

Foundations and Trends<sup>®</sup> in Machine Learning  
**Spectral Learning on Matrices and  
Tensors**

---

**Suggested Citation:** Majid Janzamin, Rong Ge, Jean Kossaifi and Anima Anandkumar (2019), Spectral Learning on Matrices and Tensors, Foundations and Trends<sup>®</sup> in Machine Learning: Vol. 12, No. 5-6, pp 393–536. DOI: 10.1561/22000000057.

**Majid Janzamin**

Twitter

majid.janzamin@gmail.com

**Rong Ge**

Duke University

rongge@cs.duke.edu

**Jean Kossaifi**

Imperial College London

jean.kossaifi@imperial.ac.uk

**Anima Anandkumar**

NVIDIA & California Institute of Technology

anima@caltech.edu

This article may be used only for the purpose of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval.

**now**

the essence of knowledge

Boston — Delft

# Contents

---

<b>1</b>	<b>Introduction</b>	<b>395</b>
1.1	Method of Moments and Moment Tensors . . . . .	397
1.2	Warm-up: Learning a Simple Model with Tensors . . . . .	398
1.3	What's Next? . . . . .	401
<b>2</b>	<b>Matrix Decomposition</b>	<b>402</b>
2.1	Low Rank Matrix Decomposition . . . . .	402
2.2	Low Rank Matrix Approximation and SVD . . . . .	406
2.3	Principal Component Analysis . . . . .	411
2.4	Whitening Transformation . . . . .	414
2.5	Canonical Correlation Analysis . . . . .	416
<b>3</b>	<b>Tensor Decomposition Algorithms</b>	<b>420</b>
3.1	Transition from Matrices to Tensors . . . . .	420
3.2	Tensor Preliminaries and Notations . . . . .	424
3.3	Uniqueness of CP decomposition . . . . .	429
3.4	Orthogonal Tensor Decomposition . . . . .	430
3.5	Tensor Power Iteration . . . . .	444
3.6	Simultaneous Diagonalization . . . . .	456
3.7	Alternating Least Squares . . . . .	459

<b>4</b>	<b>Applications of Tensor Methods</b>	<b>464</b>
4.1	Pure Topic Model Revisited . . . . .	465
4.2	Beyond Raw Moments . . . . .	467
4.3	Multi-view Models . . . . .	472
4.4	Nonlinear Model: Noisy-Or Networks . . . . .	478
4.5	Applications in Supervised Learning . . . . .	480
4.6	Other Models . . . . .	486
<b>5</b>	<b>Practical Implementations</b>	<b>488</b>
5.1	Programming Language and Framework . . . . .	488
5.2	Tensors as NumPy Arrays . . . . .	489
5.3	Basic Tensor Operations and Decomposition . . . . .	492
5.4	Example: Image Compression via Tensor Decomposition . . . . .	498
5.5	Going Further with TensorLy . . . . .	500
5.6	Scaling up with PyTorch . . . . .	501
<b>6</b>	<b>Efficiency of Tensor Decomposition</b>	<b>505</b>
6.1	Running Time and Memory Usage . . . . .	505
6.2	Sample Complexity . . . . .	509
<b>7</b>	<b>Overcomplete Tensor Decomposition</b>	<b>514</b>
7.1	Higher-order Tensors via Tensorization . . . . .	514
7.2	FOOBI Algorithm . . . . .	516
7.3	Third Order Tensors . . . . .	521
7.4	Open Problems . . . . .	524
	<b>Acknowledgements</b>	<b>526</b>
	<b>References</b>	<b>527</b>

# Spectral Learning on Matrices and Tensors

Majid Janzamin<sup>1</sup>, Rong Ge<sup>2</sup>, Jean Kossaifi<sup>3</sup> and Anima Anandkumar<sup>4</sup>

<sup>1</sup>*Twitter; majid.janzamin@gmail.com*

<sup>2</sup>*Duke University; rongge@cs.duke.edu*

<sup>3</sup>*Imperial College London; jean.kossaifi@imperial.ac.uk*

<sup>4</sup>*NVIDIA & California Institute of Technology; anima@caltech.edu*

---

## ABSTRACT

Spectral methods have been the mainstay in several domains such as machine learning, applied mathematics and scientific computing. They involve finding a certain kind of spectral decomposition to obtain basis functions that can capture important structures or directions for the problem at hand. The most common spectral method is the principal component analysis (PCA). It utilizes the principal components or the top eigenvectors of the data covariance matrix to carry out dimensionality reduction as one of its applications. This data pre-processing step is often effective in separating signal from noise.

PCA and other spectral techniques applied to *matrices* have several limitations. By limiting to only pairwise moments, they are effectively making a Gaussian approximation on the underlying data. Hence, they fail on data with hidden variables which lead to non-Gaussianity. However, in almost any data set, there are latent effects that cannot be directly observed, e.g., topics in a document corpus, or underlying causes of a disease. By extending the spectral decomposition

methods to higher order moments, we demonstrate the ability to learn a wide range of latent variable models efficiently. Higher-order moments can be represented by *tensors*, and intuitively, they can encode more information than just pairwise moment matrices. More crucially, tensor decomposition can pick up latent effects that are missed by matrix methods. For instance, tensor decomposition can uniquely identify non-orthogonal components. Exploiting these aspects turns out to be fruitful for provable unsupervised learning of a wide range of latent variable models.

We also outline the computational techniques to design efficient tensor decomposition methods. They are embarrassingly parallel and thus scalable to large data sets. Whilst there exist many optimized linear algebra software packages, efficient tensor algebra packages are also beginning to be developed. We introduce Tensorly, which has a simple python interface for expressing tensor operations. It has a flexible back-end system supporting NumPy, PyTorch, TensorFlow and MXNet amongst others. This allows it to carry out multi-GPU and CPU operations, and can also be seamlessly integrated with deep-learning functionalities.

---

# 1

---

## Introduction

---

Probabilistic models form an important area of machine learning. They attempt to model the probability distribution of the observed data, such as documents, speech and images. Often, this entails relating observed data to *latent or hidden variables*, e.g., topics for documents, words for speech and objects for images. The goal of learning is to then discover the latent variables and their relationships to the observed data.

Latent variable models have shown to be useful to provide a good explanation of the observed data, where they can capture the effect of hidden causes which are not directly observed. Learning these hidden factors is central to many applications, e.g., identifying latent diseases through observed symptoms, and identifying latent communities through observed social ties. Furthermore, latent representations are very useful in feature learning. Raw data is in general very complex and redundant and feature learning is about extracting informative features from raw data. Learning efficient and useful features is crucial for the performance of learning tasks, e.g., the classification task that we perform using the learned features.

Learning latent variable models is challenging since the latent variables cannot, by definition, be directly observed. In extreme cases, when

there are more latent variables than observations, learning is theoretically impossible because of the lack of data, unless further constraints are imposed. More generally, learning latent variable models raises several questions. How much data do we need to observe in order to uniquely determine the model's parameters? Are there efficient algorithms to effectively learn these parameters? Can we get provable guarantees on the running time of the algorithm and the number of samples required to estimate the parameters? These are all important questions about learning latent variable models that we will try to address here.

In this monograph, we survey recent progress in using spectral methods including matrix and tensor decomposition techniques to learn many popular latent variable models. With careful implementation, tensor-based methods can run efficiently in practice, and in many cases they are the only algorithms with provable guarantees on running time and sample complexity.

There exist other surveys and overviews on tensor decomposition and its applications in machine learning and beyond. Among them, the work by Kolda and Bader (2009) is very well-received in the community where they provide a comprehensive introduction to major tensor decomposition forms and algorithms and discuss some of their applications in science and engineering. More recently, Sidiropoulos *et al.* (2017) provide an overview of different types of tensor decompositions and some of their applications in signal processing and machine learning. Papalexakis *et al.* (2017) discuss several applications of tensor decompositions in data mining. Rabanser *et al.* (2017) review some basic concepts of tensor decompositions and a few applications. Debals and De Lathauwer (2017) review several tensorization techniques which had been proposed in the literature. Here, tensorization is the mapping of a vector or matrix to a tensor to enable us using tensor tools.

In contrast to the above works, our focus in this monograph is on a special type of tensor decomposition called CP decomposition (see (1.3) as an example), and we cover a wide range of algorithms to find the components of such tensor decomposition. We also discuss the usefulness of this decomposition by reviewing several probabilistic models that can be learned using such tensor methods.

## 1.1 Method of Moments and Moment Tensors

How can we learn latent variable models, even though we cannot observe the latent variables? The key lies in understanding the relationship between latent variables and observed variables. A common framework for such relationship is known as the **method of moments** which dates back to Pearson (1894).

**Pearson's 1-d Example:** The main idea of method of moments is to first estimate *moments* of the data, and use these estimates to learn the unknown parameters of the probabilistic model. For a one-dimensional random variable  $X \in \mathbb{R}$ , the  $r$ -th order moment is denoted by  $\mathbb{E}[X^r]$ , where  $r$  is a positive integer and  $\mathbb{E}[\cdot]$  is the expectation operator. Consider a simple example where  $X$  is a mixture of two Gaussian variables. More precisely, with probability  $p_1$ ,  $X$  is drawn from a Gaussian distribution with mean  $\mu_1$  and variance  $\sigma_1^2$ , and with probability  $p_2$ ,  $X$  is drawn from a Gaussian distribution with mean  $\mu_2$  and variance  $\sigma_2^2$ . Here we have  $p_1 + p_2 = 1$ . Let us consider the problem of estimating these unknown parameters given samples of  $X$ . The random variable  $X$  can be viewed as drawn from a latent variable model because given a sample of  $X$ , we do not know which Gaussian it came from. Let latent variable  $Z \in \{1, 2\}$  be a random variable with probability  $p_1$  of being 1. Then given  $Z$ ,  $X$  is just a Gaussian distribution as

$$[X|Z = z] \sim \mathcal{N}(\mu_z, \sigma_z^2).$$

As noted by Pearson (1894), even though we cannot observe  $Z$ , the moments of  $X$  are closely related to the unknown parameters (probabilities  $p_1, p_2$ , means  $\mu_1, \mu_2$ , standard deviations  $\sigma_1, \sigma_2$ ) we desire to estimate. More precisely, for the first three moments we have

$$\begin{aligned}\mathbb{E}[X] &= p_1\mu_1 + p_2\mu_2, \\ \mathbb{E}[X^2] &= p_1(\mu_1^2 + \sigma_1^2) + p_2(\mu_2^2 + \sigma_2^2), \\ \mathbb{E}[X^3] &= p_1(\mu_1^3 + 3\mu_1\sigma_1^2) + p_2(\mu_2^3 + 3\mu_2\sigma_2^2).\end{aligned}$$



The moments  $\mathbb{E}[X], \mathbb{E}[X^2], \mathbb{E}[X^3], \dots$  can be empirically estimated given observed data. Therefore, the equations above can be interpreted as a system of equations on the six unknown parameters stated above. Pearson (1894) showed that with the first 6-th moments, we have enough equations to *uniquely* determine the values of the parameters.

### Moments for Multivariate Random Variables of Higher Dimensions:


For a scalar random variable, its  $p$ -th moment is just a scalar number. However, for a random vector, higher order moments can reveal much more information. Let us consider a random vector  $X \in \mathbb{R}^d$ . The first moment of this variable is a vector  $\mu \in \mathbb{R}^d$  such that  $\mu_i = \mathbb{E}[X_i], \forall i \in [d]$ , where  $[d] := \{1, 2, \dots, d\}$ . For the second order moment, we are not only interested in the second moments of individual coordinates  $\mathbb{E}[X_i^2]$ , but also in the *correlation* between different coordinates  $\mathbb{E}[X_i X_j], i \neq j$ . Therefore, it is convenient to represent the second order moment as a  $d \times d$  symmetric matrix  $M$ , where  $M_{i,j} = \mathbb{E}[X_i X_j]$ .

This becomes more complicated when we look at higher order moments. For 3rd order moment, we are interested in the correlation between all *triplets* of variables. In order to represent this compactly, we use a 3-dimensional  $d \times d \times d$  object  $T$ , also known as a 3rd order tensor. The tensor is constructed such that  $T_{i,j,k} = \mathbb{E}[X_i X_j X_k], \forall i, j, k \in [d]$ . This tensor has  $d^3$  elements or  $\binom{d+2}{3}$  distinct entries. In general,  $p$ -th order moment can be represented as a  $p$ -th order tensor with  $d^p$  entries. These tensors are called moment tensors. Vectors and matrices are special cases of moment tensors of order 1 and 2, respectively.

In applications, it is often crucial to define what the random variable  $X$  is, and examine what moments of  $X$  we can estimate from the data. We now provide a simple example to elaborate on how to form a useful moment and defer the proposal of many more examples to Section 4.

## 1.2 Warm-up: Learning a Simple Model with Tensors

In this section, we will give a simple example to demonstrate what is a tensor decomposition, and how it can be applied to learning latent variable models. Similar ideas can be applied to more complicated models, which we will discuss in Section 4.



**Figure 1.1:** Pure Topic Model

**Pure Topic Model:** The model we consider is a very simple topic model (Papadimitriou *et al.*, 2000; Hofmann, 1999). In this model, there are  $k$  unknown topics. Each topic entails a probability distribution over words in the vocabulary. Intuitively, the probabilities represent the likelihood of using a particular word when talking about a specific topic. As an example, the word “snow” should have a high probability in the topic “weather” but not the topic “politics”. These probabilities are represented as a matrix  $A \in \mathbb{R}^{d \times k}$ , where  $d$  is the size of the vocabulary and every column represents a topic. So, the columns of matrix  $A$  correspond to the probabilities over vocabulary that each topic entails. We will use  $\mu_j \in \mathbb{R}^d, j \in [k]$  to denote these probability distribution of words given  $j$ -th topic ( $j$ -th column of matrix  $A$ ).

The model assumes each document is generated in the following way: first a topic  $h \in [k]$  is chosen with probability  $w_h$  where  $w \in \mathbb{R}^k$  is a vector of probabilities; next,  $l$  words  $x_1, x_2, \dots, x_l$  are independently sampled from the  $h$ -th topic-word probability vector  $\mu_h$ . Therefore, we finally observe words for the documents. See Figure 1.1 for a graphical illustration of this model. This is clearly a latent variable model, since we don’t observe the topics. Our goal is to learn the parameters, which include the topic probability vector  $w$  and the topic-word probability vectors  $\mu_1, \dots, \mu_k$ .

**Computing the Moments:** First, we need to identify what the interesting moments are in this case. Since all we can observe are words in documents, and documents are all generated independently at random, it is natural to consider correlations between words as moments.

We say  $x \in \mathbb{R}^d$  is an indicator vector of a word  $z$  in our size- $d$  vocabulary if the  $z$ -th coordinate of  $x$  is 1 and all other coordinates of

$x$  are 0. For each document, let  $x_1, x_2, x_3 \in \mathbb{R}^d$  be indicator vectors for the first three words. Given these word representations, the entries of the first three moments of  $x_1, x_2, x_3$  can be written as

$$\begin{aligned} M_1(i) &= \Pr[x_1 = e_i], \\ M_2(i_1, i_2) &= \Pr[x_1 = e_{i_1}, x_2 = e_{i_2}], \\ M_3(i_1, i_2, i_3) &= \Pr[x_1 = e_{i_1}, x_2 = e_{i_2}, x_3 = e_{i_3}], \end{aligned}$$

where  $e_i \in \mathbb{R}^d$  denotes the  $i$ -th basis vector in  $d$ -dimensional space. Intuitively, the first moment  $M_1$  represents the probabilities for words; the second moment  $M_2$  represents the probabilities that two words co-occur; and the third moment  $M_3$  represents the probabilities that three words co-occur.

We can empirically estimate  $M_1, M_2, M_3$  from the observed documents. Now in order to apply the method of moments, we need to represent these probabilities based on the unknown parameters of our model. We can show that

$$M_1 = \sum_{h=1}^k w_h \mu_h, \tag{1.1}$$

$$M_2 = \sum_{h=1}^k w_h \mu_h \mu_h^\top, \tag{1.2}$$

$$M_3 = \sum_{h=1}^k w_h \mu_h \otimes \mu_h \otimes \mu_h. \tag{1.3}$$

The computation follows from the law of total expectations (explained in more details in Section 4). Here, the first moment  $M_1$  is the weighted average of  $\mu_h$ ; the second moment  $M_2$  is the weighted average of outer-products  $\mu_h \mu_h^\top$ ; and the third moment  $M_3$  is the weighted average of *tensor-products*  $\mu_h \otimes \mu_h \otimes \mu_h$ . The tensor product  $\mu_h \otimes \mu_h \otimes \mu_h$  is a  $d \times d \times d$  array whose  $(i_1, i_2, i_3)$ -th entry is equal to  $\mu_h(i_1)\mu_h(i_2)\mu_h(i_3)$ . See Section 3 for more precise definition of the tensor product operator  $\otimes$ .

Note that the second moment  $M_2$  is a matrix of rank at most  $k$ , and Equation (1.2) provides a low-rank matrix decomposition of  $M_2$ .

Similarly, finding  $w_h$  and  $\mu_h$  from  $M_3$  using Equation (1.3) is a problem called *tensor decomposition*. Clearly, if we can solve this problem, and it gives a unique solution, then we have learned the parameters of the model and we are done.

### 1.3 What's Next?

In the rest of this monograph, we will discuss the properties of tensor decomposition problem, review algorithms to efficiently find the components of such decomposition, and explain how they can be applied to learn the parameters of various probabilistic models such as latent variable models.

In Section 2, we first give a brief review of some basic matrix decomposition problems, including the singular value decomposition (SVD) and canonical correlation analysis (CCA). In particular, we will emphasize why matrix decomposition is often not enough to learn all the parameters of the latent variable models.

Section 3 discusses several algorithms for tensor decomposition. We will highlight under what conditions the tensor decomposition is *unique*, which is crucial in identifying the parameters of latent variable models.

In Section 4, we give more examples on how to apply tensor decomposition to learn different latent variable models. In different situations, there are many tricks to manipulate the moments in order to get a clean equation that looks similar to (1.3).

In Section 5, we illustrate how to implement tensor operations in practice using the Python programming language. We then show how to efficiently perform tensor learning using TensorLy and scale things up using PyTorch.

Tensor decomposition and its applications in learning latent variable models are still active research directions. In the last two sections of this monograph we discuss some of the more recent results, which deals with the problem of overcomplete tensors and improves the guarantees on running time and sample complexity.

# 2

---

## Matrix Decomposition


---

In this chapter, we describe some basic applications of matrix decomposition techniques including singular value decomposition (SVD), Principle Component Analysis (PCA) and canonical correlation analysis (CCA). These techniques are widely used in data analysis, and have been covered in many previous books (see e.g., Golub and Van Loan, 1990; Horn and Johnson, 2012; Blum *et al.*, 2016).

The goal of this chapter is to give a brief overview of the matrix decomposition techniques. At the same time we try to point out connections and differences with relevant concepts in tensor decomposition. Especially, in many cases these matrix-based methods have the problem of ambiguity, and cannot be directly applied to learning parameters for latent variable models. In the next section, we will describe how these limitations can be solved by using tensor decomposition instead of matrix decomposition.

### 2.1 Low Rank Matrix Decomposition

Assuming the reader is familiar with the basics of matrix algebra, we will start with reviewing matrix decompositions and matrix rank. *Rank* is a basic property of matrices. A rank-1 matrix can be expressed as



**Figure 2.1:** Decomposition of a matrix  $M \in \mathbb{R}^{5 \times 4}$  as sum of the rank-1 components. Note that each component is the product of a column vector  $u_j$  and a row vector  $v_j^\top$ .

the outer product of two vectors as  $uv^\top$  – its  $(i, j)$ -th entry is equal to the product of the  $i$ -th entry of vector  $u$  denoted by  $u(i)$  and the  $j$ -th entry of vector  $v$  denoted by  $v(j)$ . Similarly, a matrix  $M \in \mathbb{R}^{n \times m}$  is of rank at most  $k$  if it can be written as the sum of  $k$  rank-1 matrices as

$$M = \sum_{j=1}^k u_j v_j^\top. \quad (2.1)$$

Here  $u_1, u_2, \dots, u_k \in \mathbb{R}^n$  and  $v_1, v_2, \dots, v_k \in \mathbb{R}^m$  form the rank-1 components of the matrix  $M$ . We call Equation (2.1) a *decomposition* of matrix  $M$  into rank-1 components; see Figure 2.1 for a graphical representation of this decomposition for a sample matrix  $M \in \mathbb{R}^{5 \times 4}$ .

In many practical applications, the entries of a matrix are often determined by a small number of *factors*, and each factor corresponds to a rank-1 matrix; see Equation (2.1). As a result, many matrices we observe are *close to low rank matrices* (for exact definition of closeness see Section 2.2).

We elaborate the application of low rank matrix decompositions with the following example. Psychologist Charles Spearman worked on understanding whether human intelligence is a composite of different types of measurable intelligence and analyzed that through a factor analysis (Spearman, 1904). Let us describe a highly simplified version of his method, where the hypothesis is that there are exactly two kinds of intelligence: *quantitative* and *verbal*. Suppose  $n$  students are taking  $m$  different tests on distinct subjects. We can summarize the scores that students get in different tests in a matrix  $M \in \mathbb{R}^{n \times m}$ . Each row lists the scores for a student, and each column the scores for a particular subject; see the score matrix example in Figure 2.2.

	Math.	Classics	Physics	Music		Verbal		Quantitative		
Alice	$\begin{pmatrix} 19 & 26 & 17 & 21 \\ 8 & 17 & 9 & 12 \\ 7 & 12 & 7 & 9 \\ 15 & 29 & 16 & 21 \\ 31 & 40 & 27 & 33 \end{pmatrix}$	$=$				$\begin{pmatrix} 4 \\ 3 \\ 2 \\ 5 \\ 6 \end{pmatrix}$	$\begin{pmatrix} 1 \\ 5 \\ 2 \\ 3 \end{pmatrix}^\top$	$+$	$\begin{pmatrix} 3 \\ 1 \\ 1 \\ 2 \\ 5 \end{pmatrix}$	$\begin{pmatrix} 5 \\ 2 \\ 3 \\ 3 \end{pmatrix}^\top$
Bob										
Carol										
Dave										
Eve										

**Figure 2.2:** Score Matrix  $M$  is an example for the scores of students (indexing the rows) in different tests on distinct subjects (indexing the columns). A corresponding low rank decomposition is also provided where the rank is two in this example.

According to the simplified hypothesis, each student has different quantitative and verbal strengths. Each subject also requires different levels of quantitative and verbal skills or strength. Intuitively, a student with higher strength on verbal intelligence should perform better on a test that has a high weight on verbal intelligence. Therefore, as a simplest model we can describe the relationship as a bi-linear function:

$$\begin{aligned} \text{Score}(\text{student}, \text{test}) &= \text{student}_{\text{verbal-intlg.}} \times \text{test}_{\text{verbal}} & (2.2) \\ &+ \text{student}_{\text{quant-intlg.}} \times \text{test}_{\text{quant.}} \end{aligned}$$

If we let  $u_{\text{verbal}}, u_{\text{quant.}} \in \mathbb{R}^n$  be vectors that describe the verbal/quantitative strength for each student, and let  $v_{\text{verbal}}, v_{\text{quant.}} \in \mathbb{R}^m$  be vectors that describe the requirement for each test, then we can write the score matrix  $M$  as

$$M = u_{\text{verbal}} v_{\text{verbal}}^\top + u_{\text{quant.}} v_{\text{quant.}}^\top. \tag{2.3}$$

Therefore,  $M$  is a rank 2 matrix! Here quantitative and verbal are two *factors* that influence the result of the tests. The matrix  $M$  is low rank because there are only two different factors. In general, this approach is called *factor analysis*. See Figure 2.2 for an example of matrix  $M$  and its corresponding rank 2 decomposition.

	Math.	Classics	Physics	Music		Verbal	Quantitative				
Alice	$\begin{pmatrix} 19 & 26 & 17 & 21 \\ 8 & 17 & 9 & 12 \\ 7 & 12 & 7 & 9 \\ 15 & 29 & 16 & 21 \\ 31 & 40 & 27 & 33 \end{pmatrix}$				=	$\begin{pmatrix} 4 \\ 3 \\ 2 \\ 5 \\ 6 \end{pmatrix}$	$\begin{pmatrix} 1 \\ 5 \\ 2 \\ 3 \end{pmatrix}^\top$	+	$\begin{pmatrix} 3 \\ 1 \\ 1 \\ 2 \\ 5 \end{pmatrix}$	$\begin{pmatrix} 5 \\ 2 \\ 3 \\ 3 \end{pmatrix}^\top$	
Bob											
Carol											
Dave											
Eve											
					=	$\begin{pmatrix} 1 \\ 2 \\ 1 \\ 3 \\ 1 \end{pmatrix}$	$\begin{pmatrix} 1 \\ 5 \\ 2 \\ 3 \end{pmatrix}^\top$	+	$\begin{pmatrix} 3 \\ 1 \\ 1 \\ 2 \\ 5 \end{pmatrix}$	$\begin{pmatrix} 6 \\ 7 \\ 5 \\ 6 \end{pmatrix}^\top$	

**Figure 2.3:** Two possible decompositions of the score matrix  $M$  that we originally proposed in Figure 2.2. Note that the students verbal intelligence and tests quantitative weights are different between two decompositions.

### 2.1.1 Ambiguity of Matrix Decomposition

As we described, decompositions like (2.3) are very useful as they suggest the whole  $n \times m$  matrix can be explained by a small number of components. However, if we are not only interested in the number of components, but also the exact values of the components (e.g., which student is strongest in the quantitative tasks), such decompositions are not sufficient because they are *not unique*. As an example, in Figure 2.3, we give two different decompositions of the matrix we proposed earlier.

In fact, this phenomena of non-uniqueness of matrix decomposition is very general. Consider a low rank decomposition  $M = \sum_{j=1}^k u_j v_j^\top \in \mathbb{R}^{n \times m}$ . Let  $U \in \mathbb{R}^{n \times k}$  be a matrix whose columns are  $u_j$ 's, and let  $V \in \mathbb{R}^{m \times k}$  be a matrix whose columns are  $v_j$ 's. Then we can represent  $M$  as

$$M = \sum_{j=1}^k u_j v_j^\top = UV^\top.$$



Now for any *orthonormal matrix*  $R \in \mathbb{R}^{k \times k}$  that satisfies  $RR^\top = R^\top R = I$ , we have

$$M = UV^\top = URR^\top V^\top = (UR)(VR)^\top.$$

Therefore,  $UR, VR$  defines an *equivalent* decomposition, and its components (columns of  $UR, VR$ ) can be completely different from the components in the original decomposition  $UV^\top$ . Later in Section 3.1 we will revisit this example and see why *tensor decomposition* can avoid this ambiguity.

## 2.2 Low Rank Matrix Approximation and SVD

In practice, the matrix we are working on is often not *exactly* low rank. The observed matrix can deviate from the low rank structure for many reasons including but not limited to:

- The observed values can be *noisy*.
- The factors may not interact linearly.
- There might be several prominent factors as well as many small factors.

Despite all these possible problems, the observed matrix can still be *approximately* low rank. In such cases it is beneficial to find the low rank matrix that is the *closest* to the observed matrix (in other words, that best approximates it). In this section, we describe Singular Value Decomposition (SVD) method which is an elegant way of finding the closest low rank approximation of a matrix. To do so, we first define matrix norms and provide a concrete notion of closeness in matrix approximation.

### 2.2.1 Matrix Norms

Before talking about how to find the closest matrix, we need to first define when two matrices are close. Closeness is often defined by a distance function  $d(A, B)$  for two same-size matrices  $A$  and  $B$ . For general matrices, the most popular distance functions are based on matrix norms, i.e.,  $d(A, B) = \|A - B\|$  for some matrix norm  $\|\cdot\|$ .

There are many ways to define norms of matrices. The *Frobenius* norm and *spectral/operator* norm are the most popular ones.

**Definition 2.1** (Frobenius norm). The Frobenius norm of a matrix  $M \in \mathbb{R}^{n \times m}$  is defined as

$$\|M\|_F := \sqrt{\sum_{i=1}^n \sum_{j=1}^m M_{i,j}^2}.$$

Frobenius norm is intuitive and easy to compute. However, it ignores the matrix structure and is therefore equivalent to  $\ell_2$  norm when we view the matrix as a vector. To understand the property of the matrix, we can view the matrix as a linear operator, and define its operator norm as follows.

**Definition 2.2** (Matrix spectral/operator norm). The spectral or operator norm of a matrix  $M \in \mathbb{R}^{n \times m}$  is defined as

$$\|M\| := \sup_{\|v\| \leq 1} \|Mv\|,$$

where  $\|\cdot\|$  denotes the Euclidean  $\ell_2$  norm for vectors.

The spectral norm measures how much the matrix can stretch a vector that is inside the unit sphere.

Based on the above two norms, we can now define the closest low rank matrices as

$$M_k := \arg \min_{\text{rank}(N) \leq k} \|M - N\|,$$

$$M_{k,F} := \arg \min_{\text{rank}(N) \leq k} \|M - N\|_F.$$

Both optimization problems are non-convex and may seem difficult to solve. Luckily, both of them can be solved by Singular Value Decomposition. In fact they have exactly the same solution, i.e.,  $M_k = M_{k,F}$ , as we will see in the following section.

### 2.2.2 Singular Value Decomposition

For a matrix  $M$ , the Singular Value Decomposition (SVD) is a special type of low rank decomposition where all the rank-1 components are orthogonal to each other.

**Definition 2.3** (Singular Value Decomposition(SVD), see Golub and Van Loan, 1990 2.5.3 or Horn and Johnson, 2012 7.3.1). The singular value decomposition of matrix  $M \in \mathbb{R}^{n \times m}$  is defined as

$$M = UDV^\top = \sum_{j=1}^{\min\{n,m\}} \sigma_j u_j v_j^\top,$$

where  $U := [u_1|u_2|\cdots|u_n] \in \mathbb{R}^{n \times n}$  and  $V := [v_1|v_2|\cdots|v_m] \in \mathbb{R}^{m \times m}$  are orthonormal matrices such that  $U^\top U = I$ ,  $V^\top V = I$ , and  $D \in \mathbb{R}^{n \times m}$  is a diagonal matrix whose diagonal entries are  $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_{\min\{n,m\}} \geq 0$ . The  $u_j$ 's (respectively  $v_j$ 's) are called the left (respectively right) singular vectors of  $M$  and  $\sigma_j$ 's are called the singular values of  $M$ .

Note that when  $n < m$ , we often view  $D$  as a  $n \times n$  diagonal matrix, and  $V$  as an  $m \times n$  orthonormal matrix because the extra columns of  $V$  (columns indexed by  $n < j \leq m$ ) are not relevant in the decomposition. Similarly when  $n > m$ , we often view  $U$  as a  $n \times m$  matrix.

The top singular value  $\sigma_1$  is the largest singular value that is often denoted as  $\sigma_{\max}(M)$ , and the value  $\sigma_{\min\{n,m\}}$  is the smallest singular value that is often denoted as  $\sigma_{\min}(M)$ . We now describe the optimization view-point of SVD where singular values are the maximum values of the quadratic form  $u^\top M v$  when both  $u$  and  $v$  have bounded  $\ell_2$  norms, and the corresponding components (called singular vectors) are the maximizers that are orthonormal vectors.

**Definition 2.4** (Optimization view-point of the SVD, see Horn and Johnson, 2012 7.3.10). The top singular value  $\sigma_1$  is the maximum of the quadratic form  $u^\top M v$  when  $u$  and  $v$  have bounded  $\ell_2$  norm, and the top singular vectors are the maximizers, i.e.,

$$\begin{aligned} \sigma_1 &= \max_{\|u\| \leq 1, \|v\| \leq 1} u^\top M v, \\ u_1, v_1 &= \arg \max_{\|u\| \leq 1, \|v\| \leq 1} u^\top M v. \end{aligned}$$

The remaining values/vectors are obtained by maximizing the same quadratic form, while constraining the singular vectors to be orthogonal

with all the previous ones, i.e.,

$$\begin{aligned}\sigma_j &= \max_{\|u\|\leq 1, \|v\|\leq 1, \forall i < j: u \perp u_i, v \perp v_i} u^\top M v, \\ u_j, v_j &= \arg \max_{\|u\|\leq 1, \|v\|\leq 1, \forall i < j: u \perp u_i, v \perp v_i} u^\top M v.\end{aligned}$$

As a result, we can also conclude that the spectral norm of  $M$  is equal to  $\sigma_1$ , i.e.,  $\|M\| = \sigma_1$ , since  $\|Mv\| = \max_{\|u\|\leq 1} u^\top M v$ . The singular values and singular vectors are also closely related to the eigenvalues and eigenvectors as we will demonstrate below.

**Lemma 2.1** (SVD vs. eigen-decomposition, see Horn and Johnson, 2012 7.3.5). For a matrix  $M$ , the singular values  $\sigma_j$ 's are the square roots of the eigenvalues of  $MM^\top$  or  $M^\top M$ . The left singular vectors  $u_j$ 's are eigenvectors of  $MM^\top$ , and the right singular vectors  $v_j$ 's are eigenvectors of  $M^\top M$ .

In Section 2.1.1, we described how low rank matrix decomposition is not unique under orthogonal transformation of the rank-1 components. For SVD, because of the specific structure of its singular vectors this is not necessarily the case and in most cases Singular Value Decomposition is unique.

**Theorem 2.2** (Uniqueness of Singular Value Decomposition, see Horn and Johnson, 2012 7.3.5). The SVD of matrix  $M$  defined in Definition 2.3 is unique (for the first  $\min(n, m)$  columns of  $U, V$ ) when the singular values  $\sigma_j$ 's are all distinct and nonzero.

Note that if  $n$  and  $m$  are different, say  $n < m$ , then the last  $m - n$  columns of the matrix  $V$  can be an arbitrary orthogonal basis that is orthogonal to the previous  $n$  right singular vectors, so that is never unique. On the other hand, these columns in  $V$  do not change the result of  $UDV^\top$ , so the decomposition  $\sum_{j=1}^{\min\{n, m\}} \sigma_j u_j v_j^\top$  is still unique. Following the optimization view-point of SVD in Definition 2.4, it is standard to sort the diagonal entries of  $D$  in descending order. In many applications we only care about the top- $k$  components of the SVD, which suggests the following definition of truncated SVD.

**Definition 2.5** (Truncated SVD, see Golub and Van Loan, 1990 2.5.4). Suppose  $M = UDV^\top$  is the SVD of  $M \in \mathbb{R}^{n \times m}$  and entries of  $D$  are sorted in descending order. Let  $U_{(k)} \in \mathbb{R}^{n \times k}, V_{(k)} \in \mathbb{R}^{m \times k}$  denote the matrices only including the first  $k$  columns of  $U \in \mathbb{R}^{n \times n}, V \in \mathbb{R}^{m \times m}$ , respectively, and  $D_{(k)}$  be the first  $k \times k$  submatrix of  $D$ . Then  $U_{(k)}D_{(k)}V_{(k)}^\top$  is called the top- $k$  (rank- $k$ ) truncated SVD of  $M$ .

The truncated SVD can be used to approximate a matrix, and it is optimal in both Frobenius and spectral norms as follows.

**Theorem 2.3** (Eckart-Young theorem(Eckart and Young, 1936): optimality of low rank matrix approximation). Let  $M = \sum_{j=1}^{\min\{n,m\}} \sigma_j u_j v_j^\top$  be the SVD of matrix  $M \in \mathbb{R}^{n \times m}$ , and  $M_k = \sum_{j=1}^k \sigma_j u_j v_j^\top$  be the *truncated* SVD of  $M$ . Then  $M_k$  is the best rank- $k$  approximation of  $M$  in the senses:

$$\|M - M_k\| = \sigma_{k+1} = \min_{\text{rank}(N) \leq k} \|M - N\|,$$

$$\|M - M_k\|_F = \sqrt{\sum_{j=k+1}^{\min\{n,m\}} \sigma_j^2} = \min_{\text{rank}(N) \leq k} \|M - N\|_F.$$

In addition to the above theoretical guarantees on the optimality of low rank matrix approximation, the SVD of a matrix can be computed efficiently. For general matrices the computation takes time  $O(nm \min\{n, m\})$ . The truncated SVD can usually be computed much faster, especially when the  $k$ -th singular value  $\sigma_k$  is significantly larger than the  $(k + 1)$ -th singular value (see Golub and Van Loan, 1990 Section 8.2, together with the discussions in Section 8.6).

We conclude this section by stating the application of SVD in computing the pseudo-inverse of a matrix.

**Definition 2.6** (Moore-Penrose Pseudo-inverse, see Moore, 1920; Bjerhammar, 1951; Penrose, 1955). Given a matrix  $M$  of rank  $k$ , suppose its top- $k$  truncated SVD is  $M = UDV^\top$ , then the pseudo-inverse of  $M$  is defined as  $M^\dagger = VD^{-1}U^\top$ .

Let  $P_r (= VV^\top)$  and  $P_c (= UU^\top)$  to be the projection matrix to the row-span and column-span of  $M$ , respectively; the pseudo-inverse is the only matrix that satisfies  $MM^\dagger = P_c$  and  $M^\dagger M = P_r$ .

In the next few sections we describe some of other major applications of SVD to data analysis.

