
Tensor Decompositions via Two-Mode Higher-Order SVD (HOSVD)

Miaoyan Wang
University of Pennsylvania

Yun S. Song
University of California, Berkeley

Abstract

Tensor decompositions have rich applications in statistics and machine learning, and developing efficient, accurate algorithms for the problem has received much attention recently. Here, we present a new method built on Kruskal’s uniqueness theorem to decompose symmetric, nearly orthogonally decomposable tensors. Unlike the classical higher-order singular value decomposition which unfolds a tensor along a single mode, we consider unfoldings along two modes and use rank-1 constraints to characterize the underlying components. This tensor decomposition method provably handles a greater level of noise compared to previous methods and achieves a high estimation accuracy. Numerical results demonstrate that our algorithm is robust to various noise distributions and that it performs especially favorably as the order increases.

1 Introduction

Tensor decompositions have recently drawn increased attention in statistical and machine learning applications. For example, tensor decompositions provide powerful tools to estimate parameters in various latent variable models via the method of moments [1–3] and lead to the development of efficient denoising techniques in independent component analysis [5, 6, 16].

Decomposition of higher-order tensors is fraught with challenges, since in general most computational problems regarding tensors are NP-hard [7]. However, the tensors considered in the aforementioned applications possess special structures that facilitate computation. A class of such tensors which we study in this paper

are nearly orthogonal decomposable tensors of the following form:

$$\tilde{\mathcal{T}} = \sum_{i=1}^r \lambda_i \mathbf{u}_i^{\otimes k} + \text{noise}, \quad (1)$$

where $\{\mathbf{u}_i\}_{i=1}^r$ is a set of orthonormal vectors in \mathbb{R}^d , λ_i are non-zero scalars in \mathbb{R} , and the noise is assumed to be symmetric and small (bounded by a positive constant ε under some natural measurement). Our goal is to estimate the factor pairs $\{(\mathbf{u}_i, \lambda_i) \in \mathbb{R}^d \times \mathbb{R}\}_{i=1}^r$ from the noisy observation $\tilde{\mathcal{T}}$. Often the number r of factors is also unknown and it needs to be determined from $\tilde{\mathcal{T}}$; we consider this general case at the end of the paper. In contrast to existing works that have mostly focused on order-3 tensors (i.e., $k = 3$), we here present a general framework applicable to an arbitrary order $k \geq 3$. Unless stated otherwise, we use the term “tensor” to mean a tensor of order ≥ 3 .

Current tensor decomposition methods vary in their objective function, tolerance of the noise level ε , target range of r , and runtime. The robust tensor power method (TPM) [2], which computes the fixed point of a certain iterative function, is a natural generalization of the matrix power method to the $k = 3$ case. TPM is easy to implement, but its noise tolerance is restricted, i.e., $\varepsilon \lesssim O(d^{-1})$ with random initialization. A more recent work [14] proposes using a hybrid of a square unfolding and power iterations to improve the noise tolerance. This technique is designed for single-spike models (i.e., $r = 1$) with Gaussian noise, and the sensitivity of the algorithm when some λ_i are close to each other has not been explored. Another decomposition method, orthogonal joint diagonalization (OJD) [11], seeks all factors simultaneously by joint diagonalization of a set of matrices from random contractions. This approach reaches the full potential in the full rank setting; however, applying OJD to low-rank tensors requires further modifications.

In this paper we propose a new algorithm that reduces the orthogonal decomposition of tensors to that of matrices. Specifically, we unfold the tensor along two modes and demand the left singular vectors of the resulting matrix to be close to a *Kronecker square* [8]

(i.e., $\mathbf{x}^{\otimes 2}$ for some vector \mathbf{x}). By viewing the length- d^2 vectors as elements of $\mathbb{R}^{d \times d}$, this is equivalent to imposing a “nearly-rank-1” constraint in the two-mode singular space. We show that the two-mode matrix singular value decomposition (SVD), in conjunction with the nearly-rank-1 constraint, provides an accurate approximation of $\{\mathbf{u}_i^{\otimes 2}\}$. We then estimate the underlying factors $\{\mathbf{u}_i\}$ using the dominant eigenvectors of those nearly rank-1 matrices.

Our work is inspired by and built on Kruskal’s uniqueness theorem [4, 10, 15]: when ε vanishes in (1), the set of $\{\mathbf{u}_i\}$ must be uniquely determined even in the case of degenerate λ_i s. This is in sharp contrast to the matrix case, where imposing proper constraints on the matrix SVD is needed to avoid the ambiguity caused by degenerate singular values. In fact, the notion of Kronecker-squared singular vector, denoted $\text{Vec}(\mathbf{u}_i^{\otimes 2})$, is a necessary and sufficient condition for \mathbf{u}_i to be a desired component in orthogonal tensor decomposition. Our method ensures the exact recovery in the noiseless case. Moreover, it achieves very good accuracy at a higher noise level $O(d^{-(k-2)/2})$ than TPM’s noise tolerance of $O(d^{-(k-1)/2})$ [14].

The remainder of the paper is organized as follows. In Section 2, we introduce some tensor notation and algebra. Section 3 describes the key observation that characterizes the components of orthogonal decomposable tensors. Our main algorithm and perturbation analysis are summarized in Section 4. In Section 5 we demonstrate the accuracy and efficiency of our method through simulations.

2 Preliminaries

We use $\mathcal{T} = \llbracket \mathcal{T}_{i_1 \dots i_k} \rrbracket \in \mathbb{R}^{d_1 \times \dots \times d_k}$ to denote an order- k , (d_1, \dots, d_k) -dimensional tensor with entries $\mathcal{T}_{i_1 \dots i_k} \in \mathbb{R}$ where $1 \leq i_n \leq d_n$ for $n = 1, \dots, k$. A tensor $\mathcal{T} \in \mathbb{R}^{d_1 \times \dots \times d_k}$ is called symmetric if $d_1 = \dots = d_k$ and $\mathcal{T}_{i_1 \dots i_k} = \mathcal{T}_{\pi(i_1) \dots \pi(i_k)}$ for all permutations π of $\{1, \dots, k\}$. We use $[n]$ to denote the n -set $\{1, \dots, n\}$ for $n \in \mathbb{N}_+$ and $\mathbf{S}^{d-1} = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_2 = 1\}$ to denote the $(d-1)$ -dimensional unit sphere.

