leq 0.5L$, $X_i = U_i + Z_i$ is in the same interval as U_i . \square

Proof of Theorem 7. Since $n' \geq \Omega(\frac{\log(k/\delta)}{\epsilon})$, applying Lemma 18 yields that, with probability at least $1 - \frac{\delta}{3} - n^{-\Omega(1)}$, the following holds:

- $\ell_j \leq O(kL)$ for every j .
- Let $X_i = U_i + \sigma Z_i$ for $i \in [n]$ as in Lemma 18. Then, $|Z_i| \leq 0.5L$ for every $i \in [n]$; $X_i \in I_j$ if and only if $U_i \in I_j$.

The intervals I_1, \dots, I_s are independent of every C_j and are treated as deterministic in the remaining proof. We first evaluate the expected moments of samples in C_j , conditioned on $|Z_i| \leq L' \triangleq 0.5L$. Let $X = U + \sigma Z$ where $U \sim \nu$ and $Z \sim N(0, 1)$. Then,

$$\mathbb{E}[(X - c_j)^r | X \in I_j, |Z| \leq L'] = \mathbb{E}[(X - c_j)^r | U \in I_j, |Z| \leq L'] = \mathbb{E}[(U'_j + \sigma Z)^r | |Z| \leq L'],$$

where $U'_j = U_j - c_j$, and $U_j \sim P_{U|U \in I_j}$. Since $|U'_j| \leq O(kL)$ and $L' = \Theta(\sqrt{\log n})$, the right-hand side differs from the unconditional moment by (see Lemma 32 in Appendix C)

$$|\mathbb{E}[(U'_j + \sigma Z)^r | |Z| \leq L'] - \mathbb{E}[(U'_j + \sigma Z)^r]| \leq (kL\sigma\sqrt{r})^r n^{-\Omega(1)}, \quad r = 1, \dots, 2k - 1,$$

which is less than n^{-1} when $k \leq O(\frac{\log n}{\log \log n})$. Therefore, the accuracy of empirical moments in (27), (32) and thus Theorem 1 are all applicable. Since $w_i \geq \epsilon$, with probability at least $1 - \frac{\delta}{3}$, each C_j contains $\Omega(n\epsilon)$ samples, and applying Theorem 1 yields that, with probability $1 - \frac{\delta}{3}$,

$$W_1(\hat{\nu}_j, \nu_j) \leq \begin{cases} O(Lk^{2.5} (\frac{n\epsilon}{\log(3k/\delta)})^{-\frac{1}{4k-2}}), & \sigma \text{ known,} \\ O(Lk^3 (\frac{n\epsilon}{\log(3k/\delta)})^{-\frac{1}{4k}}), & \sigma \text{ unknown,} \end{cases}$$

for every j , where ν_j denotes the distribution of U'_j and $\hat{\nu}_j$ is the estimate in Theorem 1. Using the weights threshold $\tau = \epsilon/(2k)$, and applying Lemma 19 below, we obtain that

$$d_H(\text{supp}(\hat{\nu}_j), \text{supp}(\nu_j)) \leq \frac{W_1(\hat{\nu}_j, \nu_j)}{\epsilon/(2k)}.$$

The conclusion follows. \square

Lemma 19. *Let ν be a discrete distribution whose atom has at least ϵ probability. Let S_ν and $S_{\hat{\nu}}$ denote the support sets of ν and $\hat{\nu}$, respectively. For $\hat{S} \subseteq S_{\hat{\nu}}$,*

$$d_H(S_\nu, \hat{S}) \leq \frac{W_1(\nu, \hat{\nu})}{(\min_{y \in \hat{S}} \hat{\nu}(y)) \wedge (\epsilon - \hat{\nu}(\hat{S}^c))_+}.$$

Proof. This is a generalization of Lemma 2 in the sense that the minimum weight of $\hat{\nu}$ is unknown. For any coupling P_{XY} such that $X \sim \nu$ and $Y \sim \hat{\nu}$, for any $y \in \hat{S}$,

$$\mathbb{E}|X - Y| \geq \hat{\nu}(y) \mathbb{E}[|X - Y| | Y = y] \geq \epsilon_1 \min_{x \in S_\nu} |x - y|,$$

where $\epsilon_1 = \min_{y \in \hat{S}} \hat{\nu}(y)$. Note that $\mathbb{P}[Y \notin \hat{S}, X = x] \leq \hat{\nu}(\hat{S}^c)$ and $\nu(x) \geq \epsilon$ for any $x \in S_\nu$. Then we have $\mathbb{P}[Y \in \hat{S}, X = x] \geq (\epsilon - \hat{\nu}(\hat{S}^c))_+ \triangleq \epsilon_2$, and thus

$$\mathbb{E}|X - Y| \geq \epsilon_2 \mathbb{E}[|X - Y| | X = x, Y \in \hat{S}] \geq \epsilon_2 \min_{y \in \hat{S}} |x - y|.$$

Using the definition of d_H in (23), the proof is complete. \square

8.7 Proofs for Section 5

Proof of Lemma 8. Let $U \sim \nu$ and $U' \sim \nu'$. If ν and ν' are ϵ -subgaussian, then $\text{var}[U'] \leq \epsilon^2$, and $\mathbb{E}|U|^p, \mathbb{E}|U'|^p \leq 2(\epsilon\sqrt{p/e})^p$ [BK80]. Applying the χ^2 upper bound from moment difference in Lemma 9 yields that

$$\chi^2(\nu * N(0, 1) || \nu' * N(0, 1)) \leq e^{\epsilon^2/2} \sum_{j \geq \ell+1} \frac{16\epsilon^{2j}}{\sqrt{2\pi}^j},$$

where we used Stirling's approximation $n! > \sqrt{2\pi n}(n/e)^n$. If ν and ν' are supported on $[-\epsilon, \epsilon]$, the conclusion is obtained similarly by using $\mathbb{E}|U|^p, \mathbb{E}|U'|^p \leq \epsilon^p$. \square

Proof of Proposition 7. Let ν and ν' be the optimal pair of distributions for (38). Applying Lemma 8 yields that

$$\chi^2(\nu * N(0, 1) || \nu' * N(0, 1)) \leq c \left(\frac{e\epsilon^2}{2k-1} \right)^{2k-1},$$

for some absolute constant c . The two mixing distributions satisfy (see Lemma 20 below)

$$W_1(\nu, \nu') \geq \Omega(\epsilon/\sqrt{k}).$$

The conclusion follows by choosing $\epsilon = c'\sqrt{kn}^{-\frac{1}{4k-2}}$ for some absolute constant c' and applying Le Cam's method [LC86]. \square

Lemma 20.

$$\sup\{W_1(\nu, \nu') : \mathbf{m}_\ell(\nu) = \mathbf{m}_\ell(\nu'), \nu, \nu' \text{ on } [-1, 1]\} = \Theta(\beta/(\ell + 1)).$$

Furthermore, the supremum is $\frac{\beta(\pi - o(1))}{\ell + 1}$ as $\ell \rightarrow \infty$, and is achieved by two distributions whose support sizes differ by at most one and sum up to $\ell + 2$.