### 2.3 Principal Component Analysis

In this section, we describe Principle Component Analysis (PCA) (Pearson, 1901; Hotelling, 1933)) as one of the very important and useful statistical methods for data analysis and transformation. Given data points  $x_1, x_2, \dots, x_n \in \mathbb{R}^d$  that for simplicity we assume are *centered* ( $\sum_{i=1}^n x_i = 0$ ), we are often interested in the *covariance matrix*  $M \in \mathbb{R}^{d \times d}$ :

$$M := \frac{1}{n} \sum_{i=1}^n x_i x_i^\top$$


to describe the statistical properties of the data. This matrix measures how different coordinates of the data are correlated with each other. The covariance matrix  $M$  is always positive semi-definite (PSD), and for PSD matrices the SVD always has a symmetric structure such that the left and right singular vectors are the same:

$$M = UDU^\top = \sum_{j=1}^d \sigma_j u_j u_j^\top.$$

Given the covariance matrix, we can easily compute the *variance* of the data when projected to a particular direction. Suppose  $v$  is a unit vector, then we have

$$\text{Var}[\langle v, x \rangle] = \mathbb{E}[\langle v, x \rangle^2] = \mathbb{E}[v^\top x x^\top v] = v^\top \mathbb{E}[x x^\top] v = v^\top M v.$$

Here we used the fact that matrices are linear operators, and the linearity of the expectation. From this calculation and the optimization view-point of SVD in Definition 2.4, it is immediately concluded that the top singular vector  $u_1$  is the direction where the data has largest *variance* when projected to that direction, i.e., yielding the maximum  $\frac{1}{n} \sum_{i=1}^n \langle x_i, v \rangle^2$ . This direction is usually called the *principal component*



**Figure 2.4:** Geometric representation of principle components  $u_1$  and  $u_2$  for the covariance matrix of data points  $x_i$ 's.

as it is the direction where the data is the most “spread out”. Similarly, the first  $k$  singular vectors  $u_1, u_2, \dots, u_k$  spans a subspace that has the maximum variance of all  $k$ -dimensional subspaces. Geometrically, we can view the covariance matrix of the data as an *ellipsoid*, and  $u_1$  corresponds to the longest axis; see Figure 2.4 for such geometrical representation in 2-dimensional space.

### 2.3.1 Dimensionality Reduction via PCA

Principal components can be used to form a lower dimensional subspace and project the data to that subspace. This projection simplifies the data to a much lower dimensional space, while maintaining as much variance of the data as possible as we showed earlier. PCA is the most popular tool for *dimensionality reduction* and the main *linear* technique for doing that. In the case when the data is inherently low rank (recall the test scores example in Section 2.1) but may have some noise, doing PCA can often reduce the magnitude of noise (concrete settings where this can be proved includes mixture of Gaussians, see e.g. Blum *et al.*, 2016 3.9.3).

We now formulate the dimensionality reduction problem more concretely, and provide the guarantee on the optimality of PCA. Given  $n$  data points  $x_1, x_2, \dots, x_n \in \mathbb{R}^d$ , we want to approximate these points with their projection to a lower dimensional subspace in  $\mathbb{R}^d$ . The question is what is the best  $k$ -dimensional ( $k < d$ ) affine subspace in  $\mathbb{R}^d$  for

such approximations, in the sense that the average distance between the original and approximate points is minimized, i.e.,

$$(P^*, p_0^*) := \arg \min_{\substack{P \in \mathbb{R}^{d \times d} \\ p_0 \in \mathbb{R}^d}} \frac{1}{n} \sum_{i \in [n]} \|x_i - (Px_i + p_0)\|^2, \quad (2.4)$$

s. t.  $\text{Rank}(P) = k, P^2 = P.$

Here,  $Px_i + p_0$  is the projection of  $x_i$  to the  $k$ -dimensional affine subspace in  $\mathbb{R}^d$ . This projection is specified by projection operator  $P \in \mathbb{R}^{d \times d}$  and displacement vector  $p_0 \in \mathbb{R}^d$  (Here, we assume the data points are not necessarily centered). Note that since the projection is on a  $k$ -dimensional subspace, we have  $\text{Rank}(P) = k$ . The following theorem shows that PCA is the optimal solution to this problem. This can be proved as a direct corollary of Theorem 2.3 on optimality of low rank matrix approximation using SVD.

**Theorem 2.4** (PCA is optimal solution of (2.4)). Given  $n$  data points  $x_1, x_2, \dots, x_n \in \mathbb{R}^d$ , let  $\mu \in \mathbb{R}^d$  and  $M \in \mathbb{R}^{d \times d}$  denote the corresponding mean vector and covariance matrix, respectively. Let  $M$  have SVD decomposition (the same as eigen-decomposition here)  $M = UDU^\top$ . Then, the optimal solutions of (2.4) are given by

$$P^* = U_{(k)}U_{(k)}^\top,$$

$$p_0^* = (I - P^*)\mu,$$


where  $U_{(k)} := [u_1 | u_2 | \dots | u_k] \in \mathbb{R}^{d \times k}$  is the matrix including the top  $k$  eigenvectors of  $M$ .

**Proof:** Fixing  $P$ , the  $p_0$  which minimizes cost function in (2.4) is  $p_0^* = (I - P)\mu$ . Therefore, we have

$$\begin{aligned} \sum_{i \in [n]} \|x_i - (Px_i + p_0)\|^2 &= \sum_{i \in [n]} \|(I - P)(x_i - \mu)\|^2 \\ &= \sum_{i \in [n]} \|x_i - \mu\|^2 - \sum_{i \in [n]} \|P(x_i - \mu)\|^2, \end{aligned}$$

where we used Pythagorean theorem in the last equality; see Figure 2.5 for its visualization. Therefore, the optimal solution  $P^*$  maximizes the





**Figure 2.5:** Visualization of the Pythagorean relation used in the proof of Theorem 2.4.

variance of projected points into the lower dimensional subspace as

$$\begin{aligned}
 \text{var}(PX) &= \frac{1}{n} \sum_{i \in [n]} \|P(x_i - \mu)\|^2 = \frac{1}{n} \sum_{i \in [n]} (x_i - \mu)^\top P^\top P (x_i - \mu) \\
 &= \frac{1}{n} \sum_{i \in [n]} \text{Tr} \left[ P(x_i - \mu)(x_i - \mu)^\top P^\top \right] \\
 &= \text{Tr} \left[ PMP^\top \right].
 \end{aligned}$$

From Rayleigh quotient argument, we know that for the case of  $k = 1$ , the  $P^*$  which maximizes above is  $P^* = u_1 u_1^\top$ . Similar argument can be extended to larger  $k$  which leads to  $P^* = U_{(k)} U_{(k)}^\top$ .  $\square$

From the above proof, we again see that PCA selects the lower dimensional subspace which has the maximum variance of projected points. Of course, the quality of this approximation still depends on the rank  $k$  that we choose. When the data is assumed to come from a generative model, one can often compute  $k$  by looking for a spectral gap (see e.g., Chapters 7 and 9 in Blum *et al.*, 2016). In practice, one can first choose an accuracy and then find the smallest  $k$  that achieves the desired accuracy.

## 2.4 Whitening Transformation

Another popular application of Singular Value Decomposition is to transform the data into *isotropic* position. We call a data set  $z_1, z_2, \dots, z_n \in$

$\mathbb{R}^d$  isotropic or *whitened* if the covariance matrix

$$M_z := \frac{1}{n} \sum_{i=1}^n z_i z_i^\top = I_d,$$

where  $I_d$  denotes the  $d$ -dimensional identity matrix. This basically means that the data has the *same* amount of variance in every direction. Whitening transformation has been discovered and applied in many domains (Friedman, 1987; Koivunen and Kostinski, 1999). The benefit of whitening transformation is that the result is *invariant* under linear transformations of the original data. Raw data is often not measured in the most natural way – think again about the test score example in Section 2.1, it is possible that a math exam is graded in 100 points and a writing exam has points in the range of 0 to 5. A naïve algorithm might incorrectly think that correlations with math exam is much more important because the scale is 20 times more than the writing exam. However, change of scaling is also a linear transformation, therefore applying whitening transformation to the data can avoid these misconceptions.

In order to do this, suppose the original data is  $x_1, x_2, \dots, x_n \in \mathbb{R}^d$  whose covariance matrix  $M_x := \frac{1}{n} \sum_{i=1}^n x_i x_i^\top$  is not the identity matrix. Intuitively, we would like to shrink the directions that have more variance and stretch the directions that have less variance. This can again be done by SVD as follows. Recall  $M_x$  is a PSD matrix whose SVD can be written as  $M_x = UDU^\top$ . Construct the *whitening matrix*

$$W := UD^{-1/2}, \tag{2.5}$$

and let  $z_i := W^\top x_i$ . Now we have

$$M_z := \frac{1}{n} \sum_{i=1}^n z_i z_i^\top = W^\top M_x W = D^{-1/2} U^\top (UDU^\top) UD^{-1/2} = I_d,$$

and hence, the transformed data points  $z_i$ 's are isotropic.

Note that the result of whitening transformation can be very fragile if the smallest singular value of the data is very close to 0; see the inversion in  $D^{-1/2}$ . In practice, whitening is often performed after we identify the important directions using Principle Component Analysis. Similar whitening idea is useful later for designing tensor decomposition algorithms which we describe in Section 3.4.3.

## 2.5 Canonical Correlation Analysis

All of the techniques that we have discussed so far (SVD, PCA, whitening) focus on extracting the properties of a single data set. On the other side, often in practice we would also like to understand relationships between two different sets of data. In this section, we describe the Canonical Correlation Analysis (CCA) (Hotelling, 1992), which is a very useful method to analyze the cross-covariance matrix between two different data sets. Many of the concepts introduced earlier, such as SVD and whitening, are used to describe the CCA.

Consider two sets of data points  $x_1, x_2, \dots, x_n \in \mathbb{R}^{d_1}$  and  $y_1, y_2, \dots, y_n \in \mathbb{R}^{d_2}$ . If we again use the test scores example from Section 2.1, the first set of vectors  $x_i$  would represent the performance of student  $i$  in different exams, while the second set of vectors  $y_i$  would represent other properties of the students, e.g., the student's future salary. A natural question is whether these two data sets are *correlated*. Intuitively, the correlation between two directions  $u \in \mathbb{R}^{d_1}$  and  $v \in \mathbb{R}^{d_2}$  can be defined as  $\frac{1}{n} \sum_{i=1}^n \langle u, x_i \rangle \langle v, y_i \rangle$ . However, this definition is not very robust – if we apply a linear transformation to  $x$  (say we multiply the first coordinate of  $x$  by  $10^6$ ), then the maximum correlation is likely to be changed (in this case likely to have more weight on first coordinate of  $x$ ). Therefore, to measure correlations robustly and get rid of the influence from individual data sets, Canonical Correlation Analysis tries to find the maximum correlation after *whitening* both  $x$  and  $y$ ; see Algorithm 1 for the details. In the remaining of this section, we describe how the CCA algorithm is designed and works.

Let  $M_x, M_y$  be the covariance matrices of  $\{x_i\}$ 's and  $\{y_i\}$ 's, and let  $W_x$  and  $W_y$  be the corresponding whitening matrices; see Algorithm 1 for the precise definitions. Let

$$\tilde{x}_i := W_x^\top x_i, \quad \tilde{y}_i := W_y^\top y_i$$

be the whitened data. We would like to find the most correlated directions in this pair of whitened data, i.e., we would like to find unit vectors  $\tilde{u}, \tilde{v}$  such that  $u^\top [\frac{1}{n} \sum_{i=1}^n \tilde{x}_i \tilde{y}_i^\top] v$  is maximized, i.e.,

$$\tilde{u}, \tilde{v} := \arg \max_{\|u\|=\|v\|=1} \frac{1}{n} \sum_{i=1}^n \langle u, \tilde{x}_i \rangle \langle v, \tilde{y}_i \rangle.$$

---

**Algorithm 1** Canonical Correlation Analysis (CCA)

---

**input** Two data sets  $x_1, x_2, \dots, x_n \in \mathbb{R}^{d_1}$  and  $y_1, y_2, \dots, y_n \in \mathbb{R}^{d_2}$ **output** Most cross-correlated directions between whitened pairs

- 1: Compute the covariance matrices

$$M_x := \frac{1}{n} \sum_{i=1}^n x_i x_i^\top, \quad M_y := \frac{1}{n} \sum_{i=1}^n y_i y_i^\top.$$

- 2: Use SVD to compute the whitening matrices
- $W_x, W_y$
- ; see Equation (2.5).

- 3: Compute the correlation matrix

$$M_{\tilde{x}\tilde{y}} := \frac{1}{n} \sum_{i=1}^n W_x^\top x_i y_i^\top W_y.$$

- 4: Use SVD to compute the left and right singular vectors
- $\{(\tilde{u}_j, \tilde{v}_j)\}$
- for
- $M_{\tilde{x}\tilde{y}}$
- .

- 5:
- return**
- $(W_x \tilde{u}_i, W_y \tilde{v}_i)$
- .
- 

By Definition 2.4, it is immediate to see that  $\tilde{u}, \tilde{v}$  are actually the left and right top singular vectors of the cross-covariance matrix  $M_{\tilde{x}\tilde{y}} := \frac{1}{n} \sum_{i=1}^n \tilde{x}_i \tilde{y}_i^\top$ , and this pair of directions are where the two data sets are most correlated. It is also possible to define more pairs of vectors  $(\tilde{u}_j, \tilde{v}_j)$ 's that correspond to the smaller singular vectors of the same matrix.

Often we would like to interpret the direction in the original data sets instead of the whitened ones. To do that, we would like to find a vector  $u$  such that  $\langle u, x_i \rangle = \langle \tilde{u}, \tilde{x}_i \rangle$ . That is,

$$u^\top x_i = \tilde{u}^\top \tilde{x}_i = \tilde{u}^\top W_x^\top x_i,$$

and thus, we need to have

$$u = W_x \tilde{u}.$$

On the other hand, by construction we have  $W_x^\top M_x W_x = I$ , and therefore, since  $\tilde{u}$  has unit norm,  $\tilde{u}^\top W_x^\top M_x W_x \tilde{u} = 1$ , which is to say

$u^\top M_x u = 1$  given above equality. Similarly, the fact that  $\tilde{u}_j$  and  $\tilde{u}_l$  are orthogonal if  $j \neq l$  means that  $u_j^\top M_x u_l = 0$ . In general, it is possible to define a different inner product

$$\langle u_j, u_l \rangle_{M_x} := u_j^\top M_x u_l,$$

and a corresponding vector norm  $\|u\|_{M_x} := \sqrt{u^\top M_x u}$ , and the vectors  $u_j$ 's will be orthonormal under this new inner product  $\langle \cdot, \cdot \rangle_{M_x}$ . Similarly,  $v_j$ 's should be orthonormal under the inner product  $\langle \cdot, \cdot \rangle_{M_y}$ . Using these constraints, we can describe Canonical Component Analysis more precisely as below

**Definition 2.7** (Canonical Correlation Analysis). Given two data sets  $x_1, x_2, \dots, x_n \in \mathbb{R}^{d_1}$  and  $y_1, y_2, \dots, y_n \in \mathbb{R}^{d_2}$  (without loss of generality, assume  $d_1 \leq d_2$ ), let

$$M_x := \frac{1}{n} \sum_{i=1}^n x_i x_i^\top, \quad M_y = \frac{1}{n} \sum_{i=1}^n y_i y_i^\top$$

be the corresponding covariance matrices, respectively. Canonical Correlation Analysis (CCA) finds a set of correlated directions  $u_1, u_2, \dots, u_{d_1}$  and  $v_1, v_2, \dots, v_{d_1}$  such that  $u_j^\top M_x u_j = 1$ ,  $v_j^\top M_y v_j = 1$ . The top correlated directions  $u_1, v_1$  are similar to the top singular vectors as

$$u_1, v_1 = \arg \max_{u^\top M_x u = 1, v^\top M_y v = 1} \frac{1}{n} \sum_{i=1}^n \langle u, x_i \rangle \langle v, y_i \rangle.$$

Similarly, the remaining most correlated directions are defined as the remaining singular vectors

$$u_j, v_j = \arg \max_{\substack{u^\top M_x u = 1, v^\top M_y v = 1 \\ \forall l < j \quad u^\top M_x u_l = 0, v^\top M_y v_l = 0}} \frac{1}{n} \sum_{i=1}^n \langle u, x_i \rangle \langle v, y_i \rangle.$$

The derivation of Canonical Correlation Analysis already gives an efficient algorithm as provided in Algorithm 1. It is not hard to verify the correctness of this algorithm, because after the linear transforms  $W_x$  and  $W_y$ , the objective and constraints in Definition 2.7 become exactly the same as those in Definition 2.4.

The idea of Canonical Correlation Analysis is widely used in data analysis. In particular, CCA can find directions that are “aligned” in two data sets. The same idea is also used in tensor decompositions to “align” different views of the data, see Section 3.4.3.

Using the full SVD to compute the CCA can be expensive in practice. Recently there have been several works that give efficient algorithms for computing top CCA vectors over large data sets, e.g., see Wang *et al.* (2016), Ge *et al.* (2016), Allen-Zhu and Li (2016b), and Allen-Zhu and Li (2016a).

# 3

---

## Tensor Decomposition Algorithms

---

In this chapter, we first introduce the basic concepts of tensors and state the tensor notations that we need throughout this monograph. In particular, we highlight why in many cases we need to use tensors instead of matrices and provide the guarantees on uniqueness of tensor decomposition. Then, we describe different algorithms for computing tensor decomposition.

Most of the materials in this chapter has appeared in existing literature, especially in Anandkumar *et al.*, 2014a. We do give more explanations on the whitening procedure in Section 3.4.3 and symmetrization procedure in Section 3.4.4, which were used in many previous papers but were not explicitly discussed in their general forms. We also give a new perturbation analysis for tensor power method together with whitening procedure in Section 3.5.3, which will be useful for many of the applications later in Section 4.

### 3.1 Transition from Matrices to Tensors

We can think of tensors as multi-dimensional arrays, and one of the easiest ways to get a tensor is by stacking matrices of the same dimensions resulting in third order tensors. Let us recall the test scores example

proposed in Section 2.1. Now suppose each exam has two parts – *written* and *oral*. Instead of the single score matrix  $M$  that we had before, we will now have two score matrices  $M_{\text{written}}$  and  $M_{\text{oral}}$  including the scores for written and oral exams, respectively. Similar to the earlier score matrix  $M$ , the rows of these matrices are indexed by students, and their columns are indexed by subjects/tests. When we stack these two matrices together, we get a  $n \times m \times 2$  tensor, where the third dimension is now indexed by the test format (written or oral). See the tensor in the left hand side of Figure 3.1 as the stacking of two score matrices  $M_{\text{written}}$  and  $M_{\text{oral}}$ .

Recall the simplified hypothesis states that there are two kinds of intelligence – *quantitative* and *verbal*; see Section 2.1 to review it. Different students have different strengths, and different subjects/tests also have different requirements. As a result, the score was assumed to be a bilinear function of these hidden components; see (2.2). Now with the third dimension, it is also reasonable to expect the two kind of intelligence might behave differently in different formats – intuitively, verbal skills might be slightly more important in oral exams. As a result, we can generalize the bilinear function to a *tri-linear* form as


$$\begin{aligned} \text{Score}(\text{student}, \text{test}, \text{format}) = \\ \text{student}_{\text{verbal-intlg.}} \times \text{test}_{\text{verbal}} \times \text{format}_{\text{verbal}} \\ + \text{student}_{\text{quant-intlg.}} \times \text{test}_{\text{quant.}} \times \text{format}_{\text{quant.}}, \end{aligned}$$

where  $\text{format}_{\text{verbal}}$  and  $\text{format}_{\text{quant.}}$  denote the importance of verbal and quantitative intelligence in different formats, respectively. Now similar to what we did for the matrices, we can propose the following formula as *decomposing* the tensor into the sum of two rank-1 components as

$$\begin{aligned} (M_{\text{written}}, M_{\text{oral}}) = u_{\text{verbal}} \otimes v_{\text{verbal}} \otimes w_{\text{verbal}} \\ + u_{\text{quant.}} \otimes v_{\text{quant.}} \otimes w_{\text{quant.}}. \end{aligned} \quad (3.1)$$

Here  $\otimes$  is the *tensor/outer product* operator which we will formally define in the next subsection; see (3.5).  $u_{\text{verbal}}, u_{\text{quant.}} \in \mathbb{R}^n$  and  $v_{\text{verbal}}, v_{\text{quant.}} \in \mathbb{R}^m$  are the same as in Section 2.1, and the new components  $w_{\text{verbal}}, w_{\text{quant.}} \in \mathbb{R}^2$  correspond to verbal/quantitative importance for different formats, e.g.,  $w_{\text{verbal}}(\text{oral})$  denotes the importance of verbal intelligence in tests





**Figure 3.1:** Graphical representation of CP decomposition for score tensor in (3.1). The matrix slices (oral and written) share a 2D-decomposition along the first two modes, and having the different weight factors in the third mode. The symbols ★ and ○ represent the importance weight factors for oral and written formats, respectively.

with oral format. This is a natural generalization of matrix decomposition/rank to tensors, which is commonly referred to as the CP (CANDECOMP/PARAFAC) (Hitchcock, 1927; Carroll and Chang, 1970; Harshman and Lundy, 1994) decomposition/rank of tensors; we will formally define that in (3.6). In fact, the tensor CP decomposition in (3.1) can be thought as a shared decomposition of matrices  $M_{\text{written}}$  and  $M_{\text{oral}}$  along the first two modes (corresponding to vectors  $u_{\text{verbal}}, u_{\text{quant.}}, v_{\text{verbal}}, v_{\text{quant.}}$ ) with extra weight factors which are collected in the third mode as vectors  $w_{\text{verbal}}, w_{\text{quant.}}$ . This is graphically represented in Figure 3.1.

**Why using tensors instead of matrices?** Until now going to the tensor format just seems to make things more complicated. What additional benefits do we get? One important property of tensor decomposition is *uniqueness*. When we have only one matrix of test scores, the matrix decomposition is not unique most of the time; recall Figure 2.3 where we provided an example of this situation happening. The ambiguity makes it hard to answer even some of the most basic questions such as: which student has the best quantitative strength? On the other hand, under mild conditions (see Section 3.3 for a formal discussion), the tensor decomposition is unique! Finding the unique decomposition allows us to pin down the vectors for students' strengths.

For learning latent variable models and latent representations, the *uniqueness* of tensor decomposition often translates to *identifiability*. We say a set of statistics makes the model *identifiable*, if there is only a unique set of parameters that can be consistent with what we have observed. Matrix decompositions usually correspond to pairwise correlations. Because of the ambiguities discussed earlier, for most latent variable models, pairwise correlations do not make the model identifiable. On the other hand, since tensor decompositions are unique, once we go to correlations between three or more objects, the models become identifiable. The example of learning a pure topic model was discussed in Section 1.2 and many more examples are provided in Section 4.

Tensor decomposition has also applications in many other areas such as chemometrics (Appellof and Davidson, 1981), neuroscience (Mocks, 1988), telecommunications (Sidiropoulos *et al.*, 2000), data mining (Acar *et al.*, 2005), image compression and classification (Shashua and Levin, 2001), and so on; see survey paper by Kolda and Bader (2009) for more references.

**Difficulties in Working with Tensors:** The benefit of unique decomposition comes at a cost. Although we can usually generalize notions for matrices to tensors, their counterpart in tensors are often not as well-behaved or easy to compute. In particular, tensor (CP) decomposition is much harder to compute than matrix decomposition. In fact, almost all tensor problems are NP-hard in the worst-case (Hillar and Lim, 2013). Therefore, we can only hope to find tensor decomposition in special cases. Luckily, this is usually possible when the rank of the tensor is much lower than the size of its modes which is true for many of the applications. Later in Sections 3.4–3.7, we will introduce algorithms for low rank tensor decomposition. When the rank of the tensor is high, especially when the rank is larger than the dimensions (which cannot happen for matrices), we may need more complicated techniques which we discuss in Section 7.

### 3.2 Tensor Preliminaries and Notations

In this section we describe some preliminary tensor concepts and provide formal tensor notations.


A real-valued  $p$ -th order tensor

$$T \in \bigotimes_{i=1}^p \mathbb{R}^{d_i}$$

is a member of the outer product of Euclidean spaces  $\mathbb{R}^{d_i}$ ,  $i \in [p]$ , where  $[p] := \{1, 2, \dots, p\}$ . For convenience, we restrict to the case where  $d_1 = d_2 = \dots = d_p = d$ , and simply write  $T \in \bigotimes^p \mathbb{R}^d$ .

As is the case for vectors (where  $p = 1$ ) and matrices (where  $p = 2$ ), we may identify a  $p$ -th order tensor with the  $p$ -way array of real numbers  $[T_{i_1, i_2, \dots, i_p} : i_1, i_2, \dots, i_p \in [d]]$ , where  $T_{i_1, i_2, \dots, i_p}$  is the  $(i_1, i_2, \dots, i_p)$ -th entry of  $T$  with respect to a canonical basis. A tensor is also called *symmetric* if the entry values are left unchanged by the permutation of any indices. For convenience, we provide the concepts and results only for third order tensors ( $p = 3$ ) in the rest of this section. These can be similarly extended to higher order tensors.

**Tensor modes, fibers and slices:** The different dimensions of the tensor are referred to as *modes*. For instance, for a matrix, the first mode refers to columns and the second mode refers to rows. In addition, *fibers* are higher order analogues of matrix rows and columns. A fiber is obtained by fixing all but one of the indices of the tensor and is arranged as a column vector. For instance, for a matrix, its mode-1 fiber is any matrix column while a mode-2 fiber is any row. For a third order tensor  $T \in \mathbb{R}^{d \times d \times d}$ , the mode-1 fiber is given by  $T(:, j, l)$ , mode-2 by  $T(i, :, l)$  and mode-3 by  $T(i, j, :)$  for fixed indices  $i, j, l$ . Similarly, *slices* are obtained by fixing all but two of the indices of the tensor and are represented as matrices. For example, for the third order tensor  $T \in \mathbb{R}^{d \times d \times d}$ , the slices along 3rd mode are given by  $T(:, :, l)$ . See Figure 3.2 for a graphical representation of tensor fibers and slices for a third order tensor.



**Figure 3.2:** Graphical representations of tensor fibers (left) and tensor slices (right) for a third order tensor.

**Tensor matricization:** Transforming tensors into matrices is one of the ways to work with tensors. For  $r \in \{1, 2, 3\}$ , the mode- $r$  matricization of a third order tensor  $T \in \mathbb{R}^{d \times d \times d}$ , denoted by  $\text{mat}(T, r) \in \mathbb{R}^{d \times d^2}$ , consists of all mode- $r$  fibers arranged as column vectors. For instance, the matricized version along first mode denoted by  $M \in \mathbb{R}^{d \times d^2}$  is defined such that

$$T(i, j, l) = M(i, l + (j - 1)d), \quad i, j, l \in [d]. \tag{3.2}$$

**Multilinear transformation:** We view a tensor  $T \in \mathbb{R}^{d \times d \times d}$  as a multilinear form. Consider matrices  $A, B, C \in \mathbb{R}^{d \times k}$ . Then tensor  $T(A, B, C) \in \mathbb{R}^{k \times k \times k}$  is defined such that


$$T(A, B, C)_{j_1, j_2, j_3} := \sum_{i_1, i_2, i_3 \in [d]} T_{i_1, i_2, i_3} \cdot A(i_1, j_1) \cdot B(i_2, j_2) \cdot C(i_3, j_3). \tag{3.3}$$

See Figure 3.3 for a graphical representation of multilinear form. In particular, for vectors  $u, v, w \in \mathbb{R}^d$ , we have

$$T(I, v, w) = \sum_{j, l \in [d]} v_j w_l T(:, j, l) \in \mathbb{R}^d, \tag{3.4}$$

which is a multilinear combination of the tensor mode-1 fibers. Similarly  $T(u, v, w) \in \mathbb{R}$  is a multilinear combination of the tensor entries, and  $T(I, I, w) \in \mathbb{R}^{d \times d}$  is a linear combination of the tensor slices. These multilinear forms can be similarly generalized to higher order tensors.

In the matrix case of  $M \in \mathbb{R}^{d \times d}$ , all above multilinear forms simplify



**Figure 3.3:** Tensor as a multilinear transformation and representation of Tucker decomposition of a 3rd order tensor  $T = \sum_{j_1, j_2, j_3 \in [k]} S_{j_1, j_2, j_3} \cdot a_{j_1} \otimes b_{j_2} \otimes c_{j_3} = S(A^\top, B^\top, C^\top)$

to familiar matrix-matrix and matrix-vector products such that

$$M(A, B) := A^\top M B \in \mathbb{R}^{k \times k},$$

$$M(I, v) := M v = \sum_{j \in [d]} v_j M(:, j) \in \mathbb{R}^d.$$

**Rank-1 tensor:** A 3rd order tensor  $T \in \mathbb{R}^{d \times d \times d}$  is said to be rank-1 if it can be written in the form

$$T = w \cdot a \otimes b \otimes c \Leftrightarrow T(i, j, l) = w \cdot a(i) \cdot b(j) \cdot c(l), \quad (3.5)$$

where notation  $\otimes$  represents the *outer product* and  $a \in \mathbb{R}^d, b \in \mathbb{R}^d, c \in \mathbb{R}^d$  are unit vectors (without loss of generality) and  $w \in \mathbb{R}$  is the magnitude factor.

Throughout this monograph, we also use notation  $\cdot^{\otimes 3}$  to denote

$$a^{\otimes 3} := a \otimes a \otimes a,$$

for vector  $a$ .

**Tensor CP decomposition and rank:** A tensor  $T \in \mathbb{R}^{d \times d \times d}$  is said to have a CP (CANDECOMP/PARAFAC) rank  $k \geq 1$  if  $k$  is the *minimum* number that the tensor can be written as the sum of  $k$  rank-1 tensors

$$T = \sum_{j \in [k]} w_j a_j \otimes b_j \otimes c_j, \quad w_j \in \mathbb{R}, a_j, b_j, c_j \in \mathbb{R}^d. \quad (3.6)$$



**Figure 3.4:** CP decomposition of a symmetric 3rd order tensor  $T = \sum_j a_j \otimes a_j \otimes a_j$

See Figure 3.4 for a graphical representation of CP decomposition for a symmetric 3rd order tensor. This decomposition is also closely related to the multilinear form. In particular, given  $T$  in (3.6) and vectors  $\hat{a}, \hat{b}, \hat{c} \in \mathbb{R}^d$ , we have

$$T(\hat{a}, \hat{b}, \hat{c}) = \sum_{j \in [k]} w_j \langle a_j, \hat{a} \rangle \langle b_j, \hat{b} \rangle \langle c_j, \hat{c} \rangle.$$

Consider the decomposition in equation (3.6), denote matrix  $A := [a_1 | a_2 | \dots | a_k] \in \mathbb{R}^{d \times k}$ , and similarly  $B$  and  $C$ . Without loss of generality, we assume that the matrices have normalized columns (in  $\ell_2$ -norm), since we can always rescale them, and adjust the weights  $w_j$  appropriately.

As we mentioned in the previous Section, the CP decomposition is often unique, which is very crucial to many machine learning applications. We will formally discuss that in Section 3.3.

**Tensor Tucker decomposition:** A tensor  $T \in \mathbb{R}^{d_1 \times d_2 \times d_3}$  is said to have a Tucker decomposition or Tucker representation when given core tensor  $S \in \mathbb{R}^{k_1 \times k_2 \times k_3}$  and factor matrices  $A \in \mathbb{R}^{d_1 \times k_1}$ ,  $B \in \mathbb{R}^{d_2 \times k_2}$ ,  $C \in \mathbb{R}^{d_3 \times k_3}$ , it can be written as

$$T = \sum_{j_1 \in [k_1]} \sum_{j_2 \in [k_2]} \sum_{j_3 \in [k_3]} S_{j_1, j_2, j_3} \cdot a_{j_1} \otimes b_{j_2} \otimes c_{j_3}. \quad (3.7)$$

See Figure 3.3 for a graphical representation of Tucker representation. Note that this is directly related to the multilinear form defined in (3.3) such that the R.H.S. of above equation is  $S(A^\top, B^\top, C^\top)$ . Note that the CP decomposition is a special case of the Tucker decomposition when the core tensor  $S$  is square (all modes having the same dimension) and diagonal. Unlike CP decomposition, Tucker decomposition suffers the same ambiguity problem as matrix decomposition. Therefore, we will focus on CP decomposition in this monograph. On the other hand,

Tucker decomposition can be computed efficiently, which makes it a better choice for some applications other than learning latent variable models.

**Norms:** For vector  $v \in \mathbb{R}^d$ ,

$$\|v\| := \sqrt{\sum_{i \in [d]} v_i^2}$$

denotes the Euclidean ( $\ell_2$ ) norm, and for matrix  $M \in \mathbb{R}^{d \times d}$ , the spectral (operator) norm is

$$\|M\| := \sup_{\|u\|=\|v\|=1} |M(u, v)|,$$

where  $|\cdot|$  denotes the absolute value operator.

Furthermore,  $\|T\|$  and  $\|T\|_F$  denote the spectral (operator) norm and the Frobenius norm of a tensor, respectively. In particular, for a 3rd order tensor  $T \in \mathbb{R}^{d \times d \times d}$ , we have:

$$\begin{aligned} \|T\| &:= \sup_{\|u\|=\|v\|=\|w\|=1} |T(u, v, w)|, \\ \|T\|_F &:= \sqrt{\sum_{i,j,l \in [d]} T_{i,j,l}^2}. \end{aligned}$$

We conclude this section by reviewing some additional matrix notations and operators that we need throughout this monograph.

**Matrix notations:** For a matrix  $M$  with linearly independent rows, the right pseudo-inverse denoted by  $M^\dagger$  (such that  $MM^\dagger = I$ ) is defined as

$$M^\dagger = M^\top (MM^\top)^{-1}. \quad (3.8)$$

For matrices  $A \in \mathbb{R}^{d_1 \times k}$ ,  $B \in \mathbb{R}^{d_2 \times k}$ , we introduce the following products. The *Khatri-Rao product*, also known as column-wise Kronecker product  $C := A \odot B \in \mathbb{R}^{d_1 d_2 \times k}$  is defined such that

$$C(l + (i - 1)d, j) = A(i, j) \cdot B(l, j), \quad i \in [d_1], l \in [d_2], j \in [k]. \quad (3.9)$$

Furthermore, when  $d_1 = d_2 = d$ , the *Hadamard product*  $C := A * B \in \mathbb{R}^{d \times k}$  is defined as entry-wise product such that

$$C(i, j) = A(i, j) \cdot B(i, j), \quad i \in [d], j \in [k]. \quad (3.10)$$

### 3.3 Uniqueness of CP decomposition

When we are talking about the uniqueness of tensor CP decomposition, there are still some inherent uncertainties even in the formulation of the CP decomposition. For the following decomposition

$$T = \sum_{j \in [k]} w_j a_j \otimes b_j \otimes c_j, \quad w_j \in \mathbb{R}, \quad a_j, b_j, c_j \in \mathbb{R}^d,$$

we can obviously *permute* different rank-1 components, and the result will be the same tensor. We can also *scale* vectors  $a_j, b_j, c_j$  and the weight  $w_j$  simultaneously, as long as the product of all the scalings is equal to 1 and again the result will be the same tensor. The permutation and scaling ambiguities are inherent, and can often be addressed by the particular application. In the test scores example that we have revisited throughout this work, the permutation ambiguity means we get the two rank-1 components, but we do not know which one corresponds to the quantitative and which one corresponds to the verbal factor. In this case, intuitively we know a math test should require more quantitative skill, while a writing test should require more verbal skill. Therefore it should not be hard for a human to give names to the two hidden components. The scaling ambiguity is very similar to measuring quantities using different units, and we can often choose the appropriate scaling, e.g., by enforcing the strengths of students to be within 0-100. Regardless of the scaling/units we choose, the comparison between different students/subjects still makes sense – we can still safely answer questions like which student has the best quantitative strength.

Apart from above inherent ambiguities, there are several sufficient conditions for uniqueness of tensor decomposition. The most well-known condition is formulated by Kruskal (1976) and Kruskal (1977). We first provide the definition of Kruskal rank and then state this uniqueness condition.

**Definition 3.1** (Kruskal rank). The Kruskal rank or *krank* of a matrix  $A$  denoted by  $\text{krank}(A)$  is the maximum number  $r$  such that every subset of  $r$  columns of  $A$  is linearly independent.

**Theorem 3.1** ((Kruskal, 1976; Kruskal, 1977)). The CP decomposition in (3.6) is unique (up to permutation and scaling), if we let  $A :=$



$[a_1 \ a_2 \ \dots \ a_k]$  (similarly for  $B, C$ , all with the same number of column  $k$ ) satisfy the condition

$$\text{krank}(A) + \text{krank}(B) + \text{krank}(C) \geq 2k + 2.$$

This is a mild condition when the rank of the tensor is not too high. As a comparison, matrix decomposition can be unique only when the matrix is rank 1, or we require strong conditions like orthogonality among components. In general, for non-degenerate cases when the components are in general position (with probability 1 for any continuous probability distribution on the components  $A, B, C$ ), the krank of the matrices  $A, B, C$  are equal to  $\min\{k, d\}$  (Kruskal, 1976, see a more robust version in Bhaskara *et al.*, 2014). Therefore, when  $2 \leq k \leq d$ , the Kruskal condition is always satisfied leading to unique tensor CP decomposition. Even when  $k > d$  (rank is higher than the dimension), the Kruskal condition can be satisfied as long as  $k \leq 1.5d - 1$ .

### 3.4 Orthogonal Tensor Decomposition

Tensor decomposition is in general a challenging problem. As a special and more tractable kind of decomposition, we introduce orthogonal tensor decomposition in this section. We review some useful properties of tensors that have orthogonal decomposition, and in the next section, we show how these properties lead to tensor power iteration as a natural algorithm for orthogonal tensor decomposition. It is worth mentioning here that not all tensors have orthogonal decomposition, and as we discussed in the previous section, the tensor decomposition can still be unique even when the tensor rank-1 components are not orthogonal.

We first review the spectral decomposition of symmetric matrices, and then discuss a generalization to higher-order tensors.