A symmetric tensor \mathcal{T} can be viewed as a multilinear map [12]. For any $\mathbf{x} = (x_1, \dots, x_d)^T \in \mathbb{R}^d$, define

$$\mathcal{T}(\mathbf{x}, \dots, \mathbf{x}) = \sum_{i_1=1}^d \dots \sum_{i_k=1}^d \mathcal{T}_{i_1 \dots i_k} x_{i_1} \dots x_{i_k}.$$

The *spectral norm*, or the l^2 -norm, of \mathcal{T} is defined as

$$\|\mathcal{T}\|_\sigma = \sup_{\mathbf{x} \in \mathbf{S}^{d-1}} |\mathcal{T}(\mathbf{x}, \dots, \mathbf{x})|.$$

Note that when $k = 2$, the tensor l^2 -norm reduces to the classical matrix l^2 -norm. The Frobenius norm of

the tensor \mathcal{T} is defined as

$$\|\mathcal{T}\|_F = \left(\sum_{i_1=1}^d \dots \sum_{i_k=1}^d \mathcal{T}_{i_1 \dots i_k}^2 \right)^{1/2}.$$

Given a symmetric order- k tensor $\mathcal{T} = \llbracket \mathcal{T}_{i_1 \dots i_k} \rrbracket \in \mathbb{R}^{d \times \dots \times d}$, we can map the indices from k -tuples (i_1, \dots, i_k) to ℓ -tuples, where $1 \leq \ell \leq k$, in various ways to reshape the tensor into lower-order objects. This operation is called *tensor unfolding*. For instance, an order-3 symmetric tensor \mathcal{T} can be unfolded into a d -by- d^2 matrix $\mathcal{T}_{(1)(23)}$ by defining the (a, b) -element of the matrix as $[\mathcal{T}_{(1)(23)}]_{ab} = \mathcal{T}_{i_1 i_2 i_3}$, where $a = i_1$ and $b = i_2 + (i_3 - 1)d$. The notation $(j_1 \dots j_m)$ appearing in the subscript of an unfolded tensor denotes that the modes j_1, \dots, j_m are combined into a single mode. The following three unfoldings will be of particular interest to us:

- (a) *One-mode unfolding* $\mathcal{T}_{(1)(2\dots k)}$, which reshapes \mathcal{T} into a $d \times d^{k-1}$ matrix.
- (b) *Two-mode unfolding* $\mathcal{T}_{(12)(3\dots k)}$, which reshapes \mathcal{T} into a $d^2 \times d^{k-2}$ matrix.
- (c) *Order-3 unfolding* $\mathcal{T}_{(1)(2)(3\dots k)}$, which reshapes \mathcal{T} into a $d \times d \times d^{k-2}$ cube.

We now introduce the notion of *two-mode singular space* which is central to our methods. Let $\mathcal{T} \in \mathbb{R}^{d \times \dots \times d}$ be an order- k symmetric tensor and $\mathcal{T}_{(12)(3\dots k)}$ be its two-mode unfolding. We use

$$\mathcal{T}_{(12)(3\dots k)} = \sum_i \mu_i \mathbf{a}_i \mathbf{b}_i^T$$

to denote the *two-mode higher-order SVD* (HOSVD); that is, $\mu_i \geq 0$ is the i th largest singular value of the matrix $\mathcal{T}_{(12)(3\dots k)}$, and $\mathbf{a}_i \in \mathbb{R}^{d^2}$ (respectively, $\mathbf{b}_i \in \mathbb{R}^{d^{k-2}}$) is the i th left (respectively, right) singular vector corresponding to μ_i .

Definition 2.1 (Two-Mode Singular Space). The s -truncated two-mode singular space of \mathcal{T} is defined by

$$\mathcal{L}\mathcal{S}^{(s)} = \text{Span}\{\mathbf{a}_1, \dots, \mathbf{a}_s\},$$

where \mathbf{a}_i is the i th left singular vector of $\mathcal{T}_{(12)(3\dots k)}$.

Remark 2.2. Since \mathbb{R}^{d^2} is isomorphic to $\mathbb{R}^{d \times d}$, we also write $\mathcal{L}\mathcal{S}^{(s)} = \text{Span}\{\text{Mat}(\mathbf{a}_1), \dots, \text{Mat}(\mathbf{a}_s)\}$, where $\text{Mat}(\cdot)$ denotes the matricization operation that unstacks a length- d^2 vector into a d -by- d matrix. Conversely, we use $\text{Vec}(\cdot)$ to denote the concatenation of matrix columns into a vector. Throughout the paper, we do not distinguish between these two representations of $\mathcal{L}\mathcal{S}^{(s)}$. It should be clear from the context whether we are viewing the elements of the two-mode singular space as length- d^2 vectors or $d \times d$ matrices.

Remark 2.3. If μ_s is strictly larger than μ_{s+1} , the s -truncated space $\mathcal{LS}^{(s)}$ is uniquely determined.

Tensors arising in applications (such as parameter estimation for latent variable models, independent component analysis, etc) frequently possess special structures. In this paper, we consider a class of tensors that are symmetric and (nearly) orthogonal decomposable. A tensor \mathcal{T} is called *symmetric and orthogonally decomposable* (SOD, [13]) if \mathcal{T} can be written as $\mathcal{T} = \sum_{i=1}^r \lambda_i \mathbf{u}_i^{\otimes k}$, where $\{\mathbf{u}_i\}_{i \in [r]}$ are orthonormal vectors in \mathbb{R}^d , $\{\lambda_i\}_{i \in [r]}$ are non-zero scalars in \mathbb{R} , and r is a positive integer. In this case, each \mathbf{u}_i is called a *factor* and r the *rank* of the SOD tensor \mathcal{T} .

Unlike in the matrix case where every symmetric matrix has an eigen-decomposition, symmetric tensors are not necessarily SOD. We are particularly interested in symmetric tensors within a small neighborhood of SOD tensors; i.e., $\tilde{\mathcal{T}} = \sum_{i=1}^r \lambda_i \mathbf{u}_i^{\otimes k} + \mathcal{E}$ where $\mathcal{E} \in \mathbb{R}^{d \times \dots \times d}$ is a symmetric but otherwise arbitrary tensor with $\|\mathcal{E}\|_\sigma \leq \varepsilon$. Here \mathcal{E} represents the noise tensor arising from, for example, sampling error, model misspecification, finite sample size, etc. We refer to this class of tensors as *nearly-SOD tensors* [13]. For simplicity we will assume $\lambda_i > 0$ for all $i \in [r]$ hereinafter. The case where λ_i s have arbitrary signs will be discussed at the end of Section 4.

3 Tensor vs. Matrix Decompositions

In this section, we restrict our attention to SOD tensors and relate their decompositions to matrix decompositions. The special structure of SOD tensors ($\mathcal{T} = \sum_i \lambda_i \mathbf{u}_i^{\otimes k}$, where \mathbf{u}_i are orthonormal) allows us to illustrate the main idea with a clean and simple proof. We characterize the underlying factors $\{\mathbf{u}_i\}$ of \mathcal{T} using rank-1 matrices in the two-mode singular space of \mathcal{T} . As we shall see later, this perspective plays a key role in our algorithm.

3.1 Characterization of Robust Eigenvectors

Tensor decompositions possess an interesting uniqueness property not present in matrix decompositions. Consider an SOD tensor, $\mathcal{T} = \sum_{i=1}^r \lambda_i \mathbf{u}_i^{\otimes k}$. Kruskal's theorem guarantees that the set of orthonormal vectors $\{\mathbf{u}_i\}_{i=1}^r$ is unique up to signs even when some λ_i s are degenerate. A rank $r > 1$ matrix, on the other hand, can be decomposed in multiple manners into a sum of outer-product terms in the case of degenerate λ_i s. To make the distinction explicit, we refer to these unique components $\{\mathbf{u}_i\}$ as *robust eigenvectors* [2].

Definition 3.1 (Anandkumar et al [2]). A unit vector $\mathbf{a} \in \mathbb{R}^d$ is called a robust eigenvector of \mathcal{T} if $\mathbf{a} \in \{\mathbf{u}_1, \dots, \mathbf{u}_r\}$.