Proof. It suffices to prove for $\beta = 1$. Using the dual characterization of the W_1 distance in Section 2.3, the supremum is equal to

$$\sup_{f:1\text{-Lipschitz}} \sup\{\mathbb{E}_\nu f - \mathbb{E}_{\nu'} f : \mathbf{m}_\ell(\nu) = \mathbf{m}_\ell(\nu'), \nu, \nu' \text{ on } [-\beta, \beta]\}.$$

Using the duality between moment matching and best polynomial approximation (see [WY16, Appendix E]), the optimal value is further equal to

$$2 \sup_{f:1\text{-Lipschitz}} \inf_{P:\text{degree} \leq \ell} \sup_{|x| \leq 1} |f(x) - P(x)|.$$

The above value is the best uniform approximation error over 1-Lipschitz functions, a well-studied quantity in the approximation theory (see, e.g., [Bus11, section 4.1]), and thus the optimal values in the lemma are obtained. A pair of optimal distributions are supported on the maxima and the minima of $P^* - f^*$, respectively, where f^* is the optimal 1-Lipschitz function and P^* is the best polynomial approximation for f^* . The numbers of maxima and minima differ by at most one by Chebyshev's alternating theorem (see, e.g., [Tim63, p. 54]). \square

Proof of Proposition 8. Let $\nu = N(0, \epsilon^2)$ and ν' be its k -point Gauss quadrature. Then $\mathbf{m}_{2k-1}(\nu) = \mathbf{m}_{2k-1}(\nu')$ and ν and ν' are both ϵ -subgaussian (see Lemma 21). Applying Lemma 8 yields that

$$\chi^2(\nu * N(0, 1) \| \nu' * N(0, 1)) \leq O(\epsilon^{4k}).$$

Note that $\nu * N(0, 1) = N(0, 1 + \epsilon^2)$ is a valid Gaussian mixture distribution (with single zero mean component). Between the above two mixture models, the variance parameters differ by ϵ^2 ; the mean parameters satisfy $W_1(g_k, \delta_0) \geq \Omega(\epsilon/\sqrt{k})$ (see Lemma 22). The conclusion follows by choosing $\epsilon = cn^{-\frac{1}{4k}}$ for some absolute constant c applying applying Le Cam's method [LC86]. \square

Lemma 21. *Let g_k be the k -point Gauss quadrature of $N(0, \sigma^2)$. For $j \geq 2k$, we have $m_j(g_k) \leq m_j(N(0, \sigma^2))$ when j is even, and $m_j(g_k) = m_j(N(0, \sigma^2)) = 0$ otherwise. In particular, g_k is σ -subgaussian.*

Proof. By scaling it suffices to consider $\sigma = 1$. Let $\nu = N(0, 1)$. If j is odd, $m_j(g_k) = m_j(\nu) = 0$ by symmetry. If $j \geq 2k$ and j is even, the conclusion follows from the integral representation of the error term of Gauss quadrature (see, e.g., [SB02, Theorem 3.6.24]):

$$m_j(\nu) - m_j(g_k) = \frac{f^{(2k)}(\xi)}{(2k)!} \int \pi_k^2(x) d\nu(x),$$

for some $\xi \in \mathbb{R}$; here $f(x) = x^j$, $\{x_1, \dots, x_k\}$ is the support of g_k , and $\pi_k(x) \triangleq \prod_i (x - x_i)$. Consequently, g_k is 1-subgaussian [BK80, Lemma 2]. \square

Lemma 22. *Let g_k be the k -point Gauss quadrature of $N(0, 1)$. Then*

$$\mathbb{E}_{g_k} |X| \geq (4k + 2)^{-1/2}, \quad k \geq 2.$$

Proof. Let $G_k \sim g_k$. Note that $|G_k| \leq \sqrt{4k+2}$ using the bound on the zeros of Hermite polynomials [Sze75, p. 129]. The conclusion follows from $1 = \mathbb{E}[G_k^2] \leq \mathbb{E}|G_k|\sqrt{4k+2}$. \square

Lemma 23. *Let g_k be the k -point Gauss quadrature of $N(0,1)$. Then $\mathbb{E}_{g_k}[H_j] = 0$ for $j = 1, \dots, 2k-1$, and $\mathbb{E}_{g_k}[H_{2k}] = -k!$, where H_j is the Hermite polynomial of degree j (see (25)).*

Proof. Let $Z \sim N(0,1)$ and $G_k \sim g_k$. By orthogonality of Hermite polynomials (63) we have $\mathbb{E}[H_j(Z)] = 0$ for all $j \geq 1$ and thus $\mathbb{E}[H_j(G_k)] = 0$ for $j = 1, \dots, 2k-1$. Expand $H_k^2(x)$ as

$$H_k^2(x) = H_{2k}(x) + a_{2k-1}H_{2k-1}(x) + \dots + a_1H_1(x) + a_0.$$

Since G_k is supported on the zeros of H_k , we have $0 = \mathbb{E}[H_k^2(G_k)] = \mathbb{E}[H_{2k}(G_k)] + a_0$. The conclusion follows from $k! = \mathbb{E}[H_k^2(Z)] = a_0$ (see (63)). \square

A Standard form of the semidefinite programming (24)

Given an arbitrary vector $\tilde{m} = (\tilde{m}_1, \dots, \tilde{m}_r)$, (24) computes its projection onto the moment space $\mathcal{M}_r([a, b])$. By introducing an auxiliary scalar variable t satisfying $t \geq \|x\|_2^2$, (24) is equivalent to

$$\begin{aligned} \min \quad & t - 2\langle \tilde{m}, x \rangle + \|\tilde{m}\|_2^2, \\ \text{s.t.} \quad & t \geq \|x\|_2^2, \quad x \text{ satisfies (14)}. \end{aligned}$$

This is a semidefinite programming with decision variable (x, t) , since the constraint $t \geq \|x\|_2^2$ is equivalent to $\begin{bmatrix} t & x^\top \\ x & I \end{bmatrix} \succeq 0$ using the Schur complement (see, e.g., [VB96]).

