

Efficient Algorithms for Sparse Moment Problems without Separation

Zhiyuan Fan

Jian Li

Tsinghua University

FAN-ZY19@MAILS.TSINGHUA.EDU.CN

LIJIAN83@MAIL.TSINGHUA.EDU.CN

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

We consider the sparse moment problem of learning a k -spike mixture in high-dimensional space from its noisy moment information in any dimension. We measure the accuracy of the learned mixtures using transportation distance. Previous algorithms either assume certain separation assumptions, use more recovery moments, or run in (super) exponential time. Our algorithm for the one-dimensional problem (also called the sparse Hausdorff moment problem) is a robust version of the classic Prony’s method, and our contribution mainly lies in the analysis. We adopt a global and much tighter analysis than previous work (which analyzes the perturbation of the intermediate results of Prony’s method). A useful technical ingredient is a connection between the linear system defined by the Vandermonde matrix and the Schur polynomial, which allows us to provide tight perturbation bound independent of the separation and may be useful in other contexts. To tackle the high-dimensional problem, we first solve the two-dimensional problem by extending the one-dimensional algorithm and analysis to complex numbers. Our algorithm for the high-dimensional case determines the coordinates of each spike by aligning a 1d projection of the mixture to a random vector and a set of 2d projections of the mixture. Our results have applications to learning topic models and Gaussian mixtures, implying improved sample complexity results or running time over prior work.

Keywords: Moment Problems, Topic Modeling, Gaussian Mixture Models

1. Background

We study the moment problem in which we are given the (noisy) information of the moments of a measure μ , and our goal is to recover μ up to certain accuracy. The moment problem is a classical problem studied for over a century in mathematics, statistics, computer science, physics, control theory, medical image, etc. Various versions of the problem have found numerous applications in different domains (e.g., [Schmüdgen \(2017\)](#); [Lasserre \(2009\)](#); [Ang et al. \(2002\)](#); [Pintarelli and Vericat \(2016\)](#); [Bessis et al. \(1976\)](#); [Natterer \(2001\)](#)). We study the sparse version of the problem, in which the underlying measure is a mixture of k discrete distributions over a common discrete domain $[d] = \{1, 2, \dots, d\}$. In particular, we consider mixture ϑ which is a k -spike distribution supported on the $d - 1$ dimensional simplex $\Delta_{d-1} = \{\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d \mid \sum_{i=1}^d x_i = 1, x_i \geq 0 \forall i \in [d]\}$. Each point in Δ_{d-1} represents a discrete distribution over $[d]$. We first introduce the sparse moment problem for $d = 2$ (also known as the sparse *Hausdorff moment* problem or the k -coin model ([Schmüdgen, 2017](#); [Li et al., 2015](#); [Gordon et al., 2020](#))).

The sparse Hausdorff moment problem: For $d = 2$, the mixture is a k -spike distribution supported on $\Delta_1 \cong [0, 1]$. We call the model *the k -coin model* (i.e., a mixture of k Bernoulli distributions). Denote the underlying mixture as $\vartheta = (\boldsymbol{\alpha}, \mathbf{w})$. Here, $\boldsymbol{\alpha} = \{\alpha_1, \alpha_2, \dots, \alpha_k\}$ and $\mathbf{w} = \{w_1, w_2, \dots, w_k\} \in \Delta_{k-1}$, where $\alpha_i \in [0, 1]$ specifies the i th Bernoulli distribution and

$w_i \in [0, 1]$ is the probability of mixture component α_i . For mixture $\vartheta = (\alpha, w)$, the t th moment is defined as $M_t(\vartheta) = \int_{[0,1]} \alpha^t \vartheta(d\alpha) = \sum_{i=1}^k w_i \alpha_i^t$. Our goal is to recover the unknown parameters (α, w) of the mixture ϑ given the first K noisy moment values M'_t with $|M'_t - M_t(\vartheta)|_\infty \leq \xi$ for $1 \leq t \leq K$.

We call K the *moment number* and ξ the moment precision/accuracy, both of which are essential parameters of the problem. We measure the quality of our estimation $\tilde{\vartheta}$ in terms of the transportation distance between probability distributions, i.e., $\text{Tran}(\tilde{\vartheta}, \vartheta) \leq O(\epsilon)$. Using transportation distance as the metric is advantageous for several reasons: (1) if the desired accuracy ϵ is much smaller than the minimum separation $\zeta = \min \|\alpha_i - \alpha_j\|$, the estimation $\tilde{\vartheta}$ must be a per-spike recovery since it must contain a spike that is sufficiently close to an original spike in ϑ (given the weight of the spike is lower bounded); (2) if the desired accuracy ϵ is larger than ζ or the minimum weight w_{\min} , we are allowed to be confused by two very close spikes or miss a spike with a very small weight, thus potentially avoiding the inverse dependency on ζ and w_{\min} , which are otherwise unavoidable if we must recover every spike. Generally, for this problem, we are interested in understanding the relations and trade-offs among the moment number K , the moment accuracy/precision ξ , and the accuracy ϵ of our estimation $\tilde{\vartheta}$ (i.e., transportation distance at most ϵ), and designing efficient algorithms for recovering ϑ to the desired accuracy.

The sparse moment problem in higher dimensions: For the higher dimensional case, suppose the underlying mixture is $\vartheta = (\alpha, w)$ where $\alpha = \{\alpha_1, \dots, \alpha_k\}$ ($\alpha_i \in \Delta_{d-1}$ which are d -dimensional points in a bounded domain) and $w = \{w_1, \dots, w_k\} \in \Delta_{k-1}$. For vector $\alpha_i = (\alpha_{i,1}, \dots, \alpha_{i,d})$ and multi-index $\mathbf{t} = (t_1, \dots, t_d) \in \mathbb{Z}_+^d$, we denote monomial $\alpha_i^{\mathbf{t}} = \alpha_{i,1}^{t_1} \alpha_{i,2}^{t_2} \dots \alpha_{i,d}^{t_d}$, and the \mathbf{t} -moment of ϑ as $M_{\mathbf{t}}(\vartheta) = \int_{\Delta_{d-1}} \alpha^{\mathbf{t}} \vartheta(d\alpha) = \sum_{i=1}^k w_i \alpha_i^{\mathbf{t}}$. In the sparse moment problem, we are given noisy access of \mathbf{t} -moment $M_{\mathbf{t}}(\vartheta)$ for $\|\mathbf{t}\|_1 \leq K$. Here we also call K the *moment number*. Since there are $K^{O(d)}$ different moments, when d is large, we also consider the case where we have noisy access to the moments of the projections of ϑ onto lower-dimensional subspaces (typically 1d or 2d projections). An affine transformation of the samples in concrete applications such as topic models can often obtain such noisy moments.

1.1. Applications and connections to prior work

The Hausdorff moment problem and its higher dimensional version have many applications in different areas, such as mathematics, statistics, computer science, physics, etc. In recent years, the noisy sparse moment problem has found applications to a variety of unsupervised learning scenarios, including learning topic models (Hofmann, 1999; Blei et al., 2003; Papadimitriou et al., 2000), learning Gaussian mixtures (Wu and Yang, 2020), collaborative filtering (Hofmann and Puzicha, 1999; Kleinberg and Sandler, 2008), learning a mixture of product distributions (Feldman et al., 2008; Gordon et al., 2021a) and causal inference (Gordon et al., 2021b). Now, we first discuss the connection between the sparse moment problem and learning mixture models. Then we list some applications in other areas.

Applications to learning mixture models.: Indeed, several prior results for the sparse moment problem we study in this paper were presented explicitly or implicitly in the context of learning mixture models (e.g., Rabani et al. (2014); Li et al. (2015); Kim et al. (2019); Gordon et al. (2020); Wu and Yang (2020); Doss et al. (2020)). Here we use the problem of learning topic model as an example: We are given a corpus of documents. We adopt the popular ‘‘bag of words’’ model and

take each document as an unordered multiset of words. The assumption is that a small number of k “pure” topics are distributed over the underlying vocabulary of d words. A K -word document (i.e., a string in $[d]^K$) is generated by first selecting a topic $p \in \Delta_{d-1}$ from the mixture ϑ , and then sampling K words from this topic. K is called *the snapshot number* and we call such a sample a K -*snapshot* of p . For all $\|t\|_1 \leq K$, we can obtain noisy estimates of $M_t(\vartheta)$ or $M_t(\Pi_R(\vartheta))$ (where $\Pi_R(\vartheta)$ is the projection of ϑ onto a lower dimensional subspace R) using the K -snapshots (i.e., the documents in the corpus). We aim to recover the mixture ϑ using as few samples as possible. Once we have computed the noisy moments using samples, the problem reduces to the sparse moment problem. We will discuss the applications to learning topic models in Section 6.1 and to learning Gaussian mixtures in Section 6.2.

Other Applications.: We mention a few other problems where our results are directly or potentially useful:

1. The moment problem can be used for recovering the initial condition for a class of heat diffusion equations (Ang et al., 2002): Suppose the initial condition $f(x, t)$ at $t = 0$ is a k -spike distribution (e.g., k heat sources), and we can measure $f(x, t)$ at $t = 1$ (with measurement noise), and the task is to recover the initial k -spike distribution (Ang et al., 2002). See Remark 54 for more details. In some other PDE problems, the input moment information is computed using a particular numerical algorithm (now the noise comes from the error/precision of the numerical algorithms), and the moment problem is used as a sub-procedure for solving the PDE. For example, the moment problem can be used for approximately solving a class of Poisson equations (Pintarelli and Vericat, 2016), in which the moment information can be computed using numerical integration.
2. The Hausdorff moment problem is used as a sub-procedure for bounding/locating the zeros of the partition function of the lattice Ising model (Lee-Yang theory) (Bessis et al., 1976). Here the moments can be computed exactly for some special lattices and approximated for more general lattices.
3. Our algorithm for higher dimensions is also closely related to computerized tomography, in which we would like to recover the 3d (or higher dimensional) object from its 1d or 2d projections (Natterer, 2001). Existing techniques used the information from a continuous set (or a net) of 1d or 2d projections for recovery (e.g., the classical Radon transform). Our technique (see Theorem 5 and Algorithm 3) only use only $O(d)$ 1d and 2d projections.

2. Our Contributions

The k -Coin model. We first discuss prior work on the k -coin model (ϑ is supported on $[0, 1]$), and our results. Rabani et al. (2014) showed that recovering a k -spike mixture requires at least the first $K = 2k - 1$ moments in the worst case (even without any noise). Moreover, they showed that for the topic learning problem, if one use $K = c(2k - 1)$ snapshot samples (for any constant $c \geq 1$) and wish to achieve an accuracy $O(1/k)$ in terms of transportation distance, $e^{\Omega(K)}$ samples are required (or equivalently, the moment accuracy should be at most $e^{-\Omega(K)}$). On the positive side, they solved the problem using sample complexity $\max\{(1/\zeta)^{O(k)}, (k/\epsilon)^{O(k^2)}\}$ (or moment accuracy $\min\{\zeta^{O(k)}, (\epsilon/k)^{O(k^2)}\}$), where ζ is a lower bound of the minimum separation). In fact, their algorithm is a variant of the classic Prony’s method and the (post-sampling) running time is

Reference	K	Moment Accuracy (ξ)	Running Time	Separation
Rabani et al. (2014)	$2k - 1$	$\min\{\zeta^{O(k)}, (\epsilon/k)^{O(k^2)}\}$	$\text{poly}(k)$	Required
Kim et al. (2019)	$2k - 1$	$\zeta^{O(k)} \cdot w_{\min}^2 \cdot \epsilon$	$O(k^3)$	Required
Gordon et al. (2020)	$2k$	$\zeta^{O(k)} \cdot w_{\min} \cdot \epsilon$	$O(k^{2+o(1)})$	Required
Li et al. (2015)	$2k - 1$	$(\epsilon/k)^{O(k)}$	$(k/\epsilon)^{O(k^2)}$	No Need
Wu and Yang (2020)	$2k - 1$	$(\epsilon/k)^{O(k)}$	$O(k^{2\omega})$	No Need
Theorem 3	$2k - 1$	$(\epsilon/k)^{O(k)}$	$O(k^2)$	No Need

Table 1: Algorithms for the k -coin problem where K is the moment number. The last column indicates whether the algorithm needs the separation assumption. In particular, the algorithms that require the separation assumption need to know ζ and w_{\min} where ζ is the minimum separation, i.e., $\zeta \leq \min_{i \neq j} |\alpha_i - \alpha_j| \leq 1/k$, and $w_{\min} \leq \min_i w_i$. We measure the running time in terms of the arithmetic operations. In addition, $\text{poly}(k)$ refers to $O(k^c)$ for a relatively large constant c and ω refers to the smallest real number such that two $k \times k$ matrices can be multiplied using $O(k^{\omega+o(1)})$ operations. We note that there is a trivial lower bound $\omega \geq 2$ while the best-known upper bound is $\omega \leq 2.373$ ([Alman and Williams, 2021](#)).

$\text{poly}(k)$ where their bottleneck is solving a special convex quadratic programming instance. [Li et al. \(2015\)](#) provided a linear programming-based algorithm that requires the moment accuracy $(\epsilon/k)^{O(k)}$ (or sample complexity $(k/\epsilon)^{O(k)}$), matching the lower bound in [Rabani et al. \(2014\)](#). However, after computing the noisy moments, the running time of their algorithm is super-exponential $k^{O(k^2)}$. Motivated by a problem in population genetics that could be reduced to the k -coin mixture problem, [Kim et al. \(2019\)](#) analyzed the Matrix Pencil Method, which requires moment accuracy $\zeta^{O(k)} \cdot w_{\min}^2 \cdot \epsilon$. The matrix pencil method requires solving a generalized eigenvalue problem, which needs $O(k^3)$ time. [Wu and Yang \(2020\)](#) studied the same problem in the context of learning Gaussian mixtures. In fact, they showed that learning 1d Gaussian mixture models with the same variance can be reduced to learning the k -coin model. Their algorithm achieves the optimal dependency on the moment accuracy without separation assumption. Their recovery algorithm is based on SDP and runs in time $O(k^{2\omega})$ for some $\omega < 2.373$ using the state-of-the-art SDP solver ([Huang et al., 2022](#)). Recently, [Gordon et al. \(2020\)](#) showed that under the separation assumption, it is possible to recover all k spikes up to accuracy ϵ , using the first $2k$ -moments with accuracy $\zeta^{O(k)} \cdot w_{\min} \cdot \epsilon$. Their algorithm is also a variant of Prony’s method and it runs in time $O(k^{2+o(1)})$. We summarize the prior and our results in Table 1.

Theorem 1 (*The k -Coin model, informal*) Let ϑ be an arbitrary k -spike distribution over $[0, 1]$. Suppose we have noisy moments M'_t such that $|M_t(\vartheta) - M'_t| \leq (\epsilon/k)^{O(k)}$ for $0 \leq t \leq 2k - 1$. We can obtain a mixture $\tilde{\vartheta}$ such that $\text{Tran}(\tilde{\vartheta}, \vartheta) \leq O(\epsilon)$ in $O(k^2)$ arithmetic operations.

Our techniques: Our algorithm for the k -coin model is also based on the classic Prony’s method ([de Prony, 1795](#)). Suppose the true mixture is $\vartheta = (\alpha, w)$ where $\alpha = [\alpha_1, \dots, \alpha_k]^\top$ and $w = [w_1, \dots, w_k]^\top$. In Prony’s method, if we know the moment vector M exactly, every α_i can be recovered from the roots of a polynomial whose coefficients are the entries of the eigenvector c (corresponding to eigenvalue 0) of a Hankel matrix (see e.g., the matrix \mathcal{H}_{k+1} in [Gordon et al.](#)

(2020)). However, if we only know the noisy moment vector M' , we only get a perturbation of the original Hankel matrix. Recent analyses of Prony’s method (Rabani et al., 2014; Gordon et al., 2020) aim to upper bound the error of recovery from the perturbed Hankel matrix. If α_i and α_j ($i \neq j$) are very close, or any w_i is very small, the second smallest eigenvalue of the Hankel matrix is also very close to 0. From matrix perturbation theory, we know that the eigenvector c corresponding to eigenvalue 0 is extremely sensitive to small perturbation (see e.g., Stewart (1990)). Hence, the recent analyses of Prony’s method (Rabani et al., 2014; Gordon et al., 2020) requires that α_i and α_j are separated by at least ζ , and the minimum w_i is lower bounded by w_{\min} to ensure certain stability of the eigenvector (hence the coefficients of the polynomial), and the corresponding sample complexity becomes unbounded when ζ or w_{\min} approaches to 0. Our algorithm (Algorithm 1) can be seen as a robust version of Prony’s method. Instead of computing the smallest eigenvector of the perturbed Hankel matrix, we solve a ridge regression to obtain a vector \hat{c} , which plays a similar role as c (Line 1 in Algorithm 1). However, instead of showing \hat{c} is close to c (in fact, they can be very different), we adopt a more global analysis to show that the moment vector of the estimated mixture is close to the true moment vector M , and by a moment-transportation inequality (see Section B), we can guarantee the quantity of our solution in terms of transportation distance.

Another technical challenge lies in bounding the error in recovering the weight vector w . In the noiseless case, the weight is simply the solution of $V_\alpha w = M$ where V_α is a Vandermonde matrix (see Equation (2)). It is known that Vandermonde matrices tend to be badly ill-conditioned (Gautschi and Inglese, 1987; Pan, 2016; Moitra, 2015) (with a large dependency on the inverse of the minimal separation). Hence, using standard condition number-based analysis, slight perturbations of V_α and M may result in an unbounded error.¹ However, we show interestingly, in our case, that we can bound the error independent of the separation (Lemma 32) via the connection between Vandermonde linear systems and Schur polynomial (Lemma 31). Our technology may be useful in analyzing the perturbation of Vandermonde linear systems in other contexts as well.

Higher Dimensions. For the high-dimensional case, especially when d is large, it is costly to obtain all moments $M_{\mathbf{t}}(\vartheta)$ for all \mathbf{t} with $\|\mathbf{t}\|_1 \leq K$. Indeed, most previous work (Rabani et al., 2014; Li et al., 2015; Wu and Yang, 2020; Gordon et al., 2020; Doss et al., 2020) on learning mixtures solved the problem using noisy moment information of the projections of ϑ onto some lower-dimensional subspaces (in fact, lines). In particular, if one projects ϑ to $O(k)$ directions such as in Rabani et al. (2014); Wu and Yang (2020); Gordon et al. (2020), existing techniques require the separation assumption for the recovery.² Other works without separation condition such as Li et al. (2015); Doss et al. (2020) requires to project ϑ to a net of exponentially many directions.

In our work, we assume we have noisy access to the moments of the linear projections $\Pi_R(\vartheta)$ of ϑ onto lines or 2d planes (one may refer to the formal definition of $\Pi_R(\vartheta)$ in Section 5.2). See Theorem 2 for an informal statement of our result. A comparison between our result and prior work is summarized in Table 2. One can see that we either use much fewer projections (hence better running time) or do not need the separation assumption. We note that using 2d projections is crucial

1. Note that for Vandermonde matrix $V = (a_i^j)_{0 \leq i < k, 1 \leq j \leq k}$, its determinant is $\det(V) = (-1)^{k(k-1)/2} \prod_{p < q} (a_p - a_q)$. Hence, $\det(V)^{-1}$ depends inversely on the separation.
2. It is possible to convert an algorithm with a separation assumption to an algorithm without a separation assumption by merging close-by spikes and removing spikes with small weights. However, the resulting algorithm is far from optimal. For example, sample complexity $(k/\epsilon)^{O(k)} \cdot w_{\min}^{-O(1)} \cdot \zeta^{-O(k)}$ in Gordon et al. (2020) can be converted to one without separation with sample complexity $(k/\epsilon)^{O(k^2)}$. We omit the details.

Reference	#Projections	Moment Accuracy (ξ)	Running Time	Separation
Rabani et al. (2014)	$O(k)$	$(\epsilon/k)^{O(k^2)} \cdot (w_{\min}\zeta)^{O(k^2)}$	$\text{poly}(k)$	Required
Gordon et al. (2020)	$O(k)$	$(\epsilon/k)^{O(k)} \cdot w_{\min}^{O(1)} \cdot \zeta^{O(k)}$	$\text{poly}(k)$	Required
Wu and Yang (2020)	$O(d)$	$(\epsilon/k)^{O(k)} \cdot w_{\min} \cdot \zeta$	$O(dk^{2\omega})$	Required
Li et al. (2015)	$(k/\epsilon)^{O(k)}$	$(\epsilon/k)^{O(k^2)}$	$(k/\epsilon)^{O(k^2)}$	No Need
Doss et al. (2020)	$(k/\epsilon)^{O(k)}$	$(\epsilon/k)^{O(k)}$	$(k/\epsilon)^{O(k)}$	No Need
Theorem 5	$O(d)^*$	$(\epsilon/k)^{O(k)}$	$O(dk^3)$	No Need

Table 2: Algorithms for recovering the mixture ϑ in higher dimensional space. We assume that dimension $d < k$ and we have oracle access to noisy moments of any 1d or 2d projections. The last column indicates whether the algorithm needs the separation assumption. *: The second column indicates the number of lower dimensional projections that the algorithm requires. Specifically, our algorithm uses one 1d projection and $O(d)$ 2d projections while all previous algorithms use a certain number of 1d projections.

for the improvement beyond prior work. We hope this idea of using 2d projections is helpful in solving other high-dimensional recovery problems.

Theorem 2 (*Higher dimension, informal*) *Let ϑ be an arbitrary k -spike mixture supported in Δ_{d-1} . Suppose we can access any noisy t -moments of the 1d and 2d linear projection $\Pi_R(\vartheta)$ with precision $(\epsilon/(dk))^{O(k)}$, (i.e., onto lines and 2d planes). We can construct a k -spike mixture $\tilde{\vartheta}$ such that $\text{Tran}(\tilde{\vartheta}, \vartheta) \leq O(\epsilon)$ using only $O(dk^3)$ arithmetic operations with high probability.*

Our techniques: We first solve the two-dimensional problem. This is done by extending the previous one-dimensional algorithm and its analysis to complex numbers. The real and imaginary parts of a complex location can represent the two dimensions, respectively. While most ideas are similar to the one-dimensional case, the analysis requires extra care in various places. In particular, we need to extend the moment-transportation inequality to complex numbers (Appendix B).

Our algorithm for the high-dimensional case uses the algorithms for one-dimensional and two-dimensional cases as subroutines. We first pick a random vector \mathbf{r} and learn the projection $\tilde{\phi}$ of ϑ onto \mathbf{r} using the one-dimensional algorithm. It is not difficult to show that for two spikes that are far apart, with reasonable probability, their projections on \mathbf{r} are also far apart. Hence, the remaining task is to recover the coordinates (in \mathbb{R}^d) of each spike in $\tilde{\phi}$. So, we learn the 2d-linear map $\hat{\phi}_t$ of ϑ onto the two-dimensional subspace spanned by \mathbf{r} and the t th dimensional axis for each dimension $t \in [d]$. Combining the information from the 1d projected measure $\tilde{\vartheta}$ and the 2d projected measure $\hat{\vartheta}_t$, we show that we can extract the t th coordinates of all spikes.

Application to Learning Topic Models. Our result for the high-dimensional case can be easily translated into an algorithm for topic modeling. Previously, Rabani et al. (2014) showed that it is possible to produce an estimate $\tilde{\vartheta}$ of the original mixture ϑ such that $\text{Tran}(\vartheta, \tilde{\vartheta}) \leq \epsilon$, using $\text{poly}(n, k, 1/\epsilon)$ 1- and 2-snapshot samples, and $(k/\epsilon)^{O(k^2)} \cdot (w_{\min}\zeta)^{-O(k^2)}$ K -snapshot samples, under minimal separation assumptions. Li et al. (2015) provided an LP-based learning algorithm which uses almost the same number of samples (with a slightly worse polynomial for 2-snapshot samples), but without requiring any separation assumptions. The number of the $(2k - 1)$ -snapshot samples used by both Rabani et al. (2014); Li et al. (2015) are $(k/\epsilon)^{O(k^2)}$ while that of the lower

bound (for 1d) is only $e^{\Omega(k)}$. Recently, [Gordon et al. \(2020\)](#) showed that under the minimal separation assumption, the sample complexity could be reduced to $(k/\epsilon)^{O(k)} \cdot w_{\min}^{-O(1)} \cdot \zeta^{-O(k)}$. We obtain for the first time the worst case optimal sample complexity for the high-dimensional case, improving previous work ([Rabani et al., 2014](#); [Li et al., 2015](#); [Gordon et al., 2020](#)) and matching the lower bound even for the one-dimensional case ([Rabani et al., 2014](#)). See Section 6.1 for the details.

Application to Learning Gaussian Mixtures. We also study the problem of learning the parameters of Gaussian mixtures. We assume that all Gaussian components share a variance parameter, following the setting studied by [Wu and Yang \(2020\)](#). We can leverage our algorithm for the k -coin model for the one-dimensional setting. Our algorithm achieves the same sample complexity as in [Wu and Yang \(2020\)](#), but an $O(k^2)$ post-sampling running time, improving over the SDP-based $O(k^{2\omega})$ -time algorithm developed by [Wu and Yang \(2020\)](#). See Theorem 53 for more details.

For the high-dimensional setting, we can use the dimension reduction technique in [Li et al. \(2015\)](#) or [Doss et al. \(2020\)](#) to reduce the dimension d to $O(k)$. The dimension reduction part is not a bottleneck and we assume $d = O(k)$ for the the following discussion. We show in Section G.2 that we can utilize our algorithm for the high-dimensional sparse moment problem and obtain an algorithm without any separation assumption with sample complexity $(k/\epsilon)^{O(k)}$. Note that the algorithm in [Wu and Yang \(2020\)](#) requires a sample size of $(k/\epsilon)^{O(k)} \cdot w_{\min}^{-O(1)} \cdot \zeta^{-O(k)}$, which depends on the separation parameter ζ between Gaussian distributions. Recently, [Doss et al. \(2020\)](#) removed the separation assumption and achieved the optimal sample complexity $(k/\epsilon)^{O(k)}$. Compared with [Doss et al. \(2020\)](#), the sample complexity of our algorithm is the same, but our running time is substantially better: during the sampling phase, the algorithm in [Doss et al. \(2020\)](#) requires $O(n^{5/4} \text{poly}(k))$ time where $n = (k/\epsilon)^{O(k)}$ is the number of samples (for each sample, they need to update $O(n^{1/4} \text{poly}(k))$ numbers, since their algorithm requires $O(n^{1/4})$ 1d projections) while our algorithm only needs $O(n \text{poly}(k))$ time (we only need one 1d projection and $d = O(k)$ 2d projections). The post-sampling running time of the algorithm in [Doss et al. \(2020\)](#) is exponential $(k/\epsilon)^{O(k)}$ while our algorithms runs in polynomial time $\text{poly}(k)$. For more details, see Section 6.2.

Last but not least, we argue that improving the post-sampling running time is very important, despite the exponential sample complexity for both problems. During the sampling phase, we only need to keep track of the first few moments (e.g., using basic operations such as counting or adding numbers), hence the sampling phase can be easily distributed, streamed, or implemented in inexpensive computing devices. Our mixture recovery algorithm only requires the moment information (without storing the samples) and runs in time polynomial in k (not the sample size). Moreover, one may well have other means to measure the moments to achieve the desired moment accuracy (e.g., via longer documents, prior knowledge, existing samples, etc.) Exploiting other settings and applications is an interesting further direction.

3. Preliminaries

We are given a statistical mixture ϑ of k discrete distributions over $[d] = \{1, 2, \dots, d\}$. Each discrete distribution α_i can be regarded as a point in the $(d-1)$ -simplex $\Delta_{d-1} = \{\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d \mid \sum_{i=1}^d x_i = 1, x_i \geq 0 \forall i \in [d]\}$. We use $\vartheta = (\boldsymbol{\alpha}, \mathbf{w})$ to represent the mixture where $\boldsymbol{\alpha} = \{\alpha_1, \alpha_2, \dots, \alpha_k\} \subset \Delta_{d-1}$ are the locations of spikes and $\mathbf{w} = \{w_1, w_2, \dots, w_k\} \in \Delta_{k-1}$ in which w_i is the weight of α_i . Since the dimension of Δ_{d-1} is $d-1$, we say the *dimension* of the problem is $d-1$. In addition, we use $\text{Spike}(\Delta_{d-1}, \Delta_{k-1})$ to denote the set of all such mixtures

where Δ_{d-1} indicates the domain of α_i and Δ_{k-1} indicates the domain of w . Since our algorithm may produce negative or complex weights as intermediate results, we further denote $\Sigma_{d-1}^{\mathbb{R}} = \{x = (x_1, \dots, x_d) \in \mathbb{R}^d \mid \sum_{i=1}^d x_i = 1\}$ and $\Sigma_{d-1}^{\mathbb{C}} = \{x = (x_1, \dots, x_d) \in \mathbb{C}^d \mid \sum_{i=1}^d x_i = 1\}$.

For the one-dimensional case (which is called the k -coin problem or sparse Hausdorff moment problem), we have $\vartheta = (\alpha, w) \in \text{Spike}(\Delta_1, \Delta_{k-1})$ where $\alpha = \{\alpha_1, \alpha_2, \dots, \alpha_k\}$ is a set of k discrete points in $[0, 1]$. In this scenario, for each $t \in \mathbb{N}$, we denote the t th moment as

$$M_t(\vartheta) = \int_{[0,1]} \alpha^t \vartheta(d\alpha) = \sum_{i=1}^k w_i \alpha_i^t.$$

For the higher dimensional case, we use $\mathbf{t} = (t_1, t_2, \dots, t_d) \in \mathbb{Z}_+^d$ such that $\|\mathbf{t}\|_1 = K (= 2k - 1)$ to denote a multi-index. In addition, we denote monomial $\alpha_i^{\mathbf{t}} = \alpha_{i,1}^{t_1} \alpha_{i,2}^{t_2} \dots \alpha_{i,d}^{t_d}$ for every discrete distribution $\alpha_i = (\alpha_{i,1}, \alpha_{i,2}, \dots, \alpha_{i,d}) \in \mathbb{R}^d$. The moment vector of ϑ is then defined as

$$M(\vartheta) = \{M_{\mathbf{t}}(\vartheta)\}_{\|\mathbf{t}\|_1 \leq K} \text{ where } M_{\mathbf{t}}(\vartheta) = \int_{\Delta_{d-1}} \alpha^{\mathbf{t}} \vartheta(d\alpha) = \sum_{i=1}^k w_i \alpha_i^{\mathbf{t}}. \quad (1)$$

We define the *moment distance* between two mixtures ϑ and ϑ' as the L_∞ norm of the difference between corresponding moment vectors: $\text{Mom}_K(\vartheta, \vartheta') := \max_{\|\mathbf{t}\|_1 \leq K} |M_{\mathbf{t}}(\vartheta) - M_{\mathbf{t}}(\vartheta')|$.