Decomposing an SOD tensor amounts to finding all its robust eigenvectors. Note that the tensor decomposition $\mathcal{T} = \sum_{i=1}^r \lambda_i \mathbf{u}_i^{\otimes k}$ implies a series of matrix SVDs, such as $\mathcal{T}_{(1)(2\dots k)} = \sum_{i=1}^r \lambda_i \text{Vec}(\mathbf{u}_i) \text{Vec}(\mathbf{u}_i^{\otimes k-1})^T$, $\mathcal{T}_{(12)(3\dots k)} = \sum_{i=1}^r \lambda_i \text{Vec}(\mathbf{u}_i^{\otimes 2}) \text{Vec}(\mathbf{u}_i^{\otimes k-2})^T, \dots$, and so on. This suggests a way to use matrix SVDs to recover $\{\mathbf{u}_i\}$. Despite the seeming simplicity, however, matrix SVDs may lead to spurious solutions because they are not guaranteed to be unique. For example, in the case of $\lambda_1 = \lambda_2$, $\mathbf{u}_1 + \mathbf{u}_2$ is a (unnormalized) left singular vector of $\mathcal{T}_{(1)(2\dots k)}$. Clearly, $\mathbf{u}_1 + \mathbf{u}_2$ is non-parallel to any robust eigenvector of \mathcal{T} . Similarly, $\text{Vec}(\mathbf{u}_1^{\otimes 2} + \mathbf{u}_2^{\otimes 2})$ is a (unnormalized) left singular vector of $\mathcal{T}_{(12)(3\dots k)}$ without being parallel to any $\text{Vec}(\mathbf{u}_i^{\otimes 2})$.

Fortunately, such spurious solutions can be ruled out by enforcing a certain Kronecker-product constraint on matrix singular vectors. Specifically, we have the following characterization of robust eigenvectors:

Theorem 3.2. *A unit vector $\mathbf{a} \in \mathbb{R}^d$ is a robust eigenvector of \mathcal{T} if and only if $\text{Vec}(\mathbf{a}^{\otimes 2})$ is a left singular vector of $\mathcal{T}_{(12)(3\dots k)}$ corresponding to a non-zero singular value.*

A proof of Theorem 3.2 is provided in Appendix A. Theorem 3.2 is closely related to the uniqueness property and it provides a criterion for a unit vector \mathbf{a} to be a robust eigenvector of \mathcal{T} . An earlier work [2] has shown that robust eigenvectors can be characterized using local maximizers of the objective function $\mathbf{a} \mapsto \mathcal{T}(\mathbf{a}, \dots, \mathbf{a}) / \|\mathbf{a}\|_2^k$. Our result provides a new perspective that does not require checking the gradient and/or the Hessian of the objective function.

3.2 Exact Recovery for SOD Tensors

For an SOD tensor $\mathcal{T} = \sum_{i=1}^r \lambda_i \mathbf{u}_i^{\otimes k}$, we define an r -dimensional linear space $\mathcal{LS}_0 \stackrel{\text{def}}{=} \text{Span}\{\mathbf{u}_i^{\otimes 2} : i \in [r]\}$. The elements of \mathcal{LS}_0 can be viewed as either length- d^2 vectors or d -by- d matrices. The space \mathcal{LS}_0 is exactly recovered by $\mathcal{LS}^{(r)}$, the r -truncated two-mode left singular space of \mathcal{T} , where r equals the rank of $\mathcal{T}_{(12)(3\dots k)}$. Similar to Theorem 3.2, imposing rank-1 constraints ensures the desired solutions in \mathcal{LS}_0 :

Proposition 3.3. *Every rank-1 matrix in \mathcal{LS}_0 is (up to a scalar) the Kronecker square of some robust eigenvector of \mathcal{T} .*

Proposition 3.3 implies that rank-1 matrices in \mathcal{LS}_0 are sufficient to find $\{\mathbf{u}_i\}$. Note that a matrix $\mathbf{M} \in \mathbb{R}^{d \times d}$ being rank-1 is equivalent to $\|\mathbf{M}\|_\sigma / \|\mathbf{M}\|_F = 1$, since $\|\mathbf{M}\|_\sigma \leq \|\mathbf{M}\|_F \leq \sqrt{\text{rank } \mathbf{M}} \|\mathbf{M}\|_\sigma$. Hence, Proposition 3.3 immediately suggests the following algorithm:

$$\begin{aligned} & \underset{\mathbf{M} \in \mathbb{R}^{d \times d}}{\text{maximize}} \|\mathbf{M}\|_\sigma, \\ & \text{subject to } \mathbf{M} \in \mathcal{LS}_0 \text{ and } \|\mathbf{M}\|_F = 1. \end{aligned} \quad (2)$$

Theorem 3.4 (Exact Recovery in the Noiseless Case). *The optimization problem (2) has exactly r pairs of local maximizers $\{\pm \mathbf{M}_i^* : i \in [r]\}$. Furthermore, they satisfy the following three properties:*

- (A1) $\|\mathbf{M}_i^*\|_\sigma = 1$ for all $i \in [r]$.
- (A2) $|\langle \text{Vec}(\mathbf{M}_i^*), \text{Vec}(\mathbf{M}_j^*) \rangle| = \delta_{ij}$ for all $i, j \in [r]$, where $\langle \cdot, \cdot \rangle$ denotes the inner product.
- (A3) There exists a permutation π on $[r]$ such that $\mathbf{M}_i^* = \pm \mathbf{u}_{\pi(i)}^{\otimes 2}$ for all $i \in [r]$.

A proof is provided in Appendix B.2. Since $\|\mathbf{M}\|_\sigma \leq \|\mathbf{M}\|_F$ in general and \mathbf{M} is constrained to satisfy $\|\mathbf{M}\|_F = 1$, property (A1) implies that every local maximizer \mathbf{M}_i^* is a global maximizer. Therefore, any algorithm that ensures local optimality is able to recover exactly the set of matrices $\{\mathbf{u}_i^{\otimes 2}\}$. As a by-product, the number r of factors is recovered by the number of linearly independent solutions $\{\mathbf{M}_i^*\}$. In addition, property (A3) indicates $\{\mathbf{u}_i\}$ can be extracted from the dominant eigenvectors of the optimal solutions $\{\mathbf{M}_i^*\}$.

4 Two-Mode HOSVD via Nearly Rank-1 Matrix Pursuit

In most statistical and machine learning applications, the observed SOD tensors are perturbed by noise. In this section we extend the results in Section 3 to nearly-SOD tensors, which take the form

$$\tilde{\mathcal{T}} = \sum_{i=1}^r \lambda_i \mathbf{u}_i^{\otimes k} + \mathcal{E}, \quad (3)$$

where the first part on the right hand side is SOD and $\mathcal{E} \in \mathbb{R}^{d \times \dots \times d}$ is a symmetric but otherwise arbitrary noise tensor satisfying $\|\mathcal{E}\|_\sigma \leq \varepsilon$. Our goal is to estimate the underlying pairs $\{(\mathbf{u}_i, \lambda_i) \in \mathbb{R}^d \times \mathbb{R}\}_{i \in [r]}$ from $\tilde{\mathcal{T}}$. We note that the number of factors $r \in \mathbb{N}_+$ is typically unknown and has to be determined empirically. As in most previous works on tensor decomposition, we assume that r is known in Sections 4.1 and 4.2. In Section 4.3, we describe a rule of thumb for choosing r .

