Learning Diagonal Gaussian Mixture Models and Incomplete Tensor Decompositions

Bingni Guo¹ · Jiawang Nie¹ · Zi Yang¹

Received: 8 February 2021 / Accepted: 27 May 2021 / Published online: 18 November 2021
© The Author(s) 2021

Abstract

This paper studies how to learn parameters in diagonal Gaussian mixture models. The problem can be formulated as computing incomplete symmetric tensor decompositions. We use generating polynomials to compute incomplete symmetric tensor decompositions and approximations. Then the tensor approximation method is used to learn diagonal Gaussian mixture models. We also do the stability analysis. When the first and third order moments are sufficiently accurate, we show that the obtained parameters for the Gaussian mixture models are also highly accurate. Numerical experiments are also provided.

Keywords Gaussian model · Tensor · Decomposition · Generating polynomial · Moments

Mathematics Subject Classification (2010) 15A69 · 65F99 · 65K10

1 Introduction

A Gaussian mixture model consists of several component Gaussian distributions. For given samples of a Gaussian mixture model, people often need to estimate parameters for each component Gaussian distribution [24, 32]. Consider a Gaussian mixture model with r components. For each $i \in [r] := \{1, \dots, r\}$, let ω_i be the positive probability for the i th component Gaussian to appear in the mixture model. We have each $\omega_i > 0$ and $\sum_{i=1}^r \omega_i = 1$. Suppose the i th Gaussian distribution is $\mathcal{N}(\mu_i, \Sigma_i)$, where $\mu_i \in \mathbb{R}^d$ is the

The article is dedicated to Professor Bernd Sturmfels' 60th birthday.

✉ Jiawang Nie
njw@math.ucsd.edu

Bingni Guo
b8guo@ucsd.edu

Zi Yang
ziy109@ucsd.edu

¹ Department of Mathematics, University of California San Diego, 9500 Gilman Drive, La Jolla, CA, 92093, USA

expectation (or mean) and $\Sigma_i \in \mathbb{R}^{d \times d}$ is the covariance matrix. Let $y \in \mathbb{R}^d$ be the random vector for the Gaussian mixture model and let $y_1, \dots, y_N$ be identically independent distributed (i.i.d) samples from the mixture model. Each y_j is sampled from one of the r component Gaussian distributions, associated with a label $Z_j \in [r]$ indicating the component that it is sampled from. The probability that a sample comes from the i th component is ω_i . When people observe only samples without labels, the Z_j 's are called latent variables. The density function for the random variable y is

$$f(y) := \sum_{i=1}^r \omega_i \frac{1}{\sqrt{(2\pi)^d \det \Sigma_i}} \exp \left\{ -\frac{1}{2} (y - \mu_i)^T \Sigma_i^{-1} (y - \mu_i) \right\},$$

where μ_i is the mean and Σ_i is the covariance matrix for the i th component.

Learning a Gaussian mixture model is to estimate the parameters $\omega_i, \mu_i, \Sigma_i$ for each $i \in [r]$, from given samples of y . The number of parameters in a covariance matrix grows quadratically with respect to the dimension. Due to the curse of dimensionality, the computation becomes very expensive for large d [35]. Hence, diagonal covariance matrices are preferable in applications. In this paper, we focus on learning Gaussian mixture models with diagonal covariance matrices, i.e.,

$$\Sigma_i = \text{diag}(\sigma_{i1}^2, \dots, \sigma_{id}^2), \quad i = 1, \dots, r.$$

A natural approach for recovering the unknown parameters $\omega_i, \mu_i, \Sigma_i$ is the method of moments. It estimates parameters by solving a system of multivariate polynomial equations, from moments of the random vector y . Directly solving polynomial systems may encounter non-existence or non-uniqueness of statistically meaningful solutions [57]. However, for diagonal Gaussians, the third order moment tensor can help us avoid these troubles.

Let $M_3 := \mathbb{E}(y \otimes y \otimes y)$ be the third order moment tensor for y . One can write that $y = \eta(z) + \zeta(z)$, where z is a discrete random variable such that $\text{Prob}(z = i) = \omega_i$, $\eta(i) = \mu_i \in \mathbb{R}^d$ and $\zeta(i)$ is the random variable ζ_i obeying the Gaussian distribution $\mathcal{N}(0, \Sigma_i)$. Assume all Σ_i are diagonal, then

$$\begin{aligned} M_3 &= \sum_{i=1}^r \omega_i \mathbb{E}[(\eta(i) + \zeta_i)^{\otimes 3}] \\ &= \sum_{i=1}^r \omega_i (\mu_i \otimes \mu_i \otimes \mu_i + \mathbb{E}[\mu_i \otimes \zeta_i \otimes \zeta_i] + \mathbb{E}[\zeta_i \otimes \mu_i \otimes \zeta_i] + \mathbb{E}[\zeta_i \otimes \zeta_i \otimes \mu_i]). \end{aligned}$$

The second equality holds because ζ_i has zero mean and

$$\mathbb{E}[\zeta_i \otimes \zeta_i \otimes \zeta_i] = \mathbb{E}[\mu_i \otimes \mu_i \otimes \zeta_i] = \mathbb{E}[\zeta_i \otimes \mu_i \otimes \mu_i] = \mathbb{E}[\mu_i \otimes \zeta_i \otimes \mu_i] = 0.$$

The random variable ζ_i has a diagonal covariance matrix, so $\mathbb{E}[(\zeta_i)_j (\zeta_i)_l] = 0$ for $j \neq l$. Therefore,

$$\sum_{i=1}^r \omega_i \mathbb{E}[\mu_i \otimes \zeta_i \otimes \zeta_i] = \sum_{i=1}^r \sum_{j=1}^d \omega_i \sigma_{ij}^2 \mu_i \otimes e_j \otimes e_j = \sum_{j=1}^d a_j \otimes e_j \otimes e_j,$$

where the vectors a_j are given by

$$a_j := \sum_{i=1}^r \omega_i \sigma_{ij}^2 \mu_i, \quad j = 1, \dots, d. \quad (1.1)$$

Similarly, we have

$$\sum_{i=1}^r \omega_i \mathbb{E}[\zeta_i \otimes \mu_i \otimes \zeta_i] = \sum_{j=1}^d e_j \otimes a_j \otimes e_j, \quad \sum_{i=1}^r \omega_i \mathbb{E}[\zeta_i \otimes \zeta_i \otimes \mu_i] = \sum_{j=1}^d e_j \otimes e_j \otimes a_j.$$

Therefore, we can express M_3 in terms of ω_i , μ_i , Σ_i as

$$M_3 = \sum_{i=1}^r \omega_i \mu_i \otimes \mu_i \otimes \mu_i + \sum_{j=1}^d (a_j \otimes e_j \otimes e_j + e_j \otimes a_j \otimes e_j + e_j \otimes e_j \otimes a_j). \quad (1.2)$$

We are particularly interested in the following third order symmetric tensor

$$\mathcal{F} := \sum_{i=1}^r \omega_i \mu_i \otimes \mu_i \otimes \mu_i. \quad (1.3)$$

When the labels i_1, i_2, i_3 are distinct from each other, we have

$$(M_3)_{i_1 i_2 i_3} = (\mathcal{F})_{i_1 i_2 i_3} \quad \text{for } i_1 \neq i_2 \neq i_3 \neq i_1.$$

Denote the label set

$$\Omega = \{(i_1, i_2, i_3) : i_1 \neq i_2 \neq i_3 \neq i_1, i_1, i_2, i_3 \text{ are labels for } M_3\}. \quad (1.4)$$

The tensor M_3 can be estimated from the samplings for y , so the entries $\mathcal{F}_{i_1 i_2 i_3}$ with $(i_1, i_2, i_3) \in \Omega$ can also be obtained from the estimation of M_3 . To recover the parameters ω_i , μ_i , we first find the tensor decomposition for $\mathcal{F}$, from the partially given entries $\mathcal{F}_{i_1 i_2 i_3}$ with $(i_1, i_2, i_3) \in \Omega$. Once the parameters ω_i , μ_i are known, we can determine Σ_i from the expressions of a_j as in (1.1).

The above observation leads to the incomplete tensor decomposition problem. For a third order symmetric tensor $\mathcal{F}$ whose partial entries $\mathcal{F}_{i_1 i_2 i_3}$ with $(i_1, i_2, i_3) \in \Omega$ are known, we are looking for vectors $p_1, \dots, p_r$ such that

$$\mathcal{F}_{i_1 i_2 i_3} = \left(p_1^{\otimes 3} + \dots + p_r^{\otimes 3} \right)_{i_1 i_2 i_3} \quad \text{for all } (i_1, i_2, i_3) \in \Omega.$$

The above is called an incomplete tensor decomposition for $\mathcal{F}$. To find such a tensor decomposition for $\mathcal{F}$, a straightforward approach is to do tensor completion: first find unknown tensor entries $\mathcal{F}_{i_1 i_2 i_3}$ with $(i_1, i_2, i_3) \notin \Omega$ such that the completed $\mathcal{F}$ has low rank, and then compute the tensor decomposition for $\mathcal{F}$. However, there are serious disadvantages for this approach. The theory for tensor completion or recovery, especially for symmetric tensors, is premature. Low rank tensor completion or recovery is typically not guaranteed by the currently existing methodology. Most methods for tensor completion are based on convex relaxations, e.g., the nuclear norm or trace minimization [22, 36, 41, 54, 58]. These convex relaxations may not produce low rank completions [51].

In this paper, we propose a new method for determining incomplete tensor decompositions. It is based on the generating polynomial method in [40]. The label set Ω consists of (i_1, i_2, i_3) of distinct i_1, i_2, i_3 . We can still determine some generating polynomials, from the partially given tensor entries $\mathcal{F}_{i_1 i_2 i_3}$ with $(i_1, i_2, i_3) \in \Omega$. They can be used to get the incomplete tensor decomposition. We show that this approach works very well when the rank r is roughly not more than half of the dimension d . Consequently, the parameters for the Gaussian mixture model can be recovered from the incomplete tensor decomposition of $\mathcal{F}$.

Related Work Gaussian mixture models have broad applications in machine learning problems, e.g., automatic speech recognition [30, 48, 50], hyperspectral unmixing problem [4, 34], background subtraction [32, 60] and anomaly detection [56]. They also have applications in social and biological sciences [25, 53, 59].

There exist methods for estimating unknown parameters for Gaussian mixture models. A popular method is the expectation-maximization (EM) algorithm that iteratively approximates the maximum likelihood parameter estimation [16]. This approach is widely used in applications, while its convergence property is not very reliable [49]. Dasgupta [13] introduced a method that first projects data to a randomly chosen low-dimensional subspace and then use the empirical means and covariances of low-dimensional clusters to estimate the parameters. Later, Arora and Kannan [52] extended this idea to arbitrary Gaussians. Vempala and Wong [55] introduced the spectral technique to enhance the separation condition by projecting data to principal components of the sample matrix instead of selecting a random subspace. For other subsequent work, we refer to Dasgupta and Schulman [14], Kannan et al. [29], Achlioptas et al. [1], Chaudhuri and Rao [10], Brubaker and Vempala [7] and Chaudhuri et al. [9].

Another frequently used approach is based on moments, introduced by Pearson [46]. Belkin and Sinha [3] proposed a learning algorithm for identical spherical Gaussians ($\Sigma_i = \sigma^2 I$) with arbitrarily small separation between mean vectors. It was also shown in [28] that a mixture of two Gaussians can be learned with provably minimal assumptions. Hsu and Kakade [27] provided a learning algorithm for a mixture of spherical Gaussians, i.e., each covariance matrix is a multiple of the identity matrix. This method is based on moments up to order three and only assumes non-degeneracy instead of separations. For general covariance matrices, Ge et al. [23] proposed a learning method when the dimension d is sufficiently high. More moment-based methods for general latent variable models can be found in [2].

Contributions This paper proposes a new method for learning diagonal Gaussian mixture models, based on samplings for the first and third order moments. Let $y_1, \dots, y_N$ be samples and let $\{(\omega_i, \mu_i, \Sigma_i) : i \in [r]\}$ be parameters of the diagonal Gaussian mixture model, where each covariance matrix Σ_i is diagonal. We use the samples $y_1, \dots, y_N$ to estimate the third order moment tensor M_3 , as well as the mean vector M_1 . We have seen that the tensor M_3 can be expressed as in (1.2).

For the tensor $\mathcal{F}$ in (1.3), we have $\mathcal{F}_{i_1 i_2 i_3} = (M_3)_{i_1 i_2 i_3}$ when the labels i_1, i_2, i_3 are distinct from each other. Other entries of $\mathcal{F}$ are not known, since the vectors a_j are not available. The $\mathcal{F}$ is an incompletely given tensor. We give a new method for computing the incomplete tensor decomposition of $\mathcal{F}$ when the rank r is low (roughly no more than half of the dimension d). The tensor decomposition of $\mathcal{F}$ is unique under some genericity conditions [11], so it can be used to recover parameters ω_i, μ_i . To compute the incomplete tensor decomposition of $\mathcal{F}$, we use the generating polynomial method in [40, 42]. We look for a special set of generating polynomials for $\mathcal{F}$, which can be obtained by solving linear least squares. It only requires to use the known entries of $\mathcal{F}$. The common zeros of these generating polynomials can be determined from eigenvalue decompositions. Under some genericity assumptions, these common zeros can be used to get the incomplete tensor decomposition. After this is done, the parameters ω_i, μ_i can be recovered by solving linear systems. The diagonal covariance matrices Σ_i can also be estimated by solving linear least squares. The tensor M_3 is estimated from the samples $y_1, \dots, y_N$. Typically, the tensor entries $(M_3)_{i_1 i_2 i_3}$ and $\mathcal{F}_{i_1 i_2 i_3}$, are not precisely given. We also provide a stability analysis for this case,

showing that the estimated parameters are also accurate when the entries $(M_3)_{i_1 i_2 i_3}$ have small errors.

