

Tensor Decompositions for Learning Latent Variable Models

Animashree Anandkumar

*Electrical Engineering and Computer Science
University of California, Irvine
2200 Engineering Hall
Irvine, CA 92697*

A.ANANDKUMAR@UCI.EDU

Rong Ge

*Microsoft Research
One Memorial Drive
Cambridge, MA 02142*

RONGGE@MICROSOFT.COM

Daniel Hsu

*Department of Computer Science
Columbia University
1214 Amsterdam Avenue, #0401
New York, NY 10027*

DJHSU@CS.COLUMBIA.EDU

Sham M. Kakade

*Microsoft Research
One Memorial Drive
Cambridge, MA 02142*

SKAKADE@MICROSOFT.COM

Matus Telgarsky

*Department of Statistics
Rutgers University
110 Frelinghuysen Road
Piscataway, NJ 08854*

MTELGARS@CS.UCSB.EDU

Editor: Benjamin Recht

Abstract

This work considers a computationally and statistically efficient parameter estimation method for a wide class of latent variable models—including Gaussian mixture models, hidden Markov models, and latent Dirichlet allocation—which exploits a certain tensor structure in their low-order observable moments (typically, of second- and third-order). Specifically, parameter estimation is reduced to the problem of extracting a certain (orthogonal) decomposition of a symmetric tensor derived from the moments; this decomposition can be viewed as a natural generalization of the singular value decomposition for matrices. Although tensor decompositions are generally intractable to compute, the decomposition of these specially structured tensors can be efficiently obtained by a variety of approaches, including power iterations and maximization approaches (similar to the case of matrices). A detailed analysis of a robust tensor power method is provided, establishing an analogue of Wedin’s perturbation theorem for the singular vectors of matrices. This implies a robust and computationally tractable estimation approach for several popular latent variable models.

Keywords: latent variable models, tensor decompositions, mixture models, topic models, method of moments, power method

1. Introduction

The method of moments is a classical parameter estimation technique (Pearson, 1894) from statistics which has proved invaluable in a number of application domains. The basic paradigm is simple and intuitive: (i) compute certain statistics of the data—often empirical moments such as means and correlations—and (ii) find model parameters that give rise to (nearly) the same corresponding population quantities. In a number of cases, the method of moments leads to consistent estimators which can be efficiently computed; this is especially relevant in the context of latent variable models, where standard maximum likelihood approaches are typically computationally prohibitive, and heuristic methods can be unreliable and difficult to validate with high-dimensional data. Furthermore, the method of moments can be viewed as complementary to the maximum likelihood approach; simply taking a single step of Newton-Raphson on the likelihood function starting from the moment based estimator (Le Cam, 1986) often leads to the best of both worlds: a computationally efficient estimator that is (asymptotically) statistically optimal.

The primary difficulty in learning latent variable models is that the latent (hidden) state of the data is not directly observed; rather only observed variables correlated with the hidden state are observed. As such, it is not evident the method of moments should fare any better than maximum likelihood in terms of computational performance: matching the model parameters to the observed moments may involve solving computationally intractable systems of multivariate polynomial equations. Fortunately, for many classes of latent variable models, there is rich structure in low-order moments (typically second- and third-order) which allow for this inverse moment problem to be solved efficiently (Cattell, 1944; Cardoso, 1991; Chang, 1996; Mossel and Roch, 2006; Hsu et al., 2012b; Anandkumar et al., 2012c,a; Hsu and Kakade, 2013). What is more is that these decomposition problems are often amenable to simple and efficient iterative methods, such as gradient descent and the power iteration method.

1.1 Contributions

In this work, we observe that a number of important and well-studied latent variable models—including Gaussian mixture models, hidden Markov models, and Latent Dirichlet allocation—share a certain structure in their low-order moments, and this permits certain tensor decomposition approaches to parameter estimation. In particular, this decomposition can be viewed as a natural generalization of the singular value decomposition for matrices.

While much of this (or similar) structure was implicit in several previous works (Chang, 1996; Mossel and Roch, 2006; Hsu et al., 2012b; Anandkumar et al., 2012c,a; Hsu and Kakade, 2013), here we make the decomposition explicit under a unified framework. Specifically, we express the observable moments as sums of rank-one terms, and reduce the parameter estimation task to the problem of extracting a symmetric orthogonal decomposition of a symmetric tensor derived from these observable moments. The problem can then be solved by a variety of approaches, including fixed-point and variational methods.

One approach for obtaining the orthogonal decomposition is the tensor power method of Lathauwer et al. (2000, Remark 3). We provide a convergence analysis of this method for orthogonally decomposable symmetric tensors, as well as a detailed perturbation analysis for a robust (and a computationally tractable) variant (Theorem 5.1). This perturbation analysis can be viewed as an analogue of Wedin’s perturbation theorem for singular vectors of matrices (Wedin, 1972), providing a bound on the error of the recovered decomposition in terms of the operator norm of the tensor perturbation. This analysis is subtle in at least two ways. First, unlike for matrices (where every matrix has a singular value decomposition), an orthogonal decomposition need not exist for the perturbed tensor. Our robust variant uses random restarts and deflation to extract an approximate decomposition in a computationally tractable manner. Second, the analysis of the deflation steps is non-trivial; a naïve argument would entail error accumulation in each deflation step, which we show can in fact be avoided. When this method is applied for parameter estimation in latent variable models previously discussed, improved sample complexity bounds (over previous work) can be obtained using this perturbation analysis.

Finally, we also address computational issues that arise when applying the tensor decomposition approaches to estimating latent variable models. Specifically, we show that the basic operations of simple iterative approaches (such as the tensor power method) can be efficiently executed in time linear in the dimension of the observations and the size of the training data. For instance, in a topic modeling application, the proposed methods require time linear in the number of words in the vocabulary and in the number of non-zero entries of the term-document matrix. The combination of this computational efficiency and the robustness of the tensor decomposition techniques makes the overall framework a promising approach to parameter estimation for latent variable models.

1.2 Related Work

The connection between tensor decompositions and latent variable models has a long history across many scientific and mathematical disciplines. We review some of the key works that are most closely related to ours.

1.2.1 TENSOR DECOMPOSITIONS

The role of tensor decompositions in the context of latent variable models dates back to early uses in psychometrics (Cattell, 1944). These ideas later gained popularity in chemometrics, and more recently in numerous science and engineering disciplines, including neuroscience, phylogenetics, signal processing, data mining, and computer vision. A thorough survey of these techniques and applications is given by Kolda and Bader (2009). Below, we discuss a few specific connections to two applications in machine learning and statistics, independent component analysis and latent variable models (between which there is also significant overlap).

Tensor decompositions have been used in signal processing and computational neuroscience for blind source separation and independent component analysis (ICA) (Comon and Jutten, 2010). Here, statistically independent non-Gaussian sources are linearly mixed in the observed signal, and the goal is to recover the mixing matrix (and ultimately, the original source signals). A typical solution is to locate projections of the observed signals that

correspond to local extrema of the so-called “contrast functions” which distinguish Gaussian variables from non-Gaussian variables. This method can be effectively implemented using fast descent algorithms (Hyvarinen, 1999). When using the excess kurtosis (*i.e.*, fourth-order cumulant) as the contrast function, this method reduces to a generalization of the power method for symmetric tensors (Lathauwer et al., 2000; Zhang and Golub, 2001; Kofidis and Regalia, 2002). This case is particularly important, since all local extrema of the kurtosis objective correspond to the true sources (under the assumed statistical model) (Delfosse and Loubaton, 1995); the descent methods can therefore be rigorously analyzed, and their computational and statistical complexity can be bounded (Frieze et al., 1996; Nguyen and Regev, 2009; Arora et al., 2012b).

Higher-order tensor decompositions have also been used to develop estimators for commonly used mixture models, hidden Markov models, and other related latent variable models, often using the algebraic procedure of R. Jennrich (as reported in the article of Harshman, 1970), which is based on a simultaneous diagonalization of different ways of flattening a tensor to matrices. Jennrich’s procedure was employed for parameter estimation of discrete Markov models by Chang (1996) via pair-wise and triple-wise probability tables; and it was later used for other latent variable models such as hidden Markov models (HMMs), latent trees, Gaussian mixture models, and topic models such as latent Dirichlet allocation (LDA) by many others (Mossel and Roch, 2006; Hsu et al., 2012b; Anandkumar et al., 2012c,a; Hsu and Kakade, 2013). In these contexts, it is often also possible to establish strong identifiability results, without giving an explicit estimator, by invoking the non-constructive identifiability argument of Kruskal (1977)—see the article by Allman et al. (2009) for several examples.

Related simultaneous diagonalization approaches have also been used for blind source separation and ICA (as discussed above), and a number of efficient algorithms have been developed for this problem (Bunse-Gerstner et al., 1993; Cardoso and Souloumiac, 1993; Cardoso, 1994; Cardoso and Comon, 1996; Corless et al., 1997; Ziehe et al., 2004). A rather different technique that uses tensor flattening and matrix eigenvalue decomposition has been developed by Cardoso (1991) and later by De Lathauwer et al. (2007). A significant advantage of this technique is that it can be used to estimate overcomplete mixtures, where the number of sources is larger than the observed dimension.

The relevance of tensor analysis to latent variable modeling has been long recognized in the field of algebraic statistics (Pachter and Sturmfels, 2005), and many works characterize the algebraic varieties corresponding to the moments of various classes of latent variable models (Drton et al., 2007; Sturmfels and Zwiernik, 2013). These works typically do not address computational or finite sample issues, but rather are concerned with basic questions of identifiability.

The specific tensor structure considered in the present work is the symmetric orthogonal decomposition. This decomposition expresses a tensor as a linear combination of simple tensor forms; each form is the tensor product of a vector (*i.e.*, a rank-1 tensor), and the collection of vectors form an orthonormal basis. An important property of tensors with such decompositions is that they have eigenvectors corresponding to these basis vectors. Although the concepts of eigenvalues and eigenvectors of tensors is generally significantly more complicated than their matrix counterpart—both algebraically (Qi, 2005; Cartwright and Sturmfels, 2013; Lim, 2005) and computationally (Hillar and Lim, 2013; Kofidis and

Regalia, 2002)—the special symmetric orthogonal structure we consider permits simple algorithms to efficiently and stably recover the desired decomposition. In particular, a generalization of the matrix power method to symmetric tensors, introduced by Lathauwer et al. (2000, Remark 3) and analyzed by Kofidis and Regalia (2002), provides such a decomposition. This is in fact implied by the characterization of Zhang and Golub (2001), which shows that iteratively obtaining the best rank-1 approximation of such orthogonally decomposable tensors also yields the exact decomposition. We note that in general, obtaining such approximations for general (symmetric) tensors is NP-hard (Hillar and Lim, 2013).

1.2.2 LATENT VARIABLE MODELS

This work focuses on the particular application of tensor decomposition methods to estimating latent variable models, a significant departure from many previous approaches in the machine learning and statistics literature. By far the most popular heuristic for parameter estimation for such models is the Expectation-Maximization (EM) algorithm (Dempster et al., 1977; Redner and Walker, 1984). Although EM has a number of merits, it may suffer from slow convergence and poor quality local optima (Redner and Walker, 1984), requiring practitioners to employ many additional heuristics to obtain good solutions. For some models such as latent trees (Roch, 2006) and topic models (Arora et al., 2012a), maximum likelihood estimation is NP-hard, which suggests that other estimation approaches may be more attractive. More recently, algorithms from theoretical computer science and machine learning have addressed computational and sample complexity issues related to estimating certain latent variable models such as Gaussian mixture models and HMMs (Dasgupta, 1999; Arora and Kannan, 2005; Dasgupta and Schulman, 2007; Vempala and Wang, 2004; Kannan et al., 2008; Achlioptas and McSherry, 2005; Chaudhuri and Rao, 2008; Brubaker and Vempala, 2008; Kalai et al., 2010; Belkin and Sinha, 2010; Moitra and Valiant, 2010; Hsu and Kakade, 2013; Chang, 1996; Mossel and Roch, 2006; Hsu et al., 2012b; Anandkumar et al., 2012c; Arora et al., 2012a; Anandkumar et al., 2012a). See the works by Anandkumar et al. (2012c) and Hsu and Kakade (2013) for a discussion of these methods, together with the computational and statistical hardness barriers that they face. The present work reviews a broad range of latent variables where a mild non-degeneracy condition implies the symmetric orthogonal decomposition structure in the tensors of low-order observable moments.

Notably, another class of methods, based on subspace identification (Overschee and Moor, 1996) and observable operator models/multiplicity automata (Schützenberger, 1961; Jaeger, 2000; Littman et al., 2001), have been proposed for a number of latent variable models. These methods were successfully developed for HMMs by Hsu et al. (2012b), and subsequently generalized and extended for a number of related sequential and tree Markov models (Siddiqi et al., 2010; Bailly, 2011; Boots et al., 2010; Parikh et al., 2011; Rodu et al., 2013; Balle et al., 2012; Balle and Mohri, 2012), as well as certain classes of parse tree models (Luque et al., 2012; Cohen et al., 2012; Dhillon et al., 2012). These methods use low-order moments to learn an “operator” representation of the distribution, which can be used for density estimation and belief state updates. While finite sample bounds can be given to establish the learnability of these models (Hsu et al., 2012b), the algorithms do not

actually give parameter estimates (*e.g.*, of the emission or transition matrices in the case of HMMs).

1.3 Organization

The rest of the paper is organized as follows. Section 2 reviews some basic definitions of tensors. Section 3 provides examples of a number of latent variable models which, after appropriate manipulations of their low order moments, share a certain natural tensor structure. Section 4 reduces the problem of parameter estimation to that of extracting a certain (symmetric orthogonal) decomposition of a tensor. We then provide a detailed analysis of a robust tensor power method and establish an analogue of Wedin’s perturbation theorem for the singular vectors of matrices. The discussion in Section 6 addresses a number of practical concerns that arise when dealing with moment matrices and tensors.

2. Preliminaries

We introduce some tensor notations borrowed from Lim (2005). A real p -th order tensor $A \in \bigotimes_{i=1}^p \mathbb{R}^{n_i}$ is a member of the tensor product of Euclidean spaces \mathbb{R}^{n_i} , $i \in [p]$. We generally restrict to the case where $n_1 = n_2 = \dots = n_p = n$, and simply write $A \in \bigotimes^p \mathbb{R}^n$. For a vector $v \in \mathbb{R}^n$, we use $v^{\otimes p} := v \otimes v \otimes \dots \otimes v \in \bigotimes^p \mathbb{R}^n$ to denote its p -th tensor power. 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 $[A_{i_1, i_2, \dots, i_p} : i_1, i_2, \dots, i_p \in [n]]$, where A_{i_1, i_2, \dots, i_p} is the (i_1, i_2, \dots, i_p) -th coordinate of A (with respect to a canonical basis).

We can consider A to be a multilinear map in the following sense: for a set of matrices $\{V_i \in \mathbb{R}^{n \times m_i} : i \in [p]\}$, the (i_1, i_2, \dots, i_p) -th entry in the p -way array representation of $A(V_1, V_2, \dots, V_p) \in \mathbb{R}^{m_1 \times m_2 \times \dots \times m_p}$ is

$$[A(V_1, V_2, \dots, V_p)]_{i_1, i_2, \dots, i_p} := \sum_{j_1, j_2, \dots, j_p \in [n]} A_{j_1, j_2, \dots, j_p} [V_1]_{j_1, i_1} [V_2]_{j_2, i_2} \dots [V_p]_{j_p, i_p}.$$

Note that if A is a matrix ($p = 2$), then

$$A(V_1, V_2) = V_1^\top A V_2.$$

Similarly, for a matrix A and vector $v \in \mathbb{R}^n$, we can express Av as

$$A(I, v) = Av \in \mathbb{R}^n,$$

where I is the $n \times n$ identity matrix. As a final example of this notation, observe

$$A(e_{i_1}, e_{i_2}, \dots, e_{i_p}) = A_{i_1, i_2, \dots, i_p},$$

where $\{e_1, e_2, \dots, e_n\}$ is the canonical basis for \mathbb{R}^n .

Most tensors $A \in \bigotimes^p \mathbb{R}^n$ considered in this work will be *symmetric* (sometimes called *supersymmetric*), which means that their p -way array representations are invariant to permutations of the array indices: *i.e.*, for all indices $i_1, i_2, \dots, i_p \in [n]$, $A_{i_1, i_2, \dots, i_p} = A_{i_{\pi(1)}, i_{\pi(2)}, \dots, i_{\pi(p)}}$ for any permutation π on $[p]$. It can be checked that this reduces to the usual definition of a symmetric matrix for $p = 2$.

The *rank* of a p -th order tensor $A \in \otimes^p \mathbb{R}^n$ is the smallest non-negative integer k such that $A = \sum_{j=1}^k u_{1,j} \otimes u_{2,j} \otimes \dots \otimes u_{p,j}$ for some $u_{i,j} \in \mathbb{R}^n, i \in [p], j \in [k]$, and the *symmetric rank* of a symmetric p -th order tensor A is the smallest non-negative integer k such that $A = \sum_{j=1}^k u_j^{\otimes p}$ for some $u_j \in \mathbb{R}^n, j \in [k]$.¹ The notion of rank readily reduces to the usual definition of matrix rank when $p = 2$, as revealed by the singular value decomposition. Similarly, for symmetric matrices, the symmetric rank is equivalent to the matrix rank as given by the spectral theorem. A decomposition into such rank-one terms is known as a *canonical polyadic decomposition* (Hitchcock, 1927a,b).

The notion of tensor (symmetric) rank is considerably more delicate than matrix (symmetric) rank. For instance, it is not clear *a priori* that the symmetric rank of a tensor should even be finite (Comon et al., 2008). In addition, removal of the best rank-1 approximation of a (general) tensor may increase the tensor rank of the residual (Stegeman and Comon, 2010).

Throughout, we use $\|v\| = (\sum_i v_i^2)^{1/2}$ to denote the Euclidean norm of a vector v , and $\|M\|$ to denote the spectral (operator) norm of a matrix. We also use $\|T\|$ to denote the operator norm of a tensor, which we define later.

3. Tensor Structure in Latent Variable Models

In this section, we give several examples of latent variable models whose low-order moments can be written as symmetric tensors of low symmetric rank; some of these examples can be deduced using the techniques developed in the text by McCullagh (1987). The basic form is demonstrated in Theorem 3.1 for the first example, and the general pattern will emerge from subsequent examples.

3.1 Exchangeable Single Topic Models

We first consider a simple bag-of-words model for documents in which the words in the document are assumed to be *exchangeable*. Recall that a collection of random variables x_1, x_2, \dots, x_ℓ 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_ℓ are conditionally i.i.d. given h (see Figure 1(a) for the corresponding graphical model) and the conditional distributions are identical at all the nodes.

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 $w := (w_1, w_2, \dots, w_k) \in \Delta^{k-1}$. This is modeled as a discrete random variable h such that

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

1. For even p , the definition is slightly different (Comon et al., 2008).

Given the topic h , the document's ℓ 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 ℓ words in the document by d -dimensional random vectors $x_1, x_2, \dots, x_\ell \in \mathbb{R}^d$. Specifically, we set

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

where e_1, e_2, \dots, e_d is the standard coordinate basis for \mathbb{R}^d .

One advantage of this encoding of words is that the (cross) moments of these random vectors correspond to joint probabilities over words. For instance, observe that

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

so the (i, j) -th entry of the 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 μ_j , the vector of word probabilities for topic j :

$$\mathbb{E}[x_t | h = j] = \sum_{i=1}^d \Pr[t\text{-th word} = i | h = j] e_i = \sum_{i=1}^d [\mu_j]_i e_i = \mu_j, \quad j \in [k]$$

(where $[\mu_j]_i$ is the i -th entry in the vector μ_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 :

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

This and similar calculations lead one to the following theorem.

Theorem 3.1 (Anandkumar et al., 2012c) *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

$$\begin{aligned} M_2 &= \sum_{i=1}^k w_i \mu_i \otimes \mu_i \\ M_3 &= \sum_{i=1}^k w_i \mu_i \otimes \mu_i \otimes \mu_i. \end{aligned}$$

As we will see in Section 4.3, the structure of M_2 and M_3 revealed in Theorem 3.1 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, all 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); see Section 6.1.

3.2 Beyond Raw Moments

In the single topic model above, 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. Here, we show that the desired tensor structure can be found through various manipulations of different moments.

3.2.1 SPHERICAL GAUSSIAN MIXTURES: COMMON COVARIANCE

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).

Let $w_i \in (0, 1)$ be the probability of choosing component $i \in [k]$, $\{\mu_1, \mu_2, \dots, \mu_k\} \subset \mathbb{R}^d$ be the component mean vectors, and $\sigma^2 I$ be the common covariance matrix. An observation in this model is given by

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

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

The Gaussian mixture 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_ℓ given the same fixed h (the topic of the document). In contrast, for the Gaussian mixture model, every realization of x corresponds to a different realization of h .

Theorem 3.2 (Hsu and Kakade, 2013) *Assume $d \geq k$. The variance σ^2 is the smallest eigenvalue of the covariance matrix $\mathbb{E}[x \otimes x] - \mathbb{E}[x] \otimes \mathbb{E}[x]$. Furthermore, if*

$$\begin{aligned} M_2 &:= \mathbb{E}[x \otimes x] - \sigma^2 I \\ M_3 &:= \mathbb{E}[x \otimes x \otimes x] - \sigma^2 \sum_{i=1}^d (\mathbb{E}[x] \otimes e_i \otimes e_i + e_i \otimes \mathbb{E}[x] \otimes e_i + e_i \otimes e_i \otimes \mathbb{E}[x]), \end{aligned}$$

then

$$\begin{aligned} M_2 &= \sum_{i=1}^k w_i \mu_i \otimes \mu_i \\ M_3 &= \sum_{i=1}^k w_i \mu_i \otimes \mu_i \otimes \mu_i. \end{aligned}$$

3.2.2 SPHERICAL GAUSSIAN MIXTURES: 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 $h = i$ (for some $i \in [k]$) is a multivariate Gaussian $\mathcal{N}(0, \sigma_i^2 I)$ with zero mean and spherical covariance $\sigma_i^2 I$.