Transportation Distance: For any two probability measures P and Q defined over \mathbb{R}^d , we define the L_1 -transportation distance $\text{Tran}(P, Q)$ is defined as

$$\text{Tran}(P, Q) := \inf \left\{ \int \|x - y\|_1 d\mu(x, y) : \mu \in M(P, Q) \right\}$$

where $M(P, Q)$ is the set of all joint distributions (also called coupling) on $\mathbb{R}^d \times \mathbb{R}^d$ with marginals P and Q . Transportation distance is also called Rubinstein distance, Wasserstein distance or earth mover distance in the literature. We also need to define transportation distance for signed measures and complex domains. See Appendix A.

4. An Efficient Algorithm for the k -Coin Problem

In this section, we study the k -coin model, i.e., ϑ is a k -spike distribution over $[0, 1]$. We present an efficient algorithm reconstructing the mixture from the first $K = 2k - 1$ noisy moments. Let $\vartheta := (\alpha, w) \in \text{Spike}(\Delta_1, \Delta_{k-1})$ be the ground truth mixture where $\alpha = [\alpha_1, \dots, \alpha_k]^\top \in \Delta_1^k$ and $w = [w_1, \dots, w_k]^\top \in \Delta_k$. Let $M(\vartheta) = [M_0(\vartheta), \dots, M_{2k-1}(\vartheta)]^\top$ be the ground truth moment vector containing moments of degree at most $K = 2k - 1$. We denote the noisy moment vector as $M' = [M'_0, \dots, M'_{2k-1}]$ where the error is bounded by $\|M' - M(\vartheta)\|_\infty \leq \xi$. Since ϑ is a distribution, we assume $M'_0 = M_0(\vartheta) = 1$. In the light of lower bound in Rabani et al. (2014) or Lemma 8, we further assume the noise satisfies $\xi \leq 2^{-\Omega(k)}$.

For vector $M = [M_0, \dots, M_{2k-1}]^\top$ and $\alpha = [\alpha_1, \dots, \alpha_k]^\top$, we denote

$$A_M := \begin{bmatrix} M_0 & M_1 & \cdots & M_{k-1} \\ M_1 & M_2 & \cdots & M_k \\ \vdots & \vdots & \ddots & \vdots \\ M_{k-1} & M_k & \cdots & M_{2k-2} \end{bmatrix}, \quad b_M := \begin{bmatrix} M_k \\ M_{k+1} \\ \vdots \\ M_{2k-1} \end{bmatrix}, \quad V_\alpha := \begin{bmatrix} \alpha_1^0 & \alpha_2^0 & \cdots & \alpha_k^0 \\ \alpha_1^1 & \alpha_2^1 & \cdots & \alpha_k^1 \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_1^{2k-1} & \alpha_2^{2k-1} & \cdots & \alpha_k^{2k-1} \end{bmatrix}. \quad (2)$$

Algorithm 1 Algorithm for the k-coin problem

Input: number of spikes k , noisy moments $M'(\cdot)$, noise level ξ
Output: recovered spike distribution $\tilde{\vartheta}$

- | | |
|-------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| 1: $\hat{\mathbf{c}} \leftarrow \arg \min_{\mathbf{x} \in \mathbb{R}^k} \ A_{M'} \mathbf{x} + b_{M'}\ _2^2 + \xi^2 \ \mathbf{x}\ _2^2$ | $\Rightarrow \hat{\mathbf{c}} = [\hat{c}_0, \dots, \hat{c}_{k-1}]^\top \in \mathbb{R}^k$ |
| 2: $\hat{\boldsymbol{\alpha}} \leftarrow \text{roots}(\sum_{i=0}^{k-1} \hat{c}_i x^i + x^k)$ | $\Rightarrow \hat{\boldsymbol{\alpha}} = [\hat{\alpha}_1, \dots, \hat{\alpha}_k]^\top \in \mathbb{C}^k$ |
| 3: $\bar{\boldsymbol{\alpha}} \leftarrow \text{project}_{\Delta_1}(\hat{\boldsymbol{\alpha}})$ | $\Rightarrow \bar{\boldsymbol{\alpha}} = [\bar{\alpha}_1, \dots, \bar{\alpha}_k]^\top \in \Delta_1^k$ |
| 4: $\tilde{\boldsymbol{\alpha}} \leftarrow \bar{\boldsymbol{\alpha}} + \text{Noise}(\xi)$ | $\Rightarrow \tilde{\boldsymbol{\alpha}} = [\tilde{\alpha}_1, \dots, \tilde{\alpha}_k]^\top \in \Delta_1^k$ |
| 5: $\hat{\mathbf{w}} \xleftarrow{O(1)\text{-approx}} \arg \min_{\mathbf{x} \in \mathbb{R}^k} \ V_{\tilde{\boldsymbol{\alpha}}} \mathbf{x} - M'\ _2^2$ | $\Rightarrow \hat{\mathbf{w}} = [\hat{w}_1, \dots, \hat{w}_k]^\top \in \mathbb{R}^k$ |
| 6: $\tilde{\mathbf{w}} \leftarrow \hat{\mathbf{w}} / (\sum_{i=1}^k \hat{w}_i)$ | $\Rightarrow \tilde{\mathbf{w}} = [\tilde{w}_1, \dots, \tilde{w}_k]^\top \in \Sigma_{k-1}$ |
| 7: $\tilde{\vartheta} \leftarrow (\tilde{\boldsymbol{\alpha}}, \tilde{\mathbf{w}})$ | $\Rightarrow \tilde{\vartheta} \in \text{Spike}(\Delta_1, \Sigma_{k-1})$ |
| 8: $\check{\mathbf{w}} \leftarrow \arg \min_{\mathbf{x} \in \Delta_{k-1}} \text{Tran}(\tilde{\vartheta}, (\tilde{\boldsymbol{\alpha}}, \mathbf{x}))$ | |
| 9: $\check{\vartheta} \leftarrow (\tilde{\boldsymbol{\alpha}}, \check{\mathbf{w}})$ | $\Rightarrow \check{\vartheta} \in \text{Spike}(\Delta_1, \Delta_{k-1})$ |
-

Note that A_M is a Hankel matrix and $V_{\boldsymbol{\alpha}}$ is a Vandemonde matrix.

Our algorithm is a variant of Prony's method (de Prony, 1795). The pseudocode can be found in Algorithm 1. The algorithm takes the number of spikes k , the noisy moment vector M' which $\|M' - M(\vartheta)\|_\infty \leq \xi$ and the moment accuracy ξ as the input. We describe our algorithm as follows.

Let $\mathbf{c} = [c_0, \dots, c_{k-1}]^\top \in \mathbb{R}^k$ such that $\prod_{i=1}^k (x - \alpha_i) = \sum_{i=0}^{k-1} c_i x^i + x^k$ be the characteristic vector of locations α_i . We first perform a ridge regression to obtain $\hat{\mathbf{c}}$ in Line 1. We note that $A_{M(\vartheta)} \mathbf{c} + b_{M(\vartheta)} = \mathbf{0}$. (see Lemma 24). Hence, $\hat{\mathbf{c}}$ serves a similar role as \mathbf{c} (note that $\hat{\mathbf{c}}$ is not necessarily close to \mathbf{c} without the separation assumption). From Line 2 to Line 4, we aim to obtain estimations of the positions of the spikes, i.e., α_i s. We first solve the roots of polynomial $\sum_{i=0}^{k-1} \hat{c}_i x^i + x^k$. For polynomial root findings, note that some roots may be complex without any separation assumption.³ Hence, we need to project the solutions back to Δ_1 and inject small noise, ensuring that all values are distinct and still in Δ_1 . We note that any noise of size at most ξ suffices. After recovering the positions of the spikes, we aim to recover the corresponding weights $\tilde{\mathbf{w}}$ by the linear regression defined by the Vandemonde matrix $V_{\boldsymbol{\alpha}}$ in Line 5 with weight normalization in Line 6. We note that $\tilde{\mathbf{w}}$ may still have some negative components. In Line 8, we find the closest k -spike distribution in $\text{Spike}(\Delta_1, \Delta_{k-1})$, which is our final output. The details for implementing the above steps in $O(k^2)$ time can be found in Appendix C.1. The performance guarantee of the algorithm is given by the following theorem.

Theorem 3 *Let ϑ be an arbitrary k -spike distribution over $[0, 1]$. Suppose we have noisy moments M'_t such that $|M'_t(\vartheta) - M'_t| \leq (\epsilon/k)^{O(k)}$ for $0 \leq t \leq 2k - 1$. Then Algorithm 1 outputs a mixture $\tilde{\vartheta}$ such that $\text{Tran}(\tilde{\vartheta}, \vartheta) \leq O(\epsilon)$ using $O(k^2)$ arithmetic operations.*

4.1. Proof Sketch

In this section, we present a proof sketch of Theorem 3. The complete proof of the theorem can be found in Appendix C.2. Firstly, using the relationship between $\hat{\mathbf{c}}$ and $\hat{\boldsymbol{\alpha}}$, one can show that

$$\sum_{i=1}^k w_i \prod_{j=1}^k (\alpha_i - \hat{\alpha}_j)^2 = \sum_{i=1}^k w_i \left(\sum_{j=0}^{k-1} \hat{c}_j \alpha_i^j + \alpha_i^k \right)^2 = \sum_{i=0}^{k-1} (\hat{c}_i - c_i) \left(A_{M(\vartheta)} \hat{\mathbf{c}} + b_{M(\vartheta)} \right)_i$$

3. Gordon et al. (2020) proved that all roots are real and separated under the minimal separation assumption.

Since we use the ridge regression to get $\hat{\mathbf{c}}$ (Line 1), the norms of both $A_{M'}\hat{\mathbf{c}} + b_{M'}$ and $\hat{\mathbf{c}}$ are small. Note that M' is close to $M(\boldsymbol{\vartheta})$, the norm of vector $A_{M(\boldsymbol{\vartheta})}\hat{\mathbf{c}} + b_{M(\boldsymbol{\vartheta})}$ is also small. Hence one can get

$$\sum_{i=1}^k w_i \prod_{j=1}^k (\alpha_i - \hat{\alpha}_j)^2 \leq 2^{O(k)} \cdot \zeta.$$

The above inequality shows that every ground truth spike $\alpha_i \in \boldsymbol{\alpha}$ has a close spike $\tilde{\alpha}_j \in \tilde{\boldsymbol{\alpha}}$ where the closeness depends on the weight of the spike w_i . We note that this property tolerates permutation and small spikes weights, enabling us to analyze mixtures without separation.

Next, note we can decompose $M(\boldsymbol{\vartheta}) = \sum_{i=1}^k w_i M(\alpha_i)$ where $M(\alpha_i)$ is the moment vector generated by the single-spike distribution located at α_i . By triangular inequality, the error of linear regression (Line 5) can be decomposed by $\min_{\mathbf{x}} \|V_{\tilde{\boldsymbol{\alpha}}}\mathbf{x} - M(\boldsymbol{\vartheta})\|_2 \leq \sum_{i=1}^k w_i \min_{\mathbf{x}} \|V_{\tilde{\boldsymbol{\alpha}}}\mathbf{x} - M(\alpha_i)\|_2$. Therefore, to bound $\|V_{\tilde{\boldsymbol{\alpha}}}\hat{\mathbf{w}} - M(\boldsymbol{\vartheta})\|_2 \approx \min_{\mathbf{x}} \|V_{\tilde{\boldsymbol{\alpha}}}\mathbf{x} - M(\boldsymbol{\vartheta})\|_2$, it is sufficient to prove $\min_{\mathbf{x}} \|V_{\tilde{\boldsymbol{\alpha}}}\mathbf{x} - M(\alpha_i)\|_2$ for each α_i is bounded. We use $\|V_{\tilde{\boldsymbol{\alpha}}}\mathbf{x}^* - M(\alpha_i)\|_2$ for a specific \mathbf{x}^* as its upper bound. \mathbf{x}^* is chosen to be the vector in which the first k rows of $V_{\tilde{\boldsymbol{\alpha}}}\mathbf{x}^*$ and $M(\alpha_i)$ are equal (\mathbf{x}^* is the solution of a linear system). By binomial expansion, we can compute the j th row of $V_{\tilde{\boldsymbol{\alpha}}}\mathbf{x}^*$ for large $j > k$ by

$$(V_{\tilde{\boldsymbol{\alpha}}}\mathbf{x}^*)_j = \sum_{p=1}^j \binom{j}{p} \alpha_i^{j-p} \sum_{t=1}^k x_t^* (\tilde{\alpha}_t - \alpha_i)^p.$$

Expressing \mathbf{x}^* using the adjoint matrix, using Cramer's rule, and plugging it into the right-hand side of the above equality, the resulting term appears to be difficult to bound (see the third equation in the analysis of Lemma 32). Fortunately, by leveraging a matrix equation derived from the Shur polynomial, one can get the following nice equation for the inner summation,

$$\sum_{t=1}^k x_t^* (\tilde{\alpha}_t - \alpha_i)^p = \prod_{t=1}^k (\tilde{\alpha}_t - \alpha_i) \cdot \sum_{s \subseteq [k], |s|=j-k} \prod_{t=1}^k (\tilde{\alpha}_t - \alpha_i)^{s_t},$$

where we use $(a)^b \subset \{0, \dots, b\}^a$ to denote the set that contains all vector s in which $\sum_{i=1}^a s_i = b$. Since k^{j-k} has no more than $2^{O(k)}$ elements, one can thus easily bound the approximation error by:

$$\min_{\mathbf{x}} \|V_{\tilde{\boldsymbol{\alpha}}}\mathbf{x} - M(\alpha_i)\|_2 \leq 2^{O(k)} \cdot \prod_{j=1}^k |\alpha_i - \tilde{\alpha}_j|$$

from the right-hand side of the above equation.

Note that $\tilde{\alpha}_j$ are $\hat{\alpha}_i$ close, combining all the above results gives $\|V_{\tilde{\boldsymbol{\alpha}}}\hat{\mathbf{w}} - M(\boldsymbol{\vartheta})\|_\infty \leq 2^{O(k)} \cdot \sqrt{\zeta}$, indicating moment distance between the recovered distribution and the ground truth is small. Finally, one can reach the desired statement using the moment-transportation inequality (Lemma 8), which asserts that if the moment distance between two measures is sufficiently small, the transportation distance between them is also small.

5. Efficient Algorithms for Higher Dimensions

In this section, we solve the problem in higher dimensions. We first present the algorithm for the two-dimensional problem. The algorithm for higher dimensions takes both the one-dimensional and two-dimensional algorithms as subroutines.

5.1. An Efficient Algorithm for the Two-Dimensional Problem

We first generalize the one-dimensional algorithm described in Section 4 to the dimension of two. Let $\vartheta := (\alpha, \mathbf{w}) \in \text{Spike}(\Delta_2, \Delta_{k-1})$ be the underlying mixture where $\alpha = [\alpha_1, \dots, \alpha_k]^\top$ for two-dimensional points $\alpha_i = (\alpha_{i,1}, \alpha_{i,2})$ and weight vector $\mathbf{w} = [w_1, \dots, w_k]^\top \in \Delta_{k-1}$. The true moments is defined by $M_{i,j}(\vartheta) = \sum_{t=1}^k w_t \alpha_{t,1}^i \alpha_{t,2}^j$. We assume the input noisy moments $M'_{i,j}$ satisfies $|M'_{i,j} - M_{i,j}(\vartheta)| \leq \xi$ for every $0 \leq i, j, i+j \leq 2k-1$. We further assume that $M'_{0,0} = M_{0,0}(\vartheta) = 1$ and noise level $\xi \leq 2^{-\Omega(k)}$.

The key idea for our algorithm is that, a distribution supported in \mathbb{R}^2 can be mapped to a distribution supported in the complex plane \mathbb{C} . In particular, let complex simplex $\Delta_{\mathbb{C}} = \{a + bi \mid (a, b) \in \Delta_2\}$. We denote $\beta = [\beta_1, \dots, \beta_k]^\top := [\alpha_{1,1} + \alpha_{1,2}i, \dots, \alpha_{k,1} + \alpha_{k,2}i]^\top \in \Delta_{\mathbb{C}}^k$ and define $\phi := (\beta, \mathbf{w}) \in \text{Spike}(\Delta_{\mathbb{C}}^k, \Delta_{k-1})$ to be the complex mixture corresponding to ϑ . The corresponding moments of ϕ can thus be defined as $G_{i,j}(\phi) = \sum_{t=1}^k w_t (\beta_t^\dagger)^i \beta_t^j$, where β_t^\dagger is the complex conjugate of β_t . We note $G_{i,j}$ can be computed from $M_{i,j}$ (see Lemma 37). Therefore, the task reduces to the recovery of the complex mixture ϕ given noisy moments $G'_{i,j}$. The algorithm and the proof of the following theorem are similar to that in Section 4. Handling complex numbers requires several minor changes, and we present the full details in Appendix D. For the analysis, we also need to generalize the moment-transportation inequality to the complex domains (see Appendix B.2).

Theorem 4 (*The Two-dimensional Problem*) *Let $\vartheta \in \text{Spike}(\Delta_2, \Delta_{k-1})$. Suppose we have noisy moments $M_{i,j}$ ($0 \leq i, j, i+j \leq 2k-1$) up to precision $(\epsilon/k)^{O(k)}$. We can obtain a mixture $\tilde{\vartheta}$ such that $\text{Tran}(\tilde{\vartheta}, \vartheta) \leq O(\epsilon)$ using only $O(k^3)$ arithmetic operations.*

5.2. An Efficient Algorithm for Dimension $d \geq 3$

Now, we present our algorithm for higher dimensions. Let $\vartheta := (\alpha, \mathbf{w}) \in \text{Spike}(\Delta_{d-1}, \Delta_{k-1})$ be the underlying mixture where $\alpha = [\alpha_1, \dots, \alpha_k]^\top \in \mathbb{R}^{k \times d}$ such that $\alpha_i = [\alpha_{i,1}, \dots, \alpha_{i,d}]^\top$ and $\mathbf{w} = [w_1, \dots, w_k]^\top$. We denote $\alpha_{:,i} = [\alpha_{1,i}, \dots, \alpha_{k,i}]$ as the vector of the i th coordinates of the spikes and also write $\tilde{\alpha} = [\alpha_{:,1}, \dots, \alpha_{:,k}]^\top$.

Since the number of different moments is exponential in d , we consider the setting in which one can access the noisy moments of some linear maps of ϑ onto lower-dimensional subspaces, as in previous works (Rabani et al., 2014; Li et al., 2015) (such noisy moments can be easily obtained in applications such as topic modeling, see Section 6.1, or in learning Gaussian mixtures, see Section 6.2). For some $\vartheta = (\alpha, \mathbf{w}) \in \text{Spike}(\mathbb{R}^d, \Delta_{k-1})$ and a vector $\mathbf{r} \in \mathbb{R}^d$, we denote $\Pi_{\mathbf{r}}(\vartheta) := (\alpha \mathbf{r}, \mathbf{w}) \in \text{Spike}(\mathbb{R}, \Delta_{k-1})$ as the projected distribution of ϑ along vector \mathbf{r} . For a matrix $R = [\mathbf{r}_1, \dots, \mathbf{r}_p] \in \mathbb{R}^{d \times p}$, we also denote $\Pi_R(\vartheta) := (\alpha R, \mathbf{w}) \in \text{Spike}(\mathbb{R}^p, \Delta_{k-1})$ as the projected distribution of ϑ along multiple dimensions $\mathbf{r}_1, \dots, \mathbf{r}_p$. In this section, we consider mappings to one-dimensional and two-dimensional subspaces, which suffice for our purposes.

The idea of our algorithm is as follows: We first generate a random vector \mathbf{r} . Then with some constant probability, the distance between spikes is roughly kept after the projection along \mathbf{r} . This enables us to recover a projected measure $\tilde{\phi}$ along \mathbf{r} of the discrete mixture by running the one-dimensional algorithm on the projected moments. Next, we try to assign every spike of the projected measure a coordinate in \mathbb{R}^d . This is done by running the two-dimensional algorithm over the plane spanned by \mathbf{r} and unit vector e_t for each dimension $t \in [d]$ to get $\hat{\phi}_t$, and assigning the t -th coordinate of each spike in $\tilde{\phi}$ the coordinate of the closet spike in $\hat{\phi}_t$. We use a geometric argument to show the recovered distribution is close to the ground truth ϑ in terms of the transportation distance (see Lemma 47). The algorithm and the analysis are deferred to Appendix E.

Theorem 5 *Let ϑ be an arbitrary k -spike mixture supported in Δ_{d-1} . Suppose we can access noisy \mathbf{t} -moments of $\Pi_R(\vartheta)$, with precision $(\epsilon/(dk))^{O(k)}$, where $\Pi_R(\vartheta)$ is the projected measure obtained by applying linear transformation R (with $\|R\|_\infty \leq 1$) of ϑ onto any h -dimensional subspace we choose, for all $\|\mathbf{t}\| \leq K = 2k - 1$ and $h = 1, 2$ (i.e., lines and $2d$ planes). We can construct a k -spike mixture $\tilde{\vartheta}$ such that $\text{Tran}(\tilde{\vartheta}, \vartheta) \leq O(\epsilon)$ using only $O(dk^3)$ arithmetic operations with probability at least 0.99.*

6. Applications

6.1. Applications to Topic Models

This section discusses the application of our results to topic models. Due to space limitations, the full details are deferred to Appendix F. Consider the ‘‘bag of words’’ model in which we take each document as an unordered multiset of words. Let d be the number of different words in the document. We assume there are k topics $\{\alpha_1, \alpha_2, \dots, \alpha_k\}$ that are distributed over $[d]$. Each document is generated by first selecting a topic $i \in \Delta_{d-1}$ from the mixture ϑ , and then sampling K i.i.d. words according to α_i from this topic. Here, K is referred to as *the snapshot number* and we call such a sample a K -*snapshot* of p . Our goal is to recover $\vartheta \in \text{Spike}(\Delta_{d-1}, \Delta_{k-1})$ that is a discrete distribution over k pure topics. If $d \gg k$, we can use the dimension reduction technique in Li et al. (2015) to reduce the dimension d to $O(k)$, using $\tilde{O}(\epsilon^{-6}k^3d)$ 1-snapshots and $\tilde{O}(\epsilon^{-4}k^2d)$ 2-snapshots, respectively. Then, we can apply Theorem 5 to derive the required moment accuracy. Translating the moment accuracy to the sample complexity is standard. In sum, we can obtain the following result.

Theorem 6 *There is an algorithm that can learn an arbitrary k -spike mixture supported in Δ_{d-1} for any d within L_1 transportation distance $O(\epsilon)$ with probability at least 0.99 using $\text{poly}(d, k, \frac{1}{\epsilon})$, $\text{poly}(d, k, \frac{1}{\epsilon})$, many 1-, 2-, snapshots and $(k/\epsilon)^{O(k)}$ many $(2k - 1)$ -snapshots.*

6.2. Applications to Gaussian Mixture Learning

This section shows how to leverage our results for sparse moment problems to obtain improved algorithms for learning Gaussian mixtures. The full details can be found in Appendix G. We consider the following setting studied in Wu and Yang (2020); Doss et al. (2020). We parameterise a k -Gaussian mixture in \mathbb{R}^d as $\vartheta_N = (\alpha, \mathbf{w}, \Sigma)$. Here, $\alpha = \{\alpha_1, \alpha_2, \dots, \alpha_k\}$ and $\mathbf{w} = \{w_1, w_2, \dots, w_k\} \in \Delta_{k-1}$ where $\alpha_i \in \mathbb{R}^d$ and $w_i \in [0, 1]$ represents the mean and weight of i th component. We assume all k mixture components share a common covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$ that is known in advance and the maximal eigenvalue of Σ is bounded by a constant. We focus on parameter learning, that is, to learn the parameter α and \mathbf{w} given known covariance matrix Σ and a set of i.i.d. samples from $\vartheta_N = \sum_{i=1}^k w_i N(\alpha_i, \Sigma)$. We have the following result:

Theorem 7 *There is an algorithm that can learn an arbitrary d -dimensional k -spike Gaussian mixture (with known common covariance matrix Σ) within transportation distance $O(\epsilon)$ with probability at least 0.99 using $(k/\epsilon)^{O(k)} + \text{poly}(1/\epsilon, d, k)$ samples.*

7. Concluding Remarks

We provide efficient algorithms for learning the k -mixture over discrete distributions from noisy moments. We measure the accuracy in terms of transportation distance, and our analysis is independent of the minimal separation. The techniques used in our analysis for the one-dimensional case may be useful in the perturbative analysis of other problems involving the Hankel matrix or the Vandermonde matrix (such as super-resolution).

Our problem is a special case of learning mixtures of product distribution (Gordon et al., 2021a), which is further a special case of the general problem of learning mixtures of Bayesian networks (Gordon et al., 2021b), which has important applications in causal inferences. In fact, the algorithms in Gordon et al. (2021a,b) used the algorithm for sparse Hausdorff moment problem as a subroutine. It would be interesting to see if our techniques can be useful in these problems.

8. Acknowledgement

We would like to thank Yuval Rabani and Leonard Schulman for their insightful discussions and anonymous reviewers for many helpful comments. In particular, Leonard Schulman pointed out that Lemma 31 is related to the Schur polynomial. JL also would like to thank Xuan Wu for several discussions around 2015 and a version of Lemma 8 was proved back then, independent from Wu and Yang (2020). The authors are supported in part by the National Natural Science Foundation of China Grant 62161146004.

References

- Dimitris Achlioptas and Frank McSherry. On spectral learning of mixtures of distributions. In *International Conference on Computational Learning Theory*, pages 458–469. Springer, 2005.
- Alok Aggarwal, Amotz Bar-Noy, Samir Khuller, Dina Kravets, and Baruch Schieber. Efficient minimum cost matching using quadrangle inequality. *Proceedings., 33rd Annual Symposium on Foundations of Computer Science*, pages 583–592, 1992.
- Josh Alman and Virginia Vassilevska Williams. A refined laser method and faster matrix multiplication. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 522–539. SIAM, 2021.
- Anima Anandkumar, Dean P. Foster, Daniel J. Hsu, Sham M. Kakade, and Yi-Kai Liu. Two svds suffice: Spectral decompositions for probabilistic topic modeling and latent dirichlet allocation. *ArXiv*, abs/1204.6703, 2012a.
- Animashree Anandkumar, Daniel Hsu, and Sham M Kakade. A method of moments for mixture models and hidden markov models. In *Conference on Learning Theory*, pages 33–1. JMLR Workshop and Conference Proceedings, 2012b.
- Dang Dinh Ang, Rudolf Gorenflo, Vy Khoi Le, and Dang Duc Trong. *Moment Theory and Some Inverse Problems in Potential Theory and Heat Conduction*. 2002.
- Sanjeev Arora and Ravi Kannan. Learning mixtures of separated nonspherical gaussians. *Annals of Applied Probability*, 15:69–92, 2005.

- Sanjeev Arora, Rong Ge, and Ankur Moitra. Learning topic models—going beyond svd. In *2012 IEEE 53rd annual symposium on foundations of computer science*, pages 1–10. IEEE, 2012.
- Sanjeev Arora, Rong Ge, Yoni Halpern, David Mimno, Ankur Moitra, David Sontag, Yichen Wu, and Michael Zhu. Learning topic models—provably and efficiently. *Communications of the ACM*, 61(4):85–93, 2018.
- Ainesh Bakshi, Ilias Diakonikolas, He Jia, Daniel M Kane, Pravesh K Kothari, and Santosh S Vempala. Robustly learning mixtures of k arbitrary gaussians. *Proceedings of the 54th Annual ACM symposium on Theory of Computing*, 2022.
- Edward J Barbeau. *Polynomials*. Springer Science & Business Media, 2003.
- Tuğkan Batu, Sudipto Guha, and Sampath Kannan. Inferring mixtures of Markov chains. In *International Conference on Computational Learning Theory*, pages 186–199. Springer, 2004.
- Mikhail Belkin and Kaushik Sinha. Polynomial learning of distribution families. *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 103–112, 2010.
- JD Bessis, JM Drouffe, and P Moussa. Positivity constraints for the ising ferromagnetic model. *Journal of Physics A: Mathematical and General*, 9(12):2105, 1976.
- David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- S Charles Brubaker and Santosh S Vempala. Isotropic pca and affine-invariant clustering. In *Building Bridges*, pages 241–281. Springer, 2008.
- Emmanuel J. Candès and Carlos Fernandez-Granda. Super-resolution from noisy data. *Journal of Fourier Analysis and Applications*, 19:1229–1254, 2012.
- Emmanuel J Candès and Carlos Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on pure and applied Mathematics*, 67(6):906–956, 2014.
- Kamalika Chaudhuri and Satish Rao. Learning mixtures of product distributions using correlations and independence. In *COLT*, volume 4, pages 9–20, 2008.
- Xue Chen, Daniel M. Kane, Eric Price, and Zhao Song. Fourier-sparse interpolation without a frequency gap. *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 741–750, 2016.
- Theodore S Chihara. *An introduction to orthogonal polynomials*. Courier Corporation, 2011.
- Mary Cryan, Leslie Ann Goldberg, and Paul W Goldberg. Evolutionary trees can be learned in polynomial time in the two-state general markov model. *SIAM Journal on Computing*, 31(2):375–397, 2001.
- Anirban Dasgupta, John E. Hopcroft, Jon M. Kleinberg, and Mark Sandler. On learning mixtures of heavy-tailed distributions. *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05)*, pages 491–500, 2005.

- Sanjoy Dasgupta. Learning mixtures of gaussians. In *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pages 634–644. IEEE, 1999.
- Sanjoy Dasgupta and Leonard J Schulman. A probabilistic analysis of em for mixtures of separated, spherical gaussians. *Journal of Machine Learning Research*, 8:203–226, 2007.
- Constantinos Daskalakis, Ilias Diakonikolas, and Rocco A Servedio. Learning k-modal distributions via testing. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 1371–1385. SIAM, 2012.
- R. de Prony. *Essai expérimentale et analytique*. J. Écol. Polytech. , 1(2):24–76, 1795.
- David L. Donoho. Superresolution via sparsity constraints. *Siam Journal on Mathematical Analysis*, 23:1309–1331, 1992.
- Natalie Doss, Yihong Wu, Pengkun Yang, and Harrison H Zhou. Optimal estimation of high-dimensional location gaussian mixtures. *arXiv preprint arXiv:2002.05818*, 2020.
- Richard M Dudley. *Real analysis and probability*, volume 74. Cambridge University Press, 2002.
- Jon Feldman, Rocco A Servedio, and Ryan O’Donnell. Pac learning axis-aligned mixtures of gaussians with no separation assumption. In *International Conference on Computational Learning Theory*, pages 20–34. Springer, 2006.
- Jon Feldman, Ryan O’Donnell, and Rocco A Servedio. Learning mixtures of product distributions over discrete domains. *SIAM Journal on Computing*, 37(5):1536–1564, 2008.
- Walter Gautschi and Gabriele Inglese. Lower bounds for the condition number of vandermonde matrices. *Numerische Mathematik*, 52(3):241–250, 1987.
- Spencer Gordon, Bijan Mazaheri, Leonard J Schulman, and Yuval Rabani. The sparse hausdorff moment problem, with application to topic models. *arXiv preprint arXiv:2007.08101*, 2020.
- Spencer Gordon, Bijan H Mazaheri, Yuval Rabani, and Leonard Schulman. Source identification for mixtures of product distributions. In *Conference on Learning Theory*, pages 2193–2216. PMLR, 2021a.
- Spencer L Gordon, Bijan Mazaheri, Yuval Rabani, and Leonard J Schulman. Identifying mixtures of bayesian network distributions. *arXiv preprint arXiv:2112.11602*, 2021b.
- Richard Hamming. *Numerical methods for scientists and engineers*. Courier Corporation, 2012.
- Thomas Hofmann. Probabilistic latent semantic analysis. In *UAI*, 1999.
- Thomas Hofmann and Jan Puzicha. Latent class models for collaborative filtering. In *IJCAI*, volume 99, 1999.
- Yingbo Hua and Tapan Kumar Sarker. Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise. *IEEE Trans. Acoust. Speech Signal Process.*, 38: 814–824, 1990.