The paper is organized as follows. In Section 2, we review some basic results for symmetric tensor decompositions and generating polynomials. In Section 3, we give a new algorithm for computing an incomplete tensor decomposition for $\mathcal{F}$, when only its subtensor $\mathcal{F}_\Omega$ is known. Section 4 gives the stability analysis when there are errors for the subtensor $\mathcal{F}_\Omega$. Section 5 gives the algorithm for learning Gaussian mixture models. Numerical experiments and applications are given in Section 6. We make some conclusions and discussions in Section 7.

2 Preliminary

Notation Denote $\mathbb{N}$, $\mathbb{C}$ and $\mathbb{R}$ the set of nonnegative integers, complex and real numbers respectively. Denote the cardinality of a set L as $|L|$. Denote by e_i the i th standard unit basis vector, i.e., the i th entry of e_i is one and all others are zeros. For a complex number c , $\sqrt[n]{c}$ or $c^{1/n}$ denotes the principal n th root of c . For a complex vector v , $\operatorname{Re}(v)$, $\operatorname{Im}(v)$ denotes the real part and imaginary part of v respectively. A property is said to be generic if it is true in the whole space except a subset of zero Lebesgue measure. The $\|\cdot\|$ denotes the Euclidean norm of a vector or the Frobenius norm of a matrix. For a vector or matrix, the superscript T denotes the transpose and H denotes the conjugate transpose. For $i, j \in \mathbb{N}$, $[i]$ denotes the set $\{1, 2, \dots, i\}$ and $[i, j]$ denotes the set $\{i, i+1, \dots, j\}$ if $i \leq j$. For a vector v , $v_{i_1:i_2}$ denotes the vector $(v_{i_1}, v_{i_1+1}, \dots, v_{i_2})$. For a matrix A , denote by $A_{[i_1:i_2, j_1:j_2]}$ the submatrix of A whose row labels are $i_1, i_1+1, \dots, i_2$ and whose column labels are $j_1, j_1+1, \dots, j_2$. For a tensor $\mathcal{F}$, its subtensor $\mathcal{F}_{[i_1:i_2, j_1:j_2, k_1:k_2]}$ is similarly defined.

Let $S^m(\mathbb{C}^d)$ (resp., $S^m(\mathbb{R}^d)$) denote the space of m th order symmetric tensors over the vector space $\mathbb{C}^d$ (resp., $\mathbb{R}^d$). For convenience of notation, the labels for tensors start with 0. A symmetric tensor $\mathcal{A} \in S^m(\mathbb{C}^{n+1})$ is labelled as

$$\mathcal{A} = (\mathcal{A}_{i_1 \dots i_m})_{0 \leq i_1, \dots, i_m \leq n},$$

where the entry $\mathcal{A}_{i_1 \dots i_m}$ is invariant for all permutations of $(i_1, \dots, i_m)$. The Hilbert–Schmidt norm $\|\mathcal{A}\|$ is defined as

$$\|\mathcal{A}\| := \left(\sum_{0 \leq i_1, \dots, i_m \leq n} |\mathcal{A}_{i_1 \dots i_m}|^2 \right)^{1/2}.$$

The norm of a subtensor $\|\mathcal{A}_\Omega\|$ is similarly defined. For a vector $u := (u_0, u_1, \dots, u_n) \in \mathbb{C}^d$, the tensor power $u^{\otimes m} := u \otimes \dots \otimes u$, where u is repeated m times, is defined such that

$$(u^{\otimes m})_{i_1 \dots i_m} = u_{i_1} \times \dots \times u_{i_m}.$$

For a symmetric tensor $\mathcal{F}$, its symmetric rank is

$$\operatorname{rank}_S(\mathcal{F}) := \min \left\{ r \mid \mathcal{F} = \sum_{i=1}^r u_i^{\otimes m} \right\}.$$

There are other types of tensor ranks [31, 33]. In this paper, we only deal with symmetric tensors and symmetric ranks. We refer to [12, 17, 21, 26, 31, 33] for general work about tensors and their ranks. For convenience, if $r = \operatorname{rank}_S(\mathcal{F})$, we call $\mathcal{F}$ a rank- r tensor and $\mathcal{F} = \sum_{i=1}^r u_i^{\otimes m}$ is called a rank decomposition.

For a power $\alpha := (\alpha_1, \alpha_2, \dots, \alpha_n) \in \mathbb{N}^n$ and $x := (x_1, x_2, \dots, x_n)$, denote

$$|\alpha| := \alpha_1 + \alpha_2 + \dots + \alpha_n, \quad x^\alpha := x_1^{\alpha_1} x_2^{\alpha_2} \dots x_n^{\alpha_n}, \quad x_0 := 1.$$

The monomial power set of degree m is denoted as

$$\mathbb{N}_m^n := \{\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n) \in \mathbb{N}^n : |\alpha| \leq m\}.$$

For $\alpha \in \mathbb{N}_3^n$, we can write that $x^\alpha = x_{i_1} x_{i_2} x_{i_3}$ for some $0 \leq i_1, i_2, i_3 \leq n$.

Let $\mathbb{C}[x]_m$ be the space of all polynomials in x with complex coefficients and whose degrees are no more than m . For a cubic polynomial $p \in \mathbb{C}[x]_3$ and $\mathcal{F} \in \mathbb{S}^3(\mathbb{C}^{n+1})$, we define the bilinear product (note that $x_0 = 1$)

$$\langle p, \mathcal{F} \rangle = \sum_{0 \leq i_1, i_2, i_3 \leq n} p_{i_1 i_2 i_3} \mathcal{F}_{i_1 i_2 i_3} \quad \text{for } p = \sum_{0 \leq i_1, i_2, i_3 \leq n} p_{i_1 i_2 i_3} x_{i_1} x_{i_2} x_{i_3},$$

where $p_{i_1 i_2 i_3}$ are coefficients of p . A polynomial $g \in \mathbb{C}[x]_3$ is called a *generating polynomial* for a symmetric tensor $\mathcal{F} \in \mathbb{S}^3(\mathbb{C}^{n+1})$ if

$$\langle g \cdot x^\beta, \mathcal{F} \rangle = 0 \quad \forall \beta \in \mathbb{N}_{3-\deg(g)}^n,$$

where $\deg(g)$ denotes the degree of g in x . When the order is bigger than 3, we refer to [40] for the definition of generating polynomials. They can be used to compute symmetric tensor decompositions and low rank approximations [40, 42], which are closely related to truncated moment problems and polynomial optimization [20, 37–39, 43]. There are special versions of symmetric tensors and their decompositions [19, 44, 45].

3 Incomplete Tensor Decomposition

This section discusses how to compute an incomplete tensor decomposition for a symmetric tensor $\mathcal{F} \in \mathbb{S}^3(\mathbb{C}^d)$ when only its subtensor $\mathcal{F}_\Omega$ is given, for the label set Ω in (1.4). For convenience of notation, the labels for $\mathcal{F}$ begin with zeros while a vector $u \in \mathbb{C}^d$ is still labelled as $u := (u_1, \dots, u_d)$. We set

$$n := d - 1, \quad x = (x_1, \dots, x_n), \quad x_0 := 1.$$

For a given rank r , denote the monomial sets

$$\mathcal{B}_0 := \{x_1, \dots, x_r\}, \quad \mathcal{B}_1 := \{x_i x_j : i \in [r], j \in [r+1, n]\}.$$

For a monomial power $\alpha \in \mathbb{N}^n$, by writing $\alpha \in \mathcal{B}_1$, we mean that $x^\alpha \in \mathcal{B}_1$. For each $\alpha \in \mathcal{B}_1$, one can write $\alpha = e_i + e_j$ with $i \in [r]$, $j \in [r+1, n]$. Let $\mathbb{C}^{[r] \times \mathcal{B}_1}$ denote the space of matrices labelled by the pair $(k, \alpha) \in [r] \times \mathcal{B}_1$. For each $\alpha = e_i + e_j \in \mathcal{B}_1$ and $G \in \mathbb{C}^{[r] \times \mathcal{B}_1}$, denote the quadratic polynomial in x

$$\varphi_{ij}[G](x) := \sum_{k=1}^r G(k, e_i + e_j) x_k - x_i x_j. \quad (3.1)$$

Suppose r is the symmetric rank of $\mathcal{F}$. A matrix $G \in \mathbb{C}^{[r] \times \mathcal{B}_1}$ is called a *generating matrix* of $\mathcal{F}$ if each $\varphi_{ij}[G](x)$, with $\alpha = e_i + e_j \in \mathcal{B}_1$, is a generating polynomial of $\mathcal{F}$. Equivalently, G is a generating matrix of $\mathcal{F}$ if and only if

$$\langle x_t \varphi_{ij}[G](x), \mathcal{F} \rangle = \sum_{k=1}^r G(k, e_i + e_j) \mathcal{F}_{0k t} - \mathcal{F}_{i j t} = 0, \quad t = 0, 1, \dots, n, \quad (3.2)$$

for all $i \in [r]$, $j \in [r+1, n]$. The notion *generating matrix* is motivated from the fact that the entire tensor $\mathcal{F}$ can be recursively determined by G and its first r entries (see [40]). The existence and uniqueness of the generating matrix G is shown as follows.

Theorem 3.1 Suppose $\mathcal{F}$ has the decomposition

$$\mathcal{F} = \lambda_1 \begin{bmatrix} 1 \\ u_1 \end{bmatrix}^{\otimes 3} + \cdots + \lambda_r \begin{bmatrix} 1 \\ u_r \end{bmatrix}^{\otimes 3}, \quad (3.3)$$

for vectors $u_i \in \mathbb{C}^n$ and scalars $0 \neq \lambda_i \in \mathbb{C}$. If the subvectors $(u_1)_{1:r}, \dots, (u_r)_{1:r}$ are linearly independent, then there exists a unique generating matrix $G \in \mathbb{C}^{[r] \times \mathcal{B}_1}$ satisfying (3.2) for the tensor $\mathcal{F}$.

Proof We first prove the existence. For each $i = 1, \dots, r$, denote the vectors $v_i = (u_i)_{1:r}$. Under the given assumption, $V := [v_1 \dots v_r]$ is an invertible matrix. For each $l = r+1, \dots, n$, let

$$N_l := V \cdot \text{diag}((u_1)_l, \dots, (u_r)_l) \cdot V^{-1}. \quad (3.4)$$

Then $N_l v_i = (u_i)_l v_i$ for $i = 1, \dots, r$, i.e., N_l has eigenvalues $(u_1)_l, \dots, (u_r)_l$ with corresponding eigenvectors $(u_1)_{1:r}, \dots, (u_r)_{1:r}$. We select $G \in \mathbb{C}^{[r] \times \mathcal{B}_1}$ to be the matrix such that

$$N_l = \begin{bmatrix} G(1, e_1 + e_l) & \cdots & G(r, e_1 + e_l) \\ \vdots & \ddots & \vdots \\ G(1, e_r + e_l) & \cdots & G(r, e_r + e_l) \end{bmatrix}, \quad l = r+1, \dots, n.$$

For each $s = 1, \dots, r$ and $\alpha = e_i + e_j \in \mathcal{B}_1$ with $i \in [r]$, $j \in [r+1, n]$,

$$\varphi_{ij}[G](u_s) = \sum_{k=1}^r G(k, e_i + e_j)(u_s)_k - (u_s)_i(u_s)_j = 0.$$

For each $t = 1, \dots, n$, it holds that

$$\begin{aligned} \langle x_t \varphi_{ij}[G](x), \mathcal{F} \rangle &= \left\langle \sum_{k=1}^r G(k, e_i + e_j) x_t x_k - x_t x_i x_j, \mathcal{F} \right\rangle \\ &= \left\langle \sum_{k=1}^r G(k, e_i + e_j) x_t x_k - x_t x_i x_j, \sum_{s=1}^r \lambda_s \begin{bmatrix} 1 \\ u_s \end{bmatrix}^{\otimes 3} \right\rangle \\ &= \sum_{k=1}^r G(k, e_i + e_j) \sum_{s=1}^r \lambda_s (u_s)_t (u_s)_k - \sum_{s=1}^r \lambda_s (u_s)_t (u_s)_i (u_s)_j \\ &= \sum_{s=1}^r \lambda_s (u_s)_t \left(\sum_{k=1}^r G(k, e_i + e_j)(u_s)_k - (u_s)_i (u_s)_j \right) \\ &= 0. \end{aligned}$$

When $t = 0$, we can similarly get

$$\begin{aligned} \langle \varphi_{ij}[G](x), \mathcal{F} \rangle &= \left\langle \sum_{k=1}^r G(k, e_i + e_j) x_k - x_i x_j, \mathcal{F} \right\rangle \\ &= \sum_{s=1}^r \lambda_s \left(\sum_{k=1}^r G(k, e_i + e_j)(u_s)_k - (u_s)_i (u_s)_j \right) \\ &= 0. \end{aligned}$$

Therefore, the matrix G satisfies (3.2) and it is a generating matrix for $\mathcal{F}$.

