
Estimating Mixture Models via Mixtures of Polynomials

Sida I. Wang Arun Tejasvi Chaganty Percy Liang
Computer Science Department, Stanford University, Stanford, CA, 94305
{sidaw, chaganty, pliangu}@cs.stanford.edu

Abstract

Mixture modeling is a general technique for making any simple model more expressive through weighted combination. This generality and simplicity in part explains the success of the Expectation Maximization (EM) algorithm, in which updates are easy to derive for a wide class of mixture models. However, the likelihood of a mixture model is non-convex, so EM has no known global convergence guarantees. Recently, method of moments approaches offer global guarantees for some mixture models, but they do not extend easily to the range of mixture models that exist. In this work, we present Polymom, an unifying framework based on method of moments in which estimation procedures are easily derivable, just as in EM. Polymom is applicable when the moments of a single mixture component are polynomials of the parameters. Our key observation is that the moments of the mixture model are a mixture of these polynomials, which allows us to cast estimation as a Generalized Moment Problem. We solve its relaxations using semidefinite optimization, and then extract parameters using ideas from computer algebra. This framework allows us to draw insights and apply tools from convex optimization, computer algebra and the theory of moments to study problems in statistical estimation. Simulations show good empirical performance on several models.

1 Introduction

Mixture models play a central role in machine learning and statistics, with diverse applications including bioinformatics, speech, natural language, and computer vision. The idea of mixture modeling is to explain data through a weighted combination of simple parametrized distributions [1, 2]. In practice, maximum likelihood estimation via Expectation Maximization (EM) has been the workhorse for these models, as the parameter updates are often easily derivable. However, EM is well-known to suffer from local optima. The method of moments, dating back to Pearson [3] in 1894, is enjoying a recent revival [4, 5, 6, 7, 8, 9, 10, 11, 12, 13] due to its strong global theoretical guarantees. However, current methods depend strongly on the specific distributions and are not easily extensible to new ones.

In this paper, we present a method of moments approach, which we call Polymom, for estimating a wider class of mixture models in which the moment equations are polynomial equations (Section 2). Solving general polynomial equations is NP-hard, but our key insight is that for mixture models, the moments equations are *mixtures of polynomials equations* and we can hope to solve them if the moment equations for each mixture component are *simple* polynomials equations that we can solve. Polymom proceeds as follows: First, we recover mixtures of monomials of the parameters from the data moments by solving an instance of the Generalized Moment Problem (GMP) [14, 15] (Section 3). We show that for many mixture models, the GMP can be solved with basic linear algebra and in the general case, can be approximated by an SDP in which the moment equations are linear

mixture model		sizes	
\mathbf{x}_t	data point (\mathbb{R}^D)	D	data dimensions
z_t	latent mixture component ($[K]$)	K	mixture components
$\boldsymbol{\theta}_k$	parameters of component k (\mathbb{R}^P)	P	parameters of mixture components
π_k	mixing proportion of $p(z = k)$	T	data points
$[\boldsymbol{\theta}_k]_{k=1}^K$	all model parameters	N	constraints
moments of data		$[N]$	$\{1, \dots, N\}$
$\phi_n(\mathbf{x})$	observation function	r	degree of the moment matrix
$f_n(\boldsymbol{\theta})$	observation function	$s(r)$	size of the degree r moment matrix
moments of parameters		polynomials	
\mathcal{L}_y	the Riesz linear functional	$\mathbb{R}[\boldsymbol{\theta}]$	polynomial ring in variables $\boldsymbol{\theta}$
y_α	$y_\alpha = \mathcal{L}_y(\boldsymbol{\theta}^\alpha)$, α^{th} moment	\mathbb{N}	set of non-negative integers
μ	probability measure for \mathbf{y}	$\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}$	vector of exponents (in \mathbb{N}^P or \mathbb{N}^D)
\mathbf{y}	$(y_\alpha)_\alpha$ the moment sequence	$\boldsymbol{\theta}^\alpha$	monomial $\prod_{p=1}^P \theta_p^{\alpha_p}$
$\mathbf{M}_r(\mathbf{y})$	moment matrix of degree r	$a_{n\alpha}$	coefficient of $\boldsymbol{\theta}^\alpha$ in $f_n(\boldsymbol{\theta})$

Table 1: **Notation:** We use lowercase letters (e.g., d) for indexing, and the corresponding uppercase letter to denote the upper limit (e.g., D , in “sizes”). We use lowercase letters (e.g., $\theta_{k,p}$) for scalars, lowercase bold letters (e.g., $\boldsymbol{\theta}$) for vectors, and bold capital letters (e.g., \mathbf{M}) for matrices.

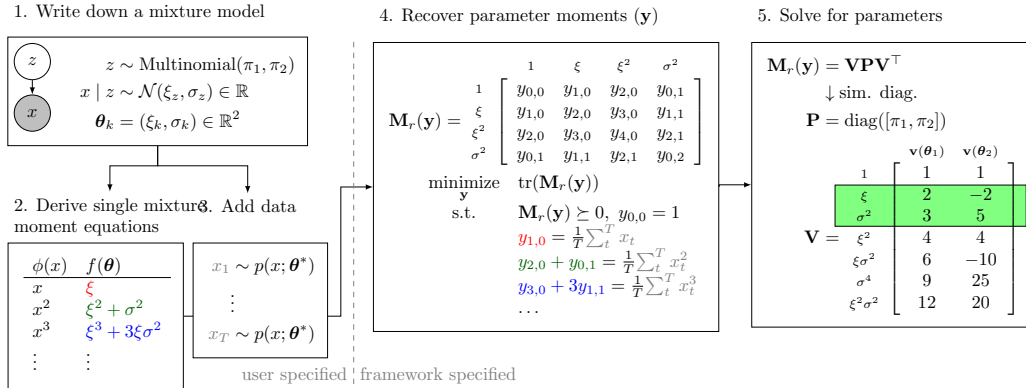


Figure 1: An overview of applying the Polymom framework.

constraints. Second, we extend multiplication matrix ideas from the computer algebra literature [16, 17, 18, 19] to extract the parameters by solving a generalized eigenvalue problem (Section 4).

Polymom improves on previous method of moments approaches in both generality and flexibility. First, while tensor factorization has been the main driver for many of the method of moments approaches for many types of mixture models, [6, 20, 9, 8, 21, 12], each model required specific adaptations which are non-trivial even for experts. In contrast, Polymom provides a unified principle for tackling new models that is as turnkey as computing gradients or EM updates. To use Polymom (Figure 1), one only needs to provide a list of observation functions (ϕ_n) and derive their expected values expressed symbolically as polynomials in the parameters of the specified model (f_n). Polymom then estimates expectations of ϕ_n and outputs parameter estimates of the specified model. Since Polymom works in an optimization framework, we can easily incorporate constraints such as non-negativity and parameter tying which is difficult to do in the tensor factorization paradigm. In simulations, we compared Polymom with EM and tensor factorization and found that Polymom performs similarly or better (Section 5).

2 Problem formulation

2.1 The method of moments estimator

In a mixture model, each data point $\mathbf{x} \in \mathbb{R}^D$ is associated with a latent component $z \in [K]$:

$$z \sim \text{Multinomial}(\pi), \quad \mathbf{x} | z \sim p(\mathbf{x}; \boldsymbol{\theta}_z^*), \quad (1)$$

where $\pi = (\pi_1, \dots, \pi_K)$ are the mixing coefficients, $\boldsymbol{\theta}_k^* \in \mathbb{R}^P$ are the true model parameters for the k^{th} mixture component, and $\mathbf{x} \in \mathbb{R}^D$ is the random variable representing data. We restrict our attention to mixtures where each component distribution comes from the same parameterized family. For example, for a mixture of Gaussians, $\boldsymbol{\theta}_k^* = (\xi_k^* \in \mathbb{R}^D, \Sigma_k^* \in \mathbb{R}^{D \times D})$ consists of the mean and covariance of component k .

We define N observation functions $\phi_n : \mathbb{R}^D \rightarrow \mathbb{R}$ for $n \in [N]$ and define $f_n(\boldsymbol{\theta})$ to be the expectation of ϕ_n over a single component with parameters $\boldsymbol{\theta}$, which we assume is a simple polynomial:

$$f_n(\boldsymbol{\theta}) := \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x}; \boldsymbol{\theta})}[\phi_n(\mathbf{x})] = \sum_{\boldsymbol{\alpha}} a_{n\boldsymbol{\alpha}} \boldsymbol{\theta}^{\boldsymbol{\alpha}}, \quad (2)$$

where $\boldsymbol{\theta}^{\boldsymbol{\alpha}} = \prod_{p=1}^P \theta_p^{\alpha_p}$. The expectation of each observation function $\mathbb{E}[\phi_n(\mathbf{x})]$ can then be expressed as a mixture of polynomials of the true parameters $\mathbb{E}[\phi_n(\mathbf{x})] = \sum_{k=1}^K \pi_k \mathbb{E}[\phi_n(\mathbf{x}) | z = k] = \sum_{k=1}^K \pi_k f_n(\boldsymbol{\theta}_k^*)$.

The method of moments for mixtures seeks parameters $[\boldsymbol{\theta}_k]_{k=1}^K$ that satisfy the *moment conditions*

$$\mathbb{E}[\phi_n(\mathbf{x})] = \sum_{k=1}^K \pi_k f_n(\boldsymbol{\theta}_k). \quad (3)$$

where $\mathbb{E}[\phi_n(\mathbf{x})]$ can be estimated from the data: $\frac{1}{T} \sum_{t=1}^T \phi_n(\mathbf{x}_t) \xrightarrow{p} \mathbb{E}[\phi_n(\mathbf{x})]$. The goal of this work is to find parameters satisfying moment conditions that can be written in the mixture of polynomial form (3). We assume that the N given observations functions ϕ_1, \dots, ϕ_N uniquely identify the model parameters (up to permutation of the components).

Example 2.1 (1-dimensional Gaussian mixture). Consider a K -mixture of 1D Gaussians with parameters $\boldsymbol{\theta}_k = [\xi_k, \sigma_k^2]$ corresponding to the mean and variance, respectively, of the k -th component (Figure 1: steps 1 and 2). We choose the observation functions, $\phi(x) = [x^1, \dots, x^6]$, which have corresponding moment polynomials, $f(\boldsymbol{\theta}) = [\xi, \xi^2 + \sigma^2, \xi^3 + 3\xi\sigma^2, \dots]$. For example, instantiating (3), $\mathbb{E}[x^2] = \sum_{k=1}^K \pi_k (\xi_k^2 + \sigma_k^2)$. Given $\phi(\mathbf{x})$ and $f(\boldsymbol{\theta}^*)$, and data, the Polynom framework can recover the parameters. Note that the 6 moments we use have been shown by [3] to be sufficient for a mixture of two Gaussians.