B Gauss quadrature algorithm

Here we describe a basic algorithm to find the Gauss quadrature in Algorithm 5 [GW69].

C Auxiliary lemmas

Lemma 24. *Let $t_1 \leq t_2 \leq \dots$ be an ordered sequence (not necessarily distinct) and $t_r < t < t_{r+1}$. Let $f(x) = \mathbf{1}_{\{x \leq t\}}$. Then*

$$f[t_i, \dots, t_j] = (-1)^{i-r} \sum_{L \in \mathcal{L}(i,j)} \prod_{(x,y) \in L} \frac{1}{t_x - t_y}, \quad i \leq r < r+1 \leq j, \quad (68)$$

where $\mathcal{L}(i, j)$ is the set of lattice paths from $(r, r+1)$ to (i, j) using steps $(0, 1)$ and $(-1, 0)$ ¹¹. Furthermore,

$$|f[t_1, \dots, t_i]| \leq \frac{\binom{i-2}{r-1}}{(t_{r+1} - t_r)^{i-1}}, \quad i \geq r+1. \quad (69)$$

¹¹ Formally, for $a, b \in \mathbb{N}^2$, a lattice path from a to b using a set of steps S is a sequence $a = x_1, x_2, \dots, x_n = b$ with all increments $x_{j+1} - x_j \in S$. In the matrix representation shown in the proof, this corresponds to a path from $a_{r,r+1}$ to $a_{i,j}$ going up and right. This path consists of entries (i, j) such that $i \leq r < r+1 \leq j$, and thus in (68) we always have $t_x \leq t_r < t_{r+1} \leq t_y$.

Algorithm 5 Quadrature rule

Input: a vector of $2k - 1$ moments (m_1, \dots, m_{2k-1}) .

Output: nodes $x = (x_1, \dots, x_k)$ and weights $w = (w_1, \dots, w_k)$.

Define the following degree- k polynomial P

$$P(x) = \det \begin{bmatrix} 1 & m_1 & \cdots & m_k \\ \vdots & \vdots & \ddots & \vdots \\ m_{k-1} & m_k & \cdots & m_{2k-1} \\ 1 & x & \cdots & x^k \end{bmatrix}.$$

Let the nodes (x_1, \dots, x_k) be the roots of the polynomial P .

Let the weights $w = (w_1, \dots, w_k)$ be

$$w = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ x_1 & x_2 & \cdots & x_k \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{k-1} & x_2^{k-1} & \cdots & x_k^{k-1} \end{bmatrix}^{-1} \begin{bmatrix} 1 \\ m_1 \\ \vdots \\ m_{k-1} \end{bmatrix}.$$

Proof. Denote by $a_{i,j} = f[t_i, \dots, t_j]$ when $i \leq j$. It is obvious that $a_{i,i} = 1$ for $i \leq r$; $a_{i,i} = 0$ for $i \geq r + 1$; $a_{i,j} = 0$ for both $i < j \leq r$ and $j > i \geq r + 1$. For $i \leq r < r + 1 \leq j$, the values can be obtained recursively by

$$a_{i,j} = \frac{a_{i,j-1} - a_{i+1,j}}{t_i - t_j}. \quad (70)$$

The above recursion can be represented in Neville's diagram as in Section 2.2. In this proof, it is equivalently represented in a upper triangular matrix as follows:

$$\begin{bmatrix} 1 & 0 & \cdots & 0 & a_{1,r+1} & \cdots \\ & 1 & \ddots & \vdots & \vdots & \\ & & 1 & 0 & a_{r-1,r+1} & \cdots \\ & & & 1 & a_{r,r+1} & \cdots \\ & & & & 0 & \cdots & 0 \\ & 0 & & & & \ddots & \vdots \\ & & & & & & 0 \end{bmatrix}.$$

In the matrix, every $a_{i,j}$ is calculated using the two values left to it and below it. The values on any path from $a_{r,r+1}$ to $a_{i,j}$ going up and right will contribute to the formula of $a_{i,j}$ in (68). The paths consist of two types: first go to $a_{i,j-1}$ and then go right; first go to $a_{i+1,j}$ and then go up. Formally, $\{L, (i, j) : L \in \mathcal{L}_{i,j-1}\} \cup \{L, (i, j) : L \in \mathcal{L}_{i+1,j}\} = \mathcal{L}_{i,j}$. This will be used in the proof of (68) by induction present next. The base cases (r^{th} row and $(r + 1)^{\text{th}}$ column) can be directly computed:

$$a_{r,j} = \prod_{v=r+1}^j \frac{1}{t_r - t_v}, \quad a_{i,r+1} = (-1)^{i-r} \prod_{v=i}^r \frac{1}{t_v - t_{r+1}}.$$

Suppose (68) holds for both $a_{i,j-1}$ and $a_{i+1,j}$. Then $a_{i,j}$ can be evaluated by

$$\begin{aligned} a_{i,j} &= \frac{(-1)^{i-r}}{t_i - t_j} \left(\sum_{L \in \mathcal{L}(i,j-1)} \prod_{(x,y) \in L} \frac{1}{t_x - t_y} + \sum_{L \in \mathcal{L}(i+1,j)} \prod_{(x,y) \in L} \frac{1}{t_x - t_y} \right) \\ &= (-1)^{i-r} \left(\sum_{L \in \mathcal{L}(i,j)} \prod_{(x,y) \in L} \frac{1}{t_x - t_y} \right). \end{aligned}$$

For the upper bound in (69), we note that $|\mathcal{L}(i,j)| \leq \binom{(r-1)+(i-(r+1))}{r-1}$ in (68), and each summand is at most $\frac{1}{(t_{r+1}-t_r)^{i-1}}$ in magnitude. \square