Second, we prove the uniqueness of such G . For each $\alpha = e_i + e_j \in \mathcal{B}_1$, let

$$F := \begin{bmatrix} \mathcal{F}_{011} & \cdots & \mathcal{F}_{0r1} \\ \vdots & \ddots & \vdots \\ \mathcal{F}_{01n} & \cdots & \mathcal{F}_{0rn} \end{bmatrix}, \quad g_{ij} := \begin{bmatrix} \mathcal{F}_{1ij} \\ \vdots \\ \mathcal{F}_{nij} \end{bmatrix}.$$

Since G satisfies (3.2), we have $F \cdot G(:, e_i + e_j) = g_{ij}$. The decomposition (3.3) implies that

$$F = [u_1 \cdots u_r] \cdot \text{diag}(\lambda_1, \dots, \lambda_r) \cdot [v_1 \cdots v_r]^T.$$

The sets $\{v_1, \dots, v_r\}$ and $\{u_1, \dots, u_r\}$ are both linearly independent. Since each $\lambda_i \neq 0$, the matrix F has full column rank. Hence, the generating matrix G satisfying $F \cdot G(:, e_i + e_j) = g_{ij}$ for all $i \in [r]$, $j \in [r+1, n]$ is unique. $\square$

The following is an example of generating matrices.

Example 3.2 Consider the tensor $\mathcal{F} \in \mathbb{S}^3(\mathbb{C}^6)$ that is given as

$$\mathcal{F} = 0.4 \cdot (1, 1, 1, 1, 1, 1)^{\otimes 3} + 0.6 \cdot (1, -1, 2, -1, 2, 3)^{\otimes 3}.$$

The rank $r = 2$, $\mathcal{B}_0 = \{x_1, x_2\}$ and $\mathcal{B}_1 = \{x_1x_3, x_1x_4, x_1x_5, x_2x_3, x_2x_4, x_2x_5\}$. We have the vectors

$$u_1 = (1, 1, 1, 1, 1), \quad u_2 = (-1, 2, -1, 2, 3), \quad v_1 = (1, 1), \quad v_2 = (-1, 2).$$

The matrices N_3, N_4, N_5 as in (3.4) are

$$\begin{aligned} N_3 &= \begin{bmatrix} 1 & -1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} 1 & -1 \\ 1 & 2 \end{bmatrix}^{-1} = \begin{bmatrix} 1/3 & 2/3 \\ 4/3 & -1/3 \end{bmatrix}, \\ N_4 &= \begin{bmatrix} 1 & -1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 1 & -1 \\ 1 & 2 \end{bmatrix}^{-1} = \begin{bmatrix} 4/3 & -1/3 \\ -2/3 & 5/3 \end{bmatrix}, \\ N_5 &= \begin{bmatrix} 1 & -1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} 1 & -1 \\ 1 & 2 \end{bmatrix}^{-1} = \begin{bmatrix} 5/3 & -2/3 \\ -4/3 & 7/3 \end{bmatrix}. \end{aligned}$$

The entries of the generating matrix G are listed as below:

$$\begin{array}{c|cccccc} k \backslash (i, j) & (1, 3) & (1, 4) & (1, 5) & (2, 3) & (2, 4) & (2, 5) \\ \hline 1 & 1/3 & 4/3 & 5/3 & 4/3 & -2/3 & -4/3 \\ 2 & 2/3 & -1/3 & -2/3 & -1/3 & 5/3 & 7/3 \end{array} \quad (3.5)$$

The generating polynomials in (3.1) are

$$\begin{aligned} \varphi_{13}[G](x) &= \frac{1}{3}x_1 + \frac{2}{3}x_2 - x_1x_3, & \varphi_{23}[G](x) &= \frac{4}{3}x_1 - \frac{1}{3}x_2 - x_2x_3, \\ \varphi_{14}[G](x) &= \frac{4}{3}x_1 - \frac{1}{3}x_2 - x_1x_4, & \varphi_{24}[G](x) &= -\frac{2}{3}x_1 + \frac{5}{3}x_2 - x_2x_4, \\ \varphi_{15}[G](x) &= \frac{5}{3}x_1 - \frac{2}{3}x_2 - x_1x_5, & \varphi_{25}[G](x) &= -\frac{4}{3}x_1 + \frac{7}{3}x_2 - x_2x_5. \end{aligned}$$

Above generating polynomials can be written in the following form

$$\begin{bmatrix} \varphi_{1j}[G](x) \\ \varphi_{2j}[G](x) \end{bmatrix} = N_j \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} - x_j \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \quad \text{for } j = 3, 4, 5.$$

For x to be a common zero of $\varphi_{1j}[G](x)$ and $\varphi_{2j}[G](x)$, it requires that (x_1, x_2) is an eigenvector of N_j with the corresponding eigenvalue x_j .

3.1 Computing the Tensor Decomposition

We show how to find an incomplete tensor decomposition (3.3) for $\mathcal{F}$ when only its sub-tensor $\mathcal{F}_\Omega$ is given, where the label set Ω is as in (1.4). Suppose that there exists the decomposition (3.3) for $\mathcal{F}$, for vectors $u_i \in \mathbb{C}^n$ and nonzero scalars $\lambda_i \in \mathbb{C}$. Assume the subvectors $(u_1)_{1:r}, \dots, (u_r)_{1:r}$ are linearly independent, so there is a unique generating matrix G for $\mathcal{F}$, by Theorem 3.1.

For each $\alpha = e_i + e_j \in \mathcal{B}_1$ with $i \in [r]$, $j \in [r+1, n]$ and for each

$$l = r+1, \dots, j-1, j+1, \dots, n,$$

the generating matrix G satisfies the equations

$$\left\langle x_l \left(\sum_{k=1}^r G(k, e_i + e_j) x_k - x_i x_j \right), \mathcal{F} \right\rangle = \sum_{k=1}^r G(k, e_i + e_j) \mathcal{F}_{0kl} - \mathcal{F}_{ijl} = 0. \quad (3.6)$$

Let the matrix $A_{ij}[\mathcal{F}] \in \mathbb{C}^{(n-r-1) \times r}$ and the vector $b_{ij}[\mathcal{F}] \in \mathbb{C}^{n-r-1}$ be such that

$$A_{ij}[\mathcal{F}] := \begin{bmatrix} \mathcal{F}_{0,1,r+1} & \cdots & \mathcal{F}_{0,r,r+1} \\ \vdots & \ddots & \vdots \\ \mathcal{F}_{0,1,j-1} & \cdots & \mathcal{F}_{0,r,j-1} \\ \mathcal{F}_{0,1,j+1} & \cdots & \mathcal{F}_{0,r,j+1} \\ \vdots & \ddots & \vdots \\ \mathcal{F}_{0,1,n} & \cdots & \mathcal{F}_{0,r,n} \end{bmatrix}, \quad b_{ij}[\mathcal{F}] := \begin{bmatrix} \mathcal{F}_{i,j,r+1} \\ \vdots \\ \mathcal{F}_{i,j,j-1} \\ \mathcal{F}_{i,j,j+1} \\ \vdots \\ \mathcal{F}_{i,j,n} \end{bmatrix}. \quad (3.7)$$

To distinguish changes in the labels of tensor entries of $\mathcal{F}$, the commas are inserted to separate labeling numbers.

The equations in (3.6) can be equivalently written as

$$A_{ij}[\mathcal{F}] \cdot G(:, e_i + e_j) = b_{ij}[\mathcal{F}]. \quad (3.8)$$

If the rank $r \leq \frac{d}{2} - 1$, then $n - r - 1 = d - r - 2 \geq r$. Thus, the number of rows is not less than the number of columns for matrices $A_{ij}[\mathcal{F}]$. If $A_{ij}[\mathcal{F}]$ has linearly independent columns, then (3.8) uniquely determines $G(:, \alpha)$. For such a case, the matrix G can be fully determined by the linear system (3.8). Let $N_{r+1}(G), \dots, N_n(G) \in \mathbb{C}^{r \times r}$ be the matrices given as

$$N_l(G) = \begin{bmatrix} G(1, e_1 + e_l) & \cdots & G(r, e_1 + e_l) \\ \vdots & \ddots & \vdots \\ G(1, e_r + e_l) & \cdots & G(r, e_r + e_l) \end{bmatrix}, \quad l = r+1, \dots, n. \quad (3.9)$$

As in the proof of Theorem 3.1, one can see that

$$N_l(G) \begin{bmatrix} (u_i)_1 \\ \vdots \\ (u_i)_r \end{bmatrix} = (u_i)_l \cdot \begin{bmatrix} (u_i)_1 \\ \vdots \\ (u_i)_r \end{bmatrix}, \quad l = r+1, \dots, n.$$

The above is equivalent to the equations

$$N_l(G)v_i = (w_i)_{l-r} \cdot v_i, \quad l = r+1, \dots, n,$$

for the vectors ($i = 1, \dots, r$)

$$v_i := (u_i)_{1:r}, \quad w_i := (u_i)_{r+1:n}. \quad (3.10)$$

Each v_i is a common eigenvector of the matrices $N_{r+1}(G), \dots, N_n(G)$ and $(w_i)_{l-r}$ is the associated eigenvalue of $N_l(G)$. These matrices may or may not have repeated eigenvalues. Therefore, we select a generic vector $\xi := (\xi_{r+1}, \dots, \xi_n)$ and let

$$N(\xi) := \xi_{r+1}N_{r+1} + \dots + \xi_n N_n. \quad (3.11)$$

The eigenvalues of $N(\xi)$ are $\xi^T w_1, \dots, \xi^T w_r$. When $w_1, \dots, w_r$ are distinct from each other and ξ is generic, the matrix $N(\xi)$ does not have a repeated eigenvalue and hence it has unique eigenvectors $v_1, \dots, v_r$, up to scaling. Let $\tilde{v}_1, \dots, \tilde{v}_r$ be unit length eigenvectors of $N(\xi)$. They are also common eigenvectors of $N_{r+1}(G), \dots, N_n(G)$. For each $i = 1, \dots, r$, let $\tilde{w}_i$ be the vector such that its j th entry $(\tilde{w}_i)_j$ is the eigenvalue of $N_{j+r}(G)$, associated to the eigenvector $\tilde{v}_i$, or equivalently,

$$\tilde{w}_i = \left(\tilde{v}_i^H N_{r+1}(G) \tilde{v}_i, \dots, \tilde{v}_i^H N_n(G) \tilde{v}_i \right), \quad i = 1, \dots, r. \quad (3.12)$$

Up to a permutation of $(\tilde{v}_1, \dots, \tilde{v}_r)$, there exist scalars γ_i such that

$$v_i = \gamma_i \tilde{v}_i, \quad w_i = \tilde{w}_i. \quad (3.13)$$

The tensor decomposition of $\mathcal{F}$ can also be written as

$$\mathcal{F} = \lambda_1 \begin{bmatrix} 1 \\ \gamma_1 \tilde{v}_1 \\ \tilde{w}_1 \end{bmatrix}^{\otimes 3} + \dots + \lambda_r \begin{bmatrix} 1 \\ \gamma_r \tilde{v}_r \\ \tilde{w}_r \end{bmatrix}^{\otimes 3}.$$

The scalars $\lambda_1, \dots, \lambda_r$ and $\gamma_1, \dots, \gamma_r$ satisfy the linear equations

$$\begin{aligned} \lambda_1 \gamma_1 \tilde{v}_1 \otimes \tilde{w}_1 + \dots + \lambda_r \gamma_r \tilde{v}_r \otimes \tilde{w}_r &= \mathcal{F}_{[0,1:r,r+1:n]}, \\ \lambda_1 \gamma_1^2 \tilde{v}_1 \otimes \tilde{v}_1 \otimes \tilde{w}_1 + \dots + \lambda_r \gamma_r^2 \tilde{v}_r \otimes \tilde{v}_r \otimes \tilde{w}_r &= \mathcal{F}_{[1:r,1:r,r+1:n]}. \end{aligned}$$

Denote the label sets

$$\begin{aligned} J_1 &:= \{(0, i_1, i_2) : i_1 \in [r], i_2 \in [r+1, n]\}, \\ J_2 &:= \{(i_1, i_2, i_3) : i_1 \neq i_2, i_1, i_2 \in [r], i_3 \in [r+1, n]\}. \end{aligned} \quad (3.14)$$

To determine the scalars λ_i, γ_i , we can solve the linear least squares

$$\min_{(\beta_1, \dots, \beta_r)} \left\| \mathcal{F}_{J_1} - \sum_{i=1}^r \beta_i \cdot \tilde{v}_i \otimes \tilde{w}_i \right\|^2, \quad (3.15)$$

$$\min_{(\theta_1, \dots, \theta_r)} \left\| \mathcal{F}_{J_2} - \sum_{k=1}^r \theta_k \cdot (\tilde{v}_k \otimes \tilde{v}_k \otimes \tilde{w}_k)_{J_2} \right\|^2. \quad (3.16)$$

Let $(\beta_1^*, \dots, \beta_r^*), (\theta_1^*, \dots, \theta_r^*)$ be minimizers of (3.15) and (3.16) respectively. Then, for each $i = 1, \dots, r$, let

$$\lambda_i := (\beta_i^*)^2 / \theta_i^*, \quad \gamma_i := \theta_i^* / \beta_i^*. \quad (3.17)$$

For the vectors ($i = 1, \dots, r$)

$$p_i := \sqrt[3]{\lambda_i} (1, \gamma_i \tilde{v}_i, \tilde{w}_i),$$

the sum $p_1^{\otimes 3} + \dots + p_r^{\otimes 3}$ is a tensor decomposition for $\mathcal{F}$. This is justified in the following theorem.

Theorem 3.3 Suppose the tensor $\mathcal{F}$ has the decomposition as in (3.3). Assume that the vectors $v_1, \dots, v_r$ are linearly independent and the vectors $w_1, \dots, w_r$ are distinct from each other, where $v_1, \dots, v_r, w_1, \dots, w_r$ are defined as in (3.10). Let ξ be a generically chosen coefficient vector and let $p_1, \dots, p_r$ be the vectors produced as above. Then, the tensor decomposition $\mathcal{F} = p_1^{\otimes 3} + \dots + p_r^{\otimes 3}$ is unique.