#### 3.4.1 Review: Matrix Decomposition

We first build intuition by reviewing the matrix setting, where the desired decomposition is the eigen-decomposition of a symmetric rank- $k$  matrix  $M = V\Lambda V^\top$ , where  $V = [v_1|v_2|\dots|v_k] \in \mathbb{R}^{d \times k}$  is the matrix with orthonormal ( $V^\top V = I$ ) eigenvectors as columns, and  $\Lambda =$

$\text{diag}(\lambda_1, \lambda_2, \dots, \lambda_k) \in \mathbb{R}^{k \times k}$  is diagonal matrix of non-zero eigenvalues. In other words,

$$M = \sum_{j=1}^k \lambda_j v_j v_j^\top = \sum_{j=1}^k \lambda_j v_j^{\otimes 2}. \quad (3.11)$$

Such a decomposition is guaranteed to exist for every symmetric matrix; see Golub and Van Loan, 1990, Chapter 8.

Recovery of the  $v_j$ 's and  $\lambda_j$ 's can be viewed in at least two ways: *fixed point* and *variational* characterizations.

### Fixed-point characterization

First, each  $v_i$  is a *fixed point* under the mapping  $u \mapsto Mu$ , up to a scaling factor  $\lambda_i$ :

$$Mv_i = \sum_{j=1}^k \lambda_j (v_j^\top v_i) v_j = \lambda_i v_i$$

as  $v_j^\top v_i = 0$  for all  $j \neq i$  by orthogonality. The  $v_i$ 's are not necessarily the only such fixed points. For instance, with the multiplicity  $\lambda_1 = \lambda_2 = \lambda$ , then any linear combination of  $v_1$  and  $v_2$  is also fixed under  $M$ . However, in this case, the decomposition in (3.11) is not unique, as  $\lambda_1 v_1 v_1^\top + \lambda_2 v_2 v_2^\top$  is equal to  $\lambda(u_1 u_1^\top + u_2 u_2^\top)$  for any pair of orthonormal vectors  $u_1$  and  $u_2$  spanning the same subspace as  $v_1$  and  $v_2$ . Nevertheless, the decomposition is unique when  $\lambda_1, \lambda_2, \dots, \lambda_k$  are distinct, whereupon the  $v_i$ 's are the only directions fixed under  $u \mapsto Mu$  up to non-trivial scaling; see Theorem 2.2.

### Variational characterization

The second view of recovery is via the *variational characterization* of the eigenvalues. Assume  $\lambda_1 > \lambda_2 > \dots > \lambda_k$ ; the case of repeated eigenvalues again leads to similar non-uniqueness as discussed above. Then the *Rayleigh quotient*

$$\frac{u^\top M u}{u^\top u}$$

is maximized over non-zero vectors by  $v_1$ . Furthermore, for any  $s \in [k]$ , the maximizer of the Rayleigh quotient, subject to being orthogonal to  $v_1, v_2, \dots, v_{s-1}$ , is  $v_s$ . Another way of obtaining this second statement is to consider the *deflated* Rayleigh quotient

$$\frac{u^\top \left( M - \sum_{j=1}^{s-1} \lambda_j v_j v_j^\top \right) u}{u^\top u},$$

and observe that  $v_s$  is the maximizer. Also see that this statement is closely related to the optimization view-point of SVD provided in Definition 2.4.

Efficient algorithms for finding these matrix decompositions are well studied (Golub and Loan, 1996, Section 8.2.3), and iterative power methods are one effective class of algorithms.

We remark that in our multilinear tensor notation, we may write the maps  $u \mapsto Mu$  and  $u \mapsto u^\top Mu / \|u\|_2^2$  as

$$u \mapsto Mu \equiv u \mapsto M(I, u), \tag{3.12}$$

$$u \mapsto \frac{u^\top Mu}{u^\top u} \equiv u \mapsto \frac{M(u, u)}{u^\top u}. \tag{3.13}$$

### 3.4.2 The Tensor Case

Decomposing general tensors is a delicate issue; tensors may not even have unique decomposition. But as we discussed earlier, tensors with orthogonal decomposition have a structure which permits a unique decomposition under a mild non-degeneracy condition.

An *orthogonal decomposition* of a symmetric tensor  $T \in \mathbb{R}^{d \times d \times d}$  is a collection of orthonormal (unit) vectors  $\{v_1, v_2, \dots, v_k\}$  together with corresponding positive scalars  $\lambda_j > 0$  such that

$$T = \sum_{j=1}^k \lambda_j v_j^{\otimes 3}. \tag{3.14}$$

In general, we say a  $p$ -th order symmetric tensor has an orthogonal decomposition if there exists a collection of orthonormal (unit) vectors  $\{v_1, v_2, \dots, v_k\}$  together with corresponding scalars  $\lambda_j$  such that

$$T = \sum_{j=1}^k \lambda_j v_j^{\otimes p}.$$

Note that for odd order tensors (especially  $p = 3$ ), we can add the requirement that the  $\lambda_j$  be positive. This convention can be followed without loss of generality since  $-\lambda_j v_j^{\otimes p} = \lambda_j (-v_j)^{\otimes p}$  whenever  $p$  is odd. Also, it should be noted that orthogonal decompositions do not necessarily exist for every symmetric tensor.

In analogy to the matrix setting, we consider two ways to view this decomposition: a fixed-point characterization and a variational characterization. Related characterizations based on optimal rank-1 approximations are given by Zhang and Golub (2001).

### Fixed-point characterization

For a tensor  $T$ , consider the vector-valued map

$$u \mapsto T(I, u, u) \tag{3.15}$$

which is the third-order generalization of (3.12).

From the definition of multilinear form in (3.4), this can be explicitly written as

$$T(I, u, u) = \sum_{i \in [d]} \sum_{j, l \in [d]} T_{i, j, l} (e_j^\top u) (e_l^\top u) e_i,$$

where  $e_i$  denotes the  $d$ -dimensional basis vector with  $i$ -th entry equal to 1 and the rest of entries being zero. Observe that (3.15) is *not* a linear map, which is a key difference compared to the matrix case where  $Mu = M(I, u)$  is a linear map of  $u$ .

An eigenvector  $u$  for a matrix  $M$  satisfies  $M(I, u) = \lambda u$ , for some scalar  $\lambda$ . We say a unit vector  $u \in \mathbb{R}^d$  is an *eigenvector* of  $T$ , with corresponding *eigenvalue*  $\lambda \in \mathbb{R}$ , if

$$T(I, u, u) = \lambda u.$$

To simplify the discussion, we assume throughout that eigenvectors have unit norm; otherwise, for scaling reasons, we replace the above equation with  $T(I, u, u) = \lambda \|u\|u$ . This concept was originally introduced by Lim

(2005) and Qi (2005)<sup>1</sup>. For orthogonally decomposable tensors  $T = \sum_{j=1}^k \lambda_j v_j^{\otimes 3}$ ,

$$T(I, u, u) = \sum_{j=1}^k \lambda_j (u^\top v_j)^2 v_j .$$

By the orthogonality of the  $v_i$ 's, it is clear that  $T(I, v_i, v_i) = \lambda_i v_i$  for all  $i \in [k]$ . Therefore, each  $(v_i, \lambda_i)$  is an eigenvector/eigenvalue pair.

There are a number of subtle differences compared to the matrix case that arise as a result of the non-linearity of (3.15). First, even with the multiplicity  $\lambda_1 = \lambda_2 = \lambda$ , a linear combination  $u := c_1 v_1 + c_2 v_2$  is not an eigenvector except in very special cases. In particular,

$$T(I, u, u) = \lambda_1 c_1^2 v_1 + \lambda_2 c_2^2 v_2 = \lambda(c_1^2 v_1 + c_2^2 v_2)$$

may not be a multiple of  $u = c_1 v_1 + c_2 v_2$ . This indicates that the issue of repeated eigenvalues does not have the same status as in the matrix case. Second, even if all the eigenvalues are distinct, it turns out that the  $v_i$ 's are not the only eigenvectors. For example, set  $u := (1/\lambda_1)v_1 + (1/\lambda_2)v_2$ . Then,

$$T(I, u, u) = \lambda_1(1/\lambda_1)^2 v_1 + \lambda_2(1/\lambda_2)^2 v_2 = u,$$

so  $u/\|u\|$  is an eigenvector with corresponding eigenvalue  $\|u\|$ . More generally, for any subset  $S \subseteq [k]$ , the vector

$$\sum_{j \in S} \frac{1}{\lambda_j} \cdot v_j$$

is an eigenvector after normalization.

As we now see, these additional eigenvectors can be viewed as spurious. We say a unit vector  $u$  is a *robust eigenvector* of  $T$  if there exists an  $\epsilon > 0$  such that for all  $\theta \in \{u' \in \mathbb{R}^d : \|u' - u\| \leq \epsilon\}$ , repeated iteration of the map

$$\bar{\theta} \mapsto \frac{T(I, \bar{\theta}, \bar{\theta})}{\|T(I, \bar{\theta}, \bar{\theta})\|} , \tag{3.16}$$

starting from  $\theta$  converges to  $u$ . Note that the map (3.16) re-scales the output to have unit Euclidean norm. Robust eigenvectors are also called attracting fixed points of (3.16); see, e.g., Kolda and Mayo, 2011.

---

<sup>1</sup>Note that there are many definitions of tensor eigenvalues and eigenvectors, see for example Qi, 2005. The definition we used here is called Z-eigenvalues/Z-eigenvectors in Qi, 2005.

The following theorem implies that if  $T$  has an orthogonal decomposition as given in (3.14), then the set of robust eigenvectors of  $T$  are precisely the set  $\{v_1, v_2, \dots, v_k\}$ , implying that the orthogonal decomposition is unique. For even order tensors, the uniqueness is true up to sign-flips of the  $v_i$ 's.

**Theorem 3.2** (Uniqueness of orthogonal tensor decomposition). Let  $T$  have an orthogonal decomposition as given in (3.14). Then,

1. The set of  $\theta \in \mathbb{R}^d$  which do not converge to some  $v_i$  under repeated iteration of (3.16) has measure zero.
2. The set of robust eigenvectors of  $T$  is equal to  $\{v_1, v_2, \dots, v_k\}$ .

See Anandkumar *et al.* (2014a) for the proof of the theorem which follows readily from simple orthogonality considerations. Note that every  $v_i$  in the orthogonal tensor decomposition is *robust*, whereas for a symmetric matrix  $M$ , for almost all initial points, the map  $\bar{\theta} \mapsto \frac{M\bar{\theta}}{\|M\bar{\theta}\|}$  converges only to an eigenvector corresponding to the largest magnitude eigenvalue. Also, since the tensor order is odd, the signs of the robust eigenvectors are fixed, as each  $-v_i$  is mapped to  $v_i$  under (3.16).

### Variational characterization

We now discuss a variational characterization of the orthogonal decomposition. The *generalized Rayleigh quotient* (Zhang and Golub, 2001) for a third-order tensor is given by

$$u \mapsto \frac{T(u, u, u)}{(u^\top u)^{3/2}},$$

which can be compared to (3.13). For an orthogonally decomposable tensor, the following theorem shows that a non-zero vector  $u \in \mathbb{R}^d$  is an *isolated local maximizer* (Nocedal and Wright, 1999) of the generalized Rayleigh quotient if and only if  $u = v_i$  for some  $i \in [k]$ .

**Theorem 3.3.** Assume  $d \geq 2$ . Let  $T$  have an orthogonal decomposition as given in (3.14), and consider the optimization problem

$$\max_{u \in \mathbb{R}^d} T(u, u, u) \quad \text{s.t.} \quad \|u\| = 1.$$

1. The stationary points are eigenvectors of  $T$ .
2. A stationary point  $u$  is an isolated local maximizer if and only if  $u = v_i$  for some  $i \in [k]$ .

See Ge *et al.* (2015a)[Section C.1] for the proof of the theorem. It is similar to local optimality analysis for ICA methods using fourth-order cumulants (Delfosse and Loubaton, 1995; Frieze *et al.*, 1996).

Again, we see similar distinctions to the matrix case. In the matrix case, the only local maximizers of the Rayleigh quotient are the eigenvectors with the largest eigenvalue (and these maximizers take on the globally optimal value). For the case of orthogonal tensor forms, the robust eigenvectors are precisely the isolated local maximizers.

An important implication of the two characterizations is that, for orthogonally decomposable tensors  $T = \sum_{j \in [k]} \lambda_j v_j^{\otimes 3}$ , (i) the local maximizers of the objective function  $T(u, u, u)/(u^\top u)^{3/2}$  correspond precisely to the vectors  $v_j$  in the decomposition, and (ii) these local maximizers can be reliably identified using a simple fixed-point iteration as in (3.16), i.e., the tensor analogue of the matrix power method. Moreover, a second-derivative test based on  $T(I, I, u)$  can be employed to test for local optimality and rule out other stationary points.

### 3.4.3 Beyond Orthogonal Tensor Decomposition

So far, we have considered tensors with orthogonal decomposition as in (3.14). We now discuss how the problem of non-orthogonal tensor decomposition can be reduced to the orthogonal tensor decomposition, and therefore, we can use the orthogonal tensor decomposition algorithms to recover the rank-1 components.

As we alluded in Section 2.4, we can pre-process the tensor using a whitening procedure, which is described in more detail in Procedure 2. This procedure orthogonalizes the components of the input tensor. After recovering the rank-1 components of the orthogonal decomposition, we apply *un-whitening* procedure proposed in Procedure 3 to recover the rank-1 components of the original non-orthogonal tensor decomposition. The whitening procedure only works when the components of the original non-orthogonal tensor are linearly independent. Luckily for

---

**Procedure 2** Whitening

---

**input** Tensor  $T = \sum_{j \in [k]} \lambda_j \cdot a_j^{\otimes 3} \in \mathbb{R}^{d \times d \times d}$ ; matrix  $M = \sum_{j \in [k]} \tilde{\lambda}_j \cdot a_j^{\otimes 2} \in \mathbb{R}^{d \times d}$ . Assume  $a_j$ 's are linearly independent.

**output** Lower dimensional tensor with orthogonal rank-1 components

- 1: Compute the rank- $k$  SVD,  $M = U \text{Diag}(\gamma) U^\top$ , where  $U \in \mathbb{R}^{d \times k}$  and  $\gamma \in \mathbb{R}^k$ .
  - 2: Compute the whitening matrix  $W := U \text{Diag}(\gamma^{-1/2}) \in \mathbb{R}^{d \times k}$ .
  - 3: **return**  $T(W, W, W) \in \mathbb{R}^{k \times k \times k}$ .
- 

many machine learning applications (such as topic models, mixtures of high dimensional Gaussians) that we will talk about in Section 4, it is natural to expect the true components to be linearly independent.

We first elaborate on the whitening step, and analyze how the proposed Procedure 2 works and orthogonalizes the components of input tensor. We then analyze the inversion of whitening operator showing how the components in the whitened space are translated back to the original space as stated in Procedure 3.

**Whitening procedure**

Consider the non-orthogonal rank- $k$  tensor

$$T = \sum_{j \in [k]} \lambda_j \cdot a_j^{\otimes 3}, \quad (3.17)$$

where the goal of whitening procedure is to reduce it to an orthogonal tensor form. To do this, we exploit a matrix  $M$  which has the same rank-1 components as  $T$  such that

$$M = \sum_{j \in [k]} \tilde{\lambda}_j \cdot a_j^{\otimes 2}. \quad (3.18)$$

In case we do not have such matrix, we can generate it as random combination of slices of  $T$  such that  $M := T(I, I, \theta) \in \mathbb{R}^{d \times d}$ , where  $\theta \sim \mathcal{N}(0, I_d)$  is a random standard Gaussian vector. It is also worth mentioning that although we refer to the rank- $k$  SVD of matrix  $M$  as  $U \text{Diag}(\gamma) U^\top$ , it might be the case that matrix  $M$  is not positive semi-definite and does not necessarily have such symmetric SVD. In that case,



$U \text{Diag}(\gamma)U^\top$  is basically the eigen-decomposition of symmetric matrix  $M$  where the entries of vector  $\gamma$  can be also negative. We can modify the whitening matrix as  $W := U \text{Diag}(|\gamma|^{-1/2})$ , where  $|\cdot|$  denotes the entry-wise absolute value, and the rest of analysis in this section would go through with minor modifications. So, in the rest of this section, we assume the entries of  $\gamma$  are all positive.

Another complication is that given the tensor  $T$  in (3.17) and  $M$  in (3.18), it is still impossible to uniquely determine  $\|a_j\|$ ,  $\tilde{\lambda}_j$  and  $\lambda_j$ . Indeed, if we scale the  $j$ -th component to  $C \cdot a_j$  using a constant  $C \neq 0$ , one just needs to scale  $\tilde{\lambda}_j$  by a factor of  $1/C^2$  and  $\lambda_j$  by a factor of  $1/C^3$  and both the tensor  $T$  and matrix  $M$  are preserved. We discuss this ambiguity in more details in Remark 3.1.

Let matrix  $W \in \mathbb{R}^{d \times k}$  denote the whitening matrix, i.e., the whitening matrix  $W$  in Procedure 2 is constructed such that  $W^\top M W = I_k$ . Applying whitening matrix  $W$  to the tensor  $T = \sum_{j \in [k]} \lambda_j \cdot a_j^{\otimes 3}$ , we have

$$\begin{aligned} T(W, W, W) &= \sum_{j \in [k]} \lambda_j \left( W^\top a_j \right)^{\otimes 3} \\ &= \sum_{j \in [k]} \frac{\lambda_j}{\tilde{\lambda}_j^{3/2}} \left( W^\top a_j \sqrt{\tilde{\lambda}_j} \right)^{\otimes 3} \\ &= \sum_{j \in [k]} \mu_j \cdot v_j^{\otimes 3}, \end{aligned} \tag{3.19}$$

where we defined

$$\mu_j := \frac{\lambda_j}{\tilde{\lambda}_j^{3/2}}, \quad v_j := W^\top a_j \sqrt{\tilde{\lambda}_j}, \quad j \in [k], \tag{3.20}$$

in the last equality. Let  $V := [v_1 | v_2 | \dots | v_k] \in \mathbb{R}^{k \times k}$  denote the factor matrix for  $T(W, W, W)$ . Then, we have

$$V := W^\top A \text{Diag}(\tilde{\lambda}^{1/2}), \tag{3.21}$$

and thus,

$$V V^\top = W^\top A \text{Diag}(\tilde{\lambda}) A^\top W = W^\top M W = I_k.$$

Since  $V$  is a square matrix, it is also concluded that  $V^\top V = I_k$ , and therefore, tensor  $T(W, W, W)$  is whitened such that its rank-1 compo-

---

**Procedure 3** Un-whitening

---

**input** Orthogonal rank-1 components  $v_j \in \mathbb{R}^k, j \in [k]$ .

**output** Rank-1 components of the original non-orthogonal tensor

- 1: Consider matrix  $M$  which was exploited for whitening in Procedure 2, and let  $\tilde{\lambda}_j, j \in [k]$  denote the corresponding coefficients as  $M = A \text{Diag}(\tilde{\lambda})A^\top$ ; see (3.18). Note that we don't know  $A$  so we need to get  $\tilde{\lambda}_j$  from other information, see Remark 3.1.
- 2: Compute the rank- $k$  SVD,  $M = U \text{Diag}(\gamma)U^\top$ , where  $U \in \mathbb{R}^{d \times k}$  and  $\gamma \in \mathbb{R}^k$ .
- 3: Compute

$$a_j = \frac{1}{\sqrt{\tilde{\lambda}_j}} U \text{Diag}(\gamma^{1/2}) v_j, \quad j \in [k].$$

- 4: **return**  $\{a_j\}_{j \in [k]}$ .
- 

nents  $v_j$ 's form an orthonormal basis. This discussion clarifies how the whitening procedure works.

**Inversion of the whitening procedure**

Let us also analyze the inversion procedure on how to transform  $v_j$ 's to  $a_j$ 's. The main step is stated in Procedure 3. According to whitening Procedure 2, let  $M = U \text{Diag}(\gamma)U^\top$ ,  $U \in \mathbb{R}^{d \times k}$ ,  $\gamma \in \mathbb{R}^k$ , denote the rank- $k$  SVD of  $M$ . Substituting whitening matrix  $W := U \text{Diag}(\gamma^{-1/2})$  in (3.21), and multiplying  $U \text{Diag}(\gamma^{1/2})$  from left, we have

$$U \text{Diag}(\gamma^{1/2})V = UU^\top A \text{Diag}(\tilde{\lambda}^{1/2}).$$

Since the column spans of  $A \in \mathbb{R}^{d \times k}$  and  $U \in \mathbb{R}^{d \times k}$  are the same (given their relations to  $M$ ),  $A$  is a fixed point for the projection operator on the subspace spanned by the columns of  $U$ . This projector operator is  $UU^\top$  (since columns of  $U$  form an orthonormal basis), and therefore,  $UU^\top A = A$ . Applying this to the above equation, we have

$$A = U \text{Diag}(\gamma^{1/2})V \text{Diag}(\tilde{\lambda}^{-1/2}),$$

i.e.,

$$a_j = \frac{1}{\sqrt{\tilde{\lambda}_j}} U \text{Diag}(\gamma^{1/2}) v_j, \quad j \in [k]. \quad (3.22)$$

As we discussed before, in general one needs some additional information to determine the coefficients  $\tilde{\lambda}_j$ 's.

**Remark 3.1.** [Scaling Ambiguity in Whitening Procedure] If one only has access to tensor  $T$  in (3.17) and matrix  $M$  in (3.18), there is no way to uniquely determine  $\|a_j\|$ ,  $\tilde{\lambda}_j$  or  $\lambda_j$ . Between these three parameters, we already have two equations for any  $j \in [k]$ : 1)  $\mu_j = \lambda_j / \tilde{\lambda}_j^{3/2}$  as in (3.20), and 2)  $a_j = \frac{1}{\sqrt{\tilde{\lambda}_j}} U \text{Diag}(\gamma^{1/2}) v_j$  as in (3.22). Note that all other variables such as  $\mu_j, v_j, \gamma, U$  can be computed from the tensor decomposition of the whitened tensor in (3.19) and the SVD decomposition of matrix  $M$ . Therefore, the three parameters still have one degree of freedom which is captured by a scaling such that if  $(a_j, \tilde{\lambda}_j, \lambda_j)$  is a set of parameters that is consistent with  $M$  and  $T$ , then for any  $C \neq 0$ ,  $(Ca_j, \tilde{\lambda}_j/C^2, \lambda_j/C^3)$  is also a set of parameters that is consistent with  $M$  and  $T$ .

There are many cases where one might have additional information to determine the exact values of  $\|a_j\|$ ,  $\tilde{\lambda}_j$  or  $\lambda_j$ . For some applications (such as topic modeling in Section 4.2.3), the components  $a_j$ 's may have unit  $\ell_1$  or  $\ell_2$  norm, in which case we should scale  $a_j$  accordingly. For some other applications such as pure topic model in Section 4.1 or mixture of Gaussians in Section 4.2.1, we know  $\lambda_j = \tilde{\lambda}_j$ , and therefore, both of them are equal to  $\mu_j^{-2}$ .

When  $\tilde{\lambda}_j$ 's are unknown at the time of running Procedure 3, one can simply choose  $\tilde{\lambda}_j = 1$ . If there is no additional information the results will give one set of parameters that are consistent with  $M$  and  $T$ . If additional information is available one can apply correct normalization afterwards.

### 3.4.4 Beyond Symmetric Tensor Decomposition

In the previous sections, we considered symmetric tensor decompositions as in (3.17). In some applications, the tensor we have access to might

be asymmetric. Consider

$$T = \sum_{j \in [k]} \lambda_j a_j \otimes b_j \otimes c_j, \quad (3.23)$$

where  $\{a_j\}, \{b_j\}, \{c_j\}$ 's are three groups of vectors that are linearly independent within the group. Here, we cannot directly apply the techniques for symmetric tensor decomposition. However, there is a way to transform this tensor to a symmetric one if we have access to some extra matrices. We discuss this process in this section which is a combination of whitening approach proposed in the previous section and the idea of CCA for matrices stated in Section 2.5. Similar to the *whitening* procedure described earlier, the symmetrization step in this section only works if the tensor components ( $\{a_j, j \in [k]\}, \{b_j, j \in [k]\}, \{c_j, j \in [k]\}$ ) are all linearly independent within their own mode. Again for many machine learning applications that requires this procedure (such as the Multi-view model and Noisy-Or networks in Section 4), it is natural to assume that the components are indeed linearly independent.

We first elaborate on the symmetrization step, and discuss how the proposed Procedure 4 works by orthogonalizing and symmetrizing the components of the input tensor. We then analyze the inversion of this process showing how the components in the whitened/symmetrized space are translated back to the original space as stated in Procedure 5.

### Symmetrization procedure

The whitening and symmetrization in Procedure 4 is adapted from whitening procedure for symmetric tensors stated in Procedure 2 with two modifications: first, the whitening is performed for an asymmetric tensor vs. a symmetric tensor in Procedure 2, and second, an extra step for symmetrization of the tensor is added. Similar to the whitening procedure, there are also additional scaling issues (as in Remark 3.1) introduced by the symmetrization procedure, we discuss that later in Remark 3.2.

In order to transform the asymmetric tensor  $T$  in (3.23) to a symmetric and orthogonal tensor, we first whiten the three modes of the tensor. Similar to the whitening argument in the previous section, let

**Procedure 4** Whitening and Symmetrization

**input** Tensor  $T = \sum_{j \in [k]} \lambda_j \cdot a_j \otimes b_j \otimes c_j \in \mathbb{R}^{d_1 \times d_2 \times d_3}$

**input** matrix  $M_a = \sum_{j \in [k]} \tilde{\lambda}_{a,j} \cdot a_j^{\otimes 2}$  (similarly  $M_b$  and  $M_c$ )

**input** matrix  $M_{a,b} = \sum_{j \in [k]} \tilde{\lambda}_{ab,j} \cdot a_j \otimes b_j$  (similarly  $M_{a,c}$ )

{Note that only  $T, M_a, M_b, M_c, M_{a,b}, M_{b,c}, M_{a,c}$  are known (usually from moment estimates), we don't observe the components  $a_j, b_j, c_j$ .}

**output** Lower dimensional symmetric tensor with orthogonal rank-1 components

- 1: Compute the rank-k SVD,  $M_a = U_a \text{Diag}(\gamma_a) U_a^\top$ , where  $U_a \in \mathbb{R}^{d_1 \times k}$  and  $\gamma \in \mathbb{R}^k$ ; and similarly for  $M_b$  and  $M_c$ .
- 2: Compute the whitening matrix  $W_a := U_a \text{Diag}(\gamma_a^{-1/2}) \in \mathbb{R}^{d_1 \times k}$ ; and similarly  $W_b$  and  $W_c$ .
- 3: Compute matrices  $R_{a,b} := W_a^\top M_{a,b} W_b$ , and  $R_{a,c} := W_a^\top M_{a,c} W_c$ .
- 4: **return**  $T(W_a, W_b R_{a,b}^\top, W_c R_{a,c}^\top) \in \mathbb{R}^{k \times k \times k}$ .

$W_a, W_b, W_c$  be the whitening matrices for different modes of the tensor; see Procedure 4 for precise definitions. Following the same calculations as in the whitening section, we have

$$T(W_a, W_b, W_c) = \sum_{j \in [k]} \hat{\mu}_j \hat{a}_j \otimes \hat{b}_j \otimes \hat{c}_j,$$

where

$$\hat{\mu}_j := \frac{\lambda_j}{\sqrt{\tilde{\lambda}_{a,j} \tilde{\lambda}_{b,j} \tilde{\lambda}_{c,j}}}, \quad \hat{a}_j := W_a^\top a_j \sqrt{\tilde{\lambda}_{a,j}}, \quad j \in [k].$$

$\hat{b}_j$  and  $\hat{c}_j$  are similarly defined. Same as before, we have transformed the tensor so that each mode now has orthogonal components, but the only difference is  $\hat{a}_j$  may not be the same as  $\hat{b}_j$  (or  $\hat{c}_j$ ), and therefore, the tensor is not symmetric yet. We will resolve this by using the cross matrices  $M_{a,b}, M_{b,c}$ . The idea is very similar to CCA stated in Section 2.5. More precisely we have:

**Claim 1.** Let  $R_{a,b} := W_a^\top M_{a,b} W_b$ , then we have

$$R_{a,b} = \sum_{j \in [k]} \tilde{\mu}_j \cdot \hat{a}_j \hat{b}_j^\top,$$

---

**Procedure 5** Inversion of Whitening and Symmetrization
 

---

**input** Orthogonal rank-1 components  $\hat{a}_j \in \mathbb{R}^k, j \in [k]$

**output** Rank-1 components of the original non-orthogonal and asymmetric tensor

1: For all  $j \in [k]$ , compute

$$\begin{aligned} a_j &= \frac{1}{\sqrt{\tilde{\lambda}_{a,j}}} U_a \text{Diag}(\gamma_a^{1/2}) \hat{a}_j, \\ b_j &= \frac{\sqrt{\tilde{\lambda}_{a,j}}}{\tilde{\lambda}_{ab,j}} U_b \text{Diag}(\gamma_b^{1/2}) R_{a,b}^\top \hat{a}_j, \\ c_j &= \frac{\sqrt{\tilde{\lambda}_{a,j}}}{\tilde{\lambda}_{ac,j}} U_c \text{Diag}(\gamma_c^{1/2}) R_{a,c}^\top \hat{a}_j, \end{aligned}$$

where the variables are the same as in Procedure 4.

2: **return**  $\{(a_j, b_j, c_j)\}_{j \in [k]}$ .

---

where  $\tilde{\mu}_j := \frac{\tilde{\lambda}_{ab,j}}{\sqrt{\tilde{\lambda}_{a,j} \tilde{\lambda}_{b,j}}}$ . In particular,  $R_{a,b} \hat{b}_j = \tilde{\mu}_j \hat{a}_j$ .

The claim follows from similar calculation as above for  $T(W_a, W_b, W_c)$ . Define  $R_{a,c} := W_a^\top M_{a,c} W_c$ , we can then use these matrices to transform between the vectors  $\hat{a}_j, \hat{b}_j$  and  $\hat{c}_j$ . More precisely

$$\begin{aligned} T(W_a, W_b R_{a,b}^\top, W_c R_{a,c}^\top) &= \sum_{j \in [k]} \lambda_j (W_a^\top a_j) \otimes (R_{a,b} W_b^\top b_j) \otimes (R_{a,c} W_c^\top c_j) \\ &= \sum_{j \in [k]} \hat{\mu}_j \cdot \hat{a}_j \otimes (R_{a,b} \hat{b}_j) \otimes (R_{a,c} \hat{c}_j) \\ &= \sum_{j \in [k]} \mu_j \cdot \hat{a}_j^{\otimes 3}, \end{aligned}$$

where  $\mu_j := \frac{\lambda_j \tilde{\lambda}_{ab,j} \tilde{\lambda}_{ac,j}}{\tilde{\lambda}_{a,j}^{3/2} \tilde{\lambda}_{b,j} \tilde{\lambda}_{c,j}}$ . We have now transformed the tensor to a symmetric and orthogonal tensor whose components are  $\{\hat{a}_j\}$ 's, and techniques for symmetric orthogonal tensors can be applied to do the decomposition.

### Inversion of the symmetrization procedure

The inversion steps are provided in Procedure 5. The analysis of the algorithm and why it works is very similar to the inversion of whitening discussed in the previous section. This technique is particularly useful for multi-view models that we will discuss in Section 4.3. As we mentioned before, there are also uncertainties about the scaling in the case of symmetrization:

**Remark 3.2.** [Scaling Ambiguity in Whitening and Symmetrization Procedure] If one only has access to tensor  $T$ , matrices  $M_a, M_b, M_c, M_{a,b}, M_{b,c}, M_{a,c}$ , there is no way to uniquely determine the 10 parameters ( $\|a_j\|, \|b_j\|, \|c_j\|, \tilde{\lambda}_{a,j}, \tilde{\lambda}_{b,j}, \tilde{\lambda}_{c,j}, \tilde{\lambda}_{ab,j}, \tilde{\lambda}_{bc,j}, \tilde{\lambda}_{ac,j}, \lambda_j$ ). The 7 known quantities  $T, M_a, M_b, M_c, M_{a,b}, M_{b,c}, M_{a,c}$  give 7 equations over these 10 parameters. The additional degrees of freedom can be described as  $(C_a\|a_j\|, C_b\|b_j\|, C_c\|c_j\|, \tilde{\lambda}_{a,j}/C_a^2, \tilde{\lambda}_{b,j}/C_b^2, \tilde{\lambda}_{c,j}/C_c^2, \tilde{\lambda}_{ab,j}/(C_aC_b), \tilde{\lambda}_{bc,j}/(C_bC_c), \tilde{\lambda}_{ac,j}/(C_aC_c), \lambda_j/(C_aC_bC_c))$ , where  $C_a, C_b, C_c$  are arbitrary nonzero constants.

As before, there are special cases where the scaling of  $a_j, b_j, c_j$  is known, which leads to three additional equations to uniquely determine all the scalings. There are also special cases where all the coefficients are the same, in which case they are all going to be equal to  $\hat{\mu}_j^{-2}$ .

When  $\tilde{\lambda}_{a,j}$ 's (and similarly,  $\tilde{\lambda}_{ab,j}$ 's and  $\tilde{\lambda}_{ac,j}$ 's) are unknown at the time of running Procedure 5, one can simply choose all of them to be equal to 1. If there is no additional information the results will give one set of parameters that are consistent with all the observed matrices and tensors. If additional information is available one can apply correct normalization afterwards.

### 3.5 Tensor Power Iteration

In the previous section, we discussed that the robust fixed-points of the tensor power iteration in (3.16)

$$\bar{\theta} \mapsto \frac{T(I, \bar{\theta}, \bar{\theta})}{\|T(I, \bar{\theta}, \bar{\theta})\|},$$

correspond to the rank-1 components of orthogonal tensor decomposition in (3.14); see Theorem 3.2. Therefore, the power iteration is a natural

and useful algorithm to recover the rank-1 components of an orthogonal tensor decomposition (Lathauwer *et al.*, 2000, Remark 3). We first state a simple convergence analysis for an orthogonally decomposable tensor  $T$ , and then discuss the analysis for approximately orthogonally decomposable tensors.

When only an approximation  $\hat{T}$  to an orthogonally decomposable tensor  $T$  is available (*e.g.*, when empirical moments are used to estimate population moments), an orthogonal decomposition need not exist for this perturbed tensor (unlike the case for matrices), and a more robust approach is required to extract the approximate decomposition. Here, we propose such a variant in Algorithm 6 and provide a detailed perturbation analysis.

### 3.5.1 Convergence analysis for orthogonally decomposable tensors

The following lemma establishes the quadratic convergence of the tensor power method, *i.e.*, repeated iteration of (3.16), for extracting a single component of the orthogonal decomposition. Note that the initial vector  $\theta_0$  determines which robust eigenvector will be the convergent point. Computation of subsequent eigenvectors can be computed with deflation, *i.e.*, by subtracting appropriate terms from  $T$ .

**Lemma 1.** [Tensor power iteration for orthogonally decomposable tensors] Let  $T \in \mathbb{R}^{d \times d \times d}$  have an orthogonal decomposition as given in (3.14). For a vector  $\theta_0 \in \mathbb{R}^d$ , suppose that the set of numbers  $|\lambda_1 v_1^\top \theta_0|, |\lambda_2 v_2^\top \theta_0|, \dots, |\lambda_k v_k^\top \theta_0|$  has a unique largest element; without loss of generality, say  $|\lambda_1 v_1^\top \theta_0|$  is this largest value and  $|\lambda_2 v_2^\top \theta_0|$  is the second largest value. For  $t = 1, 2, \dots$ , let

$$\theta_t := \frac{T(I, \theta_{t-1}, \theta_{t-1})}{\|T(I, \theta_{t-1}, \theta_{t-1})\|}.$$

Then

$$\|v_1 - \theta_t\|^2 \leq \left(2\lambda_1^2 \sum_{i=2}^k \lambda_i^{-2}\right) \cdot \left| \frac{\lambda_2 v_2^\top \theta_0}{\lambda_1 v_1^\top \theta_0} \right|^{2^{t+1}}.$$

That is, repeated iteration of (3.16) starting from  $\theta_0$  converges to  $v_1$  at a quadratic rate.



To obtain all eigenvectors, we may simply proceed iteratively using deflation, executing the power method on  $T - \sum_{j \in [s]} \lambda_j v_j^{\otimes 3}$  after having obtained robust eigenvector/eigenvalue pairs  $\{(v_j, \lambda_j), j \in [s]\}$ .

*Proof.* Let  $\bar{\theta}_0, \bar{\theta}_1, \bar{\theta}_2, \dots$  be the sequence given by

$$\bar{\theta}_0 := \theta_0, \quad \bar{\theta}_t := T(I, \theta_{t-1}, \theta_{t-1}), t \geq 1.$$

Let  $c_i := v_i^\top \theta_0$  for all  $i \in [k]$ . It is easy to check that

1.  $\theta_t = \bar{\theta}_t / \|\bar{\theta}_t\|$ ,
2.  $\bar{\theta}_t = \sum_{i=1}^k \lambda_i^{2^t-1} c_i^{2^t} v_i$ .