- Baihe Huang, Shunhua Jiang, Zhao Song, Runzhou Tao, and Ruizhe Zhang. Solving sdp faster: A robust ipm framework and efficient implementation. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 233–244. IEEE, 2022.
- Qingqing Huang and Sham M Kakade. Super-resolution off the grid. *Advances in Neural Information Processing Systems*, 28, 2015.
- 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, 2010.
- Ravindran Kannan, Hadi Salmasian, and Santosh Vempala. The spectral method for general mixture models. *SIAM Journal on Computing*, 38(3):1141–1156, 2008.
- Michael Kearns, Yishay Mansour, Dana Ron, Ronitt Rubinfeld, Robert E Schapire, and Linda Sellie. On the learnability of discrete distributions. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 273–282, 1994.
- Younhun Kim, Frederic Koehler, Ankur Moitra, Elchanan Mossel, and Govind Ramnarayan. How many subpopulations is too many? exponential lower bounds for inferring population histories. *J. Comput. Biol.*, 27:613–625, 2019.
- Jon Kleinberg and Mark Sandler. Using mixture models for collaborative filtering. *Journal of Computer and System Sciences*, 74(1):49–69, 2008.
- Jean Bernard Lasserre. *Moments, positive polynomials and their applications*, volume 1. World Scientific, 2009.
- Jian Li and Amol Deshpande. Ranking continuous probabilistic datasets. *Proceedings of the VLDB Endowment*, 3(1-2):638–649, 2010.
- Jian Li, Yuval Rabani, Leonard J Schulman, and Chaitanya Swamy. Learning arbitrary statistical mixtures of discrete distributions. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, 2015.
- Allen Liu and Ankur Moitra. Efficiently learning mixtures of mallows models. *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 627–638, 2018.
- Allen Liu and Ankur Moitra. Settling the robust learnability of mixtures of gaussians. *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, 2021.
- Raphael A Meyer, Cameron Musco, Christopher Musco, David P Woodruff, and Samson Zhou. Fast regression for structured inputs. *arXiv preprint arXiv:2203.07557*, 2022.
- Ankur Moitra. Super-resolution, extremal functions and the condition number of vandermonde matrices. *Proceedings of the forty-seventh annual ACM symposium on Theory of Computing*, 2015.
- Ankur Moitra and Gregory Valiant. Settling the polynomial learnability of mixtures of gaussians. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 93–102. IEEE, 2010.

- Elchanan Mossel and Sébastien Roch. Learning nonsingular phylogenies and hidden markov models. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 366–375, 2005.
- Frank Natterer. *The mathematics of computerized tomography*. SIAM, 2001.
- C Andrew Neff and John H Reif. An efficient algorithm for the complex roots problem. *Journal of Complexity*, 12(2):81–115, 1996.
- Victor Y Pan. How bad are vandermonde matrices? *SIAM Journal on Matrix Analysis and Applications*, 37(2):676–694, 2016.
- Christos H Papadimitriou, Prabhakar Raghavan, Hisao Tamaki, and Santosh Vempala. Latent semantic indexing: A probabilistic analysis. *Journal of Computer and System Sciences*, 61(2):217–235, 2000.
- Maria B Pintarelli and Fernando Vericat. On the numerical solution of the linear and nonlinear poisson equations seen as bi-dimensional inverse moment problems. *Journal of Interdisciplinary Mathematics*, 19(5-6):927–944, 2016.
- Yuval Rabani, Leonard J. Schulman, and Chaitanya Swamy. Learning mixtures of arbitrary distributions over large discrete domains. *Proceedings of the 5th conference on Innovations in theoretical computer science*, 2014.
- Richard H. Roy and Thomas Kailath. Esprit-estimation of signal parameters via rotational invariance techniques. *IEEE Trans. Acoust. Speech Signal Process.*, 37:984–995, 1989.
- Konrad Schmüdgen. *The moment problem*, volume 9. Springer, 2017.
- Richard P. Stanley. *Enumerative Combinatorics: Volume 1*. Cambridge University Press, USA, 2nd edition, 2011. ISBN 1107602629.
- Gilbert Stewart. *Matrix perturbation theory*, 1990.
- Santosh Vempala and Grant Wang. A spectral algorithm for learning mixtures of distributions. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 113–122. IEEE, 2002.
- Yihong Wu and Pengkun Yang. Optimal estimation of gaussian mixtures via denoised method of moments. *The Annals of Statistics*, 48(4):1981–2007, 2020.
- Wei Xu and Sanzheng Qiao. A fast symmetric svd algorithm for square hankel matrices. *Linear Algebra and its Applications*, 428:550–563, 2008.

Appendix A. Transportation Distance for Non-Probability Measures

Let 1-Lip be the set of 1-Lipschitz functions on \mathbb{R}^d , i.e., $1\text{-Lip} := \{f : \mathbb{R}^d \rightarrow \mathbb{R} \mid |f(x) - f(y)| \leq \|x - y\|_1, \forall x, y \in \mathbb{R}^d\}$. We need the following important theorem by Kantorovich and Rubinstein (see e.g., [Dudley \(2002\)](#)), which states the dual formulation of transportation distance:

$$\text{Tran}(P, Q) = \sup \left\{ \int f d(P - Q) : f \in 1\text{-Lip} \right\}. \quad (3)$$

The dual formulation (3) can be generalized to non-probability measures. For two signed measures P and Q defined over \mathbb{R}^d with $\int P dx = \int Q dx = 1$, we still can use (3) to compute its transportation distance.

We now introduce the equivalent primal formulation as follows: For some signed measure R , we define measure R^+ as its positive components that $R^+(x) = \max\{R(x), 0\}$. According to the duality formulation (3), one can see $\text{Tran}(P, Q) = \text{Tran}((P - Q)^+, (Q - P)^+)$ where both $(P - Q)^+$ and $(Q - P)^+$ are non-negative. From the primal formulation of transportation distance, we conclude another definition for generalized transportation distance:

$$\text{Tran}(P, Q) = \inf \left\{ \int \|x - y\|_1 d\mu(x, y) : \mu \in M((P - Q)^+, (Q - P)^+) \right\}. \quad (4)$$

We also need to define the transportation distance over complex measures. Denote $\mathbb{B}_{\mathbb{C}} = \{\alpha \in \mathbb{C} : |\alpha| \leq 1\}$. For some complex measure P that $\int_{\mathbb{C}} dP(x) = 1$. We denote the real and imaginary components of P as P^r and P^i respectively that $P(x) = P^r(x) + iP^i(x)$ for $P^r(x), P^i(x) \in \mathbb{R}$. We generalize transportation distance to complex weights by

$$\text{Tran}(P, Q) = \text{Tran}(P^r, Q^r) + \text{Tran}(P^i, Q^i). \quad (5)$$

We also have this transportation distance in a dual form:

$$\begin{aligned} \text{Tran}(P, Q) &= \sup \left\{ \int f^r d(P^r - Q^r) : f^r \in 1\text{-Lip} \right\} + \sup \left\{ \int f^i d(P^i - Q^i) : f^i \in 1\text{-Lip} \right\} \\ &= \sup \left\{ \text{real} \left(\int (f^r - if^i) d(P - Q) \right) : f^r, f^i \in 1\text{-Lip} \right\} \\ &\leq 2 \sup \left\{ \int f d(P - Q) : f \in 1\text{-Lip}_{\mathbb{C}} \right\} \end{aligned}$$

where $1\text{-Lip}_{\mathbb{C}} := \{f : \mathbb{C} \rightarrow \mathbb{C} \mid |f(x) - f(y)| \leq |x - y| \text{ for any } x, y \in \mathbb{C}\}$ is the complex 1-Lipschitz functions on \mathbb{C} , and the last inequality holds since we can assign $f(x) = (f^r(x) - if^i(x))/2$.

Appendix B. Moment-Transportation Inequalities

Our analysis needs to bound the transportation distance by a function of moment distance. We first present a moment-transportation inequality in one dimension ([Lemma 8](#)), and then generalize the result to complex numbers and higher dimensions. The first such inequality in one dimension is obtained in [Rabani et al. \(2014, \(Lemma 5.4\)\)](#) (with worse parameters). [Lemma 8](#) is firstly proved by [Wu and Yang \(2020\)](#) in the context of learning Gaussian mixtures. Here, we prove the theorem in our notations for completeness, and later we extend the proof to complex domains in [Section B.2](#).

B.1. A One-Dimensional Moment-Transportation Inequality

We allow more general mixtures with spikes in $[-1, 1]$ and negative weights. This slight generalization will be useful since some intermediate coordinates and weights may be negative. Recall that $\Sigma_{k-1} = \{\mathbf{w} = (w_1, \dots, w_k) \in \mathbb{R}^k \mid \sum_{i=1}^k w_i = 1\}$.

Lemma 8 (*Wu and Yang (2020)*) *For any two mixtures with k components $\vartheta, \vartheta' \in \text{Spike}([-1, 1], \Sigma_{k-1})$, it holds that*

$$\text{Tran}(\vartheta, \vartheta') \leq O(k \text{Mom}_K(\vartheta, \vartheta')^{\frac{1}{2k-1}}).$$

We begin with some notations. Let $\text{Supp} = \text{Supp}(\vartheta) \cup \text{Supp}(\vartheta')$ and let $n = |\text{Supp}|$ ($n \leq 2k$). Arrange the points in Supp as $\alpha = (\alpha_1, \dots, \alpha_n)$ where $\alpha_1 < \dots < \alpha_n$.

Definition 9 *Let $\alpha = (\alpha_1, \dots, \alpha_n)$ such that $-1 \leq \alpha_1 < \dots < \alpha_n \leq 1$. Let \mathcal{P}_α denote the set of polynomials of degree at most $n - 1$ and 1-Lipschitz over the discrete points in α , i.e.,*

$$\mathcal{P}_\alpha = \{f \mid \deg f \leq n - 1, f(\alpha_1) = 0; |f(\alpha_i) - f(\alpha_j)| \leq |\alpha_i - \alpha_j| \forall i, j \in [n]\}.$$

Note that we do not require f to be 1-Lipschitz over the entire interval $[-1, 1]$. From the dual formulation of transportation distance (see e.g., [Dudley \(2002\)](#)), we have the following proposition:

Proposition 10 $\text{Tran}(\vartheta, \vartheta') = \sup \{ \int f d(\vartheta - \vartheta') : f \in 1\text{-Lip} \} = \sup \{ \int p d(\vartheta - \vartheta') : p \in \mathcal{P}_\alpha \}.$

We can view \mathcal{P}_α as a subset in the linear space of all polynomials of degree at most $n - 1$. In fact, \mathcal{P}_α is a convex polytope (a convex combination of two polynomials being 1-Lipschitz over α is also a polynomial that is 1-Lipschitz over α). The *height* of a polynomial $\sum_i c_i x^i$ is the maximum absolute coefficient $\max_i |c_i|$. As we will see shortly, the height of a polynomial in \mathcal{P}_α is related to the required moment accuracy, and we need the height to be upper bounded by a value independent of the minimum separation ζ . However, a polynomial in \mathcal{P}_α may have a very large height (depending on the inverse of the minimum separation $1/|\alpha_{i+1} - \alpha_i|$). This can be seen from the Lagrangian interpolation formula: $p(x) := \sum_{j=1}^n p(\alpha_j) \ell_j(x)$ where ℓ_j is the Lagrange basis polynomial $\ell_j(x) := \prod_{1 \leq m \leq n, m \neq j} (x - \alpha_m) / (\alpha_j - \alpha_m)$. To remedy this, one can show that for any polynomial in \mathcal{P}_α , there is an approximate polynomial with bounded height:

Lemma 11 *For any polynomial $p(x) \in \mathcal{P}_\alpha$ and $\eta > 0$, there is a polynomial $p_\eta(x)$ of degree at most $n - 1$ such that the following properties hold:*

1. $|p_\eta(\alpha_i) - p(\alpha_i)| \leq 2\eta$, for any $i \in [n]$.
2. The height of p_η is at most $n2^{2n-2}(\frac{n}{\eta})^{n-2}$.

Before proving Lemma 11, we show Lemma 8 can be easily derived from Proposition 10 and Lemma 11.

Proof [Proof of Lemma 8] Fix any constant $\eta > 0$. Let $\xi = \text{Mom}_K(\vartheta, \vartheta')$

$$\text{Tran}(\vartheta, \vartheta') = \sup_{f \in 1\text{-Lip}} \int f d(\vartheta - \vartheta') = \sup_{p \in \mathcal{P}_\alpha} \int p d(\vartheta - \vartheta') \quad (\text{Proposition 10})$$

$$\leq 4\eta + \sup_{p \in \mathcal{P}_\alpha} \int p_\eta d(\vartheta - \vartheta') \quad (\text{Lemma 11})$$

$$\leq 4\eta + \sup_{p \in \mathcal{P}_\alpha} \sum_{i=1}^{n-1} n 2^{2n-2} \left(\frac{n}{\eta}\right)^{n-2} \text{Mom}_K(\vartheta, \vartheta') \quad (\text{Lemma 11})$$

$$\leq 4\eta + 2^{4k} (2k)^{2k} \frac{\xi}{\eta^{2k-2}} \quad (n \leq 2k)$$

By choosing $\eta = 16k\xi^{\frac{1}{2k-1}}$, we conclude that $\text{Tran}(\vartheta, \vartheta') \leq O(k\xi^{\frac{1}{2k-1}})$. ■

Next, we prove Lemma 11. We introduce the following set \mathcal{F}_α of special polynomials. They are in fact the extreme points of the polytope \mathcal{P}_α .

Definition 12 Let $\alpha = (\alpha_1, \dots, \alpha_n)$ such that $-1 \leq \alpha_1 < \dots < \alpha_n \leq 1$. Let

$$\mathcal{F}_\alpha = \{p \mid \deg p \leq n-1, p(\alpha_1) = 0; |p(\alpha_{i+1}) - p(\alpha_i)| = |\alpha_{i+1} - \alpha_i| \forall i \in [n-1]\}$$

It is easy to see that $\mathcal{F}_\alpha \subset \mathcal{P}_\alpha$ and $|\mathcal{F}_\alpha| = 2^{n-1}$. Now, we modify each polynomial $p \in \mathcal{F}_\alpha$ slightly, as follows. Consider the intervals $[\alpha_1, \alpha_2], \dots, [\alpha_{n-1}, \alpha_n]$. If $\alpha_i - \alpha_{i-1} \leq \eta/n$, we say the interval $[\alpha_{i-1}, \alpha_i]$ is a *small interval*. Otherwise, it is *large*. We first merge all consecutive small intervals. For each resulting interval, we merge it with the large interval to its right. Note that we never merge two large intervals together. Let $S = \{\alpha_{i_1} = \alpha_1, \alpha_{i_2}, \dots, \alpha_{i_m} = \alpha_n\} \subseteq \alpha$ be the endpoints of the current intervals. It is easy to see the distance between any two points in S is at least η/n (since the current interval contains exactly one original large interval). Define a continuous piecewise linear function $L : [\alpha_1, \alpha_n] \rightarrow \mathbb{R}$ as follows: (1) $L(\alpha_1) = p(\alpha_1) = 0$. (2) The breaking points are the points in S ; (3) Each linear piece of L has slope either 1 or -1 ; for two consecutive breaking points $\alpha_{i_j}, \alpha_{i_{j+1}} \in S$, if $p(\alpha_{i_j}) > p(\alpha_{i_{j+1}})$, the slope of the corresponding piece is -1 . Otherwise, it is 1.

Lemma 13 For any $\alpha_i \in \alpha$, $|L(\alpha_i) - p(\alpha_i)| \leq 2\eta$.

Proof We prove inductively that $|L(\alpha_i) - p(\alpha_i)| \leq 2i\eta/n$. The base case ($i = 1$) is trivial. Suppose the lemma holds true for all value at most i and now we show it holds for $i + 1$. There are two cases. If $\alpha_{i+1} - \alpha_i \leq \eta/n$, we have

$$\begin{aligned} |L(\alpha_{i+1}) - p(\alpha_{i+1})| &\leq |L(\alpha_{i+1}) - L(\alpha_i)| + |L(\alpha_i) - p(\alpha_i)| + |p(\alpha_i) - p(\alpha_{i+1})| \\ &= 2|\alpha_{i+1} - \alpha_i| + |L(\alpha_i) - p(\alpha_i)| \leq 2(i+1)\eta/n. \end{aligned}$$

If $\alpha_{i+1} - \alpha_i > \eta/n$, we have $\alpha_{i+1} \in S$ (by the merge procedure). Suppose α_j ($\alpha \leq i$) is the point in S right before α_{i+1} . We can see that (by the definition of p):

$$|p(\alpha_j) - p(\alpha_{i+1})| \in |\alpha_{i+1} - \alpha_j| \pm |\alpha_i - \alpha_j| \subseteq |\alpha_{i+1} - \alpha_i| \pm (j-i)\eta/n.$$

In addition, we have

$$\begin{aligned} |L(\alpha_{i+1}) - p(\alpha_{i+1})| &\leq |L(\alpha_{i+1}) - L(\alpha_j) + p(\alpha_j) - p(\alpha_{i+1})| + |L(\alpha_j) - p(\alpha_j)| \\ &\leq 2(i+1-j)\eta/n + 2j\eta/n \leq 2(i+1)\eta/n, \end{aligned}$$

where second inequality holds since $L(\alpha_{i+1}) - L(\alpha_j) = \alpha_{i-1} - \alpha_j$ when $p(\alpha_{i+1}) > p(\alpha_j)$ and $L(\alpha_{i+1}) - L(\alpha_j) = \alpha_j - \alpha_{i+1}$ when $p(\alpha_{i+1}) \leq p(\alpha_j)$. \blacksquare

Let $T_\eta(p)$ be the polynomial with degree $\leq n-1$ that interpolates $L(\alpha_i)$, (i.e. $T_\eta(p)(\alpha_i) = L(\alpha_i)$ for all $i \in [n]$.) $T_\eta(p)$ is unique by this definition and by Lemma 13, $|T_\eta(p)(\alpha_i) - p(\alpha_i)| \leq 2i \cdot \frac{\eta}{n} \leq 2\eta$. Now, we show the height of $T_\eta(p)$ can be bounded, independent of the minimum separation $|\alpha_{i+1} - \alpha_i|$, due to its special structure.

Lemma 14 *For any polynomial $p(x) \in \mathcal{F}_\alpha$, define polynomial $T_\eta(p)$ as the above. Suppose $T_\eta(p) = \sum_{i=0}^{n-1} c_i x^i$. We have $|c_t| \leq n2^{2n-2} \left(\frac{n}{\eta}\right)^{n-2}$ for $t \in [n-1]$.*

Proof The key to the proof is Newton's interpolation formula, which we briefly review here (see e.g., Hamming (2012)). Let $F[\alpha_1, \dots, \alpha_i]$ be the i th order difference of $T_\eta(p)$, which can be defined recursively:

$$\begin{aligned} F[\alpha_t, \alpha_{t+1}] &= \frac{T_\eta(p)(\alpha_{t+1}) - T_\eta(p)(\alpha_t)}{\alpha_{t+1} - \alpha_t}, \dots, \\ F[\alpha_t, \dots, \alpha_{t+i}] &= \frac{F[\alpha_{t+1}, \dots, \alpha_{t+i}] - F[\alpha_t, \dots, \alpha_{t+i-1}]}{\alpha_{t+i} - \alpha_t}. \end{aligned}$$

Then, by Newton's interpolation formula, we can write that

$$T_\eta(p)(x) = F[\alpha_1, \alpha_2](x - \alpha_1) + F[\alpha_1, \alpha_2, \alpha_3](x - \alpha_1)(x - \alpha_2) + \dots + F[\alpha_1, \dots, \alpha_n] \prod_{i=1}^{n-1} (x - \alpha_i).$$

By the definition of $T_\eta(p)$, we know that every 2nd order difference of $T_\eta(p)$ is either 1 or -1 . Now, we show inductively that the absolute value of the i th order difference absolute value is at most $2^{i-2} \left(\frac{n}{\eta}\right)^{i-2}$ for any $i = 3, \dots, n$. The base case is simple: $F[\alpha, \alpha_{t+1}] = \pm 1$. Now, we prove it for $i \geq 3$. We distinguish two cases.

1. If $\alpha_{t+i-1} - \alpha_t \leq \eta/n$, all $\alpha_t, \alpha_{t+1}, \dots, \alpha_{t+i-1}$ must belong to the same segment of L (since all intervals in this range are small, thus merge together). Therefore, we can see that $F[\alpha_t, \alpha_{t+1}] = F[\alpha_{t+1}, \alpha_{t+2}] = \dots = F[\alpha_{t+1}, \alpha_{t+i-1}]$, from which we can see any 3rd order difference (hence the 4th, 5th, up to the i th) in this interval is zero.
2. Suppose $\alpha_{t+i-1} - \alpha_t > \eta/n$. By the induction hypothesis, we have that

$$\begin{aligned} |F[\alpha_t, \dots, \alpha_{t+i-1}]| &= \left| \frac{F[\alpha_t, \dots, \alpha_{t+i-2}] - F[\alpha_{t+1}, \dots, \alpha_{t+i-1}]}{\alpha_{t+i-1} - \alpha_t} \right| \\ &\leq 2 \cdot 2^{i-3} \left(\frac{n}{\eta}\right)^{i-3} \cdot \frac{n}{\eta} \leq 2^{i-2} \left(\frac{n}{\eta}\right)^{i-2}. \end{aligned}$$

Also since $\alpha_j \in [-1, 1]$, the absolute value of the coefficient of x^t in $\prod_{j=1}^i (x - \alpha_j)$ is less than 2^n . So, finally we have $c_i \leq n2^{2n-2} \binom{n}{i}^{n-2}$. \blacksquare

Proof [Proof of Lemma 11] For every polynomial $p(x) \in \mathcal{P}_\alpha$. Let $g_i = (p(\alpha_{i+1}) - p(\alpha_i))/(\alpha_{i+1} - \alpha_i)$ for $i = [n-1]$. So $-1 \leq g_i \leq 1$, implying that the point $\mathbf{g} = (g_1, \dots, g_{n-1})$ is in the $n-1$ dimensional hypercube. Hence, by Carathéodory theorem, there are $\lambda_1, \dots, \lambda_n \geq 0, \lambda_1 + \dots + \lambda_n = 1$, such that $\mathbf{g} = \sum_{i=1}^n \lambda_i \mathbf{q}_i$ where $\mathbf{q}_1, \dots, \mathbf{q}_n$ are n vertices of the hypercube, i.e., $\mathbf{q}_i \in \{-1, +1\}^{n-1}$. For each \mathbf{q}_i , we define $p_i(x)$ be the polynomial with degree at most $n-1$ and

$$p_i(\alpha_1) = 0, p_i(\alpha_m) = \sum_{j=1}^{m-1} q_{i,j}(\alpha_{j+1} - \alpha_j).$$

where $q_{i,j}$ is the j th coordinate of \mathbf{q}_i . We can see $p_i \in \mathcal{F}_\alpha$. It is easy to verify that $p(x) = \sum_{i=1}^n \lambda_i p_i(x)$, since both the LHS and RHS take the same values at $\{\alpha_1, \dots, \alpha_n\}$. For $\eta > 0$, we define $p_\eta(x)$ as $p_\eta = \sum_{i=1}^n \lambda_i T_\eta(p_i)$. We know $|T_\eta(p_i)(\alpha_j) - p_i(\alpha_j)| \leq 2\eta$ for all $i \in [n]$ and $j \in [n]$ and the height of each $T_\eta(p_i)$ is at most $n2^{2n-2} \binom{n}{\eta}^{n-2}$. So, $p_\eta(x)$ satisfies the properties of the lemma. \blacksquare

B.2. Moment-Transportation Inequality over Complex Numbers

In this subsection, we extend Lemma 8 to complex numbers. We allow the mixture components to be complex numbers, and the mixture weight w_i can also be complex numbers. Recall in the complex domain, the definition of transportation distance is extended according to (5).

We first see that n complex numbers in \mathbb{C} can be clustered with the following guarantee.

Proposition 15 For $\alpha_1, \alpha_2, \dots, \alpha_n$ where $\alpha_i \in \mathbb{C}$ and any constant $\eta > 0$, we can partition these n points into clusters such that:

1. $|\alpha_i - \alpha_j| < \eta$ if α_i and α_j are in the same cluster.
2. $|\alpha_i - \alpha_j| > \eta/n$ if α_i and α_j are in different clusters.

Let $\text{Supp} = \text{Supp}(\vartheta) \cup \text{Supp}(\vartheta')$ and let $n = |\text{Supp}|$ ($n \leq 2k$). Arrange the points in Supp as $\alpha = (\alpha_1, \dots, \alpha_n)$ such that each cluster lies in a continuous segment. In other words, if α_i and α_j lie in the same cluster for indexes $i < j$, then $\alpha_{i'}$ and $\alpha_{j'}$ also lie in the same cluster for all $i \leq i' < j' \leq j$; if α_i and α_j lie in the different cluster for indexes $i < j$, then $\alpha_{i'}$ and $\alpha_{j'}$ also lie in the different cluster for all $i' \leq i < j \leq j'$.

Definition 16 Suppose $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n) \in \mathbb{B}_{\mathbb{C}}^n$. Let \mathcal{P}_α be the set of polynomials of degree at most $n-1$ and Lipschitz over α , i.e.,

$$\mathcal{P}_\alpha = \{P \mid \deg P \leq n-1, P(\alpha_1) = 0; |P(\alpha_i) - P(\alpha_j)| \leq |\alpha_i - \alpha_j| \forall i, j \in [n]\}.$$

From the discussion in Appendix A, we have the following proposition on the generalized transportation distance.

Proposition 17 $\text{Tran}(\vartheta, \vartheta') \leq 2 \sup \{ \int f d(\vartheta - \vartheta') : f \in 1\text{-Lip}_{\mathbb{C}} \} = 2 \sup \{ \int p d(\vartheta - \vartheta') : p \in \mathcal{P}_\alpha \}.$

Similar to the real case, we only need to focus on the extreme points of \mathcal{P}_α .

Definition 18 Let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n) \in \mathbb{B}_\mathbb{C}^n$. Let

$$\mathcal{F}_\alpha = \left\{ p \mid \deg p \leq n-1, p(\alpha_1) = 0; \frac{p(\alpha_{i+1}) - p(\alpha_i)}{\alpha_{i+1} - \alpha_i} \in \{\pm 1, \pm i\} \forall i \in [n-1] \right\}.$$

It is easy to see that $|\mathcal{F}_\alpha| = 4^{n-1}$.

Again, we modify each polynomial $p \in \mathcal{F}_\alpha$ slightly. We define a degree- $(n-1)$ polynomial $T_\eta(p) : \mathbb{C} \rightarrow \mathbb{C}$ by assigning values to all α_i for $i \in [n]$ as follows:

1. If α_i is one of the first two points in the cluster corresponding to α_i , assign $T_\eta(p)(\alpha_i) = p(\alpha_i)$.
2. Otherwise, denote α_j, α_{j+1} as the first two points of the cluster, we would assign $T_\eta(p)(\alpha_i)$ to be the linear interpolation of $T_\eta(p)(\alpha_j)$ and $T_\eta(p)(\alpha_{j+1})$. In concrete,

$$T_\eta(p)(\alpha_i) = \frac{p(\alpha_{j+1}) - p(\alpha_j)}{\alpha_{j+1} - \alpha_j} (\alpha_i - \alpha_j) + p(\alpha_j).$$

Since we have fixed n points in $n-1$ -degree polynomial $T_\eta(p)$, we can see that $T_\eta(p)$ is uniquely determined. The following lemma shows $T_\eta(p)$ is close to p over the points in α .

Lemma 19 $|T_\eta(p)(\alpha_i) - p(\alpha_i)| \leq 2\eta$ for each $\alpha_i \in \alpha$.

Proof We only need to prove the case that α_i is not the first two points in the cluster.

$$\begin{aligned} |T_\eta(p)(\alpha_i) - p(\alpha_i)| &\leq \left| \frac{p(\alpha_{j+1}) - p(\alpha_j)}{\alpha_{j+1} - \alpha_j} (\alpha_i - \alpha_j) + p(\alpha_j) - p(\alpha_i) \right| \\ &\leq \left| \frac{p(\alpha_{j+1}) - p(\alpha_j)}{\alpha_{j+1} - \alpha_j} (\alpha_i - \alpha_j) \right| + |p(\alpha_j) - p(\alpha_i)| \\ &\leq 2|\alpha_i - \alpha_j| \leq 2\eta, \end{aligned}$$

where the third inequality holds since $|p(\alpha_{j+1}) - p(\alpha_j)| = |\alpha_{j+1} - \alpha_j|$ and the last inequality holds because α_i and α_j are in the same cluster. \blacksquare

Lemma 20 Let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n) \in \mathbb{B}_\mathbb{C}$. Suppose $T_\eta(p) = \sum_{i=0}^{n-1} c_i x^i$. We have $|c_t| \leq n 2^{2n-2} \left(\frac{n}{\eta}\right)^{n-2}$.

Proof Again, we would use Newton's interpolation polynomials. Let $F[x_1, \dots, x_i]$ be the i th order difference of $T_\eta(p)$. By definition of $T_\eta(p)$, we know that every 2nd order difference of $T_\eta(p)$ is in $\{\pm 1, \pm i\}$. Now, we show inductively that the absolute value of i th order difference absolute value is at most $(2n/\eta)^{i-2}$ for any $i = 3, \dots, n$. We distinguish two cases:

1. If $|\alpha_{t+i} - \alpha_t| < \eta/n$, then all $\alpha_t, \alpha_{t+1}, \dots, \alpha_{t+i}$ lie in the same cluster (according to the assigned order). Therefore, we can see that $F[\alpha_t, \alpha_{t+1}] = F[\alpha_{t+1}, \alpha_{t+2}] = \dots = F[\alpha_{t+i-1}, \alpha_{t+i}]$, from which any 3rd order difference (hence the 4th, 5th, up to the i th) in this interval is zero.