Proof Since $v_1, \dots, v_r$ are linearly independent, the tensor decomposition (3.3) is unique, up to scalings and permutations. By Theorem 3.1, there is a unique generating matrix G for $\mathcal{F}$ satisfying (3.2). Under the given assumptions, (3.8) uniquely determines G . Note that $\xi^T w_1, \dots, \xi^T w_r$ are the eigenvalues of $N(\xi)$ and $v_1, \dots, v_r$ are the corresponding eigenvectors. When ξ is generically chosen, the values of $\xi^T w_1, \dots, \xi^T w_r$ are distinct eigenvalues of $N(\xi)$. So $N(\xi)$ has unique eigenvalue decompositions, and hence (3.13) must hold, up to a permutation of $(v_1, \dots, v_r)$. Since the coefficient matrices have full column ranks, the linear least squares problems have unique optimal solutions. Up to a permutation of $p_1, \dots, p_r$, it holds that $p_i = \sqrt[3]{\lambda_i} \begin{bmatrix} 1 \\ u_i \end{bmatrix}$. Then, the conclusion follows readily. $\square$

The following is the algorithm for computing an incomplete tensor decomposition for $\mathcal{F}$ when only its subtensor $\mathcal{F}_\Omega$ is given.

Algorithm 3.4 (Incomplete symmetric tensor decompositions)

Input: A third order symmetric subtensor $\mathcal{F}_\Omega$ and a rank $r = \text{rank}_S(\mathcal{F}) \leq \frac{d}{2} - 1$.

1. Determine the matrix G by solving (3.8) for each $\alpha = e_i + e_j \in \mathbb{B}_1$.
2. Let $N(\xi)$ be the matrix as in (3.11), for a randomly selected vector ξ . Compute the unit length eigenvectors $\tilde{v}_1, \dots, \tilde{v}_r$ of $N(\xi)$ and choose $\tilde{w}_i$ as in (3.12).
3. Solve the linear least squares (3.15) and (3.16) to get the coefficients λ_i, γ_i as in (3.17).
4. For each $i = 1, \dots, r$, let $p_i := \sqrt[3]{\lambda_i}(1, \gamma_i \tilde{v}_i, \tilde{w}_i)$.

Output: The tensor decomposition $\mathcal{F} = (p_1)^{\otimes 3} + \dots + (p_r)^{\otimes 3}$.

The following is an example of applying Algorithm 3.4.

Example 3.5 Consider the same tensor $\mathcal{F}$ as in Example 3.2. The monomial sets $\mathcal{B}_0, \mathcal{B}_1$ are the same. The matrices $A_{ij}[\mathcal{F}]$ and vectors $b_{ij}[\mathcal{F}]$ are

$$\begin{aligned} A_{13}[\mathcal{F}] &= A_{23}[\mathcal{F}] = \begin{bmatrix} -0.8 & 2.8 \\ -1.4 & 4 \end{bmatrix}, & b_{13}[\mathcal{F}] &= \begin{bmatrix} 1.6 \\ 2.2 \end{bmatrix}, & b_{23}[\mathcal{F}] &= \begin{bmatrix} -2 \\ -3.2 \end{bmatrix}, \\ A_{14}[\mathcal{F}] &= A_{24}[\mathcal{F}] = \begin{bmatrix} 1 & -0.8 \\ -1.4 & 4 \end{bmatrix}, & b_{14}[\mathcal{F}] &= \begin{bmatrix} 1.6 \\ -3.2 \end{bmatrix}, & b_{24}[\mathcal{F}] &= \begin{bmatrix} -2 \\ 7.6 \end{bmatrix}, \\ A_{15}[\mathcal{F}] &= A_{25}[\mathcal{F}] = \begin{bmatrix} 1 & -0.8 \\ -0.8 & 2.8 \end{bmatrix}, & b_{15}[\mathcal{F}] &= \begin{bmatrix} 2.2 \\ -3.2 \end{bmatrix}, & b_{25}[\mathcal{F}] &= \begin{bmatrix} -3.2 \\ 7.6 \end{bmatrix}. \end{aligned}$$

Solve (3.8) to obtain G , which is same as in (3.5). The matrices $N_3(G)$, $N_4(G)$, $N_5(G)$ are

$$N_3(G) = \begin{bmatrix} 1/3 & 2/3 \\ 4/3 & -1/3 \end{bmatrix}, \quad N_4(G) = \begin{bmatrix} 4/3 & -1/3 \\ -2/3 & 5/3 \end{bmatrix}, \quad N_5(G) = \begin{bmatrix} 5/3 & -2/3 \\ -4/3 & 7/3 \end{bmatrix}.$$

Choose a generic ξ , say, $\xi = (3, 4, 5)$, then

$$N(\xi) = \begin{bmatrix} 1/\sqrt{2} & -1/\sqrt{5} \\ 1/\sqrt{2} & 2/\sqrt{5} \end{bmatrix} \begin{bmatrix} 12 & 0 \\ 0 & 20 \end{bmatrix} \begin{bmatrix} 1/\sqrt{2} & -1/\sqrt{5} \\ 1/\sqrt{2} & 2/\sqrt{5} \end{bmatrix}^{-1}.$$

The unit length eigenvectors are

$$\tilde{v}_1 = (1/\sqrt{2}, 1/\sqrt{2}), \quad \tilde{v}_2 = (-1/\sqrt{5}, 2/\sqrt{5}).$$

As in (3.12), we get the vectors

$$w_1 = (1, 1, 1), \quad w_2 = (-1, 2, 3).$$

Solving (3.15) and (3.16), we get the scalars

$$\gamma_1 = \sqrt{2}, \quad \gamma_2 = \sqrt{5}, \quad \lambda_1 = 0.4, \quad \lambda_2 = 0.6.$$

This produces the decomposition $\mathcal{F} = \lambda_1 u_1^{\otimes 3} + \lambda_2 u_2^{\otimes 3}$ for the vectors

$$u_1 = (1, \gamma_1 v_1, w_1) = (1, 1, 1, 1, 1, 1), \quad u_2 = (1, \gamma_2 v_2, w_2) = (1, -1, 2, -1, 2, 3).$$

Remark 3.6 Algorithm 3.4 requires the value of r . This is generally a hard question. In computational practice, one can estimate the value of r as follows. Let $\text{Flat}(\mathcal{F}) \in \mathbb{C}^{(n+1) \times (n+1)^2}$ be the flattening matrix, labelled by $(i, (j, k))$ such that

$$\text{Flat}(\mathcal{F})_{i,(j,k)} = \mathcal{F}_{ijk}$$

for all $i, j, k = 0, 1, \dots, n$. The rank of $\text{Flat}(\mathcal{F})$ equals the rank of $\mathcal{F}$ when the vectors $p_1, \dots, p_r$ are linearly independent. The rank of $\text{Flat}(\mathcal{F})$ is not available since only the subtensor $(\mathcal{F})_\Omega$ is known. However, we can calculate the ranks of submatrices of $(\mathcal{F})_\Omega$ whose entries are known. If the tensor $\mathcal{F}$ as in (3.3) is such that both the sets $\{v_1, \dots, v_r\}$ and $\{w_1, \dots, w_r\}$ are linearly independent, one can see that $\sum_{i=1}^r \lambda_i v_i w_i^T$ is a known submatrix of $\text{Flat}(\mathcal{F})$ whose rank is r . This is generally the case if $r \leq \frac{d}{2} - 1$, since v_i has the length r and w_i has length $d - 1 - r \geq r$. Therefore, the known submatrices of $\text{Flat}(\mathcal{F})$ are generally sufficient to estimate $\text{rank}_S(\mathcal{F})$. For instance, we consider the case $\mathcal{F} \in S^3(\mathbb{C}^7)$. The flattening matrix $\text{Flat}(\mathcal{F})$ is

$$\begin{bmatrix} * & * & * & * & * & * & * \\ * & * & \mathcal{F}_{120} & \mathcal{F}_{130} & \mathcal{F}_{140} & \mathcal{F}_{150} & \mathcal{F}_{160} \\ * & \mathcal{F}_{210} & * & \mathcal{F}_{230} & \mathcal{F}_{240} & \mathcal{F}_{250} & \mathcal{F}_{260} \\ * & \mathcal{F}_{310} & \mathcal{F}_{320} & * & \mathcal{F}_{340} & \mathcal{F}_{350} & \mathcal{F}_{360} \\ * & \mathcal{F}_{410} & \mathcal{F}_{420} & \mathcal{F}_{430} & * & \mathcal{F}_{450} & \mathcal{F}_{460} \\ * & \mathcal{F}_{510} & \mathcal{F}_{520} & \mathcal{F}_{530} & \mathcal{F}_{540} & * & \mathcal{F}_{560} \\ * & \mathcal{F}_{610} & \mathcal{F}_{620} & \mathcal{F}_{630} & \mathcal{F}_{640} & \mathcal{F}_{650} & * \end{bmatrix},$$

where each $*$ means that entry is not given. The largest submatrices with known entries are

$$\begin{bmatrix} \mathcal{F}_{410} & \mathcal{F}_{420} & \mathcal{F}_{430} \\ \mathcal{F}_{510} & \mathcal{F}_{520} & \mathcal{F}_{530} \\ \mathcal{F}_{610} & \mathcal{F}_{620} & \mathcal{F}_{630} \end{bmatrix}, \quad \begin{bmatrix} \mathcal{F}_{140} & \mathcal{F}_{150} & \mathcal{F}_{160} \\ \mathcal{F}_{240} & \mathcal{F}_{250} & \mathcal{F}_{260} \\ \mathcal{F}_{340} & \mathcal{F}_{350} & \mathcal{F}_{360} \end{bmatrix}.$$

The rank of above matrices generally equals $\text{rank}_S(\mathcal{F})$ if $r \leq \frac{d}{2} - 1 = 2.5$.

4 Tensor Approximations and Stability Analysis

In some applications, we do not have the subtensor $\mathcal{F}_\Omega$ exactly but only have an approximation $\hat{\mathcal{F}}_\Omega$ for it. Algorithm 3.4 can still provide a good rank- r approximation for $\mathcal{F}$ when it is applied to $\hat{\mathcal{F}}_\Omega$. We define the matrix $A_{ij}[\hat{\mathcal{F}}]$ and the vector $b_{ij}[\hat{\mathcal{F}}]$ in the same way as

in (3.7), for each $\alpha = e_i + e_j \in \mathcal{B}_1$. The generating matrix G for $\mathcal{F}$ can be approximated by solving the linear least squares

$$\min_{g \in \mathbb{C}^r} \|A_{ij}[\widehat{\mathcal{F}}] \cdot g - b_{ij}[\widehat{\mathcal{F}}]\|^2, \quad (4.1)$$

for each $\alpha = e_i + e_j \in \mathbb{B}_1$. Let $\widehat{G}(\cdot, e_i + e_j)$ be the optimizer of the above and $\widehat{G}$ be the matrix consisting of all such $\widehat{G}(\cdot, e_i + e_j)$. Then $\widehat{G}$ is an approximation for G . For each $l = r + 1, \dots, n$, define the matrix $N_l(\widehat{G})$ similarly as in (3.9). Choose a generic vector $\xi = (\xi_{r+1}, \dots, \xi_n)$ and let

$$\widehat{N}(\xi) := \xi_{r+1} N_{r+1}(\widehat{G}) + \dots + \xi_n N_n(\widehat{G}). \quad (4.2)$$

The matrix $\widehat{N}(\xi)$ is an approximation for $N(\xi)$. Let $\hat{v}_1, \dots, \hat{v}_r$ be unit length eigenvectors of $\widehat{N}(\xi)$. For $k = 1, \dots, r$, let

$$\hat{w}_k := \left((\hat{v}_k)^H N_{r+1}(\widehat{G}) \hat{v}_k, \dots, (\hat{v}_k)^H N_n(\widehat{G}) \hat{v}_k \right). \quad (4.3)$$

For the label sets J_1, J_2 as in (3.14), the subtensors $\widehat{\mathcal{F}}_{J_1}, \widehat{\mathcal{F}}_{J_2}$ are similarly defined like $\mathcal{F}_{J_1}, \mathcal{F}_{J_2}$. Consider the following linear least square problems

$$\min_{(\beta_1, \dots, \beta_r)} \left\| \widehat{\mathcal{F}}_{J_1} - \sum_{k=1}^r \beta_k \cdot \hat{v}_k \otimes \hat{w}_k \right\|^2, \quad (4.4)$$

$$\min_{(\theta_1, \dots, \theta_r)} \left\| \widehat{\mathcal{F}}_{J_2} - \sum_{k=1}^r \theta_k \cdot (\hat{v}_k \otimes \hat{v}_k \otimes \hat{w}_k)_{J_2} \right\|^2. \quad (4.5)$$

Let $(\hat{\beta}_1, \dots, \hat{\beta}_r)$ and $(\hat{\theta}_1, \dots, \hat{\theta}_r)$ be their optimizers respectively. For each $k = 1, \dots, r$, let

$$\hat{\lambda}_k := (\hat{\beta}_k)^2 / \hat{\theta}_k, \quad \hat{\gamma}_k := \hat{\theta}_k / \hat{\beta}_k.$$

This results in the tensor approximation

$$\mathcal{F} \approx (\hat{p}_1)^{\otimes 3} + \dots + (\hat{p}_r)^{\otimes 3},$$

for the vectors $\hat{p}_k := \sqrt[3]{\hat{\lambda}_k} (1, \hat{\gamma}_k \hat{v}_k, \hat{w}_k)$. The above may not give an optimal tensor approximation. To get an improved one, we can use $\hat{p}_1, \dots, \hat{p}_r$ as starting points to solve the following nonlinear optimization

$$\min_{(q_1, \dots, q_r)} \left\| \left(\sum_{k=1}^r (q_k)^{\otimes 3} - \widehat{\mathcal{F}} \right)_{\Omega} \right\|^2. \quad (4.6)$$

The minimizer of the optimization (4.6) is denoted as $(p_1^*, \dots, p_r^*)$.