4.1 Algorithm

An orthogonal decomposition of a noisy tensor $\tilde{\mathcal{T}}$ does not necessarily exist, and therefore finding rank-1 matrices in the two-mode singular space of $\tilde{\mathcal{T}}$ may not be possible. Nevertheless, the formulation shown in (2) suggests a practical way to approximate $\{\mathbf{u}_i\}$. Specifically, we seek to find a “nearly” rank-1 matrix $\widehat{\mathbf{M}}$ via the following optimization:

$$\begin{aligned} & \underset{\mathbf{M} \in \mathbb{R}^{d \times d}}{\text{maximize}} \|\mathbf{M}\|_\sigma, \\ & \text{subject to } \mathbf{M} \in \mathcal{L}\mathcal{S}^{(r)} \text{ and } \|\mathbf{M}\|_F = 1, \end{aligned} \quad (4)$$

and take the dominant eigenvector $\hat{\mathbf{u}} = \arg \max_{\mathbf{x} \in \mathbf{S}^{d-1}} |\mathbf{x}^T \widehat{\mathbf{M}} \mathbf{x}|$ as an estimator of \mathbf{u}_i . Recall that $\mathcal{L}\mathcal{S}^{(r)} = \text{Span}\{\mathbf{a}_1, \dots, \mathbf{a}_r\}$, where \mathbf{a}_i is the i th left singular vector of the two-mode unfolding $\tilde{\mathcal{T}}_{(12)(3\dots k)}$. In principle, the r -truncated two-mode singular space $\mathcal{L}\mathcal{S}^{(r)}$ might not be uniquely determined; for example, this occurs when the r th and $(r+1)$ th singular values of $\tilde{\mathcal{T}}_{(12)(3\dots k)}$ are equal. We shall make the following assumption on the model (3) to ensure the uniqueness of $\mathcal{L}\mathcal{S}^{(r)}$:

Assumption 4.1 (Signal-to-Noise Ratio). In the notation of model (3), assume $\varepsilon \leq \lambda_{\min} / [c_0 d^{(k-2)/2}]$, where $\lambda_{\min} = \min_{i \in [r]} \lambda_i$ and $c_0 > 2$ is a constant that does not depend on d .

We refer to $\lambda_{\min}/\varepsilon$ as the signal-to-noise ratio (SNR) of $\tilde{\mathcal{T}}$, because λ_{\min} represents the minimum eigenvalue of the signal tensor $\sum_{i=1}^r \lambda_i \mathbf{u}_i^{\otimes k}$, whereas ε represents the maximum eigenvalue of the noise tensor \mathcal{E} . In Appendix C.2, we prove the following result:

Proposition 4.2. *Under Assumption 4.1, the two-mode singular space $\mathcal{L}\mathcal{S}^{(r)}$ is uniquely determined.*

Assumption 4.1 implies that the threshold SNR scales as $O(d^{(k-2)/2})$. Although it may appear stringent, this assumption is prevalingly made by most computationally tractable algorithms [14]. In fact, for order-3 tensors, our scaling $O(\sqrt{d})$ is less stringent than the $O(d)$ required by the power iteration approach [2].

Algorithm 1 outlines our method, which we divide into four parts for ease of reference. The algorithm consists of r successive iterations. At each iteration, we search for a nearly rank-1 matrix $\widehat{\mathbf{M}}$ in the space $\mathcal{L}\mathcal{S}^{(r)}$ and compute the top eigenvector of $\widehat{\mathbf{M}}$ (or some refined version of $\widehat{\mathbf{M}}$). We then deflate the space $\mathcal{L}\mathcal{S}^{(r)}$ and repeat the procedure until a full decomposition is obtained. In the noiseless case, Algorithm 1 guarantees the exact recovery of $\{\mathbf{u}_i\}$. Further study demonstrates that this approach also accurately approximates $\{\mathbf{u}_i\}$ in the presence of noise satisfying Assumption 4.1. Here we discuss a few critical algorithmic aspects of the four steps in Algorithm 1 and defer theoretical analyses to the next section.

Two-Mode HOSVD. We perform the r -truncated SVD on the two-mode unfolded matrix $\tilde{\mathcal{T}}_{(12)(3\dots k)} \in \mathbb{R}^{d^2 \times d^{k-2}}$. This step typically requires $O(d^k r)$ floating-point operations (flops). When $\tilde{\mathcal{T}}_{(12)(3\dots k)}$ is a “fat” matrix, the computational cost can be reduced by finding the r -truncated eigen decomposition of the d^2 -by- d^2 matrix $[\tilde{\mathcal{T}}_{(12)(3\dots k)}][\tilde{\mathcal{T}}_{(12)(3\dots k)}]^T$. The latter approach involves $O(d^k)$ flops for matrix multiplication and $O(d^4 r)$ flops for eigen decomposition. Hence, the total cost of this component is $O(d^k r)$ for $k = 3, 4$ and

Algorithm 1 Two-mode HOSVD

	Input: Noisy tensor $\tilde{\mathcal{T}}$ where $\tilde{\mathcal{T}} = \sum_{i=1}^r \lambda_i \mathbf{u}_i^{\otimes k} + \mathcal{E}$, number of factors r .
	Output: r pairs of estimators $(\hat{\mathbf{u}}_i, \hat{\lambda}_i)$.
Two-Mode HOSVD	$\left\{ \begin{array}{l} 1: \text{Reshape the tensor } \tilde{\mathcal{T}} \text{ into a } d^2\text{-by-}d^{k-2} \text{ matrix } \tilde{\mathcal{T}}_{(12)(3\dots k)}; \\ 2: \text{Find the top } r \text{ left singular vectors of } \tilde{\mathcal{T}}_{(12)(3\dots k)}, \text{ denoted } \{\mathbf{a}_1, \dots, \mathbf{a}_r\}; \\ 3: \text{Initialize } \mathcal{LS}^{(r)} = \text{Span}\{\mathbf{a}_i : i \in [r]\}; \\ 4: \text{for } i=1 \text{ to } r \text{ do} \end{array} \right.$
Nearly Rank-1 Matrix	$\left\{ \begin{array}{l} 5: \text{Solve } \widehat{\mathbf{M}}_i = \arg \max_{\mathbf{M} \in \mathcal{LS}^{(r)}, \ \mathbf{M}\ _F=1} \ \mathbf{M}\ _\sigma \text{ and } \hat{\mathbf{u}}_i = \arg \max_{\mathbf{u} \in \mathbf{S}^{d-1}} \mathbf{u}^T \widehat{\mathbf{M}}_i \mathbf{u} ; \end{array} \right.$
Post-Processing	$\left\{ \begin{array}{l} 6: \text{Update } \widehat{\mathbf{M}}_i \leftarrow \tilde{\mathcal{T}}_{(1)(2)(3\dots k)}(\mathbf{I}, \mathbf{I}, \text{Vec}(\hat{\mathbf{u}}_i^{\otimes(k-2)})) \text{ and } \hat{\mathbf{u}}_i \leftarrow \arg \max_{\mathbf{u} \in \mathbf{S}^{d-1}} \mathbf{u}^T \widehat{\mathbf{M}}_i \mathbf{u} ; \\ 7: \text{Return } (\hat{\mathbf{u}}_i, \hat{\lambda}_i) \leftarrow (\hat{\mathbf{u}}_i, \hat{\mathcal{T}}(\hat{\mathbf{u}}_i, \dots, \hat{\mathbf{u}}_i)); \end{array} \right.$
Deflation	$\left\{ \begin{array}{l} 8: \text{Set } \mathcal{LS}^{(r)} \leftarrow \mathcal{LS}^{(r)} \cap [\text{Vec}(\hat{\mathbf{u}}_i^{\otimes 2})]^\perp; \\ 9: \text{end for} \end{array} \right.$

$O(d^k)$ for $k \geq 5$.