Example 2.2 (Mixture of linear regressions). Consider a mixture of linear regressions [22, 9], where each data point $\mathbf{x} = [x, y]$ is drawn from component k by sampling x from an unknown distribution independent of k and setting $y = w_k x + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma_k^2)$. The parameters $\boldsymbol{\theta}_k = (w_k, \sigma_k^2)$ are the slope and noise variance for each component k . Let us take our observation functions to be $\phi(\mathbf{x}) = [x, xy, xy^2, x^2, \dots, x^3 y^2]$, for which the moment polynomials are $f(\boldsymbol{\theta}) = [\mathbb{E}[x], \mathbb{E}[x^2]w, \mathbb{E}[x^3]w^2 + \mathbb{E}[x]\sigma^2, \mathbb{E}[x^2], \dots]$.

In Example 2.1, the coefficients $a_{n\boldsymbol{\alpha}}$ in the polynomial $f_n(\boldsymbol{\theta})$ are just constants determined by integration. For the conditional model in Example 2.2, the coefficients depends on the data. However, we cannot handle arbitrary data dependence, see Appendix D for sufficient conditions and counterexamples.

2.2 Solving the moment conditions

Our goal is to recover model parameters $\boldsymbol{\theta}_1^*, \dots, \boldsymbol{\theta}_K^* \in \mathbb{R}^P$ for each of the K components of the mixture model that generated the data as well as their respective mixing proportions $\pi_1, \dots, \pi_K \in \mathbb{R}$. To start, let's ignore sampling noise and identifiability issues and suppose that we are given exact moment conditions as defined in (3). Each condition $f_n \in \mathbb{R}[\boldsymbol{\theta}]$ is a polynomial of the parameters $\boldsymbol{\theta}$, for $n = 1, \dots, N$.

Equation 3 is a polynomial system of N equations in the $K + K \times P$ variables $[\pi_1, \dots, \pi_K]$ and $[\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K] \in \mathbb{R}^{P \times K}$. It is natural to ask if standard polynomial solving methods can solve (3) in the case where each $f_n(\boldsymbol{\theta})$ is simple. Unfortunately, the complexity of general polynomial equation solving is lower bounded by the number of solutions, and each of the $K!$ permutations of the mixture components corresponds to a distinct solution of (3) under this polynomial system representation. While several methods can take advantage of symmetries in polynomial systems [23, 24], they still cannot be adapted to tractably solve (3) to the best of our knowledge.

The key idea of Polymom is to exploit the mixture representation of the moment equations (3). Specifically, let μ^* be a particular ‘‘mixture’’ over the component parameters $\boldsymbol{\theta}_1^*, \dots, \boldsymbol{\theta}_k^*$ (i.e. μ^* is a probability measure). Then we can express the moment conditions (3) in terms of μ^* :

$$\mathbb{E}[\phi_n(\mathbf{x})] = \int f_n(\boldsymbol{\theta}) \mu^*(d\boldsymbol{\theta}), \text{ where } \mu^*(\boldsymbol{\theta}) = \sum_{k=1}^K \pi_k \delta(\boldsymbol{\theta} - \boldsymbol{\theta}_k^*). \quad (4)$$

As a result, solving the original moment conditions (3) is equivalent to solving the following feasibility problem over μ , but where we deliberately ‘‘forget’’ the permutation of the components by using μ to represent the problem:

$$\begin{aligned} \text{find } & \mu \in \mathcal{M}_+(\mathbb{R}^P), \text{ the set of probability measures over } \mathbb{R}^P \\ \text{s.t. } & \int f_n(\boldsymbol{\theta}) \mu(d\boldsymbol{\theta}) = \mathbb{E}[\phi_n(\mathbf{x})], \quad n = 1, \dots, N \\ & \mu \text{ is } K\text{-atomic (i.e. sum of } K \text{ deltas).} \end{aligned} \quad (5)$$

If the true model parameters $[\boldsymbol{\theta}_k^*]_{k=1}^K$ can be identified by the N observed moments up to permutation, then the measure $\mu^*(\boldsymbol{\theta}) = \sum_{k=1}^K \pi_k \delta(\boldsymbol{\theta} - \boldsymbol{\theta}_k^*)$ solving **Problem 5** is also unique.

Polymom solves **Problem 5** in two steps:

1. **Moment completion (Section 3):** We show that **Problem 5** over the measure μ can be relaxed to an SDP over a certain (parameter) *moment matrix* $\mathbf{M}_r(\mathbf{y})$ whose optimal solution is $\mathbf{M}_r(\mathbf{y}^*) = \sum_{k=1}^K \pi_k \mathbf{v}_r(\boldsymbol{\theta}_k^*) \mathbf{v}_r(\boldsymbol{\theta}_k^*)^\top$, where $\mathbf{v}_r(\boldsymbol{\theta}_k^*)$ is the vector of all monomials of degree at most r .
2. **Solution extraction (Section 4):** We then take $\mathbf{M}_r(\mathbf{y})$ and construct a series of generalized eigendecomposition problems, whose eigenvalues yield $[\boldsymbol{\theta}_k^*]_{k=1}^K$.

Remark. From this point on, distributions and moments refer to μ^* which is over *parameters*, not over the data. All the structure about the data is captured in the moment conditions (3).

3 Moment completion

The first step is to reformulate **Problem 5** as an instance of the Generalized Moment Problem (GMP) introduced by [15]. A reference on the GMP, algorithms for solving GMPs, and its various extensions is [14]. We start by observing that **Problem 5** really only depends on the integrals of monomials under the measure μ : for example, if $f_n(\boldsymbol{\theta}) = 2\theta_1^3 - \theta_1^2\theta_2$, then we only need to know the integrals over the constituent monomials ($y_{3,0} := \int \theta_1^3 \mu(d\boldsymbol{\theta})$ and $y_{2,1} := \int \theta_1^2 \theta_2 \mu(d\boldsymbol{\theta})$) in order to evaluate the integral over f_n . This suggests that we can optimize over the (parameter) *moment sequence* $\mathbf{y} = (y_\alpha)_{\alpha \in \mathbb{N}^P}$, rather than the measure μ itself. We say that the moment sequence \mathbf{y} has a *representing measure* μ if $y_\alpha = \int \boldsymbol{\theta}^\alpha \mu(d\boldsymbol{\theta})$ for all α , but we do not assume that such a μ exists. The *Riesz linear functional* $\mathcal{L}_\mathbf{y} : \mathbb{R}[\boldsymbol{\theta}] \rightarrow \mathbb{R}$ is defined to be the linear map such that $\mathcal{L}_\mathbf{y}(\boldsymbol{\theta}^\alpha) := y_\alpha$ and $\mathcal{L}_\mathbf{y}(1) = 1$. For example, $\mathcal{L}_\mathbf{y}(2\theta_1^3 - \theta_1^2\theta_2 + 3) = 2y_{3,0} - y_{2,1} + 3$. If \mathbf{y} has a representing measure μ , then $\mathcal{L}_\mathbf{y}$ simply maps polynomials f to integrals of f against μ .

The key idea of the GMP approach is to convexify the problem by treating \mathbf{y} as free variables and then introduce constraints to guarantee that \mathbf{y} has a representing measure. First, let $\mathbf{v}_r(\boldsymbol{\theta}) := [\boldsymbol{\theta}^\alpha : |\alpha| \leq r] \in \mathbb{R}[\boldsymbol{\theta}]^{s(r)}$ be the vector of all $s(r)$ monomials of degree no greater than r . Then, define the *truncated moment matrix* as $\mathbf{M}_r(\mathbf{y}) := \mathcal{L}_\mathbf{y}(\mathbf{v}_r(\boldsymbol{\theta}) \mathbf{v}_r(\boldsymbol{\theta})^\top)$, where the linear functional $\mathcal{L}_\mathbf{y}$ is applied elementwise (see **Example 3.1** below). If \mathbf{y} has a representing measure μ , then $\mathbf{M}_r(\mathbf{y})$ is simply a (positive) integral over rank 1 matrices $\mathbf{v}_r(\boldsymbol{\theta}) \mathbf{v}_r(\boldsymbol{\theta})^\top$ with respect to μ , so necessarily

$\mathbf{M}_r(\mathbf{y}) \succeq 0$ holds. Furthermore, by [Theorem 1](#) [25], for \mathbf{y} to have a K -atomic representing measure, it is sufficient that $\text{rank}(\mathbf{M}_r(\mathbf{y})) = \text{rank}(\mathbf{M}_{r-1}(\mathbf{y})) = K$. So [Problem 5](#) is equivalent to

$$\begin{aligned} \text{find } & \mathbf{y} \in \mathbb{R}^N \quad (\text{or equivalently, find } \mathbf{M}(\mathbf{y})) \\ \text{s.t. } & \sum_{\alpha} a_{n\alpha} y_{\alpha} = \mathbb{E}[\phi_n(\mathbf{x})], \quad n = 1, \dots, N \\ & \mathbf{M}_r(\mathbf{y}) \succeq 0, \quad y_0 = 1 \\ & \text{rank}(\mathbf{M}_r(\mathbf{y})) = K \text{ and } \text{rank}(\mathbf{M}_{r-1}(\mathbf{y})) = K. \end{aligned} \quad (6)$$

Unfortunately, the rank constraints in [Problem 6](#) are not tractable. We use the following relaxation to obtain our final (convex) optimization problem

$$\begin{aligned} \text{minimize } & \text{tr}(\mathbf{C}\mathbf{M}_r(\mathbf{y})) \\ \text{s.t. } & \sum_{\alpha} a_{n\alpha} y_{\alpha} = \mathbb{E}[\phi_n(\mathbf{x})], \quad n = 1, \dots, N \\ & \mathbf{M}_r(\mathbf{y}) \succeq 0, \quad y_0 = 1 \end{aligned} \quad (7)$$

where $\mathbf{C} \succ 0$ is a chosen scaling matrix. A common choice is $\mathbf{C} = \mathbf{I}_{s(r)}$ corresponding to minimizing the nuclear norm of the moment matrix, the usual convex relaxation for rank. [Appendix A](#) discusses some other choices of \mathbf{C} .

Example 3.1 (moment matrix for a 1-dimensional Gaussian mixture). *Recall that the parameters $\boldsymbol{\theta} = [\xi, \sigma^2]$ are the mean and variance of a one dimensional Gaussian. Let us choose the monomials $\mathbf{v}_2(\boldsymbol{\theta}) = [1, \xi, \xi^2, \sigma^2]$. Step 4 for [Figure 1](#) shows the moment matrix when using $r = 2$. Each row and column of the moment matrix is labeled with a monomial and entry (i, j) is subscripted by the product of the monomials in row i and column j . For $\phi_2(x) = x^2$, we have $f_2(\boldsymbol{\theta}) = \xi^2 + c$, which leads to the linear constraint $y_{2,0} + y_{0,1} - \mathbb{E}[x^2] = 0$. For $\phi_3(x) = x^3$, $f_3(\boldsymbol{\theta}) = \xi^3 + 3\xi c$, leading to the constraint $y_{3,0} + 3y_{1,1} - \mathbb{E}[x^3] = 0$.*

Related work. Readers familiar with the sum of squares and polynomial optimization literature [26, 27, 28, 29] will note that [Problem 7](#) is similar to the SDP relaxation of a polynomial optimization problem. However, in typical polynomial optimization, we are only interested in solutions $\boldsymbol{\theta}^*$ that actually satisfy the given constraints, whereas here we are interested in K solutions $[\boldsymbol{\theta}_k^*]_{k=1}^K$, whose mixture satisfies constraints corresponding to the moment conditions (3). Within machine learning, generalized PCA has been formulated as a moment problem [30] and the Hankel matrix (basically the moment matrix) has been used to learn weighted automata [13]. While similar tools are used, the conceptual approach and the problems considered are different. For example, the moment matrix of this paper consists of unknown moments of the model parameters, whereas existing works considered moments of the data that are always directly observable.