Summarizing the above, we have the following algorithm for computing a tensor approximation.

Algorithm 4.1 (Incomplete symmetric tensor approximations)

Input: A third order symmetric subtensor $\widehat{\mathcal{F}}_{\Omega}$ and a rank $r \leq \frac{d}{2} - 1$.

1. Find the matrix $\widehat{G}$ by solving (4.1) for each $\alpha = e_i + e_j \in \mathbb{B}_1$.
2. Choose a generic vector and let $\widehat{N}(\xi)$ be the matrix as in (4.2). Compute unit length eigenvectors $\hat{v}_1, \dots, \hat{v}_r$ for $\widehat{N}(\xi)$ and define $\hat{w}_i$ in (4.3).
3. Solve the linear least squares (4.4), (4.5) to get the coefficients $\hat{\lambda}_i, \hat{\gamma}_i$.

4. For each $i = 1, \dots, r$, let $\hat{p}_i := \sqrt[3]{\hat{\lambda}_i}(1, \hat{\gamma}_i \hat{v}_i, \hat{w}_i)$. Then $(\hat{p}_1)^{\otimes 3} + \dots + (\hat{p}_r)^{\otimes 3}$ is a tensor approximation for $\hat{\mathcal{F}}$.
5. Use $\hat{p}_1, \dots, \hat{p}_r$ as starting points to solve the nonlinear optimization (4.6) for an optimizer $(p_1^*, \dots, p_r^*)$.

Output: The tensor approximation $(p_1^*)^{\otimes 3} + \dots + (p_r^*)^{\otimes 3}$ for $\hat{\mathcal{F}}$.

When $\hat{\mathcal{F}}$ is close to $\mathcal{F}$, Algorithm 4.1 also produces a good rank- r tensor approximation for $\mathcal{F}$. This is shown in the following.

Theorem 4.2 Suppose the tensor $\mathcal{F} = (p_1)^{\otimes 3} + \dots + (p_r)^{\otimes 3}$, with $r \leq \frac{d}{2} - 1$, satisfies the following conditions:

- (i) The leading entry of each p_i is nonzero;
- (ii) the subvectors $(p_1)_{2:r+1}, \dots, (p_r)_{2:r+1}$ are linearly independent;
- (iii) the subvectors $(p_1)_{[r+2:j, j+2:d]}, \dots, (p_r)_{[r+2:j, j+2:d]}$ are linearly independent for each $j \in [r+1, n]$;
- (iv) the eigenvalues of the matrix $N(\xi)$ in (3.11) are distinct from each other.

Let $\hat{p}_i, p_i^*$ be the vectors produced by Algorithm 4.1. If the distance $\varepsilon := \|(\hat{\mathcal{F}} - \mathcal{F})_\Omega\|$ is small enough, then there exist scalars $\hat{\tau}_i, \tau_i^*$ such that

$$(\hat{\tau}_i)^3 = (\tau_i^*)^3 = 1, \quad \|\hat{\tau}_i \hat{p}_i - p_i\| = O(\varepsilon), \quad \|\tau_i^* p_i^* - p_i\| = O(\varepsilon),$$

up to a permutation of $(p_1, \dots, p_r)$, where the constants inside $O(\cdot)$ only depend on $\mathcal{F}$ and the choice of ξ in Algorithm 4.1.

Proof The conditions (i)–(ii), by Theorem 3.1, imply that there is a unique generating matrix G for $\mathcal{F}$. The matrix G can be approximated by solving the linear least square problems (4.1). Note that

$$\|A_{ij}[\hat{\mathcal{F}}] - A_{ij}[\mathcal{F}]\| \leq \varepsilon, \quad \|b_{ij}[\hat{\mathcal{F}}] - b_{ij}[\mathcal{F}]\| \leq \varepsilon,$$

for all $\alpha = e_i + e_j \in \mathcal{B}_1$. The matrix $A_{ij}[\mathcal{F}]$ can be written as

$$A_{ij}[\mathcal{F}] = [(p_1)_{[r+2:j, j+2:d]}, \dots, (p_r)_{[r+2:j, j+2:d]}] \cdot [(p_1)_{2:r+1}, \dots, (p_r)_{2:r+1}]^T.$$

By the conditions (ii)–(iii), the matrix $A_{ij}[\mathcal{F}]$ has full column rank for each $j \in [r+1, n]$ and hence the matrix $A_{ij}[\hat{\mathcal{F}}]$ has full column rank when ε is small enough. Therefore, the linear least problems (4.1) have unique solutions and the solution $\hat{G}$ satisfies that

$$\|\hat{G} - G\| = O(\varepsilon),$$

where $O(\varepsilon)$ depends on $\mathcal{F}$ (see [15, Theorem 3.4]). For each $j = r+1, \dots, n$, $N_j(\hat{G})$ is part of the generating matrix $\hat{G}$, so

$$\|N_j(\hat{G}) - N_j(G)\| \leq \|\hat{G} - G\| = O(\varepsilon), \quad j = r+1, \dots, n.$$

This implies that $\|\hat{N}(\xi) - N(\xi)\| = O(\varepsilon)$. When ε is small enough, the matrix $\hat{N}(\xi)$ does not have repeated eigenvalues, due to the condition (iv). Thus, the matrix $N(\xi)$ has a set of unit length eigenvectors $\tilde{v}_1, \dots, \tilde{v}_r$ with eigenvalues $\tilde{w}_1, \dots, \tilde{w}_r$ respectively, such that

$$\|\hat{v}_i - \tilde{v}_i\| = O(\varepsilon), \quad \|\hat{w}_i - \tilde{w}_i\| = O(\varepsilon).$$

This follows from Proposition 4.2.1 in [8]. The constants inside the above $O(\cdot)$ depend only on $\mathcal{F}$ and ξ . The $\tilde{w}_1, \dots, \tilde{w}_r$ are scalar multiples of linearly independent vectors $(p_1)_{r+2:d}, \dots, (p_r)_{r+2:d}$ respectively, so $\tilde{w}_1, \dots, \tilde{w}_r$ are linearly independent. When ε is

small, $\hat{w}_1, \dots, \hat{w}_r$ are linearly independent as well. The scalars $\hat{\lambda}_i \hat{\gamma}_i$ and $\hat{\lambda}_i (\hat{\gamma}_i)^2$ are optimizers for the linear least square problems (4.4) and (4.5). By Theorem 3.4 in [15], we have

$$\|\hat{\lambda}_i \hat{\gamma}_i - \lambda_i \gamma_i\| = O(\varepsilon), \quad \|\hat{\lambda}_i (\hat{\gamma}_i)^2 - \lambda_i \gamma_i^2\| = O(\varepsilon).$$

The vector p_i can be written as $p_i = \sqrt[3]{\hat{\lambda}_i} (1, \gamma_i \tilde{v}_i, \tilde{w}_i)$, so we must have $\lambda_i, \gamma_i \neq 0$ due to the condition (ii). Thus, it holds that

$$\|\hat{\lambda}_i - \lambda_i\| = O(\varepsilon), \quad \|\hat{\gamma}_i - \gamma_i\| = O(\varepsilon),$$

where constants inside $O(\cdot)$ depend only on $\mathcal{F}$ and ξ . For the vectors $\tilde{p}_i := \sqrt[3]{\hat{\lambda}_i} (1, \gamma_i \tilde{v}_i, \tilde{w}_i)$, we have $\mathcal{F} = \sum_{i=1}^r \tilde{p}_i^{\otimes 3}$, by Theorem 3.3. Since $p_1, \dots, p_r$ are linearly independent by the assumption, the rank decomposition of $\mathcal{F}$ is unique up to scaling and permutation. There exist scalars $\hat{\tau}_i$ such that $(\hat{\tau}_i)^3 = 1$ and $\hat{\tau}_i \tilde{p}_i = p_i$, up to a permutation of $p_1, \dots, p_r$. For $\hat{p}_i = \sqrt[3]{\hat{\lambda}_i} (1, \hat{\gamma}_i \hat{v}_i, \hat{w}_i)$, we have $\|\hat{\tau}_i \hat{p}_i - p_i\| = O(\varepsilon)$, where the constants in $O(\cdot)$ only depend on $\mathcal{F}$ and ξ .

Since $\|\hat{\tau}_i \hat{p}_i - p_i\| = O(\varepsilon)$, we have $\|(\sum_{i=1}^r (\hat{p}_i)^{\otimes 3} - \mathcal{F})_{\Omega}\| = O(\varepsilon)$. The $(p_1^*, \dots, p_r^*)$ is a minimizer of (4.6), so

$$\left\| \left(\sum_{i=1}^r (p_i^*)^{\otimes 3} - \hat{\mathcal{F}} \right)_{\Omega} \right\| \leq \left\| \left(\sum_{i=1}^r (\hat{p}_i)^{\otimes 3} - \hat{\mathcal{F}} \right)_{\Omega} \right\| = O(\varepsilon).$$

For the tensor $\mathcal{F}^* := \sum_{i=1}^r (p_i^*)^{\otimes 3}$, we get

$$\|(\mathcal{F}^* - \mathcal{F})_{\Omega}\| \leq \|(\mathcal{F}^* - \hat{\mathcal{F}})_{\Omega}\| + \|(\hat{\mathcal{F}} - \mathcal{F})_{\Omega}\| = O(\varepsilon).$$

When Algorithm 4.1 is applied to $(\mathcal{F}^*)_{\Omega}$, Step 4 will give the exact decomposition $\mathcal{F}^* = \sum_{i=1}^r (p_i^*)^{\otimes 3}$. By repeating the previous argument, we can similarly show that $\|p_i - \tau_i^* p_i^*\| = O(\varepsilon)$ for some τ_i^* such that $(\tau_i^*)^3 = 1$, where the constants in $O(\cdot)$ only depend on $\mathcal{F}$ and ξ . $\square$

Remark 4.3 For the special case that $\varepsilon = 0$, Algorithm 4.1 is the same as Algorithm 3.4, which produces the exact rank decomposition for $\mathcal{F}$. The conditions in Theorem 4.2 are satisfied for generic vectors $p_1, \dots, p_r$, since $r \leq \frac{d}{2} - 1$. The constant in $O(\cdot)$ is not explicitly given in the proof. It is related to the condition number $\kappa(\mathcal{F})$ for tensor decomposition [5]. It was shown by Breiding and Vannieuwenhoven [5] that

$$\sqrt{\sum_{i=1}^r \|p_i^{\otimes 3} - \hat{p}_i^{\otimes 3}\|^2} \leq \kappa(\mathcal{F}) \|\mathcal{F} - \hat{\mathcal{F}}\| + c\varepsilon^2$$

for some constant c . The continuity of $\hat{G}$ in $\hat{\mathcal{F}}$ is implicitly implied by the proof. Eigenvalues and unit eigenvectors of $\hat{N}(\xi)$ are continuous in $\hat{G}$. Furthermore, $\hat{\lambda}_i, \hat{\gamma}_i$ are continuous in the eigenvalues and unit eigenvectors. All these functions are locally Lipschitz continuous. The $\hat{p}_i$ is Lipschitz continuous with respect to $\hat{\mathcal{F}}$, in a neighborhood of $\mathcal{F}$, which also implies an error bound for $\hat{p}_i$. The tensors $(p_i^*)^{\otimes 3}$ are also locally Lipschitz continuous in $\hat{\mathcal{F}}$, as illustrated in [6]. This also gives error bounds for decomposing vectors p_i^* . We refer to [5, 6] for more details about condition numbers of tensor decompositions.