Nearly Rank-1 Matrix. This step of the algorithm requires optimization, and the computational cost depends on the choice of the optimization subroutine. For $r = 1$ (i.e., single-spike model), no actual optimization is needed since $\widehat{\mathbf{M}}_1 \stackrel{\text{def}}{=} \arg \max_{\mathbf{M} \in \mathcal{LS}^{(1)}, \|\mathbf{M}\|_F=1} \|\mathbf{M}\|_\sigma$ is simply $\text{Mat}(\mathbf{a}_1)$. For $r \geq 2$, there exist various subroutines for optimizing $\|\mathbf{M}\|_\sigma$. Here we choose to use coordinate ascent with initialization $\widehat{\mathbf{M}}_i^{(0)} = \text{Mat}(\mathbf{a}_i)$ at the i th iteration. Briefly, we introduce two decision variables $\mathbf{x} \in \mathbf{S}^{d-1}$ and $\boldsymbol{\alpha} \in \mathbf{S}^{r-1}$, and rewrite (4) as $\max_{\mathbf{x} \in \mathbf{S}^{d-1}, \boldsymbol{\alpha} \in \mathbf{S}^{r-1}} H(\mathbf{x}, \boldsymbol{\alpha})$, where $H(\mathbf{x}, \boldsymbol{\alpha}) \stackrel{\text{def}}{=} \mathbf{x}^T [\alpha_1 \text{Mat}(\mathbf{a}_1) + \dots + \alpha_r \text{Mat}(\mathbf{a}_r)] \mathbf{x}$. Since both $\max_{\mathbf{x} \in \mathbf{S}^{d-1}} H(\mathbf{x}, \cdot)$ and $\max_{\boldsymbol{\alpha} \in \mathbf{S}^{r-1}} H(\cdot, \boldsymbol{\alpha})$ have closed-form solutions, we update \mathbf{x} and $\boldsymbol{\alpha}$ alternatively until convergence. The computational cost is $O(d^2) + O(d^2 r)$ for each update. Upon finding $\widehat{\mathbf{M}}_i$, we set $\hat{\mathbf{u}}_i = \arg \max_{\mathbf{u} \in \mathbf{S}^{d-1}} |\mathbf{u}^T \widehat{\mathbf{M}}_i \mathbf{u}|$.

Although the above procedure finds only a local maximum, in practice we found that the converged points generally have nearly optimal objective values. In fact, local optimality itself is not a severe issue in our context, because we seek a total of top r maximizers in $\mathcal{LS}^{(r)}$ but the order in which they are found is unimportant. When the noise is small enough, at some iterations there could be multiple close-to-optimal choices of $\widehat{\mathbf{M}}_i$, with negligible difference between their objective values. In that case, any of these choices performs equally well in estimating \mathbf{u}_i and it is thus of little interest to identify a true global optimum.

Post-Processing. In this step we update $\widehat{\mathbf{M}}_i \leftarrow \tilde{\mathcal{T}}_{(1)(2)(3\dots k)}(\mathbf{I}, \mathbf{I}, \text{Vec}(\hat{\mathbf{u}}_i^{\otimes(k-2)}))$ and take its top eigenvector as an estimator of \mathbf{u}_i . Here, \mathbf{I} denotes the d -by-

d identity matrix. This step requires $O(d^k)$ flops. The updated matrix $\widehat{\mathbf{M}}_i$ is arguably still close to rank-1 and provides a better estimation of \mathbf{u}_i . This post-processing is intended to make the algorithm more robust to higher-order noise. Intuitively, merging modes $(3, \dots, k)$ in the earlier stage of the algorithm amplifies the noise (measured by spectral norm) because of the corresponding loss in multilinear structure [17]. The updated matrix $\widehat{\mathbf{M}}_i = \tilde{\mathcal{T}}(\mathbf{I}, \mathbf{I}, \hat{\mathbf{u}}_i, \dots, \hat{\mathbf{u}}_i)$ helps to alleviate such an effect because it essentially decouples the combined modes. The theoretical analysis in the next section confirms this intuition.

Deflation. Our deflation strategy is to update the search space using the orthogonal complement of $\text{Vec}(\hat{\mathbf{u}}_i^{\otimes 2})$ in $\mathcal{LS}^{(r)}$, where $\hat{\mathbf{u}}_i$ is the estimator in the current iteration. This can be done by setting the basis vector $\mathbf{a}_j \leftarrow \mathbf{a}_j - \langle \text{Vec}(\hat{\mathbf{u}}_i^{\otimes 2}), \mathbf{a}_j \rangle \text{Vec}(\hat{\mathbf{u}}_i^{\otimes 2})$ for all $j \in [r]$, followed by normalizing $\{\mathbf{a}_j\}_{j \in [r]}$. The complexity involved is $O(d^2 r)$. An alternative deflation strategy is to update the tensor $\tilde{\mathcal{T}} \leftarrow \tilde{\mathcal{T}} - \hat{\lambda}_i \hat{\mathbf{u}}_i^{\otimes k}$ and then repeat lines 1–7 of Algorithm 1. A careful analysis reveals that both approaches effectively control the accumulated error. We adopt the former strategy because it allows us to avoid recalculating of the two-mode HOSVD.

The total cost of our algorithm is $O(d^k)$ per iteration. This is comparable to or even lower than that of competing methods. For order $k = 3$, TPM has complexity $O(d^3 M)$ per iteration (where M is the number of restarts), and OJD has complexity $O(d^3 L)$ per iteration (where L is the number of matrices). Furthermore, we have provided theoretical results for an arbitrary order k , while neither the TPM nor the OJD method provided a theoretical analysis for higher-order cases.

4.2 Theoretical Analysis

In this section we theoretically analyze the accuracy of the estimators $\{(\widehat{\mathbf{u}}_i, \widehat{\lambda}_i) \in \mathbb{R}^d \times \mathbb{R}\}_{i \in [r]}$ produced by our algorithm. We first consider extracting one pair $(\widehat{\mathbf{u}}_1, \widehat{\lambda}_1)$ by Algorithm 1 and provide the corresponding error bounds. We then show that the accumulated error caused by deflation has only a higher-order effect and is thus negligible. The stability for the full decomposition is rendered through a combination of these two results.

We develop a series of algebraic techniques based on multilinear functional analysis, Weyl's and Wedin's perturbation theorems [18, 19], some of which may be of independent interest. Here we only state the key theorems and lemmas. All proofs are deferred to Appendix C.3–C.4.

Following [14], we use a loss function that allows the sign-flip error.

Definition 4.3. Let $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$ be two unit vectors, and define

$$\text{Loss}(\mathbf{a}, \mathbf{b}) = \min(\|\mathbf{a} - \mathbf{b}\|_2, \|\mathbf{a} + \mathbf{b}\|_2).$$

If a, b are two scalars in \mathbb{R} , we define $\text{Loss}(a, b) = \min(|a - b|, |a + b|)$.

The following lemma describes the deviation of the perturbed space $\mathcal{LS}^{(r)}$ from the true space \mathcal{LS}_0 under small noise. The result can be viewed as an analogue of Wedin's perturbation in the context of two-mode HOSVD.

