ON THE ACCURACY OF HOTELLING-TYPE ASYMMETRIC TENSOR DEFLATION: A RANDOM TENSOR ANALYSIS

Mohamed El Amine Seddik[1], Maxime Guillaud[2], Alexis Decurninge[3] and José Henrique de M. Goulart[4]

[1]Technology Innovation Institute, Abu Dhabi, UAE   [2]Inria / CITI Lab, Lyon, France
[3]Huawei Technologies, Boulogne-Billancourt, France   [4]IRIT, Toulouse INP, CNRS, Toulouse, France

ABSTRACT

This work introduces an asymptotic study of Hotelling-type tensor deflation in the presence of noise, in the regime of large tensor dimensions. Specifically, we consider a low-rank asymmetric tensor model of the form $\sum_{i=1}^{r} \beta_i \mathcal{A}_i + \mathcal{W}$ where $\beta_i \geq 0$ and the $\mathcal{A}_i$'s are unit-norm rank-one tensors such that $\left| \langle \mathcal{A}_i, \mathcal{A}_j \rangle \right| \in [0,1]$ for $i \neq j$ and $\mathcal{W}$ is an additive noise term. Assuming that the dominant components are successively estimated from the noisy observation and subsequently subtracted, we leverage recent advances in random tensor theory in the regime of asymptotically large tensor dimensions to analytically characterize the estimated singular values and the alignment of estimated and true singular vectors at each step of the deflation procedure. Furthermore, this result can be used to construct estimators of the signal-to-noise ratios β_i and the alignments between the estimated and true rank-1 signal components.

Index Terms— Random tensor theory, Hotelling deflation, low-rank tensor decomposition, parameter estimation.

1. INTRODUCTION

The analysis of random tensors has attracted significant attention in the last decade since the introduction of the concept of tensor principal component analysis (PCA), which generalizes principal component analysis to high-order arrays. A simple model introduced in [1] is the so-called spiked tensor model of the form $\beta \boldsymbol{x}^{\otimes d} + \mathcal{W}/\sqrt{n}$ where $\boldsymbol{x} \in \mathbb{R}^n$ is some high-dimensional unit vector referred to as a *spike*, $\mathcal{W}$ is a *symmetric* random tensor of order d having independent Gaussian entries and $\beta \geq 0$ is a parameter controlling the signal-to-noise ratio.

Follow-up results have improved the understanding of the behavior of the spiked model and allowed to identify statistical and/or algorithmic thresholds for nontrivial signal recovery in the asymptotic (large-dimensional) regime $n \to \infty$ [2, 3, 4, 5, 6]. We briefly summarize their main findings as follows: for $d \geq 3$, it has been shown that there exists a sta-

tistical threshold $\beta_{stat} = O(1)$ below which it is information-theoretically impossible to recover or even detect the spike, while above β_{stat} recovery is theoretically possible. Moreover, the asymptotic alignment $\langle \boldsymbol{x}, \hat{\boldsymbol{x}}_{\text{ML}} \rangle$ between $\boldsymbol{x}$ and the maximum likelihood solution $\hat{\boldsymbol{x}}_{\text{ML}}$ was computed as a function of β and shown to be information-theoretically optimal [4]. Besides, since almost all tensor problems including computing $\hat{\boldsymbol{x}}_{\text{ML}}$ are NP-hard [7], many researchers were interested in identifying the algorithmic threshold for β above which recovery could be possible with a polynomial-time algorithm. In particular, the work [1] introduced a method for estimating $\boldsymbol{x}$ based on tensor unfolding and showed that spike recovery is possible above the algorithmic threshold $\beta_{\text{algo}} = O(n^{\frac{d-2}{4}})$.

These ideas were further generalized to the *asymmetric* spiked tensor model of the form $\beta \boldsymbol{x}_1 \otimes \cdots \otimes \boldsymbol{x}_d + \mathcal{W}/\sqrt{\sum_{\ell} n_{\ell}}$, where $\boldsymbol{x}_{\ell} \in \mathbb{R}^{n_{\ell}}$ are unit vectors and $\mathcal{W}$ is a random tensor with standard Gaussian i.i.d. entries. Specifically, [8] provided an analysis of the unfolding method for asymmetric tensors and determined its algorithmic threshold to be $\beta_{\text{algo}} = O(n^{\frac{d-2}{4}})$ when $n_i = n$ for all i, while the analysis of [9] suggests the existence of a statistical threshold $\beta_s = O(1)$ above which a local solution $\boldsymbol{u}_{\ell}$ of the maximum-likelihood estimator (MLE) of the dominant rank-1 component non-trivially aligns with the signal, and further quantified the asymptotic alignment $\langle \boldsymbol{x}_{\ell}, \boldsymbol{u}_{\ell} \rangle$.