Indeed,

$$\bar{\theta}_{t+1} = \sum_{i=1}^k \lambda_i (v_i^\top \bar{\theta}_t)^2 v_i = \sum_{i=1}^k \lambda_i (\lambda_i^{2^t-1} c_i^{2^t})^2 v_i = \sum_{i=1}^k \lambda_i^{2^{t+1}-1} c_i^{2^{t+1}} v_i.$$

Then

$$\begin{aligned} 1 - (v_1^\top \theta_t)^2 &= 1 - \frac{(v_1^\top \bar{\theta}_t)^2}{\|\bar{\theta}_t\|^2} = 1 - \frac{\lambda_1^{2^{t+1}-2} c_1^{2^{t+1}}}{\sum_{i=1}^k \lambda_i^{2^{t+1}-2} c_i^{2^{t+1}}} \\ &\leq \frac{\sum_{i=2}^k \lambda_i^{2^{t+1}-2} c_i^{2^{t+1}}}{\sum_{i=1}^k \lambda_i^{2^{t+1}-2} c_i^{2^{t+1}}} \\ &\leq \lambda_1^2 \sum_{i=2}^k \lambda_i^{-2} \cdot \left| \frac{\lambda_2 c_2}{\lambda_1 c_1} \right|^{2^{t+1}}. \end{aligned}$$

Since  $\lambda_1 > 0$ , we have  $v_1^\top \theta_t > 0$  and hence,

$$\|v_1 - \theta_t\|^2 = 2(1 - v_1^\top \theta_t) \leq 2(1 - (v_1^\top \theta_t)^2),$$

as required. ■

### 3.5.2 Perturbation analysis of a robust tensor power method

Now we consider the case where we have an approximation  $\hat{T}$  to an orthogonally decomposable tensor  $T$ . Here, a more robust approach is required to extract an approximate decomposition. Anandkumar *et al.*, 2014a gave such an algorithm (Algorithm 6), and provided a detailed

---

**Algorithm 6** Robust Tensor Power Method

---

**input** symmetric tensor  $\tilde{T} \in \mathbb{R}^{k \times k \times k}$ , number of iterations  $L$ ,  $N$ .  
**output** the estimated eigenvector/eigenvalue pair; the deflated tensor.

- 1: **for**  $\tau = 1$  to  $L$  **do**
- 2:   Draw  $\theta_0^{(\tau)}$  uniformly at random from the unit sphere in  $\mathbb{R}^k$ .
- 3:   **for**  $t = 1$  to  $N$  **do**
- 4:     Compute power iteration update

$$\theta_t^{(\tau)} := \frac{\tilde{T}(I, \theta_{t-1}^{(\tau)}, \theta_{t-1}^{(\tau)})}{\|\tilde{T}(I, \theta_{t-1}^{(\tau)}, \theta_{t-1}^{(\tau)})\|} \quad (3.24)$$

- 5:   **end for**
  - 6: **end for**
  - 7: Let  $\tau^* := \arg \max_{\tau \in [L]} \{\tilde{T}(\theta_N^{(\tau)}, \theta_N^{(\tau)}, \theta_N^{(\tau)})\}$ .
  - 8: Do  $N$  power iteration updates (3.24) starting from  $\theta_N^{(\tau^*)}$  to obtain  $\hat{\theta}$ , and set  $\hat{\lambda} := \tilde{T}(\hat{\theta}, \hat{\theta}, \hat{\theta})$ .
  - 9: **return** the estimated eigenvector/eigenvalue pair  $(\hat{\theta}, \hat{\lambda})$ ; the deflated tensor  $\tilde{T} - \hat{\lambda} \hat{\theta}^{\otimes 3}$ .
- 

perturbation analysis. We summarize the perturbation result here and give a generalization later in Section 3.5.3. For simplicity, we assume the tensor  $\hat{T}$  is of size  $k \times k \times k$  as per the reduction from Section 3.4.3 where whitening procedure has been applied to the original tensor. In some applications, it may be preferable to work directly with a  $d \times d \times d$  tensor of rank  $k \leq d$  (as in Lemma 1); these results apply in that setting with little modification.

Assume that the symmetric tensor  $T \in \mathbb{R}^{k \times k \times k}$  is orthogonally decomposable, and that  $\hat{T} = T + E$ , where the perturbation  $E \in \mathbb{R}^{k \times k \times k}$  is a symmetric tensor with small operator norm:

$$\|E\| := \sup_{\|\theta\|=1} |E(\theta, \theta, \theta)|.$$

In our applications that we will describe in Section 4,  $\hat{T}$  is the tensor formed by using empirical moments, while  $T$  is the orthogonally decomposable tensor derived from the population moments for the given model.

The following theorem is similar to Wedin’s perturbation theorem for singular vectors of matrices (Wedin, 1972) in that it bounds the error of the (approximate) decomposition returned by Algorithm 6 on input  $\hat{T}$  in terms of the size of the perturbation, provided that the perturbation is small enough.

**Theorem 3.4** (Anandkumar *et al.*, 2014a). Let  $\hat{T} = T + E \in \mathbb{R}^{k \times k \times k}$ , where  $T$  is a symmetric tensor with orthogonal decomposition  $T = \sum_{i=1}^k \lambda_i v_i^{\otimes 3}$  where each  $\lambda_i > 0$ ,  $\{v_1, v_2, \dots, v_k\}$  is an orthonormal basis, and  $E$  is a symmetric tensor with operator norm  $\|E\| \leq \varepsilon$ . Define  $\lambda_{\min} := \min\{\lambda_i : i \in [k]\}$ , and  $\lambda_{\max} := \max\{\lambda_i : i \in [k]\}$ . There exists universal constants  $C_1, C_2, C_3 > 0$  such that the following holds. Pick any  $\eta \in (0, 1)$ , and suppose

$$\varepsilon \leq C_1 \cdot \frac{\lambda_{\min}}{k}, \quad N \geq C_2 \cdot \left( \log(k) + \log \log \left( \frac{\lambda_{\max}}{\varepsilon} \right) \right),$$

and

$$\begin{aligned} \sqrt{\frac{\ln(L/\log_2(k/\eta))}{\ln(k)}} \cdot \left( 1 - \frac{\ln(\ln(L/\log_2(k/\eta))) + C_3}{4 \ln(L/\log_2(k/\eta))} - \sqrt{\frac{\ln(8)}{\ln(L/\log_2(k/\eta))}} \right) \\ \geq 1.02 \left( 1 + \sqrt{\frac{\ln(4)}{\ln(k)}} \right). \end{aligned}$$

(Note that the condition on  $L$  holds with  $L = \text{poly}(k) \log(1/\eta)$ .) Suppose that Algorithm 6 is iteratively called  $k$  times, where the input tensor is  $\hat{T}$  in the first call, and in each subsequent call, the input tensor is the deflated tensor returned by the previous call. Let  $(\hat{v}_1, \hat{\lambda}_1), (\hat{v}_2, \hat{\lambda}_2), \dots, (\hat{v}_k, \hat{\lambda}_k)$  be the sequence of estimated eigenvector/eigenvalue pairs returned in these  $k$  calls. With probability at least  $1 - \eta$ , there exists a permutation  $\pi$  on  $[k]$  such that

$$\|v_{\pi(j)} - \hat{v}_j\| \leq 8\varepsilon/\lambda_{\pi(j)}, \quad |\lambda_{\pi(j)} - \hat{\lambda}_j| \leq 5\varepsilon, \quad \forall j \in [k],$$

and

$$\left\| T - \sum_{j=1}^k \hat{\lambda}_j \hat{v}_j^{\otimes 3} \right\| \leq 55\varepsilon.$$

One important difference from Wedin’s theorem is that this is an algorithm dependent perturbation analysis, specific to Algorithm 6 (since the perturbed tensor need not have an orthogonal decomposition). Furthermore, note that Algorithm 6 uses multiple restarts to ensure (approximate) convergence—the intuition is that by restarting at multiple points, we eventually start at a point in which the initial contraction towards some eigenvector dominates the error  $E$  in our tensor. The proof shows that we find such a point with high probability within  $L = \text{poly}(k)$  trials. It should be noted that for large  $k$ , the required bound on  $L$  is very close to linear in  $k$ .

A final consideration is that for specific applications, it may be possible to use domain knowledge to choose better initialization points. For instance, in the topic modeling applications (*cf.* Section 4.1), the eigenvectors are related to the topic word distributions, and many documents may be primarily composed of words from just single topic. Therefore, good initialization points can be derived from these single-topic documents themselves, as these points would already be close to one of the eigenvectors.

### 3.5.3 Perturbation analysis of tensor power method with whitening

A limitation of Theorem 3.4 is that it only applies to orthogonal decompositions, while in most applications one would need to first apply the whitening transformation in Procedure 2 described in Section 3.4.3. With matrix perturbation bounds, it is possible to analyze the robustness of the combined procedure of whitening and orthogonal tensor decomposition. Variants of such analysis has appeared before in several papers, such as Anandkumar *et al.*, 2013; Janzamin *et al.*, 2014, however they are specialized to the specific setting. In this subsection we will give guarantees for such a combined procedure in the general setting.

**Theorem 3.5.** Suppose the true matrix  $M$  and tensor  $T$  have the forms

$$M = \sum_{i=1}^k \tilde{\lambda}_i a_i a_i^\top, \quad T = \sum_{i=1}^k \lambda_i a_i^{\otimes 3},$$

where  $\{a_1, a_2, \dots, a_k\}$  is not necessarily a set of orthogonal components.

Assume our algorithm only has access to noisy/perturbed versions

$$\hat{M} = M + E_M, \hat{T} = T + E_T, \quad \text{where } \|E_M\| \leq \epsilon_M, \|E_T\| \leq \epsilon_T.$$

Let  $\sigma_{\min}(M)$  be the smallest nonzero singular value of  $M$ . Suppose  $\epsilon_M \leq \sigma_{\min}(M)/4$ , let  $\Lambda_{\min} := \min\{\lambda_i \tilde{\lambda}_i^{-3/2} : i \in [k]\}$ ,  $\Lambda_{\max} := \max\{\lambda_i \tilde{\lambda}_i^{-3/2} : i \in [k]\}$ , then there exists a universal constant  $C$  such that

$$\epsilon_{TW} := C \left( \frac{\epsilon_T}{\sigma_{\min}(M)^{3/2}} + \Lambda_{\max} \frac{\epsilon_M}{\sigma_{\min}(M)} \right).$$

If  $\epsilon_{TW}$  (as  $\epsilon$ ),  $\Lambda_{\max}$  (as  $\lambda_{\max}$ ),  $\Lambda_{\min}$  (as  $\lambda_{\min}$ ),  $N, L, \eta$  satisfies the conditions in Theorem 3.4, then Algorithm 6 combined with whitening/unwhitening Procedures 2 and 3 finds pairs  $(\hat{a}_1, \hat{\Lambda}_1), (\hat{a}_2, \hat{\Lambda}_2), \dots, (\hat{a}_k, \hat{\Lambda}_k)$ , such that with probability  $\eta$  there exists a permutation  $\pi$  on  $[k]$  such that for all  $j \in [k]$ ,

$$\begin{aligned} \|\sqrt{\tilde{\lambda}_{\pi(j)}} a_{\pi(j)} - \hat{a}_j\| &\leq \frac{9\tilde{\lambda}_{\pi(j)}^{3/2} \|M\|^{1/2}}{\lambda_{\pi(j)}} \epsilon_{TW}, \\ |\lambda_{\pi(j)} \tilde{\lambda}_{\pi(j)}^{-3/2} - \hat{\Lambda}_j| &\leq 5\epsilon_{TW}. \end{aligned}$$

Note that as we discussed earlier in Remark 3.1, without additional assumptions it is impossible to determine the scaling of  $a_i$  together with  $\lambda_i$  and  $\tilde{\lambda}_i$ . The two quantities that we give perturbation bounds on ( $\lambda_{\pi(j)} \tilde{\lambda}_{\pi(j)}^{-3/2}$  and  $\sqrt{\tilde{\lambda}_{\pi(j)}} a_{\pi(j)}$ ) are the two quantities that are not effected by the scaling issue. In the special case when  $\lambda_i = \tilde{\lambda}_i$ , the pair  $(\hat{a}_j, \hat{\Lambda}_j)$  that we estimate allows us to estimate  $a_{\pi(j)} \approx \hat{\Lambda}_j \hat{a}_j$  and  $\lambda_{\pi(j)} \approx \hat{\Lambda}_j^{-2}$ .

From Theorem 3.5, it is also clear that the error comes from both the whitening process ( $\Lambda_{\max} \epsilon_M / \sigma_{\min}(M)$ ) and the estimation error in estimating the tensor ( $\epsilon_T / \sigma_{\min}(M)^{3/2}$ ). If the second order moment estimate  $M$  is not accurate enough, using this algorithm can suffer additional error. Empirically, it is often observed that alternating least squares (see Section 3.7) may perform better than using orthogonal tensor decomposition with whitening. However, we do want to emphasize that alternating least squares does not have the same provable guarantee as Theorem 3.5.

To prove Theorem 3.5, we first need to analyze the perturbation of the whitening matrix. We use Weyl's Theorem and Davis-Kahan

Theorem to do that. We state special cases of these two theorems for the setting that we are interested in. For more general forms of these theorems and other matrix perturbation inequalities, see Stewart and Sun, 1990.

**Theorem 3.6** (Weyl's Theorem (Weyl, 1912)). Let  $M \in \mathbb{R}^{d \times d}$  be a symmetric matrix, and  $E \in \mathbb{R}^{d \times d}$  be a symmetric perturbation with  $\|E\| \leq \epsilon$ . Let  $\lambda_i(M)$  be the  $i$ -th eigenvalue of  $M$ . We have

$$|\lambda_i(M) - \lambda_i(M + E)| \leq \epsilon, \quad i \in [d].$$

**Theorem 3.7** (Davis-Kahan Theorem (Davis and Kahan, 1970)). Let  $M \in \mathbb{R}^{d \times d}$  and  $\hat{M} = M + E \in \mathbb{R}^{d \times d}$  be symmetric PSD matrices with  $\|E\| \leq \epsilon$ . Suppose  $M$  is rank  $k$  and its truncated SVD is  $M = UDU^\top$ , where  $U \in \mathbb{R}^{d \times k}$  and  $D \in \mathbb{R}^{k \times k}$ . The truncated (top- $k$ ) SVD of  $\hat{M}$  is  $\hat{U}\hat{D}\hat{U}^\top$ . Let  $U^\perp \in \mathbb{R}^{d \times (d-k)}$  be the orthogonal subspace of  $U$  (that is,  $UU^\top + (U^\perp)(U^\perp)^\top = I_d$ ). Then, we have

$$\|(U^\perp)^\top \hat{U}\| \leq \|E\|/\lambda_k(\hat{M}),$$

where  $\lambda_k(\hat{M})$  denotes the  $k$ -th eigenvalue of  $\hat{M}$ .

Using these two theorems, we will prove the following guarantees for the whitening procedure.

**Lemma 3.8.** Suppose  $M \in \mathbb{R}^{d \times d}$  is a symmetric PSD matrix with rank  $k$  and  $\sigma_{\min}(M)$  denotes its smallest (nonzero) singular value. Let  $\hat{M} = M + E$  is also a symmetric matrix and  $\epsilon := \|E\| \leq \sigma_{\min}(M)/4$ . Let the truncated (top- $k$ ) SVD of  $M$  and  $\hat{M}$  be  $UDU^\top$  and  $\hat{U}\hat{D}\hat{U}^\top$ , respectively. Then, there exists an orthonormal matrix  $R \in \mathbb{R}^{k \times k}$  such that if we define  $W := UD^{-1/2}R$ ,  $\hat{W} := \hat{U}\hat{D}^{-1/2}$ ,  $B := UD^{1/2}R$ ,  $\hat{B} = \hat{U}\hat{D}^{1/2}$ , these matrices satisfy

$$\begin{aligned} \|W - \hat{W}\| &\leq \frac{5\epsilon}{\sigma_{\min}(M)^{3/2}}, \\ \|B^\top(W - \hat{W})\| &\leq \frac{3\epsilon}{\sigma_{\min}(M)}, \\ \|B - \hat{B}\| &\leq \frac{3\epsilon\sqrt{\|M\|}}{\sigma_{\min}(M)}. \end{aligned}$$

*Proof.* We first show that  $U$  and  $\hat{U}$  span similar subspace. Let  $U^\perp$  be the orthonormal subspace of  $U$  (as in Theorem 3.7). By Weyl's Theorem (Theorem 3.6), we know

$$\lambda_k(\hat{M}) \geq \sigma_{\min}(M) - \|E\| \geq 3\sigma_{\min}(M)/4. \quad (3.25)$$

Therefore, by Davis Kahan Theorem (Theorem 3.7) we have

$$\|(U^\perp)^\top \hat{U}\| \leq 4\epsilon/3\sigma_{\min}(M). \quad (3.26)$$

Now for  $W - \hat{W}$ , we have

$$\begin{aligned} \|W - \hat{W}\| &= \|(UU^\top + (U^\perp)(U^\perp)^\top)(W - \hat{W})\| \\ &\leq \|U^\top(W - \hat{W})\| + \|(U^\perp)^\top(W - \hat{W})\|. \end{aligned} \quad (3.27)$$

No matter what  $R$  is, the second term can be bounded as

$$\begin{aligned} \|(U^\perp)^\top(W - \hat{W})\| &= \|(U^\perp)^\top \hat{U} \hat{D}^{-1/2}\| \\ &\leq \lambda_k(\hat{M})^{-1/2} \|(U^\perp)^\top \hat{U}\| \\ &\leq \frac{2\epsilon}{\sigma_{\min}(M)^{3/2}}, \end{aligned} \quad (3.28)$$

where we used (3.25) and (3.26) in the last inequality<sup>2</sup>. Therefore, we only need to show that there exists an  $R$  such that the first term  $\|U^\top(W - \hat{W})\|$  is small.

Let  $\bar{M} = \hat{U} \hat{D} \hat{U}^\top$ , by Eckart-Young Theorem (Theorem 2.3), we know  $\|\bar{M} - \hat{M}\| \leq \|E\| = \epsilon$ , and thus,  $\|\bar{M} - M\| \leq 2\epsilon$ . Now for  $\hat{W}^\top(\bar{M} - M)\hat{W}$ , we have

$$\|\hat{W}^\top(\bar{M} - M)\hat{W}\| \leq \frac{2\epsilon}{\sigma_k(\hat{M})},$$

where we also used the fact that  $\|\hat{W}\| = \sigma_k(\hat{M})^{-1/2}$ . Given  $\hat{W}^\top \bar{M} \hat{W} = I_k$ , the above inequality can be rewritten as

$$\|I - \hat{W}^\top M \hat{W}\| \leq \frac{2\epsilon}{\sigma_k(\hat{M})} \leq \frac{3\epsilon}{\sigma_{\min}(M)}.$$

---

<sup>2</sup>Note that the exact constant in the last inequality is  $(\frac{4}{3})^{3/2}$ , and we replace it by 2 for simplicity. We will do similar relaxations to constants several times more later in the proof.

Let  $P := U^\top \hat{W}$ , then  $\hat{W}^\top M \hat{W} = P^\top D P$ . Since  $3\epsilon/\sigma_{\min}(M) \leq 3/4 < 1$ , by Weyl's Theorem we know the eigenvalues of  $P^\top D P$  are between  $1 \pm 3\epsilon/\sigma_{\min}(M) \in [1/4, 7/4]$ . There exists a diagonal matrix  $\Delta$  (with  $\|\Delta - I\| \leq 3\epsilon/\sigma_{\min}(M)$ ) and an orthonormal matrix  $R_1$  such that

$$P^\top D P = R_1 \Delta R_1^\top.$$

In other words, let  $R_2 = D^{1/2} P R_1 \Delta^{-1/2}$  (equivalently,  $P = D^{-1/2} R_2 \Delta^{1/2} R_1^\top$ ), we have  $R_2^\top R_2 = \Delta^{-1/2} R_1^\top (P^\top D P) R_1 D^{-1/2} = I$ , so  $R_2$  is also orthonormal.

Now we can choose  $R = R_2 R_1^\top$ , and therefore, the first term in (3.27) can be bounded as

$$\begin{aligned} \|U^\top (W - \hat{W})\| &= \|D^{-1/2} R - P\| \\ &= \|D^{-1/2} R_2 R_1^\top - D^{-1/2} R_2 \Delta^{1/2} R_1^\top\| \\ &= \|D^{-1/2} R_2 (I - \Delta^{1/2}) R_1^\top\| \\ &\leq \|D^{-1/2}\| \|I - \Delta^{1/2}\| \\ &\leq \frac{3\epsilon}{\sigma_{\min}(M)^{3/2}}. \end{aligned}$$

The last step uses the fact that  $\|I - \Delta^{1/2}\| \leq \|I - \Delta\|$  for diagonal matrix  $\Delta$ , which just follows from  $|1 - \sqrt{x}| \leq |1 - x|$  for every  $x \geq 0$ . Combining this bound with (3.27) and (3.28), we prove the first desired inequality as

$$\|W - \hat{W}\| \leq \frac{5\epsilon}{\sigma_{\min}(M)^{3/2}}.$$

With the choice of  $R$ , the second inequality is easier to prove:

$$\begin{aligned} \|B^\top (W - \hat{W})\| &= \|D^{1/2} U^\top (W - \hat{W})\| \\ &= \|R - D^{1/2} P\| \\ &= \|R_2 R_1^\top - R_2 \Delta^{1/2} R_1^\top\| \\ &\leq \|I - \Delta^{1/2}\| \\ &\leq \frac{3\epsilon}{\sigma_{\min}(M)}. \end{aligned}$$

To prove the third equation, we observe that  $B = M W$  and  $\hat{B} =$



$\bar{M}\hat{W}$ . Therefore,

$$\begin{aligned} \|B - \hat{B}\| &= \|MW - \bar{M}\hat{W}\| \\ &\leq \|M(W - \hat{W})\| + \|(M - \bar{M})\hat{W}\|. \end{aligned}$$

Here, the second term is bounded as

$$\|(M - \bar{M})\hat{W}\| \leq \|M - \bar{M}\| \|\hat{W}\| \leq \frac{3\epsilon}{\sigma_{\min}(M)^{1/2}},$$

where in the second inequality we used (3.25), the fact that  $\|M - \bar{M}\| \leq 2\epsilon$  and an upper bound on the constant term. For the first term, it can be bounded as

$$\begin{aligned} \|M(W - \hat{W})\| &= \|UD^{1/2}R - UDP\| \\ &= \|UD^{1/2}(R - D^{1/2}P)\| \\ &\leq \|UD^{1/2}\| \|R - D^{1/2}P\| \\ &\leq \frac{3\epsilon\sqrt{\|M\|}}{\sigma_{\min}(M)}. \end{aligned}$$

Note that the bound for  $\|R - D^{1/2}P\|$  is the same as the second inequality.  $\blacksquare$

Finally, we are ready to prove Theorem 3.5.

*Proof of Theorem 3.5.* We first construct the whitening matrices  $W$  and  $\hat{W}$ , and un-whitening matrices  $B$  and  $\hat{B}$  for the exact matrix  $M$  and observed matrix  $\hat{M}$  as described in Lemma 3.8. The ideal tensor that we want to perform orthogonal tensor decomposition is  $T(W, W, W)$ , however we only have access to  $\hat{T}(\hat{W}, \hat{W}, \hat{W})$ . Therefore the main part of the proof is to bound the difference between these two tensors.

Let  $v_i = \sqrt{\tilde{\lambda}_i} W^\top a_i$ . As we argued in Section 3.4.3,  $v_i$ 's are orthonormal vectors and we have (see (3.19) and (3.20))

$$T(W, W, W) = \sum_{i=1}^k \Lambda_i v_i^{\otimes 3},$$

where we defined  $\Lambda_i := \lambda_i \tilde{\lambda}_i^{-3/2}$ . Since  $T(W, W, W)$  is an orthogonal tensor, its spectral norm is equal to  $\Lambda_{\max}$ . Let  $T_W := T(W, W, W)$ ,

$\hat{T}_W := \hat{T}(\hat{W}, \hat{W}, \hat{W})$ , and  $Q := B^\top(\hat{W} - W)$ , then we have

$$\begin{aligned} \hat{T}_W - T_W &= T_W(I + Q, I + Q, I + Q) - T_W + E_T(\hat{W}, \hat{W}, \hat{W}) \\ &= T_W(Q, I, I) + T_W(I, Q, I) + T_W(I, I, Q) \\ &\quad + T_W(Q, Q, I) + T_W(Q, I, Q) + T_W(I, Q, Q) \\ &\quad + T_W(Q, Q, Q) + E_T(\hat{W}, \hat{W}, \hat{W}). \end{aligned}$$

By the second inequality of Lemma 3.8, we know  $\|Q\| \leq 3\epsilon_M/\sigma_{\min}(M) < 1$ . Thus, the first 7 terms of above equation all have spectral norm bounded by  $\|T_W\|\|Q\| \leq 3\Lambda_{\max}\epsilon_M/\sigma_{\min}(M)$ .

The last term  $E_T(\hat{W}, \hat{W}, \hat{W})$  has norm bounded by  $\|E_T\|\|\hat{W}\|^3$ . From (3.25), we know  $\|\hat{W}\| \leq 2\sigma_{\min}(M)^{-1/2}$ . Combining these bounds, we can say that for a large enough constant  $C$ , we have

$$\|\hat{T}_W - T_W\| \leq \epsilon_{T_W} := C \left( \frac{\epsilon_T}{\sigma_{\min}(M)^{3/2}} + \Lambda_{\max} \frac{\epsilon_M}{\sigma_{\min}(M)} \right).$$

By Theorem 3.4, Algorithm 6 will return a set of pairs  $\{(\hat{v}_j, \hat{\Lambda}_j) : j \in [k]\}$ , where with probability at least  $1 - \eta$ , there exists a permutation  $\pi$  on  $[k]$  such that

$$\|v_{\pi(j)} - \hat{v}_j\| \leq 8\epsilon_{T_W}/\Lambda_{\pi(j)}, \quad |\Lambda_{\pi(j)} - \hat{\Lambda}_j| \leq 5\epsilon_{T_W}, \quad j \in [k].$$

The final outputs of the algorithm are  $\hat{a}_j = \hat{B}\hat{v}_j$  and  $\hat{\Lambda}_j$ , for  $j \in [k]$ . The estimation guarantees of the eigenvalues  $\Lambda_j$ 's are already concluded above. We only need to analyze the perturbation of the unwhitening procedure for  $\hat{a}_j$ 's.

Note that  $\sqrt{\tilde{\lambda}_i}a_i = Bv_i$ . Therefore, to compare  $\sqrt{\tilde{\lambda}_{\pi(j)}}a_{\pi(j)}$  with  $\hat{a}_j$ , we only need to compare  $Bv_{\pi(j)}$  with  $\hat{B}\hat{v}_j$ .

$$\begin{aligned} \|\sqrt{\tilde{\lambda}_{\pi(j)}}a_{\pi(j)} - \hat{a}_j\| &= \|Bv_{\pi(j)} - \hat{B}\hat{v}_j\| \\ &\leq \|B(v_{\pi(j)} - \hat{v}_j)\| + \|(B - \hat{B})\hat{v}_j\| \\ &\leq \|B\|\|(v_{\pi(j)} - \hat{v}_j)\| + \|B - \hat{B}\| \\ &\leq 8\epsilon_{T_W}\sqrt{\|M\|}/\Lambda_{\pi(j)} + 3\epsilon_M\sqrt{\|M\|}/\sigma_{\min}(M) \\ &\leq 9\epsilon_{T_W}\sqrt{\|M\|}/\Lambda_{\pi(j)}, \end{aligned}$$

---

**Algorithm 7** Simultaneous Diagonalization for Tensor Decomposition

---

**input** tensor  $T = \sum_{j \in [k]} \lambda_j a_j \otimes b_j \otimes c_j \in \mathbb{R}^{d_1 \times d_2 \times d_3}$

**output** rank-1 components of tensor  $T$

- 1: Pick two random vectors  $x, y \sim \mathcal{N}(0, I_{d_3})$ .
- 2: Compute matrices

$$M_x := T(I, I, x) \in \mathbb{R}^{d_1 \times d_2}, \quad M_y := T(I, I, y) \in \mathbb{R}^{d_1 \times d_2};$$

see (3.3) for the definition of the multilinear form.

- 3: Let

- $\{(\hat{a}_j, \hat{a}_j)\}$ 's be the eigenvalues & eigenvectors of  $M_x M_y^\dagger$ .
- $\{(\hat{\beta}_j, \hat{\beta}_j)\}$ 's be the eigenvalues & eigenvectors of  $M_y^\top (M_x^\dagger)^\top$ .

Here  $\cdot^\dagger$  denotes the pseudo-inverse matrix; see Definition 2.6.

- 4: For any  $j \in [k]$ , pair  $(\hat{a}_j, \hat{\beta}_j)$  if the corresponding eigenvalues satisfy  $\hat{a}_j \hat{\beta}_j = 1$ .
  - 5: Fixing  $(\hat{a}_j, \hat{\beta}_j), j \in [k]$ , solve the linear system  $T = \sum_{j=1}^k \hat{a}_j \otimes \hat{\beta}_j \otimes \hat{c}_j$  in terms of variables  $\hat{c}_j$ 's.
  - 6: Set  $\hat{\lambda}_j = \|\hat{c}_j\|$  and  $\hat{c}_j = \hat{c}_j / \|\hat{c}_j\|$ .
  - 7: **return**  $\{(\hat{\lambda}_j; \hat{a}_j, \hat{\beta}_j, \hat{c}_j) : j \in [k]\}$
- 

where we used the fact that  $\|\hat{v}_j\| = 1$  in the second inequality.  $\|B - \hat{B}\|$  is bounded by the result of Lemma 3.8. The final step is true because  $\epsilon_{T_W} \geq C \Lambda_{\max} \epsilon_M / \sigma_{\min}(M)$ , and therefore,  $3\epsilon_M \sqrt{\|M\|} / \sigma_{\min}(M) \leq \epsilon_{T_W} \sqrt{\|M\|} / \Lambda_{\pi(j)}$  for all  $j$  as long as  $C \geq 3$ . ■

### 3.6 Simultaneous Diagonalization

In this section, we describe simultaneous diagonalization algorithm which is one of the first algorithms with provable guarantees for tensor decomposition. It was discovered in Harshman (1970) (and credited to Dr. Robert Jenrich), with generalizations in Leurgans *et al.* (1993). Simultaneous diagonalization method for tensor decomposition is provided in Algorithm 7.

Comparing to the power method, simultaneous diagonalization is

much easier to analyze, does not require the whitening procedure and can work even when the third dimension  $d_3$  is smaller than  $k$ . However, the straightforward implementation of simultaneous diagonalization is not very robust to perturbations. We give the guarantees for the simultaneous diagonalization algorithm in the noiseless setting as follows.

**Theorem 3.9.** [Simultaneous Diagonalization Guarantees in Noiseless Setting] Suppose tensor  $T$  has a rank- $k$  decomposition

$$T = \sum_{j \in [k]} \lambda_j a_j \otimes b_j \otimes c_j \in \mathbb{R}^{d_1 \times d_2 \times d_3}. \quad (3.29)$$

In addition, suppose vectors  $\{a_j\}$ 's and  $\{b_j\}$ 's are both linearly independent, and vectors  $\{c_j\}$ 's have Kruskal rank at least 2, i.e., no two  $c_i$  and  $c_j$  for  $i \neq j$  are on the same line or parallel; see Definition 3.1 for the definition of Kruskal rank. Then, with probability 1 (over the randomness of vectors  $x$  and  $y$  in the algorithm), Algorithm 7 returns a group of 4-tuples  $(\hat{\lambda}_j; \hat{a}_j, \hat{b}_j, \hat{c}_j)$  such that

$$T = \sum_{j \in [k]} \hat{\lambda}_j \hat{a}_j \otimes \hat{b}_j \otimes \hat{c}_j.$$

Furthermore,  $(\hat{\lambda}_j; \hat{a}_j, \hat{b}_j, \hat{c}_j)$  is equivalent to  $(\lambda_j; a_j, b_j, c_j)$  up to permutation and scaling.

In the rest of this section, we illustrate the ideas in different steps of the algorithm which clarifies how it decomposes rank- $k$  tensor  $T$  in (3.29), and also provides an informal proof for the above theorem.

First, we describe the structure and properties of matrices  $M_x, M_y$  (see Step 2 of the algorithm), which also clarifies why this algorithm is called simultaneous diagonalization. Following the above tensor decomposition structure for tensor  $T$ , and given the multilinear form as a linear combination of tensor slices through weight vectors  $x$  and  $y$ , we have

$$M_x = \sum_{j \in [k]} \lambda_j \langle x, c_j \rangle a_j b_j^\top = A D_x B^\top,$$

$$M_y = \sum_{j \in [k]} \lambda_j \langle y, c_j \rangle a_j b_j^\top = A D_y B^\top.$$

Here  $A \in \mathbb{R}^{d_1 \times k}$ ,  $B \in \mathbb{R}^{d_2 \times k}$  are matrices whose columns are  $\{a_j\}$ 's and  $\{b_j\}$ 's, respectively. We also define  $D_x \in \mathbb{R}^{k \times k}$  as a diagonal matrix whose  $(j, j)$ -the entry denoted by  $d_{x,j}$  is equal to  $\lambda_j \langle x, c_j \rangle$ , and similarly,  $D_y \in \mathbb{R}^{k \times k}$  as a diagonal matrix whose  $(j, j)$ -the entry denoted by  $d_{y,j}$  is equal to  $\lambda_j \langle y, c_j \rangle$ . These equations are called *diagonalizations* of  $M_x, M_y$ , and they share the *same* matrices  $A, B$ . That is why this algorithm is called *simultaneous diagonalization*.

With the above forms of  $M_x$  and  $M_y$ , we can compute the two matrices used in Step 3 as

$$\begin{aligned} M_x M_y^\dagger &= A D_x D_y^{-1} A^\dagger, \\ M_y^\top (M_x^\dagger)^\top &= B D_y D_x^{-1} B^\dagger. \end{aligned}$$

Given this, for any  $a_j$  in (3.29), we have

$$M_x M_y^\dagger a_j = A D_x D_y^{-1} A^\dagger a_j = A D_x D_y^{-1} e_j^{(k)} = d_{x,j} d_{y,j}^{-1} a_j.$$

Here in the second equality, we used the fact that  $A^\dagger A = I_k$  (with  $I_k$  denoting the  $k$ -dimensional identity matrix), and hence,  $A^\dagger a_j = e_j^{(k)}$ , where  $e_j^{(k)}$  denotes the  $j$ -th basis vector in the  $k$ -dimensional space, i.e., the  $j$ -th column of  $I_k$ . Similarly for any  $b_j$  in (3.29), we have

$$M_y^\top (M_x^\dagger)^\top b_j = d_{x,j}^{-1} d_{y,j} b_j.$$

Therefore, by the definition of matrix eigenvectors,  $\{a_j\}$ 's and  $\{b_j\}$ 's are exactly the eigenvectors of the two matrices, and the corresponding eigenvalues  $\alpha_j := d_{x,j} d_{y,j}^{-1}$  and  $\beta_j := d_{x,j}^{-1} d_{y,j}$  satisfy

$$\alpha_j \beta_j = 1, \quad j \in [k].$$

Thus, as long as the values  $\alpha_j$ 's (and inherently  $\beta_j$ 's) are unique for  $j \in [k]$ , Steps 3 and 4 of Algorithm 7 correctly find the set of  $\{(a_j, b_j)\}$ 's up to permutation. The values of  $\alpha_j$ 's (and inherently  $\beta_j$ 's) rely on the randomness of vectors  $x$  and  $y$ , and when  $\{c_j\}$ 's have Kruskal rank at least two, the values are distinct with probability 1.

Finally in Step 5, the algorithm recovers the components of the last mode of the tensor by fixing  $\{a_j\}$ 's,  $\{b_j\}$ 's, and solving a system of linear equations in terms of  $\{c_j\}$ 's. Note that the same idea appears in Alternating Least Squares algorithm that we will discuss in the next

section with more details on how to efficiently solve such a system of linear equations. When  $\{a_j\}$ 's and  $\{b_j\}$ 's are linearly independent, this step will have a unique solution. Hence, the algorithm finds the unique decomposition of tensor  $T$ .

Note that simultaneous diagonalization algorithm corresponds to a special tight case of Kruskal's condition in Theorem 3.1, where  $\text{krank}(A) = k$ ,  $\text{krank}(B) = k$ ,  $\text{krank}(C) \geq 2$ . Weakening the assumption on any of the matrices  $A$ ,  $B$  or  $C$  may make the tensor decomposition non-unique. Compared to tensor power method (with the symmetrization procedure in Section 3.4.4), the simultaneous diagonalization method does not need access to second moment matrices, and can allow one of the factors to have  $\text{krank} 2$ .

### 3.7 Alternating Least Squares

One of the most popular algorithms for tensor decomposition is *Alternating Least Squares (ALS)* method, which has been described as the “workhorse” of tensor decomposition (Kolda and Bader, 2009). This involves solving the least squares problem on a *mode* of the tensor, while keeping the other modes fixed, and alternating between the tensor modes. This becomes clearer as we describe the details of ALS as follows.

Given rank- $k$  tensor

$$T = \sum_{j \in [k]} \lambda_j a_j \otimes b_j \otimes c_j \in \mathbb{R}^{d_1 \times d_2 \times d_3},$$

the goal is to recover tensor rank-1 components  $a_j \in \mathbb{R}^{d_1}$ ,  $b_j \in \mathbb{R}^{d_2}$ ,  $c_j \in \mathbb{R}^{d_3}$ , and coefficients  $\lambda_j \in \mathbb{R}$ , for  $j \in [k]$ . As before, we assume the rank-1 components have unit norm. The problem can be formulated as the *least squares* optimization

$$\begin{aligned}
 (\lambda_j; a_j, b_j, c_j)_{j \in [k]} &:= \arg \min_{\tilde{\lambda}_j, \tilde{a}_j, \tilde{b}_j, \tilde{c}_j} \left\| T - \sum_{j \in [k]} \tilde{\lambda}_j \tilde{a}_j \otimes \tilde{b}_j \otimes \tilde{c}_j \right\|_F, \quad (3.30) \\
 \text{s.t. } \tilde{\lambda}_j &\in \mathbb{R}, \tilde{a}_j \in \mathbb{R}^{d_1}, \tilde{b}_j \in \mathbb{R}^{d_2}, \tilde{c}_j \in \mathbb{R}^{d_3},
 \end{aligned}$$

where the error between tensor  $T$  and its rank- $k$  estimation is minimized in the sense of Frobenius norm. This is a multilinear optimization program and a *non-convex* optimization problem.