Theorem 4.4 (Perturbation of \mathcal{LS}_0). *Suppose $c_0 \geq 10$ in Assumption 4.1. Then,*

$$\max_{\substack{\mathbf{M} \in \mathcal{LS}^{(r)}, \\ \|\mathbf{M}\|_F = 1}} \min_{\mathbf{M}^* \in \mathcal{LS}_0} \|\mathbf{M} - \mathbf{M}^*\|_\sigma \leq \frac{d^{\frac{k-3}{2}} \varepsilon}{\lambda_{\min}} + o(\varepsilon).$$

We choose to use the matrix spectral norm to measure the closeness of $\mathcal{LS}^{(r)}$ and \mathcal{LS}_0 . Although the Frobenius norm may seem an easier choice, it poorly reflects the multilinear nature of the tensor and leads to suboptimal results. Note that the dimension factor vanishes when $k = 3$ (i.e., order-3 tensors).

The following lemma guarantees the existence of a nearly rank-1 matrix in the perturbed space $\mathcal{LS}^{(r)}$.

Lemma 4.5 (Nearly Rank-1 Matrix). *Under Assumption 4.1, we have*

$$\max_{\substack{\mathbf{M} \in \mathcal{LS}^{(r)}, \\ \|\mathbf{M}\|_F = 1}} \|\mathbf{M}\|_\sigma \geq 1 - \frac{d^{\frac{k-2}{2}} \varepsilon}{\lambda_{\min}} + o(\varepsilon).$$

In fact, we can show that there exist at least r linearly independent matrices (here we view matrices as elements of \mathbb{R}^{d^2}) that satisfy the above inequality. More generally, at i th iteration of Algorithm 1, there are at least $r - i + 1$ linearly independent nearly rank-1 matrices in the search space.

The following lemma provides an error bound for the estimator $\widehat{\mathbf{u}}_i$ extracted from a nearly rank-1 matrix in $\mathcal{LS}^{(r)}$.

Lemma 4.6 (Error Bound for the First Component). *Let $\widehat{\mathbf{M}}_1$ and $\widehat{\mathbf{u}}_1$ respectively be the matrix and the vector defined in line 5 of Algorithm 1. Suppose $c_0 \geq 10$ in Assumption 4.1. Then, there exists $i^* \in [r]$ such that*

$$\text{Loss}(\widehat{\mathbf{u}}_1, \mathbf{u}_{i^*}) \leq \frac{d^{\frac{k-3}{2}} \varepsilon}{\lambda_{\min}} + o(\varepsilon).$$

We note that in Lemma 4.6, the perturbation bound for the estimator does not depend on the eigen-gaps (i.e., $|\lambda_i - \lambda_j|$). This feature fundamentally distinguishes tensor decompositions from matrix decompositions. Also, note that an amplification of the error (by a polynomial factor of d) is observed in the bound. The following lemma suggests a simple post-processing step which improves the error bound.

Lemma 4.7 (Post-Processing). *Let $\widehat{\mathbf{M}}_1$ and $\widehat{\mathbf{u}}_1$ respectively be the matrix and the vector defined in line 6 of Algorithm 4.1. Further, define $\widehat{\lambda}_1 = \widetilde{\mathcal{T}}(\widehat{\mathbf{u}}_1, \dots, \widehat{\mathbf{u}}_1)$ and let $i^* \in [r]$ be the index that appears in Lemma 4.6. Suppose $c_0 \geq \max\{10, \frac{3(k-2)}{2} + \frac{6\lambda_{\max}}{\lambda_{\min}}\}$ in Assumption 4.1. Then,*

$$\text{Loss}(\widehat{\mathbf{u}}_1, \mathbf{u}_{i^*}) \leq \frac{\varepsilon}{\lambda_{i^*}} + o(\varepsilon), \quad \text{Loss}(\widehat{\lambda}_1, \lambda_{i^*}) \leq \varepsilon + o(\varepsilon).$$

Compared to Lemma 4.6, the leading terms in the above error bounds no longer scale with d . Finally, a careful analysis shows that the estimation error does not amplify through deflation:

Lemma 4.8 (Deflation of the Singular Space). *Suppose $c_0 \geq \max\{10, \frac{3(k-2)}{2} + \frac{6\lambda_{\max}}{\lambda_{\min}}\}$ in Assumption 4.1, and for a fixed subset $X \subset [r]$, suppose there exists a set $\{\widehat{\mathbf{u}}_i\}_{i \in X}$ of unit vectors satisfying*

$$\text{Loss}(\widehat{\mathbf{u}}_i, \mathbf{u}_{\pi(i)}) \leq \frac{2\varepsilon}{\lambda_{\pi(i)}} + o(\varepsilon), \quad \text{for all } i \in X,$$

where π is a permutation on $[r]$. Then,

$$\max_{\substack{\mathbf{M} \in \mathcal{LS}^{(r)}(X), \\ \|\mathbf{M}\|_F = 1}} \min_{\mathbf{M}^* \in \mathcal{LS}_0(X)} \|\mathbf{M} - \mathbf{M}^*\|_\sigma \leq \frac{2d^{\frac{k-3}{2}} \varepsilon}{\lambda_{\min}} + o(\varepsilon),$$

where $\mathcal{LS}^{(r)}(X)$ and $\mathcal{LS}_0(X)$ are residual spaces defined as $\mathcal{LS}^{(r)}(X) \stackrel{\text{def}}{=} \mathcal{LS}^{(r)} \cap \text{Span}\{\widehat{\mathbf{u}}_i^{\otimes 2} : i \in X\}^\perp$ and $\mathcal{LS}_0(X) \stackrel{\text{def}}{=} \text{Span}\{\mathbf{u}_{\pi(i)}^{\otimes 2} : i \in [r] \setminus X\}$.

The above result demonstrates that the error due to deflation is bounded by a factor that does not depend on the iteration number. This implies that our proposed deflation strategy is numerically stable when extracting the subsequent components. Similarly, if we define the residual tensor, $\tilde{\mathcal{T}}(X) \stackrel{\text{def}}{=} \tilde{\mathcal{T}} - \sum_{i \in X} \hat{\lambda}_i \hat{\mathbf{u}}_i^{\otimes k}$, one can show [13] that $\left\| \tilde{\mathcal{T}}(X) - \sum_{i \in [r] \setminus X} \lambda_i \mathbf{u}_i^{\otimes k} \right\|_{\sigma} \leq c\varepsilon + o(\varepsilon)$, where $c > 0$ is a constant that does not depend on the iteration number. The above observations immediately suggest two possible ways to find the subsequent factors, by deflating either the singular space or the original tensor. Note that the deflation results for tensors are non-trivial since the analogous statement for matrices is not necessarily true [2, 13].

Applying Lemmas 4.5–4.8 successively, we obtain the main result of this section as follows.

Theorem 4.9. *Let $\tilde{\mathcal{T}} = \sum_{i=1}^r \lambda_i \mathbf{u}_i^{\otimes k} + \mathcal{E} \in \mathbb{R}^{d \times \dots \times d}$, where $\{\mathbf{u}_i\}_{i \in [r]}$ is a set of orthonormal vectors in \mathbb{R}^d , $\lambda_i > 0$ for all $i \in [r]$, and \mathcal{E} is a symmetric tensor satisfying $\|\mathcal{E}\|_{\sigma} \leq \varepsilon$. Suppose $\varepsilon \leq \lambda_{\min} / [c_0 d^{(k-2)/2}]$, where $c_0 > 0$ is a sufficiently large constant that does not depend on d . Let $\{(\hat{\mathbf{u}}_i, \hat{\lambda}_i) \in \mathbb{R}^d \times \mathbb{R}\}_{i \in [r]}$ be the output of Algorithm 1 for inputs $\tilde{\mathcal{T}}$ and r . Then, there exists a permutation π on $[r]$ such that for all $i \in [r]$,*

$$\text{Loss}(\hat{\mathbf{u}}_i, \mathbf{u}_{\pi(i)}) \leq \frac{2\varepsilon}{\lambda_{\pi(i)}} + o(\varepsilon),$$

$$\text{Loss}(\hat{\lambda}_i, \lambda_{\pi(i)}) \leq 2\varepsilon + o(\varepsilon),$$

and

$$\left\| \tilde{\mathcal{T}} - \sum_{i=1}^r \hat{\lambda}_i \hat{\mathbf{u}}_i^{\otimes k} \right\|_{\sigma} \leq C\varepsilon + o(\varepsilon),$$

where $C = C(k) > 0$ is a constant that only depends on k .