Constraints. Constraints such as non-negativity (for parameters which represent probabilities or variances) and parameter tying [31] are quite common in graphical models and are not easily addressed with existing method of moments approaches. The GMP framework allows us to incorporate some constraints using localizing matrices [32]. Consider the case of a 2D mixture of Gaussians where the mean parameters ξ_1, ξ_2 lies on the parabola $\xi_1 - \xi_2^2 = 0$ for all components. In this case, we just need to add constraints to [Problem 7](#): $y_{(1,0)+\beta} - y_{(0,2)+\beta} = 0$ for all $\beta \in \mathbb{N}^2$ up to degree $|\beta| \leq 2r - 2$. Thus, we can handle constraints during the estimation procedure rather than projecting back onto the constraint set as a post-processing step. This is necessary for models that only become identifiable by the observed moments after constraints are taken into account. By incorporating these constraints into parameter estimation, we can possibly identify the model parameters with fewer moments. We describe this method and its learning implications in [Appendix C.1](#).

Guarantees and statistical efficiency. In some circumstances, e.g. in three-view mixture models or the mixture of linear regressions, the constraints fully determine the moment matrix; we consider these cases in [Section 5](#) and [Appendix B](#). While there are no general guarantee on [Problem 7](#), the flat extension theorem tells us when the moment matrix corresponds to a unique solution (more discussions in [Appendix A](#)):

Theorem 1 (Flat extension theorem [25]). *Let \mathbf{y} be the solution to [Problem 7](#) for a particular r . If $\mathbf{M}_r(\mathbf{y}) \succeq 0$ and $\text{rank}(\mathbf{M}_{r-1}(\mathbf{y})) = \text{rank}(\mathbf{M}_r(\mathbf{y}))$ then \mathbf{y} is the optimal solution to [Problem 6](#) for $K = \text{rank}(\mathbf{M}_r(\mathbf{y}))$ and there exists a unique K -atomic supporting measure μ of $\mathbf{M}_r(\mathbf{y})$.*

Recovering $\mathbf{M}_r(\mathbf{y})$ is linearly dependent on small perturbations of the input [33], suggesting that the method has polynomial sample complexity for most models where the moments concentrate at a polynomially rate. In [Appendix C](#), we discuss a few other important considerations like noise robustness, making [Problem 7](#) more statistically efficient, and some open problems.

4 Solution extraction

Having completed the (parameter) moment matrix $\mathbf{M}_r(\mathbf{y})$ ([Section 3](#)), we now turn to the problem of extracting the model parameters $[\boldsymbol{\theta}_k^*]_{k=1}^K$. The solution extraction method we present is based on ideas from solving multivariate polynomial systems where the solutions are eigenvalues of certain multiplication matrices [16, 17, 34, 35].¹ The main advantage of the solution extraction view is that higher-order moments and structure in parameters are handled in the framework without model-specific effort.

Recall that the true moment matrix is $\mathbf{M}_r(\mathbf{y}^*) = \sum_{k=1}^K \pi_k \mathbf{v}(\boldsymbol{\theta}_k^*) \mathbf{v}(\boldsymbol{\theta}_k^*)^\top$, where $\mathbf{v}(\boldsymbol{\theta}) := [\boldsymbol{\theta}^{\alpha_1}, \dots, \boldsymbol{\theta}^{\alpha_{s(r)}}] \in \mathbb{R}[\boldsymbol{\theta}]^{s(r)}$ contains all the monomials up to degree r . We use $\boldsymbol{\theta} = [\theta_1, \dots, \theta_P]$ for variables and $[\boldsymbol{\theta}_k^*]_{k=1}^K$ for the true solutions to these variables (note the boldface). For example, $\theta_{k,p}^* := (\boldsymbol{\theta}_k^*)_p$ denotes the p^{th} value of the k^{th} component, which corresponds to a solution for the variable θ_p . Typically, $s(r) \gg K, P$ and the elements of $\mathbf{v}(\boldsymbol{\theta})$ are arranged in a degree ordering so that $\|\boldsymbol{\alpha}_i\|_1 \leq \|\boldsymbol{\alpha}_j\|_1$ for $i \leq j$. We can also write $\mathbf{M}_r(\mathbf{y}^*) = \mathbf{V}\mathbf{P}\mathbf{V}^\top$, where $\mathbf{V} := [\mathbf{v}(\boldsymbol{\theta}_1^*), \dots, \mathbf{v}(\boldsymbol{\theta}_K^*)] \in \mathbb{R}^{s(r) \times K}$ is the *canonical basis* and $\mathbf{P} := \text{diag}(\pi_1, \dots, \pi_K)$ contains the mixing proportions. At the high level, we want to factorize $\mathbf{M}_r(\mathbf{y}^*)$ to get \mathbf{V} , however we cannot simply eigen-decompose $\mathbf{M}_r(\mathbf{y}^*)$ since \mathbf{V} is not orthogonal. To overcome this challenge, we will exploit the internal structure of \mathbf{V} to construct several other matrices that share the same factors and perform simultaneous diagonalization.

Specifically, let $\mathbf{V}[\boldsymbol{\beta}_1; \dots; \boldsymbol{\beta}_K] \in \mathbb{R}^{K \times K}$ be a sub-matrix of \mathbf{V} with only the rows corresponding to monomials with exponents $\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_K \in \mathbb{N}^P$. Typically, $\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_K$ are just the first K monomials in \mathbf{v} . Now consider the exponent $\boldsymbol{\gamma}_p \in \mathbb{N}^P$ which is 1 in position p and 0 elsewhere, corresponding to the monomial $\boldsymbol{\theta}^{\boldsymbol{\gamma}_p} = \theta_p$. The key property of the canonical basis is that multiplying each column k by a monomial $\theta_{k,p}^*$ just performs a ‘‘shift’’ to another set of rows:

$$\mathbf{V}[\boldsymbol{\beta}_1; \dots; \boldsymbol{\beta}_K] \mathbf{D}_p = \mathbf{V}[\boldsymbol{\beta}_1 + \boldsymbol{\gamma}_p; \dots; \boldsymbol{\beta}_K + \boldsymbol{\gamma}_p], \quad \text{where } \mathbf{D}_p := \text{diag}(\theta_{1,p}^*, \dots, \theta_{K,p}^*). \quad (8)$$

Note that \mathbf{D}_p contains the p^{th} parameter for all K mixture components.

Example 4.1 (Shifting the canonical basis). *Let $\boldsymbol{\theta} = [\theta_1, \theta_2]$ and the true solutions be $\boldsymbol{\theta}_1^* = [2, 3]$ and $\boldsymbol{\theta}_2^* = [-2, 5]$. To extract the solution for θ_1 (which are $(\theta_{1,1}^*, \theta_{2,1}^*)$), let $\boldsymbol{\beta}_1 = (1, 0)$, $\boldsymbol{\beta}_2 = (1, 1)$, and $\boldsymbol{\gamma}_1 = (1, 0)$.*

$$\mathbf{V} = \begin{matrix} & \mathbf{v}(\boldsymbol{\theta}_1) & \mathbf{v}(\boldsymbol{\theta}_2) \\ \begin{matrix} 1 \\ \theta_1 \\ \theta_2 \\ \theta_1^2 \\ \theta_1\theta_2 \\ \theta_2^2 \\ \theta_1^2\theta_2 \end{matrix} & \begin{bmatrix} 1 & 1 \\ 2 & -2 \\ 3 & 5 \\ 4 & 4 \\ 6 & -10 \\ 9 & 25 \\ 12 & 20 \end{bmatrix} \end{matrix} \quad \underbrace{\begin{matrix} \mathbf{v}_1 & \mathbf{v}_2 \\ \theta_1 & \begin{bmatrix} 2 & -2 \\ 6 & -10 \end{bmatrix} \\ \theta_1\theta_2 & \begin{bmatrix} 2 & 0 \\ 0 & -2 \end{bmatrix} \end{matrix}}_{\mathbf{V}[\boldsymbol{\beta}_1; \boldsymbol{\beta}_2]} \underbrace{\begin{bmatrix} 2 & 0 \\ 0 & -2 \end{bmatrix}}_{\text{diag}(\theta_{1,1}^*, \theta_{2,1}^*)} = \underbrace{\begin{matrix} \mathbf{v}_1 & \mathbf{v}_2 \\ \theta_1^2 & \begin{bmatrix} 4 & 4 \\ 12 & 20 \end{bmatrix} \\ \theta_1^2\theta_2 & \begin{bmatrix} 4 & 4 \\ 12 & 20 \end{bmatrix} \end{matrix}}_{\mathbf{V}[\boldsymbol{\beta}_1 + \boldsymbol{\gamma}_1; \boldsymbol{\beta}_2 + \boldsymbol{\gamma}_1]} \quad (9)$$

While the above reveals the structure of \mathbf{V} , we don’t know \mathbf{V} . However, we recover its column space $\mathbf{U} \in \mathbb{R}^{s(r) \times K}$ from the moment matrix $\mathbf{M}_r(\mathbf{y}^*)$, for example with an SVD. Thus, we can relate \mathbf{U} and \mathbf{V} by a linear transformation: $\mathbf{V} = \mathbf{U}\mathbf{Q}$, where $\mathbf{Q} \in \mathbb{R}^{K \times K}$ is some unknown invertible matrix.

[Equation 8](#) can now be rewritten as:

$$\mathbf{U}[\boldsymbol{\beta}_1; \dots; \boldsymbol{\beta}_K] \mathbf{Q} \mathbf{D}_p = \mathbf{U}[\boldsymbol{\beta}_1 + \boldsymbol{\gamma}_p; \dots; \boldsymbol{\beta}_K + \boldsymbol{\gamma}_p] \mathbf{Q}, \quad p = 1, \dots, P, \quad (10)$$

¹ [36] is a short overview and [35] is a comprehensive treatment including numerical issues.

Table 2: Applications of the Polymom framework. See [Appendix B.2](#) for more details.