Alternating Least Squares method provides an approach to overcome the non-convexity challenge. This mainly involves modifying the optimization problem such that the optimization is performed for only one of the components while all other components are assumed to be fixed. The same step is performed by *alternating* among different components (modes). Thus, the problem in (3.30) is solved through an alternating least squares approach.

We now describe the main step of ALS by fixing the second and third mode rank-1 components, i.e., matrices  $\tilde{B} := [\tilde{b}_1 | \tilde{b}_2 | \dots | \tilde{b}_k]$  and  $\tilde{C} := [\tilde{c}_1 | \tilde{c}_2 | \dots | \tilde{c}_k]$ , and optimizing over first mode, i.e., matrix  $\tilde{A} := [\tilde{a}_1 | \tilde{a}_2 | \dots | \tilde{a}_k]$ . As the first step, since the Frobenius norm in (3.30) is an entry-wise tensor norm, we can reshape the tensor inside without changing the norm. More specifically, we can rewrite the optimization problem in (3.30) into the following equivalent form:

$$\min_{\tilde{A} \in \mathbb{R}^{d_1 \times k}} \|\text{mat}(T, 1) - \tilde{A} \cdot \text{diag}(\tilde{\lambda}) \cdot (\tilde{B} \odot \tilde{C})^\top\|_F, \tag{3.31}$$

where  $\text{mat}(T, 1)$  denotes the mode-1 matricization of  $T$ , and  $\odot$  denotes the Khatri-Rao product; see (3.2) and (3.9), respectively. Note that here we assume  $\tilde{B}$ ,  $\tilde{C}$ , and  $\tilde{\lambda}$  are fixed. We also used the following matricization property such that for vectors  $u, v, w$ , we have

$$\text{mat}(u \otimes v \otimes w, 1) = u \cdot (v \odot w)^\top.$$

The optimization problem in (3.31) is now a *linear* least squares problem, and the analysis is very similar to linear regression with the additional property that matrix  $(\tilde{B} \odot \tilde{C})^\top$  is highly-structured which is crucial for the computational efficiency of ALS as we see below. The (right) pseudo-inverse of  $(\tilde{B} \odot \tilde{C})^\top$  is (see (3.8) for the definition)

$$\left[ (\tilde{B} \odot \tilde{C})^\top \right]^\dagger = (\tilde{B} \odot \tilde{C}) \cdot \left[ (\tilde{B} \odot \tilde{C})^\top (\tilde{B} \odot \tilde{C}) \right]^{-1}.$$

Computing the inverse matrix could be the computationally-expensive part of the iterations, but the specific Khatri-Rao structure of the matrix enables us to write it as

$$(\tilde{B} \odot \tilde{C})^\top (\tilde{B} \odot \tilde{C}) = \tilde{B}^\top \tilde{B} * \tilde{C}^\top \tilde{C},$$

where  $*$  denotes the Hadamard (entry-wise) product. Thus, we only need to compute the inverse of  $k \times k$  matrix  $\tilde{B}^\top \tilde{B} * \tilde{C}^\top \tilde{C}$ , and when  $k$  is

---

**Algorithm 8** Alternating Least Squares for Tensor Decomposition
 

---

**input** tensor  $T = \sum_{j \in [k]} \lambda_j a_j \otimes b_j \otimes c_j \in \mathbb{R}^{d_1 \times d_2 \times d_3}$

**output** rank-1 components of tensor  $T$

1: Set initial estimates for  $\tilde{A} \in \mathbb{R}^{d_1 \times k}$ ,  $\tilde{B} \in \mathbb{R}^{d_2 \times k}$ ,  $\tilde{C} \in \mathbb{R}^{d_3 \times k}$ .

2: **while** converged **do**

3: Let  $\bar{A} = \text{mat}(T, 1) \cdot (\tilde{B} \odot \tilde{C}) \cdot (\tilde{B}^\top \tilde{B} * \tilde{C}^\top \tilde{C})^{-1}$ , and set

$$\tilde{\lambda}_j = \|\bar{A}_j\|, \quad \tilde{A}_j = \bar{A}_j / \tilde{\lambda}_j, \quad j \in [k].$$

4: Let  $\bar{B} = \text{mat}(T, 2) \cdot (\tilde{A} \odot \tilde{C}) \cdot (\tilde{A}^\top \tilde{A} * \tilde{C}^\top \tilde{C})^{-1}$ , and set

$$\tilde{\lambda}_j = \|\bar{B}_j\|, \quad \tilde{B}_j = \bar{B}_j / \tilde{\lambda}_j, \quad j \in [k].$$

5: Let  $\bar{C} = \text{mat}(T, 3) \cdot (\tilde{A} \odot \tilde{B}) \cdot (\tilde{A}^\top \tilde{A} * \tilde{B}^\top \tilde{B})^{-1}$ , and set

$$\tilde{\lambda}_j = \|\bar{C}_j\|, \quad \tilde{C}_j = \bar{C}_j / \tilde{\lambda}_j, \quad j \in [k].$$

6: **end while**

7: **return**  $(\tilde{\lambda}; \tilde{A}, \tilde{B}, \tilde{C})$  as the estimation of tensor decomposition components.

---

small (compared to  $d_t, t \in \{1, 2, 3\}$ ), the inverse can be computed much faster. In practice we can also compute the Hadamard product and then solve a linear system of equations rather than explicitly computing the inverse.

Finally, using the above property and imposing the unit norm constraint on the columns of  $\tilde{A}$ , we update the rank-1 components of the first mode as

$$\tilde{A} \mapsto \text{Norm} \left( \text{mat}(T, 1) \cdot (\tilde{B} \odot \tilde{C}) \cdot (\tilde{B}^\top \tilde{B} * \tilde{C}^\top \tilde{C})^{-1} \right), \quad (3.32)$$

where operator  $\text{Norm}(\cdot)$  normalizes the columns of input matrix, i.e., for vector  $v$ , we have  $\text{Norm}(v) := v / \|v\|$ .

By alternating between different modes of the tensor and with similar



calculations, we update the second and third modes as

$$\begin{aligned}\tilde{B} &\mapsto \text{Norm} \left( \text{mat}(T, 2) \cdot (\tilde{A} \odot \tilde{C}) \cdot \left( \tilde{A}^\top \tilde{A} * \tilde{C}^\top \tilde{C} \right)^{-1} \right), \\ \tilde{C} &\mapsto \text{Norm} \left( \text{mat}(T, 3) \cdot (\tilde{A} \odot \tilde{B}) \cdot \left( \tilde{A}^\top \tilde{A} * \tilde{B}^\top \tilde{B} \right)^{-1} \right).\end{aligned}$$

For the coefficient vector  $\lambda$ , we update it appropriately such that the rank-1 components have unit norm. We have summarized the ALS steps in Algorithm 8.

**ALS vs. power iteration:** The ALS updates in the rank-1 form are strongly related to the power iteration updates. Recall tensor power iteration in (3.16) which can be adapted to the asymmetric setting such that the update corresponding to the first component is (ignoring the normalization)

$$\tilde{a} \mapsto T(I, \tilde{b}, \tilde{c}).$$

The update in the right hand side can be also rewritten as

$$T(I, \tilde{b}, \tilde{c}) = \text{mat}(T, 1) \cdot (\tilde{b} \odot \tilde{c}) \propto \text{mat}(T, 1) \cdot \left( (\tilde{b} \odot \tilde{c})^\top \right)^\dagger,$$

which is basically the rank-1 form of ALS updates that we described in this section. In rank- $k$  ALS, all components are simultaneously updated at each iteration, while the rank-1 version only updates one component at a time which is basically the ALS update that we described in this section, but only for one of the components. The process then needs to be repeated for each remaining component on the deflated tensor; see Algorithm 6. By contrast, in the ALS algorithm we introduced here, all components are simultaneously updated at each iteration. Also note that if the tensor does not have an orthogonal decomposition, ALS can still work while tensor power iteration requires additional whitening step (as in Section 3.4.3) even if the components are linearly independent. However, the benefit of tensor power iteration is that we do have guarantees for it (see Section 3.5.1, while ALS is not known to converge from a random starting point even if the tensor has an orthogonal decomposition.

**Regularized ALS:** Since ALS involves solving linear least squares problems, we can also propose the regularized version of ALS. It is derived by adding a regularization term to the optimization in (3.31). The most popular form of regularization is the  $\ell_2$ -regularization which adds a term  $\alpha \|\tilde{A}\|_F^2$  to the optimization problem, where  $\alpha \geq 0$  is the regularization parameter. This leads to the ALS updates being changed as

$$\tilde{A} \mapsto \text{Norm} \left( \text{mat}(T, 1) \cdot (\tilde{B} \odot \tilde{C}) \cdot \left( \tilde{B}^\top \tilde{B} * \tilde{C}^\top \tilde{C} + \alpha I \right)^{-1} \right); \quad (3.33)$$

and similarly the updates for  $\tilde{B}$  and  $\tilde{C}$  are changed. This is specifically helpful when the non-regularized pseudo-inverse matrix is not well-behaved. We can obviously add other forms of regularization terms to the optimization problem which lead to variants of regularized ALS.

**ALS for symmetric tensors:** The ALS algorithm is naturally proposed for asymmetric tensor decomposition, where at each iteration only one component is updated while all other components are fixed. The next natural question is whether ALS can be adapted to the decomposition of symmetric tensors such as  $T = \sum_{j \in [k]} \lambda_j a_j^{\otimes 3}$ . Here we have to only estimate one matrix  $A$ . We review two heuristics to do this. Let  $\tilde{A}_t$  denote the update variable in the left hand side of (3.32) at iteration  $t$ . The first heuristic consists in, at iteration  $t$ , substituting  $\tilde{B}$  and  $\tilde{C}$  in the right hand side of (3.32) by  $\tilde{A}_{t-1}$  and  $\tilde{A}_{t-2}$ , respectively. The second heuristic consists in substituting both with  $\tilde{A}_{t-1}$ .

# 4

---

## Applications of Tensor Methods

---

In Chapter 1, we gave a few examples of latent variable models that can be learned by tensor decomposition techniques. In this chapter, we elaborate on this connection and give more examples on how to learn many probabilistic models by tensor decomposition. We cover both unsupervised and supervised settings in this chapter. We hope these examples provide a good understanding of how tensor methods are applied in the existing literature and can help in generalizing tensor decomposition techniques to learning more models. Of course, there are still many more applications of tensor decomposition techniques to learn probabilistic models in the literature, and we give a brief survey in Section 4.6.

In the unsupervised setting, we discuss models including the Gaussian mixtures, multiview mixture model, Independent Component Analysis (ICA), Latent Dirichlet Allocation (LDA) and Noisy-Or models. To this end, the observed moment is formed as a low order tensor (usually third or fourth order), and by decomposing the tensor to its rank-1 components we are able to learn the parameters of the model; see Sections 4.1-4.4 which describe this connection. The basic form is demonstrated in Theorem 4.1 for the first example, and the general

pattern will emerge from subsequent examples.

Then in Section 4.5, we show how the tensor techniques can be adapted to supervised setting, and in particular, for learning neural networks and mixtures of generalized linear models. Here, we exploit the cross-moment between the output and a specific non-linear transformation of the input. By decomposing that cross-moment into rank-1 components, we learn the parameters of the model.

## 4.1 Pure Topic Model Revisited

We start by explaining the pure topic model in more details, where it was originally introduced in Section 1.2. Recall the model is a simple bag-of-words model for documents in which the words in the document are assumed to be *exchangeable*—a collection of random variables  $x_1, x_2, \dots, x_\ell$  are exchangeable if their joint probability distribution is invariant to permutation of the indices. The well-known De Finetti’s theorem (Austin, 2008) implies that such exchangeable models can be viewed as mixture models in which there is a latent variable  $h$  such that  $x_1, x_2, \dots, x_\ell$  are conditionally i.i.d. given  $h$  (see Figure 4.1(a) for the corresponding graphical model) and the conditional distributions are identical at all the nodes, i.e., for all  $x$ ’s.

In our simplified topic model for documents, the latent variable  $h$  is interpreted as the (sole) topic of a given document, and it is assumed to take only a finite number of distinct values. Let  $k$  be the number of distinct topics in the corpus,  $d$  be the number of distinct words in the vocabulary, and  $\ell \geq 3$  be the number of words in each document. The generative process for a document is as follows: the document’s topic is drawn according to the discrete distribution specified by the probability vector  $w := (w_1, w_2, \dots, w_k) \in \Delta^{k-1}$ , where  $\Delta^{k-1} := \{v \in \mathbb{R}^k : \forall j \in [k], v_j \in [0, 1], \sum_{j \in [k]} v_j = 1\}$  denotes the probability simplex, i.e., the hidden topic  $h$  is modeled as a discrete random variable  $h$  such that

$$\Pr[h = j] = w_j, \quad j \in [k].$$

Given the topic  $h$ , the document’s  $\ell$  words are drawn independently according to the discrete distribution specified by the probability vector

$\mu_h \in \Delta^{d-1}$ . It will be convenient to represent the  $\ell$  words in the document by  $d$ -dimensional random vectors  $x_1, x_2, \dots, x_\ell \in \mathbb{R}^d$ . Specifically, we set

$$x_t = e_i \quad \text{if and only if} \quad \begin{array}{l} \text{the } t\text{-th word in the document is } i, \\ t \in [\ell], i \in [d], \end{array}$$

where  $\{e_1, e_2, \dots, e_d\}$  is the standard coordinate basis for  $\mathbb{R}^d$ . This is basically equivalent to *one-hot encoding of words* using standard basis vectors in the  $d$ -dimensional space.

As we did in Section 1.2, we will consider the *cross* moments of these vectors which means we will compute  $\mathbb{E}[x_1 x_2^\top]$  instead of  $\mathbb{E}[x_1 x_1^\top]$ . The advantage of the above encoding of words and the choice of moments is that the moments will correspond to joint probabilities over words. For instance, observe that

$$\begin{aligned} \mathbb{E}[x_1 \otimes x_2] &= \sum_{i,j \in [d]} \Pr[x_1 = e_i, x_2 = e_j] e_i \otimes e_j \\ &= \sum_{i,j \in [d]} \Pr[\text{1st word} = i, \text{2nd word} = j] e_i \otimes e_j, \end{aligned}$$

and thus, the  $(i, j)$ -the entry of the moment matrix  $\mathbb{E}[x_1 \otimes x_2]$  is  $\Pr[\text{1st word} = i, \text{2nd word} = j]$ . More generally, the  $(i_1, i_2, \dots, i_\ell)$ -th entry in the tensor  $\mathbb{E}[x_1 \otimes x_2 \otimes \dots \otimes x_\ell]$  is  $\Pr[\text{1st word} = i_1, \text{2nd word} = i_2, \dots, \ell\text{-th word} = i_\ell]$ . This means that estimating cross moments, say, of  $x_1 \otimes x_2 \otimes x_3$ , is the same as estimating joint probabilities of the first three words over all documents; recall that we assume that each document has at least three words.

The second advantage of the vector encoding of words is that the conditional expectation of  $x_t$  given  $h = j$  is simply  $\mu_j$ , the vector of word probabilities for topic  $j$ . This can be shown as

$$\begin{aligned} \mathbb{E}[x_t | h = j] &= \sum_{i \in [d]} \Pr[t\text{-th word} = i | h = j] e_i \\ &= \sum_{i \in [d]} [\mu_j]_i e_i = \mu_j, \quad j \in [k], \end{aligned}$$

where  $[\mu_j]_i$  is the  $i$ -th entry of the vector  $\mu_j$ . Because the words are conditionally independent given the topic, we can use this same property

with conditional cross moments, say, of  $x_1$  and  $x_2$ :

$$\mathbb{E}[x_1 \otimes x_2 | h = j] = \mathbb{E}[x_1 | h = j] \otimes \mathbb{E}[x_2 | h = j] = \mu_j \otimes \mu_j, \quad j \in [k].$$

Now using the law of total expectations, we know

$$\mathbb{E}[x_1 \otimes x_2] = \sum_{j=1}^k \Pr[h = j] \mathbb{E}[x_1 \otimes x_2 | h = j] = \sum_{j=1}^k w_j \mu_j \otimes \mu_j.$$

This and similar calculations lead to the following theorem.

**Theorem 4.1** (Anandkumar *et al.*, 2012c). For the above exchangeable single topic model, if

$$\begin{aligned} M_2 &:= \mathbb{E}[x_1 \otimes x_2], \\ M_3 &:= \mathbb{E}[x_1 \otimes x_2 \otimes x_3], \end{aligned}$$

then

$$\begin{aligned} M_2 &= \sum_{j \in [k]} w_j \mu_j \otimes \mu_j, \\ M_3 &= \sum_{j \in [k]} w_j \mu_j \otimes \mu_j \otimes \mu_j. \end{aligned}$$

The structure of  $M_2$  and  $M_3$  revealed in Theorem 4.1 implies that the topic vectors  $\mu_1, \mu_2, \dots, \mu_k$  can be estimated by computing a certain symmetric tensor decomposition. Moreover, due to exchangeability, any triples (resp., pairs) of words in a document—and not just the first three (resp., two) words—can be used in forming  $M_3$  (resp.,  $M_2$ ).

## 4.2 Beyond Raw Moments

In the above exchangeable single topic model, the *raw* (cross) moments of the observed words directly yield the desired symmetric tensor structure. In some other models, the raw moments do not explicitly have this form. In this section, we show that the desired tensor structure can be found through various manipulations of different moments for some other latent variable models.

### 4.2.1 Spherical Gaussian mixtures

We now consider a mixture of  $k$  Gaussian distributions with spherical covariances. We start with the simpler case where all of the covariances are identical; this probabilistic model is closely related to the (non-probabilistic)  $k$ -means clustering problem (MacQueen, 1967). We then consider the case where the spherical variances may differ.

**Common covariance.** Let  $w_j \in (0, 1)$  be the probability of choosing component  $j \in [k]$ ,  $\mu_1, \mu_2, \dots, \mu_k \in \mathbb{R}^d$  be the component mean vectors, and  $\sigma^2 I_d \in \mathbb{R}^{d \times d}$  be the common covariance matrix ( $\sigma \in \mathbb{R}$ ) for the spherical Gaussian mixtures model. Then an observation vector  $x$  in this model is given by

$$x := \mu_h + z,$$

where  $h$  is the discrete random variable with  $\Pr[h = j] = w_j$  for  $j \in [k]$  (similar to the exchangeable single topic model), and  $z \sim \mathcal{N}(0, \sigma^2 I_d)$  is an independent multivariate Gaussian random vector in  $\mathbb{R}^d$  with zero mean and spherical covariance matrix  $\sigma^2 I_d$ .

The Gaussian mixtures model differs from the exchangeable single topic model in the way observations are generated. In the single topic model, we observe multiple draws (words in a particular document)  $x_1, x_2, \dots, x_\ell$  given the same fixed  $h$  (the topic of the document). In contrast, for the Gaussian mixtures model, every realization of  $x$  corresponds to a different realization of  $h$ . The following theorem shows that how we can get the desired tensor decomposition form by modifying the raw moments.

**Theorem 4.2** (Hsu and Kakade, 2013). Assume  $d \geq k$ . The variance  $\sigma^2$  is the smallest eigenvalue of the covariance matrix  $\mathbb{E}[x \otimes x] - \mathbb{E}[x] \otimes \mathbb{E}[x]$ . Furthermore, if

$$\begin{aligned} M_2 &:= \mathbb{E}[x \otimes x] - \sigma^2 I_d, \\ M_3 &:= \mathbb{E}[x \otimes x \otimes x] \\ &\quad - \sigma^2 \sum_{i \in [d]} (\mathbb{E}[x] \otimes e_i \otimes e_i + e_i \otimes \mathbb{E}[x] \otimes e_i + e_i \otimes e_i \otimes \mathbb{E}[x]), \end{aligned}$$

then,

$$M_2 = \sum_{j \in [k]} w_j \mu_j \otimes \mu_j,$$

$$M_3 = \sum_{j \in [k]} w_j \mu_j \otimes \mu_j \otimes \mu_j.$$

**Differing covariances.** The general case is where each component may have a *different* spherical covariance. An observation in this model is again  $x = \mu_h + z$ , but now  $z \in \mathbb{R}^d$  is a random vector whose conditional distribution given  $h = j$  for some  $j \in [k]$  is a multivariate Gaussian  $\mathcal{N}(0, \sigma_j^2 I_d)$  with zero mean and spherical covariance  $\sigma_j^2 I_d$ .

**Theorem 4.3** (Hsu and Kakade, 2013). Assume  $d \geq k$ . The average variance  $\bar{\sigma}^2 := \sum_{j \in [k]} w_j \sigma_j^2$  is the smallest eigenvalue of the covariance matrix  $\mathbb{E}[x \otimes x] - \mathbb{E}[x] \otimes \mathbb{E}[x]$ . Let  $v$  be any unit norm eigenvector corresponding to the eigenvalue  $\bar{\sigma}^2$ . If

$$M_1 := \mathbb{E}[\langle v, x - \mathbb{E}[x] \rangle^2 x],$$

$$M_2 := \mathbb{E}[x \otimes x] - \bar{\sigma}^2 I_d,$$

$$M_3 := \mathbb{E}[x \otimes x \otimes x]$$

$$- \sum_{i \in [d]} (M_1 \otimes e_i \otimes e_i + e_i \otimes M_1 \otimes e_i + e_i \otimes e_i \otimes M_1),$$

where  $\langle \cdot, \cdot \rangle$  denotes the inner-product operator. Then

$$M_2 = \sum_{j \in [k]} w_j \mu_j \otimes \mu_j,$$

$$M_3 = \sum_{j \in [k]} w_j \mu_j \otimes \mu_j \otimes \mu_j.$$

As shown by Hsu and Kakade (2013),  $M_1 = \sum_{j \in [k]} w_j \sigma_j^2 \mu_j$ . Note that for the common covariance case, where  $\sigma_j^2 = \sigma^2$ , we have that  $M_1 = \sigma^2 \mathbb{E}[x]$ ; see Theorem 4.2.

#### 4.2.2 Independent component analysis (ICA)

The standard model for ICA (Comon, 1994; Cardoso and Comon, 1996; Hyvärinen and Oja, 2000; Comon and Jutten, 2010), in which



independent signals are linearly mixed and corrupted with Gaussian noise before being observed, is specified as follows. Let  $h \in \mathbb{R}^k$  be a latent random *vector* with independent coordinates,  $A \in \mathbb{R}^{d \times k}$  the mixing matrix, and  $z \in \mathbb{R}^d$  be a multivariate Gaussian random vector. The random vectors  $h$  and  $z$  are assumed to be independent. The observed random vector  $x$  in this model is given by

$$x := Ah + z.$$

Let  $\mu_j$  denote the  $j$ -th column of the mixing matrix  $A$ .

**Theorem 4.4** (Comon and Jutten, 2010). Define

$$M_4 := \mathbb{E}[x \otimes x \otimes x \otimes x] - T,$$

where  $T \in \mathbb{R}^{d \times d \times d \times d}$  is the fourth-order tensor with

$$\begin{aligned} [T]_{i_1, i_2, i_3, i_4} &:= \mathbb{E}[x_{i_1} x_{i_2}] \mathbb{E}[x_{i_3} x_{i_4}] + \mathbb{E}[x_{i_1} x_{i_3}] \mathbb{E}[x_{i_2} x_{i_4}] \\ &\quad + \mathbb{E}[x_{i_1} x_{i_4}] \mathbb{E}[x_{i_2} x_{i_3}], \quad 1 \leq i_1, i_2, i_3, i_4 \leq d, \end{aligned}$$

*i.e.*,  $T$  is the fourth derivative tensor of the function  $v \mapsto 8^{-1} \mathbb{E}[(v^\top x)^2]^2$ , and so,  $M_4$  is the fourth cumulant tensor. Let  $\kappa_j := \mathbb{E}[h_j^4] - 3$  for each  $j \in [k]$ . Then

$$M_4 = \sum_{j \in [k]} \kappa_j \mu_j \otimes \mu_j \otimes \mu_j \otimes \mu_j.$$

Note that  $\kappa_j$  corresponds to the excess kurtosis, a measure of non-Gaussianity as  $\kappa_j = 0$  if  $h_j$  is a standard normal random variable. Hence, mixing matrix  $A$  is not identifiable if  $h$  is a multivariate Gaussian.

We may derive forms similar to that of  $M_2$  and  $M_3$  in Theorem 4.1 using  $M_4$  by observing that

$$\begin{aligned} M_4(I, I, u, v) &= \sum_{j \in [k]} \kappa_j (\mu_j^\top u) (\mu_j^\top v) \mu_j \otimes \mu_j, \\ M_4(I, I, I, v) &= \sum_{j \in [k]} \kappa_j (\mu_j^\top v) \mu_j \otimes \mu_j \otimes \mu_j, \end{aligned}$$

for any vectors  $u, v \in \mathbb{R}^d$ .

### 4.2.3 Latent Dirichlet Allocation

An increasingly popular class of latent variable models are *mixed membership models*, where each datum may belong to several different latent classes simultaneously. Latent Dirichlet Allocation (LDA, Blei *et al.*, 2003) is one such model for the case of document modeling; here, each document corresponds to a mixture over topics (as opposed to just a single topic that we discussed in Section 4.1). The distribution over such topic mixtures is a Dirichlet distribution  $\text{Dir}(\alpha)$  with parameter vector  $\alpha \in \mathbb{R}_{++}^k$  with strictly positive entries; its density function over the probability simplex  $\Delta^{k-1} := \{v \in \mathbb{R}^k : \forall j \in [k], v_j \in [0, 1], \sum_{j \in [k]} v_j = 1\}$  is given by

$$p_\alpha(h) = \frac{\Gamma(\alpha_0)}{\prod_{j \in [k]} \Gamma(\alpha_j)} \prod_{j \in [k]} h_j^{\alpha_j - 1}, \quad h \in \Delta^{k-1},$$


where

$$\alpha_0 := \alpha_1 + \alpha_2 + \cdots + \alpha_k,$$

and  $\Gamma(\cdot)$  denotes the Gamma function.

As before, the  $k$  topics are specified by probability vectors  $\mu_1, \mu_2, \dots, \mu_k \in \Delta^{d-1}$  for generating words. To generate a document, we first draw the topic mixture  $h = (h_1, h_2, \dots, h_k) \sim \text{Dir}(\alpha)$ , and then conditioned on  $h$ , we draw  $\ell$  words  $x_1, x_2, \dots, x_\ell$  independently from the discrete distribution specified by the probability vector  $\sum_{j \in [k]} h_j \mu_j$ , *i.e.*, for each word  $x_t$ , we independently sample a topic  $j$  according to the topic proportion vector  $h$  and then sample  $x_t$  according to  $\mu_j$ . Again, we encode a word  $x_t$  by setting  $x_t = e_i$  if and only if the  $t$ -th word in the document is  $i$ .

The parameter  $\alpha_0$  (the sum of the “pseudo-counts”) characterizes the concentration of the distribution. As  $\alpha_0 \rightarrow 0$ , the distribution degenerates to a single topic model, *i.e.*, the limiting density has, with probability 1, exactly one entry of  $h$  being 1 and the rest are 0. At the other extreme, if  $\alpha = (c, c, \dots, c)$  for some scalar  $c > 0$ , then as  $\alpha_0 = ck \rightarrow \infty$ , the distribution of  $h$  becomes peaked around the uniform vector  $(1/k, 1/k, \dots, 1/k)$ , and furthermore, the distribution behaves like a product distribution. We are typically interested in the case where  $\alpha_0$  is small (*e.g.*, a constant independent of  $k$ ), whereupon  $h$  typically



**Figure 4.1:** Examples of latent variable models

has only a few large entries. This corresponds to the setting where the documents are mainly comprised of just a few topics.

**Theorem 4.5** (Anandkumar *et al.*, 2012a). Define

$$M_1 := \mathbb{E}[x_1],$$

$$M_2 := \mathbb{E}[x_1 \otimes x_2] - \frac{\alpha_0}{\alpha_0 + 1} M_1 \otimes M_1,$$

$$M_3 := \mathbb{E}[x_1 \otimes x_2 \otimes x_3]$$

$$- \frac{\alpha_0}{\alpha_0 + 2} \left( \mathbb{E}[x_1 \otimes x_2 \otimes M_1] + \mathbb{E}[x_1 \otimes M_1 \otimes x_2] + \mathbb{E}[M_1 \otimes x_1 \otimes x_2] \right)$$

$$+ \frac{2\alpha_0^2}{(\alpha_0 + 2)(\alpha_0 + 1)} M_1 \otimes M_1 \otimes M_1.$$

Then,

$$M_2 = \sum_{j \in [k]} \frac{\alpha_j}{(\alpha_0 + 1)\alpha_0} \mu_j \otimes \mu_j,$$

$$M_3 = \sum_{j \in [k]} \frac{2\alpha_j}{(\alpha_0 + 2)(\alpha_0 + 1)\alpha_0} \mu_j \otimes \mu_j \otimes \mu_j.$$

Note that  $\alpha_0$  needs to be known to form  $M_2$  and  $M_3$  from the raw moments. This, however, is a much weaker assumption than assuming that the entire distribution of  $h$  is known, *i.e.*, knowledge of the whole parameter vector  $\alpha$ .

### 4.3 Multi-view Models

Multi-view models (also sometimes called naïve Bayes models) are a special class of Bayesian networks in which observed variables  $x_1, x_2, \dots, x_\ell$

are conditionally independent given a latent variable  $h$ . This is similar to the exchangeable single topic model, but here we do not require the conditional distributions of the  $x_t, t \in [\ell]$ , to be identical. Techniques developed for this class can be used to handle a number of widely used models including hidden Markov models (HMMs) (Mossel and Roch, 2006; Anandkumar *et al.*, 2012c), phylogenetic tree models (Chang, 1996; Mossel and Roch, 2006), certain tree mixtures (Anandkumar *et al.*, 2012b), and certain probabilistic grammar models (Hsu *et al.*, 2012).

As before, we let  $h \in [k]$  be a discrete random variable with  $\Pr[h = j] = w_j$  for all  $j \in [k]$ . Now consider random vectors  $x_1 \in \mathbb{R}^{d_1}$ ,  $x_2 \in \mathbb{R}^{d_2}$ , and  $x_3 \in \mathbb{R}^{d_3}$  which are conditionally independent given  $h$  (see Figure 4.1(a) for the corresponding graphical model), and

$$\mathbb{E}[x_t | h = j] = \mu_{t,j}, \quad j \in [k], t \in \{1, 2, 3\},$$

where  $\mu_{t,j} \in \mathbb{R}^{d_t}$  are the conditional means of  $x_t$  given  $h = j$ . Thus, we allow the observations  $x_1, x_2, \dots, x_\ell$  to be random vectors, parameterized only by their conditional means. Importantly, these conditional distributions may be discrete, continuous, or even a mix of both.

We first note the form for the raw (cross) moments.

**Proposition 4.3.1.** We have

$$\begin{aligned} \mathbb{E}[x_t \otimes x_{t'}] &= \sum_{j \in [k]} w_j \mu_{t,j} \otimes \mu_{t',j}, \quad \{t, t'\} \subset \{1, 2, 3\}, t \neq t', \\ \mathbb{E}[x_1 \otimes x_2 \otimes x_3] &= \sum_{j \in [k]} w_j \mu_{1,j} \otimes \mu_{2,j} \otimes \mu_{3,j}. \end{aligned}$$

The cross moments do not possess a symmetric tensor form when the conditional distributions are different. We can either apply asymmetric tensor decomposition techniques to estimate conditional mean vectors  $\mu_{t,j}$ , or symmetrize the tensors by the following trick and then apply symmetric tensor decomposition techniques. Nevertheless, the moments can be “symmetrized” via a simple linear transformation of  $x_1$  and  $x_2$  (roughly speaking, this relates  $x_1$  and  $x_2$  to  $x_3$ ); this leads to an expression from which the conditional means of  $x_3$  (*i.e.*,  $\mu_{3,1}, \mu_{3,2}, \dots, \mu_{3,k}$ ) can be recovered. For simplicity, we assume  $d_1 = d_2 = d_3 = k$ ; the

general case (with  $d_t \geq k$ ) is easily handled using low-rank singular value decompositions.

**Theorem 4.6** (Anandkumar *et al.*, 2012a). Assume that  $\{\mu_{t,1}, \mu_{t,2}, \dots, \mu_{t,k}\}$  are linearly independent for each  $t \in \{1, 2, 3\}$ . Define

$$\begin{aligned}\tilde{x}_1 &:= \mathbb{E}[x_3 \otimes x_2] \mathbb{E}[x_1 \otimes x_2]^{-1} x_1, \\ \tilde{x}_2 &:= \mathbb{E}[x_3 \otimes x_1] \mathbb{E}[x_2 \otimes x_1]^{-1} x_2,\end{aligned}$$

and

$$\begin{aligned}M_2 &:= \mathbb{E}[\tilde{x}_1 \otimes \tilde{x}_2], \\ M_3 &:= \mathbb{E}[\tilde{x}_1 \otimes \tilde{x}_2 \otimes x_3].\end{aligned}$$

Then,

$$\begin{aligned}M_2 &= \sum_{j \in [k]} w_j \mu_{3,j} \otimes \mu_{3,j}, \\ M_3 &= \sum_{j \in [k]} w_j \mu_{3,j} \otimes \mu_{3,j} \otimes \mu_{3,j}.\end{aligned}$$

We now discuss three examples mostly taken from Anandkumar *et al.* (2012c) where the above observations can be applied. The first two concern mixtures of product distributions, and the last one is the time-homogeneous hidden Markov model.

### 4.3.1 Mixtures of axis-aligned Gaussians and other product distributions

The first example is a mixture of  $k$  product distributions in  $\mathbb{R}^d$  under a mild incoherence assumption (Anandkumar *et al.*, 2012c). Here, we allow each of the  $k$  component distributions to have a different product distribution (*e.g.*, Gaussian distribution with an axis-aligned covariance matrix), but require the matrix of component means  $A := [\mu_1 | \mu_2 | \dots | \mu_k] \in \mathbb{R}^{d \times k}$  to satisfy a certain (very mild) incoherence condition. The role of the incoherence condition is explained below.

For a mixture of product distributions, any partitioning of the dimensions  $[d]$  into three groups creates three (possibly asymmetric)

“views” which are conditionally independent once the mixture component is selected. However, recall that Theorem 4.6 requires that for each view, the  $k$  conditional means be linearly independent. In general, this may not be achievable; consider, for instance, the case  $\mu_i = e_i$  for each  $i \in [k]$ . Such cases, where the component means are very aligned with the coordinate basis, are precluded by the incoherence condition.

Let  $\Pi_A$  denote the orthogonal projector operator to the range of  $A$  and define  $\text{coherence}(A) := \max_{i \in [d]} \{e_i^\top \Pi_A e_i\}$  to be the largest diagonal entry of this operator, and assume  $A$  has rank  $k$ . The coherence lies between  $k/d$  and 1; it is largest when the range of  $A$  is spanned by the coordinate axes, and it is  $k/d$  when the range is spanned by a subset of the Hadamard basis of cardinality  $k$ . The incoherence condition requires that for some  $\varepsilon, \delta \in (0, 1)$ ,  $\text{coherence}(A) \leq (\varepsilon^2/6)/\ln(3k/\delta)$ . Essentially, this condition ensures that the non-degeneracy of the component means is not isolated in just a few of the  $d$  dimensions. Operationally, it implies the following.

**Proposition 4.3.2** (Anandkumar *et al.*, 2012c). Assume  $A$  has rank  $k$ , and

$$\text{coherence}(A) \leq \frac{\varepsilon^2/6}{\ln(3k/\delta)}$$

for some  $\varepsilon, \delta \in (0, 1)$ . With probability at least  $1 - \delta$ , a random partitioning of the dimensions  $[d]$  into three groups (for each  $i \in [d]$ , independently pick  $t \in \{1, 2, 3\}$  uniformly at random and put  $i$  in group  $t$ ) has the following property. For each  $t \in \{1, 2, 3\}$  and  $j \in [k]$ , let  $\mu_{t,j}$  be the entries of  $\mu_j$  put into group  $t$ , and let  $A_t := [\mu_{t,1} | \mu_{t,2} | \cdots | \mu_{t,k}]$ . Then for each  $t \in \{1, 2, 3\}$ ,  $A_t$  has full column rank, and the  $k$ -th largest singular value of  $A_t$  is at least  $\sqrt{(1-\varepsilon)/3}$  times that of  $A$ .

Therefore, three asymmetric views can be created by randomly partitioning the observed random vector  $x$  into  $x_1, x_2$ , and  $x_3$ , such that the resulting component means for each view satisfy the conditions of Theorem 4.6.

### 4.3.2 Spherical Gaussian mixtures, revisited

Consider again the case of spherical Gaussian mixtures described in Section 4.2.1. The previous analysis in Theorems 4.2 and 4.3 can be used

when the observation dimension  $d \geq k$ , and the  $k$  component means are linearly independent. We now show that when the dimension is slightly larger, say greater than  $3k$ , a different (and simpler) technique based on the multi-view structure can be used to extract the relevant structure.