Remark 4.10. With little modification, our results also apply to the case when not all λ_i are positive. In that case, Assumption 4.1 needs to be modified to $\varepsilon \leq |\lambda|_{\min} / [c_0 d^{(k-2)/2}]$ where $|\lambda|_{\min} = \min_{i \in [r]} |\lambda_i|$.

4.3 Determining the Number of Factors

Determining the rank of a general tensor is intrinsically difficult [7]. For nearly-SOD tensors, however, r can be approximated in various ways. When the number of factors r is unknown, we relax the search space $\mathcal{LS}^{(r)}$ to $\mathcal{LS}^{(n)}$ where $n = \min\{\text{Rank}(\tilde{\mathcal{T}}_{(12)(3\dots k)}), d\}$. Note that we always have $\mathcal{LS}^{(n)} \supset \mathcal{LS}^{(r)}$ under Assumption 4.1. Algorithm 2 describes a rule of thumb for choosing r .

Although Algorithm 2 is heuristic, our simulation results suggest that it provides a good approximation in most cases encountered. The use of $\hat{\lambda}_i^2$ as a guiding criterion is justified by the following observation:

Algorithm 2 Determining the Number of Factors

Input: Noisy tensor $\tilde{\mathcal{T}}$ where $\tilde{\mathcal{T}} = \sum_{i=1}^r \lambda_i \mathbf{u}_i^{\otimes k} + \mathcal{E}$;

Output: The number of factors \hat{r} .

- 1: Run Algorithm 1 with n iterations. Let $(\hat{\mathbf{M}}_i, \hat{\mathbf{u}}_i, \hat{\lambda}_i) \in \mathbb{R}^{d \times d} \times \mathbb{R}^d \times \mathbb{R}$ denote the output from the i th iteration where $i \in [n]$;
 - 2: Choose the subset $\mathcal{S} \subset [n]$ for which $\|\hat{\mathbf{M}}_i\|_{\sigma}$ (i.e., the objective value) is close to 1;
 - 3: Sort $\{\hat{\lambda}_i^2 : i \in \mathcal{S}\}$ in decreasing order.
Pick r by the “elbow” method of the scree plot (similarly as in matrix PCA).
-

$\|\tilde{\mathcal{T}} - \sum_i \hat{\lambda}_i \hat{\mathbf{u}}_i^{\otimes k}\|_F^2 \approx \|\tilde{\mathcal{T}}\|_F^2 - \sum_i \hat{\lambda}_i^2$, provided that $\{\hat{\mathbf{u}}_i\}$ are approximately orthogonal to each other.

5 Numerical Experiments

We assessed the performance of our algorithm by simulating tensors of order $k = 3, 4$, and 5 . We generated nearly-SOD tensors $\tilde{\mathcal{T}} = \sum_{i=1}^r \lambda_i \mathbf{e}_i^{\otimes k} + \mathcal{E} \in \mathbb{R}^{d \times \dots \times d}$, where \mathbf{e}_i is the i th canonical basis vector and λ_i s are i.i.d. draws from $\text{Unif}[0.8, 1.2]$. The noise tensor \mathcal{E} was generated with noise level σ under one of the following three random models:

- (a) (Gaussian) For $i_1 \leq \dots \leq i_k$, draw independent entries $\mathcal{E}_{i_1 \dots i_k} \sim \mathcal{N}(0, \sigma^2)$ uniformly at random.
- (b) (Bernoulli) For $i_1 \leq \dots \leq i_k$, draw independent entries $\mathcal{E}_{i_1 \dots i_k} = \pm \sigma$ with probability $1/2$ for each.
- (c) (Student’s t -distribution) For $i_1 \leq \dots \leq i_k$, draw independent entries $\mathcal{E}_{i_1 \dots i_k} \sim \sigma t(5)$. We use t -distribution with 5 degrees of freedom to mimic heavy-tailed data in real-world applications.

Condition on the i.i.d. entries $\{\mathcal{E}_{i_1 \dots i_k}\}_{i_1 \leq \dots \leq i_k}$, we generated the remaining entries by imposing the symmetry condition $\mathcal{E}_{i_1 \dots i_k} = \mathcal{E}_{\pi(i_1) \dots \pi(i_k)}$, where $(i_1, \dots, i_k) \in [d]^k$ and π is a permutation of $[k]$. We set the dimension $d \in \{25, 50\}$ and the number of factors $r \in \{2, 10, 25\}$, and varied σ . For each scenario, we generated 50 trials and applied Algorithm 1 to obtain the estimators $\{\hat{\mathbf{u}}_i, \hat{\lambda}_i\}_{i \in [r]}$. The estimation accuracy was assessed by the average l^2 loss $\frac{1}{r} \sum_{i=1}^r \text{Loss}(\hat{\mathbf{u}}_i, \mathbf{u}_{\pi(i)})$ across the 50 trials. We compared the two-mode HOSVD (TM-HOSVD) with the orthogonal joint diagonalization (OJD) method [11] and the tensor power method (TPM) [2] with random initialization. The software packages for TPM and OJD were originally designated for $k = 3$, so we extended them to $k \geq 4$. Since both TPM and OJD require users to specify the number of factors, we provided them with the ground-truth r .

Figures 1 and 2, and Supplementary Figure S1 (Ap-

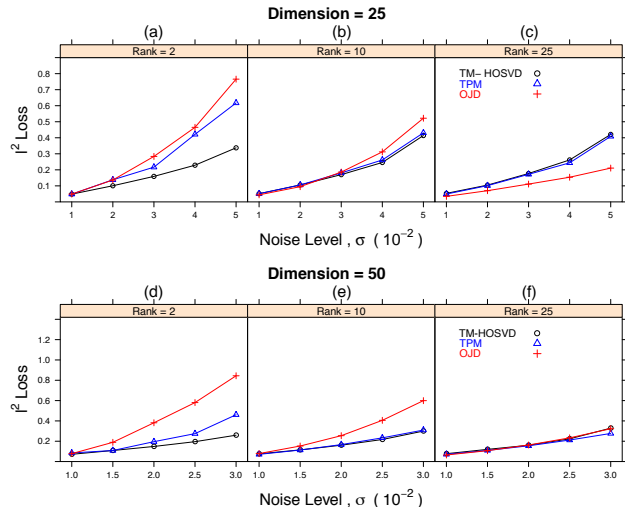


Figure 1: Average l^2 loss for decomposing order-3 nearly SOD tensors with Gaussian noise.

pendix D) demonstrate the robustness of TM-HOSVD to various types of error distributions. For $k = 3$, TM-HOSVD outperformed TPM and OJD when the rank is low to moderate. For $k = 4$ and 5, TM-HOSVD consistently achieved a higher accuracy across the full range of the rank (Figure 3 and Supplementary Figure S2, Appendix D). The empirical runtime is provided in Supplementary Table S1 (Appendix D).