Lemma 25. *Let*

$$P(x) = \prod_{i=1}^{\ell} (x - x_i) = \sum_{j=0}^{\ell} a_j x^j.$$

If $|x_i| \leq \beta$ for every i , then

$$|a_j| \leq \binom{\ell}{j} \beta^{\ell-j}.$$

Proof. P can be explicitly expanded and we obtain that

$$a_{\ell-j} = (-1)^j \sum_{\{i_1, i_2, \dots, i_j\} \subseteq [\ell]} x_{i_1} \cdot x_{i_2} \cdot \dots \cdot x_{i_j}.$$

The summation consists of $\binom{\ell}{j}$ terms, and each term is at most β^j in magnitude. \square

Lemma 26. *If $|\mathbb{E}[X^\ell] - \mathbb{E}[X'^\ell]| \leq (C\sqrt{\ell})^\ell \epsilon$ for $\ell = 1, \dots, r$, then, for γ_r in (26),*

$$|\mathbb{E}[\gamma_r(X, \sigma)] - \mathbb{E}[\gamma_r(X', \sigma)]| \leq \epsilon \left((2\sigma\sqrt{r}/e)^r + (2C\sqrt{r})^r \right).$$

Proof. Note that $|\mathbb{E}[X^\ell] - \mathbb{E}[X'^\ell]| \leq \mathbb{E}[|C\sqrt{e}Z'|^r] \epsilon$ by Lemma 27 below, where $Z' \sim N(0, 1)$. Then,

$$|\mathbb{E}[\gamma_r(X, \sigma)] - \mathbb{E}[\gamma_r(X', \sigma)]| \leq \sum_{i=0}^{\lfloor r/2 \rfloor} \frac{r! \sigma^{2i}}{i!(r-2i)! 2^i} \mathbb{E}[|C\sqrt{e}Z'|^r] \epsilon = \epsilon \cdot \mathbb{E}[(\sigma Z + |C\sqrt{e}Z'|)^r],$$

where $Z \sim N(0, 1)$ independent of Z' . Applying $(a+b)^r \leq 2^{r-1}(|a|^r + |b|^r)$ and Lemma 27 completes the proof. \square

Lemma 27.

$$(p/e)^{p/2} \leq \mathbb{E}|Z|^p \leq \sqrt{2}(p/e)^{p/2}, \quad p \geq 0.$$

Proof. Note that

$$\frac{\mathbb{E}|Z|^p}{(p/e)^{p/2}} = \frac{2^{p/2} \Gamma(\frac{p+1}{2})}{\sqrt{\pi} (p/e)^{p/2}} \triangleq f(p), \quad \forall p \geq 0.$$

Since $f(0) = 1$ and $f(\infty) = \sqrt{2}$, it suffices to show that f is increasing in $[0, \infty)$. Equivalently, $\frac{x}{2} \log \frac{2e}{x} + \log \Gamma(\frac{x+1}{2})$ is increasing, which is equivalent to $\psi(\frac{x+1}{2}) \geq \log \frac{x}{2}$ by the derivative, where $\psi(x) \triangleq \frac{d}{dx} \log \Gamma(x)$. The last inequality holds for any $x > 0$ (see, e.g., [DS16, (3)]). \square

Lemma 28. *Let $r \geq 2$. Then,*

$$\int \left(\frac{\delta}{\prod_{i=1}^r |t - x_i|} \wedge 1 \right) dt \leq 4r\delta^{\frac{1}{r}}.$$

Proof. Without loss of generality, let $x_1 \leq x_2 \leq \dots \leq x_r$. Note that

$$\int \left(\frac{\delta}{\prod_{i=1}^r |t - x_i|} \wedge 1 \right) dt = \int_{-\infty}^{x_1} + \int_{x_1}^{\frac{x_1+x_2}{2}} + \int_{\frac{x_1+x_2}{2}}^{x_2} + \dots + \int_{x_r}^{\infty}.$$

There are $2r$ terms in the summation and each term can be upper bounded by

$$\int_{x_i}^{\infty} \left(\frac{\delta}{|t - x_i|^r} \wedge 1 \right) dt = \int_0^{\infty} \left(\frac{\delta}{t^r} \wedge 1 \right) dt = \frac{r}{r-1} \delta^{\frac{1}{r}}.$$

The conclusion follows. \square

Lemma 29. *Given any $2k$ distinct points $x_1 < x_2 < \dots < x_{2k}$, there exist two distributions ν and ν' supported on $\{x_1, x_3, \dots, x_{2k-1}\}$ and $\{x_2, x_4, \dots, x_{2k}\}$, respectively, such that $\mathbf{m}_{2k-2}(\nu) = \mathbf{m}_{2k-2}(\nu')$.*

Proof. Consider the following linear equation

$$\begin{pmatrix} 1 & 1 & \dots & 1 \\ x_1 & x_2 & \dots & x_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{2k-2} & x_2^{2k-2} & \dots & x_{2k}^{2k-2} \end{pmatrix} \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_{2k} \end{pmatrix} = 0,$$

This underdetermined system has a non-zero solution. Let w be a solution with $\|w\|_1 = 2$. Since all weights sum up to zero, then positive weights in w sum up to 1 and negative weights sum up to -1 . Let one distribution be supported on x_i with weight w_i for $w_i > 0$, and the other one be supported on the remaining x_i 's with the corresponding weights $|w_i|$. Then these two distribution match the first $2k - 2$ moments.