2. Otherwise, $|\alpha_{t+i} - \alpha_t| > \eta/n$. By the induction hypothesis, we have that

$$|F[\alpha_t, \dots, \alpha_{t+i}]| = \left| \frac{F[\alpha_{t+1}, \dots, \alpha_{t+i}] - F[\alpha_t, \dots, \alpha_{t+i-1}]}{\alpha_{t+i} - \alpha_t} \right| \leq 2 \cdot (2n/\eta)^{i-3} \cdot (n/\eta) \leq (2n/\eta)^{i-2}.$$

Also since $|\alpha_j| \leq 1$, the absolute value of the coefficient of x^t in $\prod_{j=1}^i (x - \alpha_j)$ is less than 2^n . So, finally we have $c_t \leq n2^{2n-2} \left(\frac{n}{\eta}\right)^{n-2}$. \blacksquare

Lemma 21 *For any polynomial $p \in \mathcal{P}_\alpha$, there is a polynomial $p_\eta(x)$ such that the following properties hold:*

1. $|p_\eta(\alpha_i) - p(\alpha_i)| \leq 2\eta$, for any $i \in [n]$.
2. The height of g_η is at most $n2^{2n-2} \left(\frac{n}{\eta}\right)^{n-2}$.

Proof For every polynomial $p(x) \in \mathcal{P}_\alpha$, let $g = \frac{p(\alpha_{i+1}) - p(\alpha_i)}{\alpha_{i+1} - \alpha_i}$ for $i = [n-1]$. So $|g_i| \leq 1$, that means the point $\mathbf{p} = (g_1, \dots, g_{n-1})$ is in the $n-1$ dimensional complex hypercube (which has 4^{n-1} vertices). Then, we apply exactly the same argument as in the proof of Lemma 11. \blacksquare

Theorem 22 *For any two mixtures with k components $\vartheta, \vartheta' \in \text{Spike}(\mathbb{B}_\mathbb{C}, \Sigma_{k-1}^\mathbb{C})$, it holds that*

$$\text{Tran}(\vartheta, \vartheta') \leq O(k \text{Mom}_K(\vartheta, \vartheta')^{\frac{1}{2k-1}}).$$

Proof Fix any constant $\eta > 0$. Let $\xi = \text{Mom}_K(\vartheta, \vartheta')$

$$\begin{aligned} \text{Tran}(\vartheta, \vartheta') &\leq 2 \sup_{f \in 1\text{-Lip}_\mathbb{C}} \int f d(\vartheta - \vartheta') = 2 \sup_{p \in \mathcal{P}_\alpha} \int p d(\vartheta - \vartheta') \\ &\leq 2 \left(4\eta + \sup_{p \in \mathcal{P}_\alpha} \int p_\eta d(\vartheta - \vartheta') \right) && \text{(Lemma 21)} \\ &\leq 2 \left(4\eta + \sup_{p \in \mathcal{P}_\alpha} \sum_{p=0}^{2k-1} n2^{2n-2} \left(\frac{n}{\eta}\right)^{n-2} \text{Mom}_K(\vartheta, \vartheta') \right) && \text{(Lemma 21)} \\ &\leq 2 \left(4\eta + 2^{4k} (2k)^{2k} \frac{\xi}{\eta^{2k-2}} \right). && (n \leq 2k) \end{aligned}$$

By choosing $\eta = 16k\xi^{\frac{1}{2k-1}}$, we conclude that $\text{Tran}(\vartheta, \vartheta') \leq O(k\xi^{\frac{1}{2k-1}})$. \blacksquare

B.3. Moment-Transportation Inequality in Higher Dimensions

We generalize our proof for one dimension to higher dimensions. Lemma 45 shows there exists a vector \mathbf{r} such that the distance between spikes are still lower bounded (up to factors depending on k and d) after projection. This enables us to extend Lemma 8 to high dimension.

Theorem 23 *Let ϑ, ϑ' be two k -spike mixtures in $\text{Spike}(\Delta_{d-1}, \Delta_{k-1})$, and $K = 2k - 1$. Then, the following inequality holds*

$$\text{Tran}(\vartheta, \vartheta') \leq O(k^3 d \text{Mom}_K(\vartheta, \vartheta')^{\frac{1}{2k-1}}).$$

Proof Apply the argument in Lemma 45 to $\text{Supp} = \text{Supp}(\vartheta) \cup \text{Supp}(\vartheta')$. There always exists vector $\mathbf{r} \in \mathbb{S}_{d-1}$ and constant c such that $|\mathbf{r}^\top(\boldsymbol{\alpha}_i - \boldsymbol{\alpha}_j)| \geq \frac{c}{k^2 d} \|\boldsymbol{\alpha}_i - \boldsymbol{\alpha}_j\|_1$ for all $\boldsymbol{\alpha}_i, \boldsymbol{\alpha}_j \in \text{Supp}$. Note that $|\mathbf{r}^\top \boldsymbol{\alpha}_i| \leq \|\mathbf{r}\|_\infty \|\boldsymbol{\alpha}_i\|_1 \leq 1$, both $\Pi_{\mathbf{r}}(\vartheta)$ and $\Pi_{\mathbf{r}}(\vartheta')$ are in $\text{Spike}(\Delta_1, \Delta_{k-1})$. In this case,

$$\begin{aligned} \text{Tran}(\vartheta, \vartheta') &= \inf \left\{ \int \|\mathbf{x} - \mathbf{y}\|_1 d\mu(\mathbf{x}, \mathbf{y}) : \mu \in M(\vartheta, \vartheta') \right\} && \text{(Definition)} \\ &\leq \inf \left\{ \int \frac{k^2 d}{c} |\mathbf{r}^\top(\mathbf{x} - \mathbf{y})| d\mu(\mathbf{x}, \mathbf{y}) : \mu \in M(\vartheta, \vartheta') \right\} && \text{(Lemma 45)} \\ &= \frac{k^2 d}{c} \text{Tran}(\Pi_{\mathbf{r}}(\vartheta), \Pi_{\mathbf{r}}(\vartheta')) \\ &\leq O(k^3 d \text{Mom}_K(\vartheta, \vartheta')^{\frac{1}{2k-1}}). && \text{(Lemma 8)} \end{aligned}$$

■

Appendix C. Missing Details from Section 4

C.1. Implementation Details of Algorithm 1

In this subsection, we show how to implement Algorithm 1 in $O(k^2)$ time. We first perform a ridge regression to obtain $\hat{\mathbf{c}}$ in Line 1. The explicit solution of this ridge regression is $\hat{\mathbf{c}} = (A_{M'}^\top A_{M'} + \xi^2 I)^{-1} A_{M'}^\top b_{M'}$. Since $A_{M'}$ is a Hankel matrix, $\hat{\mathbf{c}} = (A_{M'} - i\xi I)^{-1} (A_{M'} + i\xi I)^{-1} A_{M'}^\top b_{M'}$ holds. Note that $\mathbf{x} = (A + \lambda I)^{-1} \mathbf{b}$ is a single step of the inverse power method, which can be computed in $O(k^2)$ time when A is a Hankel matrix (see Xu and Qiao (2008)). Hence, Line 1 can be implemented in $O(k^2)$ arithmetic operations.

In Line 2, we solve the roots of polynomial $\sum_{i=0}^{k-1} \hat{c}_i x^i + x^k$. We can use the algorithm in Neff and Reif (1996) to find the solution with ξ -additive noise in $O(k^{1+o(1)} \cdot \log \log(1/\xi))$ arithmetic operations.

Line 5 is a linear regression defined by Vandemonde matrix V_{α} . We can use the recent algorithm developed in Meyer et al. (2022) to find a constant factor multiplicative approximation, which uses $O(k^{1+o(1)})$ arithmetic operations. Note that a constant factor approximation suffices for our purpose.

In Line 8, we find the k -spike distribution in $\text{Spike}(\Delta_1, \Delta_{k-1})$ closest to $\tilde{\vartheta}$. To achieve $O(k^2)$ running time, we limit the spike distribution to have support $\tilde{\boldsymbol{\alpha}}$. In this case, the optimization problem is equivalent to finding a transportation plan from $(\tilde{\boldsymbol{\alpha}}, \tilde{\mathbf{w}}^-)$ to $(\tilde{\boldsymbol{\alpha}}, \tilde{\mathbf{w}}^+)$ where $\tilde{\mathbf{w}}^-$ and $\tilde{\mathbf{w}}^+$ are the negative components and positive components of $\tilde{\mathbf{w}}$ respectively. In concrete, $\tilde{w}_i^- =$

$\max\{0, -\tilde{w}_i\}$ and $\tilde{w}_i^+ = \max\{0, \tilde{w}_i\}$. Note that the points are in one dimension. Using the classical algorithm in [Aggarwal et al. \(1992\)](#), we can solve this transportation problem in $O(k^{1+o(1)})$ arithmetic operations.

C.2. Proof of Theorem 3

Now, we start to bound the reconstruction error of the algorithm. The following lemma is well known (see e.g., [Chihara \(2011\)](#); [Gordon et al. \(2020\)](#)). We provide a proof for completeness.

Lemma 24 *Let $\boldsymbol{\vartheta} = (\boldsymbol{\alpha}, \boldsymbol{w}) \in \text{Spike}(\Delta_1, \Delta_{k-1})$ where $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_k]^\top$ and $\boldsymbol{w} = [w_1, \dots, w_k]^\top$. Let $\boldsymbol{c} = [c_0, \dots, c_{k-1}]^\top \in \mathbb{R}^k$ such that $\prod_{i=1}^k (x - \alpha_i) = \sum_{i=0}^{k-1} c_i x^i + x^k$. Then, $\sum_{j=0}^{k-1} M_{i+j}(\boldsymbol{\vartheta}) c_j + M_{i+k}(\boldsymbol{\vartheta}) = 0$ for any $i \geq 0$. The equation can also be written as $A_{M(\boldsymbol{\vartheta})} \boldsymbol{c} + b_{M(\boldsymbol{\vartheta})} = \mathbf{0}$ in the matrix form.*

Proof By direct calculation,

$$\begin{aligned} \sum_{j=0}^{k-1} M_{i+j}(\boldsymbol{\vartheta}) c_j + M_{i+k}(\boldsymbol{\vartheta}) &= \sum_{j=0}^{k-1} \sum_{t=1}^k w_t \alpha_t^{i+j} c_j + \sum_{t=1}^k w_t \alpha_t^{i+k} = \sum_{t=1}^k w_t \alpha_t^i \left(\sum_{j=0}^{k-1} c_j \alpha_t^j + \alpha_t^k \right) \\ &= \sum_{t=1}^k w_t \alpha_t^i \prod_{j=1}^k (\alpha_t - \alpha_j) = 0. \end{aligned}$$

The i th row of $A_{M(\boldsymbol{\vartheta})} \boldsymbol{c} + b_{M(\boldsymbol{\vartheta})}$ is $\sum_{j=0}^{k-1} M_{i+j}(\boldsymbol{\vartheta}) c_j + M_{i+k}(\boldsymbol{\vartheta})$, hence the matrix form. \blacksquare

Next, the following lemma shows that the intermediate result $\hat{\boldsymbol{c}}$ a good estimation for the solution of $A_{M(\boldsymbol{\vartheta})} \boldsymbol{x} + b_{M(\boldsymbol{\vartheta})} = \mathbf{0}$ with a small norm:

Lemma 25 *Let $\boldsymbol{c} = [c_0, \dots, c_{k-1}]^\top \in \mathbb{R}^k$ be the sequence of number in which $\prod_{i=1}^k (x - \alpha_i) = \sum_{i=0}^{k-1} c_i x^i + x^k$. Suppose $\|M' - M(\boldsymbol{\vartheta})\|_\infty \leq \xi$. Let $\hat{\boldsymbol{c}} = [\hat{c}_0, \dots, \hat{c}_{k-1}]^\top \in \mathbb{R}^k$ be the intermediate result (Line 1) in Algorithm 1. Then, $\|\boldsymbol{c}\|_1 \leq 2^k$, $\|\hat{\boldsymbol{c}}\|_1 \leq 2^{O(k)}$ and $\|A_{M(\boldsymbol{\vartheta})} \hat{\boldsymbol{c}} + b_{M(\boldsymbol{\vartheta})}\|_\infty \leq 2^{O(k)} \cdot \xi$.*

Proof From Vieta's formulas (see e.g., [Barbeau \(2003\)](#)), we have $c_i = \sum_{S \in \binom{[k]}{k-i}} \prod_{j \in S} (-\alpha_j)$. Thus,

$$\|\boldsymbol{c}\|_1 = \sum_{i=0}^{k-1} |c_i| \leq \sum_{S \in 2^{[k]}} \prod_{j \in S} |\alpha_j| = \prod_{i=1}^k (1 + |\alpha_i|) \leq 2^k.$$

where the last inequality holds since $|\alpha_i| \leq 1$ for all i according to the definition of $\boldsymbol{\vartheta}$.

From Lemma 24, we can see that $\|A_{M(\boldsymbol{\vartheta})} \boldsymbol{c} + b_{M(\boldsymbol{\vartheta})}\|_\infty = 0$. Therefore, we can further get

$$\begin{aligned} \|A_{M'} \boldsymbol{c} + b_{M'}\|_\infty &\leq \|A_{M(\boldsymbol{\vartheta})} \boldsymbol{c} + b_{M(\boldsymbol{\vartheta})}\|_\infty + \|A_{M'} \boldsymbol{c} - A_{M(\boldsymbol{\vartheta})} \boldsymbol{c}\|_\infty + \|b_{M'} - b_{M(\boldsymbol{\vartheta})}\|_\infty \\ &\leq \|A_{M(\boldsymbol{\vartheta})} \boldsymbol{c} + b_{M(\boldsymbol{\vartheta})}\|_\infty + \|(A_{M'} - A_{M(\boldsymbol{\vartheta})}) \boldsymbol{c}\|_\infty + \|b_{M'} - b_{M(\boldsymbol{\vartheta})}\|_\infty \\ &\leq \|A_{M(\boldsymbol{\vartheta})} \boldsymbol{c} + b_{M(\boldsymbol{\vartheta})}\|_\infty + \|A_{M'} - A_{M(\boldsymbol{\vartheta})}\|_\infty \|\boldsymbol{c}\|_1 + \|b_{M'} - b_{M(\boldsymbol{\vartheta})}\|_\infty \\ &\leq \|A_{M(\boldsymbol{\vartheta})} \boldsymbol{c} + b_{M(\boldsymbol{\vartheta})}\|_\infty + \|M' - M(\boldsymbol{\vartheta})\|_\infty (\|\boldsymbol{c}\|_1 + 1) \\ &\leq 2^{O(k)} \cdot \xi, \end{aligned}$$

where the fourth inequality holds since $(A_{M'} - A_{M(\vartheta)})_{i,j} = (M' - M(\vartheta))_{i+j}$ and $(b_{M'} - b_{M(\vartheta)})_i = (M' - M(\vartheta))_{i+k}$. From the definition of $\hat{\mathbf{c}}$, we can see that

$$\begin{aligned} \|A_{M'}\hat{\mathbf{c}} + b_{M'}\|_2^2 + \xi^2\|\hat{\mathbf{c}}\|_2^2 &\leq \|A_{M'}\mathbf{c} + b_{M'}\|_2^2 + \xi^2\|\mathbf{c}\|_2^2 \\ &\leq k\|A_{M'}\mathbf{c} + b_{M'}\|_\infty^2 + \xi^2\|\mathbf{c}\|_1^2 \\ &\leq 2^{O(k)} \cdot \xi^2. \end{aligned}$$

The second inequality holds since $\|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_1$ and $\|\mathbf{x}\|_2 \leq \sqrt{k}\|\mathbf{x}\|_\infty$ holds for any vector $\mathbf{x} \in \mathbb{R}^k$.

Now, we can directly see that

$$\begin{aligned} \|A_{M'}\hat{\mathbf{c}} + b_{M'}\|_\infty &\leq \|A_{M'}\hat{\mathbf{c}} + b_{M'}\|_2 \leq 2^{O(k)} \cdot \xi, \\ \|\hat{\mathbf{c}}\|_1 &\leq \sqrt{k}\|\hat{\mathbf{c}}\|_2 \leq 2^{O(k)}. \end{aligned}$$

Finally, we can bound $\|A_{M(\vartheta)}\hat{\mathbf{c}} + b_{M(\vartheta)}\|_\infty$ as follows:

$$\begin{aligned} \|A_{M(\vartheta)}\hat{\mathbf{c}} + b_{M(\vartheta)}\|_\infty &\leq \|A_{M'}\hat{\mathbf{c}} + b_{M'}\|_\infty + \|A_{M'}\hat{\mathbf{c}} - A_{M(\vartheta)}\hat{\mathbf{c}}\|_\infty + \|b_{M'} - b_{M(\vartheta)}\|_\infty \\ &\leq \|A_{M'}\hat{\mathbf{c}} + b_{M'}\|_\infty + \|A_{M'} - A_{M(\vartheta)}\|_\infty\|\hat{\mathbf{c}}\|_1 + \|b_{M'} - b_{M(\vartheta)}\|_\infty \\ &\leq \|A_{M'}\hat{\mathbf{c}} + b_{M'}\|_\infty + \|M' - M(\vartheta)\|_\infty(\|\hat{\mathbf{c}}\|_1 + 1) \\ &\leq 2^{O(k)} \cdot \xi. \end{aligned}$$

■

Using this result, we are able to show that $\tilde{\alpha}$ in Line 4 is a good estimation for the ground truth spikes, α . In particular, the following lemma shows that every ground truth spike $a_i \in \alpha$ has a close spike $\tilde{\alpha}_j \in \tilde{\alpha}$ where the closeness depends on the weight of the spike w_i . We note that this property tolerates permutation and small spikes weights, enabling us to analyze mixtures without separations.

Lemma 26 *Let $\hat{\alpha} = [\hat{\alpha}_1, \dots, \hat{\alpha}_k]^\top \in \mathbb{C}^k$ be the intermediate result (Line 2) in Algorithm 1. Then, the following inequality holds:*

$$\sum_{i=1}^k w_i \prod_{j=1}^k |\alpha_i - \hat{\alpha}_j|^2 \leq 2^{O(k)} \cdot \xi.$$

Proof Since $\hat{\mathbf{c}}$ and α are real, from the definition of $\hat{\mathbf{c}}$,

$$\sum_{i=1}^k w_i \prod_{j=1}^k |\alpha_i - \hat{\alpha}_j|^2 = \sum_{i=1}^k w_i \left| \prod_{j=1}^k (\alpha_i - \hat{\alpha}_j) \right|^2 = \sum_{i=1}^k w_i \left| \sum_{j=0}^{k-1} \hat{c}_j \alpha_i^j + \alpha_i^k \right|^2 = \sum_{i=1}^k w_i \left(\sum_{j=0}^{k-1} \hat{c}_j \alpha_i^j + \alpha_i^k \right)^2.$$

By expanding the RHS, we can reach that

$$\begin{aligned}
\sum_{i=1}^k w_i \prod_{j=1}^k |\alpha_i - \hat{\alpha}_j|^2 &= \sum_{i=1}^k w_i \left(\sum_{p=0}^{k-1} \sum_{q=0}^{k-1} \hat{c}_p \hat{c}_q \alpha_i^{p+q} + 2 \sum_{p=0}^{k-1} \hat{c}_p \alpha_i^{p+k} + \alpha_i^{2k} \right) \\
&= \sum_{p=0}^{k-1} \sum_{q=0}^{k-1} \hat{c}_p \hat{c}_q \sum_{i=1}^k w_i \alpha_i^{p+q} + 2 \sum_{p=0}^{k-1} \hat{c}_p \sum_{i=1}^k w_i \alpha_i^{p+k} + \sum_{i=1}^k w_i \alpha_i^{2k} \\
&= \sum_{p=0}^{k-1} \hat{c}_p \left(\sum_{q=0}^{k-1} M_{p+q}(\boldsymbol{\vartheta}) \hat{c}_q + M_{p+k}(\boldsymbol{\vartheta}) \right) + \sum_{p=0}^{k-1} \hat{c}_p M_{p+k}(\boldsymbol{\vartheta}) + M_{2k}(\boldsymbol{\vartheta}) \\
&= \sum_{p=0}^{k-1} \hat{c}_p \left(A_{M(\boldsymbol{\vartheta})} \hat{c} + b_{M(\boldsymbol{\vartheta})} \right)_p + \sum_{p=0}^{k-1} \hat{c}_p M_{p+k}(\boldsymbol{\vartheta}) + M_{2k}(\boldsymbol{\vartheta}),
\end{aligned}$$

where the last equality holds from the definition of matrix $A_{M(\boldsymbol{\vartheta})}$ and vector $b_{M(\boldsymbol{\vartheta})}$. Moreover,

$$\begin{aligned}
\sum_{p=0}^{k-1} \hat{c}_p M_{p+k}(\boldsymbol{\vartheta}) + M_{2k}(\boldsymbol{\vartheta}) &= \sum_{p=0}^{k-1} \hat{c}_p \left(- \sum_{q=0}^{k-1} M_{p+q}(\boldsymbol{\vartheta}) c_q \right) + \left(- \sum_{q=0}^{k-1} M_{k+q}(\boldsymbol{\vartheta}) c_q \right) \\
&= - \sum_{q=0}^{k-1} c_q \left(\sum_{p=0}^{k-1} M_{p+q}(\boldsymbol{\vartheta}) \hat{c}_p + M_{k+q}(\boldsymbol{\vartheta}) \right) \\
&= - \sum_{q=0}^{k-1} c_q \left(A_{M(\boldsymbol{\vartheta})} \hat{c} + b_{M(\boldsymbol{\vartheta})} \right)_q,
\end{aligned}$$

where the first equality dues to Lemma 24 in which $M_{p+k}(\boldsymbol{\vartheta}) = - \sum_{q=0}^{k-1} M_{p+q}(\boldsymbol{\vartheta}) c_q$. Combining the above results, we finally get

$$\begin{aligned}
\sum_{i=1}^k w_i \prod_{j=1}^k |\alpha_i - \hat{\alpha}_j|^2 &= \sum_{i=0}^{k-1} (\hat{c}_i - c_i) \left(A_{M(\boldsymbol{\vartheta})} \hat{c} + b_{M(\boldsymbol{\vartheta})} \right)_i \\
&\leq (\|\hat{c}\|_1 + \|c\|_1) \|A_{M(\boldsymbol{\vartheta})} \hat{c} + b_{M(\boldsymbol{\vartheta})}\|_\infty \\
&\leq (2^k + 2^{O(k)}) \cdot 2^{O(k)} \cdot \xi \quad (\text{Lemma 25}) \\
&\leq 2^{O(k)} \cdot \xi.
\end{aligned}$$

■

We provide the following simple inequality that allows us to upper bound the impact of the injected noise.

Lemma 27 *Let $b > 0$ be some constant. For $a_1, \dots, a_k \in \mathbb{R}$ and $a'_1, \dots, a'_k \in \mathbb{R}$ such that $a_i \in [0, b]$ and $|a'_i - a_i| \leq \xi$ holds for all i . In case that $\xi \leq n^{-1}$, then,*

$$\prod_{i=1}^k a'_i \leq \prod_{i=1}^k a_i + O(b^k \cdot k\xi).$$

Proof The proof is similar to the proof of Lemma 1 in (Li and Deshpande, 2010).

$$\begin{aligned}
 \prod_{i=1}^k a'_i - \prod_{i=1}^k a_i &= \sum_{S \subseteq [k], S \neq \emptyset} \prod_{i \in S} a_i \prod_{i \notin S} (a'_i - a_i) \\
 &\leq \sum_{t=1}^k \sum_{|S|=t} \prod_{i \in S} a_i \prod_{i \notin S} (a'_i - a_i) \\
 &\leq b^k \sum_{t=1}^k \binom{k}{t} \xi^k \leq b^k \sum_{t=1}^k \frac{(k\xi)^k}{t!} \\
 &\leq b^k (e^{k\xi} - 1) = O(b^k \cdot k\xi).
 \end{aligned}$$

The third inequality holds because $\binom{k}{t} \leq \frac{k^t}{t!}$. The last inequality holds since $e^x = \sum_{i>0} \frac{x^i}{i!}$ and the fact that $e^{f(x)} = 1 + O(f(x))$ if $f(x) = O(1)$. \blacksquare

With this lemma, we are able to show that the projection (Line 3) and the injected noise (Line 4) do not introduce much extra error for our estimation of the positions of ground truth spikes.

Lemma 28 Let $\tilde{\alpha} = [\tilde{\alpha}_1, \dots, \tilde{\alpha}_k]^\top \in \mathbb{R}^k$ be the intermediate result (Line 4) in Algorithm 1. Then,

$$\sum_{i=1}^k w_i \prod_{j=1}^k |\alpha_i - \tilde{\alpha}_j|^2 \leq 2^{O(k)} \cdot \xi.$$

Proof Let $\bar{\alpha} = [\bar{\alpha}_1, \dots, \bar{\alpha}_k]^\top \in \mathbb{R}^k$ be the set of projections (Line 3). Since the projected domain Δ_1 is convex, $|\alpha_i - \bar{\alpha}_j| \leq |\alpha_i - \hat{\alpha}_j|$ always holds. This gives

$$\prod_{j=1}^k |\alpha_i - \bar{\alpha}_j|^2 \leq \prod_{j=1}^k |\alpha_i - \hat{\alpha}_j|^2.$$

Since $\tilde{\alpha}_j$ is noisy $\bar{\alpha}_j$ with additive noise of size less than ξ , we have $|\alpha_i - \tilde{\alpha}_j| \leq |\alpha_i - \bar{\alpha}_j| + \xi$. Applying Lemma 27 by regarding $|\alpha_i - \tilde{\alpha}_j|$ as a_j and $|\alpha_i - \hat{\alpha}_j|$ as a'_j , with the constrain guaranteed according to $|\alpha_i - \tilde{\alpha}_j| \leq |\alpha_i| + |\tilde{\alpha}_j| \leq 2$, we can conclude that

$$\prod_{j=1}^k |\alpha_i - \tilde{\alpha}_j|^2 \leq \prod_{j=1}^k |\alpha_i - \bar{\alpha}_j|^2 + O(2^k \cdot k\xi).$$

Combining two inequalities, we get the desired inequality:

$$\sum_{i=1}^k w_i \prod_{j=1}^k |\alpha_i - \tilde{\alpha}_j|^2 \leq \sum_{i=1}^k w_i \prod_{j=1}^k |\alpha_i - \hat{\alpha}_j|^2 + O(2^k \cdot k\xi) \leq 2^{O(k)} \cdot \xi,$$

where the last inequality is given by Lemma 26. \blacksquare

We now start to bound the error in Line 5 which is a linear regression defined by the Vandermonde matrix. We first introduce the Schur polynomial which has a strong connection to the Vandermonde matrix.

Definition 29 (Schur polynomial) Given a partition $\lambda = (\lambda_1, \dots, \lambda_k)$ where $\lambda_1 \geq \dots \geq \lambda_k \geq 0$, the Schur polynomial $s_\lambda(x_1, x_2, \dots, x_k)$ is defined by the ratio

$$s_\lambda(x_1, x_2, \dots, x_k) = \frac{\begin{vmatrix} x_1^{\lambda_1+k-1} & x_2^{\lambda_1+k-1} & \dots & x_k^{\lambda_1+k-1} \\ x_1^{\lambda_2+k-2} & x_2^{\lambda_2+k-2} & \dots & x_k^{\lambda_2+k-2} \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{\lambda_k} & x_2^{\lambda_k} & \dots & x_k^{\lambda_k} \end{vmatrix}}{\begin{vmatrix} 1 & 1 & \dots & 1 \\ x_1 & x_2 & \dots & x_k \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{k-1} & x_2^{k-1} & \dots & x_k^{k-1} \end{vmatrix}}^{-1}.$$

A Schur polynomial is a symmetric function because the numerator and denominator are both *alternating* (i.e., it changes sign if we swap two variables), and a polynomial since all alternating polynomials are divisible by the Vandermonde determinant $\prod_{i < j} (x_i - x_j)$.⁴ The following classical result (see, e.g., Stanley (2011)) shows an alternative way to calculate Schur polynomial. It plays a crucial role in bounding the error of a linear system defined by the Vandermonde matrix (Lemma 32).

Theorem 30 (e.g. Stanley (2011)) Let $SSYT(\lambda)$ be the set consists all semistandard Young tableau of shape $\lambda = (\lambda_1, \dots, \lambda_k)$. Then, Schur polynomial $s_\lambda(x_1, x_2, \dots, x_k)$ can also be computed as follows:

$$s_\lambda(x_1, x_2, \dots, x_k) = \sum_{T \in SSYT(\lambda)} \prod_{i=1}^k x_i^{T_i}.$$

where T_i counts the occurrences of the number i in T .

In a semistandard Young tableau, we allow the same number to appear more than once (or not at all), and we require the entries weakly increase along each row and strictly increase down each column. As an example, for $\lambda = (2, 1, 0)$, the list of semistandard Young tableau are:

1	1	1	1	1	2	1	2	1	3	1	3	2	2	2	3
2		3		2		3		2		3		3		3	

The corresponding Schur polynomial is

$$s_{(2,1,0)}(x_1, x_2, x_3) = x_1^2 x_2 + x_1^2 x_3 + x_1 x_2^2 + x_1 x_2 x_3 + x_1 x_3 x_2 + x_1 x_3^2 + x_2^2 x_3 + x_2 x_3^2.$$

Considering the special case where $\lambda = (j - k + 1, 1, 1, \dots, 1)$, this theorem implies the following result. We also present a self-contained proof of the lemma for completeness.

Lemma 31 For all $j \geq k$ and distinct numbers $a_1, \dots, a_k \in \mathbb{R}$ (or $a_1, \dots, a_k \in \mathbb{C}$),

$$\frac{\begin{vmatrix} a_1^j & a_2^j & \dots & a_k^j \\ a_1^1 & a_2^1 & \dots & a_k^1 \\ \vdots & \vdots & \ddots & \vdots \\ a_1^{k-1} & a_2^{k-1} & \dots & a_k^{k-1} \end{vmatrix}}{\begin{vmatrix} 1 & 1 & \dots & 1 \\ a_1 & a_2 & \dots & a_k \\ \vdots & \vdots & \ddots & \vdots \\ a_1^{k-1} & a_2^{k-1} & \dots & a_k^{k-1} \end{vmatrix}}^{-1} = \prod_{i=1}^k a_i \cdot \sum_{s \in (k)^{j-k}} \prod_{i=1}^k a_i^{s_i}$$

where $(a)^b \subset \{0, \dots, b\}^a$ denotes the set that contains all vector s in which $\sum_{i=1}^a s_i = b$.

4. It is easy to see that if a polynomial $p(x_1, \dots, x_n)$ is alternating, it must contain $(x_i - x_j)$ as a factor .