We compared the l^2 loss distribution of the three methods at the highest noise level ($\sigma = 5 \times 10^{-2}$ for order-3 tensors and $\sigma = 1.5 \times 10^{-2}$ for order-4 tensors). The results confirm our earlier conclusion: TM-HOSVD is able to tolerate a greater level of noise. In several scenarios, the loss distribution displayed a bimodal pattern, suggesting a mixture of estimates with good/bad convergence performance. This feature is particularly noticeable in Figure 2a and Figure 3d-f. Further investigation revealed that the poor convergence (e.g., oscillation or local-optimum) occurred less frequently in TM-HOSVD. This is potentially due to the fact that TM-HOSVD starts with the more informative search space $\mathcal{LS}^{(r)}$ with a lower dimension $r \leq d$, whereas TPM/OJD starts with a random direction in a d -dimensional space.

We also noticed that in the full-rank case for $k = 3$, OJD occasionally achieved the best accuracy among all three methods (Figure 1c and Figure 2c). We hypothesize that in the full-rank case the best rank- r approximation of tensors may not be achieved by successive rank-1 approximations. Both TPM and TM-HOSVD methods take a greedy approach, so the first several estimated factors tend to explain the most signal. In contrast, OJD optimizes the objective using all d factors simultaneously. Thus, the set of d estimators




Figure 2: Empirical l^2 loss distribution (plotted on a log scale for better visualization) for decomposing order-3 nearly SOD tensors with Gaussian noise, $d = 25$ and $\sigma = 5 \times 10^{-2}$.

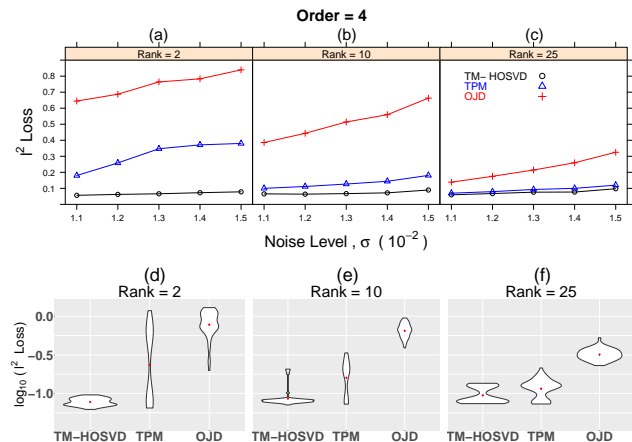


Figure 3: Top Figure: Average l^2 loss for decomposing order-4 SOD tensors with Gaussian noise, $d = 25$; Bottom Figure: Empirical distribution of l^2 loss (on a log scale) for decomposing order-4 nearly SOD tensors with Gaussian noise, $d = 25$ and $\sigma = 1.5 \times 10^{-2}$.

may explain more signal overall while the first several factors explain less. This phenomenon also explains why OJD works better for full-rank tensors than for low- and moderate-rank tensors.

6 Conclusion

We have proposed a new method for tensor decompositions based on the two-mode HOSVD. This method tolerates a higher level of noise and achieves accuracy bounds comparable to that of existing methods while displaying empirically favorable performance. Our approach extends naturally to asymmetric tensors, e.g., by using the matrix SVD in place of the eigendecomposition in lines 5–6 of Algorithm 1. In addition, recent works have shown that some non-orthogonal tensors can be converted to orthogonal tensors by an additional whitening step [2, 9]. Therefore, the two-mode HOSVD is applicable to a broad class of structured tensors. In particular, our proposed algorithm shows stable convergence and exhibits pronounced advantage especially as the order of the tensor increases.

Acknowledgments

This research is supported in part a Math+X Simons Research Grant from the Simons Foundation and a Packard Fellowship for Science and Engineering.

References

- [1] ANANDKUMAR, A., GE, R., HSU, D., AND KAKADE, S. M. A tensor spectral approach to learning mixed membership community models. *Journal of Machine Learning Research* 15, 1 (2014), 2239–2312.
- [2] ANANDKUMAR, A., GE, R., HSU, D., KAKADE, S. M., AND TELGARSKY, M. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research* 15, 1 (2014), 2773–2832.
- [3] ANANDKUMAR, A., LIU, Y.-K., HSU, D. J., FOSTER, D. P., AND KAKADE, S. M. A spectral algorithm for latent dirichlet allocation. In *Advances in Neural Information Processing Systems (NIPS)* (2012), pp. 917–925.
- [4] BHASKARA, A., CHARIKAR, M., AND VIJAYARAGHAVAN, A. Uniqueness of tensor decompositions with applications to polynomial identifiability. In *Proceedings of the 27th Conference on Learning Theory (COLT)* (2014), pp. 742–778.
- [5] COMON, P. Independent component analysis, a new concept? *Signal Processing* 36, 3 (1994), 287–314.
- [6] COMON, P., AND JUTTEN, C. *Handbook of Blind Source Separation: Independent component analysis and applications*. Academic Press, 2010.
- [7] HILLAR, C. J., AND LIM, L.-H. Most tensor problems are NP-hard. *Journal of the ACM (JACM)* 60, 6 (2013), 45.
- [8] KOFIDIS, E., AND REGALIA, P. A. On the best rank-1 approximation of higher-order supersymmetric tensors. *SIAM Journal on Matrix Analysis and Applications* 23, 3 (2002), 863–884.
- [9] KOLDA, T. G. Symmetric orthogonal tensor decomposition is trivial. *arXiv:1503.01375* (2015).
- [10] KRUSKAL, J. B. Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics. *Linear Algebra and its Applications* 18, 2 (1977), 95–138.
- [11] KULESHOV, V., CHAGANTY, A. T., AND LIANG, P. Tensor factorization via matrix factorization. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics (AISTATS)* (2015), pp. 507–516.
- [12] LIM, L.-H. Singular values and eigenvalues of tensors: a variational approach. In *Proceedings of the IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)* (2005), vol. 1, pp. 129–132.
- [13] MU, C., HSU, D., AND GOLDFARB, D. Successive rank-one approximations for nearly orthogonally decomposable symmetric tensors. *SIAM Journal on Matrix Analysis and Applications* 36, 4 (2015), 1638–1659.
- [14] RICHARD, E., AND MONTANARI, A. A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems (NIPS)* (2014), pp. 2897–2905.
- [15] SIDIROPOULOS, N. D., AND BRO, R. On the uniqueness of multilinear decomposition of N-way arrays. *Journal of Chemometrics* 14, 3 (2000), 229–239.
- [16] VEMPALA, S. S., AND XIAO, Y. Max vs min: Tensor decomposition and ICA with nearly linear sample complexity. In *Proceedings of the 28th Conference on Learning Theory (COLT)* (2015), pp. 1710–1723.
- [17] WANG, M., DAO DUC, K., FISCHER, J., AND SONG, Y. S. Operator norm inequalities between tensor unfoldings on the partition lattice. *Linear Algebra and its Applications* 520 (2017), 44–66.
- [18] WEDIN, P.-Å. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics* 12, 1 (1972), 99–111.
- [19] WEYL, H. Inequalities between the two kinds of eigenvalues of a linear transformation. *Proceedings of the National Academy of Sciences* 35, 7 (1949), 408–411.