It remains to show that the weights in any non-zero solution have alternating signs. Note that all weights are non-zero: if one w_i is zero, then the solution must be all zero since the Vandermonde matrix is of full row rank. To verify the signs of the solution, without loss generality, assume that $w_{2k} = -1$ and then

$$\begin{pmatrix} 1 & \dots & 1 \\ x_1 & \dots & x_{2k-1} \\ \vdots & \ddots & \vdots \\ x_1^{2k-2} & \dots & x_{2k-1}^{2k-2} \end{pmatrix} \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_{2k-1} \end{pmatrix} = \begin{pmatrix} 1 \\ x_{2k} \\ \vdots \\ x_{2k}^{2k-2} \end{pmatrix}.$$

The solution has an explicit formula that $w_i = P_i(x_{2k})$ where P_i is an interpolating polynomial of degree $2k - 2$ satisfying $P_i(x_j) = 1$ for $j = i$ and $P_i(x_j) = 0$ for all other $j \leq 2k - 1$. Specifically, we have $w_i = \frac{\prod_{j \neq i, j \leq 2k-1} (x_{2k} - x_j)}{\prod_{j \neq i, j \leq 2k-1} (x_i - x_j)}$, which satisfies $w_i > 0$ for odd i and $w_i < 0$ for even i . The proof is complete. \square

Lemma 30 (Non-existence of an unbiased estimator). *Let $X_1, \dots, X_m \stackrel{i.i.d.}{\sim} pN(s, \sigma^2) + (1-p)N(t, \sigma^2) = \nu * N(0, \sigma^2)$, where $\nu = p\delta_s + (1-p)\delta_t$ and p, s, t, σ are the unknown parameters. For any $r \geq 2$, unbiased estimator for the r^{th} moments of ν , namely, $ps^r + (1-p)t^r$, does not exist.*

Proof. We will derive a few necessary conditions for an unbiased estimator, denoted by $g(x_1, \dots, x_m)$, and then arrive at a contradiction. Expand the function under the Hermite basis

$$g(x_1, \dots, x_m) = \sum_{n_1, \dots, n_m \geq 0} \alpha_{n_1, \dots, n_m} \prod_i H_{n_i}(x_i),$$

and denote by $T_n(\mu, \sigma^2)$ the expected value of the Hermite polynomial $\mathbb{E}H_n(X)$ under Gaussian model $X \sim N(\mu, \sigma^2)$. Without loss of generality we may assume that the function g and the coefficients α are symmetric (permutation invariant). Then, the expected value of the function g under $\sigma^2 = 1$ is

$$\mathbb{E}[g(X_1, \dots, X_m)] = \sum_{n_1, \dots, n_m \geq 0} \alpha_{n_1, \dots, n_m} \prod_i (ps^{n_i} + (1-p)t^{n_i}), \quad (71)$$

which can be viewed as a polynomial in p , whereas the target is $ps^r + (1-p)t^r$, a linear function in p . Matching polynomial coefficients yields that

$$\sum_{n_1 + \dots + n_m \geq 0} \alpha_{n_1, \dots, n_m} t^{n_1 + \dots + n_m} = t^r, \quad (72)$$

$$\sum_{n_1 + \dots + n_m \geq 0} \alpha_{n_1, \dots, n_m} (s^{n_1} - t^{n_1}) t^{n_2 + \dots + n_m} \cdot m = s^r - t^r, \quad (73)$$

$$\sum_{n_1 + \dots + n_m \geq 0} \alpha_{n_1, \dots, n_m} \prod_{i=1}^j (s^{n_i} - t^{n_i}) t^{n_{j+1} + \dots + n_m} = 0, \quad \forall j = 2, \dots, m, \quad (74)$$

where we used the symmetry of the coefficients α . The equality (74) with $j = m$ yields that $\alpha_{n_1, \dots, n_m} \neq 0$ only if at least one n_i is zero; then (74) with $j = m - 1$ yields that $\alpha_{n_1, \dots, n_m} \neq 0$ only if at least two n_i are zero; repeating this for $j = m, m - 1, \dots, 2$, we obtain that α_{n_1, \dots, n_m} is nonzero only if at most one n_i is nonzero. Then the equality (73) implies that α_{n_1, \dots, n_m} is nonzero only if exactly one $n_i = r$ and the coefficient is necessarily $\frac{1}{m}$. Therefore, it is necessary that the symmetric function is $g(x_1, \dots, x_m) = \frac{1}{m} \sum_{i=1}^m H_r(x_i)$. However, this function is biased when $\sigma^2 \neq 1$. \square

Lemma 31. *Given a sequence $\gamma_1, \gamma_2, \dots$, let \mathbf{H}_j denote the Hankel matrix of order $j + 1$ using $1, \gamma_1, \dots, \gamma_{2j}$. Suppose $\det(\mathbf{H}_{r-1}) \neq 0$, and $\det(\mathbf{H}_r) = \det(\mathbf{H}_{r+1}) = 0$. Then,*

$$\gamma_{2r+1} = (\gamma_{r+1}, \dots, \gamma_{2r})(\mathbf{H}_{r-1})^{-1}(\gamma_r, \dots, \gamma_{2r-1})^\top.$$

Proof. The matrices \mathbf{H}_{r-1} and \mathbf{H}_r are both of rank r by their determinants. We first show that the rank of $[\mathbf{H}_r, v]$, which is the first $r + 1$ rows of \mathbf{H}_{r+1} and is of dimension $(r + 1) \times (r + 2)$, is also r , where $v \triangleq (\gamma_{r+1}, \dots, \gamma_{2r+1})^\top$. Suppose the rank is $r + 1$. Then v cannot be in the image of \mathbf{H}_r . By symmetry of the Hankel matrix, the transpose of $[\mathbf{H}_r, v]$ is the first $r + 1$ columns of \mathbf{H}_{r+1} . Those $r + 1$ columns are linearly independent when its rank is $r + 1$. Since $\det(\mathbf{H}_{r+1}) = 0$, then the last column of \mathbf{H}_{r+1} must be in the image of the first $r + 1$ columns, which is a contradiction.

Since first r columns of \mathbf{H}_{r+1} are linearly independent, and the first $r + 1$ columns of \mathbf{H}_{r+1} are of rank r . Then the $(r + 1)^{\text{th}}$ column of \mathbf{H}_{r+1} is in the image of the first r columns, and thus γ_{2r+1} is a linear combination of $\gamma_{r+1}, \dots, \gamma_{2r}$. Since \mathbf{H}_{r-1} is of full rank, the coefficients can be uniquely determined by $(\mathbf{H}_{r-1})^{-1}(\gamma_r, \dots, \gamma_{2r-1})^\top$. \square

Claim 1. *If $|x| > 1$, then*

$$|H_r(x)| \leq (\sqrt{cr}|x|)^r,$$

for some absolute constant c .