Proof We provide a self-contained proof here for completeness. According to the property of Vandermonde determinant, we have

$$\begin{vmatrix} 1 & 1 & \cdots & 1 \\ a_1^1 & a_2^1 & \cdots & a_k^1 \\ \vdots & \vdots & \ddots & \vdots \\ a_1^{k-1} & a_2^{k-1} & \cdots & a_k^{k-1} \end{vmatrix} = (-1)^{k(k-1)/2} \prod_{p < q} (a_p - a_q).$$

By expansion along first row, we have

$$\begin{aligned} \begin{vmatrix} a_1^j & a_2^j & \cdots & a_k^j \\ a_1^1 & a_2^1 & \cdots & a_k^1 \\ \vdots & \vdots & \ddots & \vdots \\ a_1^{k-1} & a_2^{k-1} & \cdots & a_k^{k-1} \end{vmatrix} &= \sum_{i=1}^k (-1)^{i+1} a_i^j \begin{vmatrix} a_1^1 & \cdots & a_{i-1}^1 & a_{i+1}^1 & \cdots & a_k^1 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ a_1^{k-1} & \cdots & a_{i-1}^{k-1} & a_{i+1}^{k-1} & \cdots & a_k^{k-1} \end{vmatrix} \\ &= \sum_{i=1}^k (-1)^{i+1} a_i^j (-1)^{(k-1)(k-2)/2} \prod_{t \neq i} a_t \prod_{p < q, p \neq i, q \neq i} (a_p - a_q) \\ &= (-1)^{k(k-1)/2} \cdot (-1)^{k+1} \sum_{i=1}^k a_i^j (-1)^{i+1} \prod_{t \neq i} a_t \prod_{p < q, p \neq i, q \neq i} (a_p - a_q) \end{aligned}$$

Thus,

$$\text{LHS}(j) = \sum_{i=1}^k a_i^j \cdot \prod_{t \neq i} \frac{a_t}{a_i - a_t}$$

Let $L(z) = \sum_{j \geq k} \text{LHS}(j) z^j$ and $R(z) = \sum_{j \geq k} \text{RHS}(j) z^j$ be the generating functions of LHS and RHS corresponding to j respectively. Then,

$$\begin{aligned} L(z) &= \sum_{i=1}^k \left(\sum_{j \geq k} a_i^j z^j \right) \prod_{t \neq i} \frac{a_t}{a_t - a_i} = z^k \sum_{i=1}^k \frac{a_i^k}{1 - a_i z} \prod_{t \neq i} \frac{a_t}{a_i - a_t} \\ R(z) &= \prod_{i=1}^k \left(\sum_{j \geq 1} a_i^j z^j \right) = z^k \prod_{t=1}^k \frac{a_t}{1 - a_t z} \end{aligned}$$

Consider the quotient between two functions,

$$\frac{L(z)}{R(z)} = \sum_{i=1}^k a_i^{k-1} \prod_{t \neq i} \frac{1 - a_t z}{a_i - a_t}.$$

For $z = a_p^{-1}$ where $p \in [1, k]$, the product in each additive term with $i \neq p$ equals 0. As a result, we can see that

$$\frac{L(a_p^{-1})}{R(a_p^{-1})} = a_p^{k-1} \prod_{t \neq p} \frac{1 - a_t a_p^{-1}}{a_i - a_t} = 1.$$

Since $L(z)/R(z)$ is polynomial of z of degree $k - 1$, from the uniqueness of polynomial, we have

$$\frac{L(z)}{R(z)} = 1 \Rightarrow \text{LHS} = \text{RHS}$$

which directly proves the statement. ■

Using this result, we can show the reconstructed moment vector $V_{\tilde{\alpha}} \hat{\boldsymbol{w}}$ is close to the ground truth $M(\boldsymbol{\vartheta})$. The next lemma shows the discrete mixture with support $\tilde{\alpha}$ can well approximate a single spike distribution at α in terms of the moments if any of the spikes in $\tilde{\alpha}$ is close to α .

Lemma 32 *Let $M(\alpha)$ be the vector $[\alpha^0, \alpha^1, \dots, \alpha^{2k-1}]^\top$. Suppose F is either $\{\mathbb{R} \text{ or } \mathbb{C}\}$. For $\alpha \in F$ with $|\alpha| \leq 1$ and $\tilde{\alpha} = [\tilde{\alpha}_1, \dots, \tilde{\alpha}_k]^\top \in F^k$ with all $\tilde{\alpha}_s$ distinct and $\|\tilde{\alpha}\|_\infty \leq 1$. We have,*

$$\min_{\boldsymbol{x} \in F^k} \|V_{\tilde{\alpha}} \boldsymbol{x} - M(\alpha)\|_2 \leq 2^{O(k)} \prod_{j=1}^k |\alpha - \tilde{\alpha}_j|.$$

Proof Let $\Delta \tilde{\alpha}_j = \tilde{\alpha}_j - \alpha$ and let $\boldsymbol{x}^* = [x_1^*, \dots, x_k^*]^\top$ be the solution to linear equation

$$\begin{bmatrix} 1 & 1 & \dots & 1 \\ \Delta \tilde{\alpha}_1^1 & \Delta \tilde{\alpha}_2^1 & \dots & \Delta \tilde{\alpha}_k^1 \\ \vdots & \vdots & \ddots & \vdots \\ \Delta \tilde{\alpha}_1^{k-1} & \Delta \tilde{\alpha}_2^{k-1} & \dots & \Delta \tilde{\alpha}_k^{k-1} \end{bmatrix} \begin{bmatrix} x_1^* \\ x_2^* \\ \vdots \\ x_k^* \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

Denote $(V_{\tilde{\alpha}} \boldsymbol{x}^* - M(\alpha))_j$ as the j th element of vector $V_{\tilde{\alpha}} \boldsymbol{x}^* - M(\alpha)$. We will show that $|(V_{\tilde{\alpha}} \boldsymbol{x}^* - M(\alpha))_j| \leq 2^k \cdot 2^{2j} \prod_{t=1}^k |\Delta \tilde{\alpha}_t|$ for any integer $j \geq 0$. It is clear that we have $\sum_{t=1}^k x_t^* \Delta \tilde{\alpha}_t^j = 0$ for any $1 \leq j < k$ according to the construction of \boldsymbol{x}^* . For larger j , we can expand the element by

$$\begin{aligned} (V_{\tilde{\alpha}} \boldsymbol{x}^* - M(\alpha))_j &= \sum_{t=1}^k x_t^* \tilde{\alpha}_t^j - \alpha^j = \sum_{t=1}^k x_t^* \sum_{p=0}^j \binom{j}{p} \Delta \tilde{\alpha}_t^p \alpha^{j-p} - \alpha^j \\ &= \sum_{p=1}^j \binom{j}{p} \alpha^{j-p} \sum_{t=1}^k x_t^* \Delta \tilde{\alpha}_t^p. \end{aligned}$$

where the second equality follows from the definition of $\Delta \tilde{\alpha}_j$. So it is sufficient to bound $\sum_{t=1}^k x_t^* \Delta \tilde{\alpha}_t^j$ for any $j \geq k$. We have that

$$\begin{aligned} \sum_{t=1}^k x_t^* \Delta \tilde{\alpha}_t^j &= \begin{vmatrix} \Delta \tilde{\alpha}_1^j & \Delta \tilde{\alpha}_2^j & \dots & \Delta \tilde{\alpha}_k^j \\ \Delta \tilde{\alpha}_1^1 & \Delta \tilde{\alpha}_2^1 & \dots & \Delta \tilde{\alpha}_k^1 \\ \vdots & \vdots & \ddots & \vdots \\ \Delta \tilde{\alpha}_1^{k-1} & \Delta \tilde{\alpha}_2^{k-1} & \dots & \Delta \tilde{\alpha}_k^{k-1} \end{vmatrix} \cdot \begin{vmatrix} 1 & 1 & \dots & 1 \\ \Delta \tilde{\alpha}_1^1 & \Delta \tilde{\alpha}_2^1 & \dots & \Delta \tilde{\alpha}_k^1 \\ \vdots & \vdots & \ddots & \vdots \\ \Delta \tilde{\alpha}_1^{k-1} & \Delta \tilde{\alpha}_2^{k-1} & \dots & \Delta \tilde{\alpha}_k^{k-1} \end{vmatrix}^{-1} \\ &= \prod_{t=1}^k \Delta \tilde{\alpha}_t \cdot \sum_{s \subseteq (k)j-k} \prod_{t=1}^k \Delta \tilde{\alpha}_t^{s_t}, \end{aligned}$$

where the first equality dues to the relation between inverse matrix and adjoint matrix, and the other equation holds from Lemma 31. Since $|\Delta\tilde{\alpha}_t| \leq |\alpha| + |\tilde{\alpha}_t| \leq 2$, $(k)^{j-k}$ has exactly $\binom{j}{k}$ terms, we can bound the summation by

$$\left| \sum_{t=1}^k x_t^* \Delta\tilde{\alpha}_t^j \right| \leq 2^k \cdot 2^{j-k} \prod_{t=1}^k |\Delta\tilde{\alpha}_t|.$$

Plugging in this result, we can reach that

$$\begin{aligned} |(V_{\tilde{\alpha}} \mathbf{x}^* - M(\alpha))_j| &= \left| \sum_{p=1}^j \binom{j}{p} \alpha^{j-p} \sum_{t=1}^k x_t^* \Delta\tilde{\alpha}_t^p \right| \leq \sum_{p=k}^j \binom{j}{p} \cdot 2^k \cdot 2^{p-k} \prod_{t=1}^k |\Delta\tilde{\alpha}_t| \\ &\leq 2^k \cdot 2^{2j} \prod_{t=1}^k |\Delta\tilde{\alpha}_t|. \end{aligned}$$

As a result, we conclude that

$$\min_{\mathbf{x} \in F^k} \|V_{\tilde{\alpha}} \mathbf{x} - M(\alpha)\|_2 \leq \|V_{\tilde{\alpha}} \mathbf{x}^* - M(\alpha)\|_2 \leq \sqrt{\sum_{j=k}^{2k-1} \left(2^k \cdot 2^{2j} \prod_{t=1}^k |\Delta\tilde{\alpha}_t| \right)^2} \leq 2^{O(k)} \prod_{j=1}^k |\alpha - \tilde{\alpha}_j|.$$

■

Lemma 28 shows that every discrete distribution in ground truth ϑ has a close spike in recovered positions $\tilde{\alpha}$ and thus can be well approximated according to the above lemma. Therefore, the mixture of the discrete distributions can also be well approximated with support $\tilde{\alpha}$. As the following lemma shows, solving linear regression in Line 5 finds a weight vector \tilde{w} in which the corresponding moment vector $V_{\tilde{\alpha}} \hat{w}$ is close to the ground truth $M(\vartheta)$.

Lemma 33 *Let $\hat{w} = [\hat{w}_1, \dots, \hat{w}_k]^\top \in \mathbb{R}^k$ be the intermediate result (Line 5) in Algorithm 1. Then,*

$$\|V_{\tilde{\alpha}} \hat{w} - M(\vartheta)\|_\infty \leq 2^{O(k)} \cdot \xi^{\frac{1}{2}}.$$

Proof Firstly, we have

$$\begin{aligned} \|V_{\tilde{\alpha}} \hat{w} - M(\vartheta)\|_2 &\leq \|V_{\tilde{\alpha}} \hat{w} - M'\|_2 + \|M(\vartheta) - M'\|_2 \\ &\leq O(1) \cdot \min_{\mathbf{x} \in \mathbb{R}^k} \|V_{\tilde{\alpha}} \mathbf{x} - M'\|_2 + \|M(\vartheta) - M'\|_2 \\ &\leq O(1) \cdot \min_{\mathbf{x} \in \mathbb{R}^k} \|V_{\tilde{\alpha}} \mathbf{x} - M(\vartheta)\|_2 + O(1) \cdot \|M(\vartheta) - M'\|_2 \end{aligned}$$

where the first and third inequalities hold due to triangle inequality, and the second inequality holds since $\|V_{\tilde{\alpha}} \hat{w} - M'\|_2^2 \leq O(1) \cdot \min_{\mathbf{x} \in \mathbb{R}^k} \|V_{\tilde{\alpha}} \mathbf{x} - M'\|_2^2$ (we only find an $O(1)$ -approximation in this step).

For the first term, we can see that

$$\begin{aligned}
\min_{\mathbf{x} \in \mathbb{R}^k} \|V_{\tilde{\alpha}} \mathbf{x} - M(\boldsymbol{\vartheta})\|_2 &= \min_{\mathbf{x}_1, \dots, \mathbf{x}_k \in \mathbb{R}^k} \left\| \sum_{i=1}^k w_i V_{\tilde{\alpha}} \mathbf{x}_i - \sum_{i=1}^k w_i M(\alpha_i) \right\|_2 \\
&\leq \sum_{i=1}^k w_i \min_{\mathbf{x} \in \mathbb{R}^k} \|V_{\tilde{\alpha}} \mathbf{x} - M(\alpha_i)\|_2 \\
&\leq 2^{O(k)} \sum_{i=1}^k w_i \prod_{j=1}^k |\alpha_i - \tilde{\alpha}_j| && \text{(Lemma 32)} \\
&\leq 2^{O(k)} \sqrt{\sum_{i=1}^k w_i \prod_{j=1}^k |\alpha_i - \tilde{\alpha}_j|^2} && \text{(AM-QM Inequality, } \mathbf{w} \in \Delta_{k-1}\text{)} \\
&\leq 2^{O(k)} \cdot \sqrt{\xi}. && \text{(Lemma 28).}
\end{aligned}$$

For the second term, $\|M(\boldsymbol{\vartheta}) - M'\|_2 \leq \sqrt{k} \|M(\boldsymbol{\vartheta}) - M'\|_\infty \leq \sqrt{k} \cdot \xi$. As a result,

$$\|V_{\tilde{\alpha}} \hat{\mathbf{w}} - M(\boldsymbol{\vartheta})\|_\infty \leq \|V_{\tilde{\alpha}} \hat{\mathbf{w}} - M(\boldsymbol{\vartheta})\|_2 \leq O(1) \cdot 2^{O(k)} \cdot \sqrt{\xi} + O(1) \cdot \sqrt{k} \cdot \xi \leq 2^{O(k)} \cdot \zeta^{\frac{1}{2}}.$$

■

The moment-transportation inequality (Lemma 8) requires the input to be signed measure. So, we normalize the recovered spikes in Line 6. The following lemma shows the moment vector would not change too much after normalization.

Lemma 34 *Let $\tilde{\boldsymbol{\vartheta}} = (\tilde{\alpha}, \tilde{\mathbf{w}})$ be the intermediate result (Line 7) in Algorithm 1. Then, we have*

$$\text{Mom}_K(\tilde{\boldsymbol{\vartheta}}, \boldsymbol{\vartheta}) \leq 2^{O(k)} \cdot \xi^{\frac{1}{2}}.$$

Proof According to the definition of $V_{\tilde{\alpha}}$ and $M(\boldsymbol{\vartheta})$, we have $(V_{\tilde{\alpha}} \hat{\mathbf{w}} - M(\boldsymbol{\vartheta}))_1 = \sum_{i=1}^k \hat{w}_i - 1$. Therefore, $|\sum_{i=1}^k \hat{w}_i - 1| \leq |(V_{\tilde{\alpha}} \hat{\mathbf{w}} - M(\boldsymbol{\vartheta}))_1| \leq \|V_{\tilde{\alpha}} \hat{\mathbf{w}} - M(\boldsymbol{\vartheta})\|_\infty \leq 2^{O(k)} \cdot \xi^{\frac{1}{2}}$.

Note that $\|\tilde{\mathbf{w}} - \hat{\mathbf{w}}\|_1 = \|\tilde{\mathbf{w}}\|_1 - \|\hat{\mathbf{w}}\|_1 = |(\sum_{i=1}^k \hat{w}_i)^{-1} - 1|$. For $2^{O(k)} \cdot \xi^{\frac{1}{2}} \leq 1/2$, we can conclude that $\|\tilde{\mathbf{w}} - \hat{\mathbf{w}}\|_1 = |(\sum_{i=1}^k \hat{w}_i)^{-1} - 1| \leq 2|\sum_{i=1}^k \hat{w}_i - 1| \leq 2^{O(k)} \cdot \xi^{\frac{1}{2}}$.

Thus, we can reach that

$$\begin{aligned}
\text{Mom}_K(\tilde{\boldsymbol{\vartheta}}, \boldsymbol{\vartheta}) &= \|M(\tilde{\boldsymbol{\vartheta}}) - M(\boldsymbol{\vartheta})\|_\infty = \|V_{\tilde{\alpha}} \tilde{\mathbf{w}} - M(\boldsymbol{\vartheta})\|_\infty \\
&\leq \|V_{\tilde{\alpha}} \hat{\mathbf{w}} - M(\boldsymbol{\vartheta})\|_\infty + \|V_{\tilde{\alpha}} \tilde{\mathbf{w}} - V_{\tilde{\alpha}} \hat{\mathbf{w}}\|_\infty \\
&\leq \|V_{\tilde{\alpha}} \hat{\mathbf{w}} - M(\boldsymbol{\vartheta})\|_\infty + \|V_{\tilde{\alpha}}\|_\infty \|\tilde{\mathbf{w}} - \hat{\mathbf{w}}\|_1 \\
&\leq 2^{O(k)} \cdot \xi^{\frac{1}{2}}.
\end{aligned}$$

where the first inequality holds due to triangle inequality and the third holds since $\|V_{\tilde{\alpha}}\|_\infty \leq 1$. ■

Moreover, the weight can still be negative after normalization. To find a reconstruction in the original space, we want to find a mixture of discrete distributions in $\text{Spike}(\Delta_1, \Delta_{k-1})$ that is close to $\boldsymbol{\vartheta}$. However, this step can cast a huge impact on the moment. Thus, we directly estimate the influence in terms of the transportation distance instead of the moment distance. We note the transportation distance in the next lemma is defined possibly for non-probability measures (see Equation (4) in Appendix A for details).

Lemma 35 *Let \mathcal{C} be a compact convex set in \mathbb{R}^d and let $\tilde{\vartheta} = (\tilde{\alpha}, \tilde{w}) \in \text{Spike}(\mathcal{C}, \Sigma_{k-1})$ be a k -spike distribution over support \mathcal{C} . Then, we have*

$$\min_{\tilde{\vartheta} \in \text{Spike}(\tilde{\alpha}, \Delta_{k-1})} \text{Tran}(\tilde{\vartheta}, \tilde{\vartheta}) \leq 2 \min_{\bar{\vartheta} \in \text{Spike}(\mathcal{C}, \Delta_{k-1})} \text{Tran}(\tilde{\vartheta}, \bar{\vartheta}).$$

Note that the minimization on the left is over support $\tilde{\alpha}$.

Proof Consider any $\bar{\vartheta} = (\bar{\alpha}, \bar{w}) \in \text{Spike}(\mathcal{C}, \Delta_{k-1})$ such that $\bar{\alpha} = [\bar{\alpha}_1, \dots, \bar{\alpha}_k]^\top$ and $\bar{w} = [\bar{w}_1, \dots, \bar{w}_k]^\top$. Let $\vartheta' = \arg \min_{\vartheta' \in \text{Spike}(\tilde{\alpha}, \Sigma_{k-1})} \text{Tran}(\vartheta', \bar{\vartheta})$. Since $\tilde{\vartheta} \in \text{Spike}(\tilde{\alpha}, \Sigma_{k-1})$, we have $\text{Tran}(\vartheta', \bar{\vartheta}) \leq \text{Tran}(\tilde{\vartheta}, \bar{\vartheta})$. From triangle inequality, we have

$$\text{Tran}(\tilde{\vartheta}, \vartheta') \leq \text{Tran}(\tilde{\vartheta}, \bar{\vartheta}) + \text{Tran}(\bar{\vartheta}, \vartheta') \leq 2 \text{Tran}(\tilde{\vartheta}, \bar{\vartheta}).$$

Now, we start to show $\vartheta' \in \text{Spike}(\tilde{\alpha}, \Delta_{k-1})$, i.e., the weight of spikes in the closest distribution to $\bar{\vartheta}$ with support $\tilde{\alpha}$ should be all non-negative.

Towards a contradiction, assume the optimal distribution ϑ' contains at least one negative spike, that is, let $(\vartheta')^- = (\alpha, \max\{-w, \mathbf{0}\})$ be the negative spikes of ϑ' , we have $(\vartheta')^- \neq \mathbf{0}$. In this case, since $\bar{\vartheta}$ is a probability distribution, $\text{Tran}(\bar{\vartheta}, \vartheta') = \text{Tran}(\bar{\vartheta} + (\vartheta')^-, (\vartheta')^+)$ where $(\vartheta')^+ = (\alpha, \max\{w, \mathbf{0}\})$ is positive spikes of ϑ' .

Let μ' be the optimal matching distribution corresponding to $\text{Tran}(\bar{\vartheta} + (\vartheta')^-, (\vartheta')^+)$. From its definition, μ' is non-negative. Let μ'' be the matching distribution in μ' that maps $\bar{\vartheta}$ and $(\vartheta')^+$. Let ϑ'' be the marginal distribution of μ'' other than $\bar{\vartheta}$. In this case, μ'' is a valid matching distribution for $\text{Tran}(\bar{\vartheta}, \vartheta'')$. Moreover, μ'' gives less transportation distance than μ' does since the cost on the eliminated terms is non-negative. Hence, $\text{Tran}(\bar{\vartheta}, \vartheta'') < \text{Tran}(\bar{\vartheta}, \vartheta')$, which is a contradiction.

As a result, $\vartheta' \in \text{Spike}(\tilde{\alpha}, \Delta_{k-1})$. Therefore, $\min_{\tilde{\vartheta} \in \text{Spike}(\tilde{\alpha}, \Delta_{k-1})} \text{Tran}(\tilde{\vartheta}, \tilde{\vartheta}) \leq \text{Tran}(\tilde{\vartheta}, \vartheta')$.

By taking minimization over $\bar{\vartheta}$, we can conclude that

$$\min_{\tilde{\vartheta} \in \text{Spike}(\tilde{\alpha}, \Delta_{k-1})} \text{Tran}(\tilde{\vartheta}, \tilde{\vartheta}) \leq 2 \min_{\bar{\vartheta} \in \text{Spike}(\mathcal{C}, \Delta_{k-1})} \text{Tran}(\tilde{\vartheta}, \bar{\vartheta}).$$

■

Finally, we are ready to bound the transportation distance error of the reconstructed k -spike distribution:

Lemma 36 *Let $\tilde{\vartheta} = (\tilde{\alpha}, \tilde{w})$ be the final result (Line 9) in Algorithm 1. Then, we have*

$$\text{Tran}(\tilde{\vartheta}, \vartheta) \leq O(k\xi^{\frac{1}{4k-2}}).$$

Proof Combining Lemma 8 and Lemma 34, we can see

$$\text{Tran}(\tilde{\vartheta}, \vartheta) \leq O(k \cdot (2^{O(k)} \cdot \xi^{\frac{1}{2}})^{\frac{1}{2k-1}}) \leq O(k\xi^{\frac{1}{4k-2}}).$$

Moreover, from Lemma 35, we have

$$\text{Tran}(\tilde{\vartheta}, \tilde{\vartheta}) \leq 2 \min_{\bar{\vartheta} \in \text{Spike}(\Delta_1, \Delta_{k-1})} \text{Tran}(\tilde{\vartheta}, \bar{\vartheta}) \leq 2 \text{Tran}(\tilde{\vartheta}, \vartheta) = O(k\xi^{\frac{1}{4k-2}}).$$

Finally, by triangle inequality, $\text{Tran}(\tilde{\vartheta}, \vartheta) \leq \text{Tran}(\tilde{\vartheta}, \tilde{\vartheta}) + \text{Tran}(\tilde{\vartheta}, \vartheta) \leq O(k\xi^{\frac{1}{4k-2}})$. ■

By choosing $\xi = (\epsilon/k)^{O(k)}$, the previous lemma directly implies Theorem 3.

Appendix D. An Algorithm for the Two-Dimensional Problem

We can generalize the 1-dimension algorithm described in Section 4 to 2-dimension. Let $\vartheta := (\boldsymbol{\alpha}, \boldsymbol{w}) \in \text{Spike}(\Delta_2, \Delta_{k-1})$ be the underlying mixture where $\boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_k\}$ for $\alpha_i = (\alpha_{i,1}, \alpha_{i,2})$ and $\boldsymbol{w} = [w_1, \dots, w_k]^\top \in \Delta_{k-1}$. The true moments can be computed according to $M_{i,j}(\vartheta) = \sum_{t=1}^k w_t \alpha_{t,1}^i \alpha_{t,2}^j$. The input is the noisy moments $M'_{i,j}$ in $0 \leq i, j, i+j \leq 2k-1$ such that $|M'_{i,j} - M_{i,j}(\vartheta)| \leq \xi$. We further assume that $M'_{0,0} = M_{0,0}(\vartheta) = 1$ and $\xi \leq 2^{-\Omega(k)}$.

The key idea is simple: a distribution supported in \mathbb{R}^2 can be mapped to a distribution supported in the complex plane \mathbb{C} . In particular, we define the complex set $\Delta_{\mathbb{C}} = \{a + bi \mid (a, b) \in \Delta_2\}$. Moreover, we denote $\boldsymbol{\beta} = [\beta_1, \dots, \beta_k]^\top := [\alpha_{1,1} + \alpha_{1,2}i, \dots, \alpha_{k,1} + \alpha_{k,2}i]^\top \in \Delta_{\mathbb{C}}^k$, and define $\phi := (\boldsymbol{\beta}, \boldsymbol{w}) \in \text{Spike}(\Delta_{\mathbb{C}}^k, \Delta_{k-1})$ to be the complex mixture corresponding to ϑ . The corresponding moments of ϕ can thus be defined as $G_{i,j}(\phi) = \sum_{t=1}^k w_t (\beta_t^\dagger)^i \beta_t^j$.

For $G = [G_{i,j}]_{0 \leq i \leq k; 0 \leq j \leq k-1}^\top$, denote

$$A_G := \begin{bmatrix} G_{0,0} & G_{0,1} & \cdots & G_{0,k-1} \\ G_{1,0} & G_{1,1} & \cdots & G_{1,k-1} \\ \vdots & \vdots & \ddots & \vdots \\ G_{k-1,0} & G_{k-1,1} & \cdots & G_{k-1,k-1} \end{bmatrix}, b_G := \begin{bmatrix} G_{0,k} \\ G_{1,k} \\ \vdots \\ G_{k-1,k} \end{bmatrix}, M_G := \begin{bmatrix} G_{0,0} \\ G_{0,1} \\ \vdots \\ G_{0,k-1} \end{bmatrix}.$$

Algorithm 2 Reconstruction Algorithm in Two-Dimensional Problem

Input: number of spikes k , noisy moments $M'(\cdot)$, noise level ξ

Output: recovered spike distribution $\check{\vartheta}$

- | | |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| 1: $G' \leftarrow [G'_{i,j} := \sum_{p=0}^i \sum_{q=0}^j \binom{i}{p} \binom{j}{q} (-i)^{i-p} i^{j-q} M'_{p+q, i+j-p-q}]_{0 \leq i \leq k; 0 \leq j \leq k-1}^\top$ | |
| 2: $\hat{\boldsymbol{c}} \leftarrow \arg \min_{\boldsymbol{x} \in \mathbb{C}^k} \ A_{G'} \boldsymbol{x} + b_{G'}\ _2^2 + \xi^2 \ \boldsymbol{x}\ _2^2$ | $\Rightarrow \hat{\boldsymbol{c}} = [\hat{c}_0, \dots, \hat{c}_{k-1}]^\top \in \mathbb{C}^k$ |
| 3: $\hat{\boldsymbol{\beta}} \leftarrow \text{roots}(\sum_{i=0}^{k-1} \hat{c}_i x^i + x^k)$ | $\Rightarrow \hat{\boldsymbol{\beta}} = [\hat{\beta}_1, \dots, \hat{\beta}_k]^\top \in \mathbb{C}^k$ |
| 4: $\bar{\boldsymbol{\beta}} \leftarrow \text{project}_{\Delta_{\mathbb{C}}}(\hat{\boldsymbol{\beta}})$ | $\Rightarrow \bar{\boldsymbol{\beta}} = [\bar{\beta}_1, \dots, \bar{\beta}_k]^\top \in \Delta_{\mathbb{C}}^k$ |
| 5: $\tilde{\boldsymbol{\beta}} \leftarrow \bar{\boldsymbol{\beta}} + \text{Noise}(\xi)$ | $\Rightarrow \tilde{\boldsymbol{\beta}} = [\tilde{\beta}_1, \dots, \tilde{\beta}_k]^\top \in \Delta_{\mathbb{C}}^k$ |
| 6: $\hat{\boldsymbol{w}} \leftarrow \arg \min_{\boldsymbol{x} \in \mathbb{C}^k} \ V_{\tilde{\boldsymbol{\beta}}} \boldsymbol{x} - M_{G'}\ _2^2$ | $\Rightarrow \hat{\boldsymbol{w}} = [\hat{w}_1, \dots, \hat{w}_k]^\top \in \mathbb{C}^k$ |
| 7: $\tilde{\boldsymbol{w}} \leftarrow \hat{\boldsymbol{w}} / (\sum_{i=1}^k \hat{w}_i)$ | $\Rightarrow \tilde{\boldsymbol{w}} = [\tilde{w}_1, \dots, \tilde{w}_k]^\top \in \Sigma_{k-1}^{\mathbb{C}}$ |
| 8: $\tilde{\boldsymbol{\phi}} \leftarrow (\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{w}})$ | $\Rightarrow \tilde{\boldsymbol{\phi}} \in \text{Spike}(\Delta_{\mathbb{C}}, \Sigma_{k-1}^{\mathbb{C}})$ |
| 9: $\tilde{\boldsymbol{\vartheta}} \leftarrow ([\text{real}(\tilde{\boldsymbol{\beta}}), \text{imag}(\tilde{\boldsymbol{\beta}})], \text{real}(\tilde{\boldsymbol{w}}))$ | $\Rightarrow \tilde{\boldsymbol{\vartheta}} \in \text{Spike}(\Delta_2, \Sigma_{k-1})$ |
| 10: $\check{\boldsymbol{w}} \leftarrow \arg \min_{\boldsymbol{x} \in \Delta_{k-1}} \text{Tran}(\tilde{\boldsymbol{\vartheta}}, (\tilde{\boldsymbol{\alpha}}, \boldsymbol{x}))$ | |
| 11: $\check{\boldsymbol{\vartheta}} \leftarrow (\tilde{\boldsymbol{\alpha}}, \check{\boldsymbol{w}})$ | $\Rightarrow \check{\boldsymbol{\vartheta}} \in \text{Spike}(\Delta_2, \Delta_{k-1})$ |
-

The pseudocode can be found in Algorithm 2. The algorithm takes the number of spikes k and the error bound ξ to reconstruct the original spike distribution using empirical moments M' . Now, we describe the implementation details of our algorithm.

We first calculate the empirical complete moments G' of ϕ in Line 1. Since the process follows relationship between $G(\phi)$ and $M(\vartheta)$ (see Lemma 37), G' would be an estimation of ground truth $G(\phi)$. Since it is a two-dimensional convolution, this step can be implemented in $O(k^3)$ arithmetic operations.

Then, we perform a ridge regression to obtain \mathbf{c} in Line 2. We note that $A_{G(\phi)}\mathbf{c} + b_{G(\phi)} = \mathbf{0}$. (see Lemma 38). Hence, $\hat{\mathbf{c}}$ can be seen as an approximation of \mathbf{c} . The explicit solution of this ridge regression is $\hat{\mathbf{c}} = (A_{G'}^\top A_{G'} + \xi^2 I)^{-1} A_{G'}^\top b_{G'}$ which can be computed in $O(k^3)$ time.