Example 4.4 We consider the same tensor $\mathcal{F}$ as in Example 3.2. The subtensor $(\mathcal{F})_{\Omega}$ is perturbed to $(\hat{\mathcal{F}})_{\Omega}$. The perturbation is randomly generated from the Gaussian distribution $\mathcal{N}(0, 0.01)$. For neatness of the paper, we do not display $(\hat{\mathcal{F}})_{\Omega}$ here. We use Algorithm 4.1

to compute the incomplete tensor approximation. The matrices $A_{ij}[\widehat{\mathcal{F}}]$ and vectors $b_{ij}[\widehat{\mathcal{F}}]$ are given as follows:

$$\begin{aligned} A_{13}[\widehat{\mathcal{F}}] &= A_{23}[\widehat{\mathcal{F}}] = \begin{bmatrix} -0.8135 & 2.7988 \\ -1.3697 & 4.0149 \end{bmatrix}, & b_{13}[\widehat{\mathcal{F}}] &= \begin{bmatrix} 1.5980 \\ 2.1879 \end{bmatrix}, & b_{23}[\widehat{\mathcal{F}}] &= \begin{bmatrix} -2.0047 \\ -3.2027 \end{bmatrix}, \\ A_{14}[\widehat{\mathcal{F}}] &= A_{24}[\widehat{\mathcal{F}}] = \begin{bmatrix} 1.0277 & -0.8020 \\ -1.3697 & 4.0149 \end{bmatrix}, & b_{14}[\widehat{\mathcal{F}}] &= \begin{bmatrix} 1.5920 \\ -3.2013 \end{bmatrix}, & b_{24}[\widehat{\mathcal{F}}] &= \begin{bmatrix} -2.0059 \\ 7.5915 \end{bmatrix}, \\ A_{15}[\widehat{\mathcal{F}}] &= A_{25}[\widehat{\mathcal{F}}] = \begin{bmatrix} 1.0277 & -0.8020 \\ -0.8135 & 2.7988 \end{bmatrix}, & b_{15}[\widehat{\mathcal{F}}] &= \begin{bmatrix} 2.1993 \\ -3.2020 \end{bmatrix}, & b_{25}[\widehat{\mathcal{F}}] &= \begin{bmatrix} -3.1917 \\ 7.6153 \end{bmatrix}. \end{aligned}$$

The linear least square problems (4.1) are solved to obtain $\widehat{G}$ and $N_3(\widehat{G})$, $N_4(\widehat{G})$, $N_5(\widehat{G})$, which are

$$\begin{aligned} N_3(\widehat{G}) &= \begin{bmatrix} 0.5156 & 0.7208 \\ 1.6132 & -0.2474 \end{bmatrix}, & N_4(\widehat{G}) &= \begin{bmatrix} 1.2631 & -0.3665 \\ -0.6489 & 1.6695 \end{bmatrix}, \\ N_5(\widehat{G}) &= \begin{bmatrix} 1.6131 & -0.6752 \\ -1.2704 & 2.3517 \end{bmatrix}. \end{aligned}$$

For $\xi = (3, 4, 5)$, the eigendecomposition of the matrix $\widehat{N}(\xi)$ in (4.2) is

$$\widehat{N}(\xi) = \begin{bmatrix} -0.7078 & 0.4470 \\ -0.7064 & -0.8945 \end{bmatrix} \begin{bmatrix} 12.0343 & 0 \\ 0 & 20.0786 \end{bmatrix} \begin{bmatrix} -0.7524 & 0.4499 \\ -0.6588 & -0.8931 \end{bmatrix}^{-1}.$$

It has eigenvectors $\widehat{v}_1 = (-0.7078, -0.7064)$, $\widehat{v}_2 = (0.4470, -0.8945)$. The vectors $\widehat{w}_1$, $\widehat{w}_2$ obtained as in (4.3) are

$$\widehat{w}_1 = (1.2021, 0.9918, 0.9899), \quad \widehat{w}_2 = (-1.0389, 2.0145, 3.0016).$$

By solving (4.4) and (4.5), we got the scalars

$$\widehat{\gamma}_1 = -1.1990, \quad \widehat{\gamma}_2 = -2.1458, \quad \widehat{\lambda}_1 = 0.4521, \quad \widehat{\lambda}_2 = 0.6232.$$

Finally, we got the decomposition $\widehat{\lambda}_1 \widehat{u}_1^{\otimes 3} + \widehat{\lambda}_2 \widehat{u}_2^{\otimes 3}$ with

$$\begin{aligned} \widehat{u}_1 &= (1, \widehat{\gamma}_1 \widehat{v}_1, \widehat{w}_1) = (1, 0.8477, 0.8479, 1.2021, 0.9918, 0.9899), \\ \widehat{u}_2 &= (1, \widehat{\gamma}_2 \widehat{v}_2, \widehat{w}_2) = (1, -0.9776, 1.9102, -1.0389, 2.0145, 3.0016). \end{aligned}$$

They are pretty close to the decomposition of $\mathcal{F}$.

5 Learning Diagonal Gaussian Mixture

We use the incomplete tensor decomposition or approximation method to learn parameters for Gaussian mixture models. Algorithms 3.4 and 4.1 can be applied to do that.

Let y be the random variable of dimension d for a Gaussian mixture model, with r components of Gaussian distribution parameters $(\omega_i, \mu_i, \Sigma_i)$, $i = 1, \dots, r$. We consider the case that $r \leq \frac{d}{2} - 1$. Let $y_1, \dots, y_N$ be samples drawn from the Gaussian mixture model. The sample average

$$\widehat{M}_1 := \frac{1}{N}(y_1 + \dots + y_N)$$

is an estimation for the mean $M_1 := \mathbb{E}[y] = \omega_1 \mu_1 + \dots + \omega_r \mu_r$. The symmetric tensor

$$\widehat{M}_3 := \frac{1}{N}(y_1^{\otimes 3} + \dots + y_N^{\otimes 3})$$

is an estimation for the third order moment tensor $M_3 := \mathbb{E}[y^{\otimes 3}]$. Recall that $\mathcal{F} = \sum_{i=1}^r \omega_i \mu_i^{\otimes 3}$. When all the covariance matrices Σ_i are diagonal, we have shown in (1.2) that

$$M_3 = \mathcal{F} + \sum_{j=1}^d (a_j \otimes e_j \otimes e_j + e_j \otimes a_j \otimes e_j + e_j \otimes e_j \otimes a_j).$$

If the labels i_1, i_2, i_3 are distinct from each other, $(M_3)_{i_1 i_2 i_3} = (\mathcal{F})_{i_1 i_2 i_3}$. Recall the label set Ω in (1.4). It holds that

$$(M_3)_\Omega = (\mathcal{F})_\Omega.$$

Note that $(\widehat{M}_3)_\Omega$ is only an approximation for $(M_3)_\Omega$ and $(\mathcal{F})_\Omega$, due to sampling errors. If the rank $r \leq \frac{d}{2} - 1$, we can apply Algorithm 4.1 with the input $(\widehat{M}_3)_\Omega$, to compute a rank- r tensor approximation for $\mathcal{F}$. Suppose the tensor approximation produced by Algorithm 4.1 is

$$\mathcal{F} \approx (p_1^*)^{\otimes 3} + \cdots + (p_r^*)^{\otimes 3}.$$

The computed $p_1^*, \dots, p_r^*$ may not be real vectors, even if $\mathcal{F}$ is real. When the error $\varepsilon := \|(\mathcal{F} - \widehat{M}_3)_\Omega\|$ is small, by Theorem 4.2, we know

$$\|\tau_i^* p_i^* - \sqrt[3]{\omega_i} \mu_i\| = O(\varepsilon)$$

where $(\tau_i^*)^3 = 1$. In computation, we can choose τ_i^* such that $(\tau_i^*)^3 = 1$ and the imaginary part vector $\text{Im}(\tau_i^* p_i^*)$ has the smallest norm. It can be done by checking the imaginary part of $\tau_i^* p_i^*$ one by one for

$$\tau_i^* = 1, \quad -\frac{1}{2} + \frac{\sqrt{-3}}{2}, \quad -\frac{1}{2} - \frac{\sqrt{-3}}{2}.$$

Then we get the real vector

$$\hat{q}_i := \text{Re}(\tau_i^* p_i^*).$$

It is expected that $\hat{q}_i \approx \sqrt[3]{\omega_i} \mu_i$. Since

$$M_1 = \omega_1 \mu_1 + \cdots + \omega_r \mu_r \approx \omega_1^{2/3} \hat{q}_1 + \cdots + \omega_r^{2/3} \hat{q}_r,$$

the scalars $\omega_1^{2/3}, \dots, \omega_r^{2/3}$ can be obtained by solving the nonnegative linear least squares

$$\min_{(\beta_1, \dots, \beta_r) \in \mathbb{R}_+^r} \left\| \widehat{M}_1 - \sum_{i=1}^r \beta_i \hat{q}_i \right\|^2. \quad (5.1)$$

Let $(\beta_1^*, \dots, \beta_r^*)$ be an optimizer for the above, then $\hat{\omega}_i := (\beta_i^*)^{3/2}$ is a good approximation for ω_i and the vector

$$\hat{\mu}_i := \hat{q}_i / \sqrt[3]{\hat{\omega}_i}$$

is a good approximation for μ_i . We may use

$$\hat{\mu}_i, \quad \left(\sum_{j=1}^r \hat{\omega}_j \right)^{-1} \hat{\omega}_i, \quad i = 1, \dots, r$$

as starting points to solve the nonlinear optimization

$$\begin{cases} \min_{(\omega_1, \dots, \omega_r, \mu_1, \dots, \mu_r)} & \|\sum_{i=1}^r \omega_i \mu_i - \widehat{M}_1\|^2 + \|\sum_{i=1}^r \omega_i (\mu_i^{\otimes 3})_\Omega - (\widehat{M}_3)_\Omega\|^2 \\ \text{subject to} & \omega_1 + \cdots + \omega_r = 1, \omega_1, \dots, \omega_r \geq 0, \end{cases} \quad (5.2)$$

for getting improved approximations. Suppose an optimizer of the above is

$$(\omega_1^*, \dots, \omega_r^*, \mu_1^*, \dots, \mu_r^*).$$

Now, we discuss how to estimate the diagonal covariance matrices Σ_i . Let

$$\mathcal{A} := M_3 - \mathcal{F}, \quad \widehat{\mathcal{A}} := \widehat{M}_3 - (\widehat{q}_1)^{\otimes 3} - \cdots - (\widehat{q}_r)^{\otimes 3}. \quad (5.3)$$

By (1.2), we know that

$$\mathcal{A} = \sum_{j=1}^d (a_j \otimes e_j \otimes e_j + e_j \otimes a_j \otimes e_j + e_j \otimes e_j \otimes a_j), \quad (5.4)$$

where $a_j = \sum_{i=1}^r \omega_i \sigma_{ij}^2 \mu_i$ for $j = 1, \dots, d$. The equation (5.4) implies that

$$(a_j)_j = \frac{1}{3} \mathcal{A}_{jjj}, \quad (a_j)_i = \mathcal{A}_{jij},$$

for $i, j = 1, \dots, d$ and $i \neq j$. So we choose vectors $\hat{a}_j \in \mathbb{R}^d$ such that

$$(\hat{a}_j)_j = \frac{1}{3} \widehat{\mathcal{A}}_{jjj}, \quad (\hat{a}_j)_i = \widehat{\mathcal{A}}_{jij} \quad \text{for } i \neq j. \quad (5.5)$$

Since $\hat{a}_j \approx \sum_{i=1}^r \omega_i \sigma_{ij}^2 \mu_i$, the covariance matrices $\Sigma_i = \text{diag}(\sigma_{i1}^2, \dots, \sigma_{id}^2)$ can be estimated by solving the nonnegative linear least squares ($j = 1, \dots, d$)

$$\begin{cases} \min_{(\beta_{1j}, \dots, \beta_{rj})} \|\hat{a}_j - \sum_{i=1}^r \omega_i^* \mu_i^* \beta_{ij}\|^2 \\ \text{subject to } \beta_{1j} \geq 0, \dots, \beta_{rj} \geq 0. \end{cases} \quad (5.6)$$

For each j , let $(\beta_{1j}^*, \dots, \beta_{rj}^*)$ be the optimizer for the above. When $(\widehat{M}_3)_\Omega$ is close to $(M_3)_\Omega$, it is expected that β_{ij}^* is close to $(\sigma_{ij})^2$. Therefore, we can estimate the covariance matrices Σ_i as follows

$$\Sigma_i^* := \text{diag}(\beta_{i1}^*, \dots, \beta_{id}^*), \quad (\sigma_{ij}^*)^2 := \beta_{ij}^*. \quad (5.7)$$

The following is the algorithm for learning Gaussian mixture models.

Algorithm 5.1 (Learning diagonal Gaussian mixture models)

Input: Samples $\{y_1, \dots, y_N\} \subseteq \mathbb{R}^d$ drawn from a Gaussian mixture model and the number r of component Gaussian distributions.

- Step 1. Compute the sample averages $\widehat{M}_1 := \frac{1}{N} \sum_{i=1}^N y_i$ and $\widehat{M}_3 := \frac{1}{N} \sum_{i=1}^N y_i^{\otimes 3}$.
- Step 2. Apply Algorithm 4.1 to the subtensor $(\widehat{\mathcal{F}})_\Omega := (\widehat{M}_3)_\Omega$. Let $(p_1^*)^{\otimes 3} + \cdots + (p_r^*)^{\otimes 3}$ be the obtained rank- r tensor approximation for $\widehat{\mathcal{F}}$. For each $i = 1, \dots, r$, let $\hat{q}_i := \text{Re}(\tau_i p_i^*)$ where τ_i is the cube root of 1 that minimizes the imaginary part vector norm $\|\text{Im}(\tau_i p_i^*)\|$.
- Step 3. Solve (5.1) to get $\hat{\omega}_1, \dots, \hat{\omega}_r$ and $\hat{\mu}_i = q_i / \sqrt[3]{\hat{\omega}_i}, i = 1, \dots, r$.
- Step 4. Use the above $\hat{\omega}_i, \hat{q}_i$ as initial points to solve the nonlinear optimization (5.2) for the optimal $\omega_i^*, \mu_i^*, i = 1, \dots, r$.
- Step 5. Get vectors $\hat{a}_1, \dots, \hat{a}_d$ as in (5.5). Solve the optimization (5.6) to get optimizers β_{ij}^* and then choose Σ_i^* as in (5.7).

Output: Component Gaussian distribution parameters $(\mu_i^*, \Sigma_i^*, \omega_i^*), i = 1, \dots, r$.

The sample averages $\widehat{M}_1, \widehat{M}_3$ can typically be used as good estimates for the true moments M_1, M_3 . When the value of r is not known, it can be determined as in Remark 3.6. The performance of Algorithm 5.1 is analyzed as follows.

Theorem 5.2 Consider the d -dimensional diagonal Gaussian mixture model with parameters $\{(\omega_i, \mu_i, \Sigma_i) : i \in [r]\}$ and $r \leq \frac{d}{2} - 1$. Let $\{(\omega_i^*, \mu_i^*, \Sigma_i^*) : i \in [r]\}$ be produced by Algorithm 5.1. If the distance $\varepsilon := \max(\|M_3 - \widehat{M}_3\|, \|M_1 - \widehat{M}_1\|)$ is small enough and the tensor $\mathcal{F} = \sum_{i=1}^r \omega_i \mu_i^{\otimes 3}$ satisfies conditions of Theorem 4.2, then

$$\|\mu_i - \mu_i^*\| = O(\varepsilon), \quad \|\omega_i - \omega_i^*\| = O(\varepsilon), \quad \|\Sigma_i - \Sigma_i^*\| = O(\varepsilon),$$

where the above constants inside $O(\cdot)$ only depend on parameters $\{(\omega_i, \mu_i, \Sigma_i) : i \in [r]\}$ and the choice of ξ in Algorithm 5.1.