Theorem 3.3 (Hsu and Kakade, 2013) *Assume $d \geq k$. The average variance $\bar{\sigma}^2 := \sum_{i=1}^k w_i \sigma_i^2$ 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*

$$\begin{aligned} M_1 &:= \mathbb{E}[x(v^\top(x - \mathbb{E}[x]))^2] \\ M_2 &:= \mathbb{E}[x \otimes x] - \bar{\sigma}^2 I \\ M_3 &:= \mathbb{E}[x \otimes x \otimes x] - \sum_{i=1}^d (M_1 \otimes e_i \otimes e_i + e_i \otimes M_1 \otimes e_i + e_i \otimes e_i \otimes M_1), \end{aligned}$$

then

$$\begin{aligned} M_2 &= \sum_{i=1}^k w_i \mu_i \otimes \mu_i \\ M_3 &= \sum_{i=1}^k w_i \mu_i \otimes \mu_i \otimes \mu_i. \end{aligned}$$

As shown by Hsu and Kakade (2013), $M_1 = \sum_{i=1}^k w_i \sigma_i^2 \mu_i$. Note that for the common covariance case, where $\sigma_i^2 = \sigma^2$, we have that $M_1 = \sigma^2 \mathbb{E}[x]$ (cf. Theorem 3.2).

3.2.3 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 be a multivariate Gaussian random vector. The random vectors h and z are assumed to be independent. The observed random vector is

$$x := Ah + z.$$

Let μ_i denote the i -th column of the mixing matrix A .

Theorem 3.4 (Comon and Jutten, 2010) *Define*

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

where T is the fourth-order tensor with

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

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

$$M_4 = \sum_{i=1}^k \kappa_i \mu_i \otimes \mu_i \otimes \mu_i \otimes \mu_i.$$

Note that κ_i corresponds to the excess kurtosis, a measure of non-Gaussianity as $\kappa_i = 0$ if h_i is a standard normal random variable. Furthermore, note that A is not identifiable if h is a multivariate Gaussian.

We may derive forms similar to that of M_2 and M_3 from Theorem 3.1 using M_4 by observing that

$$M_4(I, I, u, v) = \sum_{i=1}^k \kappa_i (\mu_i^\top u) (\mu_i^\top v) \mu_i \otimes \mu_i,$$

$$M_4(I, I, I, v) = \sum_{i=1}^k \kappa_i (\mu_i^\top v) \mu_i \otimes \mu_i \otimes \mu_i$$

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

3.2.4 LATENT DIRICHLET ALLOCATION (LDA)

An increasingly popular class of latent variable models are *mixed membership models*, where each datum may belong to several different latent classes simultaneously. LDA 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). 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 over the probability simplex $\Delta^{k-1} := \{v \in \mathbb{R}^k : v_i \in [0, 1] \forall i \in [k], \sum_{i=1}^k v_i = 1\}$ is given by

$$p_\alpha(h) = \frac{\Gamma(\alpha_0)}{\prod_{i=1}^k \Gamma(\alpha_i)} \prod_{i=1}^k h_i^{\alpha_i - 1}, \quad h \in \Delta^{k-1}$$

where

$$\alpha_0 := \alpha_1 + \alpha_2 + \dots + \alpha_k.$$

As before, the k topics are specified by probability vectors $\mu_1, \mu_2, \dots, \mu_k \in \Delta^{d-1}$. 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 ℓ words x_1, x_2, \dots, x_ℓ independently from the discrete distribution specified by the probability vector $\sum_{i=1}^k h_i \mu_i$ (i.e., for each x_t , we independently sample a topic j according to h and then sample x_t according to μ_j). Again, we encode a word x_t by setting $x_t = e_i$ iff the t -th word in the document is i .

The parameter α_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 $(1/k, 1/k, \dots, 1/k)$ (furthermore, the distribution behaves like a product distribution). We are typically interested in




Figure 1: Examples of latent variable models.

the case where α_0 is small (*e.g.*, a constant independent of k), whereupon h typically has only a few large entries. This corresponds to the setting where the documents are mainly comprised of just a few topics.

Theorem 3.5 (Anandkumar et al., 2012a) *Define*

$$\begin{aligned}
 M_1 &:= \mathbb{E}[x_1] \\
 M_2 &:= \mathbb{E}[x_1 \otimes x_2] - \frac{\alpha_0}{\alpha_0 + 1} M_1 \otimes M_1 \\
 M_3 &:= \mathbb{E}[x_1 \otimes x_2 \otimes x_3] \\
 &\quad - \frac{\alpha_0}{\alpha_0 + 2} \left(\mathbb{E}[x_1 \otimes x_2 \otimes M_1] + \mathbb{E}[x_1 \otimes M_1 \otimes x_2] + \mathbb{E}[M_1 \otimes x_1 \otimes x_2] \right) \\
 &\quad + \frac{2\alpha_0^2}{(\alpha_0 + 2)(\alpha_0 + 1)} M_1 \otimes M_1 \otimes M_1.
 \end{aligned}$$

Then

$$\begin{aligned}
 M_2 &= \sum_{i=1}^k \frac{\alpha_i}{(\alpha_0 + 1)\alpha_0} \mu_i \otimes \mu_i \\
 M_3 &= \sum_{i=1}^k \frac{2\alpha_i}{(\alpha_0 + 2)(\alpha_0 + 1)\alpha_0} \mu_i \otimes \mu_i \otimes \mu_i.
 \end{aligned}$$

Note that α_0 needs to be known to form M_2 and M_3 from the raw moments. This, however, is a much weaker than assuming that the entire distribution of h is known (*i.e.*, knowledge of the whole parameter vector α).

3.3 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_ℓ 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 (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., 2012a).

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 , and

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

where the $\mu_{t,j} \in \mathbb{R}^{d_t}$ are the conditional means of the x_t given $h = j$. Thus, we allow the observations x_1, x_2, \dots, x_ℓ 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 3.1 *We have that:*

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

The cross moments do not possess a symmetric tensor form when the conditional distributions are different. 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 3.6 (Anandkumar et al., 2012a) *Assume that $\{\mu_{v,1}, \mu_{v,2}, \dots, \mu_{v,k}\}$ are linearly independent for each $v \in \{1, 2, 3\}$. Define*

$$\begin{aligned} \tilde{x}_1 &:= \mathbb{E}[x_3 \otimes x_2] \mathbb{E}[x_1 \otimes x_2]^{-1} x_1 \\ \tilde{x}_2 &:= \mathbb{E}[x_3 \otimes x_1] \mathbb{E}[x_2 \otimes x_1]^{-1} x_2 \\ M_2 &:= \mathbb{E}[\tilde{x}_1 \otimes \tilde{x}_2] \\ M_3 &:= \mathbb{E}[\tilde{x}_1 \otimes \tilde{x}_2 \otimes x_3]. \end{aligned}$$

Then

$$\begin{aligned} M_2 &= \sum_{i=1}^k w_i \mu_{3,i} \otimes \mu_{3,i} \\ M_3 &= \sum_{i=1}^k w_i \mu_{3,i} \otimes \mu_{3,i} \otimes \mu_{3,i}. \end{aligned}$$

We now discuss three examples (taken mostly 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.

3.3.1 MIXTURES OF AXIS-ALIGNED GAUSSIANS AND OTHER PRODUCT DISTRIBUTIONS

The first example is a mixture of k product distributions in \mathbb{R}^n 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 $A := [\mu_1 | \mu_2 | \cdots | \mu_k] \in \mathbb{R}^{n \times k}$ 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 $[n]$ into three groups creates three (possibly asymmetric) “views” which are conditionally independent once the mixture component is selected. However, recall that Theorem 3.6 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.

Define $\text{coherence}(A) := \max_{i \in [n]} \{e_i^\top \Pi_A e_i\}$ to be the largest diagonal entry of the orthogonal projector to the range of A , and assume A has rank k . The coherence lies between k/n and 1; it is largest when the range of A is spanned by the coordinate axes, and it is k/n when the range is spanned by a subset of the Hadamard basis of cardinality k . The incoherence condition requires, 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 n dimensions. Operationally, it implies the following.

Proposition 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 $[n]$ into three groups (for each $i \in [n]$, 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 μ_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 3.6.

3.3.2 SPHERICAL GAUSSIAN MIXTURES, REVISITED

Consider again the case of spherical Gaussian mixtures (cf. Section 3.2). As we shall see in Section 4.3, the previous techniques (based on Theorem 3.2 and Theorem 3.3) lead to estimation procedures when the dimension of x is k or greater (and when 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 $x \in \mathbb{R}^n$ into three views $\tilde{x}_1, \tilde{x}_2, \tilde{x}_3 \in \mathbb{R}^d$ for $d := n/3$. 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 $A := [\mu_1 | \mu_2 | \dots | \mu_k] \in \mathbb{R}^{n \times k}$ has rank k and $n \geq 3k$. To see this, observe that a random rotation in \mathbb{R}^n followed by a restriction to d coordinates is simply a random projection from \mathbb{R}^n to \mathbb{R}^d , and that a random projection of a linear subspace of dimension k to \mathbb{R}^d is almost surely injective as long as $d \geq k$. Applying this observation to the range of A implies the following.

Proposition 3.3 (Hsu and Kakade, 2013) *Assume A has rank k and that $n \geq 3k$. Let $R \in \mathbb{R}^{n \times n}$ be chosen uniformly at random among all orthogonal $n \times n$ matrices, and set $\tilde{x} := Rx \in \mathbb{R}^n$ and $\tilde{A} := RA = [R\mu_1 | R\mu_2 | \dots | R\mu_k] \in \mathbb{R}^{n \times k}$. Partition $[n]$ 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 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 3.2. With this approach, it is sufficient to require $n = O(k \log k)$ (for constant ε and δ), 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.

3.3.3 HIDDEN MARKOV MODELS

Our last example is the time-homogeneous HMM 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 observation x_t at time t (a random vector taking values in \mathbb{R}^d) is independent of all other observations and hidden states. See Figure 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: 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 ,

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

Proposition 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 $j \in [k]$,

$$\begin{aligned}\mathbb{E}[x_1|h = j] &= O \operatorname{diag}(\pi)T^\top \operatorname{diag}(w)^{-1}e_j \\ \mathbb{E}[x_2|h = j] &= Oe_j \\ \mathbb{E}[x_3|h = j] &= OTe_j.\end{aligned}$$

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) π and $T\pi$ have positive entries.

4. Orthogonal Tensor Decompositions

We now show how recovering the μ_i 's in our aforementioned problems reduces to the problem of finding a certain orthogonal tensor decomposition of a symmetric tensor. We start by reviewing the spectral decomposition of symmetric matrices, and then discuss a generalization to the higher-order tensor case. Finally, we show how orthogonal tensor decompositions can be used for estimating the latent variable models from the previous section.

4.1 Review: The Matrix Case

We first build intuition by reviewing the matrix setting, where the desired decomposition is the eigendecomposition of a symmetric rank- k matrix $M = V\Lambda V^\top$, where $V = [v_1|v_2|\dots|v_k] \in \mathbb{R}^{n \times k}$ is the matrix with orthonormal eigenvectors as columns, and $\Lambda = \operatorname{diag}(\lambda_1, \lambda_2, \dots, \lambda_k) \in \mathbb{R}^{k \times k}$ is diagonal matrix of non-zero eigenvalues. In other words,

$$M = \sum_{i=1}^k \lambda_i v_i v_i^\top = \sum_{i=1}^k \lambda_i v_i^{\otimes 2}. \quad (1)$$

Such a decomposition is guaranteed to exist for every symmetric matrix.

Recovery of the v_i 's and λ_i 's can be viewed at least two ways. First, each v_i is fixed under the mapping $u \mapsto Mu$, up to a scaling factor λ_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 similarly fixed under M . However, in this case, the decomposition in (1) 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_j 's are the only directions fixed under $u \mapsto Mu$ up to non-trivial scaling.

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*

$$u \mapsto \frac{u^\top Mu}{u^\top u}$$

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

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

and observe that v_s is the maximizer.

Efficient algorithms for finding these matrix decompositions are well studied (Golub and van 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

$$\begin{aligned} u \mapsto Mu &\equiv u \mapsto M(I, u), & (2) \\ u \mapsto \frac{u^\top Mu}{u^\top u} &\equiv u \mapsto \frac{M(u, u)}{u^\top u}. & (3) \end{aligned}$$

4.2 The Tensor Case

Decomposing general tensors is a delicate issue; tensors may not even have unique decompositions. Fortunately, the orthogonal tensors that arise in the aforementioned models have a structure which permits a unique decomposition under a mild non-degeneracy condition. We focus our attention to the case $p = 3$, *i.e.*, a third order tensor; the ideas extend to general p with minor modifications.

An *orthogonal decomposition* of a symmetric tensor $T \in \otimes^3 \mathbb{R}^n$ is a collection of orthonormal (unit) vectors $\{v_1, v_2, \dots, v_k\}$ together with corresponding positive scalars $\lambda_i > 0$ such that

$$T = \sum_{i=1}^k \lambda_i v_i^{\otimes 3}. \tag{4}$$

Note that since we are focusing on odd-order tensors ($p = 3$), we have added the requirement that the λ_i be positive. This convention can be followed without loss of generality since $-\lambda_i v_i^{\otimes p} = \lambda_i (-v_i)^{\otimes p}$ 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).

4.2.1 FIXED-POINT CHARACTERIZATION

For a tensor T , consider the vector-valued map

$$u \mapsto T(I, u, u) \tag{5}$$

which is the third-order generalization of (2). This can be explicitly written as

$$T(I, u, u) = \sum_{i=1}^d \sum_{1 \leq j, l \leq d} T_{i,j,l} (e_j^\top u) (e_l^\top u) e_i.$$

Observe that (5) is *not* a linear map, which is a key difference compared to the matrix case.

An eigenvector u for a matrix M satisfies $M(I, u) = \lambda u$, for some scalar λ . We say a unit vector $u \in \mathbb{R}^n$ 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). For orthogonally decomposable tensors $T = \sum_{i=1}^k \lambda_i v_i^{\otimes 3}$,

$$T(I, u, u) = \sum_{i=1}^k \lambda_i (u^\top v_i)^2 v_i .$$

By the orthogonality of the v_i , it is clear that $T(I, v_i, v_i) = \lambda_i v_i$ for all $i \in [k]$. Therefore each (v_i, λ_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 (5). First, even with the multiplicity $\lambda_1 = \lambda_2 = \lambda$, a linear combination $u := c_1 v_1 + c_2 v_2$ may *not* be an eigenvector. 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 $c_1 v_1 + c_2 v_2$. 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. More generally, for any subset $S \subseteq [k]$, the vector

$$\sum_{i \in S} \frac{1}{\lambda_i} \cdot v_i$$

is (proportional to) an eigenvector.

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 $\theta \in \{u' \in \mathbb{R}^n : \|u' - u\| \leq \epsilon\}$, repeated iteration of the map

$$\bar{\theta} \mapsto \frac{T(I, \bar{\theta}, \bar{\theta})}{\|T(I, \bar{\theta}, \bar{\theta})\|}, \tag{6}$$

starting from θ converges to u . Note that the map (6) rescales the output to have unit Euclidean norm. Robust eigenvectors are also called attracting fixed points of (6) (see, *e.g.*, Kolda and Mayo, 2011).

The following theorem implies that if T has an orthogonal decomposition as given in (4), 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 .)

Theorem 4.1 *Let T have an orthogonal decomposition as given in (4).*

1. *The set of $\theta \in \mathbb{R}^n$ which do not converge to some v_i under repeated iteration of (6) has measure zero.*
2. *The set of robust eigenvectors of T is equal to $\{v_1, v_2, \dots, v_k\}$.*

The proof of Theorem 4.1 is given in Appendix A.1, and 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 (6).

4.2.2 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

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

which can be compared to (3). For an orthogonally decomposable tensor, the following theorem shows that a non-zero vector $u \in \mathbb{R}^n$ 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 4.2 *Let T have an orthogonal decomposition as given in (4), and consider the optimization problem*

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

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 $i \in [k]$.*

The proof of Theorem 4.2 is given in Appendix A.2. It is similar to local optimality analysis for ICA methods using fourth-order cumulants (*e.g.*, 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 , (i) the local maximizers of the objective function $u \mapsto T(u, u, u)/(u^\top u)^{3/2}$ correspond precisely to the vectors v_i in the decomposition, and (ii) these local maximizers can be reliably identified using a simple fixed-point iteration (*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.

4.3 Estimation via Orthogonal Tensor Decompositions

We now demonstrate how the moment tensors obtained for various latent variable models in Section 3 can be reduced to an orthogonal form. For concreteness, we take the specific form from the exchangeable single topic model (Theorem 3.1):

$$M_2 = \sum_{i=1}^k w_i \mu_i \otimes \mu_i,$$

$$M_3 = \sum_{i=1}^k w_i \mu_i \otimes \mu_i \otimes \mu_i.$$

(The more general case allows the weights w_i in M_2 to differ in M_3 , but for simplicity we keep them the same in the following discussion.) We now show how to reduce these forms to an orthogonally decomposable tensor from which the w_i and μ_i can be recovered. See Appendix D for a discussion as to how previous approaches (Mossel and Roch, 2006; Anandkumar et al., 2012c,a; Hsu and Kakade, 2013) achieved this decomposition through a certain simultaneous diagonalization method.

Throughout, we assume the following non-degeneracy condition.

Condition 4.1 (Non-degeneracy) *The vectors $\mu_1, \mu_2, \dots, \mu_k \in \mathbb{R}^d$ are linearly independent, and the scalars $w_1, w_2, \dots, w_k > 0$ are strictly positive.*

Observe that Condition 4.1 implies that $M_2 \succeq 0$ is positive semidefinite and has rank k . This is often a mild condition in applications. When this condition is not met, learning is conjectured to be generally hard for both computational (Mossel and Roch, 2006) and information-theoretic reasons (Moitra and Valiant, 2010). As discussed by Hsu et al. (2012b) and Hsu and Kakade (2013), when the non-degeneracy condition does not hold, it is often possible to combine multiple observations using tensor products to increase the rank of the relevant matrices. Indeed, this observation has been rigorously formulated in very recent works of Bhaskara et al. (2014) and Anderson et al. (2014) using the framework of smoothed analysis (Spielman and Teng, 2009).

4.3.1 THE REDUCTION

First, let $W \in \mathbb{R}^{d \times k}$ be a linear transformation such that

$$M_2(W, W) = W^\top M_2 W = I$$

where I is the $k \times k$ identity matrix (*i.e.*, W whitens M_2). Since $M_2 \succeq 0$, we may for concreteness take $W := UD^{-1/2}$, where $U \in \mathbb{R}^{d \times k}$ is the matrix of orthonormal eigenvectors of M_2 , and $D \in \mathbb{R}^{k \times k}$ is the diagonal matrix of positive eigenvalues of M_2 . Let

$$\tilde{\mu}_i := \sqrt{w_i} W^\top \mu_i.$$

Observe that

$$M_2(W, W) = \sum_{i=1}^k W^\top (\sqrt{w_i} \mu_i) (\sqrt{w_i} \mu_i)^\top W = \sum_{i=1}^k \tilde{\mu}_i \tilde{\mu}_i^\top = I,$$

so the $\tilde{\mu}_i \in \mathbb{R}^k$ are orthonormal vectors.

Now define $\widetilde{M}_3 := M_3(W, W, W) \in \mathbb{R}^{k \times k \times k}$, so that

$$\widetilde{M}_3 = \sum_{i=1}^k w_i (W^\top \mu_i)^{\otimes 3} = \sum_{i=1}^k \frac{1}{\sqrt{w_i}} \tilde{\mu}_i^{\otimes 3}.$$

As the following theorem shows, the orthogonal decomposition of \widetilde{M}_3 can be obtained by identifying its robust eigenvectors, upon which the original parameters w_i and μ_i can be recovered. For simplicity, we only state the result in terms of robust eigenvector/eigenvalue pairs; one may also easily state everything in variational form using Theorem 4.2.

Theorem 4.3 *Assume Condition 4.1 and take \widetilde{M}_3 as defined above.*

1. *The set of robust eigenvectors of \widetilde{M}_3 is equal to $\{\tilde{\mu}_1, \tilde{\mu}_2, \dots, \tilde{\mu}_k\}$.*
2. *The eigenvalue corresponding to the robust eigenvector $\tilde{\mu}_i$ of \widetilde{M}_3 is equal to $1/\sqrt{w_i}$, for all $i \in [k]$.*
3. *If $B \in \mathbb{R}^{d \times k}$ is the Moore-Penrose pseudoinverse of W^\top , and (v, λ) is a robust eigenvector/eigenvalue pair of \widetilde{M}_3 , then $\lambda Bv = \mu_i$ for some $i \in [k]$.*

The theorem follows by combining the above discussion with the robust eigenvector characterization of Theorem 4.1. Recall that we have taken as convention that eigenvectors have unit norm, so the μ_i are exactly determined from the robust eigenvector/eigenvalue pairs of \widetilde{M}_3 (together with the pseudoinverse of W^\top); in particular, the scale of each μ_i is correctly identified (along with the corresponding w_i). Relative to previous works on moment-based estimators for latent variable models (*e.g.*, Anandkumar et al., 2012c,a; Hsu and Kakade, 2013), Theorem 4.3 emphasizes the role of the special tensor structure, which in turn makes transparent the applicability of methods for orthogonal tensor decomposition.

4.3.2 LOCAL MAXIMIZERS OF (CROSS MOMENT) SKEWNESS

The variational characterization provides an interesting perspective on the robust eigenvectors for these latent variable models. Consider the exchangeable single topic models (Theorem 3.1), and the objective function

$$u \mapsto \frac{\mathbb{E}[(x_1^\top u)(x_2^\top u)(x_3^\top u)]}{\mathbb{E}[(x_1^\top u)(x_2^\top u)]^{3/2}} = \frac{M_3(u, u, u)}{M_2(u, u)^{3/2}}.$$

In this case, every local maximizer u^* satisfies $M_2(I, u^*) = \sqrt{w_i} \mu_i$ for some $i \in [k]$. The objective function can be interpreted as the (cross moment) skewness of the random vectors x_1, x_2, x_3 along direction u .