Mixture of linear regressions	
<p>Model $\mathbf{x} = [x, v]$ is observed where $x \in \mathbb{R}^D$ is drawn from an unspecified distribution and $v \sim \mathcal{N}(\mathbf{w} \cdot x, \sigma^2 I)$, and σ^2 is known. The parameters are $\theta_k^* = (\mathbf{w}_k) \in \mathbb{R}^D$.</p>	<p>Observation functions $\phi_{\alpha,b}(\mathbf{x}) = x^\alpha v^b$ for $0 \leq \alpha \leq 3, b \in [2]$.</p> <p>Moment polynomials $f_{\alpha,1}(\theta) = \sum_{p=1}^P \mathbb{E}[x^{\alpha+\gamma_p}] w_p$ $f_{\alpha,2}(\theta) = \mathbb{E}[x^\alpha] \sigma^2 + \sum_{p,q=1}^P \mathbb{E}[x^\alpha x_p x_q] w_p w_q$, where the $\gamma_p \in \mathbb{N}^P$ is 1 in position p and 0 elsewhere.</p>
Mixture of Gaussians	
<p>Model $\mathbf{x} \in \mathbb{R}^D$ is observed where \mathbf{x} is drawn from a Gaussian with diagonal covariance: $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\xi}, \text{diag}(\mathbf{c}))$. The parameters are $\theta_k^* = (\boldsymbol{\xi}_k, \mathbf{c}_k) \in \mathbb{R}^{D+D}$.</p>	<p>Observation functions $\phi_\alpha(\mathbf{x}) = \mathbf{x}^\alpha$ for $0 \leq \alpha \leq 4$.</p> <p>Moment polynomials $f_\alpha(\theta) = \prod_{d=1}^D h_{\alpha_d}(\xi_d, c_d)$.²</p>
Multiview mixtures	
<p>Model With 3 views, $\mathbf{x} = [x^{(1)}, x^{(2)}, x^{(3)}]$ is observed where $x^{(1)}, x^{(2)}, x^{(3)} \in \mathbb{R}^D$ and $x^{(\ell)}$ is drawn from an unspecified distribution with mean $\boldsymbol{\xi}^{(\ell)}$ for $\ell \in [3]$. The parameters are $\theta_k^* = (\boldsymbol{\xi}_k^{(1)}, \boldsymbol{\xi}_k^{(2)}, \boldsymbol{\xi}_k^{(3)}) \in \mathbb{R}^{D+D+D}$.</p>	<p>Observation functions $\phi_{ijk}(\mathbf{x}) = x_i^{(1)} x_j^{(2)} x_k^{(3)}$ where $1 \leq i, j, k \leq D$.</p> <p>Moment polynomials $f_{ijk}(\theta) = \xi_i^{(1)} \xi_j^{(2)} \xi_k^{(3)}$.</p>

which is a generalized eigenvalue problem where \mathbf{D}_p are the eigenvalues and \mathbf{Q} are the eigenvectors. Crucially, the eigenvalues, $\mathbf{D}_p = \text{diag}(\theta_{1,p}^*, \dots, \theta_{K,p}^*)$ give us solutions to our parameters. Note that for any choice of β_1, \dots, β_K and $p \in [P]$, we have generalized eigenvalue problems that share eigenvectors \mathbf{Q} , though their eigenvalues \mathbf{D}_p may differ. Corresponding eigenvalues (and hence solutions) can be obtained by solving a simultaneous generalized eigenvalue problem, e.g., by using random projections like Algorithm B of [4] or more robust [37] simultaneous diagonalization algorithms [38, 39, 40].

We describe one approach to solve (10), which is similar to Algorithm B of [4]. The idea is to take P random weighted combinations of the equations (10) and solve the resulting (generalized) eigendecomposition problems. Let $R \in \mathbb{R}^{P \times P}$ be a random matrix whose entries are drawn from $\mathcal{N}(0, 1)$. Then solve $\mathbf{U}[\beta_1; \dots; \beta_K]^{-1} \left(\sum_{p=1}^P R_{q,p} \mathbf{U}[\beta_1 + \gamma_p; \dots; \beta_K + \gamma_p] \right) \mathbf{Q} = \mathbf{Q} \mathbf{D}_q$ for each $q = 1, \dots, Q$. The resulting eigenvalues can be collected in $\Lambda \in \mathbb{R}^{P \times K}$, where $\Lambda_{q,k} = \mathbf{D}_{q,k,k}$. Note that by definition $\Lambda_{q,k} = \sum_{p=1}^P R_{q,p} \theta_{k,p}^*$, so we can simply invert to obtain $[\theta_1^*, \dots, \theta_K^*] = R^{-1} \Lambda$. Although this simple approach does not have great numerical properties, these eigenvalue problems are solvable if the eigenvalues $[\lambda_{q,1}, \dots, \lambda_{q,K}]$ are distinct for all q , which happens with probability 1 as long as the parameters θ_k^* are different from each other. In [Appendix B.1](#), we show how the tensor decomposition algorithm from [4] can be seen as solving (10) for a particular instantiation of β_1, \dots, β_K .

5 Applications

Let us now look at some applications of Polymom. [Table 2](#) presents several models with corresponding observation functions and moment polynomials. It is fairly straightforward to write down observation functions for a given model. The moment polynomials can then be derived by computing expectations under the model, a computation comparable to deriving updates for EM.

² $h_\alpha(\xi, c) = \sum_{i=0}^{\lfloor \alpha/2 \rfloor} a_{\alpha, \alpha-2i} \xi^{\alpha-2i} c^i$ and $a_{\alpha,i}$ be the absolute value of the coefficient of the degree i term of the α^{th} (univariate) Hermite polynomial. For example, the first few are $h_1(\xi, c) = \xi$, $h_2(\xi, c) = \xi^2 + c$, $h_3(\xi, c) = \xi^3 + 3\xi c$, $h_4(\xi, c) = \xi^4 + 6\xi^2 c + 3c^2$.

	Method.	EM	TF	Poly	EM	TF	Poly	EM	TF	Poly
Gaussians	K, D	$T = 10^3$			$T = 10^4$			$T = 10^5$		
spherical	2, 2	0.37	2.05	0.58	0.24	0.73	0.29	0.19	0.36	0.14
diagonal	2, 2	0.44	2.15	0.48	0.48	4.03	0.40	0.38	2.46	0.35
constrained	2, 2	0.49	7.52	0.38	0.47	2.56	0.30	0.34	3.02	0.29
Others	K, D	$T = 10^4$			$T = 10^5$			$T = 10^6$		
3-view	3, 3	0.38	0.51	0.57	0.31	0.33	0.26	0.36	0.16	0.12
lin. reg.	2, 2	-	-	3.51	-	-	2.60	-	-	2.52

Table 3: T is the number of samples, and the error metric is defined above. **Methods:** EM: sklearn initialized with k-means using 5 random restarts; TF: tensor power method implemented in Python; Poly: Polymom by solving [Problem 7](#). **Models:** for mixture of Gaussians, we have $\sigma \approx 2\|\mu_1 - \mu_2\|_2$; ‘spherical’ and ‘diagonal’ describes the type of covariance matrix. The mean parameters of constrained Gaussians satisfies $\mu_1 + \mu_2 = 1$. The best result is **bolded**. TF only handles spherical variance, but it was of interest to see what TF does if the data is drawn from mixture of Gaussians with diagonal covariance, these results are in ~~strikeout~~.

We implemented Polymom for several mixture models in Python.³ We used CVXOPT to handle the SDP and the random projections algorithm from [to](#) to extract solutions. In [Table 3](#), we show the relative error $\max_k \|\theta_k - \theta_k^*\|_2 / \|\theta_k^*\|_2$ averaged over 10 random models of each class.

Guarantees. In the rest of this section, we will discuss guarantees on parameter recovery for each of these models. In summary, we match many of the existing results in the literature for the mixture of linear regressions and multiview mixtures when $K \leq D$. In these case the moment matrix is fully determined by the linear constraints and [Problem 7](#) is just a linear solve. More discussions can be found in [Appendix B.2](#).

In addition, we can obtain *per-instance* guarantees in the following sense. Recall that Polymom involves solving an SDP relaxation and performing solution extraction. If the SDP solution has a flat extension ([Theorem 1](#)) at the true number of components K (a checkable assumption), then we have solved the moment completion problem exactly, and since solution extraction always works, we are guaranteed to obtain the true parameters. On the other hand, if the SDP solution has a higher rank $K' > K$, then as a consolation prize, we have found a K' -mixture model that matches the moments (that we observed) of the true K -mixture model.

6 Conclusion

We presented an unifying framework for learning many types of mixture models via the method of moments. For example, for the mixture of Gaussians, we can apply the same algorithm to both mixtures in 1D needing higher-order moments [3, 11] and mixtures in high dimensions where lower-order moments suffice [6]. The Generalized Moment Problem [15, 14] and its semidefinite relaxation hierarchies is what gives us the generality, although we rely heavily on the ability of nuclear norm minimization to recover the underlying rank. As a result, while we always obtain parameters satisfying the moment conditions, we do not have formal guarantees on consistent estimation in general, although we do have guarantees for several model families. The second main tool is solution extraction, which characterizes a more general structure of mixture models compared the tensor structure observed by [6, 4]. This view draws connections to the literature on solving polynomial systems, where many techniques might be useful [35, 18, 19]. Finally, through the connections we’ve drawn, it is our hope that Polymom can make the method of moments as turnkey as EM on more latent-variable models, and provide a way to improve the statistical efficiency of method of moments procedures.

Acknowledgments. This work was supported by a Microsoft Faculty Research Fellowship to the third author and a NSERC PGS-D fellowship for the first author.

References

- [1] D. M. Titterton, A. F. Smith, and U. E. Makov. *Statistical analysis of finite mixture distributions*, volume 7. Wiley New York, 1985.