Proof For the vectors $p_i := \sqrt[3]{\omega_i} \mu_i$, we have $\mathcal{F} = \sum_{i=1}^r p_i^{\otimes 3}$. Since

$$\|(\mathcal{F} - \widehat{\mathcal{F}})_{\Omega}\| = \|(M_3 - \widehat{M}_3)_{\Omega}\| \leq \varepsilon$$

and $\mathcal{F}$ satisfies conditions of Theorem 4.2, we know $\|\tau_i^* p_i^* - p_i\| = O(\varepsilon)$ for some $(\tau_i^*)^3 = 1$, by Theorem 4.2. The constants inside $O(\varepsilon)$ depend on parameters of the Gaussian model and ξ . Then, we have $\|\text{Im}(\tau_i^* p_i^*)\| = O(\varepsilon)$ since the vectors p_i are real. When ε is small enough, such τ_i^* is the τ in Step 2 of Algorithm 5.1 that minimizes $\|\text{Im}(\tau_i p_i^*)\|$, so we have

$$\|\hat{q}_i - p_i\| \leq \|\tau_i p_i^* - p_i\| = O(\varepsilon),$$

where $\hat{q}_i = \text{Re}(\tau_i p_i^*)$ is from Step 2. The vectors $\hat{q}_1, \dots, \hat{q}_r$ are linearly independent when ε is small. Thus, the problem (5.1) has a unique solution and the weights $\hat{\omega}_i$ can be found by solving (5.1). Since $\|M_1 - \widehat{M}_1\| \leq \varepsilon$ and $\|\hat{q}_i - p_i\| = O(\varepsilon)$, we have $\|\omega_i - \hat{\omega}_i\| = O(\varepsilon)$ (see [15, Theorem 3.4]). The mean vectors $\hat{\mu}_i$ are obtained by $\hat{\mu}_i = \hat{q}_i / \sqrt[3]{\hat{\omega}_i}$, so the approximation error is

$$\|\mu_i - \hat{\mu}_i\| = \|p_i / \sqrt[3]{\omega_i} - \hat{q}_i / \sqrt[3]{\hat{\omega}_i}\| = O(\varepsilon).$$

The constants inside the above $O(\varepsilon)$ depend on parameters of the Gaussian mixture model and ξ .

The problem (5.2) is solved to obtain ω_i^* and μ_i^* , so

$$\left\| \widehat{M}_1 - \sum_{i=1}^r \omega_i^* \mu_i^* \right\| + \left\| \widehat{\mathcal{F}} - \sum_{i=1}^r \omega_i^* (\mu_i^*)^{\otimes 3} \right\| = O(\varepsilon).$$

Let $\mathcal{F}^* := \sum_{i=1}^r \omega_i^* (\mu_i^*)^{\otimes 3} = \sum_{i=1}^r (\sqrt[3]{\omega_i^*} \mu_i^*)^{\otimes 3}$, then

$$\|\mathcal{F} - \mathcal{F}^*\| \leq \|\mathcal{F} - \widehat{\mathcal{F}}\| + \|\widehat{\mathcal{F}} - \mathcal{F}^*\| = O(\varepsilon).$$

Theorem 4.2 implies $\|p_i - \sqrt[3]{\omega_i^*} \mu_i^*\| = O(\varepsilon)$. In addition, we have

$$\left\| \widehat{M}_1 - \sum_{i=1}^r \omega_i^* \mu_i^* \right\| = \left\| \widehat{M}_1 - \sum_{i=1}^r (\omega_i^*)^{2/3} \sqrt[3]{\omega_i^*} \mu_i^* \right\| = O(\varepsilon).$$

The first order moment is $M_1 = \sum_{i=1}^r (\omega_i)^{2/3} p_i$. Since $\|M_1 - \widehat{M}_1\| = O(\varepsilon)$ and $\|p_i - \sqrt[3]{\omega_i^*} \mu_i^*\| = O(\varepsilon)$, it holds that $\|\omega_i^{2/3} - (\omega_i^*)^{2/3}\| = O(\varepsilon)$ by [15, Theorem 3.4]. This implies that $\|\omega_i - \omega_i^*\| = O(\varepsilon)$, so

$$\|\mu_i - \mu_i^*\| = \|p_i / \sqrt[3]{\omega_i} - (\sqrt[3]{\omega_i^*} \mu_i^*) / \sqrt[3]{\omega_i^*}\| = O(\varepsilon).$$

The constants inside the above $O(\cdot)$ only depend on parameters $\{(\omega_i, \mu_i, \Sigma_i) : i \in [r]\}$ and ξ .

The covariance matrices Σ_i are recovered by solving the linear least squares (5.6). In the least square problems, it holds that $\|\omega_i \mu_i - \omega_i^* \mu_i^*\| = O(\varepsilon)$ and

$$\|\mathcal{A} - \widehat{\mathcal{A}}\| \leq \|M_3 - \widehat{M}_3\| + \|\mathcal{F} - \sum_{i=1}^r \widehat{q}_i^{\otimes 3}\| = O(\varepsilon),$$

where tensors $\mathcal{A}$, $\widehat{\mathcal{A}}$ are defined in (5.3). When the error ε is small, the vectors $\omega_i^* \mu_1^*, \dots, \omega_i^* \mu_r^*$ are linearly independent and hence (5.6) has a unique solution for each j . By [15, Theorem 3.4], we have

$$\|(\sigma_{ij})^2 - (\sigma_{ij}^*)^2\| = O(\varepsilon).$$

It implies that $\|\Sigma_i - \Sigma_i^*\| = O(\varepsilon)$, where the constants inside $O(\cdot)$ only depend on parameters $\{(\omega_i, \mu_i, \Sigma_i) : i \in [r]\}$ and ξ . $\square$

6 Numerical Simulations

This section gives numerical experiments for our proposed methods. The computation is implemented in MATLAB R2019b, on an Alienware personal computer with Intel(R)Core(TM)i7-9700K CPU@3.60GHz and RAM 16.0G. The MATLAB function `lsqnonlin` is used to solve (4.6) in Algorithm 4.1 and the MATLAB function `fmincon` is used to solve (5.2) in Algorithm 5.1. We compare our method with the classical EM algorithm, which is implemented by the MATLAB function `fitgmdist` (MaxIter is set to be 100 and RegularizationValue is set to be 0.001).

First, we show the performance of Algorithm 4.1 for computing incomplete symmetric tensor approximations. For a range of dimension d and rank r , we get the tensor $\mathcal{F} = (p_1)^{\otimes 3} + \dots + (p_r)^{\otimes 3}$, where each p_i is randomly generated according to the Gaussian distribution in MATLAB. Then, we apply the perturbation $(\widehat{\mathcal{F}})_{\Omega} = (\mathcal{F})_{\Omega} + \mathcal{E}_{\Omega}$, where $\mathcal{E}$ is a randomly generated tensor, also according to the Gaussian distribution in MATLAB, with the norm $\|\mathcal{E}_{\omega}\|_{\Omega} = \varepsilon$. After that, Algorithm 4.1 is applied to the subtensor $(\widehat{\mathcal{F}})_{\Omega}$ to find the rank- r tensor approximation. The approximation quality is measured by the absolute error and the relative error

$$\text{abs-error} := \|(\mathcal{F}^* - \mathcal{F})_{\Omega}\|, \quad \text{rel-error} := \frac{\|(\mathcal{F}^* - \widehat{\mathcal{F}})_{\Omega}\|}{\|(\mathcal{F} - \widehat{\mathcal{F}})_{\Omega}\|},$$

where $\mathcal{F}^*$ is the output of Algorithm 4.1. For each case of (d, r, ε) , we generate 100 random instances. The min, average, and max relative errors for each dimension d and rank r are reported in the Table 1. The results show that Algorithm 4.1 performs very well for computing tensor approximations.

Second, we explore the performance of Algorithm 5.1 for learning diagonal Gaussian mixture models. We compare it with the classical EM algorithm, for which the MATLAB function `fitgmdist` is used (MaxIter is set to be 100 and RegularizationValue is set to be 0.0001). The dimensions $d = 20, 30, 40, 50, 60$ are tested. Three values of r are tested for each case of d . We generate 100 random instances of $\{(\omega_i, \mu_i, \Sigma_i) : i = 1, \dots, r\}$ for $d \in \{20, 30, 40\}$, and 20 random instances for $d \in \{50, 60\}$, because of the relatively more computational time for the latter case. For each instance, 10000 samples are generated. To generate the weights $\omega_1, \dots, \omega_r$, we first use the MATLAB function `randi` to generate a random 10000-dimensional integer vector of entries from $[r]$, then the occurring frequency of i in $[r]$ is used as the weight ω_i . For each diagonal covariance matrix Σ_i ,

Table 1 The performance of Algorithm 4.1

<i>d</i>	<i>r</i>	ε	rel-error			abs-error			time
			min	average	max	min	average	max	
20	3	0.1	0.9610	0.9731	0.9835	0.0141	0.0268	0.0556	0.2687
	5	0.01	0.9634	0.9700	0.9742	0.0019	0.0032	0.0068	0.2392
	7	0.001	0.9148	0.9373	0.9525	$2.3 \cdot 10^{-4}$	$3.8 \cdot 10^{-4}$	$6.6 \cdot 10^{-4}$	0.2638
30	4	0.1	0.9816	0.9854	0.9890	0.0094	0.0174	0.0533	0.4386
	8	0.01	0.9634	0.9700	0.9742	0.0015	0.0024	0.0060	0.7957
	11	0.001	0.9501	0.9587	0.9667	$1.8 \cdot 10^{-4}$	$3.0 \cdot 10^{-4}$	$5.7 \cdot 10^{-4}$	0.8954
40	6	0.1	0.9853	0.9877	0.9904	0.0099	0.0146	0.0359	1.7779
	10	0.01	0.9761	0.9795	0.9820	0.0013	0.0020	0.0045	2.6454
	15	0.001	0.9653	0.9690	0.9734	$1.7 \cdot 10^{-4}$	$2.6 \cdot 10^{-4}$	$4.8 \cdot 10^{-4}$	3.6785
50	7	0.1	0.9887	0.9911	0.9925	0.0081	0.0128	0.0294	4.9774
	13	0.01	0.9812	0.9831	0.9854	0.0011	0.0018	0.0045	8.7655
	18	0.001	0.9739	0.9767	0.9792	$1.5 \cdot 10^{-4}$	$2.2 \cdot 10^{-4}$	$4.1 \cdot 10^{-4}$	11.6248

its diagonal vector is set to be the square of a random vector generated by the MATLAB function randn. Each sample is generated from one of r component Gaussian distributions, so they are naturally separated into r groups. Algorithm 5.1 and the EM algorithm

Table 2 Comparison between Algorithm 5.1 and EM for simulations

<i>d</i>	<i>r</i>	Accuracy		Time	
		Algorithm 5.1	EM	Algorithm 5.1	EM
20	3	0.9861	0.9763	0.8745	0.1649
	5	0.9740	0.9400	2.3476	0.3852
	7	0.9659	0.9252	3.4352	0.6777
30	4	0.9965	0.9684	4.5266	0.2959
	8	0.9923	0.9277	8.5494	0.8525
	11	0.9895	0.9219	17.2091	1.4106
40	6	0.9990	0.9117	18.9160	0.6273
	10	0.9981	0.8931	28.4161	1.2617
	15	0.9971	0.9111	69.8013	2.0627
50	7	0.9997	0.8997	40.6810	0.8314
	13	0.9995	0.9073	104.7927	1.7867
	18	0.9993	0.9038	163.2711	2.6862
60	8	0.9999	0.8874	93.9836	1.1266
	15	0.9998	0.8632	234.0331	2.6435
	22	0.9995	0.8929	497.9371	3.5527

Table 3 Classification results on 8 textures

Accuracy	Algorithm 5.1	EM
Bark.0000	0.5376	0.8413
Bark.0009	0.5107	0.7150
Flowers.0001	0.8137	0.6315
Tile.0000	0.8219	0.7239
Paintings.11.0001	0.8047	0.7350
Grass.0001	0.9841	0.9068
Brick.0004	0.9406	0.8854
Fabric.0013	0.9220	0.9048

are applied to fit the Gaussian mixture model to the 10000 samples for each instance. For each sample, we calculate the likelihood of the sample to each component Gaussian distribution in the estimated Gaussian mixture model. A sample is classified to the i th group if its likelihood for the i th component is maximum. The classification accuracy is the rate that samples are classified to the correct group. In Table 2, for each pair (d, r) , we report the accuracy of Algorithm 5.1 in the first row and the accuracy of the EM algorithm in the second row. As one can see, Algorithm 5.1 performs better than EM algorithm, and its accuracy isn't affected when the dimensions and ranks increase. Indeed, as the difference between the dimension d and the rank r increases, Algorithm 5.1 becomes more and more accurate. This is opposite to the EM algorithm. The reason is that the difference between the number of rows and the number of columns of $A_{ij}[\mathcal{F}]$ in (3.7) increases as $d - r$ becomes bigger, which makes Algorithm 5.1 more robust.