Proof. For $|x| > 1$,

$$\begin{aligned} |H_r(x)| &\leq r! \sum_{j=0}^{\lfloor r/2 \rfloor} \frac{(1/2)^j}{j!(r-2j)!} |x|^r = |x|^r |H_n(\mathbf{i})| = |x|^r |\mathbb{E}(\mathbf{i} + \mathbf{i}Z)^r| \\ &= |x|^r |\mathbb{E}(1 + Z)^r| \leq (\sqrt{cr}|x|)^r, \end{aligned}$$

for some absolute constant c , where $\mathbf{i} = \sqrt{-1}$ and $Z \sim N(0, 1)$. \square

Claim 2. Let $Z \sim N(0, 1)$.

$$\mathbb{P}[Z > M] \leq e^{-\frac{M^2}{2}}.$$

Proof. Applying Chernoff bound yields that

$$\mathbb{P}[Z > M] \leq \exp(-\sup_t (tM - t^2/2)) = \exp(-M^2/2). \quad \square$$

Claim 3. For r even, and $M \geq 1$,

$$\mathbb{E}[Z^r \mathbf{1}_{\{|Z| > M\}}] \leq r(O(\sqrt{r}))^r \left(M^{r-1} e^{-\frac{M^2}{2}} \right).$$

Proof. Applying integral by parts yields that

$$\int_M^\infty x^r e^{-\frac{x^2}{2}} dx = M^{r-1} e^{-\frac{M^2}{2}} + (r-1)M^{r-3} e^{-\frac{M^2}{2}} + (r-1)(r-3)M^{r-5} e^{-\frac{M^2}{2}} + \dots + (r-1)!! \int_M^\infty e^{-\frac{x^2}{2}} dx.$$

Applying Claim 2 and $(r-1)!! \leq (O(\sqrt{r}))^r$, the conclusion follows. \square

Lemma 32. For $M \geq 1$,

$$0 \leq \mathbb{E}[Z^r] - \mathbb{E}[Z^r | |Z| \leq M] \leq r(O(\sqrt{r}))^r \left(M^{r-1} e^{-\frac{M^2}{2}} \right).$$

Proof. For r odd, we have $\mathbb{E}[Z^r] - \mathbb{E}[Z^r | |Z| \leq M] = 0$. For r even, the left inequality is immediate since $x \mapsto x^r$ is increasing. For the right inequality,

$$\mathbb{E}[Z^r] - \mathbb{E}[Z^r | |Z| \leq M] = \mathbb{E}[Z^r] - \frac{\mathbb{E}[Z^r \mathbf{1}_{\{|Z| \leq M\}}]}{\mathbb{P}[|Z| \leq M]} \leq \frac{\mathbb{E}[Z^r] - \mathbb{E}[Z^r \mathbf{1}_{\{|Z| \leq M\}}]}{\mathbb{P}[|Z| \leq M]},$$

and the conclusion follows from Claim 3. \square

Lemma 33 (Distribution of random projection). Let X be uniformly distributed over the unit sphere S^{d-1} . For any $a \in S^{d-1}$ and $r > 0$,

$$\mathbb{P}[|\langle a, X \rangle| < r] < r\sqrt{d}.$$

Proof. Denote the surface area of the d -dimensional unit sphere by $S_{d-1} = \frac{2\pi^{d/2}}{\Gamma(d/2)}$. By symmetry,

$$\mathbb{P}[|\langle a, X \rangle| < r] = \mathbb{P}[|X_1| < r] = \frac{\int_{-r}^r (\sqrt{1-x^2})^{d-2} S_{d-2} \sqrt{1-x^2} dx}{S_{d-1}} = \frac{2S_{d-2}}{S_{d-1}} \int_0^r (1-x^2)^{\frac{d-3}{2}} dx < r\sqrt{d},$$

where X_1 is the first coordinate of X . \square

Lemma 34 (Accuracy of the spectral method). *Let $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} \frac{1}{2}N(-\theta, I_d) + \frac{1}{2}N(\theta, I_d)$, where $\theta \in \mathbb{R}^d$. Let λ_S be the largest eigenvalue of $S - I_d$, where $S = \frac{1}{n} \sum_i X_i X_i^\top$ denotes the sample covariance matrix, and \hat{v} the corresponding normalized eigenvector, where we decree that $\theta^\top \hat{v} \geq 0$. Let $\hat{s} = \sqrt{(\lambda_S)_+}$ and $\hat{\theta} = \hat{s}\hat{v}$. If $n > d$, then, with high probability,*

$$\|\theta - \hat{\theta}\|_2 \leq O(d/n)^{1/4}.$$

Proof. The samples can be represented in a matrix form $X = \theta\varepsilon^\top + Z \in \mathbb{R}^{d \times n}$, where $\varepsilon \in \mathbb{R}^n$ is a vector of independent Rademacher random variables, and Z has independent standard normal entries. Using $\varepsilon^\top \varepsilon = n$, we have

$$S - I_d = \theta\theta^\top + B + C,$$

where $B = \frac{1}{n}ZZ^\top - I_d$ and $C = \frac{1}{n}(\theta\varepsilon^\top Z^\top + Z\varepsilon\theta^\top)$ are both symmetric. With high probability, the largest eigenvalue of B is at most $d/n + 2\sqrt{d/n}$ (see [DS01, Theorem II.13]), which is $O(\sqrt{d/n})$ when $n > d$, and the spectral norm of C is also $O(\sqrt{d/n})$. Then, $|\lambda_S - \|\theta\|_2^2| \leq O(\sqrt{d/n})$ by Weyl's inequality, and thus $|\hat{s} - \|\theta\|_2| \leq O(d/n)^{1/4}$. Since \hat{v} maximizes $\|u^\top(S - I_d)u\|$ among all unit vectors $u \in \mathbb{R}^d$, including the direction of θ , then we obtain that $(\theta^\top \hat{v})^2 \geq \|\theta\|_2^2 - O(\sqrt{d/n})$, and consequently,

$$\|\theta - \|\theta\|_2 \hat{v}\|_2^2 \leq O(\sqrt{d/n}).$$

The conclusion follows from the triangle inequality. \square

Lemma 35. *The boundary of the space of the first $2k - 1$ moments of all distributions on \mathbb{R} corresponds to distributions with fewer than k atoms, while the interior corresponds to exactly k atoms.*

Proof. Given $m = (m_1, \dots, m_{2k-1})$ that corresponds to a distribution of exactly k atoms, by [Lin89, Theorem 2A], the moment matrix \mathbf{M}_{k-1} is positive definite. For any vector m' in a sufficiently small ball around m , the corresponding moment matrix \mathbf{M}'_{k-1} is still positive definite. Consequently, the matrix \mathbf{M}'_{k-1} is of full rank, and thus m' is a legitimate moment vector by [Las09, Theorem 3.4] (or [CF91, Theorem 3.1]). If m corresponds to a distribution with exactly $r < k$ atoms, by [Lin89, Theorem 2A], \mathbf{M}_{r-1} is positive definite while \mathbf{M}_r is rank deficient. Then, m is no longer in the moment space if m_{2r} is decreased. \square

Acknowledgment

We are grateful to Philippe Rigollet for bringing [HK15] to our attention and Harry Zhou for pointing out [Che95]. We also thank Sivaraman Balakrishnan for helpful discussions on [BWY17, Ngu13].