We again use a randomized reduction. Specifically, we create three views by (i) applying a random rotation to  $x$ , and then (ii) partitioning  $x \in \mathbb{R}^d$  into three views  $\tilde{x}_1, \tilde{x}_2, \tilde{x}_3 \in \mathbb{R}^{\tilde{d}}$  for  $\tilde{d} := d/3$ . By the rotational invariance of the multivariate Gaussian distribution, the distribution of  $x$  after random rotation is still a mixture of spherical Gaussians (*i.e.*, a mixture of product distributions), and thus  $\tilde{x}_1, \tilde{x}_2, \tilde{x}_3$  are conditionally independent given  $h$ . What remains to be checked is that, for each view  $t \in \{1, 2, 3\}$ , the matrix of conditional means of  $\tilde{x}_t$  for each view has full column rank. This is true with probability 1 as long as the matrix of conditional means  $A := [\mu_1 | \mu_2 | \dots | \mu_k] \in \mathbb{R}^{d \times k}$  has rank  $k$  and  $d \geq 3k$ . To see this, observe that a random rotation in  $\mathbb{R}^d$  followed by a restriction to  $\tilde{d}$  coordinates is simply a random projection from  $\mathbb{R}^d$  to  $\mathbb{R}^{\tilde{d}}$ , and that a random projection of a linear subspace of dimension  $k$  to  $\mathbb{R}^{\tilde{d}}$  is almost surely injective as long as  $\tilde{d} \geq k$ . Applying this observation to the range of  $A$  implies the following.

**Proposition 4.3.3** (Hsu and Kakade, 2013). Assume  $A$  has rank  $k$  and that  $d \geq 3k$ . Let  $R \in \mathbb{R}^{d \times d}$  be chosen uniformly at random among all orthogonal  $d \times d$  matrices, and set  $\tilde{x} := Rx \in \mathbb{R}^d$  and  $\tilde{A} := RA = [R\mu_1 | R\mu_2 | \dots | R\mu_k] \in \mathbb{R}^{d \times k}$ . Partition  $[d]$  into three groups of sizes  $d_1, d_2, d_3$  with  $d_t \geq k$  for each  $t \in \{1, 2, 3\}$ . Furthermore, for each  $t$ , define  $\tilde{x}_t \in \mathbb{R}^{d_t}$  (respectively,  $\tilde{A}_t \in \mathbb{R}^{d_t \times k}$ ) to be the subvector of  $\tilde{x}$  (resp., submatrix of  $\tilde{A}$ ) obtained by selecting the  $d_t$  entries (resp., rows) in the  $t$ -th group. Then  $\tilde{x}_1, \tilde{x}_2, \tilde{x}_3$  are conditionally independent given  $h$ ;  $\mathbb{E}[\tilde{x}_t | h = j] = \tilde{A}_t e_j$  for each  $j \in [k]$  and  $t \in \{1, 2, 3\}$ ; and with probability 1, the matrices  $\tilde{A}_1, \tilde{A}_2, \tilde{A}_3$  have full column rank.

It is possible to obtain a quantitative bound on the  $k$ -th largest singular value of each  $A_t$  in terms of the  $k$ -th largest singular value of  $A$  (analogous to Proposition 4.3.2). One avenue is to show that a random rotation in fact causes  $\tilde{A}$  to have low coherence, after which we can apply Proposition 4.3.2. With this approach, it is sufficient to require  $n = O(k \log k)$  (for constant  $\varepsilon$  and  $\delta$ ), which results in the  $k$ -th largest

singular value of each  $A_t$  being a constant fraction of the  $k$ -th largest singular value of  $A$ . We conjecture that, in fact,  $n \geq c \cdot k$  for some  $c > 3$  suffices.

### 4.3.3 Hidden Markov models

Our next example is the time-homogeneous Hidden Markov models (HMM) (Baum and Petrie, 1966) for sequences of vector-valued observations  $x_1, x_2, \dots \in \mathbb{R}^d$ . Consider a Markov chain of discrete hidden states  $y_1 \rightarrow y_2 \rightarrow y_3 \rightarrow \dots$  over  $k$  possible states  $[k]$ ; given a state  $y_t$  at time  $t$ , the random observation  $x_t \in \mathbb{R}^d$  at time  $t$  is independent of all other observations and hidden states. See Figure 4.1(b).

Let  $\pi \in \Delta^{k-1}$  be the initial state distribution (*i.e.*, the distribution of  $y_1$ ), and  $T \in \mathbb{R}^{k \times k}$  be the stochastic transition matrix for the hidden state Markov chain such that for all times  $t$ ,

$$\Pr[y_{t+1} = i | y_t = j] = T_{i,j}, \quad i, j \in [k].$$

Finally, let  $O \in \mathbb{R}^{d \times k}$  be the matrix whose  $j$ -th column is the conditional expectation of  $x_t$  given  $y_t = j$ : for all times  $t$ ,

$$\mathbb{E}[x_t | y_t = j] = Oe_j, \quad j \in [k].$$

**Proposition 4.3.4** (Anandkumar *et al.*, 2012c). Define  $h := y_2$ , where  $y_2$  is the second hidden state in the Markov chain. Then

- $x_1, x_2, x_3$  are conditionally independent given  $h$ ;
- the distribution of  $h$  is given by the vector  $w := T\pi \in \Delta^{k-1}$ ;
- for all  $j \in [k]$ ,

$$\mathbb{E}[x_1 | h = j] = O \operatorname{diag}(\pi) T^\top \operatorname{diag}(w)^{-1} e_j$$

$$\mathbb{E}[x_2 | h = j] = Oe_j$$

$$\mathbb{E}[x_3 | h = j] = OTe_j.$$

Note the matrix of conditional means of  $x_t$  has full column rank, for each  $t \in \{1, 2, 3\}$ , provided that: (i)  $O$  has full column rank, (ii)  $T$  is invertible, and (iii)  $\pi$  and  $T\pi$  have positive entries. Using the result of this proposition, we can formulate the problem as a multi-view mixture model and apply Theorem 4.6.



#### 4.4 Nonlinear Model: Noisy-Or Networks

The models we stated in the previous sections are all *linear* for the purpose of tensor decomposition; in particular, the observed moment tensors  $T$  have an exact decomposition with the rank-1 components as the desired parameters to be learned. This behavior is fairly common if given hidden components, the conditional expectation of the observation is a linear combination of different components, e.g., in the Latent Dirichlet Allocation model, if the document has a mixture of topics, the probabilities of observing different words are also linear mixtures.

In more complicated models, the observation may not be linear. In this section, we consider the *noisy-or* model, which is among the first non-linear models that can be learned by tensor decomposition.

The noisy-or model is a Bayes network with binary latent variables  $h \in \{0, 1\}^k$ , and binary observed variables  $x \in \{0, 1\}^d$ . The hidden variables are independent Bernoulli variables with parameter  $\rho$ , i.e.,  $\Pr[h_j = 1] = \rho, j \in [k]$ . The conditional distribution  $\Pr[x|h]$  is parameterized by a non-negative weight matrix  $W \in \mathbb{R}^{d \times k}$ . Conditioned on  $h$ , the observations  $x_1, \dots, x_d$  are independent with distribution

$$\Pr[x_i = 0 | h] = \prod_{j=1}^k \exp(-W_{ij}h_j) = \exp(-\langle W^i, h \rangle), \quad i \in [d], \quad (4.1)$$

where  $W^i$  denotes the  $i$ -th row of  $W$ . This model is often used to model the relationship between diseases and symptoms, as in the classical human-constructed tool for medical diagnosis called *Quick Medical Reference* (QMR-DT) by Shwe *et al.* (1991). In this case, the latent variables  $h_j$ 's are diseases and observed variables  $x_i$ 's are symptoms. We see that  $1 - \exp(-W_{ij}h_j)$  can be thought of as the probability that disease  $h_j$  activates symptom  $x_i$ , and  $x_i$  is activated if one of  $h_j$ 's activates it. This also explains the name of the model, noisy-or.

Given (4.1) and the independence of different  $x_i$ 's given  $h$ , we have

$$\Pr[x | h] = \prod_{i=1}^d \left(1 - \exp(-\langle W^i, h \rangle)\right)^{x_i} \left(\exp(-\langle W^i, h \rangle)\right)^{1-x_i}.$$

Contrasting with the linear models in the previous sections, we see that

under this model when a patient has multiple diseases, the expectation of the symptoms  $x$  is *not* a linear combination of different components.

**Point-wise Mutual Information:** Since the conditional probability is a product of  $d$  terms, it is natural to consider taking log in order to convert it into a summation. This motivates the use of Point-wise Mutual Information (PMI), which is a common metric for the correlations between two events. Given events  $X$  and  $Y$ , the PMI is defined as

$$\text{PMI}(X, Y) := \log \frac{\Pr[X, Y]}{\Pr[X] \Pr[Y]}.$$

Intuitively, if  $X, Y$  are independent, then  $\text{PMI} = 0$ ; if they are positively correlated, then  $\text{PMI} > 0$ ; if they are negatively correlated, then  $\text{PMI} < 0$ . This can be also generalized to three random variables as

$$\text{PMI3}(X, Y, Z) := \log \frac{\Pr[X, Y] \Pr[Y, Z] \Pr[X, Z]}{\Pr[X, Y, Z] \Pr[X] \Pr[Y] \Pr[Z]}.$$

For the noisy-or networks, we use PMI and PMI3 as the (generalized) moments that we observe. More precisely, we define the following PMI matrix  $M \in \mathbb{R}^{d \times d}$  and PMI tensor  $T \in \mathbb{R}^{d \times d \times d}$  as

$$\begin{aligned} M_{i_1, i_2} &:= \text{PMI}(1 - x_{i_1}, 1 - x_{i_2}), \quad i_1, i_2 \in [d]. \\ T_{i_1, i_2, i_3} &:= \text{PMI3}(1 - x_{i_1}, 1 - x_{i_2}, 1 - x_{i_3}), \quad i_1, i_2, i_3 \in [d]. \end{aligned}$$

These tabulate the correlations among all pairs and triples of symptoms; more specifically, they incorporate indicator random variable for the symptom being absent.

As before, we would like to have a low rank decomposition for these observed matrix and tensor. This is almost true except for some small perturbations as follows. For convenience, we define  $F, G \in \mathbb{R}^{d \times k}$  as

$$\begin{aligned} F &:= 1 - \exp(-W) \\ G &:= 1 - \exp(-2W). \end{aligned}$$

Using these quantities we can approximately represent the PMI matrix and tensor in low rank forms.

**Proposition 4.1** (Proposition 2.1 in Arora *et al.*, 2017). Let  $F_j, G_j \in \mathbb{R}^d$  denote the  $j$ -th columns of the above matrices  $F, G$ , respectively. Then we have

$$M \approx \rho \left( FF^\top + \rho GG^\top \right) = \rho \sum_{j=1}^k F_j F_j^\top + \rho^2 \sum_{j=1}^k G_j G_j^\top$$

$$T \approx \rho \sum_{j=1}^k F_j \otimes F_j \otimes F_j + \rho^2 \sum_{j=1}^k G_j \otimes G_j \otimes G_j.$$

The approximation in both equations are due to higher order terms in the Taylor expansions and are dominated by the term with  $G$ . Recall that  $\rho$  is the probability of any disease  $h_j$  being present, and therefore, for this application we expect  $\rho$  to be small. Hence, the terms with  $F$  are much larger than the terms with  $G$  and we can say that applying tensor decomposition approximately recovers columns of  $F$ . Several ideas and many more details are required in analyzing the effect of the perturbation  $G$  (since  $G$  is not as small as required in Theorem 3.4); we refer interested readers to Arora *et al.*, 2017.

## 4.5 Applications in Supervised Learning

In this section, we describe how tensor methods can be also used in supervised learning applications contrasting with the unsupervised problems that we have described so far. In particular, we focus on neural networks to elaborate on this application of tensor methods. This is a very interesting extension given the vast applications of neural networks that have significantly improved predictive performance across multiple domains such as computer vision and speech recognition with rapidly growing influence in many other areas. Although we focus on neural networks in this monograph, similar tensor techniques are also applied to learning mixtures of generalized linear models in Sedghi *et al.* (2016). Most of the discussions and results in the rest of this section are borrowed from Janzamin *et al.* (2015).

In previous sections, we discussed the application of tensor methods for learning latent variable models and latent representations which was performed in an unsupervised manner. Thus, when considering supervised learning tasks such as training neural networks, the first natural

and major question that we have to answer is how to adapt these tensor methods to supervised learning. To answer this, we incorporate a generative approach in the problem and propose non-linear transformation of the input which is basically new features extracted from the input. We refer to this new transformation as *score function* of the input. These new extracted features enable us to formulate the problem of training neural networks as the tensor decomposition problem. More concretely, we show that the cross-moment between output and the score function of the input has information about the weight parameters of the neural network in its rank-1 components. Before providing more details, we first elaborate more on exploiting a generative model, and in particular, the score functions.

**Generative vs. discriminative models:** Generative models incorporate a joint distribution  $p(x, y)$  over both the input  $x$  and label  $y$ . On the other hand, discriminative models such as neural networks only incorporate the conditional distribution  $p(y|x)$ . While training neural networks for general input  $x$  is NP-hard, does knowledge about the input distribution  $p(x)$  make learning tractable?

Here, we assume knowledge of the input density  $p(x)$  which can be any continuous differentiable function. While unsupervised learning problem of estimation of density  $p(x)$  is itself a hard problem for general models, here we investigate how  $p(x)$  can be exploited to make training of neural networks tractable. The knowledge of  $p(x)$  is naturally available in the *experimental design* framework, where the person designing the experiments has the ability to choose the input distribution. Examples include conducting polling, carrying out drug trials, collecting survey information, and so on.

We utilize the knowledge about the input density  $p(x)$  (up to normalization)<sup>1</sup> to obtain certain (non-linear) transformations of the input, given by the class of score functions. *Score functions* are normalized derivatives of the input pdf; see (4.3). If the input is a vector (the typical case), the first order score function (i.e., the first derivative) is a

---

<sup>1</sup>We do not require the knowledge of the normalizing constant or the partition function, which is  $\#P$  hard to compute (Wainwright and Jordan, 2008).

vector, the second order score is a matrix, and the higher order scores are tensors.

### 4.5.1 Moment tensor structure in neural networks

We consider a neural network with one hidden layer of dimension  $k$ . Let the output  $\tilde{y} \in \{0, 1\}$  be the binary label, and  $x \in \mathbb{R}^d$  be the feature (input) vector; see Janzamin *et al.* (2015) for generalization to higher dimensional output (multi-label and multi-class), and also the continuous output case. We consider the label generating model

$$\tilde{f}(x) := \mathbb{E}[\tilde{y}|x] = \langle a_2, \sigma(A_1^\top x + b_1) \rangle + b_2, \quad (4.2)$$


where  $\sigma(\cdot)$  is a (linear/nonlinear) element-wise function named as activation function; see Figure 4.2 for a schematic representation of label-function in (4.2) in the general case of vector output  $\tilde{y}$ .

In this section, we only focus on the *realizable* setting. In this setting, the goal is to learn the parameters of the neural network specified in (4.2), i.e., to learn the weight matrices (vectors)  $A_1 \in \mathbb{R}^{d \times k}$ ,  $a_2 \in \mathbb{R}^k$  and bias vectors  $b_1 \in \mathbb{R}^k$ ,  $b_2 \in \mathbb{R}$ , given labeled data samples  $\{(x_i, \tilde{y}_i)\}$ . This only involves the estimation analysis where we have a label-function  $\tilde{f}(x)$  specified in (4.2) with fixed unknown parameters  $A_1, b_1, a_2, b_2$ , and we would like to learn these parameters and finally bound the overall function estimation error  $\mathbb{E}_x[|\tilde{f}(x) - \hat{f}(x)|^2]$ , where  $\hat{f}(x)$  is the estimation of fixed neural network  $\tilde{f}(x)$  given finite samples. The approximation<sup>2</sup> analysis and consequently the risk bound is out of the focus of this monograph and the interested reader is referred to read Janzamin *et al.* (2015) for details.

We are now ready to explain how learning the parameters of two-layer neural network in Figure 4.2 can be characterized as a tensor decomposition algorithm. Note that we only provide the tensor decomposition algorithm for learning the parameters of first layer and as described in Janzamin *et al.* (2015), the bias parameter in first layer is learned using a Fourier method and the parameters of second layer are learned using linear regression. These parts are not within the focus of

---

<sup>2</sup>Here by approximation we mean how accurate the neural network can approximate any arbitrary function  $f(x)$ .



**Figure 4.2:** Graphical representation of a neural network,  $\mathbb{E}[\hat{y}|x] = A_2^\top \sigma(A_1^\top x + b_1) + b_2$ .

this monograph and we refer the reader to Janzamin *et al.* (2015) for more details. Note that most of the unknown parameters (compare the dimensions of matrix  $A_1$ , vectors  $a_2$ ,  $b_1$ , and scalar  $b_2$ ) are estimated in the first part, and thus, the tensor decomposition method for estimating  $A_1$  is the main part of the learning algorithm.

In order to provide the tensor structure, we first define the score functions as follows.

**Score function**

The  $m$ -th order score function  $\mathcal{S}_m(x) \in \otimes^m \mathbb{R}^d$  is defined as (Janzamin *et al.*, 2014)

$$\mathcal{S}_m(x) := (-1)^m \frac{\nabla_x^{(m)} p(x)}{p(x)}, \tag{4.3}$$

where  $p(x)$  is the probability density function of random vector  $x \in \mathbb{R}^d$ , and  $\nabla_x^{(m)}$  denotes the  $m$ -th order derivative operator defined as follows. For function  $g(x) : \mathbb{R}^d \rightarrow \mathbb{R}$  with vector input  $x \in \mathbb{R}^d$ , the  $m$ -th order derivative w.r.t. variable  $x$  is denoted by  $\nabla_x^{(m)} g(x) \in \otimes^m \mathbb{R}^d$  (which is a  $m$ -th order tensor) such that

$$\left[ \nabla_x^{(m)} g(x) \right]_{i_1, \dots, i_m} := \frac{\partial g(x)}{\partial x_{i_1} \partial x_{i_2} \dots \partial x_{i_m}}, \quad i_1, \dots, i_m \in [d]. \tag{4.4}$$

When it is clear from the context, we drop the subscript  $x$  and write the derivative as  $\nabla^{(m)} g(x)$ .

The main property of score functions as yielding differential operators that enables us to estimate the weight matrix  $A_1$  via tensor decomposition is discussed in next subsection; see Equation (4.5).

Note that in this framework, we assume access to a sufficiently good approximation of the input pdf  $p(x)$  and the corresponding score functions  $\mathcal{S}_2(x)$ ,  $\mathcal{S}_3(x)$ . Indeed, estimating these quantities in general is a hard problem, but there exist numerous instances where this becomes tractable. Examples include spectral methods for learning latent variable models such as Gaussian mixtures, topic or admixture models, independent component analysis (ICA) and so on that we discussed in the previous sections. Moreover, there have been recent advances in non-parametric score matching methods (Sriperumbudur *et al.*, 2013) for density estimation in infinite dimensional exponential families with guaranteed convergence rates. These methods can be used to estimate the input pdf in an unsupervised manner. Below, we discuss more about score function estimation methods. In this work, we focus on how we can use the input generative information to make training of neural networks tractable. We refer the interested reader to Janzamin *et al.* (2015) for more discussions on this.

**Estimation of score function** There are various efficient methods for estimating the score function. The framework of score matching is popular for parameter estimation in probabilistic models (Hyvärinen, 2005; Swersky *et al.*, 2011), where the criterion is to fit parameters based on matching the data score function. Swersky *et al.* (2011) analyze the score matching for latent energy-based models. In deep learning, the framework of auto-encoders attempts to find encoding and decoding functions which minimize the reconstruction error under added noise; the so-called Denoising Auto-Encoders (DAE). This is an unsupervised framework involving only unlabeled samples. Alain and Bengio (2012) argue that the DAE approximately learns the first order score function of the input, as the noise variance goes to zero. Sriperumbudur *et al.* (2013) propose non-parametric score matching methods for density estimation in infinite dimensional exponential families with guaranteed convergence rates. Therefore, we can use any of these methods for

estimating  $\mathcal{S}_1(x)$  and use the recursive form (Janzamin *et al.*, 2014)

$$\mathcal{S}_m(x) = -\mathcal{S}_{m-1}(x) \otimes \nabla_x \log p(x) - \nabla_x \mathcal{S}_{m-1}(x)$$

to estimate higher order score functions. Despite the existence of these techniques, there still exist so much room for proposing methods to efficiently estimate score functions.

### Tensor form of the moment

The score functions are new representations (extracted features) of input data  $x$  that can be used for training neural networks. The score functions have the property of yielding differential operators with respect to the input distribution. More precisely, for label-function  $\tilde{f}(x) := \mathbb{E}[\tilde{y}|x]$ , Janzamin *et al.* (2014) show that

$$\mathbb{E}[\tilde{y} \cdot \mathcal{S}_3(x)] = \mathbb{E}[\nabla_x^{(3)} \tilde{f}(x)]. \quad (4.5)$$

Now for the neural network output in (4.2), note that the function  $\tilde{f}(x)$  is a non-linear function of both input  $x$  and weight matrix  $A_1$ . The expectation operator  $\mathbb{E}[\cdot]$  averages out the dependency on  $x$ , and the derivative acts as a *linearization operator* as follows. In the neural network output (4.2), we observe that the columns of weight vector  $A_1$  are the linear coefficients involved with input variable  $x$ . When taking the derivative of this function, by the chain rule, these linear coefficients show up in the final form. With this intuition, we are now ready to provide the precise form of the moment where we show how the cross-moment between label and score function as  $\mathbb{E}[\tilde{y} \cdot \mathcal{S}_3(x)]$  leads to a tensor decomposition form for estimating weight matrix  $A_1$ :

**Lemma 4.7** ((Janzamin *et al.*, 2015)). For the two-layer neural network specified in (4.2), we have

$$\mathbb{E}[\tilde{y} \cdot \mathcal{S}_3(x)] = \sum_{j \in [k]} \lambda_j \cdot (A_1)_j \otimes (A_1)_j \otimes (A_1)_j, \quad (4.6)$$

where  $(A_1)_j \in \mathbb{R}^d$  denotes the  $j$ -th column of  $A_1$ , and

$$\lambda_j = \mathbb{E}[\sigma'''(z_j)] \cdot a_2(j), \quad (4.7)$$

for vector  $z := A_1^\top x + b_1$  as the input to the nonlinear operator  $\sigma(\cdot)$ .



This is proved by the main property of score functions as yielding differential operators that was described earlier. This lemma shows that by decomposing the cross-moment tensor  $\mathbb{E}[\tilde{y} \cdot \mathcal{S}_3(x)]$ , we can recover the columns of  $A_1$ . This clarifies how the score function acts as a linearization operator while the final output is nonlinear in terms of  $A_1$ .

## 4.6 Other Models

Tensor decompositions have been applied to learn many other models. Several ideas we introduced in this section originated from more complicated settings, and can be applied to more models. Here we give hints to more examples, but the list is by no means complete.

The idea of manipulating moments was well-known in the ICA (Independent Component Analysis) literature, where cumulants are used frequently instead of moments. For other distributions, it was used for the Latent Dirichlet Allocation (Anandkumar *et al.*, 2012a), and widely applied in all the models where the hidden variables are not categorical.

The multi-view model was first used in Mossel and Roch (2006) to learn Hidden Markov Models and Phylogeny Tree Reconstruction. The original technique in Mossel and Roch (2006) was based on spectral algorithms and was not viewed as tensor decomposition, however it is very similar to the simultaneous diagonalization algorithm we introduced in Section 3.6.

Tensor methods can be also applied to learning more complicated mixtures of Gaussians, where each component may have a different, non-spherical component (Ge *et al.*, 2015b). The covariance matrix creates many technical problems which is beyond the scope of this monograph. The idea of creating different views as we discussed in Section 4.3 can be also applied to learning community models (Anandkumar *et al.*, 2013).

Tensor decomposition is also particularly useful in the context of deep neural networks, most notably with the aim to speed up computation. One way to do so is to apply tensor factorization to the kernel of convolutional layers (Tai *et al.*, 2016). In particular, by applying CP decomposition to the convolutional kernel of a pre-trained network,

not only is it possible to reduce the number of parameters, but it also gives a way of re-expressing the convolution in terms of a series of smaller and more efficient convolutions. Both ALS (Lebedev *et al.*, 2015) and tensor power method (Astrid and Lee, 2017) have been considered. This process typically results in a performance deterioration which is restored by fine-tuning. A similar result can be obtained using Tucker decomposition (Kim *et al.*, 2016). It is possible to go further and jointly parameterize multiple layers or whole networks, resulting in large parameter space savings without loss of performance (Kossaifi *et al.*, 2019a).

We can also preserve the multi-linear structure in the activation tensor, using tensor contraction (Kossaifi *et al.*, 2017), or by removing fully connected layers and flattening layers altogether and replacing with tensor regression layers (Kossaifi *et al.*, 2018). Adding a stochastic regularization on the rank of the decomposition can also help render the models more robustly (Kolbeinsson *et al.*, 2019). *Tensorization* can be also leveraged by applying it to the weight matrix of fully-connected layers (Novikov *et al.*, 2015).

# 5

---

## Practical Implementations

---

We have so far covered many aspects of tensors including tensor decomposition and how they are useful in learning different machine learning models in both supervised and unsupervised settings. In this section, we discuss practical implementation of tensor operations using Python programming language. We first motivate our choice and introduce some actual code to perform tensor operations and tensor decomposition. We then briefly show how to perform more advanced tensor operations using TensorLy (Kossaifi *et al.*, 2019b), a library for tensor learning in Python. Finally, we show how to scale up our algorithms using the PyTorch deep learning framework (Paszke *et al.*, 2017) as a backend for TensorLy.

### 5.1 Programming Language and Framework

Throughout this section, we present the implementations in Python language. Python is a multi-purpose and powerful programming language that is emerging as the prime choice for Machine Learning and data science. Its readability allows us to focus on the underlying concepts we are implementing without getting distracted by low-level considerations such as memory handling or obscure syntax. Its huge popularity

means that good libraries exist to solve most of our computational needs. In particular, NumPy (Walt *et al.*, 2011) is an established and robust library for numerical computation. It offers a high performance structure for manipulating multi-dimensional arrays. TensorLy builds on top of this and provides a simple API for fast and easy tensor manipulation. TensorLy has a system of backends that allows you to switch transparently from NumPy to PyTorch, MXNet, TensorFlow, etc. This means you can perform any of the operations seamlessly on all these frameworks. In particular, using a deep learning framework such as PyTorch as backend, it is easy to scale operations to GPUs and multi-machines.

### 5.1.1 Pre-requisite

In order to run the codes presented in this section, you will need a working installation of Python 3.0, along with NumPy (for the numerical array structure), SciPy (Jones *et al.*, 2001) (for scientific python), and optionally Matplotlib (Hunter, 2007) for visualization.

The easiest way to get all these is to install the Anaconda distribution (<https://anaconda.org/>) which comes with all the above bundled and pre-compiled so you do not have to do anything else!

## 5.2 Tensors as NumPy Arrays

You may recall from Section 3.1 that tensors can be identified as multi-dimensional arrays. Therefore, we represent tensors as NumPy arrays, which are multi-dimensional arrays.

Let's take as an example a tensor  $T \in \mathbb{R}^{3 \times 4 \times 2}$ , defined by the following frontal slices:

$$T(:, :, 1) = \begin{bmatrix} 0 & 2 & 4 & 6 \\ 8 & 10 & 12 & 14 \\ 16 & 18 & 20 & 22 \end{bmatrix}$$

and

$$T(:, :, 2) = \begin{bmatrix} 1 & 3 & 5 & 7 \\ 9 & 11 & 13 & 15 \\ 17 & 19 & 21 & 23 \end{bmatrix}$$

In NumPy we can instantiate new arrays from nested lists of values. For instance, matrices are represented as a list of rows, where each row is itself a list. Let’s define the slices of above tensor  $T$  as 2-D NumPy arrays:

---

```

1 # We first import numpy
2 import numpy as np
3
4 # First frontal slice
5 T1 = np.array([[ 0.,  2.,  4.,  6.],
6               [ 8., 10., 12., 14.],
7               [16., 18., 20., 22.]])
8
9 # Second frontal slice
10 T2 = np.array([[ 1.,  3.,  5.,  7.],
11               [ 9., 11., 13., 15.],
12               [17., 19., 21., 23.]])

```

---

Let’s now write a function that stacks these frontal slices into a third order tensor:

---

```

1 def tensor_from_frontal_slices(*matrices):
2     """Creates a tensor from its frontal slices
3
4     Parameters
5     Parameters
6     matrices : 2D-Numpy arrays
7
8     Returns
9     Returns
10    tensor : 3D-Numpy arrays
11            its frontal slices are the matrices passed as input
12    """
13    return np.concatenate([matrix[:, :, np.newaxis]\
14                           for matrix in matrices], axis=-1)

```

---

We can then build the full tensor  $T$  from its frontal slices  $T1$  and  $T2$  created above:

---

```

1 T = tensor_from_frontal_slices(T1, T2)

```

---

We can inspect the frontal slices naturally using almost the same notation as in the math. To do so we fix the last index while iterating over other modes (using ':').

---

```

1 >>> T[:, :, 0]
2 array([[ 0.,  2.,  4.,  6.],
3        [ 8., 10., 12., 14.],
4        [16., 18., 20., 22.]])
5 >>> T[:, :, 1]
6 array([[ 1.,  3.,  5.,  7.],
7        [ 9., 11., 13., 15.],
8        [17., 19., 21., 23.]])

```

---

Remember that in NumPy (and generally, in Python), like in C, indexing starts at zero. In the same way, you can also inspect the horizontal slices (by fixing the first index) and lateral slices (by fixing the second index).

Similarly, we can easily inspect the fibers which, as you may recall, are higher-order analogues to column and rows. We can obtain the fibers of  $T$  by fixing all indices but one:

---

```

1 # First column (mode-1 fiber)
2 >>> T[:, 0, 0]
3 array([ 0.,  8., 16.])
4
5 # First row (mode-2 fiber)
6 >>> T[0, :, 0]
7 array([ 0.,  2.,  4.,  6.])
8
9 # First tube (mode-3 fiber)
10 >>> T[0, 0, :]
11 array([ 0.,  1.])

```

---

Finally, you can access the size of a tensor via its *shape*, which indicates the size of the tensor along each of its modes. For instance, our tensor  $T$  has shape  $(3, 4, 2)$ :

---

```

1 >>> T.shape
2 (3, 4, 2)

```

---

---

**Procedure 9** Tensor unfolding
 

---

**input** Tensor  $T$  of shape  $(d_1, d_2, \dots, d_n)$ ; unfolding mode  $m$ .

**output** Mode- $m$  matricization (unfolding)

- 1: Move the  $m^{\text{th}}$  dimension to the first position.
  - 2: Reshape into a matrix  $M$  of shape  $(d_m, \prod_{k \neq m} d_k)$ .
  - 3: **return**  $M$ .
- 

### 5.3 Basic Tensor Operations and Decomposition

Tensor matricization, or unfolding, as introduced in Equation (3.2) and described in Procedure 9 naturally translates into Python. One important consideration when implementing algorithms that manipulate tensors is the way elements are organised in memory. You can think of the memory as one long vector of numbers. Because of the way CPU and GPU operate, it matters how these elements are layered in the memory. To store a matrix, for instance, we can either organise the elements row-after-row (also called C-ordering) or column-after-column (also called Fortran ordering). In NumPy, elements are organised by default in row-order, same for PyTorch. It so happens that the definition of the unfolding we use is adapted for such ordering, thus avoiding expensive reordering of the data.

As a result, matricization (or unfolding) of a tensor along a given mode simplifies to moving that mode to the front and reshaping into a matrix as also described in Procedure 9.

---

```

1 def unfold(tensor, mode):
2     """Returns unfolding of a tensor — modes starting at 0.
3
4     Parameters
5     tensor : ndarray
6     mode : int (default is 0), mode along which to unfold
7
8     Returns
9     ndarray
10    unfolded_tensor of shape
11    """
12    return np.reshape(np.moveaxis(tensor, mode, 0),
13                      (tensor.shape[mode], -1))
14
15

```

---

Folding the tensor is done by performing the inverse operations: we first reshape the matrix into a tensor and move back the first dimension to its original position.

---

```

1 def fold(unfolded_tensor, mode, shape):
2     """Refolds the unfolded tensor into a full tensor.
3         In other words, refolds the n-mode unfolded tensor
4         into the original tensor of the specified shape.
5
6     Parameters
7     -----
8     unfolded_tensor : ndarray
9         unfolded tensor of shape ``(shape[mode], -1)``
10    mode : int
11        the mode of the unfolding
12    shape : tuple
13        shape of the original tensor before unfolding
14
15    Returns
16    -----
17    ndarray
18        folded_tensor of shape `shape`
19    """
20    full_shape = list(shape)
21    mode_dim = full_shape.pop(mode)
22    full_shape.insert(0, mode_dim)
23    return np.moveaxis(np.reshape(
24        unfolded_tensor, full_shape), 0, mode)

```

---

### 5.3.1 CP decomposition

Now that we know how to manipulate tensors using NumPy arrays, we are ready to implement a simple version of the CP decomposition via Alternating Least Squares, as explained in Section 3.7. We will start by writing the auxiliary functions we need in the main algorithm.

CP decomposition expresses its input tensor as a sum of outer products of vectors; see Equation (3.6) for the definition. Taking the unfolded expression, there is a useful equivalent formulation that uses the Khatri-Rao product which we used in Equation (3.31). In particular, for vectors  $u, v, w$ , we have

$$\text{mat}(u \otimes v \otimes w, 1) = u \cdot (v \odot w)^\top.$$

Note that here,  $\text{mat}(\cdot, 1)$  corresponds to unfolding along mode 0 in our code.



Let's first write a function to take the Khatri-Rao product of two matrices, as defined in equation (3.9). A naive, literal implementation of that equation could be as follows:

---

```

1 def naive_khatri_rao(A, B):
2     # Both matrices must have the same number of columns k
3     d1, k = A.shape
4     d2, k = B.shape
5
6     # The khatri-rao product has size d1d2 x k
7     C = np.zeros((d1*d2, k))
8     for i in range(d1):
9         for l in range(d2):
10            for j in range(k):
11                # Indexing starts at 0!
12                C[l + i*d1, j] = A[i, j]*B[l, j]
13     return C

```

---

However, loops are typically slow in Python and this naive implementation is as a result extremely slow. By contrast, we can use the built-in `einsum` function from NumPy, which uses Einstein's notation to define the operation, to write a *vectorized* version. This results in a much more efficient function:

---

```

1 def khatri_rao(matrix1, matrix2):
2     """Returns the khatri-rao product of matrix1 and matrix2
3     """
4     n_columns = matrix1.shape[1]
5     result = np.einsum('ij,lj->ilj', matrix1, matrix2)
6     return result.reshape((-1, n_columns))

```

---

Recall that the khatri-rao takes a column-wise Kronecker product of two matrices with the same number of columns. The `einsum` function here expresses this idea in terms of indices, where  $A$  is indexed by  $i$  and  $j$  and  $B$  is indexed by  $l$  and  $j$ . The output is of size  $ilj$  and we simply have to reshape it into a matrix of the appropriate size.

Then, given a third order tensor in its Kruskal form (i.e., a decomposed tensor, expressed as a series of factors  $A, B$  and  $C$  with unit norm and the associated vector of coefficients  $\lambda$  implying the norms), we need a method to return the reconstruction  $T = \sum_{j \in [k]} \lambda_j a_j \otimes b_j \otimes c_j$ . Using the above matricization property, this reconstruction can also be written in its unfolded form as  $mat(T, 1) = \tilde{A} \cdot \text{diag}(\tilde{\lambda}) \cdot (\tilde{B} \odot \tilde{C})^\top$ ,

resulting in the following function:

---

```

1 def kruskal_to_tensor(weights, A, B, C):
2     """Converts the kruskal form into a tensor
3     """
4     full_shape = (A.shape[0], B.shape[0], C.shape[0])
5     # Reconstruct in unfolded form
6     unfolded_tensor = np.dot(A.dot(np.diag(weights)),
7                             khatri_rao(B, C).T)
8     # Fold back to a tensor
9     return fold(unfolded_tensor, 0, full_shape)

```

---

To measure convergence, we can use, for instance, the Frobenius norm of the reconstruction error. Recall that the Frobenius norm is simply the square root of the sum of the squared elements of the tensor. This can be written in NumPy as :

---

```

1 def frobenius_norm(tensor):
2     """Frobenius norm of the tensor
3     """
4     return np.sqrt(np.sum(tensor**2))

```

---

We are now ready to implement the Alternating Least Squares method for Tensor Decomposition described in Algorithm 8.