From Line 3 to Line 5, we aim to estimate the positions of the spikes of complex correspondence, i.e., β_i s. Similar to the 1-dimensional case, we solve the roots of polynomial $\sum_{i=0}^{k-1} \hat{c}_i x^i + x^k$. Note that the roots we found may locate outside $\Delta_{\mathbb{C}}$, which is the support of the ground truth. Thus, we use the description of Δ_2 to project the solutions back to $\Delta_{\mathbb{C}}$ and inject small noise to ensure that all values are distinct and $\tilde{\beta}$ are still in $\Delta_{\mathbb{C}}$. Any noise of size at most ξ suffices here. We note that from the definition of complex correspondence, the realized spikes in the original space are $[\text{real}(\tilde{\beta}), \text{imag}(\tilde{\beta})]$. For implementation, this step can be done using the numerical root-finding algorithm allowing ξ additive noise in $O(k^{1+o(1)} \cdot \log \log(1/\xi))$ arithmetic operations.

After that, we aim to recover the weight of the spikes. Line 6 is a linear regression defined by Vandemonde matrix V_{β} . Since β may be complex numbers, we would calculate over complex space. Again, we would apply moment inequality to bound the recovered parameter error. Hence, we normalize $\hat{\mathbf{w}}$ and get $\tilde{\mathbf{w}}$ in Line 9. Note from our definition of transportation distance, the real components and imaginary components are considered separately. Hence, we can discard the imaginary parts of $\tilde{\mathbf{w}}$ and reconstruct the k -spike distribution in Line 7. Using linear regression, this step can be done in $O(k^3)$ time.

The remaining thing is to deal with the negative weights of $\tilde{\mathbf{w}}$. We find a close k -spike distribution in $\text{Spike}(\Delta_2, \Delta_{k-1})$ in Line 7. The optimization problem is equivalent to finding a transportation from $(\tilde{\alpha}, \tilde{\mathbf{w}}^-)$ to $(\tilde{\alpha}, \tilde{\mathbf{w}}^+)$ where $\tilde{\mathbf{w}}^-$ and $\tilde{\mathbf{w}}^+$ are the negative components and positive components of $\tilde{\mathbf{w}}$ respectively, i.e., $\tilde{w}_i^- = \max\{0, -\tilde{w}_i\}$ and $\tilde{w}_i^+ = \max\{0, \tilde{w}_i\}$. This transportation can be found using the standard network flow technique, which takes $O(k^3)$ time.

Since the noise satisfies $\xi \leq 2^{-\Omega(k)}$, the whole algorithm requires $O(k^3)$ arithmetic operations.

D.1. Error Analysis

From now, we bound the reconstruction error of the algorithm. The following lemma presents the relationship between the moments of the original k -spike distribution and its complex correspondence.

Lemma 37 *Let $\vartheta = (\alpha, \mathbf{w}) \in \text{Spike}(\Delta_2, \Delta_{k-1})$ and let $\phi = (\beta, \mathbf{w}) \in \text{Spike}(\Delta_{\mathbb{C}}, \Delta_{k-1})$ be its complex correspondence. Then, the complete moment satisfies*

$$G_{i,j}(\phi) = \sum_{p=0}^i \sum_{q=0}^j \binom{i}{p} \binom{j}{q} (-i)^{i-p} i^{j-q} M_{p+q, i+j-p-q}(\vartheta).$$

Proof According to the definition,

$$\begin{aligned} G_{i,j}(\phi) &= \sum_{t=1}^k w_t (\beta_t^\dagger)^i \beta_t^j = \sum_{t=1}^k w_t (\alpha_{t,1} - \alpha_{t,2}i)^i (\alpha_{t,1} + \alpha_{t,2}i)^j \\ &= \sum_{t=1}^k w_t \sum_{p=0}^i \binom{i}{p} \alpha_{t,1}^p \alpha_{t,2}^{i-p} (-i)^{i-p} \sum_{q=0}^j \binom{j}{q} \alpha_{t,1}^q \alpha_{t,2}^{j-q} i^{j-q} \\ &= \sum_{p=0}^i \sum_{q=0}^j \binom{i}{p} \binom{j}{q} (-i)^{i-p} i^{j-q} M_{p+q, i+j-p-q}(\vartheta). \end{aligned}$$

■

The following lemma is a complex extension of Lemma 24.

Lemma 38 *Let $\phi = (\boldsymbol{\beta}, \boldsymbol{w}) \in \text{Spike}(\Delta_{\mathbb{C}}, \Delta_{k-1})$ where $\boldsymbol{\beta} = [\beta_1, \dots, \beta_k]^\top$ and $\boldsymbol{w} = [w_1, \dots, w_k]^\top$. Let $\boldsymbol{c} = [c_0, \dots, c_{k-1}]^\top \in \mathbb{C}^k$ such that $\prod_{i=1}^k (x - \alpha_i) = \sum_{i=0}^{k-1} c_i x^i + x^k$. Then,*

$$\sum_{j=0}^{k-1} G_{i,j}(\boldsymbol{\vartheta}) c_j + G_{i,k}(\boldsymbol{\vartheta}) = \sum_{j=0}^{k-1} G_{j,i}(\boldsymbol{\vartheta}) c_j^\dagger + G_{k,i}(\boldsymbol{\vartheta}) = 0,$$

for all $i \geq 0$. In matrix form, the equation can be written as:

$$A_{G(\boldsymbol{\vartheta})} \boldsymbol{c} + b_{G(\boldsymbol{\vartheta})} = \mathbf{0}.$$

Proof According to the definition of G , it holds that

$$\begin{aligned} \sum_{j=0}^{k-1} G_{i,j}(\phi) c_j + G_{k,i}(\phi) &= \sum_{j=0}^{k-1} \sum_{t=1}^k w_t (\beta_t^\dagger)^i \beta_t^j c_j + \sum_{t=1}^k w_t (\beta_t^\dagger)^i \beta_t^k \\ &= \sum_{t=1}^k w_t (\beta_t^\dagger)^i \left(\sum_{j=0}^{k-1} c_j \beta_t^j + \beta_t^k \right) \\ &= \sum_{t=1}^k w_t (\beta_t^\dagger)^i \prod_{j=1}^k (\beta_t - \beta_j) = 0, \end{aligned}$$

$$\begin{aligned} \sum_{j=0}^{k-1} G_{j,i}(\phi) c_j^\dagger + G_{k,i}(\phi) &= \sum_{j=0}^{k-1} \sum_{t=1}^k w_t (\beta_t^\dagger)^j \beta_t^i c_j^\dagger + \sum_{t=1}^k w_t (\beta_t^\dagger)^k \beta_t^i \\ &= \left(\sum_{j=0}^{k-1} \sum_{t=1}^k w_t (\beta_t^\dagger)^i \beta_t^j c_j + \sum_{t=1}^k w_t (\beta_t^\dagger)^i \beta_t^k \right)^\dagger = 0. \end{aligned}$$

■

Similar to Lemma 25, we present some useful properties of $\hat{\boldsymbol{c}}$.

Lemma 39 *Let $\hat{\boldsymbol{c}} = [\hat{c}_0, \dots, \hat{c}_{k-1}]^\top \in \mathbb{R}^k$ be the intermediate result (Line 2) in the algorithm. Then, $\|G' - G(\phi)\|_\infty \leq 2^{2k} \cdot \xi$, $\|\boldsymbol{c}\|_1 \leq 2^k$, $\|\hat{\boldsymbol{c}}\|_1 \leq 2^{O(k)}$ and $\|A_{G(\phi)} \hat{\boldsymbol{c}} + b_{G(\phi)}\|_\infty \leq 2^{O(k)} \cdot \xi$.*

Proof From Vieta's formulas, we have $c_i = \sum_{S \in \binom{[k]}{k-i}} \prod_{j \in S} (-\alpha_j)$. Thus,

$$\|\boldsymbol{c}\|_1 = \sum_{i=0}^{k-1} |c_i| \leq \sum_{S \in 2^{[k]}} \prod_{j \in S} |\beta_j| = \prod_{i=1}^k (1 + |\beta_i|) \leq 2^k.$$

where the last inequality holds because $|\beta_i| = \sqrt{\alpha_{i,1}^2 + \alpha_{i,2}^2} \leq 1$ for all i .

According to Lemma 37,

$$|G'_{i,j} - G_{i,j}(\phi)| \leq \sum_{p=0}^i \sum_{q=0}^j \left| \binom{i}{p} \binom{j}{q} (-i)^{i-p} i^{j-q} \right| \|M' - M(\vartheta)\|_\infty \leq 2^{i+j} \cdot \xi.$$

This shows that $\|G' - G(\phi)\|_\infty \leq 2^{2k} \cdot \xi$.

From Lemma 24, we can see that $\|A_{M(\vartheta)}\mathbf{c} + b_{M(\vartheta)}\|_\infty = 0$. Therefore,

$$\begin{aligned} \|A_{G'}\mathbf{c} + b_{G'}\|_\infty &\leq \|A_{G(\phi)}\mathbf{c} + b_{G(\phi)}\|_\infty + \|A_{G'}\mathbf{c} - A_{G(\phi)}\mathbf{c}\|_\infty + \|b_{G'} - b_{G(\phi)}\|_\infty \\ &\leq \|A_{G(\phi)}\mathbf{c} + b_{G(\phi)}\|_\infty + \|(A_{G'} - A_{G(\phi)})\mathbf{c}\|_\infty + \|b_{G'} - b_{G(\phi)}\|_\infty \\ &\leq \|A_{G(\phi)}\mathbf{c} + b_{G(\phi)}\|_\infty + \|A_{G'} - A_{G(\phi)}\|_\infty \|\mathbf{c}\|_1 + \|b_{G'} - b_{G(\phi)}\|_\infty \\ &\leq \|A_{G(\phi)}\mathbf{c} + b_{G(\phi)}\|_\infty + \|G' - G(\phi)\|_\infty (\|\mathbf{c}\|_1 + 1) \\ &\leq 2^{O(k)} \cdot \xi. \end{aligned}$$

The fourth inequality holds since $(A_{G'} - A_{G(\phi)})_{i,j} = (G' - G(\phi))_{i,j}$ and $(b_{G'} - b_{G(\phi)})_i = (G' - G(\phi))_{i,k}$.

From the definition of $\hat{\mathbf{c}}$, we can see that

$$\begin{aligned} \|A_{G'}\hat{\mathbf{c}} + b_{G'}\|_2^2 + \xi^2 \|\hat{\mathbf{c}}\|_2^2 &\leq \|A_{G'}\mathbf{c} + b_{G'}\|_2^2 + \xi^2 \|\mathbf{c}\|_2^2 \\ &\leq k \|A_{G'}\mathbf{c} + b_{G'}\|_\infty^2 + \xi^2 \|\mathbf{c}\|_1^2 \\ &\leq 2^{O(k)} \cdot \xi^2. \end{aligned}$$

The second inequality holds since $\|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_1$ and $\|\mathbf{x}\|_2 \leq \sqrt{k} \|\mathbf{x}\|_\infty$ holds for any vector $\mathbf{x} \in \mathbb{R}^k$.

Now, we get that

$$\begin{aligned} \|A_{G'}\hat{\mathbf{c}} + b_{G'}\|_\infty &\leq \|A_{G'}\hat{\mathbf{c}} + b_{G'}\|_2 \leq 2^{O(k)} \cdot \xi, \\ \|\hat{\mathbf{c}}\|_1 &\leq \sqrt{k} \|\hat{\mathbf{c}}\|_2 \leq 2^{O(k)}. \end{aligned}$$

Finally, we can bound $\|A_{G(\phi)}\hat{\mathbf{c}} + b_{G(\phi)}\|_\infty$ according to

$$\begin{aligned} \|A_{G(\phi)}\hat{\mathbf{c}} + b_{G(\phi)}\|_\infty &\leq \|A_{G'}\hat{\mathbf{c}} + b_{G'}\|_\infty + \|A_{G'}\hat{\mathbf{c}} - A_{G(\phi)}\hat{\mathbf{c}}\|_\infty + \|b_{G'} - b_{G(\phi)}\|_\infty \\ &\leq \|A_{G'}\hat{\mathbf{c}} + b_{G'}\|_\infty + \|A_{G'} - A_{G(\phi)}\|_\infty \|\hat{\mathbf{c}}\|_1 + \|b_{G'} - b_{G(\phi)}\|_\infty \\ &\leq \|A_{G'}\hat{\mathbf{c}} + b_{G'}\|_\infty + \|G' - G(\phi)\|_\infty (\|\hat{\mathbf{c}}\|_1 + 1) \\ &\leq 2^{O(k)} \cdot \xi, \end{aligned}$$

which finishes the proof. ■

We then show $\tilde{\mathbf{w}}$ obtained in Line 7 is close to the ground truth β .

Lemma 40 *Let $\hat{\beta} = [\hat{\beta}_1, \dots, \hat{\beta}_k]^\top \in \mathbb{C}^k$ be the intermediate result (Line 3) in Algorithm 2. Then,*

$$\sum_{i=1}^k w_i \prod_{j=1}^k |\beta_i - \hat{\beta}_j|^2 \leq 2^{O(k)} \cdot \xi.$$

Proof Consider

$$\begin{aligned}
& \sum_{i=1}^k w_i \prod_{j=1}^k |\beta_i - \hat{\beta}_j|^2 = \sum_{i=1}^k w_i \left| \prod_{j=1}^k (\beta_i - \hat{\beta}_j) \right|^2 = \sum_{i=1}^k w_i \left| \sum_{j=0}^{k-1} \hat{c}_j \beta_i^t + \beta_i^k \right|^2 \\
&= \sum_{i=1}^k w_i \left(\sum_{j=0}^{k-1} \hat{c}_j \beta_i^t + \beta_i^k \right)^\dagger \left(\sum_{j=0}^{k-1} \hat{c}_j \beta_i^t + \beta_i^k \right) \\
&= \sum_{i=1}^k w_i \left(\sum_{p=0}^{k-1} \sum_{q=0}^{k-1} \hat{c}_p^\dagger \hat{c}_q (\beta_i^\dagger)^p \beta_i^q + \sum_{p=0}^{k-1} \hat{c}_p^\dagger (\beta_i^\dagger)^p \beta_i^k + \sum_{p=0}^{k-1} \hat{c}_p \beta_i^p (\beta_i^\dagger)^k + (\beta_i^\dagger)^k \beta_i^k \right) \\
&= \sum_{p=0}^{k-1} \sum_{q=0}^{k-1} \hat{c}_p^\dagger \hat{c}_q \sum_{i=1}^k w_i (\beta_i^\dagger)^p \beta_i^q + \sum_{p=0}^{k-1} \hat{c}_p^\dagger \sum_{i=1}^k w_i (\beta_i^\dagger)^p \beta_i^k + \sum_{p=0}^{k-1} \hat{c}_p \sum_{i=1}^k w_i \beta_i^p (\beta_i^\dagger)^k + \sum_{i=1}^k w_i (\beta_i^\dagger)^k \beta_i^k \\
&= \sum_{p=0}^{k-1} \hat{c}_p^\dagger \left(\sum_{q=0}^{k-1} G_{p,q}(\boldsymbol{\vartheta}) \hat{c}_q + G_{p,k}(\boldsymbol{\vartheta}) \right) + \sum_{p=0}^{k-1} \hat{c}_p G_{k,p}(\boldsymbol{\vartheta}) + G_{k,k}(\boldsymbol{\vartheta}) \\
&= \sum_{p=0}^{k-1} \hat{c}_p^\dagger (A_{G(\boldsymbol{\vartheta})} \hat{c} + b_{G(\boldsymbol{\vartheta})})_p + \sum_{p=0}^{k-1} \hat{c}_p G_{k,p}(\boldsymbol{\vartheta}) + G_{k,k}(\boldsymbol{\vartheta}),
\end{aligned}$$

where the last equality holds from the definition of matrix $A_{G(\boldsymbol{\vartheta})}$ and vector $b_{G(\boldsymbol{\vartheta})}$. According to Lemma 38, $G_{k,p}(\boldsymbol{\vartheta}) = -\sum_{q=0}^{k-1} G_{q,p}(\boldsymbol{\vartheta}) c_q^\dagger$, so

$$\begin{aligned}
\sum_{p=0}^{k-1} \hat{c}_p G_{k,p}(\boldsymbol{\vartheta}) + G_{k,k}(\boldsymbol{\vartheta}) &= \sum_{p=0}^{k-1} \hat{c}_p \left(-\sum_{q=0}^{k-1} G_{q,p}(\boldsymbol{\vartheta}) c_q^\dagger \right) + \left(-\sum_{q=0}^{k-1} G_{q,k}(\boldsymbol{\vartheta}) c_q^\dagger \right) \\
&= -\sum_{q=0}^{k-1} c_q^\dagger \left(\sum_{p=0}^{k-1} G_{q,p}(\boldsymbol{\vartheta}) \hat{c}_p + G_{k,q}(\boldsymbol{\vartheta}) \right) \\
&= -\sum_{q=0}^{k-1} c_q^\dagger (A_{G(\boldsymbol{\vartheta})} \hat{c} + b_{G(\boldsymbol{\vartheta})})_q.
\end{aligned}$$

Therefore,

$$\begin{aligned}
\sum_{i=1}^k w_i \prod_{j=1}^k |\beta_i - \hat{\beta}_j|^2 &= \sum_{i=0}^{k-1} (\hat{c}_i^\dagger - c_i^\dagger) (A_{G(\boldsymbol{\vartheta})} \hat{c} + b_{G(\boldsymbol{\vartheta})})_i \\
&\leq (\|\hat{c}\|_1 + \|c\|_1) \|A_{G(\boldsymbol{\vartheta})} \hat{c} + b_{G(\boldsymbol{\vartheta})}\|_\infty \\
&\leq (2^k + 2^{O(k)}) \cdot 2^{O(k)} \cdot \xi \quad (\text{Lemma 39}) \\
&\leq 2^{O(k)} \cdot \xi.
\end{aligned}$$

■

Similar to Lemma 28, the following lemma shows the error still can be bounded after projection and injecting noise.

Lemma 41 Let $\tilde{\beta} = [\tilde{\beta}_1, \dots, \tilde{\beta}_k]^\top \in \mathbb{C}^k$ be the intermediate result (Line 5) in Algorithm 2. Then,

$$\sum_{i=1}^k w_i \prod_{j=1}^k |\beta_i - \tilde{\beta}_j|^2 \leq 2^{O(k)} \cdot \xi.$$

Proof Let $\bar{\beta} = [\bar{\beta}_1, \dots, \bar{\beta}_k]^\top \in \mathbb{R}^k$ be the set of projections (Line 4). Since Δ_2 is convex, $\Delta_{\mathbb{C}}$ is also convex. From $\beta_i \in \Delta_{\mathbb{C}}$, $|\beta_i - \bar{\beta}_j| \leq |\beta_i - \hat{\beta}_j|$ holds. Thus,

$$\prod_{j=1}^k |\beta_i - \bar{\beta}_j|^2 \leq \prod_{j=1}^k |\beta_i - \hat{\beta}_j|^2.$$

Recall that $\tilde{\beta}_j$ is obtained from $\bar{\beta}_j$ by adding of size noise no more than ξ . We have $|\beta_i - \tilde{\beta}_j| \leq |\beta_i - \bar{\beta}_j| + \xi$. Apply Lemma 27 by regarding $|\beta_i - \tilde{\beta}_j|$ as a_j and $|\beta_i - \hat{\beta}_j|$ as a'_j . From $|\beta_i - \tilde{\beta}_j| \leq |\beta_i - \bar{\beta}_j| + \xi \leq |\beta_i - \hat{\beta}_j| + \xi$, we can conclude that

$$\prod_{j=1}^k |\beta_i - \tilde{\beta}_j|^2 \leq \prod_{j=1}^k |\beta_i - \bar{\beta}_j|^2 + O(2^k \cdot k\xi).$$

Combining two inequalities,

$$\sum_{i=1}^k w_i \prod_{j=1}^k |\beta_i - \tilde{\beta}_j|^2 \leq \sum_{i=1}^k w_i \prod_{j=1}^k |\beta_i - \hat{\beta}_j|^2 + 2k \cdot \xi \leq 2^{O(k)} \cdot \xi.$$

■

Similar to Lemma 33, we can bound the error of approximating ground truth over recovered spikes.

Lemma 42 Let $\hat{\mathbf{w}} = [\hat{w}_1, \dots, \hat{w}_k]^\top \in \mathbb{R}^k$ be the intermediate result (Line 6) in Algorithm 2. Then,

$$\|V_{\tilde{\beta}} \hat{\mathbf{w}} - M(\phi)\|_\infty \leq 2^{O(k)} \cdot \xi^{\frac{1}{2}}.$$

Proof Firstly,

$$\begin{aligned} \|V_{\tilde{\beta}} \hat{\mathbf{w}} - M(\phi)\|_2 &\leq \|V_{\tilde{\beta}} \hat{\mathbf{w}} - M_{G'}\|_2 + \|M(\phi) - M_{G'}\|_2 \\ &\leq \min_{\mathbf{x} \in \mathbb{C}^k} \|V_{\tilde{\beta}} \mathbf{x} - M_{G'}\|_2 + \|M(\phi) - M_{G'}\|_2 \\ &\leq \min_{\mathbf{x} \in \mathbb{C}^k} \|V_{\tilde{\beta}} \mathbf{x} - M(\phi)\|_2 + \|M(\phi) - M_{G'}\|_2 \end{aligned}$$

where the first and third inequalities hold due to triangle inequality, and the second inequality holds since $\|V_{\tilde{\beta}} \hat{\mathbf{w}} - M_{G'}\|_2^2 \leq \min_{\mathbf{x} \in \mathbb{R}^k} \|V_{\tilde{\beta}} \mathbf{x} - M_{G'}\|_2^2$.

For the first term, we can see that

$$\begin{aligned}
\min_{\mathbf{x} \in \mathbb{C}^k} \|V_{\tilde{\beta}} \mathbf{x} - M(\phi)\|_2 &= \min_{\mathbf{x}_1, \dots, \mathbf{x}_k \in \mathbb{C}^k} \left\| \sum_{i=1}^k w_i V_{\tilde{\beta}} \mathbf{x}_i - \sum_{i=1}^k w_i M(\beta_i) \right\|_2 \\
&\leq \sum_{i=1}^k w_i \min_{\mathbf{x} \in \mathbb{C}^k} \|V_{\tilde{\beta}} \mathbf{x} - M(\beta_i)\|_2 \\
&\leq 2^{O(k)} \sum_{i=1}^k w_i \prod_{j=1}^k |\beta_i - \tilde{\beta}_j| && \text{(Lemma 32)} \\
&\leq 2^{O(k)} \sqrt{\sum_{i=1}^k w_i \prod_{j=1}^k |\beta_i - \tilde{\beta}_j|^2} && \text{(AM-QM Inequality, } \mathbf{w} \in \Delta_{k-1}\text{)} \\
&\leq 2^{O(k)} \cdot \sqrt{\xi}. && \text{(Lemma 41).}
\end{aligned}$$

For the second term,

$$\begin{aligned}
\|M(\phi) - M_{G'}\|_2 &\leq \sqrt{k} \|M(\phi) - M_{G'}\|_\infty \\
&\leq \sqrt{k} \|G' - G(\phi)\|_\infty \\
&\leq 2^{O(k)} \cdot \xi. && \text{(Lemma 39)}
\end{aligned}$$

As a result,

$$\|V_{\tilde{\beta}} \hat{\mathbf{w}} - M(\phi)\|_\infty \leq \|V_{\tilde{\beta}} \hat{\mathbf{w}} - M(\phi)\|_2 \leq 2^{O(k)} \cdot \sqrt{\xi} + 2^{O(k)} \cdot \xi \leq 2^{O(k)} \cdot \xi^{\frac{1}{2}}.$$

■

Similar to Lemma 34, we can prove that $\tilde{\phi}$ is also a good estimation of ϕ .

Lemma 43 *Let $\tilde{\phi} = (\tilde{\beta}, \tilde{\mathbf{w}})$ be the intermediate result (Line 8) in Algorithm 2. Then,*

$$\text{Mom}_K(\tilde{\phi}, \phi) \leq 2^{O(k)} \cdot \xi^{\frac{1}{2}}.$$

Proof According to the definition of $V_{\tilde{\beta}}$ and $M(\phi)$, we have $(V_{\tilde{\beta}} \hat{\mathbf{w}} - M(\phi))_1 = \sum_{i=1}^k \hat{w}_i - 1$. Therefore, $|\sum_{i=1}^k \hat{w}_i - 1| \leq |(V_{\tilde{\beta}} \hat{\mathbf{w}} - M(\phi))_1| \leq \|V_{\tilde{\beta}} \hat{\mathbf{w}} - M(\phi)\|_\infty \leq 2^{O(k)} \cdot \xi^{\frac{1}{2}}$.

Note that $\|\tilde{\mathbf{w}} - \hat{\mathbf{w}}\|_1 = \|\tilde{\mathbf{w}}\|_1 - \|\hat{\mathbf{w}}\|_1 = |(\sum_{i=1}^k \hat{w}_i)^{-1} - 1|$. For $2^{O(k)} \cdot \xi^{\frac{1}{2}} \leq 1/2$, we can conclude that $\|\tilde{\mathbf{w}} - \hat{\mathbf{w}}\|_1 = |(\sum_{i=1}^k \hat{w}_i)^{-1} - 1| \leq 2 |\sum_{i=1}^k \hat{w}_i - 1| \leq 2^{O(k)} \cdot \xi^{\frac{1}{2}}$.

Thus,

$$\begin{aligned}
\text{Mom}_K(\tilde{\phi}, \phi) &= \|M(\tilde{\phi}) - M(\phi)\|_\infty = \|V_{\tilde{\beta}} \tilde{\mathbf{w}} - M(\phi)\|_\infty \\
&\leq \|V_{\tilde{\beta}} \hat{\mathbf{w}} - M(\phi)\|_\infty + \|V_{\tilde{\beta}}\|_\infty \|\tilde{\mathbf{w}} - \hat{\mathbf{w}}\|_1 \\
&\leq 2^{O(k)} \cdot \xi^{\frac{1}{2}}.
\end{aligned}$$

where the first inequality holds because of the triangle inequality. ■

Now everything is in place to show the final bound.

Lemma 44 *Let $\check{\boldsymbol{\vartheta}} = (\check{\boldsymbol{\alpha}}, \check{\boldsymbol{w}})$ be the final result (Line 11) in Algorithm 2. Then, we have*

$$\text{Tran}(\check{\boldsymbol{\vartheta}}, \boldsymbol{\vartheta}) \leq O(k\xi^{\frac{1}{4k-2}}).$$

Proof Apply Lemma 43 to the result of Lemma 22. We have

$$\text{Tran}(\check{\boldsymbol{\phi}}, \boldsymbol{\phi}) \leq O(k \cdot (2^{O(k)} \cdot \xi^{\frac{1}{2}})^{\frac{1}{2k-1}}) \leq O(k\xi^{\frac{1}{4k-2}}).$$

Denote $\check{\boldsymbol{\phi}}' = (\check{\boldsymbol{\beta}}, \text{real}(\check{\boldsymbol{w}}))$ as the result that discards all imaginary components from $\check{\boldsymbol{\phi}}$. From the definition of transportation distance for complex weights, the imaginary components is independent from the real components. Moreover, $\boldsymbol{\phi}$ has no imaginary components. Thus, discarding all imaginary components from $\check{\boldsymbol{\phi}}$ reduces the transportation distance to $\boldsymbol{\phi}$. That is,

$$\text{Tran}(\check{\boldsymbol{\phi}}', \boldsymbol{\phi}) \leq \text{Tran}(\check{\boldsymbol{\phi}}, \boldsymbol{\phi}).$$

Moreover, from any $\beta_i, \beta_j \in \mathbb{C}$, we have

$$\begin{aligned} \|[\text{real}(\beta_i), \text{imag}(\beta_i)] - [\text{real}(\beta_j), \text{imag}(\beta_j)]\|_1 &\leq 2\|[\text{real}(\beta_i), \text{imag}(\beta_i)] - [\text{real}(\beta_j), \text{imag}(\beta_j)]\|_2 \\ &= 2|\beta_i - \beta_j|. \end{aligned}$$

This shows the distance over \mathbb{C} is smaller than 2 times of the distance over \mathbb{R}^2 . As a result,

$$\text{Tran}(\check{\boldsymbol{\vartheta}}, \boldsymbol{\vartheta}) \leq 2 \text{Tran}(\check{\boldsymbol{\vartheta}}, \boldsymbol{\vartheta}) \leq O(k\xi^{\frac{1}{4k-2}}).$$

Moreover, from Lemma 35, we have

$$\text{Tran}(\check{\boldsymbol{\vartheta}}, \check{\boldsymbol{\vartheta}}) \leq 2 \min_{\check{\boldsymbol{\vartheta}} \in \text{Spike}(\Delta_2, \Delta_{k-1})} \text{Tran}(\check{\boldsymbol{\vartheta}}, \bar{\boldsymbol{\vartheta}}) \leq 2 \text{Tran}(\check{\boldsymbol{\vartheta}}, \boldsymbol{\vartheta}) = O(k\xi^{\frac{1}{4k-2}}).$$

Finally, from triangle inequality, $\text{Tran}(\check{\boldsymbol{\vartheta}}, \boldsymbol{\vartheta}) \leq \text{Tran}(\check{\boldsymbol{\vartheta}}, \check{\boldsymbol{\vartheta}}) + \text{Tran}(\check{\boldsymbol{\vartheta}}, \boldsymbol{\vartheta}) \leq O(k\xi^{\frac{1}{4k-2}})$. \blacksquare

Appendix E. An Algorithm for Higher-dimensional Problem

We assume the moment information is revealed by a noisy moment oracle. A moment oracle corresponding to k -spike distribution $\boldsymbol{\vartheta}$ is defined to be a function $M'(\cdot)$ that takes a matrix (or a vector) R with $\|R\|_\infty \leq 1$ and generates a noisy vector $M'(R)$ such that $\|M'(R) - M(\Pi_R(\boldsymbol{\vartheta}))\|_\infty \leq \xi$ and $M'_0(R) = M_0(\Pi_R(\boldsymbol{\vartheta})) = 1$ where $\xi \leq 2^{-\Omega(k)}$ is the noise level and M is the moment vector defined in (1).

Let $e_t \in \mathbb{R}^d$ be the unit vector that takes value 1 in only the t -th coordinate and 0 in the remaining coordinates. Denote $\mathbf{1}_d = (1, \dots, 1) \in \mathbb{R}^d$. Let \mathbb{S}_{d-1} be the unit sphere in \mathbb{R}^d .

The pseudocode can be found in Algorithm 3. The algorithm's input consists of the number of spikes k , the moment error bound ξ , and the noisy moment oracle $M'(\cdot)$ mentioned above. Now, we describe the implementation details of our algorithm.