5. Tensor Power Method

In this section, we consider the tensor power method of Lathauwer et al. (2000, Remark 3) for orthogonal tensor decomposition. We first state a simple convergence analysis for an orthogonally decomposable tensor T .

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 for the case of matrices), and a more robust approach is required to extract the approximate decomposition. Here, we propose such a variant in Algorithm 1 and provide a detailed perturbation analysis. We note that alternative approaches such as simultaneous diagonalization can also be employed (see Appendix D).

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 (6)—for extracting a single component of the orthogonal decomposition. Note that the initial vector θ_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 5.1 *Let $T \in \otimes^3 \mathbb{R}^n$ have an orthogonal decomposition as given in (4). For a vector $\theta_0 \in \mathbb{R}^n$, 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 (6) starting from θ_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 $T - \sum_j \lambda_j v_j^{\otimes 3}$ after having obtained robust eigenvector / eigenvalue pairs $\{(v_j, \lambda_j)\}$.

Proof Let $\bar{\theta}_0, \bar{\theta}_1, \bar{\theta}_2, \dots$ be the sequence given by $\bar{\theta}_0 := \theta_0$ and $\bar{\theta}_t := T(I, \theta_{t-1}, \theta_{t-1})$ for $t \geq 1$. Let $c_i := v_i^\top \theta_0$ for all $i \in [k]$. It is easy to check that (i) $\theta_t = \bar{\theta}_t / \|\bar{\theta}_t\|$, and (ii) $\bar{\theta}_t = \sum_{i=1}^k \lambda_i^{2^t-1} c_i^{2^t} v_i$. (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. ■

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. We propose such an algorithm in Algorithm 1, and provide a detailed perturbation analysis. For simplicity, we assume the tensor \hat{T} is of size $k \times k \times k$ as per the reduction from Section 4.3. In some applications, it may be preferable to work directly with a $n \times n \times n$ tensor of rank $k \leq n$ (as in Lemma 5.1); our results apply in that setting with little modification.

Algorithm 1 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

$$\theta_t^{(\tau)} := \frac{\tilde{T}(I, \theta_{t-1}^{(\tau)}, \theta_{t-1}^{(\tau)})}{\|\tilde{T}(I, \theta_{t-1}^{(\tau)}, \theta_{t-1}^{(\tau)})\|} \tag{7}$$

- 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 (7) 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 $\tilde{T} - \hat{\lambda} \hat{\theta}^{\otimes 3}$.
-

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 latent variable model applications, \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. In the context of parameter estimation (as in Section 4.3), E must account for any error amplification throughout the reduction, such as in the whitening step (see, *e.g.*, Hsu and Kakade, 2013, for such an analysis).

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 1 on input \hat{T} in terms of the size of the perturbation, provided that the perturbation is small enough.

Theorem 5.1 *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 $\|E\| \leq \epsilon$. 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*

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

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 1 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 π on $[k]$ such that

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

and

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

The proof of Theorem 5.1 is given in Appendix B.

One important difference from Wedin’s theorem is that this is an algorithm dependent perturbation analysis, specific to Algorithm 1 (since the perturbed tensor need not have an

orthogonal decomposition). Furthermore, note that Algorithm 1 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 .

We note that it is also possible to design a variant of Algorithm 1 that instead uses a stopping criterion to determine if an iterate has (almost) converged to an eigenvector. For instance, if $\tilde{T}(\theta, \theta, \theta) > \max\{\|\tilde{T}\|_F/\sqrt{2r}, \|\tilde{T}(I, I, \theta)\|_F/1.05\}$, where $\|\tilde{T}\|_F$ is the tensor Frobenius norm (vectorized Euclidean norm), and r is the expected rank of the unperturbed tensor ($r = k - \#$ of deflation steps), then it can be shown that θ must be close to one of the eigenvectors, provided that the perturbation is small enough. Using such a stopping criterion can reduce the number of random restarts when a good initial point is found early on. See Appendix C for details.

In general, it is possible, when run on a general symmetric tensor (*e.g.*, \hat{T}), for the tensor power method to exhibit oscillatory behavior (Kofidis and Regalia, 2002, Example 1). This is not in conflict with Theorem 5.1, which effectively bounds the amplitude of these oscillations; in particular, if $\hat{T} = T + E$ is a tensor built from empirical moments, the error term E (and thus the amplitude of the oscillations) can be driven down by drawing more samples. The practical value of addressing these oscillations and perhaps stabilizing the algorithm is an interesting direction for future research (Kolda and Mayo, 2011).

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 3.1), 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.

6. Discussion

In this section, we discuss some practical and application-oriented issues related to the tensor decomposition approach to learning latent variable models.

6.1 Practical Implementation Considerations

A number of practical concerns arise when dealing with moment matrices and tensors. Below, we address two issues that are especially pertinent to topic modeling applications (Anandkumar et al., 2012c,a) or other settings where the observations are sparse.

6.1.1 EFFICIENT MOMENT REPRESENTATION FOR EXCHANGEABLE MODELS

In an exchangeable bag-of-words model, it is assumed that the words x_1, x_2, \dots, x_ℓ in a document are conditionally i.i.d. given the topic h . This allows one to estimate p -th order moments using just p words per document. The estimators obtained via Theorem 3.1 (single topic model) and Theorem 3.5 (LDA) use only up to third-order moments, which suggests that each document only needs to have three words.

In practice, one should use all of the words in a document for efficient estimation of the moments. One way to do this is to average over all $\binom{\ell}{3} \cdot 3!$ ordered triples of words in a document of length ℓ . At first blush, this seems computationally expensive (when ℓ is large), but as it turns out, the averaging can be done implicitly, as shown by Zou et al. (2013). Let $c \in \mathbb{R}^d$ be the word count vector for a document of length ℓ , so c_i is the number of occurrences of word i in the document, and $\sum_{i=1}^d c_i = \ell$. Note that c is a sufficient statistic for the document. Then, the contribution of this document to the empirical third-order moment tensor is given by

$$\frac{1}{\binom{\ell}{3}} \cdot \frac{1}{3!} \cdot \left(c \otimes c \otimes c + 2 \sum_{i=1}^d c_i (e_i \otimes e_i \otimes e_i) - \sum_{i=1}^d \sum_{j=1}^d c_i c_j (e_i \otimes e_i \otimes e_j) - \sum_{i=1}^d \sum_{j=1}^d c_i c_j (e_i \otimes e_j \otimes e_i) - \sum_{i=1}^d \sum_{j=1}^d c_i c_j (e_i \otimes e_j \otimes e_j) \right). \quad (8)$$

It can be checked that this quantity is equal to

$$\frac{1}{\binom{\ell}{3}} \cdot \frac{1}{3!} \cdot \sum_{\text{ordered word triple } (x, y, z)} e_x \otimes e_y \otimes e_z$$

where the sum is over all ordered word triples in the document. A similar expression is easily derived for the contribution of the document to the empirical second-order moment matrix:

$$\frac{1}{\binom{\ell}{2}} \cdot \frac{1}{2!} \cdot \left(c \otimes c - \text{diag}(c) \right). \quad (9)$$

Note that the word count vector c is generally a sparse vector, so this representation allows for efficient multiplication by the moment matrices and tensors in time linear in the size of the document corpus (*i.e.*, the number of non-zero entries in the term-document matrix).

6.1.2 DIMENSIONALITY REDUCTION

Another serious concern regarding the use of tensor forms of moments is the need to operate on multidimensional arrays with $\Omega(d^3)$ values (it is typically not exactly d^3 due to symmetry). When d is large (*e.g.*, when it is the size of the vocabulary in natural language applications), even storing a third-order tensor in memory can be prohibitive. Sparsity is one factor that alleviates this problem. Another approach is to use efficient linear dimensionality reduction. When this is combined with efficient techniques for matrix and tensor multiplication that avoid explicitly constructing the moment matrices and tensors (such as the procedure described above), it is possible to avoid any computational scaling more than linear in the dimension d and the training sample size.

Consider for concreteness the tensor decomposition approach for the exchangeable single topic model as discussed in Section 4.3. Using recent techniques for randomized linear algebra computations (*e.g.*, Halko et al., 2011), it is possible to efficiently approximate the whitening matrix $W \in \mathbb{R}^{d \times k}$ from the second-moment matrix $M_2 \in \mathbb{R}^{d \times d}$. To do this, one first multiplies M_2 by a random matrix $R \in \mathbb{R}^{d \times k'}$ for some $k' \geq k$, and then computes the top k singular vectors of the product $M_2 R$. This provides a basis $U \in \mathbb{R}^{d \times k}$ whose span

is approximately the range of M_2 . From here, an approximate SVD of $U^\top M_2 U$ is used to compute the approximate whitening matrix W . Note that both matrix products $M_2 R$ and $U^\top M_2 U$ may be performed via implicit access to M_2 by exploiting (9), so that M_2 need not be explicitly formed. With the whitening matrix W in hand, the third-moment tensor $\widetilde{M}_3 = M_3(W, W, W) \in \mathbb{R}^{k \times k \times k}$ can be implicitly computed via (8). For instance, the core computation in the tensor power method $\theta' := \widetilde{M}_3(I, \theta, \theta)$ is performed by (i) computing $\eta := W\theta$, (ii) computing $\eta' := M_3(I, \eta, \eta)$, and finally (iii) computing $\theta' := W^\top \eta'$. Using the fact that M_3 is an empirical third-order moment tensor, these steps can be computed with $O(dk + N)$ operations, where N is the number of non-zero entries in the term-document matrix (Zou et al., 2013).

6.2 Computational Complexity

It is interesting to consider the computational complexity of the tensor power method in the dense setting where $T \in \mathbb{R}^{k \times k \times k}$ is orthogonally decomposable but otherwise unstructured. Each iteration requires $O(k^3)$ operations, and assuming at most $k^{1+\delta}$ random restarts for extracting each eigenvector (for some small $\delta > 0$) and $O(\log(k) + \log \log(1/\epsilon))$ iterations per restart, the total running time is $O(k^{5+\delta}(\log(k) + \log \log(1/\epsilon)))$ to extract all k eigenvectors and eigenvalues.

An alternative approach to extracting the orthogonal decomposition of T is to reorganize T into a matrix $M \in \mathbb{R}^{k \times k^2}$ by flattening two of the dimensions into one. In this case, if $T = \sum_{i=1}^k \lambda_i v_i^{\otimes 3}$, then $M = \sum_{i=1}^k \lambda_i v_i \otimes \text{vec}(v_i \otimes v_i)$. This reveals the singular value decomposition of M (assuming the eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_k$ are distinct), and therefore can be computed with $O(k^4)$ operations. Therefore it seems that the tensor power method is less efficient than a pure matrix-based approach via singular value decomposition. However, it should be noted that this matrix-based approach fails to recover the decomposition when eigenvalues are repeated, and can be unstable when the gap between eigenvalues is small—see Appendix D for more discussion.

It is worth noting that the running times differ by roughly a factor of $\Theta(k^{1+\delta})$, which can be accounted for by the random restarts. This gap can potentially be alleviated or removed by using a more clever method for initialization. Moreover, using special structure in the problem (as discussed above) can also improve the running time of the tensor power method.

6.3 Sample Complexity Bounds

Previous work on using linear algebraic methods for estimating latent variable models crucially rely on matrix perturbation analysis for deriving sample complexity bounds (Mossel and Roch, 2006; Hsu et al., 2012b; Anandkumar et al., 2012c,a; Hsu and Kakade, 2013). The learning algorithms in these works are plug-in estimators that use empirical moments in place of the population moments, and then follow algebraic manipulations that result in the desired parameter estimates. As long as these manipulations can tolerate small perturbations of the population moments, a sample complexity bound can be obtained by exploiting the convergence of the empirical moments to the population moments via the law of large numbers. As discussed in Appendix D, these approaches do not directly lead to practical

algorithms due to a certain amplification of the error (a polynomial factor of k , which is observed in practice).

Using the perturbation analysis for the tensor power method, improved sample complexity bounds can be obtained for all of the examples discussed in Section 3. The underlying analysis remains the same as in previous works (*e.g.*, Anandkumar et al., 2012a; Hsu and Kakade, 2013), the main difference being the accuracy of the orthogonal tensor decomposition obtained via the tensor power method. Relative to the previously cited works, the sample complexity bound will be considerably improved in its dependence on the rank parameter k , as Theorem 5.1 implies that the tensor estimation error (*e.g.*, error in estimating \widetilde{M}_3 from Section 4.3) is not amplified by any factor explicitly depending on k (there is a requirement that the error be smaller than some factor depending on k , but this only contributes to a lower-order term in the sample complexity bound). See Appendix D for further discussion regarding the stability of the techniques from these previous works.

6.4 Other Perspectives

The tensor power method is simply one approach for extracting the orthogonal decomposition needed in parameter estimation. The characterizations from Section 4.2 suggest that a number of fixed point and variational techniques may be possible (and Appendix D provides yet another perspective based on simultaneous diagonalization). One important consideration is that the model is often misspecified, and therefore approaches with more robust guarantees (*e.g.*, for convergence) are desirable. Our own experience with the tensor power method (as applied to exchangeable topic modeling) is that while model misspecification does indeed affect convergence, the results can be very reasonable even after just a dozen or so iterations (Anandkumar et al., 2012a). Nevertheless, robustness is likely more important in other applications, and thus the stabilization approaches (Kofidis and Regalia, 2002; Regalia and Kofidis, 2003; Erdogan, 2009; Kolda and Mayo, 2011) may be advantageous.

Acknowledgments

We thank Boaz Barak, Dean Foster, Jon Kelner, and Greg Valiant for helpful discussions. We are also grateful to Hanzhang Hu, Drew Bagnell, and Martial Hebert for alerting us of an issue with Theorem 4.2 and suggesting a simple fix. This work was completed while DH was a postdoctoral researcher at Microsoft Research New England, and partly while AA, RG, and MT were visiting the same lab. AA is supported in part by the NSF Award CCF-1219234, AFOSR Award FA9550-10-1-0310 and the ARO Award W911NF-12-1-0404.

Appendix A. Fixed-Point and Variational Characterizations of Orthogonal Tensor Decompositions

We give detailed proofs of Theorems 4.1 and 4.2 in this section for completeness.

A.1 Proof of Theorem 4.1

Theorem A.1 *Let T have an orthogonal decomposition as given in (4).*

1. The set of $\theta \in \mathbb{R}^n$ which do not converge to some v_i under repeated iteration of (6) has measure zero.
2. The set of robust eigenvectors of T is $\{v_1, v_2, \dots, v_k\}$.

Proof For a random choice of $\theta \in \mathbb{R}^n$ (under any distribution absolutely continuous with respect to Lebesgue measure), the values $|\lambda_1 v_1^\top \theta|, |\lambda_2 v_2^\top \theta|, \dots, |\lambda_k v_k^\top \theta|$ will be distinct with probability 1. Therefore, there exists a unique largest value, say $|\lambda_i v_i^\top \theta|$ for some $i \in [k]$, and by Lemma 5.1, we have convergence to v_i under repeated iteration of (6). Thus the first claim holds.

We now prove the second claim. First, we show that every v_i is a robust eigenvector. Pick any $i \in [k]$, and note that for a sufficiently small ball around v_i , we have that for all θ in this ball, $\lambda_i v_i^\top \theta$ is strictly greater than $\lambda_j v_j^\top \theta$ for $j \in [k] \setminus \{i\}$. Thus by Lemma 5.1, v_i is a robust eigenvector. Now we show that the v_i are the only robust eigenvectors. Suppose there exists some robust eigenvector u not equal to v_i for any $i \in [k]$. Then there exists a positive measure set around u such that all points in this set converge to u under repeated iteration of (6). This contradicts the first claim. ■

A.2 Proof of Theorem 4.2

Theorem A.2 *Let T have an orthogonal decomposition as given in (4), and consider the optimization problem*

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

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 $i \in [k]$.

Proof Consider the Lagrangian form of the corresponding constrained maximization problem over unit vectors $u \in \mathbb{R}^n$:

$$\mathcal{L}(u, \lambda) := T(u, u, u) - \frac{3}{2}\lambda(u^\top u - 1).$$

Since

$$\nabla_u \mathcal{L}(u, \lambda) = \nabla_u \left(\sum_{i=1}^k \lambda_i (v_i^\top u)^3 - \frac{3}{2}\lambda(u^\top u - 1) \right) = 3 \left(T(I, u, u) - \lambda u \right),$$

the stationary points $u \in \mathbb{R}^n$ (with $\|u\| \leq 1$) satisfy

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

for some $\lambda \in \mathbb{R}$, *i.e.*, (u, λ) is an eigenvector/eigenvalue pair of T .

Now we characterize the isolated local maximizers. Observe that if $u \neq 0$ and $T(I, u, u) = \lambda u$ for $\lambda < 0$, then $T(u, u, u) < 0$. Therefore $u' = (1 - \delta)u$ for any $\delta \in (0, 1)$ satisfies $T(u', u', u') = (1 - \delta)^3 T(u, u, u) > T(u, u, u)$. So such a u cannot be a local maximizer.

Moreover, if $\|u\| < 1$ and $T(I, u, u) = \lambda u$ for $\lambda > 0$, then $u' = (1 + \delta)u$ for a small enough $\delta \in (0, 1)$ satisfies $\|u'\| \leq 1$ and $T(u', u', u') = (1 + \delta)^3 T(u, u, u) > T(u, u, u)$. Therefore a local maximizer must have $T(I, u, u) = \lambda u$ for some $\lambda \geq 0$, and $\|u\| = 1$ whenever $\lambda > 0$.

Extend $\{v_1, v_2, \dots, v_k\}$ to an orthonormal basis $\{v_1, v_2, \dots, v_n\}$ of \mathbb{R}^n . Now pick any stationary point $u = \sum_{i=1}^n c_i v_i$ with $\lambda := T(u, u, u) = u^\top T(I, u, u)$. Then

$$\lambda_i c_i^2 = \lambda_i (u^\top v_i)^2 = v_i^\top T(I, u, u) = \lambda v_i^\top u = \lambda c_i \geq 0, \quad i \in [k],$$

and thus

$$\nabla_u^2 \mathcal{L}(u, \lambda) = 6 \sum_{i=1}^k \lambda_i c_i v_i v_i^\top - 3\lambda I = 3\lambda \left(2 \sum_{i \in \Omega} v_i v_i^\top - I \right)$$

where $\Omega := \{i \in [k] : c_i \neq 0\}$. This implies that for any unit vector $w \in \mathbb{R}^n$,

$$w^\top \nabla_u^2 \mathcal{L}(u, \lambda) w = 3\lambda \left(2 \sum_{i \in \Omega} (v_i^\top w)^2 - 1 \right).$$

The point u is an isolated local maximum if the above quantity is strictly negative for all unit vectors w orthogonal to u . We now consider three cases depending on the cardinality of Ω and the sign of λ .

- Case 1: $|\Omega| = 1$ and $\lambda > 0$. This means $u = v_i$ for some $i \in [k]$ (as $u = -v_i$ implies $\lambda = -\lambda_i < 0$). In this case,

$$w^\top \nabla_u^2 \mathcal{L}(u, \lambda) w = 3\lambda_i (2(v_i^\top w)^2 - 1) = -3\lambda_i < 0$$

for all $w \in \mathbb{R}^n$ satisfying $(u^\top w)^2 = (v_i^\top w)^2 = 0$. Hence u is an isolated local maximizer.

- Case 2: $|\Omega| \geq 2$ and $\lambda > 0$. Since $|\Omega| \geq 2$, we may pick a strict non-empty subset $S \subsetneq \Omega$ and set

$$w := \frac{1}{Z} \left(\frac{1}{Z_S} \sum_{i \in S} c_i v_i - \frac{1}{Z_{S^c}} \sum_{i \in \Omega \setminus S} c_i v_i \right)$$

where $Z_S := \sum_{i \in S} c_i^2$, $Z_{S^c} := \sum_{i \in \Omega \setminus S} c_i^2$, and $Z := \sqrt{1/Z_S + 1/Z_{S^c}}$. It is easy to check that $\|w\|^2 = \sum_{i \in \Omega} (v_i^\top w)^2 = 1$ and $u^\top w = 0$. Consider any open neighborhood U of u , and pick $\delta > 0$ small enough so that $\tilde{u} := \sqrt{1 - \delta^2} u + \delta w$ is contained in U . Set $u_0 := \sqrt{1 - \delta^2} u$. By Taylor's theorem, there exists $\epsilon \in [0, \delta]$ such that, for

$\tilde{u} := u_0 + \epsilon w$, we have

$$\begin{aligned}
 T(\tilde{u}, \tilde{u}, \tilde{u}) &= T(u_0, u_0, u_0) + \nabla_u T(u, u, u)^\top (\tilde{u} - u_0) \Big|_{u=u_0} \\
 &\quad + \frac{1}{2} (\tilde{u} - u_0)^\top \nabla_u^2 T(u, u, u) (\tilde{u} - u_0) \Big|_{u=\tilde{u}} \\
 &= (1 - \delta^2)^{3/2} \lambda + \delta(1 - \delta^2) \lambda u^\top w + \frac{1}{2} \delta^2 w^\top \nabla_u^2 T(u, u, u) w \Big|_{u=\tilde{u}} \\
 &= (1 - \delta^2)^{3/2} \lambda + 0 + 3\delta^2 \sum_{i=1}^k \lambda_i (v_i^\top (u_0 + \epsilon w)) (v_i^\top w)^2 \\
 &= (1 - \delta^2)^{3/2} \lambda + 3\delta^2 \sqrt{1 - \delta^2} \sum_{i=1}^k \lambda_i c_i (v_i^\top w)^2 + 3\delta^2 \epsilon \sum_{i=1}^k \lambda_i (v_i^\top w)^3 \\
 &= (1 - \delta^2)^{3/2} \lambda + 3\delta^2 \sqrt{1 - \delta^2} \lambda \sum_{i \in \Omega} (v_i^\top w)^2 + 3\delta^2 \epsilon \sum_{i=1}^k \lambda_i (v_i^\top w)^3 \\
 &= (1 - \delta^2)^{3/2} \lambda + 3\delta^2 \sqrt{1 - \delta^2} \lambda + 3\delta^2 \epsilon \sum_{i=1}^k \lambda_i (v_i^\top w)^3 \\
 &= \left(1 - \frac{3}{2} \delta^2 + O(\delta^4)\right) \lambda + 3\delta^2 \sqrt{1 - \delta^2} \lambda + 3\delta^2 \epsilon \sum_{i=1}^k \lambda_i (v_i^\top w)^3.
 \end{aligned}$$

Since $\epsilon \leq \delta$, for small enough δ , the RHS is strictly greater than λ . This implies that u is not an isolated local maximizer.