References

- [ADV13] M Andersen, Joachim Dahl, and Lieven Vandenbergh. CVXOPT: A Python package for convex optimization. 2013. abel.ee.ucla.edu/cvxopt.
- [AGH⁺14] Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15:2773–2832, 2014.

- [AK01] Sanjeev Arora and Ravi Kannan. Learning mixtures of arbitrary Gaussians. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 247–257. ACM, 2001.
- [Akh65] Naum Ilich Akhiezer. *The classical moment problem: and some related questions in analysis*, volume 5. Oliver & Boyd, 1965.
- [AM74] David F Andrews and Colin L Mallows. Scale mixtures of normal distributions. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 99–102, 1974.
- [AM05] Dimitris Achlioptas and Frank McSherry. On spectral learning of mixtures of distributions. In *International Conference on Computational Learning Theory*, pages 458–469. Springer, 2005.
- [ARS16] Carlos Améndola, Kristian Ranestad, and Bernd Sturmfels. Algebraic identifiability of Gaussian mixtures. *International Mathematics Research Notices*, 2016.
- [AS64] Milton Abramowitz and Irene A Stegun. *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*. Courier Corporation, 1964.
- [Ash65] Robert B. Ash. *Information Theory*. Dover Publications Inc., New York, NY, 1965.
- [Atk08] Kendall E Atkinson. *An introduction to numerical analysis*. John Wiley & Sons, 2008.
- [Ber80] Joseph Berkson. Minimum chi-square, not maximum likelihood! (with discussion). *The Annals of Statistics*, pages 457–487, 1980.
- [BK80] V. V Buldygin and Y. V. Kozachenko. Sub-Gaussian random variables. *Ukrainian Mathematical Journal*, 32(6):483–489, 1980.
- [BS10] Mikhail Belkin and Kaushik Sinha. Polynomial learning of distribution families. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 103–112. IEEE, 2010.
- [Bus11] Jorge Bustamante. *Algebraic Approximation: A Guide to Past and Current Solutions*. Springer Science & Business Media, 2011.
- [BV08] Spencer Charles Brubaker and Santosh Vempala. Isotropic PCA and affine-invariant clustering. In *IEEE 49th Annual IEEE Symposium on Foundations of Computer Science, 2008.*, pages 551–560. IEEE, 2008.
- [BWY17] Sivaraman Balakrishnan, Martin J Wainwright, and Bin Yu. Statistical guarantees for the em algorithm: From population to sample-based analysis. *The Annals of Statistics*, 45(1):77–120, 2017.
- [CF91] Raúl E Curto and Lawrence A Fialkow. Recursiveness, positivity, and truncated moment problems. *Houston Journal of Mathematics*, 17(4):603–635, 1991.
- [Cha10] Pierre Chaussé. Computing generalized method of moments and generalized empirical likelihood with R. *Journal of Statistical Software*, 34(11):1–35, 2010.
- [Che95] Jiahua Chen. Optimal rate of convergence for finite mixture models. *The Annals of Statistics*, pages 221–233, 1995.

- [Chr77] Elwin Bruno Christoffel. Sur une classe particulière de fonctions entières et de fractions continues. *Annali di Matematica Pura ed Applicata (1867-1897)*, 8(1):1–10, 1877.
- [CL11] T.T. Cai and M. G. Low. Testing composite hypotheses, Hermite polynomials and optimal estimation of a nonsmooth functional. *The Annals of Statistics*, 39(2):1012–1041, 2011.
- [Das99] Sanjoy Dasgupta. Learning mixtures of Gaussians. In *Foundations of computer science, 1999. 40th annual symposium on*, pages 634–644. IEEE, 1999.
- [DB16] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- [Dia87] Persi Diaconis. Application of the method of moments in probability and statistics. In *Moments in mathematics*, volume 37, pages 125–139. Amer. Math. Soc.: Providence, RI, 1987.
- [DLR77] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the royal statistical society. Series B (methodological)*, pages 1–38, 1977.
- [DS01] Kenneth R Davidson and Stanislaw J Szarek. Local operator theory, random matrices and Banach spaces. *Handbook of the geometry of Banach spaces*, 1(317-366):131, 2001.
- [DS16] Harold G Diamond and Armin Straub. Bounds for the logarithm of the Euler gamma function and its derivatives. *Journal of Mathematical Analysis and Applications*, 433(2):1072–1083, 2016.
- [FS06] Sylvia Frühwirth-Schnatter. *Finite mixture and Markov switching models*. Springer Science & Business Media, 2006.
- [Gau04] Walter Gautschi. *Orthogonal polynomials: computation and approximation*. Oxford University Press on Demand, 2004.
- [GR07] I. S. Gradshteyn and I. M. Ryzhik. *Table of Integrals Series and Products*. Academic, New York, NY, seventh edition, 2007.
- [GVDV01] Subhashis Ghosal and Aad W Van Der Vaart. Entropies and rates of convergence for maximum likelihood and Bayes estimation for mixtures of normal densities. *Annals of Statistics*, pages 1233–1263, 2001.
- [GW69] Gene H Golub and John H Welsch. Calculation of Gauss quadrature rules. *Mathematics of computation*, 23(106):221–230, 1969.
- [GW00] C. R. Genovese and L. Wasserman. Rates of convergence for the Gaussian mixture sieve. *Annals of Statistics*, 28(4):1105–1127, 2000.
- [Hal05] Alastair R Hall. *Generalized method of moments*. Oxford University Press, 2005.
- [Han82] Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, pages 1029–1054, 1982.
- [HJ12] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 2 edition, 2012.