³<https://github.com/sidaw/polymom>

- [2] G. McLachlan and D. Peel. *Finite mixture models*. John Wiley & Sons, 2004.
- [3] K. Pearson. Contributions to the mathematical theory of evolution. *Philosophical Transactions of the Royal Society of London. A*, 185:71–110, 1894.
- [4] A. Anandkumar, D. Hsu, and S. M. Kakade. A method of moments for mixture models and hidden Markov models. In *Conference on Learning Theory (COLT)*, 2012.
- [5] A. Anandkumar, D. P. Foster, D. Hsu, S. M. Kakade, and Y. Liu. Two SVDs suffice: Spectral decompositions for probabilistic topic modeling and latent Dirichlet allocation. In *Advances in Neural Information Processing Systems (NIPS)*, 2012.
- [6] A. Anandkumar, R. Ge, D. Hsu, S. M. Kakade, and M. Telgarsky. Tensor decompositions for learning latent variable models. *arXiv*, 2013.
- [7] D. Hsu, S. M. Kakade, and P. Liang. Identifiability and unmixing of latent parse trees. In *Advances in Neural Information Processing Systems (NIPS)*, 2012.
- [8] D. Hsu and S. M. Kakade. Learning mixtures of spherical Gaussians: Moment methods and spectral decompositions. In *Innovations in Theoretical Computer Science (ITCS)*, 2013.
- [9] A. Chaganty and P. Liang. Spectral experts for estimating mixtures of linear regressions. In *International Conference on Machine Learning (ICML)*, 2013.
- [10] A. T. Kalai, A. Moitra, and G. Valiant. Efficiently learning mixtures of two Gaussians. In *Symposium on Theory of Computing (STOC)*, pages 553–562, 2010.
- [11] M. Hardt and E. Price. Sharp bounds for learning a mixture of two Gaussians. *arXiv preprint arXiv:1404.4997*, 2014.
- [12] R. Ge, Q. Huang, and S. M. Kakade. Learning mixtures of Gaussians in high dimensions. *arXiv preprint arXiv:1503.00424*, 2015.
- [13] B. Balle, X. Carreras, F. M. Luque, and A. Quattoni. Spectral learning of weighted automata - a forward-backward perspective. *Machine Learning*, 96(1):33–63, 2014.
- [14] J. B. Lasserre. *Moments, Positive Polynomials and Their Applications*. Imperial College Press, 2011.
- [15] J. B. Lasserre. A semidefinite programming approach to the generalized problem of moments. *Mathematical Programming*, 112(1):65–92, 2008.
- [16] H. J. Stetter. Multivariate polynomial equations as matrix eigenproblems. *WSSIA*, 2:355–371, 1993.
- [17] H. M. Möller and H. J. Stetter. Multivariate polynomial equations with multiple zeros solved by matrix eigenproblems. *Numerische Mathematik*, 70(3):311–329, 1995.
- [18] B. Sturmfels. *Solving systems of polynomial equations*. American Mathematical Society, 2002.
- [19] D. Henrion and J. Lasserre. Detecting global optimality and extracting solutions in GloptiPoly. In *Positive polynomials in control*, pages 293–310, 2005.
- [20] A. Anandkumar, R. Ge, D. Hsu, and S. Kakade. A tensor spectral approach to learning mixed membership community models. In *Conference on Learning Theory (COLT)*, pages 867–881, 2013.
- [21] A. Anandkumar, R. Ge, and M. Janzamin. Provable learning of overcomplete latent variable models: Semi-supervised and unsupervised settings. *arXiv preprint arXiv:1408.0553*, 2014.
- [22] K. Viele and B. Tong. Modeling with mixtures of linear regressions. *Statistics and Computing*, 12(4):315–330, 2002.
- [23] B. Sturmfels. *Algorithms in invariant theory*. Springer Science & Business Media, 2008.
- [24] R. M. Corless, K. Gatermann, and I. S. Kotsireas. Using symmetries in the eigenvalue method for polynomial systems. *Journal of Symbolic Computation*, 44(11):1536–1550, 2009.
- [25] R. E. Curto and L. A. Fialkow. *Solution of the truncated complex moment problem for flat data*, volume 568. American Mathematical Society, 1996 1996.
- [26] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.
- [27] M. Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging applications of algebraic geometry*, pages 157–270, 2009.
- [28] P. A. Parrilo and B. Sturmfels. Minimizing polynomial functions. *Algorithmic and quantitative real algebraic geometry, DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 60:83–99, 2003.
- [29] P. A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical programming*, 96(2):293–320, 2003.
- [30] N. Ozay, M. Sznajder, C. M. Lagoa, and O. I. Camps. GPCA with denoising: a moments-based convex approach. In *Computer Vision and Pattern Recognition (CVPR)*, pages 3209–3216, 2010.
- [31] D. Koller and N. Friedman. *Probabilistic graphical models: principles and techniques*. MIT Press, 2009.
- [32] R. Curto and L. Fialkow. The truncated complex K-moment problem. *Transactions of the American mathematical society*, 352(6):2825–2855, 2000.

- [33] R. W. Freund and F. Jarre. A sensitivity result for semidefinite programs. *Operations Research Letters*, 32:126–132, 2004.
- [34] R. M. Corless, P. M. Gianni, B. M. Trager, and S. M. Watt. The singular value decomposition for polynomial systems. In *International Symposium on Symbolic and Algebraic Computation*, pages 195–207, 1995.
- [35] H. J. Stetter. *Numerical polynomial algebra*. Siam, 2004.
- [36] P. Dreesen, K. Batselier, and B. D. Moor. Back to the roots: Polynomial system solving, linear algebra, systems theory. In *IFAC Symposium on System Identification (SYSID)*, pages 1203–1208, 2012.
- [37] V. Kuleshov, A. Chaganty, and P. Liang. Simultaneous diagonalization: the asymmetric, low-rank, and noisy settings. *arXiv*, 2015.
- [38] J. Cardoso and A. Souloumiac. Jacobi angles for simultaneous diagonalization. *SIAM Journal on Matrix Analysis and Applications*, 17(1):161–164, 1996.
- [39] A. Bunse-Gerstner, R. Byers, and V. Mehrmann. Numerical methods for simultaneous diagonalization. *SIAM Journal on Matrix Analysis and Applications*, 14(4):927–949, 1993.
- [40] B. Afsari. Simple LU and QR based non-orthogonal matrix joint diagonalization. In *Independent Component Analysis and Blind Signal Separation*, pages 1–7, 2006.

A Theory of the Moment Completion Problem

For solution extraction, we assumed that moments of all monomials are observed but for many models only polynomials of parameters can be estimated from the data. For example, in a Gaussian mixture the 2^{nd} moment observable function $\phi(x) = \xi^2 + c$ is a polynomial, but solution extraction requires moments of monomials like ξ^2 and c . Furthermore, we assumed in [Section 4](#) that there exists underlying true parameters $[\theta_k^*]_{k=1}^K$ while an arbitrary moment sequence of the parameters \mathbf{y} and its corresponding moment matrix $\mathbf{M}(\mathbf{y})$ may not correspond to any parameters (i.e. no representing measure). In [Section 5](#), we give an example of how moment completion can be done with just linear algebra for multiview models, and we focus on the harder case of having to solve the SDP [Problem 7](#) in this section.

While we do not have a complete answer since the rank constrained [Problem 6](#) cannot be solved, we analyze the relaxed moment completion problem ([Problem 7](#)) and give some sufficient conditions for solution extraction and sufficient conditions for parameter recovery.

A.1 Conditions for solution extraction

Under the exact setting, and assuming a moment matrix generated by true parameters θ^* , In [Section 4](#), we showed that simple conditions based only on the column space basis is sufficient for solution extraction to be successful. However, to further investigate consistency and noise, we need to address a few more important issues. First consider the noiseless setting, we may not have enough moment constraints to guarantee a unique solution (identifiability). Even if we assume that we have enough constraints for identifying a K mixture, we still do not know if solving the relaxed [Problem 7](#) that relaxed the rank = K constraint can recover the true parameters. Second, under noise, there may not exist a rank K basis of the moment matrix and even when a rank K basis exists, it may not correspond to any true parameters.

We now address the problem of when some moment matching parameters can be extracted. We first discuss the concept of a *flat extension* which is the same as conditions in [Section 4](#) where “ $\mathbf{B}_p := \mathbf{U}[\gamma_p + [\beta_1, \dots, \beta_K]]$ is observed” and $\mathbf{U}[\beta_1; \dots; \beta_K]$ is a column space basis of $\mathbf{M}_r(\mathbf{y})$. Let the highest degree monomial of $\mathbf{U}[\beta_1; \dots; \beta_K]$ be of degree $r - 1 = \deg(\theta^{\beta_K}) = |\beta_K|$, and the highest degree monomial of $\mathbf{B}_p := \mathbf{U}[\gamma_p + [\beta_1, \dots, \beta_K]]$ be of degree $r = |\gamma_p + \beta_K| = |\deg(\beta_K)| + 1$. Since $\mathbf{U}[\beta_1; \dots; \beta_K]$ is a basis of $\text{col}(\mathbf{M}_r(\mathbf{y}))$

$$\text{rank}(\mathbf{M}_{r-1}(\mathbf{y})) = \text{rank}(\mathbf{U}[\beta_1; \dots; \beta_K]) = K \quad (11)$$

$$= \text{rank}(\mathbf{M}_r(\mathbf{y})) \geq \text{rank}(\mathbf{U}[\gamma_p + [\beta_1, \dots, \beta_K]]). \quad (12)$$

If we got this basis from the moment matrix, then we say that the moment matrix $\mathbf{M}_{r-1}(\mathbf{y})$ corresponding to $\mathbf{U}[\beta_1; \dots; \beta_K]$ has a flat extension, because $\mathbf{M}_{r-1}(\mathbf{y})$ can be extended to a moment matrix $\mathbf{M}_r(\mathbf{y})$ with higher degree monomials without an increase in rank. The concept of flat extension and its consequences are of central importance for the truncated moment problem, which is quite relevant to our problem and studied extensively by [1, 2, 3, 4]. Next, we reproduce the simplest flat extension theorem:

Theorem A.1 ([1]: flat extension theorem). *Suppose $\mathbf{M}_{r-1}(\mathbf{y}) \succeq 0$ and there exists $\mathbf{M}_r(\mathbf{y})$ so that $\text{rank}(\mathbf{M}_r(\mathbf{y})) = \text{rank}(\mathbf{M}_{r-1}(\mathbf{y}))$ (i.e. a flat extension), then there exists an unique $\text{rank}(\mathbf{M}_r(\mathbf{y}))$ -atomic representing measure μ of $\mathbf{M}_r(\mathbf{y})$.*

Here the first column of $\mathbf{M}_r(\mathbf{y})$ contains every monomial of degree up to r so that $\deg(\mathbf{v}_r(\theta)) = r$. However, several generalizations of the flat extension theorem are also useful for estimation of mixture models where sparse monomials are handled [5, 6] or where constraints are handled [4].

The conceptual importance is that [Theorem A.1](#) allows us to work with just the moment matrix satisfying constraints from possibly noisy observations, without assuming the moment matrix is generated by some true parameters. Of course, it also provides a checkable criterion for when solutions can be extracted [7]. We still do not know if solving [Problem 7](#) provides a flat extension in a finite number of steps. [8, 9, 10] investigated this issue very recently and showed that linear optimization over the cone of moments have finite convergence under generic conditions (theorem 4.2 of [10]).

Still, our issue is not fully resolved as representing measures under linear constraints may not be unique, and as a result even a flat moment matrix may not correspond to the true parameters. For parameter fitting, we’d like to find the solution with minimal rank or otherwise optimal in some way. We explore this issue next but unfortunately we can only give some partial answers.

Proposition A.2 (existence of \mathbf{C}). *In the noiseless setting, there exist \mathbf{C} so that minimizing $\mathbf{C} \bullet \mathbf{M}_r(\mathbf{y}) = \mathbf{c} \cdot \mathbf{y}$ will give the right solution.*

Proof. Let $\mathbf{M}_r(\mathbf{y}) = \mathbf{U}\Sigma\mathbf{U}^T$ be the SVD with $\mathbf{U} \in \mathbb{R}^{s(r) \times K}$ and $\Sigma \in \mathbb{R}^{K \times K}$. Let $\mathbf{U}_\perp \in \mathbb{R}^{s(r) \times (s(r)-K)}$ be the orthogonal compliment of \mathbf{U} , then any $\mathbf{C} = \mathbf{U}_\perp \mathbf{D} \mathbf{U}_\perp^T$ is fine for arbitrary diagonal matrix $\mathbf{D} \in \mathbb{R}^{(s(r)-K) \times (s(r)-K)}$.