---

```

1 def parafac(tensor, rank, l2_reg=1e-3, n_iter_max=200, tol=1e-10):
2     """CANDECOMP/PARAFAC decomposition via ALS
3
4     Parameters
5     -----
6     tensor : ndarray
7     rank   : int
8             number of components
9     l2_reg : float, default is 0.1
10            regularization parameter (\alpha)
11     n_iter_max : int
12                maximum number of iterations
13     tol : float, optional
14           tolerance: the algorithm stops when the variation in
15           the reconstruction error is less than the tolerance
16     verbose : int, optional
17              level of verbosity
18
19     Returns
20     -----
21     weights, A, B, C : weights, factors of the decomposition
22     """
23     # Initialize the factors of the decomposition randomly
24     A = np.random.random_sample((tensor.shape[0], rank))
25     B = np.random.random_sample((tensor.shape[1], rank))

```

```

26     C = np.random.random_sample((tensor.shape[2], rank))
27
28     # Norm of the input tensor
29     norm_tensor = frobenius_norm(tensor)
30     error = None
31
32     # Initialize the weights to 1
33     weights = np.ones(rank)
34
35     # Avoid division by zero
36     eps = 1e-12
37
38     # Regularization term \alpha*I
39     regularization = np.eye(rank)*l2_reg
40
41     for iteration in range(n_iter_max):
42         # Update A
43         prod = B.T.dot(B)*C.T.dot(C) + regularization
44         factor = unfold(tensor, 0).dot(khatri_rao(B, C))
45         A = np.linalg.solve(prod.T, factor.T).T
46         # Normalization (of the columns) of A
47         weights = np.linalg.norm(A, ord=2, axis=0)
48         A /= (weights[None, :] + eps)
49
50         # Update B
51         prod = A.T.dot(A)*C.T.dot(C) + regularization
52         factor = unfold(tensor, 1).dot(khatri_rao(A, C))
53         B = np.linalg.solve(prod.T, factor.T).T
54         # Normalization of B
55         weights = np.linalg.norm(B, ord=2, axis=0)
56         B /= (weights[None, :] + eps)
57
58         # Update C
59         prod = A.T.dot(A)*B.T.dot(B) + regularization
60         factor = unfold(tensor, 2).dot(khatri_rao(A, B))
61         C = np.linalg.solve(prod.T, factor.T).T
62         # Normalization of C
63         weights = np.linalg.norm(C, ord=2, axis=0)
64         C /= (weights[None, :] + eps)
65
66         # Compute the reconstruction error
67         prev_error = error
68         rec = kruskal_to_tensor(weights, A, B, C)
69         error = frobenius_norm(tensor - rec) / norm_tensor
70
71         if iteration > 1:
72             if tol and abs(prev_error - error) < tol:
73                 print('converged in {} iterations.'.format(
74                     iteration))
75                 break
76
77     return weights, A, B, C

```

Using our previously introduced tensor  $T$  as an example, we can verify that our algorithm indeed does what it is supposed to:

---

```

1 # decompose T into factors using CP
2 weights, A, B, C = parafac(T, 3)

```

---

```

3
4 # reconstruct the full tensor from these
5 rec = kruskal_to_tensor(weights, A, B, C)
6
7 # verify that the reconstruction is correct
8 np.testing.assert_array_equal(np.round(rec), T)

```

---

Let's now go over some aspects of the algorithm we just wrote, in particular, how we integrated unit-norm constraints on the columns of the factor, as well as  $\ell_2$  regularization.

**Normalization:** Within the CP decomposition method, after updating each factor, we further normalize it by dividing each column by its norm, as also done in equation (3.32). For example, for the first factor matrix A, we have:

---

```

1 # First, update the factor as previously
2 A = np.linalg.solve(prod.T, factor.T).T
3
4 # Normalize the columns
5 weights = np.linalg.norm(A, ord=2, axis=0)
6 A /= (weights[None, :] + eps)

```

---

We do similar normalization for the other two factor matrices B and C. Note that we have also added a tiny value `eps` to the normalization, where `eps` is defined as  $10^{-12}$ , which is close to machine precision. This additional term is used to avoid any division by zero. Note that, here, we are using float64, which has a machine epsilon of about  $10^{-15}$ , this would have to be adapted when changing the data type (e.g. to float32).

**Broadcasting:** In the last line of the update of A, the expression `weights[None, :]` is equivalent to `weights[np.newaxis, :]`. In other words, we add a dimension (of 1) to `weights`, and consider it as a matrix of size (1, `rank`) rather than a vector of length `rank`. This allows us to use *broadcasting*: `weight` is broadcasted to the same shape as the factor without actually duplicating the memory. This results in an efficient *vectorized* operation which divides each element of each column of the factor by the norm of that column.

This concept of broadcasting can also be used to simplify our

`kruskal_to_tensor` by replacing the matrix multiplication of the first factor `A` and `diag(weights)` with a simple element-wise multiplication:

---

```

1 def kruskal_to_tensor(weights, A, B, C):
2     """Converts the kruskal form into a tensor
3     """
4     full_shape = (A.shape[0], B.shape[0], C.shape[0])
5     # The main difference: we incorporate the weights
6     unfolded_tensor = np.dot(A*weights[np.newaxis, :],
7                               khatri_rao(B, C).T)
8     return fold(unfolded_tensor, 0, full_shape)

```

---

**Regularization:** In section 3.7, we also introduced an  $\ell_2$  regularized version of the ALS. The difference with the unregularized version is an additional term in the pseudo-inverse in the ALS updates; see Equation (3.33). Considering a regularization parameter  $\alpha = \text{12\_reg}$ , the update for factor `A` changes by the addition of a weighted identity matrix  $\alpha I$  to the product  $(\tilde{B} \odot \tilde{C})^\top (\tilde{B} \odot \tilde{C})$  and similarly for `B` and `C`. In the code, `np.eye(rank)` is the identity matrix of size `rank`  $\times$  `rank`.

## 5.4 Example: Image Compression via Tensor Decomposition

We now use our function to compress an image. We use as an example an image of a raccoon that comes shipped in with the SciPy library.

---

```

1 from scipy.misc import face
2
3 # Load the face
4 image = face()
5
6 # Convert it to a tensor of floats
7 image = np.array(image, dtype=np.float64)
8
9 # Check the size of the image
10 print(image.shape)
11 # (768, 1024, 3)

```

---

Our image is a third order tensor of shape (height, width, 3), the last mode corresponding to the RGB channels (Red, Green, Blue), the way colors are encoded on your computer. You can see the original image in Figure 5.1, in this case with a height of 768 and a width of



**Figure 5.1:** Our beautiful guinea pig, which happens to be a raccoon. On the left, the original image, and on the right, the reconstructed image from the factors of the decomposition, with a rank 50.

1024.

To visualize the tensor, we need a helper function to convert tensors of floats (typically stored into 64 bits) into an image, which consists of values stored into 8 bits. Here, a simple conversion suffices since the image already has a dynamic range between 0 and 255 as it was originally stored in 8 bits. If the image had a high dynamic range (higher than 255) then a more complex transformation (*tone mapping*) such as histogram equalization would be needed.

---

```
1 def to_image(tensor):
2     """convert a tensor of float values into an image
3     """
4     tensor -= tensor.min()
5     tensor /= tensor.max()
6     tensor *= 255
7     return tensor.astype(np.uint8)
```

---

This type of conversion, called tone mapping, can be much more complex than this simple conversion. Since we have a dynamic range between 0 and 255, it is appropriate here, but in general, when converting an image from 32 bits to just 8, we might want to use more complex techniques such as histogram normalization.

Now that the image is loaded in memory, we can apply our decomposition method to it, and build a reconstructed image from the compressed version, i.e., the factors of the decomposition,

---

```

1 # Apply CP decomposition
2 weights, A, B, C = factors = parafac(image, rank=50, tol=10e-6)
3
4 # Reconstruct the full image
5 rec = kruskal_to_tensor(weights, A, B, C)

```

---

If you want to visualise the result, you can do so easily with Matplotlib:

---

```

1 rec_image = to_image(rec)
2
3 #Import matplotlib to plot the image
4 import matplotlib.pyplot as plt
5 plt.imshow(rec_image)
6 plt.show()

```

---

The original image, shown in Figure 5.1, has height  $\times$  width  $\times$  #channels =  $768 \times 1024 \times 3 = 2,359,296$  elements. The decomposition, on the other hand, expresses the image as a series of factors A, B and C containing respectively height  $\times$  rank =  $768 \times 50$ , width  $\times$  rank =  $1024 \times 50$  and #channels  $\times$  rank =  $3 \times 50$ . In addition, we have to count the elements of `weights`, which is a vector of length equal to the rank. In total, the decomposition only has a total of less than 90,000 parameters, or approximately  $26\times$  less than the original image. Yet, as you can see in Figure 5.1, the reconstructed image looks visually similar to the uncompressed image.

Note that the CP decomposition is not the best fit here, since the same rank is used for all modes, including the RGB channels. This is a case where a Tucker decomposition would be more adapted as we can select the *Tucker rank* (or *multi-linear rank*) to more closely match that of the input tensor.

## 5.5 Going Further with TensorLy

We have so far shown how to implement some basic tensor manipulation functions as well as a CP decomposition algorithm based on Alternating Least Squares method. However, in practice, we want well-tested and robust algorithms that work at scale. This already exists in the TensorLy

library, which implements the methods presented in this section, and several more including Tucker decomposition, Robust Tensor PCA, low-rank tensor regression, etc.

The easiest way is to install TensorLy with pip (by simply typing `pip install tensorly` in the console). You can also install it directly from source at <https://github.com/tensorly/tensorly>.

When you have it installed, the usage is similar to what we have introduced above:

---

```
1 import tensorly as tl
2 import numpy as np
3
4 # Create a random tensor:
5 T = tl.tensor(np.random.random((10, 10, 10)))
6
7 # unfold the tensor:
8 unfolding = tl.unfold(T, mode=0)
9
10 # fold it back into a tensor
11 tl.fold(unfolding, mode=0, shape=tl.shape(T))
```

---

Decompositions are already implemented and can be readily applied to an input tensor:

---

```
1 from tensorly.decomposition import parafac, tucker
2
3 # CP decomposition
4 weights, factors = parafac(T, rank=3, normalize_factors=True)
5
6 # Tucker decomposition returns a core tensor and factor matrices
7 core, factors = tucker(T, ranks=[3, 2, 4])
```

---

You can also easily perform tensor regression using TensorLy, with a similar API that scikit-learn (Pedregosa *et al.*, 2011) offers. Refer to the website for a detailed tutorial<sup>1</sup> and API guide.

## 5.6 Scaling up with PyTorch

All the examples we have presented so far used small tensors that fit nicely in the memory of most commodity laptops and could be

---

<sup>1</sup><https://tensorly.github.io/dev/>



run quickly on their CPUs. However, as the size of the data and the complexity of the algorithms grow, we need highly-optimized functions that run on both GPU and CPU and on several machines in parallel. Running in multi-machines setup introduces the challenge of distributed inference and training. These can be incredibly complex to implement correctly. Fortunately, libraries exist that take care of it for you and let you focus on the logic of your model. One notable such framework is PyTorch (Paszke *et al.*, 2017).

By default, TensorLy uses NumPy as its backend. However, you can easily switch to PyTorch, a deep learning framework optimized for running large scale methods. Once you have installed PyTorch, you can easily use it as a backend for TensorLy and have all the operations run transparently on multiple machines and GPU. While CPUs performs operations on tensors in a mostly sequential way, GPUs accelerate operations by running them efficiently in parallel: modern CPUs typically contain up to 16 cores, while a GPU has thousands of them.

---

```

1 import tensorly as tl
2 import torch
3
4 # Use PyTorch as the backend
5 tl.set_backend('pytorch')
6
7 # Create a random tensor:
8 T = tl.tensor(np.random.random((10, 10, 10)))
9
10 type(T) # torch.Tensor!
11
12 # You can also specify where the tensor lives:
13 T = tl.tensor(np.random.random((10, 10, 10)), device='cuda:0')
```

---

Now, not only do all the algorithms in TensorLy run on GPU and CPU, you can also interface it easily with PyTorch and Deep Learning algorithms:

---

```

1 # Let's create a random number generator (rng)
2 random_state = 1234
3 rng = tl.random.check_random_state(random_state)
4
5 # You can put your tensor on cpu or gpu
6 device = 'cpu' # Or 'cuda:0'
7
```

---

---

```

8 # Create a random tensor
9 shape = [5, 5, 5]
10 tensor = tl.tensor(rng.random_sample(shape), device=device)

```

---

We have created a random tensor, which we will try to decompose in the Tucker form. This time, however, we will optimize the factors using gradient descent.

And this is where the magic happens: we can attach gradients to the tensors, using `requires_grad` parameter.

---

```

1 #Initialize a random Tucker decomposition of that tensor
2
3 # We choose a rank for the decomposition
4 rank = [5, 5, 5]
5
6 # We initialize a random Tucker core
7 core = tl.tensor(rng.random_sample(rank), requires_grad=True,
8                 device=device)
9
10 # We create a list of random factors
11 factors = []
12 for i in range(tl.ndim(tensor)):
13     factor = tl.tensor(rng.random_sample(
14                       (tensor.shape[i], rank[i])),
15                       requires_grad=True, device=device)
16     factors.append(factor)
17
18 #Let's use the simplest possible learning method: SGD
19 def SGD(params, lr):
20     for param in params:
21         # Gradient update
22         param.data -= lr * param.grad.data
23         # Reset the gradients
24         param.grad.data.zero_()

```

---

Now we can iterate through the training loop using gradient back-propagation:

---

```

1 n_iter = 7000
2 lr = 0.01
3 penalty = 0.1
4
5 for i in range(1, n_iter + 1):
6     # Reconstruct the tensor from the decomposed form
7     rec = tl.tucker_to_tensor((core, factors))
8
9     # l2 loss
10    loss = tl.norm(rec - tensor, 2)
11
12    # l2 penalty on the factors of the decomposition
13    for f in factors:

```

```

14         loss = loss + penalty * tl.norm(f, 2)
15
16     loss.backward()
17     SGD([core] + factors, lr)
18
19     if i % 100 == 0:
20         rec_error = tl.norm(rec - tensor, 2)/tl.norm(tensor, 2)
21         print("Epoch {}, Rec. error: {}".format(i, rec_error))
22
23     if i % 3000 == 0:
24         # Learning rate decay every 3000 iterations
25         lr /= 10

```

---

You will see the loss gradually go down as the approximation improves. You can verify that the relative reconstruction error is indeed small (we compute the error within a `no_grad` context as we do not want to compute gradients here):

```

1 with torch.no_grad():
2     # reconstruct the full tensor from these
3     rec = tl.tucker_to_tensor((core, factors))
4     relative_error = tl.norm(rec - tensor)/tl.norm(tensor)
5
6     print(relative_error)

```

---

To conclude, we have demonstrated in this section how to go from theory to a working implementation of tensor methods. These are powerful tools that can be efficiently leveraged using TensorLy. Using PyTorch and TensorLy together, you can easily combine tensor methods and Deep Learning, and run your model at scale across several machines and GPUs on millions of data samples. Next, we will discuss further practical considerations of tensor decomposition such as running time, memory usage, and sample complexity.

# 6

---

## Efficiency of Tensor Decomposition

---

In this section, we discuss the running time, memory usage and sample complexity for algorithms based on tensor decomposition.

Tensors are objects of very high dimensions; even a 3rd order  $d \times d \times d$  tensor with  $d = 10,000$  is already huge and hard to fit into memory of a single machine. A common misconception about tensor decomposition algorithms is that they need to use at least  $\Theta(d^3)$  memory, running time and number of samples, because the intuition is one needs at least one unit of resource for each entry in the tensor. This is in fact far from truth and the requirements on these three resources can be much smaller than  $O(d^3)$ . Furthermore, many tensor algorithms can be naturally parallelized and some of them can be run in an online fashion which greatly reduces the amount of memory required.

### 6.1 Running Time and Memory Usage

Storing a tensor explicitly as a high dimensional array and directly performing the computations on the explicit tensor can be very expensive. However, when applied to learning latent variable models and more generally when the tensor has an intrinsic lower dimensional structure, the tensor decomposition algorithms can often be made efficient.

**Number of Components:** Latent variable models represent observed variables using hidden variables, e.g., Gaussian mixture model with hidden Gaussian components, topic models with hidden topics and many other models that we described in Section 4. The good news is in most of the cases, the number of hidden components  $k$  is often much smaller than the dimension  $d$  of observed variables. For example, in topic modeling, the dimension  $d$  is equal to the number of words in vocabulary, which is at least in the order of thousands, while the number of topics can be  $k = 100$  in many applications. In these cases, after applying the Whitening Procedure proposed in Procedure 2 we only need to work with a  $k \times k \times k$  tensor which is easy to store in memory and allow for efficient computations.

### 6.1.1 Online Tensor Decomposition

Even when the number of components is large, it is still possible to run many tensor decomposition algorithms without explicitly constructing the tensor. This is because in most of the algorithms we only need to consider the effect of the tensor applied to vectors/matrices and not the whole tensor itself.

**Tensor Power Method:** It is very straightforward to convert each iteration of tensor power method to an online algorithm. In many cases, given samples  $x^{(1)}, \dots, x^{(n)}$ , the empirical tensor that we estimate can be represented as  $\frac{1}{n} \sum_{i=1}^n S(x^{(i)})$  where  $S(\cdot)$  is a function that maps a sample to a tensor. As an example, consider the multi-view model as explained in Section 4.3. Each sample  $x$  consists of three views  $(x_1, x_2, x_3)$ . Let  $S(x) = x_1 \otimes x_2 \otimes x_3$ , then we desire to estimate the mean tensor  $\mathbb{E}[S(x)]$ . Given  $n$  samples  $\{(x_1^{(i)}, x_2^{(i)}, x_3^{(i)}), i \in [n]\}$ , then the estimated empirical tensor is

$$\hat{T} = \frac{1}{n} \sum_{i=1}^n S(x^{(i)}) = \frac{1}{n} \sum_{i=1}^n x_1^{(i)} \otimes x_2^{(i)} \otimes x_3^{(i)}. \quad (6.1)$$

In tensor power method, the main iteration in (3.16) involves applying the tensor  $\hat{T}$  to vectors  $u, v$ , which can be easily done as

$$\hat{T}(u, v, I) = \frac{1}{n} \sum_{i=1}^n \langle x_1^{(i)}, u \rangle \langle x_2^{(i)}, v \rangle x_3^{(i)}. \quad (6.2)$$

Clearly, using this formula we only need to compute two inner-products for each sample, and the algorithm never needs to store more than a constant number of vectors.

**Claim 2** (Online Tensor Power Iteration). In many settings, one iteration of tensor power method can be done in time  $O(nd)$ , where  $n$  is the number of samples and  $d$  is the dimension. If number of samples is large enough, the algorithm is guaranteed to find an accurately estimated component in  $O(nd \log d)$  time with high probability.

**Alternating Least Squares:** ALS method relies on repeatedly solving least square problems; see Algorithm 8 for the details. To simplify the discussion, we focus on one step of the algorithm, where we are given matrices  $A, B$ , eigenvalues  $\lambda$  and want to find  $C$  such that  $\sum_{j=1}^k \lambda_j a_j \otimes b_j \otimes c_j$  is as close to the empirical tensor  $\hat{T}$  as possible; this is what Step 5 in Algorithm 8 does. All other steps are symmetric and can be computed similarly.

First, we observe that the problem can be decoupled into  $d$  sub-problems – one for finding each row of  $C$ . Consider the variant of Equation (3.31) for updating matrix  $C$  (when  $A, B$  and  $\lambda$  are fixed), and pick the  $i$ -th row of  $\text{mat}(T, 3) \in \mathbb{R}^{d \times d^2}$  and matricize it to a  $d \times d$  matrix. This leads to the following set of sub-problems to solve for different rows of matrix  $C$  denoted by  $C^{(i)}$ ,

$$\min_{C^{(i)}} \left\| \hat{T}(I, I, e_i) - \sum_{j=1}^k \lambda_j C_{i,j} a_j b_j^\top \right\|_F, \quad i \in [d].$$

These  $d$  sub-problems can be solved in parallel which makes it faster to run ALS.

Furthermore, we can use efficient gradient-based methods in the context of online learning even without exploiting parallelization as

above. Recently there has been a lot of research on using online gradient-based algorithms to solve least square problems (Shalev-Shwartz and Zhang, 2013; Johnson and Zhang, 2013), and they can all be applied here. A common assumption in these works is that the objective function can be decomposed into the sum of  $n$  terms, where the gradient for each term can be computed efficiently. More precisely, the optimization should be of the form

$$\min_C \sum_{i=1}^n f(C, x^{(i)}). \tag{6.3}$$

The guarantee for these online algorithms can be stated in the following informal statement.

**Claim 3.** Suppose the objective function in (6.3) is well-conditioned and  $n$  is large enough, and the time for computing the gradient for a single  $f$  is  $T$ . Then, there exist algorithms that can find the optimal solution with accuracy  $\epsilon$  in time  $O(Tn \log n/\epsilon)$ .

In other words, when the problem is well-conditioned, the algorithms only need a few passes on the data set to find an accurate solution. Having these results for gradient-based methods, we now convert the objective function of ALS to a form similar to the one in (6.3). Again suppose we are in the setting that the empirical tensor can be computed as average of  $S(x^{(i)})$ 's; see Equation (6.1). Recall the original objective function for ALS is

$$\min_C \left\| \sum_{j=1}^k \lambda_j a_j \otimes b_j \otimes c_j - \mathbb{E}_{i \in [n]} S(x^{(i)}) \right\|_F^2,$$

where  $\mathbb{E}_{i \in [n]}(\cdot) := \frac{1}{n} \sum_{i \in [n]} \cdot^{(i)}$ . For any random variable  $X$ , we know

$$(a - \mathbb{E}[X])^2 = \mathbb{E}[(a - X)^2] - \mathbb{E}[(X - \mathbb{E}[X])^2].$$

Therefore, we can rewrite the objective function as

$$\min_C \frac{1}{n} \sum_{i=1}^n \left\| \sum_{j=1}^k \lambda_j a_j \otimes b_j \otimes c_j - S(x^{(i)}) \right\|_F^2 - \frac{1}{n} \sum_{i=1}^n \|S(x^{(i)}) - \hat{T}\|_F^2,$$

where  $\hat{T} := \frac{1}{n} \sum_{i \in [n]} S(x^{(i)})$ . The second term does not depend on  $C$ , so it can be ignored in the optimization problem. Let

$$f(C, x^{(i)}) := \left\| \sum_{j=1}^k \lambda_j a_j \otimes b_j \otimes c_j - S(x^{(i)}) \right\|_F^2,$$

and thus, we have rewritten the objective function as  $\frac{1}{n} \sum_{i=1}^n f(C, x^{(i)})$ , which is exactly the form required in (6.3). The gradient of  $f$  functions w.r.t. to the columns of matrix  $C$  denoted by  $C_t$  can also be computed as

$$\frac{\partial}{\partial C_t} f(C, x^{(i)}) = 2\lambda_t \sum_{j=1}^k \lambda_j \langle a_j, a_t \rangle \langle b_j, b_t \rangle C_j - 2\lambda_t S(x^{(i)})(a_t, b_t, I).$$

Computing this stochastic gradient for all the entries of matrix  $C$ , i.e., all  $C_{i,j}$ 's only take  $\Theta(k^2 d)$  time. Then, combined with Claim 3 allows the least squares problem to be solved efficiently. However, from an arbitrary initialization, we do not have any theoretical bounds on the condition number of these least-squares problems, or the number of iterations it takes ALS to converge. Theoretical analysis of ALS algorithm is still an open problem.

## 6.2 Sample Complexity

One major drawback of tensor decomposition algorithms is that they often require a fairly large number of samples. A large number of samples may be hard to get in practice, and can also slow down the algorithms – as we just saw, many of the tensor decomposition algorithms can be implemented so that they only need to go through the data set small number of times.

A misleading intuition argues that in order to estimate every entry of an  $d \times d \times d$  tensor to an accuracy of  $\epsilon$ , one would need  $d^3/\epsilon^2$  samples, which is often too large to be practical. However, this argument is based on the incorrect assumptions that 1) each sample is highly noisy and only provide a small amount of information; 2) the tensor decomposition algorithms require every entry of the tensor to be estimated accurately. The real number of samples required is distinct for different applications, and is far from well-understood.



### 6.2.1 Tensor Concentration Bounds

In tensor decompositions, often we do not need to estimate every entry of the tensor. Instead, we would like to approximate the tensor in a certain norm, e.g., spectral norm, Frobenius norm and other norms based on Sum-of-Squares relaxations are often used. For a specific norm, tensor concentration bounds give estimates on how many samples we need in order to estimate the tensor within some error  $\epsilon$ .

When the norm is the Frobenius norm, or the spectral norm of some unfolded version of the tensor (matricized version), the problem can be reduced to vector concentration bounds or matrix concentration bounds. There has been a lot of research on matrix concentrations, many popular bounds can be found in Tropp, 2012.

**Tensors with Independent Entries:** For the spectral norm of the tensor, one of the first concentration bounds is by Latała (2005), and later generalized in Nguyen *et al.* (2010). They consider the case when there is a random tensor  $T$  whose entries are independent random variables with zero mean. For simplicity, we state the following corollary to give a flavor on what they provide.

**Corollary 6.1** (Corollary 3 of Nguyen *et al.*, 2010). Suppose order- $p$  tensor  $T \in \mathbb{R}^{d_1 \times d_2 \times \dots \times d_p}$  has i.i.d. standard Gaussian entries. Then for every  $p$ , there exists a constant  $C_p > 0$  such that with high probability,

$$\|T\|^2 \leq C_p \max\{d_1, d_2, \dots, d_p\}.$$

This shows the spectral norm of a Gaussian tensor only depends on its largest dimension. More specifically, for a  $d \times d \times d \times d$  tensor, its spectral norm is still with high probability  $O(\sqrt{d})$ , which is much smaller than its Frobenius norm  $\Theta(d^2)$  or the spectral norm of an unfolded matricization  $\Theta(d)$ . The technique used in these papers is called the “entropy-concentration”. The key idea is to argue about linear forms  $T(v_1, \dots, v_p)$  separately for vectors  $v_i$ ’s that are sparse (low entropy) and dense.

**Tensors from Latent Variable Models:** When the tensor is constructed from a latent variable model (see Section 4 for many examples), the

coordinates of the tensor are often not independent. A case-by-case analysis is required. Anandkumar *et al.* (2014b) analyzed the number of samples required for multi-view model and independent component analysis. The ideas used are again similar to the entropy concentration approach, except a vector is considered “sparse” if it has large correlation only with a few components. We provide the guarantee in a simple multi-view model; refer to Anandkumar *et al.* (2014b) for more detailed results.

**Corollary 6.2** (Spectral Norm Bound for Multi-view Model by Anandkumar *et al.* (2014b)). Consider a simple multi-view model where the latent variable has  $k$  possibilities. Each sample is generated by first picking a hidden variable  $h \in [k]$ , and then observing  $x_1 = a_h + \zeta_1$ ,  $x_2 = b_h + \zeta_2$ ,  $x_3 = c_h + \zeta_3$ . Here for  $h \in [k]$ ,  $a_h$ 's,  $b_h$ 's,  $c_h$ 's are  $d$ -dimensional conditional means and assumed to be random unit vectors, and  $\zeta_1$ ,  $\zeta_2$ ,  $\zeta_3$  are independent random Gaussian noise vectors whose variance is 1 in each coordinate. Given  $n$  samples  $\{(x_1^{(i)}, x_2^{(i)}, x_3^{(i)}) : i \in [n]\}$ , let

$$\hat{T} := \frac{1}{n} \sum_{i=1}^n x_1^{(i)} \otimes x_2^{(i)} \otimes x_3^{(i)},$$

$$T := \frac{1}{n} \sum_{i=1}^n h_i a_{h_i} \otimes b_{h_i} \otimes c_{h_i},$$

where  $h_i$  denotes the true hidden value for the  $i$ -th sample. Then with high probability,

$$\|\hat{T} - T\| \leq O(\sqrt{d/n} \cdot \text{poly log } n).$$

Note that in the above model, the noise is extremely high where the total norm of the noise is  $O(\sqrt{d})$  compared to the norm of the signal  $\|a_h\| = \|b_h\| = \|c_h\| = 1$ . Even in this high-noise regime, it only takes  $d \cdot \text{poly log } d$  samples to estimate the tensor with constant accuracy in spectral norm. The result is tight up to poly log factors.

### 6.2.2 Case Study: Tensor PCA and Tensor Completion

Given a tensor  $\hat{T} = T + E$  where  $E$  is a perturbation tensor, concentration bounds give us tools to bound the norm of error  $E$ . However,

different tensor decomposition algorithms may have different requirement on  $E$ . Finding the “most robust” tensor decomposition algorithm is still an open problem. In this section, we will describe recent progress in some specific problems.

**Tensor PCA:** The model of tensor PCA is very simple. There is an unknown signal  $v \in \mathbb{R}^d$  with  $\|v\| = 1$ . Now suppose we are given tensor

$$\hat{T} = \tau \cdot v \otimes v \otimes v + E,$$

where  $E \in \mathbb{R}^{d \times d \times d}$  is a noise tensor whose entries are independent standard Gaussians, and  $\tau \in \mathbb{R}$  is a scalar. The goal is to find a vector that is within a small constant distance to  $v$  when  $\hat{T}$  is given. The parameter  $\tau$  determines the signal-to-noise ratio, and the problem is easier when  $\tau$  is larger. This problem was originally proposed by Richard and Montanari, 2014 as a simple statistical model for tensor PCA.

If the algorithm can take exponential time, then the best solution is to find the unit vector  $u$  that maximizes  $\hat{T}(u, u, u)$ . By Corollary 6.1 we know the spectral norm of  $E$  is bounded by  $O(\sqrt{d})$ , and therefore, as long as  $\tau \geq C\sqrt{d}$  for some universal constant  $C$ , the optimal direction  $u$  has to be close to  $v$ .

However, when the algorithm is required to run in polynomial time, the problem becomes harder. The best known result is from Hopkins *et al.* (2015) as follows.

**Theorem 6.3** (Hopkins *et al.*, 2015). If  $\tau = Cd^{3/4}$  for some universal constant  $C$ , then there is an efficient algorithm that finds a vector  $u$  such that with high probability  $\|u - v\| \leq 0.1$ . Moreover, no Sum-of-Squares algorithm of degree at most 4 can do better.

The term  $d^{3/4}$  in the above theorem is between the information theoretic limit  $\sqrt{d}$  and the trivial solution that treats the tensor as a  $d \times d^2$  matrix which gives  $\Theta(d)$  bound for  $\tau$ . The problem can also be solved more efficiently using a homotopy optimization approach (Anandkumar *et al.*, 2017). However, it seems there are some fundamental difficulties in going below  $d^{3/4}$ .

**Tensor Completion:** A very closely related problem is called tensor completion. In this problem, we observe a random subset of entries of a low-rank tensor  $T$ , and the goal is to recover the original full low-rank tensor. Barak and Moitra (2016) provide a tight bound on the number of samples required to recover  $T$ .

**Theorem 6.4** (Barak and Moitra, 2016, informal). Suppose tensor  $T \in \mathbb{R}^{d \times d \times d}$  has rank  $k$ , given  $n = d^{1.5}k$  poly  $\log d$  random observations of the entries of the tensor, there exists an algorithm that recovers  $T$  up to a lower order error term.

For small  $k$ , the term  $d^{1.5}k$  poly  $\log d$  in the above guarantee is again between information theoretic limit  $\Theta(dk \log d)$  and the trivial solution that considers the tensor as a  $d \times d^2$  matrix which gives  $\Theta(d^2k \log d)$ . Barak and Moitra, 2016 showed the  $d^{1.5}$  dependency which is likely to be tight because improving this bound will also give a better algorithm for refuting random 3-XOR clauses (which is a conjectured hard problem (Feige, 2002; Grigoriev, 2001; Schoenebeck, 2008)). The algorithm is again based on Sum-of-Squares. Recently there were also several improvements in the recovery guarantees, see e.g., Potechin and Steurer (2017) and references therein.

# 7

---

## Overcomplete Tensor Decomposition

---

Unlike matrices, the rank of a tensor can be higher than its dimension. We call such tensors *overcomplete*. Overcomplete tensors can still have a unique decomposition; recall Theorem 3.1 for 3rd order tensors and see Sidiropoulos and Bro (2000) for higher order tensors. This is useful in the application of learning latent variable models: it is possible to learn a model with more components than the number of dimensions, e.g., a mixture of 100 Gaussians in 50 dimensions.

However, finding a CP decomposition for an overcomplete tensor is much harder than the undercomplete case (when the rank is at most the dimension). In this section we will describe a few techniques for decomposing overcomplete tensors.

### 7.1 Higher-order Tensors via Tensorization

For higher order tensors, the most straightforward approach to handle overcomplete decomposition is to convert them to lower order tensors but in higher dimension. We call this approach as *tensorization* and describe it in this section. For simplicity, we restrict our attention to symmetric tensors, but what we discuss here also applies to asymmetric

**Algorithm 10** Decomposing Overcomplete Tensors via Tensorization**input** tensor  $T = \sum_{j=1}^k a_j^{\otimes 6}$ **output** rank-1 components  $a_j, j \in [k]$ 

- 1: Reshape the tensor  $T$  to  $\hat{T} \in \mathbb{R}^{d^2 \times d^2 \times d^2}$ , each mode of  $\hat{T}$  is indexed by  $[d] \times [d]$ , and  $\hat{T}_{(i_1, i_2), (i_3, i_4), (i_5, i_6)} = T_{i_1, i_2, i_3, i_4, i_5, i_6}$ .
- 2: Use Tensor Power Method (Algorithm 6) to decompose  $\hat{T} = \sum_{j=1}^k b_j^{\otimes 3}$  and recover  $b_j$ 's.
- 3: For each  $b_j \in \mathbb{R}^{d^2}, j \in [k]$ , reshape it as a  $d \times d$  matrix, and let  $(\lambda_j, v_j)$  be its top singular value-vector pair.
- 4: **return**  $\sqrt{\lambda_j} v_j, j \in [k]$ .

tensors. Consider a 6th order tensor

$$T = \sum_{j=1}^k \lambda_j a_j^{\otimes 6},$$

with  $d$ -dimensional rank-1 components  $a_j \in \mathbb{R}^d, j \in [k]$ , and real weights  $\lambda_j \in \mathbb{R}, j \in [k]$ , where  $k \gg d$ . We can reshape this tensor as a 3rd order tensor  $\hat{T} \in \mathbb{R}^{d^2 \times d^2 \times d^2}$  as follows. Let

$$b_j := a_j \odot a_j \in \mathbb{R}^{d^2},$$

where  $\odot$  denotes the Khatri-Rao product defined in (3.9); note that with vector inputs, this works the same as Kronecker product. Then we have

$$\hat{T} = \sum_{j=1}^k \lambda_j b_j^{\otimes 3}.$$

We call this process of reshaping the tensor to a different order as tensorization. Now for the 3rd order tensor  $\hat{T}$ , if the rank-1 components  $b_j$ 's are linearly independent, we can use the tensor decomposition algorithms in Section 3 to recover its rank-1 components  $b_j$ 's, and since  $b_j = a_j \odot a_j$ , the original rank-1 components  $a_j$  is computed as the top singular vector of the matricized version of  $b_j$ . This whole approach is provided in Algorithm 10.

If  $k \leq \binom{d+1}{2}$  and the vectors  $a_j$ 's are in general position, then the vectors  $b_j$ 's are going to be linearly independent. Recent work by Bhaskara

*et al.* (2014) shows that if  $a_j$ 's are perturbed by a random Gaussian noise, the smallest singular value of matricized  $b_j$ 's are lower bounded (where the lower bound depends polynomially on the magnitude of the noise and the dimension). As a result this algorithm is robust to small amount of noise.

In Algorithm 10, we used tensor power iteration as the core tensor decomposition algorithm. It is worth mentioning that we can also use other tensor decomposition algorithms. In particular, if we use simultaneous diagonalization algorithm as proposed in Algorithm 7 instead of tensor power method, then it suffices to have a 5th order input tensor  $T = \sum_{j=1}^k \lambda_j a_j^{\otimes 5}$ . Again we can reshape the tensor as

$$\hat{T} = \sum_{j=1}^k \lambda_j b_j \otimes b_j \otimes a_j.$$

Even though the third mode still only has  $d$  dimensions which is smaller than rank  $k$  in the overcomplete regime, simultaneous diagonalization only requires the rank to be less or equal to the dimension of the first two modes, and thus, the algorithm can work.

The same idea can be also generalized to even higher order tensors. In general, if vectors  $a_j^{\otimes r}$  are linearly independent, then we can apply simultaneous diagonalization algorithm to a  $(2r + 1)$ -th order tensor and compute the unique tensor decomposition.

This algorithm can be applied to learning several latent variable models, as long as we have access to higher order tensors. In some applications such as pure topic models, this is fairly straightforward as we only need to form the moment using the correlations of  $p$  words instead of 3. In other applications this may require careful manipulations of the moments. In both cases, working with higher order tensors can potentially increase the sample complexity and running time of the algorithm.

## 7.2 FOBI Algorithm

In practice, working with high-order tensors is often too expensive in terms of both sample complexity and running time. Therefore, it is useful to design algorithms that can handle overcomplete tensors when

the order of the tensor is low, e.g., 3rd or 4th order tensors. For 4th order tensors, De Lathauwer *et al.* (2007) proposed an algorithm called FOABI (Fourth-Order-Only Blind Identification) that can work up to rank  $k = Cd^2$  for some fixed constant  $C > 0$ . For simplicity, we again describe the algorithm for symmetric tensors

$$T = \sum_{j=1}^k \lambda_j a_j^{\otimes 4},$$

and we will keep the notation  $b_j := a_j \odot a_j$  as we had in the previous section. We will also assume  $\lambda_j > 0$  and the components are real-valued. All these requirements can be removed for this algorithm and interested readers are encouraged to check the original paper. We provide the FOABI method in Algorithm 11. Intuitively, the algorithm has three main parts

1. Finding the span of vectors  $\{b_j : j \in [k]\}$ .
2. Finding the  $b_j$ 's.
3. Computing the  $a_j$ 's.

Step 1 is done using a SVD operation and step 3 is achieved the same as what we discussed in the previous section for higher order tensors. The magic happens in step 2 where the algorithm uses a quadratic operator to detect rank-1 matrices. In the rest of this section, we describe these steps in more details.

### 7.2.1 Finding Span of $\{b_j\}$ 's

In the first step, we try to find the span of the vectors  $b_j$ 's. This is very simple as we can matricize the tensor as

$$M := T_{\{1,2\},\{3,4\}} = \sum_{j=1}^k \lambda_j b_j b_j^\top \in \mathbb{R}^{d^2 \times d^2}, \quad (7.1)$$

where  $b_j := a_j \odot a_j \in \mathbb{R}^{d^2}$ . Therefore, we just need to compute the column span (or row span) of  $M$ , and it would correspond to the span of vectors  $\{b_j\}$ 's. In order to make the algorithm more robust to noise, we use singular value decomposition to find the top singular values, and drop all the singular values that are very close to 0.