- [HK13] 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.
- [HK15] Philippe Heinrich and Jonas Kahn. Optimal rates for finite mixture estimation. *arXiv:1507.04313*, 2015.
- [HP15] Moritz Hardt and Eric Price. Tight bounds for learning a mixture of two gaussians. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 753–760. ACM, 2015.
- [Jew82] Nicholas P Jewell. Mixtures of exponential distributions. *The annals of statistics*, pages 479–484, 1982.
- [KMV10] 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.
- [Kos07] Michael R Kosorok. *Introduction to empirical processes and semiparametric inference*. Springer Science & Business Media, 2007.
- [Kra32] Michel Krawtchouk. Sur le problème de moments. In *ICM Proceedings*, pages 127–128, 1932. Available at <https://www.mathunion.org/fileadmin/ICM/Proceedings/ICM1932.2/ICM1932.2.ocr.pdf>.
- [KS53] Samuel Karlin and Lloyd S Shapley. *Geometry of moment spaces*. Number 12. American Mathematical Soc., 1953.
- [KSV05] 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.
- [KW56] Jack Kiefer and Jacob Wolfowitz. Consistency of the maximum likelihood estimator in the presence of infinitely many incidental parameters. *The Annals of Mathematical Statistics*, pages 887–906, 1956.
- [KX03] Dimitris Karlis and Evdokia Xekalaki. Choosing initial values for the EM algorithm for finite mixtures. *Computational Statistics & Data Analysis*, 41(3):577–590, 2003.
- [KX05] Dimitris Karlis and Evdokia Xekalaki. Mixed Poisson distributions. *International Statistical Review*, 73(1):35–58, 2005.
- [Las09] Jean Bernard Lasserre. *Moments, positive polynomials and their applications*, volume 1. World Scientific, 2009.
- [LC86] L. Le Cam. *Asymptotic methods in statistical decision theory*. Springer-Verlag, New York, NY, 1986.
- [Lin89] Bruce G Lindsay. Moment matrices: applications in mixtures. *The Annals of Statistics*, pages 722–740, 1989.
- [Lin95] Bruce G Lindsay. Mixture models: theory, geometry and applications. In *NSF-CBMS regional conference series in probability and statistics*, pages i–163. JSTOR, 1995.

- [Llo82] Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982.
- [Mor82] Carl N Morris. Natural exponential families with quadratic variance functions. *The Annals of Statistics*, pages 65–80, 1982.
- [MV10] Ankur Moitra and Gregory Valiant. Settling the polynomial learnability of mixtures of Gaussians. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 93–102. IEEE, 2010.
- [Ngu13] XuanLong Nguyen. Convergence of latent mixing measures in finite and infinite mixture models. *The Annals of Statistics*, 41(1):370–400, 2013.
- [Pea94] Karl Pearson. Contributions to the mathematical theory of evolution. *Philosophical Transactions of the Royal Society of London. A*, 185:71–110, 1894.
- [PL01] Ramani S Pilla and Bruce G Lindsay. Alternative EM methods for nonparametric finite mixture models. *Biometrika*, 88(2):535–550, 2001.
- [PSWS03] Javier Portilla, Vasily Strela, Martin J Wainwright, and Eero P Simoncelli. Image denoising using scale mixtures of Gaussians in the wavelet domain. *IEEE Transactions on Image processing*, 12(11):1338–1351, 2003.
- [RW84] Richard A Redner and Homer F Walker. Mixture densities, maximum likelihood and the EM algorithm. *SIAM review*, 26(2):195–239, 1984.
- [SB02] J. Stoer and R. Bulirsch. *Introduction to Numerical Analysis*. Springer-Verlag, New York, NY, 3rd edition, 2002.
- [Sch17] Konrad Schmüdgen. *The moment problem*. Springer, 2017.
- [SMA00] Wilfried Seidel, Karl Mosler, and Manfred Alker. A cautionary note on likelihood ratio tests in mixture models. *Annals of the Institute of Statistical Mathematics*, 52(3):481–487, 2000.
- [ST43] James Alexander Shohat and Jacob David Tamarkin. *The problem of moments*. Number 1. American Mathematical Soc., 1943.
- [Sze75] G. Szegő. *Orthogonal polynomials*. American Mathematical Society, Providence, RI, 4th edition, 1975.
- [Tim63] Aleksandr Filippovich Timan. *Theory of approximation of functions of a real variable*. Pergamon Press, 1963.
- [Usp37] James Victor Uspensky. *Introduction to mathematical probability*. McGraw-Hill, 1937.
- [VB96] Lieven Vandenberghe and Stephen Boyd. Semidefinite programming. *SIAM review*, 38(1):49–95, 1996.
- [VdV00] Aad W. Van der Vaart. *Asymptotic statistics*. Cambridge university press, Cambridge, United Kingdom, 2000.
- [Vil03] C. Villani. *Topics in optimal transportation*. American Mathematical Society, Providence, RI, 2003.

- [Vil08] C. Villani. *Optimal Transport: Old and New*. Springer Verlag, Berlin, 2008.
- [VW04] Santosh Vempala and Grant Wang. A spectral algorithm for learning mixture models. *Journal of Computer and System Sciences*, 68(4):841–860, 2004.
- [WS00] Martin J Wainwright and Eero P Simoncelli. Scale mixtures of Gaussians and the statistics of natural images. In *Advances in neural information processing systems*, pages 855–861, 2000.
- [WSV12] Henry Wolkowicz, Romesh Saigal, and Lieven Vandenbergh. *Handbook of semidefinite programming: theory, algorithms, and applications*, volume 27. Springer Science & Business Media, 2012.
- [WV10] Yihong Wu and Sergio Verdú. The impact of constellation cardinality on Gaussian channel capacity. In *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, pages 620–628. IEEE, 2010.
- [WY15] Yihong Wu and Pengkun Yang. Chebyshev polynomials, moment matching, and optimal estimation of the unseen. *arXiv:1504.01227*, 2015.
- [WY16] Yihong Wu and Pengkun Yang. Minimax rates of entropy estimation on large alphabets via best polynomial approximation. *IEEE Transactions on Information Theory*, 62(6):3702–3720, 2016.
- [XJ96] Lei Xu and Michael I Jordan. On convergence properties of the EM algorithm for Gaussian mixtures. *Neural computation*, 8(1):129–151, 1996.