The convex iteration algorithm [11] is one way to reduce rank that sometimes works for us empirically, where if the convex iteration algorithm converges to 0, then the moment matrix has rank K .

B Examples

B.1 Tensor factorization as solution extraction

Example B.1 (Tensor decomposition as solution extraction). *Many latent variable models have been tackled via tensor decomposition [12], and symmetric, undercomplete tensor decomposition can be framed as a solution extraction problem. Suppose we observe the tensor $\mathbf{T} := \sum_{k=1}^K \theta_k^{*\otimes 3} \in \mathbb{R}^{P \times P \times P}$. We would like to recover the components θ_k^* . For us, the inputs are constraints $\theta_r \theta_s \theta_t - T_{rst} = 0$ for all $r, s, t = 1, \dots, P$. Choose $\mathbf{v}(\boldsymbol{\theta}) = [1, \theta_1, \dots, \theta_P, \theta_1^2, \theta_1 \theta_2, \dots, \theta_P^2] = [1, \boldsymbol{\theta}, \text{vecs}(\boldsymbol{\theta} \otimes \boldsymbol{\theta})]$, where $\text{vecs} : \mathbb{R}^{P \times P} \rightarrow \mathbb{R}^{P^2}$ just flattens the matrix. In the simplest case, suppose $P = K$ and $\text{rank}(\mathbf{U}) = K$. Then the fully observed \mathbf{U} is*

$$\mathbf{U} = \begin{array}{c} \text{size} \\ 1 \\ P \\ P^2 \end{array} \begin{array}{c} P \\ \left[\begin{array}{c} \mathbf{U}_1 \\ \mathbf{U}_2 \\ \mathbf{U}_3 \end{array} \right] \end{array} = \begin{array}{c} \text{terms} \\ 1 \\ \boldsymbol{\theta} \\ \text{vecs}(\boldsymbol{\theta} \otimes \boldsymbol{\theta}) \end{array} \begin{array}{c} \boldsymbol{\theta} \\ \left[\begin{array}{c} \mathcal{L}_y(\boldsymbol{\theta}) \\ \mathcal{L}_y(\boldsymbol{\theta} \otimes \boldsymbol{\theta}) \\ \mathcal{L}_y(\text{vecs}(\boldsymbol{\theta} \otimes \boldsymbol{\theta}) \otimes \boldsymbol{\theta}) \end{array} \right] \end{array} \quad (13)$$

where the linear functional \mathcal{L}_y applies elementwise. One choice of basis is just all the variables $\mathbf{U}[\beta_1; \dots; \beta_K] = \mathbf{U}_2$ and the eigenvalue problem we are required to solve is the generalized Hermitian eigenvalue problem $\mathbf{U}_2 \mathbf{Q} \mathbf{D} = \left(\sum_{p=1}^P \eta_p \mathcal{L}_y(\theta_p \boldsymbol{\theta} \otimes \boldsymbol{\theta}) \right) \mathbf{Q}$. [13] proposed an algorithm that is procedurally identical, where, in their notation $\text{Pairs} := \mathbf{U}_2$ and $\text{Triples}(\eta) := \left(\sum_{p=1}^P \eta_p \mathcal{L}_y(\theta_p \boldsymbol{\theta} \otimes \boldsymbol{\theta}) \right)$, and the algorithm proposed needed to solve the eigenvalue problem $B(\eta) = \text{Pairs}^{-1} \text{Triples}(\eta)$.

Typically, β_1, \dots, β_K are just the first K monomials in \mathbf{v} (i.e. the K monomials of the smallest degree).

Under this formulation, generalization to the fully-observed overcomplete tensor decomposition case $K \geq D = P$ is clear if we observe enough moments to have enough basis vectors such that $\text{rank}(\mathbf{U}[\beta_1; \dots; \beta_K]) = K$:

Proposition B.2. *If $K \leq 1 + P + P^2 + \dots + P^r = \frac{P^{r+1}-1}{P-1}$, then solution extraction succeeds if we observe moments up to order $2r + 1$ and monomials vectors of the true parameters $\mathbf{v}_r(\boldsymbol{\theta}_1), \dots, \mathbf{v}_r(\boldsymbol{\theta}_K)$ are linearly independent.*

Proof. To get the theoretical result, it suffices to consider higher-order moments:

$$\mathbf{U} = \begin{array}{c} \text{terms} \\ \text{vecs}(\boldsymbol{\theta}^{\otimes r}) \\ \text{vecs}(\boldsymbol{\theta}^{\otimes r+1}) \end{array} \begin{array}{c} \text{vecs}(\boldsymbol{\theta}^{\otimes r}) \\ \left[\begin{array}{c} \mathcal{L}_y(\text{vecs}(\boldsymbol{\theta}^{\otimes r}) \otimes \text{vecs}(\boldsymbol{\theta}^{\otimes r})) \\ \mathcal{L}_y(\text{vecs}(\boldsymbol{\theta}^{\otimes r+1}) \otimes \text{vecs}(\boldsymbol{\theta}^{\otimes r})) \end{array} \right] \end{array} \quad (14)$$

where we can take the $\mathbf{U}[\beta_1; \dots; \beta_K]$ from the top block, and $\mathbf{U}[\beta_1 + \gamma_q; \dots; \beta_K + \gamma_q]$ belongs to the bottom block for all q . So $2r + 1$ order moments is needed if $K \leq P^r$ and this result is comparable to [14]. In practice, we would take all moments $\text{vecs}(\boldsymbol{\theta}^{\otimes 1}), \dots, \text{vecs}(\boldsymbol{\theta}^{\otimes r+1})$. We may use lower order moments as well:

$$\mathbf{U} = \begin{array}{c} \text{terms} \\ \text{vecs}(\boldsymbol{\theta}^{\otimes 1}) \\ \text{vecs}(\boldsymbol{\theta}^{\otimes 2}) \\ \vdots \\ \text{vecs}(\boldsymbol{\theta}^{\otimes r+1}) \end{array} \begin{array}{c} \text{vecs}(\boldsymbol{\theta}^{\otimes 1}) \quad \text{vecs}(\boldsymbol{\theta}^{\otimes 2}) \quad \dots \quad \text{vecs}(\boldsymbol{\theta}^{\otimes r}) \\ \left[\begin{array}{c} \vdots \\ \dots \mathcal{L}_y(\text{vecs}(\boldsymbol{\theta}^{\otimes l}) \otimes \text{vecs}(\boldsymbol{\theta}^{\otimes m})) \dots \\ \vdots \end{array} \right] \end{array} \quad (15)$$

where the entry of this matrix at block l, m is $\mathcal{L}_y(\text{vecs}(\boldsymbol{\theta}^{\otimes l}) \otimes \text{vecs}(\boldsymbol{\theta}^{\otimes m}))$ as expected. While this still requires observing $2r + 1^{\text{th}}$ order moments, lower order moments are more accurate and can result in better parameter estimates. \square

B.2 Moment completion for specific models

B.2.1 Mixture of Linear Regressions

In [Example 2.2](#), we described the mixture of linear regressions model in 1-dimension with parameters $\theta_k^* = (w_k, \sigma_k^2)$. Let us now consider the D -dimensional extension: we observe $\mathbf{x} = [x, v]^4$ where $x := [x_1, \dots, x_D]$

⁴ We use v here since y is reserved for the parameter moments.

is drawn from an unspecified distribution and $v = w \cdot x + \epsilon$ with $\epsilon \sim \mathcal{N}(0, \sigma^2)$ for a known σ . The parameters are $\theta_k^* = (w_k)$ for $1 \leq k \leq K$. Next, we choose observation functions $\phi_{\alpha,b}(\mathbf{x}) = x^\alpha v^b$ for $\alpha : 0 \leq |\alpha| \leq 3$ and $0 \leq b \leq 3$, with corresponding moment polynomials: $f_{\alpha,b}(\theta, \mathbf{x}) = x^\alpha \mathbb{E}_{\epsilon \sim \mathcal{N}(0, \sigma^2)} \left[(w \cdot x + \epsilon)^b \right]$. These polynomials can be expressed in closed form using Hermite polynomials (see [Section B.2.2](#)). For example, $f_{0,2}(\theta, \mathbf{x}) = ((w \cdot x)^2 + \sigma^2)$.

Given these observation functions and moment polynomials, and data, the Polymom framework solves the moment completion problem ([Problem 6](#)) followed by solution extraction ([Section 4](#)) to recover the parameters. Further, we can guarantee that Polymom can recover parameters for this model when $K \leq D$ by showing that [Problem 6](#) can be solved exactly. Note that while no entry of the moment matrix is directly observed, each observation gives us a linear constraint on the entries of the moment matrix. Let $\gamma_p \in \mathbb{N}^P$ be the vector with value 1 at position p and 0 elsewhere, then $\mathcal{L}_y(f_{\alpha,1}(\theta)) = \sum_{p=1}^P \mathbb{E}[x^{\alpha+\gamma_p}] y_{\gamma_p}$, and $\mathcal{L}_y(f_{\alpha,2}(\theta)) = \left(\mathbb{E}[x^\alpha] \sigma^2 + \sum_{p,q=1}^P \mathbb{E}[x^{\alpha+\gamma_p+\gamma_q}] y_{\gamma_p+\gamma_q} \right)$, etc. When $K \leq D$, there are enough equations that this system admits a unique solution for \mathbf{y} .

Note that [15] recover parameters for this model by solving a series of low-rank tensor recovery problems, which ultimately requires the computation of the same moments described above. In contrast, the Polymom framework makes the dependence on moments upfront and takes care of the heavy-lifting in a problem-agnostic manner. Furthermore, we can even obtain parameters outside the regime of [15]: with the above observation functions and moment polynomials, we can recover parameters (with a certificate).

B.2.2 Mixture of Gaussians

We now look at D -dimensional extensions to [Example 2.1](#). Let the data be drawn from Gaussians with diagonal covariance, $x|z \sim \mathcal{N}(\xi_z, \text{diag}(c_z))$. The parameters of this model are $\theta_k^* = (\xi_k, c_k) \in \mathbb{R}^{2D}$. The observable functions are $\phi_\alpha(\mathbf{x}) := \mathbf{x}^\alpha$, and the moment polynomials are $f_\alpha(\theta) = \mathbb{E}[\mathbf{x}^\alpha] = \prod_{d=1}^D h_{\alpha[d]}(\xi[d], c[d])$, where $h_\alpha(\xi, c) = \sum_{i=0}^{\lfloor \alpha/2 \rfloor} a_{\alpha, \alpha-2i} \xi^{\alpha-2i} c^i$ and $a_{\alpha,i}$ be the absolute value of the coefficient of the degree i term of the α^{th} (univariate) Hermite polynomial. The first few are $h_1(\xi, c) = \xi$, $h_2(\xi, c) = \xi^2 + c$, $h_3(\xi, c) = \xi^3 + 3\xi c$, $h_4(\xi, c) = \xi^4 + 6\xi^2 c + 3c^2$.