- Case 3: $|\Omega| = 0$ or $\lambda = 0$. Note that if $|\Omega| = 0$, then $\lambda = 0$, so we just consider $\lambda = 0$. Consider any open neighborhood U of u , and pick $j \in [n]$ and $\delta > 0$ small enough so that $\tilde{u} := \sqrt{1 - \delta^2} u + \delta v_j$ is contained in U . Then

$$T(\tilde{u}, \tilde{u}, \tilde{u}) = (1 - \delta^2)^{3/2} T(u, u, u) + 3\lambda_j (1 - \delta^2) \delta c_j^2 + 3\lambda_j \sqrt{1 - \delta^2} \delta^2 c_j + \delta^3 > 0 = \lambda$$

for sufficiently small δ . Thus u is not an isolated local maximizer.

From these exhaustive cases, we conclude that a stationary point u is an isolated local maximizer if and only if $u = v_i$ for some $i \in [k]$. ■

We are grateful to Hanzhang Hu, Drew Bagnell, and Martial Hebert for alerting us of an issue with our original statement of Theorem 4.2 and its proof, and for suggesting a simple fix. The original statement used the optimization constraint $\|u\| = 1$ (rather than $\|u\| \leq 1$), but the characterization of the decomposition with this constraint is then only given by isolated local maximizers u with the additional constraint that $T(u, u, u) > 0$ —that is, there can be isolated local maximizers with $T(u, u, u) \leq 0$ that are not vectors in the decomposition. The suggested fix of Hu, Bagnell, and Herbert is to relax to $\|u\| \leq 1$, which eliminates isolated local maximizers with $T(u, u, u) \leq 0$; this way, the characterization of the decomposition is simply the isolated local maximizers under the relaxed constraint.

Appendix B. Analysis of Robust Power Method

In this section, we prove Theorem 5.1. The proof is structured as follows. In Appendix B.1, we show that with high probability, at least one out of L random vectors will be a good initializer for the tensor power iterations. An initializer is good if its projection onto an eigenvector is noticeably larger than its projection onto other eigenvectors. We then analyze in Appendix B.2 the convergence behavior of the tensor power iterations. Relative to the proof of Lemma 5.1, this analysis is complicated by the tensor perturbation. We show that there is an initial slow convergence phase (linear rate rather than quadratic), but as soon as the projection of the iterate onto an eigenvector is large enough, it enters the quadratic convergence regime until the perturbation dominates. Finally, we show how errors accrue due to deflation in Appendix B.3, which is rather subtle and different from deflation with matrix eigendecompositions. This is because when some initial set of eigenvectors and eigenvalues are accurately recovered, the additional errors due to deflation are effectively only lower-order terms. These three pieces are assembled in Appendix B.4 to complete the proof of Theorem 5.1.

B.1 Initialization

Consider a set of non-negative numbers $\tilde{\lambda}_1, \tilde{\lambda}_2, \dots, \tilde{\lambda}_k \geq 0$. For $\gamma \in (0, 1)$, we say a unit vector $\theta_0 \in \mathbb{R}^k$ is γ -separated relative to $i^* \in [k]$ if

$$\tilde{\lambda}_{i^*} |\theta_{i^*,0}| - \max_{i \in [k] \setminus \{i^*\}} \tilde{\lambda}_i |\theta_{i,0}| \geq \gamma \tilde{\lambda}_{i^*} |\theta_{i^*,0}|$$

(the dependence on $\tilde{\lambda}_1, \tilde{\lambda}_2, \dots, \tilde{\lambda}_k$ is implicit).

The following lemma shows that for any constant γ , with probability at least $1 - \eta$, at least one out of $\text{poly}(k) \log(1/\eta)$ i.i.d. random vectors (uniformly distributed over the unit sphere S^{k-1}) is γ -separated relative to $\arg \max_{i \in [k]} \tilde{\lambda}_i$. (For small enough γ and large enough k , the polynomial is close to linear in k .)

Lemma B.1 *There exists an absolute constant $c > 0$ such that if positive integer $L \geq 2$ satisfies*

$$\sqrt{\frac{\ln(L)}{\ln(k)}} \cdot \left(1 - \frac{\ln(\ln(L)) + c}{4 \ln(L)} - \sqrt{\frac{\ln(8)}{\ln(L)}} \right) \geq \frac{1}{1 - \gamma} \cdot \left(1 + \sqrt{\frac{\ln(4)}{\ln(k)}} \right), \quad (10)$$

the following holds. With probability at least $1/2$ over the choice of L i.i.d. random vectors drawn uniformly distributed over the unit sphere S^{k-1} in \mathbb{R}^k , at least one of the vectors is γ -separated relative to $\arg \max_{i \in [k]} \tilde{\lambda}_i$. Moreover, with the same c, L , and for any $\eta \in (0, 1)$, with probability at least $1 - \eta$ over $L \cdot \log_2(1/\eta)$ i.i.d. uniform random unit vectors, at least one of the vectors is γ -separated.

Proof Without loss of generality, assume $\arg \max_{i \in [k]} \tilde{\lambda}_i = 1$. Consider a random matrix $Z \in \mathbb{R}^{k \times L}$ whose entries are independent $\mathcal{N}(0, 1)$ random variables; we take the j -th column of Z to be comprised of the random variables used for the j -th random vector (before normalization). Specifically, for the j -th random vector,

$$\theta_{i,0} := \frac{Z_{i,j}}{\sqrt{\sum_{i'=1}^k Z_{i',j}^2}}, \quad i \in [n].$$

It suffices to show that with probability at least $1/2$, there is a column $j^* \in [L]$ such that

$$|Z_{1,j^*}| \geq \frac{1}{1-\gamma} \max_{i \in [k] \setminus \{1\}} |Z_{i,j^*}|.$$

Since $\max_{j \in [L]} |Z_{1,j}|$ is a 1-Lipschitz function of L independent $\mathcal{N}(0, 1)$ random variables, it follows that

$$\Pr \left[\left| \max_{j \in [L]} |Z_{1,j}| - \text{median} \left[\max_{j \in [L]} |Z_{1,j}| \right] \right| > \sqrt{2 \ln(8)} \right] \leq 1/4.$$

Moreover,

$$\text{median} \left[\max_{j \in [L]} |Z_{1,j}| \right] \geq \text{median} \left[\max_{j \in [L]} Z_{1,j} \right] =: m.$$

Observe that the cumulative distribution function of $\max_{j \in [L]} Z_{1,j}$ is given by $F(z) = \Phi(z)^L$, where Φ is the standard Gaussian CDF. Since $F(m) = 1/2$, it follows that $m = \Phi^{-1}(2^{-1/L})$. It can be checked that

$$\Phi^{-1}(2^{-1/L}) \geq \sqrt{2 \ln(L)} - \frac{\ln(\ln(L)) + c}{2\sqrt{2 \ln(L)}}$$

for some absolute constant $c > 0$. Also, let $j^* := \arg \max_{j \in [L]} |Z_{1,j}|$.

Now for each $j \in [L]$, let $|Z_{2:k,j}| := \max\{|Z_{2,j}|, |Z_{3,j}|, \dots, |Z_{k,j}|\}$. Again, since $|Z_{2:k,j}|$ is a 1-Lipschitz function of $k-1$ independent $\mathcal{N}(0, 1)$ random variables, it follows that

$$\Pr \left[|Z_{2:k,j}| > \mathbb{E} \left[|Z_{2:k,j}| \right] + \sqrt{2 \ln(4)} \right] \leq 1/4.$$

Moreover, by a standard argument,

$$\mathbb{E} \left[|Z_{2:k,j}| \right] \leq \sqrt{2 \ln(k)}.$$

Since $|Z_{2:k,j}|$ is independent of $|Z_{1,j}|$ for all $j \in [L]$, it follows that the previous two displayed inequalities also hold with j replaced by j^* .

Therefore we conclude with a union bound that with probability at least $1/2$,

$$|Z_{1,j^*}| \geq \sqrt{2 \ln(L)} - \frac{\ln(\ln(L)) + c}{2\sqrt{2 \ln(L)}} - \sqrt{2 \ln(8)} \quad \text{and} \quad |Z_{2:k,j^*}| \leq \sqrt{2 \ln(k)} + \sqrt{2 \ln(4)}.$$

Since L satisfies (10) by assumption, in this event, the j^* -th random vector is γ -separated. \blacksquare

B.2 Tensor Power Iterations

Recall the update rule used in the power method. Let $\theta_t = \sum_{i=1}^k \theta_{i,t} v_i \in \mathbb{R}^k$ be the unit vector at time t . Then

$$\theta_{t+1} = \sum_{i=1}^k \theta_{i,t+1} v_i := \tilde{T}(I, \theta_t, \theta_t) / \|\tilde{T}(I, \theta_t, \theta_t)\|.$$

In this subsection, we assume that \tilde{T} has the form

$$\tilde{T} = \sum_{i=1}^k \tilde{\lambda}_i v_i^{\otimes 3} + \tilde{E} \quad (11)$$

where $\{v_1, v_2, \dots, v_k\}$ is an orthonormal basis, and, without loss of generality,

$$\tilde{\lambda}_1 |\theta_{1,t}| = \max_{i \in [k]} \tilde{\lambda}_i |\theta_{i,t}| > 0.$$

Also, define

$$\tilde{\lambda}_{\min} := \min\{\tilde{\lambda}_i : i \in [k], \tilde{\lambda}_i > 0\}, \quad \tilde{\lambda}_{\max} := \max\{\tilde{\lambda}_i : i \in [k]\}.$$

We further assume the error \tilde{E} is a symmetric tensor such that, for some constant $p > 1$,

$$\|\tilde{E}(I, u, u)\| \leq \tilde{\epsilon}, \quad \forall u \in S^{k-1}, \quad (12)$$

$$\|\tilde{E}(I, u, u)\| \leq \tilde{\epsilon}/p, \quad \forall u \in S^{k-1} \text{ s.t. } (u^\top v_1)^2 \geq 1 - (3\tilde{\epsilon}/\tilde{\lambda}_1)^2. \quad (13)$$

In the next two propositions (Propositions B.1 and B.2) and next two lemmas (Lemmas B.2 and B.3), we analyze the power method iterations using \tilde{T} at some arbitrary iterate θ_t using only the property (12) of \tilde{E} . But throughout, the quantity $\tilde{\epsilon}$ can be replaced by $\tilde{\epsilon}/p$ if θ_t satisfies $(\theta_t^\top v_1)^2 \geq 1 - (3\tilde{\epsilon}/\tilde{\lambda}_1)^2$ as per property (13).

Define

$$\begin{aligned} R_\tau &:= \left(\frac{\theta_{1,\tau}^2}{1 - \theta_{1,\tau}^2} \right)^{1/2}, & r_{i,\tau} &:= \frac{\tilde{\lambda}_1 \theta_{1,\tau}}{\tilde{\lambda}_i |\theta_{i,\tau}|}, \\ \gamma_\tau &:= 1 - \frac{1}{\min_{i \neq 1} |r_{i,\tau}|}, & \delta_\tau &:= \frac{\tilde{\epsilon}}{\tilde{\lambda}_1 \theta_{1,\tau}^2}, & \kappa &:= \frac{\tilde{\lambda}_{\max}}{\tilde{\lambda}_1} \end{aligned} \quad (14)$$

for $\tau \in \{t, t+1\}$.

Proposition B.1

$$\min_{i \neq 1} |r_{i,t}| \geq \frac{R_t}{\kappa}, \quad \gamma_t \geq 1 - \frac{\kappa}{R_t}, \quad \theta_{1,t}^2 = \frac{R_t^2}{1 + R_t^2}.$$

Proposition B.2

$$r_{i,t+1} \geq r_{i,t}^2 \cdot \frac{1 - \delta_t}{1 + \kappa \delta_t r_{i,t}^2} = \frac{1 - \delta_t}{\frac{1}{r_{i,t}^2} + \kappa \delta_t}, \quad i \in [k], \quad (15)$$

$$R_{t+1} \geq R_t \cdot \frac{1 - \delta_t}{1 - \gamma_t + \delta_t R_t} \geq \frac{1 - \delta_t}{\frac{\kappa}{R_t^2} + \delta_t}. \quad (16)$$

Proof Let $\check{\theta}_{t+1} := \tilde{T}(I, \theta_t, \theta_t)$, so $\theta_{t+1} = \check{\theta}_{t+1} / \|\check{\theta}_{t+1}\|$. Since $\check{\theta}_{i,t+1} = \tilde{T}(v_i, \theta_t, \theta_t) = T(v_i, \theta_t, \theta_t) + E(v_i, \theta_t, \theta_t)$, we have

$$\check{\theta}_{i,t+1} = \tilde{\lambda}_i \theta_{i,t}^2 + E(v_i, \theta_t, \theta_t), \quad i \in [k].$$

Using the triangle inequality and the fact $\|E(v_i, \theta_t, \theta_t)\| \leq \tilde{\epsilon}$, we have

$$\check{\theta}_{i,t+1} \geq \tilde{\lambda}_i \theta_{i,t}^2 - \tilde{\epsilon} \geq |\theta_{i,t}| \cdot \left(\tilde{\lambda}_i |\theta_{i,t}| - \tilde{\epsilon}/|\theta_{i,t}| \right) \quad (17)$$

and

$$|\check{\theta}_{i,t+1}| \leq |\tilde{\lambda}_i \theta_{i,t}^2| + \tilde{\epsilon} \leq |\theta_{i,t}| \cdot \left(\tilde{\lambda}_i |\theta_{i,t}| + \tilde{\epsilon}/|\theta_{i,t}| \right) \quad (18)$$

for all $i \in [k]$. Combining (17) and (18) gives

$$r_{i,t+1} = \frac{\tilde{\lambda}_1 \theta_{1,t+1}}{\tilde{\lambda}_i |\theta_{i,t+1}|} = \frac{\tilde{\lambda}_1 \check{\theta}_{1,t+1}}{\tilde{\lambda}_i |\check{\theta}_{i,t+1}|} \geq r_{i,t}^2 \cdot \frac{1 - \delta_t}{1 + \frac{\tilde{\epsilon}}{\tilde{\lambda}_i \theta_{i,t}^2}} = r_{i,t}^2 \cdot \frac{1 - \delta_t}{1 + (\tilde{\lambda}_i/\tilde{\lambda}_1) \delta_t r_{i,t}^2} \geq r_{i,t}^2 \cdot \frac{1 - \delta_t}{1 + \kappa \delta_t r_{i,t}^2}.$$

Moreover, by the triangle inequality and Hölder's inequality,

$$\begin{aligned} \left(\sum_{i=2}^n [\check{\theta}_{i,t+1}]^2 \right)^{1/2} &= \left(\sum_{i=2}^n \left(\tilde{\lambda}_i \theta_{i,t}^2 + E(v_i, \theta_t, \theta_t) \right)^2 \right)^{1/2} \\ &\leq \left(\sum_{i=2}^n \tilde{\lambda}_i^2 \theta_{i,t}^4 \right)^{1/2} + \left(\sum_{i=2}^n E(v_i, \theta_t, \theta_t)^2 \right)^{1/2} \\ &\leq \max_{i \neq 1} \tilde{\lambda}_i |\theta_{i,t}| \left(\sum_{i=2}^n \theta_{i,t}^2 \right)^{1/2} + \tilde{\epsilon} \\ &= (1 - \theta_{1,t}^2)^{1/2} \cdot \left(\max_{i \neq 1} \tilde{\lambda}_i |\theta_{i,t}| + \tilde{\epsilon}/(1 - \theta_{1,t}^2)^{1/2} \right). \end{aligned} \quad (19)$$

Combining (17) and (19) gives

$$\frac{|\theta_{1,t+1}|}{(1 - \theta_{1,t+1}^2)^{1/2}} = \frac{|\check{\theta}_{1,t+1}|}{\left(\sum_{i=2}^n [\check{\theta}_{i,t+1}]^2 \right)^{1/2}} \geq \frac{|\theta_{1,t}|}{(1 - \theta_{1,t}^2)^{1/2}} \cdot \frac{\tilde{\lambda}_1 |\theta_{1,t}| - \tilde{\epsilon}/|\theta_{1,t}|}{\max_{i \neq 1} \tilde{\lambda}_i |\theta_{i,t}| + \tilde{\epsilon}/(1 - \theta_{1,t}^2)^{1/2}}.$$

In terms of R_{t+1} , R_t , γ_t , and δ_t , this reads

$$R_{t+1} \geq \frac{1 - \delta_t}{(1 - \gamma_t) \left(\frac{1 - \theta_{1,t}^2}{\theta_{1,t}^2} \right)^{1/2} + \delta_t} = R_t \cdot \frac{1 - \delta_t}{1 - \gamma_t + \delta_t R_t} = \frac{1 - \delta_t}{\frac{1 - \gamma_t}{R_t} + \delta_t} \geq \frac{1 - \delta_t}{\frac{\kappa}{R_t^2} + \delta_t}$$

where the last inequality follows from Proposition B.1. ■

Lemma B.2 Fix any $\rho > 1$. Assume

$$0 \leq \delta_t < \min \left\{ \frac{1}{2(1 + 2\kappa\rho^2)}, \frac{1 - 1/\rho}{1 + \kappa\rho} \right\}$$

and $\gamma_t > 2(1 + 2\kappa\rho^2)\delta_t$.

1. If $r_{i,t}^2 \leq 2\rho^2$, then $r_{i,t+1} \geq |r_{i,t}| \left(1 + \frac{\gamma_t}{2} \right)$.

2. If $\rho^2 < r_{i,t}^2$, then $r_{i,t+1} \geq \min\{r_{i,t}^2/\rho, \frac{1-\delta_t-1/\rho}{\kappa\delta_t}\}$.
3. $\gamma_{t+1} \geq \min\{\gamma_t, 1 - 1/\rho\}$.
4. If $\min_{i \neq 1} r_{i,t}^2 > (\rho(1 - \delta_t) - 1)/(\kappa\delta_t)$, then $R_{t+1} > \frac{1-\delta_t-1/\rho}{\kappa\delta_t} \cdot \frac{\tilde{\lambda}_{\min}}{\tilde{\lambda}_1} \cdot \frac{1}{\sqrt{k}}$.
5. If $R_t \leq 1 + 2\kappa\rho^2$, then $R_{t+1} \geq R_t(1 + \frac{\gamma_t}{3})$, $\theta_{1,t+1}^2 \geq \theta_{1,t}^2$, and $\delta_{t+1} \leq \delta_t$.

Proof Consider two (overlapping) cases depending on $r_{i,t}^2$.

- Case 1: $r_{i,t}^2 \leq 2\rho^2$. By (15) from Proposition B.2,

$$r_{i,t+1} \geq r_{i,t}^2 \cdot \frac{1 - \delta_t}{1 + \kappa\delta_t r_{i,t}^2} \geq |r_{i,t}| \cdot \frac{1}{1 - \gamma_t} \cdot \frac{1 - \delta_t}{1 + 2\kappa\rho^2\delta_t} \geq |r_{i,t}| \left(1 + \frac{\gamma_t}{2}\right)$$

where the last inequality uses the assumption $\gamma_t > 2(1 + 2\kappa\rho^2)\delta_t$. This proves the first claim.

- Case 2: $\rho^2 < r_{i,t}^2$. We split into two sub-cases. Suppose $r_{i,t}^2 \leq (\rho(1 - \delta_t) - 1)/(\kappa\delta_t)$. Then, by (15),

$$r_{i,t+1} \geq r_{i,t}^2 \cdot \frac{1 - \delta_t}{1 + \kappa\delta_t r_{i,t}^2} \geq r_{i,t}^2 \cdot \frac{1 - \delta_t}{1 + \kappa\delta_t \frac{\rho(1-\delta_t)-1}{\kappa\delta_t}} = \frac{r_{i,t}^2}{\rho}.$$

Now suppose instead $r_{i,t}^2 > (\rho(1 - \delta_t) - 1)/(\kappa\delta_t)$. Then

$$r_{i,t+1} \geq \frac{1 - \delta_t}{\frac{\kappa\delta_t}{\rho(1-\delta_t)-1} + \kappa\delta_t} = \frac{1 - \delta_t - 1/\rho}{\kappa\delta_t}. \quad (20)$$

Observe that if $\min_{i \neq 1} r_{i,t}^2 \leq (\rho(1 - \delta_t) - 1)/(\kappa\delta_t)$, then $r_{i,t+1} \geq |r_{i,t}|$ for all $i \in [k]$, and hence $\gamma_{t+1} \geq \gamma_t$. Otherwise we have $\gamma_{t+1} > 1 - \frac{\kappa\delta_t}{1-\delta_t-1/\rho} > 1 - 1/\rho$. This proves the third claim.