We first generate a random vector \boldsymbol{r} in Line 2 by sampling from unit sphere $\mathbb{S}_{d-1} = \{\boldsymbol{r} \in \mathbb{R}^d \mid \|\boldsymbol{r}\|_2 = 1\}$. Note that by a probabilistic argument, the distance between spikes is roughly kept after the projection along \boldsymbol{r} (see Lemma 45). Then, we aim to use the one-dimension algorithm to

Algorithm 3 Reconstruction Algorithm in High Dimension

1: **input:** number of spikes k , noisy moments $M'(\cdot)$, noise level ξ
output: recovered spike distribution $\tilde{\vartheta}$

2: $\mathbf{r} \leftarrow \mathbb{S}_{d-1}$

3: $\tilde{\phi}' \leftarrow \text{OneDimension}(k, M'(\frac{\mathbf{r}+1}{4}), \xi) \quad \Rightarrow \tilde{\phi}' = (\tilde{\mathbf{y}}', \tilde{\mathbf{w}}') \in \text{Spike}(\Delta_1, \Delta_{k-1})$

4: $\tilde{\phi} \leftarrow ([4\tilde{\mathbf{y}}' - \mathbf{1}_d], \tilde{\mathbf{w}}') \quad \Rightarrow \tilde{\phi} = (\tilde{\mathbf{y}}, \tilde{\mathbf{w}})$

5: **for** $t \in [1, d]$ **do**

6: $\hat{\phi}'_t \leftarrow \text{TwoDimension}(k, M'([\frac{\mathbf{r}+1}{4}, \frac{e_t}{2}]), \xi) \quad \Rightarrow \hat{\phi}'_t = ([\hat{\mathbf{y}}'_{:,t}, \hat{\alpha}'_{:,t}], \hat{\mathbf{w}}'_{:,t}) \in \text{Spike}(\Delta_2, \Delta_{k-1})$

7: $\hat{\phi}_t \leftarrow ([4\mathbf{y}'_{:,t} - \mathbf{1}_d, 2\hat{\alpha}'_{:,t}], \hat{\mathbf{w}}'_{:,t}) \quad \Rightarrow \hat{\phi}_t = ([\hat{\mathbf{y}}_{:,t}, \hat{\alpha}_{:,t}], \hat{\mathbf{w}}_{:,t})$

8: **for** $j \in [1, k]$ **do**

9: $s_{j,t}^* \leftarrow \arg \min_{s: \hat{w}_{s,t} > \sqrt{\epsilon}} |\tilde{y}_j - \hat{y}_{s,t}|$

10: $\tilde{\alpha}_{j,t} \leftarrow \hat{\alpha}_{s_{j,t}^*, t}$

11: **end for**

12: $\tilde{\alpha}_{:,t} \leftarrow [\tilde{\alpha}_{1,t}, \dots, \tilde{\alpha}_{k,t}]^\top \quad \Rightarrow \tilde{\alpha} \in \mathbb{R}^k$

13: **end for**

14: $\tilde{\alpha} \leftarrow [\tilde{\alpha}_{:,1}, \dots, \tilde{\alpha}_{:,d}]^\top \quad \Rightarrow \tilde{\alpha} = [\tilde{\alpha}_1, \dots, \tilde{\alpha}_k] \in \mathbb{R}^{d \times k}$

15: $\tilde{\vartheta} \leftarrow (\tilde{\alpha}, \tilde{\mathbf{w}}) \quad \Rightarrow \tilde{\vartheta} \in \text{Spike}(\mathbb{R}^d, \Delta_{k-1})$

16: $\tilde{\alpha} \leftarrow \text{project}_{\Delta_{d-1}}(\tilde{\alpha}) \quad \Rightarrow \tilde{\alpha} = [\tilde{\alpha}_1, \dots, \tilde{\alpha}_k] \in \Delta_{d-1}^k$

17: $\tilde{\vartheta} \leftarrow (\tilde{\alpha}, \tilde{\mathbf{w}}) \quad \Rightarrow \tilde{\vartheta} \in \text{Spike}(\Delta_{d-1}, \Delta_{k-1})$

recover $\Pi_{\mathbf{r}}(\vartheta)$. However, the support of $\Pi_{\mathbf{r}}(\vartheta)$ is contained in $[-1, 1]$ but not $[0, 1]$. In Line 3, we apply the 1-dimension algorithm (Algorithm 1) to the noisy moments of a shifted and scaled map $\Pi_{(\mathbf{r}+1_d)/4}(\vartheta)$ whose support is in $[0, 1]$. Then we scale and shift back to obtain the result $\tilde{\phi}$ in Line 4. Intuitively, $\tilde{\phi}$ is close to $\Pi_{\mathbf{r}}(\vartheta)$.

Now, we try to recover the coordinates of the spikes. In particular, we would like to find, for each spike of $\tilde{\phi}$, the coordinates of the corresponding location in Δ_{d-1} . This is done by considering each dimension separately. For each coordinate $t \in [d]$, we run the two-dimension algorithm (i.e., Algorithm 2) to recovery the projection $\Pi_{[\mathbf{r}, e_t]}(\vartheta)$, that is a linear map of ϑ onto the subspace spanned by vector \mathbf{r} and axis e_t . However, $[\mathbf{r}, e_t]\alpha_i$ is in $[-1, 1] \times [0, 1]$. Again, in Line 6, we apply the algorithm to the noisy moments of a shifted and scaled projection $\Pi_{B_t}(\vartheta)$ where $B_t = [\frac{\mathbf{r}+1_d}{4}, \frac{e_t}{2}]$, and scale the result back in Line 7. The result is denoted as $\hat{\phi}_t$.

Next, we assemble the 1d projection $\tilde{\phi}$ and 2-d projections $\hat{\phi}_t, t \in [d]$. Due to the noise, $\tilde{\phi}_t$ and $\tilde{\phi}$ may have different supports along \mathbf{r} . So we assign the coordinates of $\tilde{\phi}$ according to the closest spike of $\hat{\phi}_t$ with relatively large weights in Line 7 to Line 12. Due to the error, the support of the reconstructed distribution $\tilde{\vartheta}$ may not locate in Δ_{d-1} . Thus the final step in Line 16 is to project each individual spike back to Δ_{d-1} .

We note the bottleneck of this algorithm is Line 6, where we have to solve the two-dimensional problem d times. Hence, the total running time of the algorithm is $O(dk^3)$.

E.1. Error Analysis

First, we show the random vector \mathbf{r} has a critical property that two distant spikes are still distant after the projection with high probability. This result is standard and similar statements are known in the literature (e.g., Lemma 37 in [Wu and Yang \(2020\)](#)).

Lemma 45 *Let $\text{Supp} = (\alpha_1, \dots, \alpha_k) \subset \Delta_{d-1}$ such that $d > 3$. For every constant $0 < \eta < 1$ and a random vector $\mathbf{r} \leftarrow \mathbb{S}_{d-1}$, with probability at least $1 - \eta$, for any $\alpha_i, \alpha_j \in \text{Supp}$, we have that*

$$|\mathbf{r}^\top(\alpha_i - \alpha_j)| \geq \frac{\eta}{k^2 d} \|\alpha_i - \alpha_j\|_1.$$

Proof Let $0 < \epsilon < 1$ be some constant depending on k and d which we fix later. Assume $\mathbf{r} = [r_1, \dots, r_d]^\top$. For a fixed pair (α_i, α_j) , we have

$$\Pr_{\mathbf{r} \sim \mathbb{S}_{d-1}} [|\mathbf{r}^\top(\alpha_i - \alpha_j)| < \epsilon \|\alpha_i - \alpha_j\|_2] = \Pr_{\mathbf{r} \sim \mathbb{S}_{d-1}} [|r_1| < \epsilon].$$

Suppose \mathbf{r} is obtained by first sampling a vector from unit ball \mathbb{B}_d and then normalizing it to the unit sphere \mathbb{S}_{d-1} . We can see that

$$\Pr_{\mathbf{r} \sim \mathbb{S}_{d-1}} [|r_1| < \epsilon] \leq \Pr_{\mathbf{r} \sim \mathbb{B}_d} [|r_1| < \epsilon]$$

where the RHS is equal to the probability that \mathbf{r} lies in the slice $[-\epsilon, \epsilon]$.

Let V_d be the volume of d -dimensional ball. We have

$$\Pr_{\mathbf{r} \sim \mathbb{B}_d} [|r_1| < \epsilon] = \frac{\int_{-\epsilon}^{\epsilon} V_{d-1}(1-x^2)^{\frac{d-1}{2}} dx}{V_d} \leq \epsilon \sqrt{d}$$

where the last inequality holds since $V_d/V_{d-1} \geq 2/\sqrt{d}$.

By union bound over all pairs (α_i, α_j) , we can see that the failure probability can be bounded by

$$\Pr_{\mathbf{r} \sim \mathbb{S}_{d-1}} [\exists \alpha_i, \alpha_j \in \text{Supp} : |\mathbf{r}^\top(A_i - A_j)| < \epsilon \|A_i - A_j\|_2] < k^2 \epsilon \sqrt{d}.$$

Notice that $\|\alpha\|_1 \leq \sqrt{d} \|\alpha\|_2$ for any $\alpha \in \mathbb{R}^d$, and let $\epsilon = \frac{\eta}{k^2 \sqrt{d}}$. We conclude that

$$\Pr_{\mathbf{r} \sim \mathbb{S}_{d-1}} [\forall \alpha_i, \alpha_j \in \text{Supp} : |\mathbf{r}^\top(A_i - A_j)| \geq \frac{\eta}{k^2 d} \|A_i - A_j\|_1] > 1 - \eta.$$

■

Next, we show that we can recover and reconstruct the higher dimension mixture ϑ , using its 1d projection $\tilde{\phi} \sim \Pi_{\mathbf{r}}(\vartheta)$, by assigning each spike $(\tilde{y}_i, \tilde{w}_i)$ in the one dimension distribution a location in \mathbb{R}^d . In a high level, consider a clustering of the spikes in $\Pi_{\mathbf{r}}(\vartheta)$. The spikes in one cluster still form a cluster in the original ϑ by Lemma 45 since the projection approximately keeps the distance. Thus assigning the clusters a location in \mathbb{R}^d can produce a good estimation for ϑ . The following lemma formally proves that we can reconstruct the original mixture using two-dimensional projections.

Lemma 46 Let $\tilde{\vartheta} \in \text{Spike}(\mathbb{R}^d, \Delta_{k-1})$ be the intermediate result (Line 15) in Algorithm 3. Let ϵ be the smallest number such that $\text{Tran}(\tilde{\phi}, \Pi_r(\vartheta)) \leq \epsilon$ and $\text{Tran}(\hat{\phi}_t, \Pi_{[r, e_t]}(\vartheta)) \leq \epsilon$ for all t . If Lemma 45 holds for ϑ , we have that

$$\text{Tran}(\tilde{\vartheta}, \vartheta) \leq O(k^3 d^2 \sqrt{\epsilon}).$$

Proof Let C_1, \dots, C_m be a partition of $[k]$ for ϑ such that:

- $|\mathbf{r}^\top \alpha_i - \mathbf{r}^\top \alpha_j| \leq k\sqrt{\epsilon}$ if $i \in C_l, j \in C_l$ for some $l \in [m]$.
- $|\mathbf{r}^\top \alpha_i - \mathbf{r}^\top \alpha_j| > \sqrt{\epsilon}$ if $i \in C_l, j \in C_{l'}$ for $l \neq l'$

For every cluster C_p , we construct another set \tilde{C}_p for $\tilde{\vartheta}$ containing every index $j \in [k]$ such that $\min_{i \in C_p} |\tilde{y}_j - \mathbf{r}^\top \alpha_i| < \frac{\sqrt{\epsilon}}{2}$. From the construction of the partition, for each spike \tilde{y}_j , there exists at most one $C_{p'}$ such that $\min_{i \in C_{p'}} |\tilde{y}_j - \mathbf{r}^\top \alpha_i| < \frac{\sqrt{\epsilon}}{2}$. Thus, we have all $\tilde{C}_1, \dots, \tilde{C}_m$ disjoint.

In this case, for every pair of $i \notin C_p$ and $j \in \tilde{C}_p$, we have $|\tilde{y}_j - \mathbf{r}^\top \alpha_i| > \frac{\sqrt{\epsilon}}{2}$. Since the total transportation distance between $\tilde{\phi}$ and $\Pi_r(\vartheta)$ is no more than ϵ . Hence, a portion of at most $\frac{\epsilon}{\sqrt{\epsilon}/2}$ of the weight in C_p can be matched out of \tilde{C}_p . This gives

$$\left| \sum_{i \in C_p} w_i - \sum_{j \in \tilde{C}_p} \tilde{w}_j \right| < 2\sqrt{\epsilon}.$$

Recall the transportation distance $\text{Tran}(\vartheta, \tilde{\vartheta})$ is the solution of the following program over coupling distribution μ .

$$\begin{aligned} \min \quad & \sum_{i=1}^k \sum_{j=1}^k \mu_{i,j} \|\tilde{\alpha}_j - \alpha_i\|_1 \\ \text{s.t.} \quad & w_i = \sum_{j=1}^k \mu_{i,j}, \tilde{w}_j = \sum_{i=1}^k \mu_{i,j}, \mu \geq \mathbf{0}. \end{aligned}$$

Consider a specific distribution μ generated by the following two-stage procedure: Initially, set $\mu_{i,j} = 0$ for every pair of (i, j) . In the first stage, increase $\mu_{i,j}$ for every $i \in C_p$ and $j \in \tilde{C}_p$ whenever $w_i \geq \sum_{j=1}^k \mu_{i,j}, \tilde{w}_j \geq \sum_{i=1}^k \mu_{i,j}$. In the second stage, increase $\mu_{i,j}$ arbitrarily for every pair of (i, j) to satisfy every constraint. Since the restriction set is convex, C_p are disjoint, and \tilde{C}_p are disjoint, for some specific p , this two-stage process ensures

$$\sum_{i \in C_p} \sum_{j \in \tilde{C}_p} \mu_{i,j} = \min \left\{ \sum_{i \in C_p} w_i, \sum_{j \in \tilde{C}_p} \tilde{w}_j \right\}.$$

As a result, the coupling distribution satisfies

$$\sum_{i \in C_p} \sum_{j \notin \tilde{C}_p} \mu_{i,j} \leq \left| \sum_{i \in C_p} w_i - \sum_{j \in \tilde{C}_p} \tilde{w}_j \right| \leq 2\sqrt{\epsilon}.$$

Now we start to show μ is a good coupling distribution for $\text{Tran}(\boldsymbol{\vartheta}, \tilde{\boldsymbol{\vartheta}})$. Since $\{C_p\}$ is a partition of $[k]$, we can decompose the transportation distance into:

$$\text{Tran}(\boldsymbol{\vartheta}, \tilde{\boldsymbol{\vartheta}}) \leq \sum_{i=1}^k \sum_{j=1}^k \mu_{i,j} \|\tilde{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_i\|_1 = \sum_p \sum_{i \in C_p} \sum_{j=1}^k \mu_{i,j} \|\tilde{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_i\|_1.$$

Next, we bound the RHS by considering each cluster C_p . If it holds $\sum_{i \in C_p} \sum_{j=1}^k \mu_{i,j} \leq 2k\sqrt{\epsilon}$, from $\|\tilde{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_i\|_1 \leq 1$, we can directly get

$$\sum_{i \in C_p} \sum_{j=1}^k \mu_{i,j} \|\tilde{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_i\|_1 \leq 2k\sqrt{\epsilon}.$$

Otherwise, we have $\sum_{i \in C_p} \sum_{j=1}^k \mu_{i,j} > 2k\sqrt{\epsilon}$. We can decompose the last summation into:

$$\sum_{i \in C_p} \sum_{j=1}^k \mu_{i,j} \|\tilde{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_i\|_1 = \sum_{i \in C_p} \sum_{j \notin \tilde{C}_p} \mu_{i,j} \|\tilde{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_i\|_1 + \sum_{i \in C_p} \sum_{j \in \tilde{C}_p} \mu_{i,j} \|\tilde{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_i\|_1.$$

For the first term, we have $\sum_{i \in C_p} \sum_{j \notin \tilde{C}_p} \mu_{i,j} \leq 2\sqrt{\epsilon}$ as we had discussed above. Together with $\|\tilde{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_i\|_1 \leq 1$, we can bound the first term by $2\sqrt{\epsilon}$. For the second term, we have $\sum_{i \in C_p} \sum_{j \in \tilde{C}_p} \mu_{i,j} \leq 1$ according to the property of coupling distribution and the following lemma shows $\|\tilde{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_i\|_1$ is small. We delay the proof for readability.

Lemma 47 *If $\sum_{i \in C_p} w_i > 2k\sqrt{\epsilon}$, then $\|\tilde{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_i\|_1 \leq O(k^3 d^2 \sqrt{\epsilon})$ for all $i \in C_p$ and $j \in \tilde{C}_p$.*

As a result, we can see that

$$\sum_{i \in C_p} \sum_{j=1}^k \mu_{i,j} \|\tilde{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_i\|_1 \leq O(k^3 d^2 \sqrt{\epsilon}).$$

Hence, we can conclude that $\text{Tran}(\boldsymbol{\vartheta}, \tilde{\boldsymbol{\vartheta}}) \leq O(k^3 d^2 \sqrt{\epsilon})$. ■

Proof [Proof of Lemma 47] For some fixed cluster C_p with $i \in C_p$ and $j \in \tilde{C}_p$, we will prove the statement by showing $\boldsymbol{\alpha}_i$ and $\tilde{\boldsymbol{\alpha}}_j$ are close on every dimension t .

Firstly, we claim that for the cluster C_p , there exist $s \in [k]$ and $i_1 \in C_p$ in which $\hat{w}_{s,t} \geq \sqrt{\epsilon}$ and $|\hat{y}_{s,t} - \mathbf{r}^\top \boldsymbol{\alpha}_{i_1}| + |\hat{\alpha}_{s,t} - \alpha_{i_1,t}| \leq \sqrt{\epsilon}$ in which $\hat{\phi}_t$ contains a spike at $(\hat{y}_{s,t}, \hat{\alpha}_{s,t})$ of weight $\hat{w}_{s,t}$ and $\alpha_{i_1,t}$ is the dimension t of $\boldsymbol{\alpha}_{i_1}$. Towards a contradiction, each spike in $\hat{\phi}_t$ is either of weight less than $\sqrt{\epsilon}$ or is of distance at least $\sqrt{\epsilon}$ from the projected spikes in C_p along $[\mathbf{r}, e_t]$. Then, each unit of the weights in C_p suffers a transportation cost of at least $\sqrt{\epsilon}$ after the first $k \times \sqrt{\epsilon}$ weights since it is not sufficient to cover spikes in C_p using nearby spikes. Note the total weight in C_p is at least $2k\sqrt{\epsilon}$. So in this case, the transportation distance between $\hat{\phi}_t$ and $\Pi_{[\mathbf{r}, e_t]}(\boldsymbol{\vartheta})$ is at least $\text{Tran}(\hat{\phi}_t, \Pi_{[\mathbf{r}, e_t]}(\boldsymbol{\vartheta})) \geq (2k\sqrt{\epsilon} - \sqrt{\epsilon} \times k) \times \sqrt{\epsilon} > k\epsilon$ which contradicts to the definition of ϵ .

Next, we take some $i_2 \in C_p$ for j such that $|\tilde{y}_j - \mathbf{r}^\top \boldsymbol{\alpha}_{i_2}| \leq \frac{\sqrt{\epsilon}}{2}$. This always exists according to the definition of \tilde{C}_p . From the clustering, it holds that $|\mathbf{r}^\top \boldsymbol{\alpha}_i - \mathbf{r}^\top \boldsymbol{\alpha}_{i_2}| \leq k\sqrt{\epsilon}$ and $|\mathbf{r}^\top \boldsymbol{\alpha}_{i_1} - \mathbf{r}^\top \boldsymbol{\alpha}_{i_2}| \leq k\sqrt{\epsilon}$. With the triangle inequality, we have

$$|\tilde{y}_j - \hat{y}_{s,t}| \leq |\hat{y}_{s,t} - \mathbf{r}^\top \boldsymbol{\alpha}_{i_1}| + |\mathbf{r}^\top \boldsymbol{\alpha}_{i_1} - \mathbf{r}^\top \boldsymbol{\alpha}_{i_2}| + |\tilde{y}_j - \mathbf{r}^\top \boldsymbol{\alpha}_{i_2}| \leq (1.5 + k)\sqrt{\epsilon}.$$

Since $\widehat{w}_{s,t} \geq \sqrt{\epsilon}$ which fits the condition in Line 9 of Algorithm 3, the minimization thus ensures $|\widetilde{y}_j - \widehat{y}_{s_{j,t}^*,t}| \leq (1.5 + k)\sqrt{\epsilon}$.

In addition, since the weight of spike $(\widehat{y}_{s_{j,t}^*,t}, \widehat{\alpha}_{s_{j,t}^*,t})$ satisfies $\widehat{w}_{s_{j,t}^*,t} > \sqrt{\epsilon}$ and the total transportation distance satisfies $\text{Tran}(\widehat{\phi}_t, \Pi_{[r,ct]}(\boldsymbol{\vartheta})) \leq \epsilon$. There must exist some $i_3 \in [k]$ such that $|\mathbf{r}^\top \boldsymbol{\alpha}_{i_3} - \widehat{y}_{s_{j,t}^*,t}| + |\alpha_{i_3,t} - \widehat{\alpha}_{s_{j,t}^*,t}| < \frac{\epsilon}{\sqrt{\epsilon}} = \sqrt{\epsilon}$, that is, $(\widehat{y}_{s_{j,t}^*,t}, \widehat{\alpha}_{s_{j,t}^*,t})$ is close to spike $(\mathbf{r}^\top \boldsymbol{\alpha}_{i_3}, \alpha_{i_3,t})$. From the triangle inequality, we have

$$\begin{aligned} & |\mathbf{r}^\top \boldsymbol{\alpha}_i - \mathbf{r}^\top \boldsymbol{\alpha}_{i_3}| + |\alpha_{i_3,t} - \widehat{\alpha}_{s_{j,t}^*,t}| \\ & \leq |\mathbf{r}^\top \boldsymbol{\alpha}_i - \mathbf{r}^\top \boldsymbol{\alpha}_{i_2}| + |\widetilde{y}_j - \mathbf{r}^\top \boldsymbol{\alpha}_{i_2}| + |\widetilde{y}_j - \widehat{y}_{s_{j,t}^*,t}| + |\mathbf{r}^\top \boldsymbol{\alpha}_{i_3} - \widehat{y}_{s_{j,t}^*,t}| + |\alpha_{i_3,t} - \widehat{\alpha}_{s_{j,t}^*,t}| \\ & \leq (3 + 2k)\sqrt{\epsilon}. \end{aligned}$$

Applying Lemma 45 for $\boldsymbol{\alpha}_i$ s, we have that $|\mathbf{r}^\top \boldsymbol{\alpha}_i - \mathbf{r}^\top \boldsymbol{\alpha}_{i_3}| \geq \frac{\eta}{k^2 d} \|\boldsymbol{\alpha}_i - \boldsymbol{\alpha}_{i_3}\|_1 \geq \frac{\eta}{k^2 d} |\alpha_{i,t} - \alpha_{i_3,t}|$ under the good event. Thus,

$$\begin{aligned} |\widetilde{\alpha}_{j,t} - \alpha_{i,t}| &= |\widehat{\alpha}_{s_{j,t}^*,t} - \alpha_{i,t}| \\ &\leq |\alpha_{i,t} - \alpha_{i_3,t}| + |\alpha_{i_3,t} - \widehat{\alpha}_{s_{j,t}^*,t}| \\ &\leq \frac{k^2 d}{\eta} (|\mathbf{r}^\top \boldsymbol{\alpha}_i - \mathbf{r}^\top \boldsymbol{\alpha}_{i_3}| + |\alpha_{i_3,t} - \widehat{\alpha}_{s_{j,t}^*,t}|) \\ &\leq \frac{k^2 d}{\eta} \cdot (3 + 2k)\sqrt{\epsilon} \leq O(k^3 d \sqrt{\epsilon}) \end{aligned}$$

where the equality holds due to the assignment of $\widetilde{\alpha}_{j,t} \leftarrow \widehat{\alpha}_{s_{j,t}^*,t}$ in Line 10 of Algorithm 3 and the second inequality holds since $\eta < 1$. Taking summation over t , we conclude that

$$\|\widetilde{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_i\|_1 \leq \sum_{t=1}^d |\widetilde{\alpha}_{j,t} - \alpha_{i,t}| \leq O(k^3 d^2 \sqrt{\epsilon}).$$

■

Note that we can bound ϵ using Lemma 36 and Lemma 44. As a result, we can provide a performance guarantee of our algorithm for the high-dimensional case.

Theorem 48 *Let $\check{\boldsymbol{\vartheta}} = (\check{\boldsymbol{\alpha}}, \check{\boldsymbol{w}})$ be the final result (Line 17) in Algorithm 3. Then with probability at least $1 - \eta$ for any fixed constant $0 < \eta < 1$,*

$$\text{Tran}(\check{\boldsymbol{\vartheta}}, \boldsymbol{\vartheta}) \leq (kd\xi^{\frac{1}{k}})^{O(1)}.$$

Proof Since $\check{\boldsymbol{\phi}}'$ is generated by one dimension algorithm, according to Lemma 36, we have

$$\text{Tran}(\check{\boldsymbol{\phi}}', \Pi_{\frac{r+1}{4}}(\boldsymbol{\vartheta})) \leq O(k\xi^{\frac{1}{4k-2}}).$$

Note that $\Pi_{\frac{r+1}{4}}(\boldsymbol{\vartheta})$ can be transformed to $\Pi_r(\boldsymbol{\vartheta})$ by transforming the coordinates by $\alpha \rightarrow 4\alpha - 1$. We have

$$\text{Tran}(\check{\boldsymbol{\vartheta}}, \Pi_r(\boldsymbol{\vartheta})) \leq 4 \text{Tran}(\check{\boldsymbol{\phi}}', \Pi_{\frac{r+1}{4}}(\boldsymbol{\vartheta})) \leq O(k\xi^{\frac{1}{4k-2}}).$$

Similarly, according to Lemma 44, we can see $\text{Tran}(\hat{\phi}_t, \Pi_{[r, \epsilon_t]}(\boldsymbol{\vartheta})) \leq O(k\xi^{\frac{1}{4k-2}})$. Thus, from the definition of ϵ , we have $\epsilon \leq O(k\xi^{\frac{1}{4k-2}})$. With Lemma 46, we have

$$\text{Tran}(\tilde{\boldsymbol{\vartheta}}, \boldsymbol{\vartheta}) \leq O(k^3 d^2 \sqrt{\epsilon}) \leq (kd\xi^{\frac{1}{k}})^{O(1)}.$$

Note that Δ_{d-1} is a convex set and $\check{\alpha}_j$ is the projection of $\tilde{\alpha}_j$ to the same set. Thus, $\|\alpha_i - \check{\alpha}_j\|_1 \leq \|\alpha_i - \tilde{\alpha}_j\|_1$ holds for all $A_i \in \Delta_{d-1}$ and j . Hence, the projection to Δ_{d-1} does not increase the transportation distance. As a result, $\text{Tran}(\check{\boldsymbol{\vartheta}}, \tilde{\boldsymbol{\vartheta}}) \leq \text{Tran}(\tilde{\boldsymbol{\vartheta}}, \boldsymbol{\vartheta})$. Finally, by triangle inequality, we have

$$\text{Tran}(\check{\boldsymbol{\vartheta}}, \boldsymbol{\vartheta}) \leq \text{Tran}(\tilde{\boldsymbol{\vartheta}}, \boldsymbol{\vartheta}) + \text{Tran}(\check{\boldsymbol{\vartheta}}, \tilde{\boldsymbol{\vartheta}}) \leq (kd\xi^{\frac{1}{k}})^{O(1)}.$$

■

By choosing $\xi = (\epsilon/(dk))^{O(k)}$, the previous theorem directly implies Theorem 5.

Appendix F. Applications to Topic Models

In this section, we discuss the application of our results to topic models. Here, the underlying vocabulary consists of d words. We are given a corpus of documents. We adopt the popular ‘‘bag of words’’ model and take each document as an unordered multiset of words. The assumption is that there is a small number k of ‘‘pure’’ topics, where each topic is a distribution over $[d]$. A K -word document (i.e., a string in $[d]^K$) is generated by first selecting a topic $i \in \Delta_{d-1}$ from the mixture $\boldsymbol{\vartheta}$, and then sampling K words i.i.d. according to α_i from this topic. Here, K is referred to as *the snapshot number* and we call such a sample a K -*snapshot* of p . Our goal here is to recover $\boldsymbol{\vartheta} \in \text{Spike}(\Delta_{d-1}, \Delta_{k-1})$, which is a discrete distribution over k pure topics. Again, we measure the accuracy in terms of L_1 -transportation distance.

F.1. An Algorithm for $d = O(k)$

In this section, we first directly handle the high-dimensional case using Algorithm 3. We will perform dimension reduction when $d \gg k$, in Section F.2.

Lemma 49 *There is an algorithm that can learn an arbitrary k -spike mixture of discrete distributions supported in Δ_{d-1} for any d within L_1 transportation distance ϵ with probability at least 0.99 using $(kd/\epsilon)^{O(k)}$ many $(2k-1)$ -snapshots.*

The following lemma is a high-dimensional version for empirical error (see e.g. Rabani et al. (2014); Li et al. (2015); Gordon et al. (2020)).

Lemma 50 *For matrix $R = [r_1, \dots, r_p] \in \mathbb{R}^{d \times p}$ with $\|R\|_\infty \leq 1$, using $2^{O(kp)} \cdot \xi^{-2} \cdot \log(1/\eta)$ samples, we can obtain transformed empirical moments $M'(R)$ satisfying*

$$\max_{|\mathbf{t}| \leq K} |M'_t(R) - M_t(\Pi_R(\boldsymbol{\vartheta}))| \leq \xi$$

with probability at least $1 - \eta$.

Proof Assume the words of a document are x_1, \dots, x_K . If the document is from some topic i , we have $\mathbb{E}[x_j = c] = \alpha_{p,t}$ for all i, j, c . For each word x_j , we generate a i.i.d. random vector $\beta_j = [\beta_{1,j}, \dots, \beta_{p,j}]^\top \in \{0, 1\}^p$ such that $\mathbb{E}[\beta_{c,j}] = r_{c,x_j}$. For $\mathbf{t} = (t_1, \dots, t_p)$ with $|\mathbf{t}| \leq K$, we define observable $h_{\mathbf{t}} = \mathbf{1}[\forall c: \sum_{j=1}^K \beta_{c,j} = t_c]$.

Note that $h_{\mathbf{t}}$ is the normalized histogram for projected distribution $\Pi_R(\vartheta)$. Hence, from the relation between normalized histogram and standard moments, we can conclude that

$$M(\Pi_R(\vartheta)) = \left(\bigotimes_{t=1}^p \text{Pas} \right) h$$

where $\text{Pas} \in \mathbb{R}^{(K+1) \times (K+1)}$ such that $\text{Pas}_{i,j} = [j \geq i] \binom{j}{i} \binom{K+1}{i}$ and \otimes is the tensor product. Since $\|\text{Pas}\| \leq 2^{O(k)}$, we have $(\bigotimes_{t=1}^p \text{Pas}) \leq 2^{O(kp)}$.