Using this set of ϕ_α and f_α , Polymom will attempt to solve the SDP in [Problem 7](#) and recover the parameters. In this case however, the moment conditions are non-trivial and we cannot guarantee recovery of the true parameters. However, Polymom is guaranteed to recover parameters that match the moments and that minimizes nuclear norm.

We can modify this model by introducing constraints: consider the case of 2D mixture where the mean parameters for all components lies on a parabola $\xi_1 - \xi_2^2 = 0$. In this case, we just need to add constraints to [Problem 7](#): $y_{(1,0)+\beta} - y_{(0,2)+\beta} = 0$ for all $\beta \in \mathbb{N}^2$ up to degree $|\beta| \leq 2r - 2$.

By incorporating these constraints at estimation time, we can possibly identify the model parameters with less moments. See [Section C](#) for more details.

B.2.3 Multiview Mixtures

Here we consider the three-view mixture model which has been well studied in [12, section 3.3]. We will show that we can solve the model without explicit whitening, a transformation that has been shown to introduce noise[16]. The model is a mixture of three conditionally independent arbitrary distributions parameterized by their conditional means: we have $z \sim \text{Multinomial } \pi$, $\mathbf{x}_l|z \sim p_l(\xi_z^{(l)})$ where $p_l(\xi_z^{(l)})$ is such that $E_{\mathbf{x}_l|z}[\mathbf{x}_l] = \xi$. The parameters are $\theta_k = [\xi^{(1)}, \xi^{(2)}, \xi^{(3)}]$. Using the observation functions $\phi = [x^{(1)}, x^{(2)}, x^{(3)}, x^{(1)} \otimes x^{(2)}, \dots, x^{(1)} \otimes x^{(2)} \otimes x^{(3)}]$, we have the following moment polynomials, $f = [\xi^{(1)}, \xi^{(2)}, \xi^{(3)}, \xi^{(1)} \otimes \xi^{(2)}, \dots, \xi^{(1)} \otimes \xi^{(2)} \otimes \xi^{(3)}]$.

The multiview mixture model is another model for which we can guarantee parameter recovery when $K \leq D$. To prove this is the case, we will again show that [Problem 7](#) can be solved exactly. It suffices to consider just the first P columns of the moment matrix \mathbf{M}_2 , which are almost directly observable. As before, $\text{vecs}(\cdot)$ just flattens a matrix into a vector.

$$\mathbf{M}_2^\top = \begin{matrix} & \xi_1 & \xi_2 & \xi_3 & \text{vecs}(\xi_1 \otimes \xi_2) & \text{vecs}(\xi_1 \otimes \xi_3) & \text{vecs}(\xi_2 \otimes \xi_3) \\ \begin{matrix} \xi_1 \\ \xi_2 \\ \xi_3 \end{matrix} & \begin{bmatrix} Z_{2,0,0} & \mathbf{Y}_{1,1,0} & \mathbf{Y}_{1,0,1} & Z_{2,1,0} & Z_{2,0,1} & \mathbf{Y}_{1,1,1} \\ \mathbf{Y}_{1,1,0} & Z_{0,2,0} & \mathbf{Y}_{0,1,1} & Z_{1,2,0} & \mathbf{Y}_{1,1,1} & Z_{0,2,1} \\ \mathbf{Y}_{1,0,1} & \mathbf{Y}_{0,1,1} & Z_{0,0,2} & \mathbf{Y}_{1,1,1} & Z_{1,0,2} & Z_{0,1,2} \end{bmatrix} \end{matrix} \quad (16)$$

where $Y_{\alpha_1, \alpha_3, \alpha_3}$ and $Z_{\alpha_1, \alpha_3, \alpha_3}$ are both equal to $\mathcal{L}_y(\xi_1^{\otimes \alpha_1} \otimes \xi_2^{\otimes \alpha_2} \otimes \xi_3^{\otimes \alpha_3})$, but are used to respectively denote observed and unknown variables. However, this equation is only partially true as both sides contain the

same set of values but the precise arrangements depends on where the minor matrix appears in the moment matrix. We ignore this problem as it should be clear from the row and column labels. In the undercomplete case, it is assumed that $\text{rank}(\mathbf{U}) = K \leq \min(P_1, P_2, P_3)$, thus we can easily complete this matrix using simple linear algebra in the exact case by repeatedly applying [Lemma B.3](#) below. Generally, we may try to complete the moment matrix by solving [Problem 7](#) from these partial observations, provided that optimizing with the nuclear norm recovers the true rank.

Lemma B.3 (low rank completion of missing corner). *For any matrix $\Gamma = \begin{bmatrix} A & B \\ C & X \end{bmatrix}$ with a missing block X , where $\text{rank}(\Gamma) = \text{rank}(A) = \text{rank}(B) = K$ and $A \in \mathbb{R}^{K \times K}$, $X = CA^{-1}B$ uniquely completes Γ .*

Proof. Because A contains the entire K elements basis, there exists unique $Y, Z \in \mathbb{R}^{K \times K}$ so that $B = AY$ and $C = ZA$. Similarly, $X = ZB = CA^{-1}B$.

C Extensions

C.1 Constraints on parameters

Constraints on parameters is a common and important consideration in applications. While constraints can often be addressed in maximum likelihood or *maximum a posteriori* learning using EM [17, see shared parameters], it is less clear how to address constraints under the tensor decomposition approach because of its reliance on special tensor structure and it is well-known that MME generally can give us parameters outside of the parameter space even in the well-specified case.

Example C.1. *Examples of constraints on parameters* **Some parameters are known:** Gaussian with sparse covariance matrix where we already know that some dimensions are uncorrelated; to solve a substitution cipher using an HMM, the transitions matrix is a language model that is given.

Parameters are tied: transitions in an HMM might only depend on the relatively difference between states if the states are ordered i.e. the transition matrix is Toeplitz.

Polytope constraints: some of our parameters might be probabilities: $\boldsymbol{\theta} = [\pi_1, \dots, \pi_P, \xi_1, \dots]$, $\pi_p \geq 0$, $\sum_{p=1}^P \pi_p = 1$ e.g. prior parameter in LDA.

Semialgebraic constraints: For some polynomial $g \in \mathbb{R}[\boldsymbol{\theta}]$, $g_i(\boldsymbol{\theta}_k^*) \geq 0, i = 1, \dots, I$. Quite powerful, can express things like discrete sets $\boldsymbol{\theta}_i \in \{0, 1\}$, ellipsoids.

The obvious attempt is to project to the feasible set after computing an unconstrained estimation with MME. But this approach has several serious issues. First, some constrained models are only identifiable after the constraints are taken into account, which happens when the model has a lot of parameters and we cannot observe correspondingly more moments. In this case, unconstrained estimation is useful only if we can characterize the entire subset of the parameters space satisfying moment conditions, which is generally not possible in the tensor decomposition approach. Second, we need to determine what projection to use. In the case of two equal parameters, if one estimate is much more noisy than the other, it can be better to just ignore the more noisy estimate than to project under the wrong metric (see [Example C.3](#)). Third and strangely, even in the case when the first two issues are handled, it was observed by [18] for probabilities parameters, that clipping to 0 is empirically inferior compared to heuristics like taking the absolute value, which is not a projection.

Under the Polymom formulation, we can take constraints into account during estimation. The technique of localizing matrix [3] in moment theory allows us to deal with semialgebraic constraints. Of course, the computational complexity increases if the constraints are themselves complicated and high degree. Next, we define the localization matrix, give an example, and then give a constrained version of the flat extension theorem.

Example C.2 (localizing matrix for an inequality constraint). *Let $\boldsymbol{\theta} = [c, \xi]$, so that $\boldsymbol{\theta}^\alpha = c^{\alpha_1} \xi^{\alpha_2}$ and $\mathcal{L}_y(\boldsymbol{\theta}^\alpha) = y_\alpha$, and chose the monomials $\mathbf{v}_2(\boldsymbol{\theta}) = [1, c, \xi, c^2, c\xi, \xi^2]$. Suppose that c is the variance and we want to have constraint that $c - 1 \geq 0$, then*

$$\mathbf{M}_1((c-1)\mathbf{y}) = \begin{matrix} & & 1 & & c & & \xi \\ & 1 & & & & & \\ & c & & & & & \\ & \xi & & & & & \end{matrix} \begin{bmatrix} y_{1,0} - 1 & y_{2,0} - y_{1,0} & y_{1,1} - y_{1,0} \\ y_{2,0} - y_{1,0} & y_{3,0} - y_{2,0} & y_{2,1} - y_{2,0} \\ y_{1,1} - y_{0,1} & y_{2,1} - y_{1,1} & y_{1,2} - y_{0,2} \end{bmatrix} \quad (17)$$

it is clear that a necessary condition for extracted solutions to satisfy the constraint $c - 1 \geq 0$ is that $\mathbf{M}_1((c-1)\mathbf{y}) \succeq 0$ since $\mathbf{f}^\top \mathbf{M}_1((c-1)\mathbf{y}) \mathbf{f} = \mathcal{L}_y(f(\boldsymbol{\theta})^2(c-1)) \geq 0$.

C.2 Noise and statistical efficiency

In the presense of noise [Problem 7](#) may not be feasible and even if it was, it may not be ideal to exactly match noisy moments. Furthermore, it is argued that higher order moments are too noisy to be useful, but there are also more of them and they do contain more information about the model parameters as long as we can model how noisy they are. We consider the problem with slack ϵ and a weighting matrix $\mathbf{W} \succ 0 \in \mathbb{R}^{N \times N}$ modelling how much noise is present in each constraint function. This effect is fairly well-known, and here is a very simple example which shows that even much more noisy measurements can improve efficiency.