If $\min_{i \neq 1} r_{i,t}^2 > (\rho(1 - \delta_t) - 1)/(\kappa\delta_t)$, then we may apply the inequality (20) from the second sub-case of Case 2 above to get

$$R_{t+1} = \frac{1}{\left(\sum_{i \neq 1} (\tilde{\lambda}_1/\tilde{\lambda}_i)^2 / r_{i,t+1}^2\right)^{1/2}} > \left(\frac{1 - \delta_t - 1/\rho}{\kappa\delta_t}\right) \cdot \frac{\tilde{\lambda}_{\min}}{\tilde{\lambda}_1} \cdot \frac{1}{\sqrt{k}}.$$

This proves the fourth claim.

Finally, for the last claim, if $R_t \leq 1 + 2\kappa\rho^2$, then by (16) from Proposition B.2 and the assumption $\gamma_t > 2(1 + 2\kappa\rho^2)\delta_t$,

$$R_{t+1} \geq R_t \cdot \frac{1 - \delta_t}{1 - \gamma_t + \delta_t R_t} \geq R_t \cdot \frac{1 - \frac{\gamma_t}{2(1+2\kappa\rho^2)}}{1 - \gamma_t/2} \geq R_t \left(1 + \gamma_t \cdot \frac{\kappa\rho^2}{1 + 2\kappa\rho^2}\right) \geq R_t \left(1 + \frac{\gamma_t}{3}\right).$$

This in turn implies that $\theta_{1,t+1}^2 \geq \theta_{1,t}^2$ via Proposition B.1, and thus $\delta_{t+1} \leq \delta_t$. \blacksquare

Lemma B.3 Assume $0 \leq \delta_t < 1/2$ and $\gamma_t > 0$. Pick any $\beta > \alpha > 0$ such that

$$\frac{\alpha}{(1+\alpha)(1+\alpha^2)} \geq \frac{\tilde{\epsilon}}{\gamma_t \tilde{\lambda}_1}, \quad \frac{\alpha}{2(1+\alpha)(1+\beta^2)} \geq \frac{\tilde{\epsilon}}{\tilde{\lambda}_1}.$$

1. If $R_t \geq 1/\alpha$, then $R_{t+1} \geq 1/\alpha$.
2. If $1/\alpha > R_t \geq 1/\beta$, then $R_{t+1} \geq \min\{R_t^2/(2\kappa), 1/\alpha\}$.

Proof Observe that for any $c > 0$,

$$R_t \geq \frac{1}{c} \Leftrightarrow \theta_{1,t}^2 \geq \frac{1}{1+c^2} \Leftrightarrow \delta_t \leq \frac{(1+c^2)\tilde{\epsilon}}{\tilde{\lambda}_1}. \tag{21}$$

Now consider the following cases depending on R_t .

- Case 1: $R_t \geq 1/\alpha$. In this case, we have

$$\delta_t \leq \frac{(1+\alpha^2)\tilde{\epsilon}}{\tilde{\lambda}_1} \leq \frac{\alpha\gamma_t}{1+\alpha}$$

by (21) (with $c = \alpha$) and the condition on α . Combining this with (16) from Proposition B.2 gives

$$R_{t+1} \geq \frac{1-\delta_t}{\frac{1-\gamma_t}{R_t} + \delta_t} \geq \frac{1-\frac{\alpha\gamma_t}{1+\alpha}}{(1-\gamma_t)\alpha + \frac{\alpha\gamma_t}{1+\alpha}} = \frac{1}{\alpha}.$$

- Case 2: $1/\beta \leq R_t < 1/\alpha$. In this case, we have

$$\delta_t \leq \frac{(1+\beta^2)\tilde{\epsilon}}{\tilde{\lambda}_1} \leq \frac{\alpha}{2(1+\alpha)}$$

by (21) (with $c = \beta$) and the conditions on α and β . If $\delta_t \geq 1/(2 + R_t^2/\kappa)$, then (16) implies

$$R_{t+1} \geq \frac{1-\delta_t}{\frac{\kappa}{R_t^2} + \delta_t} \geq \frac{1-2\delta_t}{2\delta_t} \geq \frac{1-\frac{\alpha}{1+\alpha}}{\frac{\alpha}{1+\alpha}} = \frac{1}{\alpha}.$$

If instead $\delta_t < 1/(2 + R_t^2/\kappa)$, then (16) implies

$$R_{t+1} \geq \frac{1-\delta_t}{\frac{\kappa}{R_t^2} + \delta_t} > \frac{1-\frac{1}{2+R_t^2/\kappa}}{\frac{\kappa}{R_t^2} + \frac{1}{2+R_t^2/\kappa}} = \frac{R_t^2}{2\kappa}.$$

■

B.2.1 APPROXIMATE RECOVERY OF A SINGLE EIGENVECTOR

We now state the main result regarding the approximate recovery of a single eigenvector using the tensor power method on \tilde{T} . Here, we exploit the special properties of the error \tilde{E} —both (12) and (13).

Lemma B.4 *There exists a universal constant $C > 0$ such that the following holds. Let $i^* := \arg \max_{i \in [k]} \tilde{\lambda}_i |\theta_{i,0}|$. If*

$$\tilde{\epsilon} < \frac{\gamma_0}{2(1+8\kappa)} \cdot \tilde{\lambda}_{\min} \cdot \theta_{i^*,0}^2 \quad \text{and} \quad N \geq C \cdot \left(\frac{\log(k\kappa)}{\gamma_0} + \log \log \frac{p\tilde{\lambda}_{i^*}}{\tilde{\epsilon}} \right),$$

then after $t \geq N$ iterations of the tensor power method on tensor \tilde{T} as defined in (11) and satisfying (12) and (13), the final vector θ_t satisfies

$$\theta_{i^*,t} \geq \sqrt{1 - \left(\frac{3\tilde{\epsilon}}{p\tilde{\lambda}_{i^*}} \right)^2}, \quad \|\theta_t - v_{i^*}\| \leq \frac{4\tilde{\epsilon}}{p\tilde{\lambda}_{i^*}}, \quad |\tilde{T}(\theta_t, \theta_t, \theta_t) - \tilde{\lambda}_{i^*}| \leq \left(27\kappa \left(\frac{\tilde{\epsilon}}{p\tilde{\lambda}_{i^*}} \right)^2 + 2 \right) \frac{\tilde{\epsilon}}{p}.$$

Proof Assume without loss of generality that $i^* = 1$. We consider three phases: (i) iterations before the first time t such that $R_t > 1 + 2\kappa\rho^2 = 1 + 8\kappa$ (using $\rho := 2$), (ii) the subsequent iterations before the first time t such that $R_t \geq 1/\alpha$ (where α will be defined below), and finally (iii) the remaining iterations.

We begin by analyzing the first phase, *i.e.*, the iterates in $T_1 := \{t \geq 0 : R_t \leq 1 + 2\kappa\rho^2 = 1 + 8\kappa\}$. Observe that the condition on $\tilde{\epsilon}$ implies

$$\delta_0 = \frac{\tilde{\epsilon}}{\tilde{\lambda}_1 \theta_{1,0}^2} < \frac{\gamma_0}{2(1+8\kappa)} \cdot \frac{\tilde{\lambda}_{\min}}{\tilde{\lambda}_1} \leq \min \left\{ \frac{\gamma_0}{2(1+2\kappa\rho^2)}, \frac{1-1/\rho}{2(1+2\kappa\rho^2)} \right\},$$

and hence the preconditions on δ_t and γ_t of Lemma B.2 hold for $t = 0$. For all $t \in T_1$ satisfying the preconditions, Lemma B.2 implies that $\delta_{t+1} \leq \delta_t$ and $\gamma_{t+1} \geq \min\{\gamma_t, 1-1/\rho\}$, so the next iteration also satisfies the preconditions. Hence by induction, the preconditions hold for all iterations in T_1 . Moreover, for all $i \in [k]$, we have

$$|r_{i,0}| \geq \frac{1}{1-\gamma_0};$$

and while $t \in T_1$: (i) $|r_{i,t}|$ increases at a linear rate while $r_{i,t}^2 \leq 2\rho^2$, and (ii) $|r_{i,t}|$ increases at a quadratic rate while $\rho^2 \leq r_{i,t}^2 \leq \frac{1-\delta_t-1/\rho}{\kappa\delta_t}$. (The specific rates are given, respectively, in Lemma B.2, claims 1 and 2.) Since $\frac{1-\delta_t-1/\rho}{\kappa\delta_t} \leq \frac{\tilde{\lambda}_1}{2\kappa\tilde{\epsilon}}$, it follows that $\min_{i \neq 1} r_{i,t}^2 \leq \frac{1-\delta_t-1/\rho}{\kappa\delta_t}$ for at most

$$\frac{2}{\gamma_0} \ln \left(\frac{\sqrt{2\rho^2}}{1-\gamma_0} \right) + \ln \left(\frac{\ln \frac{\tilde{\lambda}_1}{2\kappa\tilde{\epsilon}}}{\ln \sqrt{2}} \right) = O \left(\frac{1}{\gamma_0} + \log \log \frac{\tilde{\lambda}_1}{\tilde{\epsilon}} \right) \quad (22)$$

iterations in T_1 . As soon as $\min_{i \neq 1} r_{i,t}^2 > \frac{1-\delta_t-1/\rho}{\kappa\delta_t}$, we have that in the next iteration,

$$R_{t+1} > \frac{1-\delta_t-1/\rho}{\kappa\delta_t} \cdot \frac{\tilde{\lambda}_{\min}}{\tilde{\lambda}_1} \cdot \frac{1}{\sqrt{k}} \geq \frac{7}{\sqrt{k}};$$

and all the while R_t is growing at a linear rate (given in Lemma B.2, claim 5). Therefore, there are at most an additional

$$1 + \frac{3}{\gamma_0} \ln\left(\frac{1+8\kappa}{7/\sqrt{k}}\right) = O\left(\frac{\log(k\kappa)}{\gamma_0}\right) \quad (23)$$

iterations in T_1 over that counted in (22). Therefore, by combining the counts in (22) and (23), we have that the number of iterations in the first phase satisfies

$$|T_1| = O\left(\log \log \frac{\tilde{\lambda}_1}{\tilde{\epsilon}} + \frac{\log(k\kappa)}{\gamma_0}\right).$$

We now analyze the second phase, *i.e.*, the iterates in $T_2 := \{t \geq 0 : t \notin T_1, R_t < 1/\alpha\}$. Define

$$\alpha := \frac{3\tilde{\epsilon}}{\tilde{\lambda}_1}, \quad \beta := \frac{1}{1+2\kappa\rho^2} = \frac{1}{1+8\kappa}.$$

Note that for the initial iteration $t' := \min T_2$, we have that $R_{t'} \geq 1+2\kappa\rho^2 = 1+8\kappa = 1/\beta$, and by Proposition B.1, $\gamma_{t'} \geq 1 - \kappa/(1+8\kappa) > 7/8$. It can be checked that $\delta_t, \gamma_t, \alpha$, and β satisfy the preconditions of Lemma B.3 for this initial iteration t' . For all $t \in T_2$ satisfying these preconditions, Lemma B.3 implies that $R_{t+1} \geq \min\{R_t, 1/\alpha\}$, $\theta_{1,t+1}^2 \geq \min\{\theta_{1,t}^2, 1/(1+\alpha^2)\}$ (via Proposition B.1), $\delta_{t+1} \leq \max\{\delta_t, (1+\alpha)^2\tilde{\epsilon}/\tilde{\lambda}_1\}$ (using the definition of δ_t), and $\gamma_{t+1} \geq \min\{\gamma_t, 1 - \alpha\kappa\}$ (via Proposition B.1). Hence the next iteration $t+1$ also satisfies the preconditions, and by induction, so do all iterations in T_2 . To bound the number of iterations in T_2 , observe that R_t increases at a quadratic rate until $R_t \geq 1/\alpha$, so

$$|T_2| \leq \ln\left(\frac{\ln(1/\alpha)}{\ln((1/\beta)/(2\kappa))}\right) < \ln\left(\frac{\ln \frac{\tilde{\lambda}_1}{3\tilde{\epsilon}}}{\ln 4}\right) = O\left(\log \log \frac{\tilde{\lambda}_1}{\tilde{\epsilon}}\right). \quad (24)$$

Therefore the total number of iterations before $R_t \geq 1/\alpha$ is

$$O\left(\frac{\log(k\kappa)}{\gamma_0} + \log \log \frac{\tilde{\lambda}_1}{\tilde{\epsilon}}\right).$$

After $R_{t''} \geq 1/\alpha$ (for $t'' := \max(T_1 \cup T_2) + 1$), we have

$$\theta_{1,t''}^2 \geq \frac{1/\alpha^2}{1+1/\alpha^2} \geq 1 - \alpha^2 \geq 1 - \left(\frac{3\tilde{\epsilon}}{\tilde{\lambda}_1}\right)^2.$$

Therefore, the vector $\theta_{t''}$ satisfies the condition for property (13) of \tilde{E} to hold. Now we apply Lemma B.3 using $\tilde{\epsilon}/p$ in place of $\tilde{\epsilon}$, including in the definition of δ_t (which we call $\bar{\delta}_t$):

$$\bar{\delta}_t := \frac{\tilde{\epsilon}}{p\tilde{\lambda}_1\theta_{1,t}^2};$$

we also replace α and β with $\bar{\alpha}$ and $\bar{\beta}$, which we set to

$$\bar{\alpha} := \frac{3\tilde{\epsilon}}{p\tilde{\lambda}_1}, \quad \bar{\beta} := \frac{3\tilde{\epsilon}}{\tilde{\lambda}_1}.$$

It can be checked that $\bar{\delta}_{t''} \in (0, 1/2)$, $\gamma_{t''} \geq 1 - 3\tilde{\epsilon}\kappa/\lambda_1 > 0$,

$$\frac{\bar{\alpha}}{(1+\bar{\alpha})(1+\bar{\alpha}^2)} \geq \frac{\tilde{\epsilon}}{p(1-3\tilde{\epsilon}\kappa/\lambda_1)\tilde{\lambda}_1} \geq \frac{\tilde{\epsilon}}{p\gamma_{t''}\tilde{\lambda}_1}, \quad \frac{\bar{\alpha}}{2(1+\bar{\alpha})(1+\bar{\beta}^2)} \geq \frac{\tilde{\epsilon}}{p\tilde{\lambda}_1}.$$

Therefore, the preconditions of Lemma B.3 are satisfied for the initial iteration t'' in this final phase, and by the same arguments as before, the preconditions hold for all subsequent iterations $t \geq t''$. Initially, we have $R_{t''} \geq 1/\alpha \geq 1/\bar{\beta}$, and by Lemma B.3, we have that R_t increases at a quadratic rate in this final phase until $R_t \geq 1/\bar{\alpha}$. So the number of iterations before $R_t \geq 1/\bar{\alpha}$ can be bounded as

$$\ln\left(\frac{\ln(1/\bar{\alpha})}{\ln((1/\bar{\beta})/(2\kappa))}\right) = \ln\left(\frac{\ln\frac{p\tilde{\lambda}_1}{3\tilde{\epsilon}}}{\ln\left(\frac{\lambda_1}{3\tilde{\epsilon}} \cdot \frac{1}{2\kappa}\right)}\right) \leq \ln \ln \frac{p\tilde{\lambda}_1}{3\tilde{\epsilon}} = O\left(\log \log \frac{p\tilde{\lambda}_1}{\tilde{\epsilon}}\right).$$

Once $R_t \geq 1/\bar{\alpha}$, we have

$$\theta_{1,t}^2 \geq 1 - \left(\frac{3\tilde{\epsilon}}{p\tilde{\lambda}_1}\right)^2.$$

Since $\text{sign}(\theta_{1,t}) = r_{1,t} \geq r_{1,t-1}^2 \cdot (1 - \bar{\delta}_{t-1}) / (1 + \kappa\bar{\delta}_{t-1}r_{1,t-1}^2) = (1 - \bar{\delta}_{t-1}) / (1 + \kappa\bar{\delta}_{t-1}) > 0$ by Proposition B.2, we have $\theta_{1,t} > 0$. Therefore we can conclude that

$$\|\theta_t - v_1\| = \sqrt{2(1 - \theta_{1,t})} \leq \sqrt{2\left(1 - \sqrt{1 - (3\tilde{\epsilon}/(p\tilde{\lambda}_1))^2}\right)} \leq 4\tilde{\epsilon}/(p\tilde{\lambda}_1).$$

Finally,

$$\begin{aligned} |\tilde{T}(\theta_t, \theta_t, \theta_t) - \tilde{\lambda}_1| &= \left| \tilde{\lambda}_1(\theta_{1,t}^3 - 1) + \sum_{i=2}^k \tilde{\lambda}_i \theta_{i,t}^3 + \tilde{E}(\theta_t, \theta_t, \theta_t) \right| \\ &\leq \tilde{\lambda}_1 |\theta_{1,t}^3 - 1| + \sum_{i=2}^k \tilde{\lambda}_i |\theta_{i,t}| \theta_{i,t}^2 + \|\tilde{E}(I, \theta_t, \theta_t)\| \\ &\leq \tilde{\lambda}_1 (1 - \theta_{1,t} + |\theta_{1,t}(1 - \theta_{1,t}^2)|) + \max_{i \neq 1} \tilde{\lambda}_i |\theta_{i,t}| \sum_{i=2}^k \theta_{i,t}^2 + \|\tilde{E}(I, \theta_t, \theta_t)\| \\ &\leq \tilde{\lambda}_1 (1 - \theta_{1,t} + |\theta_{1,t}(1 - \theta_{1,t}^2)|) + \max_{i \neq 1} \tilde{\lambda}_i \sqrt{1 - \theta_{1,t}^2} \sum_{i=2}^k \theta_{i,t}^2 + \|\tilde{E}(I, \theta_t, \theta_t)\| \\ &= \tilde{\lambda}_1 (1 - \theta_{1,t} + |\theta_{1,t}(1 - \theta_{1,t}^2)|) + \max_{i \neq 1} \tilde{\lambda}_i (1 - \theta_{1,t}^2)^{3/2} + \|\tilde{E}(I, \theta_t, \theta_t)\| \\ &\leq \tilde{\lambda}_1 \cdot 3 \left(\frac{3\tilde{\epsilon}}{p\tilde{\lambda}_1}\right)^2 + \kappa \tilde{\lambda}_1 \cdot \left(\frac{3\tilde{\epsilon}}{p\tilde{\lambda}_1}\right)^3 + \frac{\tilde{\epsilon}}{p} \\ &\leq \frac{(27\kappa \cdot (\tilde{\epsilon}/p\tilde{\lambda}_1)^2 + 2)\tilde{\epsilon}}{p}. \end{aligned}$$

■

B.3 Deflation

Lemma B.5 Fix some $\tilde{\epsilon} \geq 0$. Let $\{v_1, v_2, \dots, v_k\}$ be an orthonormal basis for \mathbb{R}^k , and $\lambda_1, \lambda_2, \dots, \lambda_k \geq 0$ with $\lambda_{\min} := \min_{i \in [k]} \lambda_i$. Also, let $\{\hat{v}_1, \hat{v}_2, \dots, \hat{v}_k\}$ be a set of unit vectors in \mathbb{R}^k (not necessarily orthogonal), $\hat{\lambda}_1, \hat{\lambda}_2, \dots, \hat{\lambda}_k \geq 0$ be non-negative scalars, and define

$$\mathcal{E}_i := \lambda_i v_i^{\otimes 3} - \hat{\lambda}_i \hat{v}_i^{\otimes 3}, \quad i \in [k].$$

Pick any $t \in [k]$. If

$$\begin{aligned} |\hat{\lambda}_i - \lambda_i| &\leq \tilde{\epsilon}, \\ \|\hat{v}_i - v_i\| &\leq \min\{\sqrt{2}, 2\tilde{\epsilon}/\lambda_i\} \end{aligned}$$

for all $i \in [t]$, then for any unit vector $u \in S^{k-1}$,

$$\begin{aligned} \left\| \sum_{i=1}^t \mathcal{E}_i(I, u, u) \right\|_2^2 &\leq \left(4(5 + 11\tilde{\epsilon}/\lambda_{\min})^2 + 128(1 + \tilde{\epsilon}/\lambda_{\min})^2(\tilde{\epsilon}/\lambda_{\min})^2 \right) \tilde{\epsilon}^2 \sum_{i=1}^t (u^\top v_i)^2 \\ &\quad + 64(1 + \tilde{\epsilon}/\lambda_{\min})^2 \tilde{\epsilon}^2 \sum_{i=1}^t (\tilde{\epsilon}/\lambda_i)^2 + 2048(1 + \tilde{\epsilon}/\lambda_{\min})^2 \tilde{\epsilon}^2 \left(\sum_{i=1}^t (\tilde{\epsilon}/\lambda_i)^3 \right)^2. \end{aligned}$$

In particular, for any $\Delta \in (0, 1)$, there exists a constant $\Delta' > 0$ (depending only on Δ) such that $\tilde{\epsilon} \leq \Delta' \lambda_{\min}/\sqrt{k}$ implies

$$\left\| \sum_{i=1}^t \mathcal{E}_i(I, u, u) \right\|_2^2 \leq \left(\Delta + 100 \sum_{i=1}^t (u^\top v_i)^2 \right) \tilde{\epsilon}^2.$$

Proof For any unit vector u and $i \in [t]$, the error term

$$\mathcal{E}_i(I, u, u) = \lambda_i (u^\top v_i)^2 v_i - \hat{\lambda}_i (u^\top \hat{v}_i)^2 \hat{v}_i$$

lives in $\text{span}\{v_i, \hat{v}_i\}$; this space is the same as $\text{span}\{v_i, \hat{v}_i^\perp\}$, where

$$\hat{v}_i^\perp := \hat{v}_i - (v_i^\top \hat{v}_i) v_i$$

is the projection of \hat{v}_i onto the subspace orthogonal to v_i . Since $\|\hat{v}_i - v_i\|^2 = 2(1 - v_i^\top \hat{v}_i)$, it follows that

$$c_i := v_i^\top \hat{v}_i = 1 - \|\hat{v}_i - v_i\|^2/2 \geq 0$$

(the inequality follows from the assumption $\|\hat{v}_i - v_i\| \leq \sqrt{2}$, which in turn implies $0 \leq c_i \leq 1$). By the Pythagorean theorem and the above inequality for c_i ,

$$\|\hat{v}_i^\perp\|^2 = 1 - c_i^2 \leq \|\hat{v}_i - v_i\|^2.$$

Later, we will also need the following bound, which is easily derived from the above inequalities and the triangle inequality:

$$|1 - c_i^3| = |1 - c_i + c_i(1 - c_i^2)| \leq 1 - c_i + |c_i(1 - c_i^2)| \leq 1.5\|\hat{v}_i - v_i\|^2.$$