Let $\hat{h}_{\mathbf{t}}$ be the empirical average of $h_{\mathbf{t}}$. Since $h_{\mathbf{t}}$ is a Bernoulli variable, from Hoeffding's inequality, we have $\Pr[|\hat{h}_{\mathbf{t}} - \mathbb{E}[h_{\mathbf{t}}]| < \epsilon] \leq 2 \exp(-2\epsilon^2 s)$. By applying union bound on \mathbf{t} , we can see that using $s > 2^{O(kp)} \cdot \xi^{-2} \cdot \log(1/\eta)$ samples, we have $|\hat{h}_{\mathbf{t}} - \mathbb{E}[h_{\mathbf{t}}]| < \xi \cdot 2^{-\Omega(kp)}$ with probability at least $1 - \eta$. Moreover, if we calculate M' from \hat{h} using the relationship between M and h , this leads to $\max_{|\mathbf{t}| \leq K} |M'_{\mathbf{t}}(R) - M_{\mathbf{t}}(\Pi_R(\vartheta))| \leq \xi$. ■

Proof [Proof of Lemma 49] Let $\xi = (\epsilon/kd)^{O(k)}$. From Theorem 48, the recovered k -spike distribution $\check{\vartheta}$ satisfies $\text{Tran}(\check{\vartheta}, \vartheta) \leq (kd\xi^{\frac{1}{k}})^{O(1)} = O(\epsilon)$. According to Lemma 50, we can construct such a noisy moment oracle $M'(\cdot)$ using $(kd/\epsilon)^{O(k)}$ samples with high probability. ■

F.2. Dimension Reduction When $d \gg k^{\Omega(1)}$

In topic modeling, the number of words is typically much larger than the number of pure topics (i.e., $d \gg k$). In this case, we face another challenge in recovering the mixture. First, there are $\binom{d}{1} + \binom{d}{2} + \dots + \binom{d}{k} = O(d^k)$ many different moments. Obtaining all empirical moments accurately enough would require a huge number of $2k - 1$ -snapshot samples. So if $d \gg k$, we reduce the dimension from d to $O(k)$. Now, we prove the following theorem. It improves on the result in Rabani et al. (2014); Li et al. (2015), which uses more than $(k/\epsilon)^{O(k^2)}$ $(2k - 1)$ -snapshots, and Gordon et al. (2020) which use $(k/\epsilon)^{O(k)} \cdot (w_{\min}\zeta^k)^{-O(1)}$ $2k$ -snapshots and requires the minimum separation assumption.

Let ϑ be the target k -spike mixture in Δ_{d-1} ($d \gg k$). Suppose there is a learning algorithm \mathcal{A} satisfying the assumption in the theorem. We show how to apply \mathcal{A} to the projection of ϑ to a subspace of dimension at most k in \mathbb{R}^d . However, an arbitrary subspace of dimension at most k is not enough, even the one spanning the k spikes of ϑ . This is because L_1 distance is not rotationally invariant and not preserved under projection. For example, the L_1 distance between $(1/d, \dots, 1/d)$ and $(2/d, \dots, 2/d, 0, \dots, 0)$ is 1 in \mathbb{R}^d but only $\sqrt{1/d}$ in the line spanned by the two points. Hence, an accurate estimation of the projected measure $\vartheta_B = \Pi_B(\vartheta)$ may not translate to an accurate estimation in the ambient space \mathbb{R}^d . Here, we can use the dimension reduction method developed by Li et al. (2015), which shows that it suffices to project the mixture to a special subspace B , such that a unit L_1 -ball in B (L_1 measured in \mathbb{R}^d) is close to being an $\tilde{O}(d^{-1/2})$ L_2 -ball (\tilde{O} hides factors depending only on k and ϵ , but not d).

Lemma 51 (Li et al., 2015) *There is an algorithm that requires $\text{poly}(d, k, \frac{1}{\epsilon})$ many 1-, 2- snapshots to construct a subspace B ($\dim(B) \leq k$) with the following useful properties:*

1. Suppose $\{b_1, b_2, \dots, b_m\}$ is an orthonormal basis of B where $m = \dim(B) \leq k$ (such a basis can be produced by the above algorithm as well). For any $i \in [m]$, $\|b_i\|_\infty \leq O(k^{3/2}\epsilon^{-2}d^{-1/2})$.
2. Suppose we can learn an approximation $\tilde{\vartheta}_B$, supported on $\text{Span}(B)$, of the projected measure $\vartheta_B = \Pi_{\text{Span}(B)}(\vartheta)$ such that $\text{Tran}(\vartheta_B, \tilde{\vartheta}_B) \leq \epsilon_1 = \text{poly}(1/k, \epsilon)$ (here L_1 is measured in \mathbb{R}^d , not in the subspace) using $N_1(d)$, $N_2(d)$ and $N_K(d)$ 1-, 2-, and K -snapshot samples. Then there is an algorithm for learning a mixture $\tilde{\vartheta}$ such that $\text{Tran}(\vartheta, \tilde{\vartheta}) \leq \epsilon$ using $O(N_1(d/\epsilon) + d \log d/\epsilon^3)$, $O(N_2(d/\epsilon) + O(k^4 d^3 \log n/\epsilon^6))$ and $O(N_K(d/\epsilon))$ 1-, 2-, and K -snapshot samples respectively.

Proof [Proof of Theorem 6] In light of Lemma 51, we only need to show how to learn the projection $\vartheta_B = \Pi_{\text{Span}(B)}(\vartheta)$ using the algorithm developed in Section F.1.

First, we show how to translate a K -snapshots from ϑ to a K -snapshot in some new k -dimensional mixture ϑ_{new} (then we apply the algorithm in Section F.1 for dimension k). Let $L = \sum_{i=1}^m \|b_i\|_\infty$. So $L \leq O(k^{5/2}\epsilon^{-2}d^{-1/2})$. Let $f : [-L, L] \rightarrow [0, 1]$ defined as $f(x) = \frac{x}{2mL} + \frac{1}{2m}$. For each word in the original K -snapshot, say $i \in [d]$, we translate it to j with probability $q_{i,j} = f(b_{j,i})$, for each $j \in [m]$, where $b_{j,i}$ denotes the i th coordinate of b_j , and translate it into $m+1$ with probability $q_{i,m+1} = 1 - \sum_{j=1}^m q_{i,j}$. Since we have $\sum_{j=1}^m f(b_{j,i}) = \frac{1}{2mL} \sum_{j=1}^m b_{j,i} + \frac{1}{2}$, we know that $q_{i,1}, q_{i,2}, \dots, q_{i,m+1} \in [0, 1]$ and $\sum_{j=1}^{k+1} q_{i,j} = 1$. So this is a well defined k -spike mixture ϑ_{new} and its K snapshots. Then we apply our learning algorithm \mathcal{A} for dimension k to obtain some ϑ'_{new} such that $\text{Tran}(\vartheta_{\text{new}}, \vartheta'_{\text{new}}) \leq \epsilon_1 = \text{poly}(1/k, \epsilon)$ with $\text{poly}(k, 1/\epsilon_1) = \text{poly}(k, 1/\epsilon)$ many $2k-1$ -snapshots.

Next we can see ϑ_{new} (in \mathbb{R}^k) and ϑ_B (in $\text{Span}(B)$) are related by an affine transform. For a spike $\alpha_i \in \text{Supp}(\vartheta)$, one can see that $\Pi_B(\alpha_i) = \sum_{j=1}^m a_{i,j} b_j$, where $a_{i,j} = \langle \alpha_i, b_j \rangle$. α_i produces a new spike β_i in ϑ_{new} and the mapping is as follows: for each $j \in [k]$, we have

$$\beta_{i,j} = \sum_{t=1}^d \alpha_{i,t} f(b_{j,t}) = \sum_{t=1}^d \alpha_{i,t} \left(\frac{b_{j,t}}{2mL} + \frac{1}{2m} \right) = \frac{a_{i,j}}{2mL} + \frac{1}{2m}$$

($\beta_{i,j}$ denotes the j th coordinate of β_i). So $a_{i,j} = 2mL(\beta_{i,j} - \frac{1}{2m})$ and $g(x) = 2mLx - L$ is the affine transformation. Now we can translate ϑ'_{new} into a mixture in $\text{Span}(B)$ by applying $g(\cdot)$ in each coordinate of each $\beta_i \in \text{Supp}(\vartheta'_{\text{new}})$, and obtain an estimation of ϑ_B , say ϑ'_B .

It remains to show $\text{Tran}(\vartheta'_B, \vartheta_B) \leq \epsilon_1$. If we let Tran_B denote the transportation distance in $\text{Span}(B)$ (where if $a = \sum_{j=1}^k a_j b_j, c = \sum_{j=1}^k c_j b_j$ then $\text{Tran}_B(a, c) = \sum_{j=1}^k |a_j - c_j|$). Then $\text{Tran}_B(\vartheta'_B, \vartheta_B) \leq 2mL \text{Tran}(\vartheta'_{\text{new}}, \vartheta_{\text{new}}) = \text{poly}(\epsilon, \frac{1}{k})/\sqrt{d}$. Hence, we have

$$\text{Tran}(\vartheta'_B, \vartheta_B) \leq d \max_i \|b_i\|_\infty \text{Tran}_B(\vartheta'_B, \vartheta_B) \leq \epsilon_1.$$

where the first inequality holds since $\text{Tran}(a, c) = \|\sum_{j=1}^k (a_j - c_j) b_j\|_1 \leq d \max_i \|b_i\|_\infty \sum_{j=1}^k (a_j - c_j)$ for any two point $a, c \in \text{Span}(B)$. \blacksquare

Appendix G. Applications to Gaussian Mixture Learning

In this section, we show how to leverage our results for sparse moment problem to obtain improved algorithms for learning Gaussian mixtures. We consider the following setting studied in Wu and

Yang (2020); Doss et al. (2020). A k -Gaussian mixture in \mathbb{R}^d can be parameterised as $\vartheta_N = (\boldsymbol{\alpha}, \boldsymbol{w}, \Sigma)$. Here, $\boldsymbol{\alpha} = \{\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2, \dots, \boldsymbol{\alpha}_k\}$ and $\boldsymbol{w} = \{w_1, w_2, \dots, w_k\} \in \Delta_{k-1}$ where $\boldsymbol{\alpha}_i \in \mathbb{R}^d$ specifies the mean of i th component and $w_i \in [0, 1]$ is the corresponding weight. $\Sigma \in \mathbb{R}^{d \times d}$ is the common covariance matrix for all k mixture components and we assume that Σ is known in advance. We further assume $\|\boldsymbol{\alpha}_i\|_2 \leq 1$ and the maximal eigenvalue $\|\Sigma\|_2$ is bounded by a constant. For k -Gaussian mixture $\vartheta_N = (\boldsymbol{\alpha}, \boldsymbol{w}, \Sigma)$, each observation is distributed as

$$\vartheta_N \sim \sum_{i=1}^k w_i N(\boldsymbol{\alpha}_i, \Sigma).$$

We consider the parameter learning problem, that is, to learn the parameter $\boldsymbol{\alpha}$ and \boldsymbol{w} given known covariance matrix Σ and a set of i.i.d. samples from ϑ_N . The model is also called *Gaussian location mixture model* (Wu and Yang, 2020; Doss et al., 2020).⁵

G.1. Efficient Algorithm for $d = 1$

In the 1-dimensional case, we denote the known variance by σ which is upper bounded by some constant. As shown by Wu and Yang (2020), the moments of Gaussian mixture have a close connection with the moments of corresponding discrete distributions $\vartheta = (\boldsymbol{\alpha}, \boldsymbol{w})$. For $x \sim N(\mu, 1)$, we have $\mathbb{E}[H_t(x)] = \mu^t$ for Hermite polynomial $H_t(x)$ which is defined as

$$H_t(x) = \sum_{i=0}^t h_{t,i} x^i = t! \sum_{j=0}^{\lfloor \frac{t}{2} \rfloor} \frac{(-1/2)^j}{j!(t-2j)!} x^{t-2j}. \quad (6)$$

For 1-dimensional Gaussian mixture $\vartheta_N = (\boldsymbol{\alpha}, \boldsymbol{w}, \sigma)$ with variance σ , the t th moment of the discrete distribution $\vartheta = (\boldsymbol{\alpha}, \boldsymbol{w})$ satisfies

$$M_t(\vartheta) = \mathbb{E}_{x \sim \vartheta_N} \left[\sum_{i=0}^t h_{t,i} \sigma^{t-i} x^i \right].$$

The following lemma guarantees the performance of estimating the moment by sampling.

Lemma 52 (Lemma 5 of Wu and Yang (2020), restated) *Let $x_1, \dots, x_n \sim \vartheta_N$ be n independent samples.*

$$\widetilde{M}_t(\vartheta) = \frac{1}{n} \sum_{j=1}^n \left[\sum_{i=0}^t h_{t,i} \sigma^{t-i} x_j^i \right]$$

is an unbiased estimator for $M_t(\vartheta)$, and the variance of this estimator can be bounded by

$$\text{Var}[\widetilde{M}_t(\vartheta)] \leq \frac{1}{n} (\sigma t)^{O(t)}.$$

Thus we can compute the moment of the original distribution and use our Algorithm 1 for recovering the parameter of ϑ . More concretely, we can replace the last two lines (an SDP) in Algorithm 2 of Wu and Yang (2020) by our Algorithm 1, and obtain the following theorem. The post-sampling time is improved from from $O(k^{2\omega})$ in Wu and Yang (2020) to $O(k^2)$.

5. Wu and Yang (2020) also studied the problem with unknown Σ . We leave it as a future direction.

Theorem 53 *Let ϑ_N be an arbitrary k -Gaussian mixture over \mathbb{R} with means $\alpha_1, \dots, \alpha_n$ and known variance σ bounded by some constant. There is an algorithm that can learn the parameter $\vartheta = (\alpha, w)$ within transportation distance $O(\epsilon)$, with probability at least 0.99, using $(k/\epsilon)^{O(k)}$ samples from the mixture ϑ_N . Moreover, once we obtain the estimation of the moments of ϑ_N , our algorithm only uses $O(k^2)$ arithmetic operations.*

Proof Let c be the constant to be determined. With $n = (k/\epsilon)^{ck}$ samples, the variance of the empirical moment can be bounded by

$$\text{Var}[M'_t(\vartheta)] \leq \frac{1}{(k/\epsilon)^{ck}} \cdot k^{O(k)} \leq (\epsilon/k)^{\Omega(k)}$$

where the first inequality holds due to Lemma 52 and σ is bounded by a constant, and the second inequality holds by selecting a large enough constant $c > 0$. According to Chebyshev's inequality, for each t , with probability at least $1 - 0.01k^{-2}$,

$$|M'_t(\vartheta) - M_t(\vartheta)| \leq 0.1k \cdot (\epsilon/k)^{\Omega(k)} \leq (\epsilon/k)^{\Omega(k)}.$$

By taking union bound, the probability that the above inequality holds for all $0 \leq t \leq 2k - 1$ is greater than 0.99. We can conclude the result by applying Theorem 3 directly. \blacksquare

Remark 54 (*Connection between heat equations and Gaussian mixtures*) Consider the heat equation

$$\frac{\partial u(x, t)}{\partial t} = \Delta_x u(x, t), \quad \text{where} \quad \Delta_x u = \sum_i \frac{\partial^2 u}{\partial x_i^2}.$$

Suppose at $t = 0$, $u(x, 0)$ is a k -spike distribution $u(x, 0) = \sum_{i=1}^k w_i \delta_{\alpha_i}(x)$. It is well known that for $t > 0$, $u(x, t) = \sum_{i=1}^k w_i N(\alpha_i, 2t)$ (see e.g., [Ang et al. \(2002\)](#); [Moitra and Valiant \(2010\)](#)). Hence, if we know $u(\cdot, t)$ at time $t > 0$, we can compute the moments and recovering the heat source $u(\cdot, 0)$ at $t = 0$ is equivalent to recovering the means of the Gaussian mixture components (i.e., the spikes) from the moment information, hence can be solved by the moment problem.

G.2. Efficient Algorithm for $d > 1$

For higher dimensional Gaussian mixture, we can reduce the problem to learning the discrete mixture $\vartheta = (\alpha, w)$, and leverage Algorithm 3 to solve the problem. Assume the locations of all Gaussians are in the unit ball. The error can be bounded easily according to Theorem 5.

Firstly, as shown in [Doss et al. \(2020\)](#), one can use SVD to reduce a d -dimensions problem to a k -dimension problem using $\text{poly}(1/\epsilon, d, k)$ samples. Thus, we only need to consider settings with $d \leq k$. Similar to the 1d case, we can transform the problem of learning Gaussian mixture $\vartheta_N = (\alpha, w, \Sigma)$ to the problem of learning the discrete mixture $\vartheta = (\alpha, w)$. In particular, as we show below, we can estimate the moments of the projection of ϑ from the samples. After obtaining the noisy moment information, we can apply Algorithm 3 to recover the parameter (α, w) . We mention that we need to modify Algorithm 3 slightly: we change to domain from Δ_{d-1} to the unit ball (it suffices to change the projected domain in Line 16 in Algorithm 3 from the simplex to the unit ball). We note this does not affect the proof to Theorem 5 since it only requires the projected space to be a convex domain.

The remaining task is to show how to estimate the projected moments for 1d and 2d projections. To estimate the moments of 1-dimensional projection, we can use the estimator in the last section. In particular, for the projected measure $\Pi_{\mathbf{r}}(\boldsymbol{\vartheta})$ where \mathbf{r} is an arbitrary unit vector, the t th moment can be computed by

$$M_t(\Pi_{\mathbf{r}}(\boldsymbol{\vartheta})) = \mathbb{E}_{\mathbf{x} \sim \boldsymbol{\vartheta}_N} \left[\sum_{i=0}^t h_{t,i}(\mathbf{r}^\top \Sigma \mathbf{r})^{t-i} (\mathbf{r}^\top \mathbf{x})^i \right].$$

We can use sample average to estimate $M_t(\Pi_{\mathbf{r}}(\boldsymbol{\vartheta}))$, and we denote the estimation as $\widetilde{M}_t(\Pi_{\mathbf{r}}(\boldsymbol{\vartheta}))$.

Now, we consider the problem of estimating the moments for 2-dimensional projection along \mathbf{r}_1 and \mathbf{r}_2 . First, we compute the vector $\mathbf{r}'_2 = \mathbf{r}_2 - \frac{\mathbf{r}_1^\top \Sigma \mathbf{r}_2}{\mathbf{r}_1^\top \Sigma \mathbf{r}_1} \mathbf{r}_1$. According to Gram–Schmidt process, \mathbf{r}'_2 is Σ -orthogonal to \mathbf{r}_1 , i.e., $\mathbf{r}_1^\top \Sigma \mathbf{r}'_2 = 0$. Moreover, $\mathbf{r}_1^\top \mathbf{x}$ and $\mathbf{r}'_2{}^\top \mathbf{x}$ are independent Gaussian distributions for variable $\mathbf{x} \sim N(0, \Sigma)$. Note that for independent Gaussian variables $x_1 \sim N(\mu_1, 1)$ and $x_2 \sim N(\mu_2, 1)$, we have $\mathbb{E}[H_{t_1}(x_1)H_{t_2}(x_2)] = \mathbb{E}[H_{t_1}(x_1)]\mathbb{E}[H_{t_2}(x_2)] = \mu_1^{t_1} \mu_2^{t_2}$, where H is the Hermite polynomial defined in (6). As a result, the projected moments along \mathbf{r}_1 and \mathbf{r}'_2 can be computed by

$$M_{(t_1, t_2)}(\Pi_{[\mathbf{r}_1, \mathbf{r}'_2]}(\boldsymbol{\vartheta})) = \mathbb{E}_{\mathbf{x} \sim \boldsymbol{\vartheta}_N} \left[\sum_{i_1=0}^{t_1} \sum_{i_2=0}^{t_2} h_{t_1, i_1} h_{t_2, i_2} (\mathbf{r}_1^\top \Sigma \mathbf{r}_1)^{t_1-i_1} (\mathbf{r}'_2{}^\top \Sigma \mathbf{r}'_2)^{t_2-i_2} (\mathbf{r}_1^\top \mathbf{x})^{i_1} (\mathbf{r}'_2{}^\top \mathbf{x})^{i_2} \right].$$

Moreover, the projected moments along \mathbf{r}_1 and \mathbf{r}_2 can be computed by

$$M_{(t_1, t_2)}(\Pi_{[\mathbf{r}_1, \mathbf{r}_2]}(\boldsymbol{\vartheta})) = \sum_{i=0}^{t_2} \binom{t_2}{i} \left(\frac{\mathbf{r}_1^\top \Sigma \mathbf{r}_2}{\mathbf{r}_1^\top \Sigma \mathbf{r}_1} \right)^i M_{(t_1+i, t_2-i)}(\Pi_{[\mathbf{r}_1, \mathbf{r}'_2]}(\boldsymbol{\vartheta}))$$

since $\mathbf{r}_2^\top \mathbf{x} = \mathbf{r}'_2{}^\top \mathbf{x} + \frac{\mathbf{r}_1^\top \Sigma \mathbf{r}_2}{\mathbf{r}_1^\top \Sigma \mathbf{r}_1} \mathbf{r}_1^\top \mathbf{x}$ which implies $(\mathbf{r}_2^\top \mathbf{x})^{t_2} = \sum_{i=0}^{t_2} \binom{t_2}{i} \left(\frac{\mathbf{r}_1^\top \Sigma \mathbf{r}_2}{\mathbf{r}_1^\top \Sigma \mathbf{r}_1} \right)^i (\mathbf{r}_1^\top \mathbf{x})^i (\mathbf{r}'_2{}^\top \mathbf{x})^{t_2-i}$.

We use $\widetilde{M}_{(t_1, t_2)}(\Pi_{[\mathbf{r}_1, \mathbf{r}_2]}(\boldsymbol{\vartheta}))$ to denote the sample average estimation of $M_{(t_1, t_2)}(\Pi_{[\mathbf{r}_1, \mathbf{r}_2]}(\boldsymbol{\vartheta}))$. We are ready to state the performance guarantee:

Proof [Proof of Theorem 7] We first consider the problems of dimension $d \leq k$. The only thing we need to show is that the above estimators $\widetilde{M}_t(\Pi_{\mathbf{r}}(\boldsymbol{\vartheta}))$ and $\widetilde{M}_t(\Pi_{[\mathbf{r}_1, \mathbf{r}_2]}(\boldsymbol{\vartheta}))$ have sufficient accuracy with high probability using $n = (k/\epsilon)^{\Theta(k)}$ samples. Clearly, both estimators are unbiased. So it is sufficient to bound the variance of the estimators.

For 1-dimensional projected moments, we can use the same argument as Theorem 53. In particular, we can show that with probability at least 0.999, for all $0 \leq t \leq 2k - 1$,

$$|\widetilde{M}_t(\Pi_{\mathbf{r}}(\boldsymbol{\vartheta})) - M_t(\Pi_{\mathbf{r}}(\boldsymbol{\vartheta}))| \leq (k/\epsilon)^{O(k)}.$$

For 2-dimensional projected moments, we first bound the variance of the estimator along \mathbf{r}_1 and \mathbf{r}'_2 . With the same argument as Lemma 52, we have

$$\text{Var}[\widetilde{M}_{(t_1, t_2)}(\Pi_{[\mathbf{r}_1, \mathbf{r}'_2]}(\boldsymbol{\vartheta}))] \leq (\epsilon/k)^{\Omega(k)}.$$

for all $t_1 + t_2 \leq 2k$ with probability at least 0.999. Moreover, since we randomly chooses the same \mathbf{r}_1 over a sphere of radius $\Theta(1)$ in Line 6 of Algorithm 3, we have $\mathbf{r}_1^\top \Sigma \mathbf{r}_1 \geq \Omega(\frac{\mu}{k^2} \|\mathbf{r}_1\|_2^2 \|\Sigma\|_2)$ with probability at least $1 - \mu$. In this case, we have

$$\frac{\mathbf{r}_1^\top \Sigma \mathbf{r}_2}{\mathbf{r}_1^\top \Sigma \mathbf{r}_1} \leq \frac{\|\mathbf{r}_1\|_2 \|\Sigma\|_2 \|\mathbf{r}_2\|_2}{\Omega(\frac{\mu}{k^2} \|\mathbf{r}_1\|_2^2 \|\Sigma\|_2)} \leq O(k^2/\mu)$$

where the last inequality holds because \mathbf{r}_1 and \mathbf{r}_2 are of norm $\Theta(1)$. Since $\sqrt{\text{Var}[x + y]} \leq \sqrt{\text{Var}[x]} + \sqrt{\text{Var}[y]}$ for any random variables x and y ,

$$\sqrt{\text{Var}[\widetilde{M}_{(t_1, t_2)}(\Pi_{[\mathbf{r}_1, \mathbf{r}_2]}(\boldsymbol{\vartheta}))]} \leq \sum_{i=0}^{t_2} \binom{t_2}{i} \left(\frac{\mathbf{r}_1^\top \Sigma \mathbf{r}_2}{\mathbf{r}_1^\top \Sigma \mathbf{r}_1} \right)^i \sqrt{\text{Var}[\widetilde{M}_{(t_1+i, t_2-i)}(\Pi_{[\mathbf{r}_1, \mathbf{r}_2]}(\boldsymbol{\vartheta}))]}.$$

As a result, by choosing $\mu = 0.001$, we have

$$\text{Var}[\widetilde{M}_{(t_1, t_2)}(\Pi_{[\mathbf{r}_1, \mathbf{r}_2]}(\boldsymbol{\vartheta}))] \leq (\epsilon/k)^{\Omega(k)}$$

holds for all \mathbf{r}_2 with probability at least 0.999. Conditioning on that this event holds, according to Chebyshev's inequality and a union bound, with probability at least $1 - 0.001d$, for all $0 \leq t_1, t_2 \leq 2k - 1$,

$$|\widetilde{M}_{(t_1, t_2)}(\Pi_{[\mathbf{r}_1, \mathbf{r}_2]}(\boldsymbol{\vartheta})) - M_{(t_1, t_2)}(\Pi_{[\mathbf{r}_1, \mathbf{r}_2]}(\boldsymbol{\vartheta}))| \leq (\epsilon/k)^{\Omega(k)}.$$

Recall that the algorithm uses 1d projection for one \mathbf{r} , 2d projection for one \mathbf{r}_1 and d different \mathbf{r}_2 and the algorithm succeeds with high probability statement for random \mathbf{r} . Hence, by union bound over all these events, the projected moment estimations have sufficient accuracy with a probability at least 0.995.

If the dimension $d > k$, we can apply the dimension reduction in [Doss et al. \(2020\)](#), which requires an extra sample complexity of $\text{poly}(1/\epsilon, d, k)$. The total sample complexity is $(k/\epsilon)^{O(k)} + \text{poly}(1/\epsilon, d, k)$. \blacksquare

Finally, we discuss the running time during the sampling phase and post-sampling phase. The dimension reduction requires $\text{poly}(1/\epsilon, d, k)$ samples and $O(d^3)$ time ([Doss et al., 2020](#)). For the recovery problem in dimension k , each 1-d or 2-d moment oracle can be computed in $O(nk^3)$ time where $n = (k/\epsilon)^{O(k)}$ is the number of samples. Since we only require $O(k)$ moment oracles, the sampling time can be bounded by $O(nk^4)$. This improves the $O(n^{5/4} \text{poly}(k))$ sampling time in [Doss et al. \(2020\)](#) (their algorithm requires 1d projections to $O(n^{1/4})$ many directions. For the post-sampling running time, our Algorithm 3 runs in time $O(k^4)$ (since $d \leq k$ by dimension reduction). This improves the $O(n^{1/2} \text{poly}(k))$ post-sampling time in [Doss et al. \(2020\)](#).

Appendix H. Other Related Work

The problem we study can be seen as a sparse version of the classic *moment problem* in which our goal is to invert the mapping that takes a measure to the sequences of moments ([Schmüdgen, 2017](#); [Lasserre, 2009](#)). When the measure is supported on a finite interval, the problem is known as the Hausdorff moment problem.

Learning statistical mixture models has been studied extensively for the past two decades. A central problem in this area was the problem of learning a mixture of high-dimensional Gaussians, even robustly ([Dasgupta, 1999](#); [Dasgupta and Schulman, 2007](#); [Arora and Kannan, 2005](#); [Vempala and Wang, 2002](#); [Kannan et al., 2008](#); [Achlioptas and McSherry, 2005](#); [Feldman et al., 2006](#); [Brubaker and Vempala, 2008](#); [Kalai et al., 2010](#); [Belkin and Sinha, 2010](#); [Moitra and Valiant, 2010](#); [Liu and Moitra, 2021](#); [Bakshi et al., 2022](#); [Wu and Yang, 2020](#); [Doss et al., 2020](#)). Many other structured mixture models have also been studied (see e.g., [Kearns et al. \(1994\)](#); [Cryan et al. \(2001\)](#); [Batu et al. \(2004\)](#); [Mossel and Roch \(2005\)](#); [Dasgupta et al. \(2005\)](#); [Feldman et al. \(2008\)](#); [Kannan](#)

et al. (2008); Chaudhuri and Rao (2008); Daskalakis et al. (2012); Liu and Moitra (2018)). Our problem is closely related to topic models which have also been studied extensively recently (Arora et al., 2012; Anandkumar et al., 2012a,b; Arora et al., 2018). Most work make certain assumptions on the structure of the mixture, such as the pure topics being separated, Dirichlet prior, the existence of anchor words, or certain rank conditions. Some assumptions (such as Arora et al. (2012); Anandkumar et al. (2012a)) allow one to use documents of constant length (independent of the number of pure topics k and the desired accuracy ϵ) for recovering all topics.

Our problem is also related to the super-resolution problem in which each measurement takes the form $v_\ell = \int e^{2\pi\ell t} d\vartheta(t) + \xi_\ell$ where ξ_ℓ is the noise of the measurement. We can observe the first few v_ℓ for $|\ell| \leq K$ where K is called *cutoff frequency*, and the goal is to recover the original signal ϑ . There is a long history of the estimation problem. The noiseless version can be solved by Prony’s method (de Prony, 1795), ESPRIT algorithm (Roy and Kailath, 1989) or matrix pencil method (Hua and Sarkar, 1990), if $K \geq k$ (i.e., we have $K = 2k + 1$ measurements). The noisy case has also been studied extensively and a central goal is to understand the relations between the cutoff frequency, the size of measure noises and minimum separation (Donoho, 1992; Candès and Fernandez-Granda, 2014, 2012; Moitra, 2015; Huang and Kakade, 2015; Chen et al., 2016). Various properties, such as the condition number of the Vandermonde matrix, also play essential roles in this line of study (Moitra, 2015). The relation between the Vandermonde matrix and Schur polynomial is also exploited in Chen et al. (2016).