---

**Algorithm 11** FOOBI for Decomposing Overcomplete Tensors (De Lathauwer *et al.*, 2007)

---

**input** tensor  $T = \sum_{j \in [k]} \lambda_j a_j^{\otimes 4}$

**output** rank-1 components  $\{(\lambda_j, a_j)\}$ 's

- 1: Reshape the tensor to a matrix

$$M = \sum_{j=1}^k \lambda_j (a_j \odot a_j)(a_j \odot a_j)^\top \in \mathbb{R}^{d^2 \times d^2}.$$

- 2: Compute the SVD of  $M$  as  $M = UDU^\top$ .

- 3: Let  $L : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^{d^4}$  be a quadratic operator such that

$$L(A)(i_1, i_2, j_1, j_2) = \det \begin{pmatrix} A_{i_1, j_1} & A_{i_1, j_2} \\ A_{i_2, j_1} & A_{i_2, j_2} \end{pmatrix}.$$

Let  $\tilde{L} : \mathbb{R}^{d^2 \times d^2} \rightarrow \mathbb{R}^{d^4}$  be the unique linear operator that satisfy  $\tilde{L}(A \otimes A) = L(A)$ .

- 4: Construct matrix  $Z = \tilde{L}[(UD^{1/2}) \otimes (UD^{1/2})]$ .

- 5: Let  $y_1, y_2, \dots, y_k \in \mathbb{R}^{k^2}$  be the  $k$  least right singular vectors of  $Z$ .

- 6: Pick random vectors  $u, v \in \text{span}(y_1, \dots, y_k)$ , reshape them as  $k \times k$  matrices  $U, V$ .

- 7: Use Simultaneous Diagonalization (see Algorithm 7 in Section 3.6) to express  $U = QD_UQ^\top$  and  $V = QD_VQ^\top$ , where  $D_U, D_V$  are  $k \times k$  diagonal matrices and  $Q \in \mathbb{R}^{k \times k}$  is shared between  $U, V$ .

- 8: Let  $x_1, \dots, x_k \in \mathbb{R}^k$  be the columns of  $Q$ ,  $b_j = UD^{1/2}x_j$  for all  $j \in [k]$ .

- 9: For  $j \in [k]$ , reshape  $b_j \in \mathbb{R}^{d^2}$  to  $d \times d$  matrix and let  $(\delta_j, v_j)$  be its top singular value-vector pair.

- 10: **return**  $\{(\delta_j^2, v_j) : j \in [k]\}$ .
- 

### 7.2.2 Finding $\{b_j\}$ 's

In the second step of the algorithm, we can view the vectors  $b_j$ 's as reshaped  $d \times d$  matrices. For vector  $u \in \mathbb{R}^{d^2}$ , the matricized version

$\text{mat}(u) \in \mathbb{R}^{d \times d}$  is defined as

$$\text{mat}(u)_{i,j} = u(d(j-1) + i), \quad i, j \in [d],$$

which is formed by stacking the entries of  $u$  in the columns of the matrix. Given the definition of  $b_j$ 's, we have  $\text{mat}(b_j) = a_j a_j^\top$  which are rank-1 matrices. In addition, from the previous step, we know the linear subspace spanned by these matrices. Using the key observation that  $\text{mat}(b_j)$  are all rank-1 matrices, we hope to recover  $b_j$ 's as follows.

Suppose the SVD of  $M$  in (7.1) is denoted by  $UDU^\top$ , where  $U \in \mathbb{R}^{d^2 \times k}$  be an orthonormal matrix that represents the span of the vectors  $\{b_j\}$ 's. Since we know

$$M = (UD^{1/2})(UD^{1/2})^\top = \sum_{j=1}^k (\sqrt{\lambda_j} b_j)(\sqrt{\lambda_j} b_j)^\top,$$

there exists an orthogonal matrix  $R$  such that the columns of  $UD^{1/2}R$  are equal to  $\{\sqrt{\lambda_j} b_j\}$ 's. In order to find the vectors  $b_j$ 's, we need to find the columns of this orthogonal matrix  $R$  denoted by  $x_j$  such that  $\text{mat}(UD^{1/2}x_j)$ 's are rank-1 matrices; recall the above discussion that the matricized versions of  $b_j$ 's are rank-1 matrices.

Finding these  $x_j$  directions is not an easy task. De Lathauwer *et al.* (2007) show that it is possible to do this using a very interesting rank-1 detector.

### 7.2.3 Rank-1 Detector

For a matrix  $A \in \mathbb{R}^{d \times d}$ , we know  $A$  is rank at most 1 if and only if determinants of all  $2 \times 2$  submatrices of  $A$  are equal to 0. In particular, for a symmetric matrix  $A$ , we can define a mapping  $L(A)$  that maps  $A$  to a  $\mathcal{D} \approx d^4/8$  dimensional space, where each entry in  $L(A)$  corresponds to the value of the determinant of a unique  $2 \times 2$  submatrix of  $A$ . The exact number of dimensions is  $\mathcal{D} = \binom{d}{2}^{+1}$ , because that is the number of 4-tuples  $(i_1, i_2, j_1, j_2)$  where  $i_1 < i_2, j_1 < j_2$  and  $(i_1, i_2) \leq (j_1, j_2)$ .

**Definition 7.1** (rank-1 detector). Function  $L$  maps  $d \times d$  symmetric matrices to  $\mathcal{D} = \binom{d}{2}^{+1}$  dimensional space indexed by  $(i_1, i_2, j_1, j_2)$

where  $i_1 < i_2, j_1 < j_2$  and  $(i_1, i_2) \leq (j_1, j_2)$ , where

$$L(A)(i_1, i_2, j_1, j_2) = \det \begin{pmatrix} A_{i_1, j_1} & A_{i_1, j_2} \\ A_{i_2, j_1} & A_{i_2, j_2} \end{pmatrix}.$$

It is easy to prove that this rank-1 detector indeed works.

**Claim 4.** Symmetric matrix  $A \in \mathbb{R}^{d \times d}$  is of rank at most 1, if and only if  $L(A) = 0$ .

The mapping  $L(A)$  is quadratic in the entries of  $A$ . Therefore, if we apply  $L$  to the matrix  $Vx$ , then  $L(Vx)$  is also quadratic in the variable  $x$ . Naïvely,  $L(Vx) = 0$  would give a set of quadratic equations, and solving a system of quadratic equations is again hard in general. Luckily, we have a very large number of equations –  $\mathcal{D} \approx d^4/8$ . This allows us to use a *linearization* approach: instead of treating  $L(Vx) = 0$  as a system of quadratic equations over  $x$ , we will *lift* the variables to  $X = x \odot x$  and view  $X$  as a  $\binom{d+1}{2}$  dimensional vector. Now in our problem,  $L(UD^{1/2}x)$  is equal to a linear operator  $\tilde{L}$  applied to  $[(UD^{1/2}) \otimes (UD^{1/2})]X$ , i.e.,

$$L(UD^{1/2}x) = \tilde{L}([(UD^{1/2}) \otimes (UD^{1/2})]X).$$

**Definition 7.2** (Linearized detector). Linearized rank-1 detector  $\tilde{L}$  maps a  $d^2 \times d^2$  matrix such that

$$\tilde{L}(A \otimes A) = L(A).$$

More precisely, for two matrices  $A, B \in \mathbb{R}^{d \times d}$  we have

$$\tilde{L}(A \otimes B)(i_1, i_2, j_1, j_2) = A_{i_1, j_1} B_{i_2, j_2} - A_{i_1, j_2} B_{i_2, j_1},$$

which is the determinant of the  $2 \times 2$  submatrix of  $A$  when  $A = B$ .

Since  $\tilde{L}$  is a linear operator, we can represent it as a matrix; let  $Z = \tilde{L}([(UD^{1/2}) \otimes (UD^{1/2})])$ , and try to solve the system of linear equations  $ZX = 0$ . Of course, in general doing this ignores the structure in  $X$  (that it has the form of  $X = x \odot x$ ), and will not work for general quadratic equations. In this specific case, De Lathauwer *et al.* (2007) were able to show that the only solutions of this equation are linear combinations of the desired solution.

**Theorem 7.1** ((De Lathauwer *et al.*, 2007)). Let  $a_j$ 's be in general positions and  $k \leq Cd^2$  for some universal constant  $C > 0$ . Let  $x_1, x_2, \dots, x_k$  be  $k$  vectors such that  $Ux_j = b_j$ , and  $X_j = x_j \odot x_j$  for  $j \in [k]$ . Then the solution of  $\tilde{L}[(UD^{1/2}) \otimes (UD^{1/2})]X = 0$  (the null space of  $Z$ ) is exactly equal to the span of  $X_j$ 's.

### 7.2.4 Finding the Rank-1 Components

In this final step, we are given a subspace which is equal to the span of  $\{x_j \odot x_j\}$ 's, and we are trying to find  $\{x_j\}$ 's. At a first glance, this might look exactly the same as the problem we were facing in the previous step: we were given the span of  $\{a_j \odot a_j\}$  and trying to find  $\{a_j\}$ 's. Luckily these two problems are actually very different – in both cases we are looking for  $k$  vectors, but previously we were given a span of  $d \times d$  matrices (and  $k \gg d$ ) and now we have a span of  $k \times k$  matrices. The vectors  $a_j$ 's cannot be linearly independent, while the vectors  $x_j$ 's are usually linearly independent. Now, to find  $\{x_j\}$ 's, the key observation is that every matrix in the span of  $\{x_j \odot x_j\}$ 's can be *simultaneously diagonalized*, and the vectors  $x_j$ 's are the only way to do that. Therefore the last step of the algorithm is very similar to the simultaneous diagonalization algorithm for undercomplete tensor decomposition. In Algorithm 11, for simplicity we just applied simultaneous diagonalization on two random matrices  $U, V$  in the subspace found in the last step (as we discussed in Section 3.6 this can be easily done by eigendecomposition of  $UV^{-1}$ ). To ensure numerical stability, the original FOABI algorithm requires a simultaneous diagonalization of all the  $y_j$ 's. As a result, we find  $k$  vectors  $x_1, \dots, x_k$  such that  $UD^{1/2}x_j = \sqrt{\lambda_j}b_j$ . The rest of the algorithm is simply recovering  $a_j$ 's from  $b_j$ 's.

## 7.3 Third Order Tensors

Algorithms like FOABI can work with tensors with order at least 4. That leaves only third order tensors. We still don't have any algorithms for overcomplete third order tensors when the components are only guaranteed to be in general position.

Third order tensor is very special and might be fundamentally more

difficult to decompose. As an example, it is very easy to construct an explicit 4-th order tensor that has rank at least  $d^2$ , because the rank of the tensor is at least as large as its matricizations. However, for a  $d \times d \times d$  third order tensor, all matricizations can have rank at most  $d$ ; note that even the most balanced matricizations have dimensions  $d^2 \times d$  or  $d \times d^2$ . It is still an open problem to construct an explicit third order tensor whose rank is super-linear in dimension, and in fact doing so will lead to circuit lower bounds that were open for decades (Raz, 2013).

Because of these difficulties, researchers have focused on the simpler setting where the components  $\{a_j\}$ 's are chosen from a random distribution. Even in this simple case, the only provable algorithm relies on complicated algorithms called Sum-of-Squares Hierarchies. We refer the readers to the survey by Barak and Steurer (2014). In the rest of this section, we give some intuitions on how to handle overcomplete third order tensors without going into the details.

### 7.3.1 Lifting the Tensor

A key technique in handling third order tensor is to *lift* the tensor into a higher order tensor. This can either be done explicitly, or implicitly using Sum-of-Squares framework. Here we show a simple transformation that lifts a third order tensor to a 4th order tensor. Again for simplicity, we only work with symmetric tensors in this section.

**Definition 7.3** (Lifted tensor). Given a tensor  $T \in \mathbb{R}^{d \times d \times d}$ , we can construct a lifted tensor  $M(T) \in \mathbb{R}^{d \times d \times d \times d}$  as

$$M(T)_{i_1, i_2, i_3, i_4} := \sum_{i=1}^d T_{i, i_1, i_2} T_{i, i_3, i_4}.$$

Note that for a rank-1 tensor  $T = a \otimes a \otimes a$ , the lifted tensor  $M(T) = \|a\|^2 a^{\otimes 4}$  is also rank-1. It is easier to interpret the lifted tensor using the multilinear form, in particular, we have

$$M(T)(x, x, x, x) = \|T(:, x, x)\|^2.$$

As a result, if  $T$  has decomposition  $T = \sum_{j=1}^k a_j \otimes a_j \otimes a_j$ , we can represent  $M(T)$  as a low rank tensor plus noise. This is formulated as follows.

**Theorem 7.2** (Ge and Ma, 2015). Suppose

$$T = \sum_{j \in [k]} a_j \otimes a_j \otimes a_j,$$

then the lifted tensor can be represented as

$$M(T) = \sum_{j \in [k]} \|a_j\|^2 a_j^{\otimes 4} + M',$$

where

$$M' = \sum_{i \neq j \in [k]} \langle a_i, a_j \rangle a_i \otimes a_i \otimes a_j \otimes a_j.$$

Furthermore, suppose  $a_j$ 's are chosen according to Gaussian distribution with expected square norm 1. Then the norm  $\|M'_{\{1,3\},\{2,4\}}\|$  is bounded by  $o(1)$  when  $k \leq d^{3/2}/\text{poly} \log(d)$ , where  $M'_{\{1,3\},\{2,4\}} \in \mathbb{R}^{d^2 \times d^2}$  denotes the matricization of tensor  $M'$  such that the 1st and 3rd modes are stacked along the rows, and the 2nd and 4th modes are stacked along the columns of the matrix.

Intuitively, this theorem shows that after the lifting operation, we get a 4th order rank- $k$  tensor with noise  $M'$ . The norm of  $M'$  is small compared to the true components in the low rank decomposition. Therefore, it is possible to find  $a_j$ 's as long as we can decompose 4th order tensors under such kind and amount of noise. Ge and Ma (2015) gave a quasi-polynomial time algorithm to do this. Later, Ma *et al.* (2016) showed it is also possible to do this within polynomial time.

### 7.3.2 Robust 4th Order Tensor Decomposition

In order to solve the 4th order tensor decomposition problem, we might want to use the FOABI algorithm described earlier. However, the noise term  $M'$  here has spectral norm  $o(1)$ , and the FOABI algorithm is not known to be robust to such perturbations.

Using Sum-of-Squares techniques, Ma *et al.* (2016) gave an algorithm that can decompose a 4th order tensor even under significant noise.

**Theorem 7.3** (Ma *et al.*, 2016). Let  $T \in \mathbb{R}^{d \times d \times d \times d}$  be a symmetric 4th order tensor and  $a_1, \dots, a_k \in \mathbb{R}^d$  be a set of vectors. Define perturbation

tensor  $E := T - \sum_{j \in [k]} a_j^{\otimes 4}$ , and define  $A$  as the matrix with columns  $a_j^{\otimes 2}$ . If the (matricized) perturbation norm  $\|E_{\{1,2\},\{3,4\}}\| \leq \delta \cdot \sigma_k(AA^\top)$  for some  $\delta > 0$ , then there is an algorithm that outputs a set of vectors  $\hat{a}_j$ , and there is a permutation  $\pi : [k] \rightarrow [k]$  such that for every  $j \in [k]$ , we have

$$\min\{\|a_j - \hat{a}_{\pi(j)}\|, \|a_j + \hat{a}_{\pi(j)}\|\} \leq O(\delta \|A\| / \sigma_k(A)) \|a_j\|.$$

Note that we cannot directly combine this theorem with Theorem 7.1 to get a complete algorithm for decomposing overcomplete 3rd order tensors. There are a few technical issues: 1. The norms of  $\|a_j\|$ 's are not exactly 1, but they are very close to 1 by concentration; 2. we need to reshape the tensor so that the matricization  $T_{\{1,2\},\{3,4\}}$  has small spectral norm; 3. when  $a_j$ 's are random,  $\|A\|/\sigma_k(A)$  is usually  $\sqrt{k/d}$  which is bigger than 1. The first two problems are easy to handle, while the third problem requires more work.

Ma *et al.* (2016) also give a direct analysis for overcomplete 3rd order tensors using Sum-of-Squares, and that analysis do not rely on the explicit lifting.

## 7.4 Open Problems

Despite the algorithms we provided, decomposing an overcomplete tensor is still a very difficult problem. The algorithms often require access to high-order tensors, which is often expensive in both sample complexity and running time. Sum-of-Squares algorithms can tolerate more noise and therefore, work with fewer samples, but the running time is prohibitive. Finding a provable overcomplete tensor decomposition algorithm that is efficient in practice is still a major problem.

Although the current provable algorithms are quite complicated, in practice algorithms like Alternating Least Squares or Power Method (see Section 3) work surprisingly well even when the tensor is overcomplete. For a random 3rd order tensor with dimension 100 and rank 1000, Alternating Least Squares with random initialization almost always converges to the right answer within 10 iterations. This is very surprising and we do not yet know how to prove it works. When the components are not randomly generated, people have observed Alternating Least

Squares can be sometimes slow (Comon, [2002](#)). How to handle and analyze these kind of tensors is also widely open.



## Acknowledgements

---

The authors are grateful to anonymous reviewers for valuable comments that have significantly improved the manuscript.

## References

---

- Acar, E., S. A. Çamtepe, M. S. Krishnamoorthy, and B. Yener. 2005. “Modeling and multiway analysis of chatroom tensors”. In: *Intelligence and Security Informatics*. Springer. 256–268.
- Alain, G. and Y. Bengio. 2012. “What regularized auto-encoders learn from the data generating distribution”. *arXiv preprint arXiv:1211.4246*.
- Allen-Zhu, Z. and Y. Li. 2016a. “Doubly accelerated methods for faster CCA and generalized eigendecomposition”. *arXiv preprint arXiv:1607.06017*.
- Allen-Zhu, Z. and Y. Li. 2016b. “First Efficient Convergence for Streaming k-PCA: a Global, Gap-Free, and Near-Optimal Rate”. *arXiv preprint arXiv:1607.07837*.
- Anandkumar, A., R. Ge, D. Hsu, and S. M. Kakade. 2013. “A Tensor Spectral Approach to Learning Mixed Membership Community Models”. In: *Conference on Learning Theory (COLT)*.
- Anandkumar, A., R. Ge, D. Hsu, S. M. Kakade, and M. Telgarsky. 2014a. “Tensor Methods for Learning Latent Variable Models”. *J. of Machine Learning Research*. 15: 2773–2832.
- Anandkumar, A., Y. Deng, R. Ge, and H. Mobahi. 2017. “Homotopy Analysis for Tensor PCA”. In: *Conference on Learning Theory*.
- Anandkumar, A., D. P. Foster, D. Hsu, S. M. Kakade, and Y.-K. Liu. 2012a. “A Spectral Algorithm for Latent Dirichlet Allocation”. In: *Advances in Neural Information Processing Systems 25*.

- Anandkumar, A., R. Ge, and M. Janzamin. 2014b. “Sample Complexity Analysis for Learning Overcomplete Latent Variable Models through Tensor Methods”. *arXiv preprint arXiv:1408.0553*. Aug.
- Anandkumar, A., D. Hsu, F. Huang, and S. M. Kakade. 2012b. “Learning Mixtures of Tree Graphical Models”. In: *Advances in Neural Information Processing Systems 25*.
- Anandkumar, A., D. Hsu, and S. M. Kakade. 2012c. “A method of moments for mixture models and hidden Markov models”. In: *COLT*.
- Appelhof, C. J. and E. Davidson. 1981. “Strategies for analyzing data from video fluorometric monitoring of liquid chromatographic effluents”. *Analytical Chemistry*. 53(13): 2053–2056.
- Arora, S., R. Ge, T. Ma, and A. Risteski. 2017. “Provable learning of Noisy-or Networks”. In: *Proceedings of the forty-ninth annual ACM symposium on Theory of computing*.
- Astrid, M. and S. Lee. 2017. “CP-decomposition with Tensor Power Method for Convolutional Neural Networks Compression”. *CoRR*. abs/1701.07148.
- Austin, T. 2008. “On exchangeable random variables and the statistics of large graphs and hypergraphs”. *Probab. Survey*. 5: 80–145.
- Barak, B. and A. Moitra. 2016. “Noisy tensor completion via the sum-of-squares hierarchy”. In: *Conference on Learning Theory*. 417–445.
- Barak, B. and D. Steurer. 2014. “Sum-of-squares proofs and the quest toward optimal algorithms”. *arXiv preprint arXiv:1404.5236*.
- Baum, L. E. and T. Petrie. 1966. “Statistical inference for probabilistic functions of finite state Markov chains”. *The annals of mathematical statistics*. 37(6): 1554–1563.
- Bhaskara, A., M. Charikar, A. Moitra, and A. Vijayaraghavan. 2014. “Smoothed analysis of tensor decompositions”. In: *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*. ACM. 594–603.
- Bjerhammar, A. 1951. *Application of calculus of matrices to method of least squares: with special reference to geodetic calculations*. Elander.
- Blei, D. M., A. Y. Ng, and M. I. Jordan. 2003. “Latent dirichlet allocation”. *Journal of machine Learning research*. 3(Jan): 993–1022.
- Blum, A., J. Hopcroft, and R. Kannan. 2016. “Foundations of data science”. *Vorabversion eines Lehrbuchs*.

- Cardoso, J.-F. and P. Comon. 1996. "Independent Component Analysis, A Survey Of Some Algebraic Methods". In: *IEEE International Symposium on Circuits and Systems*. 93–96.
- Carroll, J. D. and J.-J. Chang. 1970. "Analysis of individual differences in multidimensional scaling via an N-way generalization of "Eckart-Young" decomposition". *Psychometrika*. 35(3): 283–319.
- Chang, J. T. 1996. "Full reconstruction of Markov models on evolutionary trees: Identifiability and consistency". *Mathematical Biosciences*. 137: 51–73.
- Comon, P. 1994. "Independent Component Analysis, a new concept?". *Signal Processing*. 36(3): 287–314.
- Comon, P. 2002. "Tensor decompositions". *Mathematics in Signal Processing V*: 1–24.
- Comon, P. and C. Jutten. 2010. *Handbook of Blind Source Separation: Independent Component Analysis and Applications*. Academic Press. Elsevier.
- Davis, C. and W. M. Kahan. 1970. "The rotation of eigenvectors by a perturbation. III". *SIAM Journal on Numerical Analysis*. 7(1): 1–46.
- De Lathauwer, L., J. Castaing, and J.-F. Cardoso. 2007. "Fourth-order cumulant-based blind identification of underdetermined mixtures". *Signal Processing, IEEE Transactions on*. 55(6): 2965–2973.
- Debals, O. and L. De Lathauwer. 2017. "The concept of tensorization". *Tech. rep.* Technical Report 17–99, ESAT–STADIUS, KU Leuven, Belgium.
- Delfosse, N. and P. Loubaton. 1995. "Adaptive blind separation of independent sources: a deflation approach". *Signal processing*. 45(1): 59–83.
- Eckart, C. and G. Young. 1936. "The approximation of one matrix by another of lower rank". *Psychometrika*. 1(3): 211–218.
- Feige, U. 2002. "Relations between average case complexity and approximation complexity". In: *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*. ACM. 534–543.
- Friedman, J. H. 1987. "Exploratory projection pursuit". *Journal of the American statistical association*. 82(397): 249–266.

- Frieze, A., M. Jerrum, and R. Kannan. 1996. "Learning linear transformations". In: *Foundations of Computer Science, 1996. Proceedings., 37th Annual Symposium on*. IEEE. 359–368.
- Ge, R., F. Huang, C. Jin, and Y. Yuan. 2015a. "Escaping from saddle points—online stochastic gradient for tensor decomposition". In: *Conference on Learning Theory*. 797–842.
- Ge, R., Q. Huang, and S. M. Kakade. 2015b. "Learning mixtures of gaussians in high dimensions". In: *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*. ACM. 761–770.
- Ge, R., C. Jin, P. Netrapalli, A. Sidford, *et al.* 2016. "Efficient algorithms for large-scale generalized eigenvector computation and canonical correlation analysis". In: *International Conference on Machine Learning*. 2741–2750.
- Ge, R. and T. Ma. 2015. "Decomposing overcomplete 3rd order tensors using sum-of-squares algorithms". In: *RANDOM*.
- Golub, G. H. and C. F. van Loan. 1996. *Matrix Computations*. Johns Hopkins University Press.
- Golub, G. H. and C. F. Van Loan. 1990. "Matrix computations".
- Grigoriev, D. 2001. "Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity". *Theoretical Computer Science*. 259(1-2): 613–622.
- Harshman, R. A. 1970. "Foundations of the PARAFAC procedure: models and conditions for an "explanatory" multimodal factor analysis".
- Harshman, R. A. and M. E. Lundy. 1994. "PARAFAC: Parallel factor analysis". *Computational Statistics & Data Analysis*. 18(1): 39–72.
- Hillar, C. J. and L.-H. Lim. 2013. "Most tensor problems are NP-hard". *Journal of the ACM (JACM)*. 60(6): 45.
- Hitchcock, F. L. 1927. "The expression of a tensor or a polyadic as a sum of products". *Journal of Mathematics and Physics*. 6(1-4): 164–189.
- Hofmann, T. 1999. "Probabilistic latent semantic analysis". In: *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc. 289–296.
- Hopkins, S. B., J. Shi, and D. Steurer. 2015. "Tensor principal component analysis via sum-of-square proofs". In: *Conference on Learning Theory*. 956–1006.

- Horn, R. A. and C. R. Johnson. 2012. *Matrix analysis*. Cambridge university press.
- Hotelling, H. 1933. “Analysis of a complex of statistical variables into principal components.” *Journal of educational psychology*. 24(6): 417.
- Hotelling, H. 1992. “Relations between two sets of variates”. In: *Breakthroughs in statistics*. Springer. 162–190.
- Hsu, D. and S. M. Kakade. 2013. “Learning mixtures of spherical Gaussians: moment methods and spectral decompositions”. In: *Fourth Innovations in Theoretical Computer Science*.
- Hsu, D., S. M. Kakade, and P. Liang. 2012. “Identifiability and unmixing of latent parse trees”. In: *Advances in Neural Information Processing Systems 25*.
- Hunter, J. D. 2007. “Matplotlib: A 2D Graphics Environment”. *Computing in Science Engineering*. 9(3): 90–95.
- Hyvärinen, A. and E. Oja. 2000. “Independent component analysis: algorithms and applications”. *Neural Networks*. 13(4–5): 411–430.
- Hyvärinen, A. 2005. “Estimation of non-normalized statistical models by score matching”. In: *Journal of Machine Learning Research*. 695–709.
- Janzamin, M., H. Sedghi, and A. Anandkumar. 2014. “Score Function Features for Discriminative Learning: Matrix and Tensor Frameworks”. *arXiv preprint arXiv:1412.2863*. Dec.
- Janzamin, M., H. Sedghi, and A. Anandkumar. 2015. “Beating the perils of non-convexity: Guaranteed training of neural networks using tensor methods”. *arXiv preprint arXiv:1506.08473*.
- Johnson, R. and T. Zhang. 2013. “Accelerating stochastic gradient descent using predictive variance reduction”. In: *Advances in neural information processing systems*. 315–323.
- Jones, E., T. Oliphant, P. Peterson, *et al.* 2001. “SciPy: Open source scientific tools for Python”. [Online; accessed 2016-10-21]. URL: <http://www.scipy.org/>.
- Kim, Y., E. Park, S. Yoo, T. Choi, L. Yang, and D. Shin. 2016. “Compression of Deep Convolutional Neural Networks for Fast and Low Power Mobile Applications”. *ICLR*. May.

- Koivunen, A. and A. Kostinski. 1999. “The feasibility of data whitening to improve performance of weather radar”. *Journal of Applied Meteorology*. 38(6): 741–749.
- Kolbeinsson, A., J. Kossaifi, Y. Panagakis, A. Anandkumar, I. Tzoulaki, and P. Matthews. 2019. “Stochastically Rank-Regularized Tensor Regression Networks”. *CoRR*. abs/1902.10758.
- Kolda, T. G. and J. R. Mayo. 2011. “Shifted Power Method for Computing Tensor Eigenpairs”. *SIAM Journal on Matrix Analysis and Applications*. 32(4): 1095–1124.
- Kolda, T. G. and B. W. Bader. 2009. “Tensor decompositions and applications”. *SIAM review*. 51(3): 455–500.
- Kossaifi, J., A. Bulat, G. Tzimiropoulos, and M. Pantic. 2019a. “T-Net: Parametrizing Fully Convolutional Nets with a Single High-Order Tensor”. In: *CVPR*. 7822–7831.
- Kossaifi, J., A. Khanna, Z. Lipton, T. Furlanello, and A. Anandkumar. 2017. “Tensor contraction layers for parsimonious deep nets”. In: *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*. IEEE. 1940–1946.
- Kossaifi, J., Z. C. Lipton, A. Khanna, T. Furlanello, and A. Anandkumar. 2018. “Tensor Regression Networks”. *CoRR*. abs/1707.08308.
- Kossaifi, J., Y. Panagakis, A. Anandkumar, and M. Pantic. 2019b. “TensorLy: Tensor Learning in Python”. *Journal of Machine Learning Research*. 20(26): 1–6. URL: <http://jmlr.org/papers/v20/18-277.html>.
- Kruskal, J. 1976. “More factors than subjects, tests and treatments: an indeterminacy theorem for canonical decomposition and individual differences scaling”. *Psychometrika*. 41(3): 281–293.
- Kruskal, J. 1977. “Three-way arrays: Rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics”. *Linear algebra and its applications*. 18(2): 95–138.
- Latała, R. 2005. “Some estimates of norms of random matrices”. *Proceedings of the American Mathematical Society*. 133(5): 1273–1282.
- Lathauwer, L. D., B. D. Moor, and J. Vandewalle. 2000. “On the Best rank-1 and Rank- $(R_1, R_2, \dots, R_N)$  Approximation and Applications of Higher-Order Tensors”. *SIAM J. Matrix Anal. Appl.* 21(4): 1324–1342.

- Lebedev, V., Y. Ganin, M. Rakhuba, I. V. Oseledets, and V. S. Lempitsky. 2015. “Speeding-up Convolutional Neural Networks Using Fine-tuned CP-Decomposition”. In: *ICLR*.
- Leurgans, S., R. Ross, and R. Abel. 1993. “A decomposition for three-way arrays”. *SIAM Journal on Matrix Analysis and Applications*. 14(4): 1064–1083.
- Lim, L.-H. 2005. “Singular values and eigenvalues of tensors: a variational approach”. *Proceedings of the IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP '05)*. 1: 129–132.
- Ma, T., J. Shi, and D. Steurer. 2016. “Polynomial-time tensor decompositions with sum-of-squares”. In: *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*. IEEE. 438–446.
- MacQueen, J. B. 1967. “Some Methods for Classification and Analysis of Multivariate Observations”. In: *Proceedings of the fifth Berkeley Symposium on Mathematical Statistics and Probability*. Vol. 1. University of California Press. 281–297.
- Mocks, J. 1988. “Topographic components model for event-related potentials and some biophysical considerations”. *IEEE transactions on biomedical engineering*. 6(35): 482–484.
- Moore, E. H. 1920. “On the reciprocal of the general algebraic matrix”. *Bull. Am. Math. Soc.* 26: 394–395.
- Mossel, E. and S. Roch. 2006. “Learning Nonsingular Phylogenies and Hidden Markov Models”. *Annals of Applied Probability*. 16(2): 583–614.
- Nguyen, N. H., P. Drineas, and T. D. Tran. 2010. “Tensor sparsification via a bound on the spectral norm of random tensors”. *arXiv preprint arXiv:1005.4732*.
- Nocedal, J. and S. J. Wright. 1999. *Numerical Optimization*. Springer.
- Novikov, A., D. Podoprikin, A. Osokin, and D. Vetrov. 2015. “Tensorizing Neural Networks”. In: *Proceedings of the 28th International Conference on Neural Information Processing Systems. NIPS'15*. Montreal, Canada. 442–450.
- Papadimitriou, C. H., P. Raghavan, H. Tamaki, and S. Vempala. 2000. “Latent semantic indexing: A probabilistic analysis”. *Journal of Computer and System Sciences*. 61(2): 217–235.



- Papalexakis, E. E., C. Faloutsos, and N. D. Sidiropoulos. 2017. “Tensors for data mining and data fusion: Models, applications, and scalable algorithms”. *ACM Transactions on Intelligent Systems and Technology (TIST)*. 8(2): 16.
- Paszke, A., S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. 2017. “Automatic Differentiation in PyTorch”. In: *NIPS Autodiff Workshop*.
- Pearson, K. 1894. “Contributions to the mathematical theory of evolution”. *Philosophical Transactions of the Royal Society of London. A*. 185: 71–110.
- Pearson, K. 1901. “LIII. On lines and planes of closest fit to systems of points in space”. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*. 2(11): 559–572.
- Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. “Scikit-learn: Machine Learning in Python”. *Journal of Machine Learning Research*. 12: 2825–2830.
- Penrose, R. 1955. “A generalized inverse for matrices”. In: *Mathematical proceedings of the Cambridge philosophical society*. Vol. 51. No. 3. Cambridge University Press. 406–413.
- Potechin, A. and D. Steurer. 2017. “Exact tensor completion with sum-of-squares”. *arXiv preprint arXiv:1702.06237*.
- Qi, L. 2005. “Eigenvalues of a real supersymmetric tensor”. *Journal of Symbolic Computation*. 40(6): 1302–1324.
- Rabanser, S., O. Shchur, and S. Günnemann. 2017. “Introduction to tensor decompositions and their applications in machine learning”. *arXiv preprint arXiv:1711.10781*.
- Raz, R. 2013. “Tensor-rank and lower bounds for arithmetic formulas”. *Journal of the ACM (JACM)*. 60(6): 40.
- Richard, E. and A. Montanari. 2014. “A statistical model for tensor PCA”. In: *Advances in Neural Information Processing Systems*. 2897–2905.
- Schoenebeck, G. 2008. “Linear level Lasserre lower bounds for certain k-CSPs”. In: *Foundations of Computer Science, 2008. FOCS’08. IEEE 49th Annual IEEE Symposium on*. IEEE. 593–602.

- Sedghi, H., M. Janzamin, and A. Anandkumar. 2016. "Provable tensor methods for learning mixtures of generalized linear models". In: *Artificial Intelligence and Statistics*. 1223–1231.
- Shalev-Shwartz, S. and T. Zhang. 2013. "Stochastic dual coordinate ascent methods for regularized loss minimization". *Journal of Machine Learning Research*. 14(Feb): 567–599.
- Shashua, A. and A. Levin. 2001. "Linear image coding for regression and classification using the tensor-rank principle". In: *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*. Vol. 1. IEEE. I–42.
- Shwe, M. A., B. Middleton, D. Heckerman, M. Henrion, E. Horvitz, H. Lehmann, and G. Cooper. 1991. "Probabilistic diagnosis using a reformulation of the INTERNIST-1/QMR knowledge base". *Methods of information in Medicine*. 30(4): 241–255.
- Sidiropoulos, N. D. and R. Bro. 2000. "On the uniqueness of multilinear decomposition of N-way arrays". *Journal of Chemometrics*. 14(3): 229–239.
- Sidiropoulos, N. D., R. Bro, and G. B. Giannakis. 2000. "Parallel factor analysis in sensor array processing". *Signal Processing, IEEE Transactions on*. 48(8): 2377–2388.
- Sidiropoulos, N. D., L. De Lathauwer, X. Fu, K. Huang, E. E. Papalexakis, and C. Faloutsos. 2017. "Tensor decomposition for signal processing and machine learning". *IEEE Transactions on Signal Processing*. 65(13): 3551–3582.
- Spearman, C. 1904. "General Intelligence," Objectively Determined and Measured". *The American Journal of Psychology*. 15(2): 201–292.
- Sriperumbudur, B., K. Fukumizu, R. Kumar, A. Gretton, and A. Hyvärinen. 2013. "Density estimation in infinite dimensional exponential families". *arXiv preprint arXiv:1312.3516*.
- stewart, G. and J.-G. Sun. 1990. *Matrix perturbation theory*. Academic Press.
- Swersky, K., D. Buchman, N. D. Freitas, B. M. Marlin, *et al.* 2011. "On autoencoders and score matching for energy based models". In: *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*. 1201–1208.

- Tai, C., T. Xiao, X. Wang, and W. E. 2016. “Convolutional neural networks with low-rank regularization”. *ICLR*.
- Tropp, J. A. 2012. “User-friendly tail bounds for sums of random matrices”. *Foundations of computational mathematics*. 12(4): 389–434.
- Wainwright, M. and M. Jordan. 2008. “Graphical models, exponential families, and variational inference”. *Foundations and Trends® in Machine Learning*. 1(1-2): 1–305.
- Walt, S. van der, S. C. Colbert, and G. Varoquaux. 2011. “The NumPy Array: A Structure for Efficient Numerical Computation”. *Computing in Science Engineering*. 13(2): 22–30.
- Wang, W., J. Wang, and N. Srebro. 2016. “Globally convergent stochastic optimization for canonical correlation analysis”. *Advances in Neural Information Processing Systems*.
- Wedin, P. 1972. “Perturbation bounds in connection with singular value decomposition”. *BIT Numerical Mathematics*. 12(1): 99–111.
- Weyl, H. 1912. “Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen (mit einer Anwendung auf die Theorie der Hohlraumstrahlung)”. *Mathematische Annalen*. 71(4): 441–479.
- Zhang, T. and G. Golub. 2001. “Rank-one approximation to high order tensors”. *SIAM Journal on Matrix Analysis and Applications*. 23: 534–550.