We now express $\mathcal{E}_i(I, u, u)$ in terms of the coordinate system defined by v_i and \hat{v}_i^\perp , depicted below. Define

$$a_i := u^\top v_i \quad \text{and} \quad b_i := u^\top (\hat{v}_i^\perp / \|\hat{v}_i^\perp\|).$$

(Note that the part of u living in $\text{span}\{v_i, \hat{v}_i^\perp\}^\perp$ is irrelevant for analyzing $\mathcal{E}_i(I, u, u)$.) We have

$$\begin{aligned} \mathcal{E}_i(I, u, u) &= \lambda_i (u^\top v_i)^2 v_i - \hat{\lambda}_i (u^\top \hat{v}_i)^2 \hat{v}_i \\ &= \lambda_i a_i^2 v_i - \hat{\lambda}_i (a_i c_i + \|\hat{v}_i^\perp\| b_i)^2 (c_i v_i + \hat{v}_i^\perp) \\ &= \lambda_i a_i^2 v_i - \hat{\lambda}_i (a_i^2 c_i^2 + 2\|\hat{v}_i^\perp\| a_i b_i c_i + \|\hat{v}_i^\perp\|^2 b_i^2) c_i v_i - \hat{\lambda}_i (a_i c_i + \|\hat{v}_i^\perp\| b_i)^2 \hat{v}_i^\perp \\ &= \underbrace{\left((\lambda_i - \hat{\lambda}_i c_i^3) a_i^2 - 2\hat{\lambda}_i \|\hat{v}_i^\perp\| a_i b_i c_i^2 - \hat{\lambda}_i \|\hat{v}_i^\perp\|^2 b_i^2 c_i \right)}_{=: A_i} v_i - \underbrace{\hat{\lambda}_i \|\hat{v}_i^\perp\| (a_i c_i + \|\hat{v}_i^\perp\| b_i)^2}_{=: B_i} (\hat{v}_i^\perp / \|\hat{v}_i^\perp\|) \\ &= A_i v_i - B_i (\hat{v}_i^\perp / \|\hat{v}_i^\perp\|). \end{aligned}$$

The overall error can also be expressed in terms of the A_i and B_i :

$$\begin{aligned} \left\| \sum_{i=1}^t \mathcal{E}_i(I, u, u) \right\|_2^2 &= \left\| \sum_{i=1}^t A_i v_i - \sum_{i=1}^t B_i (\hat{v}_i^\perp / \|\hat{v}_i^\perp\|) \right\|_2^2 \\ &\leq 2 \left\| \sum_{i=1}^t A_i v_i \right\|_2^2 + 2 \left\| \sum_{i=1}^t B_i (\hat{v}_i^\perp / \|\hat{v}_i^\perp\|) \right\|_2^2 \\ &\leq 2 \sum_{i=1}^t A_i^2 + 2 \left(\sum_{i=1}^t |B_i| \right)^2 \end{aligned} \tag{25}$$

where the first inequality uses the fact $(x + y)^2 \leq 2(x^2 + y^2)$ and the triangle inequality, and the second inequality uses the orthonormality of the v_i and the triangle inequality.

It remains to bound A_i^2 and $|B_i|$ in terms of $|a_i|$, λ_i , and $\tilde{\epsilon}$. The first term, A_i^2 , can be bounded using the triangle inequality and the various bounds on $|\lambda_i - \hat{\lambda}_i|$, $\|\hat{v}_i - v_i\|$, $\|\hat{v}_i^\perp\|$, and c_i :

$$\begin{aligned} |A_i| &\leq (|\lambda_i - \hat{\lambda}_i| c_i^3 + \lambda_i |c_i^3 - 1|) a_i^2 + 2(\lambda_i + |\lambda_i - \hat{\lambda}_i|) \|\hat{v}_i^\perp\| |a_i b_i| c_i^2 + (\lambda_i + |\lambda_i - \hat{\lambda}_i|) \|\hat{v}_i^\perp\|^2 b_i^2 c_i \\ &\leq (|\lambda_i - \hat{\lambda}_i| + 1.5\lambda_i \|\hat{v}_i - v_i\| + 2(\lambda_i + |\lambda_i - \hat{\lambda}_i|) \|\hat{v}_i - v_i\|) |a_i| + (\lambda_i + |\lambda_i - \hat{\lambda}_i|) \|\hat{v}_i - v_i\|^2 \\ &\leq (\tilde{\epsilon} + 7\tilde{\epsilon}^2/\lambda_i + 4\tilde{\epsilon} + 4\tilde{\epsilon}^2/\lambda_i) |a_i| + 4\tilde{\epsilon}^2/\lambda_i + \tilde{\epsilon}^3/\lambda_i^2 \\ &= (5 + 11\tilde{\epsilon}/\lambda_i) \tilde{\epsilon} |a_i| + 4(1 + \tilde{\epsilon}/\lambda_i) \tilde{\epsilon}^2/\lambda_i, \end{aligned}$$

and therefore (via $(x + y)^2 \leq 2(x^2 + y^2)$)

$$A_i^2 \leq 2(5 + 11\tilde{\epsilon}/\lambda_i)^2 \tilde{\epsilon}^2 a_i^2 + 32(1 + \tilde{\epsilon}/\lambda_i)^2 \tilde{\epsilon}^4/\lambda_i^2.$$

The second term, $|B_i|$, is bounded similarly:

$$\begin{aligned} |B_i| &\leq 2(\lambda_i + |\lambda_i - \hat{\lambda}_i|) \|\hat{v}_i^\perp\|^2 (a_i^2 + \|\hat{v}_i^\perp\|^2) \\ &\leq 2(\lambda_i + |\lambda_i - \hat{\lambda}_i|) \|\hat{v}_i - v_i\|^2 (a_i^2 + \|\hat{v}_i - v_i\|^2) \\ &\leq 8(1 + \tilde{\epsilon}/\lambda_i) (\tilde{\epsilon}^2/\lambda_i) a_i^2 + 32(1 + \tilde{\epsilon}/\lambda_i) \tilde{\epsilon}^4/\lambda_i^3. \end{aligned}$$

Therefore, using the inequality from (25) and again $(x + y)^2 \leq 2(x^2 + y^2)$,

$$\begin{aligned}
 \left\| \sum_{i=1}^t \mathcal{E}_i(I, u, u) \right\|_2^2 &\leq 2 \sum_{i=1}^t A_i^2 + 2 \left(\sum_{i=1}^t |B_i| \right)^2 \\
 &\leq 4(5 + 11\tilde{\epsilon}/\lambda_{\min})^2 \tilde{\epsilon}^2 \sum_{i=1}^t a_i^2 + 64(1 + \tilde{\epsilon}/\lambda_{\min})^2 \tilde{\epsilon}^2 \sum_{i=1}^t (\tilde{\epsilon}/\lambda_i)^2 \\
 &\quad + 2 \left(8(1 + \tilde{\epsilon}/\lambda_{\min})(\tilde{\epsilon}^2/\lambda_{\min}) \sum_{i=1}^t a_i^2 + 32(1 + \tilde{\epsilon}/\lambda_{\min})\tilde{\epsilon} \sum_{i=1}^t (\tilde{\epsilon}/\lambda_i)^3 \right)^2 \\
 &\leq 4(5 + 11\tilde{\epsilon}/\lambda_{\min})^2 \tilde{\epsilon}^2 \sum_{i=1}^t a_i^2 + 64(1 + \tilde{\epsilon}/\lambda_{\min})^2 \tilde{\epsilon}^2 \sum_{i=1}^t (\tilde{\epsilon}/\lambda_i)^2 \\
 &\quad + 128(1 + \tilde{\epsilon}/\lambda_{\min})^2 (\tilde{\epsilon}/\lambda_{\min})^2 \tilde{\epsilon}^2 \sum_{i=1}^t a_i^2 \\
 &\quad + 2048(1 + \tilde{\epsilon}/\lambda_{\min})^2 \tilde{\epsilon}^2 \left(\sum_{i=1}^t (\tilde{\epsilon}/\lambda_i)^3 \right)^2 \\
 &= \left(4(5 + 11\tilde{\epsilon}/\lambda_{\min})^2 + 128(1 + \tilde{\epsilon}/\lambda_{\min})^2 (\tilde{\epsilon}/\lambda_{\min})^2 \right) \tilde{\epsilon}^2 \sum_{i=1}^t a_i^2 \\
 &\quad + 64(1 + \tilde{\epsilon}/\lambda_{\min})^2 \tilde{\epsilon}^2 \sum_{i=1}^t (\tilde{\epsilon}/\lambda_i)^2 + 2048(1 + \tilde{\epsilon}/\lambda_{\min})^2 \tilde{\epsilon}^2 \left(\sum_{i=1}^t (\tilde{\epsilon}/\lambda_i)^3 \right)^2.
 \end{aligned}$$

■

B.4 Proof of the Main Theorem

Theorem B.1 *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 has operator norm $\epsilon := \|E\|$. 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*

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

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 1 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 π on $[k]$ such that

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

and

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

Proof We prove by induction that for each $i \in [k]$ (corresponding to the i -th call to Algorithm 1), with probability at least $1 - i\eta/k$, there exists a permutation π on $[k]$ such that the following assertions hold.

1. For all $j \leq i$, $\|v_{\pi(j)} - \hat{v}_j\| \leq 8\epsilon/\lambda_{\pi(j)}$ and $|\lambda_{\pi(j)} - \hat{\lambda}_j| \leq 12\epsilon$.
2. The error tensor

$$\tilde{E}_{i+1} := \left(\hat{T} - \sum_{j \leq i} \hat{\lambda}_j \hat{v}_j^{\otimes 3} \right) - \sum_{j \geq i+1} \lambda_{\pi(j)} v_{\pi(j)}^{\otimes 3} = E + \sum_{j \leq i} \left(\lambda_{\pi(j)} v_{\pi(j)}^{\otimes 3} - \hat{\lambda}_j \hat{v}_j^{\otimes 3} \right)$$

satisfies

$$\|\tilde{E}_{i+1}(I, u, u)\| \leq 56\epsilon, \quad \forall u \in S^{k-1}; \quad (26)$$

$$\|\tilde{E}_{i+1}(I, u, u)\| \leq 2\epsilon, \quad \forall u \in S^{k-1} \text{ s.t. } \exists j \geq i+1 \bullet (u^\top v_{\pi(j)})^2 \geq 1 - (168\epsilon/\lambda_{\pi(j)})^2. \quad (27)$$

We actually take $i = 0$ as the base case, so we can ignore the first assertion, and just observe that for $i = 0$,

$$\tilde{E}_1 = \hat{T} - \sum_{j=1}^k \lambda_j v_j^{\otimes 3} = E.$$

We have $\|\tilde{E}_1\| = \|E\| = \epsilon$, and therefore the second assertion holds.

Now fix some $i \in [k]$, and assume as the inductive hypothesis that, with probability at least $1 - (i-1)\eta/k$, there exists a permutation π such that two assertions above hold for $i-1$ (call this Event_{i-1}). The i -th call to Algorithm 1 takes as input

$$\tilde{T}_i := \hat{T} - \sum_{j \leq i-1} \hat{\lambda}_j \hat{v}_j^{\otimes 3},$$

which is intended to be an approximation to

$$T_i := \sum_{j \geq i} \lambda_{\pi(j)} v_{\pi(j)}^{\otimes 3}.$$

Observe that

$$\tilde{T}_i - T_i = \tilde{E}_i,$$

which satisfies the second assertion in the inductive hypothesis. We may write $T_i = \sum_{l=1}^k \tilde{\lambda}_l v_l^{\otimes 3}$ where $\tilde{\lambda}_l = \lambda_l$ whenever $\pi^{-1}(l) \geq i$, and $\tilde{\lambda}_l = 0$ whenever $\pi^{-1}(l) \leq i-1$. This form is used when referring to \tilde{T} or the $\tilde{\lambda}_i$ in preceding lemmas (in particular, Lemma B.1 and Lemma B.4).

By Lemma B.1, with conditional probability at least $1 - \eta/k$ given Event_{i-1} , at least one of $\theta_0^{(\tau)}$ for $\tau \in [L]$ is γ -separated relative to $\pi(j_{\max})$, where $j_{\max} := \arg \max_{j \geq i} \lambda_{\pi(j)}$, (for $\gamma = 0.01$; call this Event'_i ; note that the application of Lemma B.1 determines C_3). Therefore $\Pr[\text{Event}_{i-1} \cap \text{Event}'_i] = \Pr[\text{Event}'_i | \text{Event}_{i-1}] \Pr[\text{Event}_{i-1}] \geq (1 - \eta/k)(1 - (i-1)\eta/k) \geq 1 - i\eta/k$. It remains to show that $\text{Event}_{i-1} \cap \text{Event}'_i \subseteq \text{Event}_i$; so henceforth we condition on $\text{Event}_{i-1} \cap \text{Event}'_i$.

Set

$$C_1 := \min \{ (56 \cdot 9 \cdot 102)^{-1}, (100 \cdot 168)^{-1}, \Delta' \text{ from Lemma B.5 with } \Delta = 1/50 \}. \quad (28)$$

For all $\tau \in [L]$ such that $\theta_0^{(\tau)}$ is γ -separated relative to $\pi(j_{\max})$, we have (i) $|\theta_{j_{\max},0}^{(\tau)}| \geq 1/\sqrt{k}$, and (ii) that by Lemma B.4 (using $\tilde{\epsilon}/p := 2\epsilon$, $\kappa := 1$, and $i^* := \pi(j_{\max})$, and providing C_2),

$$|\tilde{T}_i(\theta_N^{(\tau)}, \theta_N^{(\tau)}, \theta_N^{(\tau)}) - \lambda_{\pi(j_{\max})}| \leq 5\epsilon$$

(notice by definition that $\gamma \geq 1/100$ implies $\gamma_0 \geq 1/(1+\gamma) \geq 1/101$, thus it follows from the bounds on the other quantities that $\tilde{\epsilon} = 2p\epsilon \leq 56C_1 \cdot \frac{\lambda_{\min}}{k} < \frac{\gamma_0}{2(1+8\kappa)} \cdot \tilde{\lambda}_{\min} \cdot \theta_{i^*,0}^2$ as necessary). Therefore $\theta_N := \theta_N^{(\tau^*)}$ must satisfy

$$\tilde{T}_i(\theta_N, \theta_N, \theta_N) = \max_{\tau \in [L]} \tilde{T}_i(\theta_N^{(\tau)}, \theta_N^{(\tau)}, \theta_N^{(\tau)}) \geq \max_{j \geq i} \lambda_{\pi(j)} - 5\epsilon = \lambda_{\pi(j_{\max})} - 5\epsilon.$$

On the other hand, by the triangle inequality,

$$\begin{aligned} \tilde{T}_i(\theta_N, \theta_N, \theta_N) &\leq \sum_{j \geq i} \lambda_{\pi(j)} \theta_{\pi(j),N}^3 + |\tilde{E}_i(\theta_N, \theta_N, \theta_N)| \\ &\leq \sum_{j \geq i} \lambda_{\pi(j)} |\theta_{\pi(j),N}| \theta_{\pi(j),N}^2 + 56\epsilon \\ &\leq \lambda_{\pi(j^*)} |\theta_{\pi(j^*),N}| + 56\epsilon \end{aligned}$$

where $j^* := \arg \max_{j \geq i} \lambda_{\pi(j)} |\theta_{\pi(j),N}|$. Therefore

$$\lambda_{\pi(j^*)} |\theta_{\pi(j^*),N}| \geq \lambda_{\pi(j_{\max})} - 5\epsilon - 56\epsilon \geq \frac{4}{5} \lambda_{\pi(j_{\max})}.$$

Squaring both sides and using the fact that $\theta_{\pi(j^*),N}^2 + \theta_{\pi(j),N}^2 \leq 1$ for any $j \neq j^*$,

$$\begin{aligned} (\lambda_{\pi(j^*)} \theta_{\pi(j^*),N})^2 &\geq \frac{16}{25} (\lambda_{\pi(j_{\max})} \theta_{\pi(j^*),N})^2 + \frac{16}{25} (\lambda_{\pi(j_{\max})} \theta_{\pi(j),N})^2 \\ &\geq \frac{16}{25} (\lambda_{\pi(j^*)} \theta_{\pi(j^*),N})^2 + \frac{16}{25} (\lambda_{\pi(j)} \theta_{\pi(j),N})^2 \end{aligned}$$

which in turn implies

$$\lambda_{\pi(j)}|\theta_{\pi(j),N}| \leq \frac{3}{4}\lambda_{\pi(j^*)}|\theta_{\pi(j^*),N}|, \quad j \neq j^*.$$

This means that θ_N is $(1/4)$ -separated relative to $\pi(j^*)$. Also, observe that

$$|\theta_{\pi(j^*),N}| \geq \frac{4}{5} \cdot \frac{\lambda_{\pi(j_{\max})}}{\lambda_{\pi(j^*)}} \geq \frac{4}{5}, \quad \frac{\lambda_{\pi(j_{\max})}}{\lambda_{\pi(j^*)}} \leq \frac{5}{4}.$$

Therefore by Lemma B.4 (using $\tilde{\epsilon}/p := 2\epsilon$, $\gamma := 1/4$, and $\kappa := 5/4$), executing another N power iterations starting from θ_N gives a vector $\hat{\theta}$ that satisfies

$$\|\hat{\theta} - v_{\pi(j^*)}\| \leq \frac{8\epsilon}{\lambda_{\pi(j^*)}}, \quad |\hat{\lambda} - \lambda_{\pi(j^*)}| \leq 5\epsilon.$$

Since $\hat{v}_i = \hat{\theta}$ and $\hat{\lambda}_i = \hat{\lambda}$, the first assertion of the inductive hypothesis is satisfied, as we can modify the permutation π by swapping $\pi(i)$ and $\pi(j^*)$ without affecting the values of $\{\pi(j) : j \leq i-1\}$ (recall $j^* \geq i$).

We now argue that \tilde{E}_{i+1} has the required properties to complete the inductive step. By Lemma B.5 (using $\tilde{\epsilon} := 5\epsilon$ and $\Delta := 1/50$, the latter providing one upper bound on C_1 as per (28)), we have for any unit vector $u \in S^{k-1}$,

$$\left\| \left(\sum_{j \leq i} \left(\lambda_{\pi(j)} v_{\pi(j)}^{\otimes 3} - \hat{\lambda}_j \hat{v}_j^{\otimes 3} \right) \right) (I, u, u) \right\| \leq \left(1/50 + 100 \sum_{j=1}^i (u^\top v_{\pi(j)})^2 \right)^{1/2} 5\epsilon \leq 55\epsilon. \quad (29)$$

Therefore by the triangle inequality,

$$\|\tilde{E}_{i+1}(I, u, u)\| \leq \|E(I, u, u)\| + \left\| \left(\sum_{j \leq i} \left(\lambda_{\pi(j)} v_{\pi(j)}^{\otimes 3} - \hat{\lambda}_j \hat{v}_j^{\otimes 3} \right) \right) (I, u, u) \right\| \leq 56\epsilon.$$

Thus the bound (26) holds.

To prove that (27) holds, pick any unit vector $u \in S^{k-1}$ such that there exists $j' \geq i+1$ with $(u^\top v_{\pi(j')})^2 \geq 1 - (168\epsilon/\lambda_{\pi(j')})^2$. We have, via the second bound on C_1 in (28) and the corresponding assumed bound $\epsilon \leq C_1 \cdot \frac{\lambda_{\min}}{k}$,

$$100 \sum_{j=1}^i (u^\top v_{\pi(j)})^2 \leq 100 \left(1 - (u^\top v_{\pi(j')})^2 \right) \leq 100 \left(\frac{168\epsilon}{\lambda_{\pi(j')}} \right)^2 \leq \frac{1}{50},$$

and therefore

$$\left(1/50 + 100 \sum_{j=1}^i (u^\top v_{\pi(j)})^2 \right)^{1/2} 5\epsilon \leq (1/50 + 1/50)^{1/2} 5\epsilon \leq \epsilon.$$

By the triangle inequality, we have $\|\tilde{E}_{i+1}(I, u, u)\| \leq 2\epsilon$. Therefore (27) holds, so the second assertion of the inductive hypothesis holds. Thus $\text{Event}_{i-1} \cap \text{Event}'_i \subseteq \text{Event}_i$, and $\Pr[\text{Event}_i] \geq \Pr[\text{Event}_{i-1} \cap \text{Event}'_i] \geq 1 - i\eta/k$. We conclude that by the induction principle,

there exists a permutation π such that two assertions hold for $i = k$, with probability at least $1 - \eta$.

From the last induction step ($i = k$), it is also clear from (29) that $\|T - \sum_{j=1}^k \hat{\lambda}_j \hat{v}_j^{\otimes 3}\| \leq 55\epsilon$ (in $\text{Event}_{k-1} \cap \text{Event}'_k$). This completes the proof of the theorem. \blacksquare

Appendix C. Variant of Robust Power Method that uses a Stopping Condition

In this section we analyze a variant of Algorithm 1 that uses a stopping condition. The variant is described in Algorithm 2. The key difference is that the inner for-loop is repeated until a stopping condition is satisfied (rather than explicitly L times). The stopping condition ensures that the power iteration is converging to an eigenvector, and it will be satisfied within $\text{poly}(k)$ random restarts with high probability. The condition depends on one new quantity, r , which should be set to $r := k - \#$ deflation steps so far (*i.e.*, the first call to Algorithm 2 uses $r = k$, the second call uses $r = k - 1$, and so on).