Example C.3 (efficient estimation). *Suppose $X \sim \mathcal{N}([\xi, \xi], \text{diag}[\sigma^2, c\sigma^2])$ and we would like to estimate the mean parameter ξ by matching moments. Any estimators of the form $\hat{\xi} = \frac{1}{T} \sum_{t=1}^T (\gamma x_{t,1} + (1-\gamma)x_{t,2})$ are consistent and has risk*

$$R = \mathbb{E} [(\hat{\xi} - \xi)^2] = \mathbb{E} \left[\left(\gamma \sum_{t=1}^T \mathbf{x}_{t,2} - \gamma\xi + (1-\gamma) \sum_{t=1}^T \mathbf{x}_{t,2} - (1-\gamma)\xi \right)^2 \right] \quad (18)$$

$$= \mathbb{E} \left[\gamma^2 \left(\xi - \frac{1}{T} \sum_{t=1}^T \mathbf{x}_{t,1} \right)^2 + (1-\gamma)^2 \left(\xi - \frac{1}{T} \sum_{t=1}^T \mathbf{x}_{t,2} \right)^2 \right] \quad (19)$$

$$= \frac{1}{T} (\gamma^2 \sigma^2 + (1-\gamma)^2 c\sigma^2) \quad (20)$$

under the squared loss, and the efficient estimator would have $\gamma = \frac{c-1}{c}$ and a risk of $\frac{\sigma^2}{T} \frac{c^2-c+1}{c^2}$. For $c = 10$, the risk for efficient estimation is $0.91 \frac{\sigma^2}{T}$ whereas for $\gamma = 0.5$, the risk is $2.75 \frac{\sigma^2}{T}$.

$$\begin{aligned} & \underset{\mathbf{g}, \mathbf{y}}{\text{minimize}} && \mathbf{C} \bullet \mathbf{M}(\mathbf{y}) \\ & \text{s.t.} && g_n = \sum_{\alpha} a_n \alpha y_{\alpha} - \mathbb{E}[\phi_n(\mathbf{x})] \\ & && \mathbf{g}^T \mathbf{W} \mathbf{g} \leq \epsilon \\ & && \mathbf{M}(\mathbf{y}) \succeq 0 \end{aligned} \quad (21)$$

In the simplest case when $\mathbf{W} = \mathbf{I}_N$, and $\epsilon = 0$, [Problem 21](#) is the same as [Problem 7](#).

$$\begin{aligned} & \underset{\mathbf{g}, \mathbf{y}}{\text{minimize}} && \mathbf{C} \bullet \mathbf{M}(\mathbf{y}) \\ & \text{s.t.} && g_n = \sum_{\alpha} a_n \alpha y_{\alpha} - \mathbb{E}[\phi_n(\mathbf{x})] \\ & && \mathbf{W} \bullet \mathbf{F} \leq \epsilon \\ & && \mathbf{M}(\mathbf{y}) \succeq 0 \\ & && \begin{bmatrix} 1 & \mathbf{g}^T \\ \mathbf{g} & \mathbf{F} \end{bmatrix} \succeq 0 \end{aligned} \quad (22)$$

A good weighting matrix \mathbf{W} should put more weights on moment conditions that can be estimated more precisely. The asymptotically efficient weighting matrix suggested by the *Generalized Method of Moments* [19] is

$$\mathbf{W}^{-1} = \mathbb{E} \left[\mathbf{g}([\boldsymbol{\theta}_k]_{k=1}^K, \mathbf{x}) \mathbf{g}([\boldsymbol{\theta}_k]_{k=1}^K, \mathbf{x})^T \right] \approx \frac{1}{T} \sum_{t=1}^T \mathbf{g}([\boldsymbol{\theta}_k]_{k=1}^K, \mathbf{x}) \mathbf{g}([\boldsymbol{\theta}_k]_{k=1}^K, \mathbf{x})^T \quad (23)$$

Theorem 2 (Gen.MM is asymptotically efficient [19]). *Let $g_n(\boldsymbol{\theta}, \mathbf{X}) := \sum_k f_n(\boldsymbol{\theta}_k) - h_n(\mathbf{X})$ so that $\mathbb{E}[h_n(\mathbf{X})] = \mathbb{E}[\phi_n(\mathbf{x})]$. Let $\mathbf{W}^{-1} = \mathbb{E}[\mathbf{g}(\boldsymbol{\theta}, \mathbf{X}) \mathbf{g}(\boldsymbol{\theta}, \mathbf{X})^T] \approx \frac{1}{T} \sum_{t=1}^T \mathbf{g}(\boldsymbol{\theta}, \mathbf{X}_t) \mathbf{g}(\boldsymbol{\theta}, \mathbf{X}_t)^T$ Iterative Gen.MM is efficient with this weighting matrix \mathbf{W} .*

D Separability

For conditional models, the coefficients of the moment polynomials can depend on the data but such dependence can sometimes break the process of converting from component moment constraints to mixture moment constraints. In this section, we define separability, which is a sufficient condition on what dependence is allowed under Polymom and then we give some counterexamples.

Consider a mixture of linear regressions [20, 15], where the parameters $\boldsymbol{\theta}_k = (w_k, \sigma_k^2)$ are the slope and noise variance for each component k . Then each data point $\mathbf{x} = [x, y]$ is drawn from component k by sampling x from an unknown distribution independent of k and setting $y = w_k x + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma_k^2)$. If we take

observation function $\phi_{b,c}(\mathbf{x}) = x^b y^c$, then the corresponding $f_{b,c}(\boldsymbol{\theta})$ depends on the unknown distribution of x : for example, $f_{1,2}(\boldsymbol{\theta}) = \mathbb{E}[x^3]w + \mathbb{E}[x]\sigma^2$. In contrast, for the mixture of Gaussians, we had $f_2(\boldsymbol{\theta}) = \mu^2 + \sigma^2$, which only depends on the parameters.

However, not all is lost, since the key thing is that $f_{1,2}(\boldsymbol{\theta})$ depends only on the distribution of x , which is independent of the component k and furthermore can be estimated from data. More generally, we allow f_n to depend on \mathbf{x} but in a restricted way. We say that $f_n(\boldsymbol{\theta}, \mathbf{x})$ is *separable* if $\mathbb{E}[f_n(\boldsymbol{\theta}, \mathbf{x})]$ does not depend on the parameters $\{\boldsymbol{\theta}_k\}_{k=1}^K$ of the mixture generating \mathbf{x} . In other words,

$$\mathbb{E}[\phi_n(\mathbf{x})] = \mathbb{E}[f_n(\boldsymbol{\theta}, \mathbf{x})] \text{ where for all } k : \mathbb{E}_{\mathbf{x}|z=k}[f(\boldsymbol{\theta}; \mathbf{x})] = f(\boldsymbol{\theta}). \mathbb{E}[f_n(\boldsymbol{\theta}, \mathbf{x}) | z = k] = \mathbb{E}[f_n(\boldsymbol{\theta}, \mathbf{x})] \in \mathbb{R}[\boldsymbol{\theta}]. \quad (24)$$

In this case, we can define $f_n(\boldsymbol{\theta}) := \mathbb{E}[f_n(\boldsymbol{\theta}, \mathbf{x})]$, and (3) is still valid. For the mixture of linear regressions, we would define $f_{b,c}(\boldsymbol{\theta}, \mathbf{x}) = x^b \mathbb{E}_{\epsilon \sim \mathcal{N}(0, \sigma^2)}[(wx + \epsilon)^c]$. In this more general setup, the approximate moment equations on T data points is $\frac{1}{T} \sum_{t=1}^T [f_n(\boldsymbol{\theta}, \mathbf{x}_t)] = \frac{1}{T} \sum_{t=1}^T \phi_n(\mathbf{x}_t)$.

An example of non-separability is a mixture of linear regressions where the variance is not a parameter and is different across mixture components: $\boldsymbol{\theta} = (w)$ and $\mathbf{x} = (x, y)$. Recall that $\mathbb{E}[xy^2] = \sum_{k=1}^K \pi_k (\mathbb{E}[x^3]w_k^2 + \mathbb{E}[x]\sigma_k^2)$, but $\mathbb{E}[x^3]w_k^2 + \mathbb{E}[x]\sigma_k^2$ cannot be written as $\mathbb{E}[f_n(w_k, \mathbf{x})]$ for any f_n , since it depends on σ_k^2 . Thus, this example falls outside our framework. In the simplest case, we can make $f_n(w, \mathbf{x})$ separable by introducing σ_k as a parameter, but this is not always possible if the noise distribution is unknown or if $\sigma_k(x)$ depends on x . For example, if we have heteroskedastic noise, $\mathbb{E}[x^a(y - w \cdot x)] = 0$ are valid moment constraints for individual components, but it is not clear how to convert this to the mixture case.

References

- [1] R. E. Curto and L. A. Fialkow. *Solution of the truncated complex moment problem for flat data*, volume 568. American Mathematical Society, 1996 1996.
- [2] R. E. Curto and L. A. Fialkow. *Flat extensions of positive moment matrices: Recursively generated relations*, volume 648. American Mathematical Society, 1998.
- [3] R. Curto and L. Fialkow. The truncated complex K-moment problem. *Transactions of the American mathematical society*, 352(6):2825–2855, 2000.
- [4] R. E. Curto and L. A. Fialkow. Truncated K-moment problems in several variables. *arXiv preprint arXiv:math/0507067*, 2005.
- [5] M. Laurent. A sparse flat extension theorem for moment matrices. *arXiv preprint arXiv:0812.2563*, 2008.
- [6] M. Laurent and B. Mourrain. A generalized flat extension theorem for moment matrices. *Archiv der Mathematik*, 93(1):87–98, 2009.
- [7] J. Nie. Certifying convergence of lasserres hierarchy via flat truncation. *Mathematical Programming*, 142(1):485–510, 2013.
- [8] J. Nie. Optimality conditions and finite convergence of lasserres hierarchy. *Mathematical programming*, 146(1):97–121, 2014.
- [9] J. Nie. The A-truncated K-moment problem. *Foundations of Computational Mathematics*, 14(6):1243–1276, 2014.
- [10] J. Nie. Linear optimization with cones of moments and nonnegative polynomials. *Mathematical Programming*, pages 1–28, 2013.
- [11] J. Dattorro. *Convex Optimization and Euclidean Distance Geometry*. Meboo, 2005 2005.
- [12] A. Anandkumar, R. Ge, D. Hsu, S. M. Kakade, and M. Telgarsky. Tensor decompositions for learning latent variable models. *arXiv*, 2013.
- [13] A. Anandkumar, D. Hsu, and S. M. Kakade. A method of moments for mixture models and hidden Markov models. In *Conference on Learning Theory (COLT)*, 2012.
- [14] A. Anandkumar, R. Ge, and M. Janzamin. Sample complexity analysis for learning overcomplete latent variable models through tensor methods. *arXiv preprint arXiv:1408.0553*, 2014.
- [15] A. Chaganty and P. Liang. Spectral experts for estimating mixtures of linear regressions. In *International Conference on Machine Learning (ICML)*, 2013.
- [16] V. Kuleshov, A. Chaganty, and P. Liang. Tensor factorization via matrix factorization. In *Artificial Intelligence and Statistics (AISTATS)*, 2015.
- [17] D. Koller and N. Friedman. *Probabilistic graphical models: principles and techniques*. MIT Press, 2009.

- [18] S. B. Cohen, K. Stratos, M. Collins, D. P. Foster, and L. H. Ungar. Experiments with spectral learning of latent-variable pcfgs. In *North American Association for Computational Linguistics (NAACL)*, pages 148–157, 2013.
- [19] L. P. Hansen. Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, 50:1029–1054, 1982.
- [20] K. Viele and B. Tong. Modeling with mixtures of linear regressions. *Statistics and Computing*, 12(4):315–330, 2002.