Last, we apply Algorithm 5.1 to do texture classifications. We select 8 textured images of 512×512 pixels from the VisTex database (Fig. 1). We use the MATLAB function `rgb2gray` to convert them into grayscale version since we only need their structure and

Fig. 1 Textures from VisTex

texture information. Each image is divided into subimages of 32×32 pixels. We perform the discrete cosine transformation (DCT) on each block of size 16×16 pixels with overlap of 8 pixels. Each component of ‘Wavelet-like’ DCT feature is the sum of the absolute value of the DCT coefficients in the corresponding sub-block. So the dimension d of the feature vector extracted from each subimage is 13. We use blocks extracted from the first 160 subimages for training and those from the rest 96 subimages for testing. We refer to [47] for more details. For each image, we apply Algorithm 5.1 and the EM algorithm to fit a Gaussian mixture model to the image. We choose the number of components r according to Remark 3.6. To classify the test data, we follow the Bayesian decision rule that assigns each block to the texture which maximizes the posteriori probability, where we assume a uniform prior over all classes [18]. The classification accuracy is the rate that a subimage is correctly classified, which is shown in Table 3. Algorithm 5.1 outperforms the classical EM algorithm for the accuracy rates for six of the images.

7 Conclusions and Future Work

This paper gives a new algorithm for learning Gaussian mixture models with diagonal covariance matrices. We first give a method for computing incomplete symmetric tensor decompositions. It is based on generating polynomials. The method is described in Algorithm 3.4. When the input subtensor has small errors, we can similarly compute the incomplete symmetric tensor approximation, which is given by Algorithm 4.1. We have shown in Theorem 4.2 that if the input subtensor is sufficiently close to a low rank one, the produced tensor approximation is highly accurate. Then unknown parameters for Gaussian mixture models can be recovered by using the incomplete tensor decomposition method. It is described in Algorithm 5.1. When the estimations of M_1 and M_3 are accurate, the parameters recovered by Algorithm 5.1 are also accurate. The computational simulations demonstrate the good performance of the proposed method.

The proposed method deals with the case that the number of Gaussian components is less than one half of the dimension. How do we compute incomplete symmetric tensor decompositions when the set Ω is not like (1.4)? How can we learn parameters for Gaussian mixture models with more components? How can we do that when the covariance matrices are not diagonal? They are important and interesting topics for future research work.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Achlioptas, D., McSherry, F.: On spectral learning of mixtures of distributions. In: Auer, P., Meir, R. (eds.) *Learning Theory. International Conference on Computational Learning Theory. Lecture Notes in Computer Science*, vol. 3559, pp. 458–469. Springer, Berlin, Heidelberg (2005)
2. Anandkumar, A., Ge, R., Hsu, D., Kakade, S., Telgarsky, M.: Tensor decompositions for learning latent variable models. *J. Mach. Learn. Res.* **15**, 2773–2832 (2014)

3. Belkin, M., Sinha, K.: Toward learning Gaussian mixtures with arbitrary separation. In: The 23rd Annual Conference on Learning Theory (COLT 2010), Haifa, June 27–29 (2010)
4. Bioucas-Dias, J.M., Plaza, A., Dobigeon, N., Parente, M., Du, Q., Gader, P., Chanussot, J.: Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **5**, 354–379 (2012)
5. Breiding, P., Vannieuwenhoven, N.: The condition number of join decompositions. *SIAM J. Matrix Anal. Appl.* **39**, 287–309 (2018)
6. Breiding, P., Vannieuwenhoven, N.: The condition number of Riemannian approximation problems. *SIAM J. Optim.* **31**, 1049–1077 (2021)
7. Brubaker, S.C., Vempala, S.S.: Isotropic PCA and affine-invariant clustering. In: Grötschel, M., Katona, G.O.H., Sági, G. (eds.) *Building Bridges*. Bolyai Society Mathematical Studies, vol. 19, pp. 241–281. Springer, Berlin, Heidelberg (2008)
8. Chatelin, F.: *Eigenvalues of Matrices: Revised Edition*. SIAM, Philadelphia (2012)
9. Chaudhuri, K., Kakade, S., Livescu, K., Sridharan, K.: Multi-view clustering via canonical correlation analysis. In: *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 129–136 (2009)
10. Chaudhuri, K., Rao, S.: Learning mixtures of product distributions using correlations and independence. In: *The 21st Annual Conference on Learning Theory (COLT 2008)*, 9–12 July, 2008, Helsinki, pp. 9–20 (2008)
11. Chiantini, L., Ottaviani, G., Vannieuwenhoven, N.: On generic identifiability of symmetric tensors of subgeneric rank. *Trans. Amer. Math. Soc.* **369**, 4021–4042 (2017)
12. Comon, P., Lim, L.-H., Qi, Y., Ye, K.: Topology of tensor ranks. *Adv. Math.* **367**, 107128 (2020)
13. Dasgupta, S.: Learning mixtures of Gaussians. In: *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pp. 634–644. IEEE (1999)
14. Dasgupta, S., Schulman, L.: A two-round variant of EM for Gaussian mixtures. [arXiv:1301.3850](https://arxiv.org/abs/1301.3850) (2013)
15. Demmel, J.: *Applied Numerical Linear Algebra*. SIAM, Philadelphia (1997)
16. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Ser. B (Methodol.)* **39**, 1–22 (1977)
17. de Silva, V., Lim, L.-H.: Tensor rank and the ill-posedness of the best low-rank approximation problem. *SIAM J. Matrix Anal. Appl.* **30**, 1084–1127 (2008)
18. Dixit, M., Rasiwasia, N., Vasconcelos, N.: Adapted Gaussian models for image classification. *CVPR 2011*, pp 937–943. IEEE (2011)
19. Dressler, M., Nie, J., Yang, Z.: Separability of Hermitian tensors and PSD decompositions. [arXiv:2011.08132](https://arxiv.org/abs/2011.08132) (2020)
20. Fialkow, L., Nie, J.: The truncated moment problem via homogenization and flat extensions. *J. Funct. Anal.* **263**, 1682–1700 (2012)
21. Friedland, S.: Remarks on the symmetric rank of symmetric tensors. *SIAM J. Matrix Anal. Appl.* **37**, 320–337 (2016)
22. Friedland, S., Lim, L.-H.: Nuclear norm of higher-order tensors. *Math. Comput.* **87**, 1255–1281 (2018)
23. Ge, R., Huang, Q., Kakade, S.M.: Learning mixtures of Gaussians in high dimensions. In: *Proceedings of the forty-seventh annual ACM symposium on Theory of Computing*, pp. 761–770 (2015)
24. Ge, F., Ju, Y., Qi, Z., Lin, Y.: Parameter estimation of a Gaussian mixture model for wind power forecast error by Riemann L-BFGS optimization. *IEEE Access* **6**, 38892–38899 (2018)
25. Haas, M., Mittnik, S., Paolella, M.S.: Modelling and predicting market risk with Laplace–Gaussian mixture distributions. *Appl. Financ. Econ.* **16**, 1145–1162 (2006)
26. Hillar, C.J., Lim, L.-H.: Most tensor problems are NP-hard. *J. ACM* **60**, 45 (2013)
27. Hsu, D., Kakade, S.M.: Learning mixtures of spherical Gaussians: moment methods and spectral decompositions. In: *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pp. 11–20 (2013)
28. Kalai, A.T., Moitra, A., Valiant, G.: Efficiently learning mixtures of two Gaussians. In: *Proceedings of the forty-second ACM symposium on Theory of computing*, pp. 553–562 (2010)
29. Kannan, R., Salmasian, H., Vempala, S.: The spectral method for general mixture models. In: Auer, P., Meir, R. (eds.) *Learning Theory. International Conference on Computational Learning Theory. Lecture Notes in Computer Science*, vol. 3559, pp. 444–457. Springer, Berlin, Heidelberg (2005)
30. Karpagavalli, S., Chandra, E.: A review on automatic speech recognition architecture and approaches. *Int. J. Signal Process. Image Process. Pattern Recognit.* **9**, 393–404 (2016)

31. Landsberg, J.M.: *Tensors: Geometry and Applications*. Graduate Studies in Mathematics, vol. 128. American Mathematical Society, Providence, RI (2012)
32. Lee, D.-S.: Effective Gaussian mixture learning for video background subtraction. *IEEE Trans. Pattern Anal. Mach. Intell.* **27**, 827–832 (2005)
33. Lim, L.-H.: Tensors and hypermatrices. In: Hogben, L. (ed.) *Handbook of Linear Algebra*, 2nd edn., pp. 15-1–15-30. CRC Press, Boca Raton, FL (2013)
34. Ma, Y., Jin, Q., Mei, X., Dai, X., Fan, F., Li, H., Huang, J.: Hyperspectral unmixing with Gaussian mixture model and low-rank representation. *Remote Sens.* **11**, 911 (2019)
35. Magdon-Ismaïl, M., Purnell, J.T.: Approximating the covariance matrix of GMMs with low-rank perturbations. In: Eyfe, C., et al. (eds.) *Intelligent Data Engineering and Automated Learning – IDEAL 2010*. International Conference on Intelligent Data Engineering and Automated Learning. Lecture Notes in Computer Science, vol. 6283, pp. 300–307. Springer, Berlin, Heidelberg (2010)
36. Mu, C., Huang, B., Wright, J., Goldfarb, D.: Square deal: lower bounds and improved relaxations for tensor recovery. In: *Proceeding of the 31st International Conference on Machine Learning (PMLR)*, vol. 32, pp. 73–81 (2014)
37. Nie, J., Sturmfels, B.: Matrix cubes parameterized by eigenvalues. *SIAM J. matrix Anal. Appl.* **31**, 755–766 (2009)
38. Nie, J.: The hierarchy of local minimums in polynomial optimization. *Math. Program.* **151**, 555–583 (2015)
39. Nie, J.: Linear optimization with cones of moments and nonnegative polynomials. *Math. Program.* **153**, 247–274 (2015)
40. Nie, J.: Generating polynomials and symmetric tensor decompositions. *Found. Comput. Math.* **17**, 423–465 (2017)
41. Nie, J.: Symmetric tensor nuclear norms. *SIAM J. Appl. Algebra Geom.* **1**, 599–625 (2017)
42. Nie, J.: Low rank symmetric tensor approximations. *SIAM J. Matrix Anal. Appl.* **38**, 1517–1540 (2017)
43. Nie, J.: Tight relaxations for polynomial optimization and Lagrange multiplier expressions. *Math. Program.* **178**, 1–37 (2019)
44. Nie, J., Ye, K.: Hankel tensor decompositions and ranks. *SIAM J. Matrix Anal. Appl.* **40**, 486–516 (2019)
45. Nie, J., Yang, Z.: Hermitian tensor decompositions. *SIAM J. Matrix Anal. Appl.* **41**, 1115–1144 (2020)
46. Pearson, K.: Contributions to the mathematical theory of evolution. *Philos. Trans. R. Soc. Lond. A* **185**, 71–110 (1894)
47. Permuter, H., Francos, J., Jermyn, I.: A study of Gaussian mixture models of color and texture features for image classification and segmentation. *Pattern Recognit.* **39**, 695–706 (2006)
48. Povey, D., Burget, L., Agarwal, M., Akyazi, P., Kai, F., Ghoshal, A., Glembek, O., Goel, N., Karafiát, M., Rastrow, A., Rose, R.C., Schwarz, P., Thomas, S.: The subspace Gaussian mixture model—A structured model for speech recognition. *Comput. Speech Lang.* **25**, 404–439 (2011)
49. Redner, R.A., Walker, H.F.: Mixture densities, maximum likelihood and the EM algorithm. *SIAM Rev.* **26**, 195–239 (1984)
50. Reynolds, D.A.: Speaker identification and verification using Gaussian mixture speaker models. *Speech Commun.* **17**, 91–108 (1995)
51. Romera-Paredes, B., Pontil, M.: A new convex relaxation for tensor completion. In: *Advances in Neural Information Processing Systems 26*, pp. 2967–2975 (2013)
52. Sanjeev, A., Kannan, R.: Learning mixtures of arbitrary Gaussians. In: *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pp. 247–257 (2001)
53. Shekoffeh, Y., Jafari, S., Sprott, J.C., Golpayegani, S.M.R.H., Almasganj, F.: A Gaussian mixture model based cost function for parameter estimation of chaotic biological systems. *Commun. Nonlinear Sci. Numer. Simul.* **20**, 469–481 (2015)
54. Tang, G., Shah, P.: Guaranteed tensor decomposition: a moment approach. In: *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pp. 1491–1500 (2015)
55. Vempala, S., Wang, G.: A spectral algorithm for learning mixture models. *J. Comput. Syst. Sci.* **68**, 841–860 (2004)
56. Veracini, T., Matteoli, S., Diani, M., Corsini, G.: Fully unsupervised learning of Gaussian mixtures for anomaly detection in hyperspectral imagery. In: *2009 Ninth International Conference on Intelligent Systems Design and Applications*, pp. 596–601. IEEE (2009)
57. Wu, Y., Yang, P.: Optimal estimation of Gaussian mixtures via denoised method of moments. *Ann. Statist.* **48**, 1981–2007 (2020)
58. Yuan, M., Zhang, C.-H.: On tensor completion via nuclear norm minimization. *Found. Comput. Math.* **16**, 1031–1068 (2016)

59. Zhang, H., Giles, C.L., Foley, H.C., Yen, J.: Probabilistic community discovery using hierarchical latent Gaussian mixture model. In: AAAI-07/IAAI-07 Proceedings: 22nd AAAI Conference on Artificial Intelligence and the 19th Innovative Applications of Artificial Intelligence Conference, vol. 1, pp. 663–668 (2007)
60. Zivkovic, Z.: Improved adaptive Gaussian mixture model for background subtraction. In: Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004, pp. 28–31. IEEE (2004)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.