Algorithm 2 Robust tensor power method with stopping condition

input symmetric tensor $\tilde{T} \in \mathbb{R}^{k \times k \times k}$, number of iterations N , expected rank r .

output the estimated eigenvector/eigenvalue pair; the deflated tensor.

- 1: **repeat**
- 2: Draw θ_0 uniformly at random from the unit sphere in \mathbb{R}^k .
- 3: **for** $t = 1$ to N **do**
- 4: Compute power iteration update

$$\theta_t := \frac{\tilde{T}(I, \theta_{t-1}, \theta_{t-1})}{\|\tilde{T}(I, \theta_{t-1}, \theta_{t-1})\|} \quad (30)$$

- 5: **end for**
- 6: **until** the following stopping condition is satisfied:

$$|\tilde{T}(\theta_N, \theta_N, \theta_N)| \geq \max \left\{ \frac{1}{2\sqrt{r}} \|\tilde{T}\|_F, \frac{1}{1.05} \|\tilde{T}(I, I, \theta_N)\|_F \right\}.$$

- 7: Do N power iteration updates (30) starting from θ_N to obtain $\hat{\theta}$, and set $\hat{\lambda} := \tilde{T}(\hat{\theta}, \hat{\theta}, \hat{\theta})$.
 - 8: **return** the estimated eigenvector/eigenvalue pair $(\hat{\theta}, \hat{\lambda})$; the deflated tensor $\tilde{T} - \hat{\lambda} \hat{\theta}^{\otimes 3}$.
-

C.1 Stopping Condition Analysis

For a matrix A , we use $\|A\|_F := (\sum_{i,j} A_{i,j}^2)^{1/2}$ to denote its Frobenius norm. For a third-order tensor A , we use $\|A\|_F := (\sum_i \|A(I, I, e_i)\|_F^2)^{1/2} = (\sum_i \|A(I, I, v_i)\|_F^2)^{1/2}$.

Define \tilde{T} as before in (11):

$$\tilde{T} := \sum_{i=1}^k \tilde{\lambda}_i v_i^{\otimes 3} + \tilde{E}.$$

We assume \tilde{E} is a symmetric tensor such that, for some constant $p > 1$,

$$\begin{aligned} \|\tilde{E}(I, u, u)\| &\leq \tilde{\epsilon}, \quad \forall u \in S^{k-1}; \\ \|\tilde{E}(I, u, u)\| &\leq \tilde{\epsilon}/p, \quad \forall u \in S^{k-1} \text{ s.t. } (u^\top v_1)^2 \geq 1 - (3\tilde{\epsilon}/\tilde{\lambda}_1)^2; \\ \|\tilde{E}\|_F &\leq \tilde{\epsilon}_F. \end{aligned}$$

Assume that not all $\tilde{\lambda}_i$ are zero, and define

$$\begin{aligned} \tilde{\lambda}_{\min} &:= \min\{\tilde{\lambda}_i : i \in [k], \tilde{\lambda}_i > 0\}, & \tilde{\lambda}_{\max} &:= \max\{\tilde{\lambda}_i : i \in [k]\}, \\ \ell &:= |\{i \in [k] : \tilde{\lambda}_i > 0\}|, & \tilde{\lambda}_{\text{avg}} &:= \left(\frac{1}{\ell} \sum_{i=1}^k \tilde{\lambda}_i^2\right)^{1/2}. \end{aligned}$$

We show in Lemma C.1 that if the stopping condition is satisfied by a vector θ , then it must be close to an eigenvector of \tilde{T} . Then in Lemma C.2, we show that the stopping condition is satisfied by θ_N when θ_0 is a good starting point (as per the conditions of Lemma B.4).

Lemma C.1 *Fix any vector $\theta = \sum_{i=1}^k \theta_i v_i$, and let $i^* := \arg \max_{i \in [k]} \tilde{\lambda}_i |\theta_i|$. Assume that $\ell \geq 1$ and that for some $\alpha \in (0, 1/20)$ and $\beta \geq 2\alpha/\sqrt{k}$,*

$$\tilde{\epsilon} \leq \alpha \cdot \frac{\tilde{\lambda}_{\min}}{\sqrt{k}}, \quad \tilde{\epsilon}_F \leq \sqrt{\ell} \left(\frac{1}{2} - \frac{\alpha}{\beta\sqrt{k}}\right) \cdot \tilde{\lambda}_{\text{avg}}.$$

If the stopping condition

$$|\tilde{T}(\theta, \theta, \theta)| \geq \max\left\{\frac{\beta}{\sqrt{\ell}} \|\tilde{T}\|_F, \frac{1}{1+\alpha} \|\tilde{T}(I, I, \theta)\|_F\right\} \quad (31)$$

holds, then

1. $\tilde{\lambda}_{i^*} \geq \beta \tilde{\lambda}_{\text{avg}}/2$ and $\tilde{\lambda}_{i^*} |\theta_{i^*}| > 0$;
2. $\max_{i \neq i^*} \tilde{\lambda}_i |\theta_i| \leq \sqrt{7\alpha} \cdot \tilde{\lambda}_{i^*} |\theta_{i^*}|$;
3. $\theta_{i^*} \geq 1 - 2\alpha$.

Proof Without loss of generality, assume $i^* = 1$. First, we claim that $\tilde{\lambda}_1 |\theta_1| > 0$. By the triangle inequality,

$$|\tilde{T}(\theta, \theta, \theta)| \leq \sum_{i=1}^k \tilde{\lambda}_i \theta_i^3 + |\tilde{E}(\theta, \theta, \theta)| \leq \sum_{i=1}^k \tilde{\lambda}_i |\theta_i| \theta_i^2 + \tilde{\epsilon} \leq \tilde{\lambda}_1 |\theta_1| + \tilde{\epsilon}.$$

Moreover,

$$\begin{aligned}
 \|\tilde{T}\|_F &\geq \left\| \sum_{i=1}^k \tilde{\lambda}_i v_i^{\otimes 3} \right\|_F - \|\tilde{E}\|_F \\
 &= \left(\sum_{j=1}^k \left\| \sum_{i=1}^k \tilde{\lambda}_i v_i v_i^\top (v_i^\top v_j) \right\|_F^2 \right)^{1/2} - \|\tilde{E}\|_F \\
 &= \left(\sum_{j=1}^k \left\| \tilde{\lambda}_j v_j v_j^\top \right\|_F^2 \right)^{1/2} - \|\tilde{E}\|_F \\
 &= \left(\sum_{j=1}^k \tilde{\lambda}_j^2 \right)^{1/2} - \|\tilde{E}\|_F \\
 &\geq \sqrt{\ell} \tilde{\lambda}_{\text{avg}} - \tilde{\epsilon}_F.
 \end{aligned}$$

By assumption, $|\tilde{T}(\theta, \theta, \theta)| \geq (\beta/\sqrt{\ell})\|\tilde{T}\|_F$, so

$$\tilde{\lambda}_1 |\theta_1| \geq \beta \tilde{\lambda}_{\text{avg}} - \frac{\beta}{\sqrt{\ell}} \tilde{\epsilon}_F - \tilde{\epsilon} \geq \beta \tilde{\lambda}_{\text{avg}} - \beta \left(\frac{1}{2} - \frac{\alpha}{\beta \sqrt{k}} \right) \tilde{\lambda}_{\text{avg}} - \frac{\alpha}{\sqrt{k}} \tilde{\lambda}_{\text{min}} \geq \frac{\beta}{2} \tilde{\lambda}_{\text{avg}}$$

where the second inequality follows from the assumptions on $\tilde{\epsilon}$ and $\tilde{\epsilon}_F$. Since $\beta > 0$, $\tilde{\lambda}_{\text{avg}} > 0$, and $|\theta_1| \leq 1$, it follows that

$$\tilde{\lambda}_1 \geq \frac{\beta}{2} \tilde{\lambda}_{\text{avg}}, \quad \tilde{\lambda}_1 |\theta_1| > 0.$$

This proves the first claim.

Now we prove the second claim. Define $\tilde{M} := \tilde{T}(I, I, \theta) = \sum_{i=1}^k \tilde{\lambda}_i \theta_i v_i v_i^\top + \tilde{E}(I, I, \theta)$ (a symmetric $k \times k$ matrix), and consider its eigenvalue decomposition

$$\tilde{M} = \sum_{i=1}^k \phi_i u_i u_i^\top$$

where, without loss of generality, $|\phi_1| \geq |\phi_2| \geq \dots \geq |\phi_k|$ and $\{u_1, u_2, \dots, u_k\}$ is an orthonormal basis. Let $M := \sum_{i=1}^k \tilde{\lambda}_i \theta_i v_i v_i^\top$, so $\tilde{M} = M + \tilde{E}(I, I, \theta)$. Note that the $\tilde{\lambda}_i |\theta_i|$ and $|\phi_i|$ are the singular values of M and \tilde{M} , respectively. We now show that the assumption on $|\tilde{T}(\theta, \theta, \theta)|$ implies that almost all of the energy in M is contained in its top singular component.

By Weyl's theorem,

$$|\phi_1| \leq \tilde{\lambda}_1 |\theta_1| + \|\tilde{M} - M\| \leq \tilde{\lambda}_1 |\theta_1| + \tilde{\epsilon}.$$

Next, observe that the assumption $\|\tilde{T}(I, I, \theta)\|_F \leq (1 + \alpha)\tilde{T}(\theta, \theta, \theta)$ is equivalent to $(1 + \alpha)\theta^\top \tilde{M} \theta \geq \|\tilde{M}\|_F$. Therefore, using the fact that $|\phi_1| = \max_{u \in S^{k-1}} |u^\top \tilde{M} u|$, the triangle

inequality, and the fact $\|A\|_F \leq \sqrt{k}\|A\|$ for any matrix $A \in \mathbb{R}^{k \times k}$,

$$\begin{aligned}
 (1 + \alpha)|\phi_1| &\geq (1 + \alpha)\theta^\top \tilde{M}\theta \geq \|\tilde{M}\|_F \\
 &\geq \left\| \sum_{i=1}^k \tilde{\lambda}_i \theta_i v_i v_i^\top \right\|_F - \|\tilde{E}(I, I, \theta)\|_F \\
 &\geq \left(\sum_{i=1}^k \tilde{\lambda}_i^2 \theta_i^2 \right)^{1/2} - \sqrt{k} \|\tilde{E}(I, I, \theta)\| \\
 &\geq \left(\sum_{i=1}^k \tilde{\lambda}_i^2 \theta_i^2 \right)^{1/2} - \sqrt{k} \tilde{\epsilon}.
 \end{aligned} \tag{32}$$

Combining these bounds on $|\phi_1|$ gives

$$\tilde{\lambda}_1 |\theta_1| + \tilde{\epsilon} \geq \frac{1}{1 + \alpha} \left[\left(\sum_{i=1}^k \tilde{\lambda}_i^2 \theta_i^2 \right)^{1/2} - \sqrt{k} \tilde{\epsilon} \right]. \tag{33}$$

The assumption $\tilde{\epsilon} \leq \alpha \tilde{\lambda}_{\min} / \sqrt{k}$ implies that

$$\sqrt{k} \tilde{\epsilon} \leq \alpha \tilde{\lambda}_{\min} \leq \alpha \left(\sum_{i=1}^k \tilde{\lambda}_i^2 \theta_i^2 \right)^{1/2}.$$

Moreover, since $\tilde{\lambda}_1 |\theta_1| > 0$ (by the first claim) and $\tilde{\lambda}_1 |\theta_1| = \max_{i \in [k]} \tilde{\lambda}_i |\theta_i|$, it follows that

$$\tilde{\lambda}_1 |\theta_1| \geq \tilde{\lambda}_{\min} \max_{i \in [k]} |\theta_i| \geq \frac{\tilde{\lambda}_{\min}}{\sqrt{k}}, \tag{34}$$

so we also have

$$\tilde{\epsilon} \leq \alpha \tilde{\lambda}_1 |\theta_1|.$$

Applying these bounds on $\tilde{\epsilon}$ to (33), we obtain

$$\tilde{\lambda}_1 |\theta_1| \geq \frac{1 - \alpha}{(1 + \alpha)^2} \left(\sum_{i=1}^k \tilde{\lambda}_i^2 \theta_i^2 \right)^{1/2} \geq \frac{1 - \alpha}{(1 + \alpha)^2} \left(\tilde{\lambda}_1^2 \theta_1^2 + \max_{i \neq 1} \tilde{\lambda}_i^2 \theta_i^2 \right)^{1/2}$$

which in turn implies (for $\alpha \in (0, 1/20)$)

$$\max_{i \neq 1} \tilde{\lambda}_i^2 \theta_i^2 \leq \left(\frac{(1 + \alpha)^4}{(1 - \alpha)^2} - 1 \right) \cdot \tilde{\lambda}_1^2 \theta_1^2 \leq 7\alpha \cdot \tilde{\lambda}_1^2 \theta_1^2.$$

Therefore $\max_{i \neq 1} \tilde{\lambda}_i |\theta_i| \leq \sqrt{7\alpha} \cdot \tilde{\lambda}_1 |\theta_1|$, proving the second claim.

Now we prove the final claim. This is done by (i) showing that θ has a large projection onto u_1 , (ii) using an SVD perturbation argument to show that $\pm u_1$ is close to v_1 , and (iii) concluding that θ has a large projection onto v_1 .

We begin by showing that $(u_1^\top \theta)^2$ is large. Observe that from (32), we have $(1 + \alpha)^2 \phi_1^2 \geq \|\tilde{M}\|_F^2 \geq \phi_1^2 + \max_{i \neq 1} \phi_i^2$, and therefore

$$\max_{i \neq 1} |\phi_i| \leq \sqrt{2\alpha + \alpha^2} \cdot |\phi_1|.$$

Moreover, by the triangle inequality,

$$\begin{aligned} |\theta^\top \tilde{M}\theta| &\leq \sum_{i=1}^k |\phi_i| (u_i^\top \theta)^2 \\ &\leq |\phi_1| (u_1^\top \theta)^2 + \max_{i \neq 1} |\phi_i| (1 - (u_1^\top \theta)^2) \\ &= (u_1^\top \theta)^2 (|\phi_1| - \max_{i \neq 1} |\phi_i|) + \max_{i \neq 1} |\phi_i|. \end{aligned}$$

Using (32) once more, we have $|\theta^\top \tilde{M}\theta| \geq \|\tilde{M}\|_F / (1 + \alpha) \geq |\phi_1| / (1 + \alpha)$, so

$$(u_1^\top \theta)^2 \geq \frac{\frac{1}{1+\alpha} - \max_{i \neq 1} \frac{|\phi_i|}{|\phi_1|}}{1 - \max_{i \neq 1} \frac{|\phi_i|}{|\phi_1|}} = 1 - \frac{\alpha}{(1 + \alpha) \left(1 - \max_{i \neq 1} \frac{|\phi_i|}{|\phi_1|}\right)} \leq 1 - \frac{\alpha}{(1 + \alpha)(1 - \sqrt{2\alpha + \alpha^2})}.$$

Now we show that $(u_1^\top v_1)^2$ is also large. By the second claim, the assumption on $\tilde{\epsilon}$, and (34),

$$\tilde{\lambda}_1 |\theta_1| - \max_{i \neq 1} \tilde{\lambda}_i |\theta_i| > (1 - \sqrt{7\alpha}) \cdot \tilde{\lambda}_1 |\theta_1| \geq (1 - \sqrt{7\alpha}) \cdot \tilde{\lambda}_{\min} / \sqrt{k}.$$

Combining this with Weyl's theorem gives

$$|\phi_1| - \max_{i \neq 1} \tilde{\lambda}_i |\theta_i| \geq \tilde{\lambda}_1 |\theta_1| - \tilde{\epsilon} - \max_{i \neq 1} \tilde{\lambda}_i |\theta_i| \geq (1 - (\alpha + \sqrt{7\alpha})) \cdot \tilde{\lambda}_{\min} / \sqrt{k},$$

so we may apply Wedin's theorem to obtain

$$(u_1^\top v_1)^2 \geq 1 - \left(\frac{\|\tilde{E}(I, I, \theta)\|}{|\phi_1| - \max_{i \neq 1} \tilde{\lambda}_i |\theta_i|} \right)^2 \geq 1 - \left(\frac{\alpha}{1 - (\alpha + \sqrt{7\alpha})} \right)^2.$$

It remains to show that $\theta_1 = v_1^\top \theta$ is large. Indeed, by the triangle inequality, Cauchy-Schwarz, and the above inequalities on $(u_1^\top v_1)^2$ and $(u_1^\top \theta)^2$,

$$\begin{aligned} |v_1^\top \theta| &= \left| \sum_{i=1}^k (u_i^\top v_1) (u_i^\top \theta) \right| \\ &\geq |u_1^\top v_1| |u_1^\top \theta| - \sum_{i=2}^k |u_i^\top v_1| |u_i^\top \theta| \\ &\geq |u_1^\top v_1| |u_1^\top \theta| - \left(\sum_{i=2}^k (u_i^\top v_1)^2 \right)^{1/2} \left(\sum_{i=2}^k (u_i^\top \theta)^2 \right)^{1/2} \\ &= |u_1^\top v_1| |u_1^\top \theta| - \left((1 - (u_1^\top v_1)^2) (1 - (u_1^\top \theta)^2) \right)^{1/2} \\ &\geq \left(\left(1 - \frac{\alpha}{(1 + \alpha)(1 - \sqrt{2\alpha + \alpha^2})} \right) \left(1 - \left(\frac{\alpha}{1 - (\alpha + \sqrt{7\alpha})} \right)^2 \right) \right)^{1/2} \\ &\quad - \left(\frac{\alpha}{(1 + \alpha)(1 - \sqrt{2\alpha + \alpha^2})} \cdot \left(\frac{\alpha}{1 - (\alpha + \sqrt{7\alpha})} \right)^2 \right)^{1/2} \\ &\geq 1 - 2\alpha \end{aligned}$$

for $\alpha \in (0, 1/20)$. Moreover, by assumption we have $\tilde{T}(\theta, \theta, \theta) \geq 0$, and

$$\begin{aligned}
 \tilde{T}(\theta, \theta, \theta) &= \sum_{i=1}^k \tilde{\lambda}_i \theta_i^3 + \tilde{E}(\theta, \theta, \theta) \\
 &= \tilde{\lambda}_1 \theta_1^3 + \sum_{i=2}^k \tilde{\lambda}_i \theta_i^3 + \tilde{E}(\theta, \theta, \theta) \\
 &\leq \tilde{\lambda}_1 \theta_1^3 + \max_{i \neq 1} \tilde{\lambda}_i |\theta_i| \sum_{i=2}^k \theta_i^2 + \tilde{\epsilon} \\
 &\leq \tilde{\lambda}_1 \theta_1^3 + \sqrt{7\alpha} \tilde{\lambda}_1 |\theta_1| (1 - \theta_1^2) + \tilde{\epsilon} \quad (\text{by the second claim}) \\
 &\leq \tilde{\lambda}_1 |\theta_1|^3 \left(\text{sign}(\theta_1) + \frac{\sqrt{7\alpha}}{(1 - 2\alpha)^2} - \sqrt{7\alpha} + \frac{\alpha}{(1 - 2\alpha)^3} \right) \quad (\text{since } |\theta_1| \geq 1 - 2\alpha) \\
 &< \tilde{\lambda}_1 |\theta_1|^3 (\text{sign}(\theta_1) + 1)
 \end{aligned}$$

so $\text{sign}(\theta_1) > -1$, meaning $\theta_1 > 0$. Therefore $\theta_1 = |\theta_1| \geq 1 - 2\alpha$. This proves the final claim. \blacksquare

Lemma C.2 Fix $\alpha, \beta \in (0, 1)$. Assume $\tilde{\lambda}_{i^*} = \max_{i \in [k]} \tilde{\lambda}_i$ and

$$\tilde{\epsilon} \leq \min \left\{ \frac{\alpha}{5\sqrt{k} + 7}, \frac{1 - \beta}{7} \right\} \cdot \tilde{\lambda}_{i^*}, \quad \tilde{\epsilon}_F \leq \sqrt{\ell} \cdot \frac{1 - \beta}{2\beta} \cdot \tilde{\lambda}_{i^*}.$$

To the conclusion of Lemma B.4, it can be added that the stopping condition (31) is satisfied by $\theta = \theta_t$.

Proof Without loss of generality, assume $i^* = 1$. By the triangle inequality and Cauchy-Schwarz,

$$\begin{aligned}
 \|\tilde{T}(I, I, \theta_t)\|_F &\leq \tilde{\lambda}_1 |\theta_{1,t}| + \sum_{i \neq 1} \lambda_i |\theta_{i,t}| + \|\tilde{E}(I, I, \theta_t)\|_F \leq \tilde{\lambda}_1 |\theta_{1,t}| + \tilde{\lambda}_1 \sqrt{k} \left(\sum_{i \neq 1} \theta_{i,t}^2 \right)^{1/2} + \sqrt{k} \tilde{\epsilon} \\
 &\leq \tilde{\lambda}_1 |\theta_{1,t}| + \frac{3\sqrt{k} \tilde{\epsilon}}{p} + \sqrt{k} \tilde{\epsilon}.
 \end{aligned}$$

where the last step uses the fact that $\theta_{1,t}^2 \geq 1 - (3\tilde{\epsilon}/(p\tilde{\lambda}_1))^2$. Moreover,

$$\tilde{T}(\theta_t, \theta_t, \theta_t) \geq \tilde{\lambda}_1 - \left(27 \left(\frac{\tilde{\epsilon}}{p\tilde{\lambda}_1} \right)^2 + 2 \right) \frac{\tilde{\epsilon}}{p}.$$

Combining these two inequalities with the assumption on $\tilde{\epsilon}$ implies that

$$\tilde{T}(\theta_t, \theta_t, \theta_t) \geq \frac{1}{1 + \alpha} \|\tilde{T}(I, I, \theta_t)\|_F.$$

Using the definition of the tensor Frobenius norm, we have

$$\frac{1}{\sqrt{\ell}} \|\tilde{T}\|_F \leq \frac{1}{\sqrt{\ell}} \left\| \sum_{i=1}^k \tilde{\lambda}_i v_i^{\otimes 3} \right\|_F + \frac{1}{\sqrt{\ell}} \|\tilde{E}\|_F = \tilde{\lambda}_{\text{avg}} + \frac{1}{\sqrt{\ell}} \|\tilde{E}\|_F \leq \tilde{\lambda}_{\text{avg}} + \frac{1}{\sqrt{\ell}} \tilde{\epsilon}_F.$$

Combining this with the above inequality implies

$$\tilde{T}(I, I, \theta_t) \geq \frac{\beta}{\sqrt{\ell}} \|\tilde{T}\|_F.$$

Therefore the stopping condition (31) is satisfied. ■

C.2 Sketch of Analysis of Algorithm 2

The analysis of Algorithm 2 is very similar to the proof of Theorem 5.1 for Algorithm 1, so here we just sketch the essential differences.

First, the guarantee afforded to Algorithm 2 is somewhat different than Theorem 5.1. Specifically, it is of the following form: (i) under appropriate conditions, upon termination, the algorithm returns an accurate decomposition, and (ii) the algorithm terminates after $\text{poly}(k)$ random restarts with high probability.

The conditions on ϵ and N are the same (but for possibly different universal constants C_1, C_2). In Lemma C.1 and Lemma C.2, there is reference to a condition on the Frobenius norm of E , but we may use the inequality $\|E\|_F \leq k\|E\| \leq k\epsilon$ so that the condition is subsumed by the ϵ condition.

Now we outline the differences relative to the proof of Theorem 5.1. The basic structure of the induction argument is the same. In the induction step, we argue that (i) if the stopping condition is satisfied, then by Lemma C.1 (with $\alpha = 0.05$ and $\beta = 1/2$), we have a vector θ_N such that, for some $j^* \geq i$,

1. $\lambda_{\pi(j^*)} \geq \lambda_{\pi(j_{\max})}/(4\sqrt{k})$;
2. θ_N is $(1/4)$ -separated relative to $\pi(j^*)$;
3. $\theta_{\pi(j^*), N} \geq 4/5$;

and (ii) the stopping condition is satisfied within $\text{poly}(k)$ random restarts (via Lemma B.1 and Lemma C.2) with high probability. We now invoke Lemma B.4 to argue that executing another N power iterations starting from θ_N gives a vector $\hat{\theta}$ that satisfies

$$\|\hat{\theta} - v_{\pi(j^*)}\| \leq \frac{8\epsilon}{\lambda_{\pi(j^*)}}, \quad |\hat{\lambda} - \lambda_{\pi(j^*)}| \leq 5\epsilon.$$

The main difference here, relative to the proof of Theorem 5.1, is that we use $\kappa := 4\sqrt{k}$ (rather than $\kappa = O(1)$), but this ultimately leads to the same guarantee after taking into consideration the condition $\epsilon \leq C_1 \lambda_{\min}/k$. The remainder of the analysis is essentially the same as the proof of Theorem 5.1.

Appendix D. Simultaneous Diagonalization for Tensor Decomposition

As discussed in the introduction, another standard approach to certain tensor decomposition problems is to simultaneously diagonalize a collection of similar matrices obtained from the given tensor. We now examine this approach in the context of our latent variable models, where

$$M_2 = \sum_{i=1}^k w_i \mu_i \otimes \mu_i$$

$$M_3 = \sum_{i=1}^k w_i \mu_i \otimes \mu_i \otimes \mu_i.$$

Let $V := [\mu_1 | \mu_2 | \dots | \mu_k]$ and $D(\eta) := \text{diag}(\mu_1^\top \eta, \mu_2^\top \eta, \dots, \mu_k^\top \eta)$, so

$$M_2 = V \text{diag}(w_1, w_2, \dots, w_k) V^\top$$

$$M_3(I, I, \eta) = V \text{diag}(w_1, w_2, \dots, w_k) D(\eta) V^\top$$

Thus, the problem of determining the μ_i can be cast as a simultaneous diagonalization problem: find a matrix X such that $X^\top M_2 X$ and $X^\top M_3(I, I, \eta) X$ (for all η) are diagonal. It is easy to see that if the μ_i are linearly independent, then the solution $X^\top = V^\dagger$ is unique up to permutation and rescaling of the columns.

With exact moments, a simple approach is as follows. Assume for simplicity that $d = k$, and define

$$M(\eta) := M_3(I, I, \eta) M_2^{-1} = V D(\eta) V^{-1}.$$

Observe that if the diagonal entries of $D(\eta)$ are distinct, then the eigenvectors of $M(\eta)$ are the columns of V (up to permutation and scaling). This criterion is satisfied almost surely when η is chosen randomly from a continuous distribution over \mathbb{R}^k .

The above technique (or some variant thereof) was previously used to give the efficient learnability results, where the computational and sample complexity bounds were polynomial in relevant parameters of the problem, including the rank parameter k (Mossel and Roch, 2006; Anandkumar et al., 2012c,a; Hsu and Kakade, 2013). However, the specific polynomial dependence on k was rather large due to the need for the diagonal entries of $D(\eta)$ to be well-separated. This is because with finite samples, $M(\eta)$ is only known up to some perturbation, and thus the sample complexity bound depends inversely in (some polynomial of) the separation of the diagonal entries of $D(\eta)$. With η drawn uniformly at random from the unit sphere in \mathbb{R}^k , the separation was only guaranteed to be roughly $1/k^{2.5}$ (Anandkumar et al., 2012c) (while this may be a loose estimate, the instability is observed in practice). In contrast, using the tensor power method to approximately recover V (and hence the model parameters μ_i and w_i) requires only a mild, *lower-order* dependence on k .

It should be noted, however, that the use of a single random choice of η is quite restrictive, and it is easy to see that a simultaneous diagonalization of $M(\eta)$ for several choices of η can be beneficial. While the uniqueness of the eigendecomposition of $M(\eta)$ is only guaranteed when the diagonal entries of $D(\eta)$ are distinct, the *simultaneous diagonalization* of $M(\eta^{(1)}), M(\eta^{(2)}), \dots, M(\eta^{(m)})$ for vectors $\eta^{(1)}, \eta^{(2)}, \dots, \eta^{(m)}$ is unique as long as the

columns of

$$\begin{bmatrix} \mu_1^\top \eta^{(1)} & \mu_2^\top \eta^{(1)} & \cdots & \mu_k^\top \eta^{(1)} \\ \mu_1^\top \eta^{(2)} & \mu_2^\top \eta^{(2)} & \cdots & \mu_k^\top \eta^{(2)} \\ \vdots & \vdots & \ddots & \vdots \\ \mu_1^\top \eta^{(m)} & \mu_2^\top \eta^{(m)} & \cdots & \mu_k^\top \eta^{(m)} \end{bmatrix}$$

are distinct (*i.e.*, for each pair of column indices i, j , there exists a row index r such that the (r, i) -th and (r, j) -th entries are distinct). This is a much weaker requirement for uniqueness, and therefore may translate to an improved perturbation analysis. In fact, using the techniques discussed in Section 4.3, we may even reduce the problem to an orthogonal simultaneous diagonalization, which may be easier to obtain. Furthermore, a number of robust numerical methods for (approximately) simultaneously diagonalizing collections of matrices have been proposed and used successfully in the literature (*e.g.*, Bunse-Gerstner et al., 1993; Cardoso and Souloumiac, 1993; Cardoso, 1994; Cardoso and Comon, 1996; Ziehe et al., 2004). Another alternative and a more stable approach compared to full diagonalization is a Schur-like method which finds a unitary matrix U which simultaneously triangularizes the respective matrices (Corless et al., 1997). It is an interesting open question whether these techniques can yield similar improved learnability results and also enjoy the attractive computational properties of the tensor power method.

References

- D. Achlioptas and F. McSherry. On spectral learning of mixtures of distributions. In *Eighteenth Annual Conference on Learning Theory*, pages 458–469, 2005.
- E. S. Allman, C. Matias, and J. A. Rhodes. Identifiability of parameters in latent structure models with many observed variables. *The Annals of Statistics*, 37(6A):3099–3132, 2009.
- A. Anandkumar, D. P. Foster, D. Hsu, S. M. Kakade, and Y.-K. Liu. A spectral algorithm for latent Dirichlet allocation. In *Advances in Neural Information Processing Systems 25*, 2012a.
- A. Anandkumar, D. Hsu, F. Huang, and S. M. Kakade. Learning mixtures of tree graphical models. In *Advances in Neural Information Processing Systems 25*, 2012b.
- A. Anandkumar, D. Hsu, and S. M. Kakade. A method of moments for mixture models and hidden Markov models. In *Twenty-Fifth Annual Conference on Learning Theory*, volume 23, pages 33.1–33.34, 2012c.
- J. Anderson, M. Belkin, N. Goyal, L. Rademacher, and J. Voss. The more, the merrier: the blessing of dimensionality for learning large Gaussian mixtures. In *Twenty-Seventh Annual Conference on Learning Theory*, 2014.
- S. Arora and R. Kannan. Learning mixtures of separated nonspherical Gaussians. *The Annals of Applied Probability*, 15(1A):69–92, 2005.
- S. Arora, R. Ge, and A. Moitra. Learning topic models — going beyond SVD. In *Fifty-Third IEEE Annual Symposium on Foundations of Computer Science*, pages 1–10, 2012a.

- S. Arora, R. Ge, A. Moitra, and S. Sachdeva. Provable ICA with unknown Gaussian noise, and implications for Gaussian mixtures and autoencoders. In *Advances in Neural Information Processing Systems 25*, 2012b.
- T. Austin. On exchangeable random variables and the statistics of large graphs and hypergraphs. *Probab. Survey*, 5:80–145, 2008.
- R. Bailly. Quadratic weighted automata: Spectral algorithm and likelihood maximization. *Journal of Machine Learning Research*, 2011.
- B. Balle and M. Mohri. Spectral learning of general weighted automata via constrained matrix completion. In *Advances in Neural Information Processing Systems 25*, 2012.
- B. Balle, A. Quattoni, and X. Carreras. Local loss optimization in operator models: A new insight into spectral learning. In *Twenty-Ninth International Conference on Machine Learning*, 2012.
- M. Belkin and K. Sinha. Polynomial learning of distribution families. In *Fifty-First Annual IEEE Symposium on Foundations of Computer Science*, pages 103–112, 2010.
- A. Bhaskara, M. Charikar, A. Moitra, and A. Vijayaraghavan. Smoothed analysis of tensor decompositions. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, 2014.
- B. Boots, S. M. Siddiqi, and G. J. Gordon. Closing the learning-planning loop with predictive state representations. In *Proceedings of the Robotics Science and Systems Conference*, 2010.
- S. C. Brubaker and S. Vempala. Isotropic PCA and affine-invariant clustering. In *Forty-Ninth Annual IEEE Symposium on Foundations of Computer Science*, 2008.
- A. Bunse-Gerstner, R. Byers, and V. Mehrmann. Numerical methods for simultaneous diagonalization. *SIAM Journal on Matrix Analysis and Applications*, 14(4):927–949, 1993.
- J.-F. Cardoso. Super-symmetric decomposition of the fourth-order cumulant tensor. blind identification of more sources than sensors. In *Acoustics, Speech, and Signal Processing, 1991. ICASSP-91., 1991 International Conference on*, pages 3109–3112. IEEE, 1991.
- J.-F. Cardoso. Perturbation of joint diagonalizers. Technical Report 94D027, Signal Department, Télécom Paris, 1994.
- J.-F. Cardoso and P. Comon. Independent component analysis, a survey of some algebraic methods. In *IEEE International Symposium on Circuits and Systems*, pages 93–96, 1996.
- J.-F. Cardoso and A. Souloumiac. Blind beamforming for non Gaussian signals. *IEE Proceedings-F*, 140(6):362–370, 1993.
- D. Cartwright and B. Sturmfels. The number of eigenvalues of a tensor. *Linear Algebra Appl.*, 438(2):942–952, 2013.

- R. B. Cattell. Parallel proportional profiles and other principles for determining the choice of factors by rotation. *Psychometrika*, 9(4):267–283, 1944.
- J. T. Chang. Full reconstruction of Markov models on evolutionary trees: Identifiability and consistency. *Mathematical Biosciences*, 137:51–73, 1996.
- K. Chaudhuri and S. Rao. Learning mixtures of product distributions using correlations and independence. In *Twenty-First Annual Conference on Learning Theory*, pages 9–20, 2008.
- S. B. Cohen, K. Stratos, M. Collins, D. P. Foster, and L. Ungar. Spectral learning of latent-variable PCFGs. In *Fiftieth Annual Meeting of the Association for Computational Linguistics*, 2012.
- P. Comon. Independent component analysis, a new concept? *Signal Processing*, 36(3):287–314, 1994.
- P. Comon and C. Jutten. *Handbook of Blind Source Separation: Independent Component Analysis and Applications*. Academic Press. Elsevier, 2010.
- P. Comon, G. Golub, L.-H. Lim, and B. Mourrain. Symmetric tensors and symmetric tensor rank. *SIAM Journal on Matrix Analysis Appl.*, 30(3):1254–1279, 2008.
- R. M. Corless, P. M. Gianni, and B. M. Trager. A reordered Schur factorization method for zero-dimensional polynomial systems with multiple roots. In *Proceedings of the 1997 International Symposium on Symbolic and Algebraic Computation*, pages 133–140. ACM, 1997.
- S. Dasgupta. Learning mixtures of Gaussians. In *Fortieth Annual IEEE Symposium on Foundations of Computer Science*, pages 634–644, 1999.
- S. Dasgupta and L. Schulman. A probabilistic analysis of EM for mixtures of separated, spherical Gaussians. *Journal of Machine Learning Research*, 8(Feb):203–226, 2007.
- L. De Lathauwer, J. Castaing, and J.-F. Cardoso. Fourth-order cumulant-based blind identification of underdetermined mixtures. *Signal Processing, IEEE Transactions on*, 55(6):2965–2973, 2007.
- N. Delfosse and P. Loubaton. Adaptive blind separation of independent sources: a deflation approach. *Signal processing*, 45(1):59–83, 1995.
- A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum-likelihood from incomplete data via the EM algorithm. *J. Royal Statist. Soc. Ser. B*, 39:1–38, 1977.
- P. Dhillon, J. Rodu, M. Collins, D. P. Foster, and L. Ungar. Spectral dependency parsing with latent variables. In *Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, 2012.
- M. Drton, B. Sturmfels, and S. Sullivant. Algebraic factor analysis: tetrads, pentads and beyond. *Probability Theory and Related Fields*, 138(3):463–493, 2007.

- A. T. Erdogan. On the convergence of ICA algorithms with symmetric orthogonalization. *IEEE Transactions on Signal Processing*, 57:2209–2221, 2009.
- A. M. Frieze, M. Jerrum, and R. Kannan. Learning linear transformations. In *Thirty-Seventh Annual Symposium on Foundations of Computer Science*, pages 359–368, 1996.
- G. H. Golub and C. F. van Loan. *Matrix Computations*. Johns Hopkins University Press, 1996.
- N. Halko, P.-G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2), 2011.
- R. Harshman. Foundations of the PARAFAC procedure: model and conditions for an ‘explanatory’ multi-mode factor analysis. Technical report, UCLA Working Papers in Phonetics, 1970.
- C. J. Hillar and L.-H. Lim. Most tensor problems are NP-hard. *J. ACM*, 60(6):45:1–45:39, November 2013. ISSN 0004-5411. doi: 10.1145/2512329.
- F. L. Hitchcock. The expression of a tensor or a polyadic as a sum of products. *Journal of Mathematics and Physics*, 6:164–189, 1927a.
- F. L. Hitchcock. Multiple invariants and generalized rank of a p-way matrix or tensor. *Journal of Mathematics and Physics*, 7:39–79, 1927b.
- D. Hsu and S. M. Kakade. Learning mixtures of spherical Gaussians: moment methods and spectral decompositions. In *Fourth Innovations in Theoretical Computer Science*, 2013.
- D. Hsu, S. M. Kakade, and P. Liang. Identifiability and unmixing of latent parse trees. In *Advances in Neural Information Processing Systems 25*, 2012a.
- D. Hsu, S. M. Kakade, and T. Zhang. A spectral algorithm for learning hidden Markov models. *Journal of Computer and System Sciences*, 78(5):1460–1480, 2012b.
- A. Hyvarinen. Fast and robust fixed-point algorithms for independent component analysis. *Neural Networks, IEEE Transactions on*, 10(3):626–634, 1999.
- A. Hyvärinen and E. Oja. Independent component analysis: algorithms and applications. *Neural Networks*, 13(4–5):411–430, 2000.
- H. Jaeger. Observable operator models for discrete stochastic time series. *Neural Comput.*, 12(6), 2000.
- A. T. Kalai, A. Moitra, and G. Valiant. Efficiently learning mixtures of two Gaussians. In *Forty-second ACM Symposium on Theory of Computing*, pages 553–562, 2010.
- R. Kannan, H. Salmasian, and S. Vempala. The spectral method for general mixture models. *SIAM Journal on Computing*, 38(3):1141–1156, 2008.

- E. Kofidis and P. A. Regalia. On the best rank-1 approximation of higher-order supersymmetric tensors. *SIAM Journal on Matrix Analysis and Applications*, 23(3):863–884, 2002.
- T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455, 2009.
- T. G. Kolda and J. R. Mayo. Shifted power method for computing tensor eigenpairs. *SIAM Journal on Matrix Analysis and Applications*, 32(4):1095–1124, October 2011.
- J. B. Kruskal. Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics. *Linear Algebra and Appl.*, 18(2):95–138, 1977.
- L. D. Lathauwer, B. D. Moor, and J. Vandewalle. 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, 2000.
- L. Le Cam. *Asymptotic Methods in Statistical Decision Theory*. Springer, 1986.
- L.-H. Lim. Singular values and eigenvalues of tensors: a variational approach. *Proceedings of the IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, 1:129–132, 2005.
- M. Littman, R. Sutton, and S. Singh. Predictive representations of state. In *Advances in Neural Information Processing Systems 14*, pages 1555–1561, 2001.
- F. M. Luque, A. Quattoni, B. Balle, and X. Carreras. Spectral learning for non-deterministic dependency parsing. In *Conference of the European Chapter of the Association for Computational Linguistics*, 2012.
- J. B. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 281–297. University of California Press, 1967.
- P. McCullagh. *Tensor Methods in Statistics*. Chapman and Hall, 1987.
- A. Moitra and G. Valiant. Settling the polynomial learnability of mixtures of Gaussians. In *Fifty-First Annual IEEE Symposium on Foundations of Computer Science*, pages 93–102, 2010.
- E. Mossel and S. Roch. Learning nonsingular phylogenies and hidden Markov models. *Annals of Applied Probability*, 16(2):583–614, 2006.
- P. Q. Nguyen and O. Regev. Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. *Journal of Cryptology*, 22(2):139–160, 2009.
- J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 1999.
- P. V. Overschee and B. D. Moor. *Subspace Identification of Linear Systems*. Kluwer Academic Publishers, 1996.

- L. Pachter and B. Sturmfels. *Algebraic Statistics for Computational Biology*, volume 13. Cambridge University Press, 2005.
- A. Parikh, L. Song, and E. P. Xing. A spectral algorithm for latent tree graphical models. In *Twenty-Eighth International Conference on Machine Learning*, 2011.
- K. Pearson. Contributions to the mathematical theory of evolution. *Philosophical Transactions of the Royal Society, London, A.*, page 71, 1894.
- L. Qi. Eigenvalues of a real supersymmetric tensor. *Journal of Symbolic Computation*, 40(6):1302–1324, 2005.
- R. A. Redner and H. F. Walker. Mixture densities, maximum likelihood and the EM algorithm. *SIAM Review*, 26(2):195–239, 1984.
- P. A. Regalia and E. Kofidis. Monotonic convergence of fixed-point algorithms for ICA. *IEEE Transactions on Neural Networks*, 14:943–949, 2003.
- S. Roch. A short proof that phylogenetic tree reconstruction by maximum likelihood is hard. *IEEE/ACM Trans. Comput. Biol. Bioinformatics*, 3(1), 2006.
- J. Rodu, D. P. Foster, W. Wu, and L. H. Ungar. Using regression for spectral estimation of HMMs. In *Statistical Language and Speech Processing*, pages 212–223, 2013.
- M. P. Schützenberger. On the definition of a family of automata. *Inf. Control*, 4:245–270, 1961.
- S. M. Siddiqi, B. Boots, and G. J. Gordon. Reduced-rank hidden Markov models. In *Thirteenth International Conference on Artificial Intelligence and Statistics*, 2010.
- D. A. Spielman and S. H. Teng. Smoothed analysis: An attempt to explain the behavior of algorithms in practice. *Communications of the ACM*, pages 76–84, 2009.
- A. Stegeman and P. Comon. Subtracting a best rank-1 approximation may increase tensor rank. *Linear Algebra and Its Applications*, 433:1276–1300, 2010.
- B. Sturmfels and P. Zwiernik. Binary cumulant varieties. *Ann. Comb.*, (17):229–250, 2013.
- S. Vempala and G. Wang. A spectral algorithm for learning mixtures models. *Journal of Computer and System Sciences*, 68(4):841–860, 2004.
- P. Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111, 1972.
- T. Zhang and G. Golub. Rank-one approximation to high order tensors. *SIAM Journal on Matrix Analysis and Applications*, 23:534–550, 2001.
- A. Ziehe, P. Laskov, G. Nolte, and K. R. Müller. A fast algorithm for joint diagonalization with non-orthogonal transformations and its application to blind source separation. *Journal of Machine Learning Research*, 5:777–800, 2004.
- J. Zou, D. Hsu, D. Parkes, and R. P. Adams. Contrastive learning using spectral methods. In *Advances in Neural Information Processing Systems 26*, 2013.