Contribution: In this paper, we address the extension of these ideas to a more general setting, namely asymmetric *low-rank* spiked tensors of the form $\sum_{i=1}^{r} \beta_i \boldsymbol{x}_{i,1} \otimes \cdots \otimes \boldsymbol{x}_{i,d} + \mathcal{W}/\sqrt{\sum_{\ell} n_{\ell}}$ where $\boldsymbol{x}_{i,\ell} \in \mathbb{R}^{n_{\ell}}$ are unit vectors. Specifically, we consider the study of a simple deflation procedure, first introduced by Hotelling in the context of matrix principal component analysis [10] and still used in modern applications such as the recent AlphaTensor model [11], which consists in iterated rank-one approximation followed by subtraction of the estimated rank-one component.

Differently from the matrix PCA, the rank-one terms $\beta_i \boldsymbol{x}_{i,1} \otimes \cdots \otimes \boldsymbol{x}_{i,d}$ in our model are *not* pairwise orthogonal, and therefore in general deflation cannot possibly recover them exactly, even in the absence of noise [12]. Our main goal is to precisely compute how well deflation can find an *approximate* decomposition, in terms of the alignments of the

estimated terms with the true ones, and as a function of the cross-term alignments $\langle \boldsymbol{x}_{i,\ell}, \boldsymbol{x}_{j,\ell} \rangle \in (0,1)$ for $i \neq j$, while the analysis in the orthogonal case (i.e. when $\langle \boldsymbol{x}_{i,\ell}, \boldsymbol{x}_{j,\ell} \rangle = 0$ for $i \neq j$) trivially boils down to the rank-one model studied in [5, 9]. Precisely, we characterize this behavior by computing the alignments $\langle \boldsymbol{x}_{i,\ell}, \boldsymbol{u}_{j,\ell} \rangle$ in the asymmetric case, in the high-dimensional regime when $n_\ell \to \infty$, using the random matrix approach developed in [6, 9].

2. NOTATION AND BACKGROUND

Scalars (resp. vectors) are denoted by lowercase (resp. bold lowercase) letters. Tensors are denoted as in $\mathcal{A}$, $\mathcal{T}$. We set $[n] \equiv \{1, \ldots, n\}$, and $\xrightarrow{\text{a.s.}}$ denotes almost sure convergence.

2.1. TENSOR NOTATION AND CONTRACTIONS

In this section, we provide the main tensor notation and definitions used throughout the paper.

Order-d tensors: The set of order-d tensors of size $(n_1, \ldots, n_d)$ is denoted $\mathbb{R}^{\prod_{\ell=1}^{d} n_\ell}$. The scalar $T_{p_1 \cdots p_d}$ or $[\mathcal{T}]_{p_1 \cdots p_d}$ denotes the $(p_1, \ldots, p_d)$ entry of $\mathcal{T} \in \mathbb{R}^{\prod_{\ell=1}^{d} n_\ell}$.

Rank-r tensors: A tensor $\mathcal{T}$ is said to be of rank-one if it can be represented as the outer product of d real-valued vectors $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_d) \in \mathbb{R}^{n_1} \times \cdots \times \mathbb{R}^{n_d}$. In this case, we write $\mathcal{T} = \boldsymbol{x}_1 \otimes \cdots \otimes \boldsymbol{x}_d$, where the outer product is defined so that $[\boldsymbol{x}_1 \otimes \cdots \otimes \boldsymbol{x}_d]_{p_1, \ldots, p_d} = \prod_{\ell=1}^{d} [\boldsymbol{x}_\ell]_{p_\ell}$. More generally, a tensor $\mathcal{T} \in \mathbb{R}^{\prod_{\ell=1}^{d} n_\ell}$ is of rank r if it can be expressed as the sum of r rank-one terms but not less, and is written as $\mathcal{T} = \sum_{i=1}^{r} \boldsymbol{x}_{i,1} \otimes \cdots \otimes \boldsymbol{x}_{i,d}$, where $\boldsymbol{x}_{i,\ell} \in \mathbb{R}^{n_\ell}$ for all $(i, \ell) \in [r] \times [d]$.

Tensor contractions: For any vectors $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_d$, contractions of a tensor $\mathcal{A}$ are denoted by $\mathcal{A}(\boldsymbol{u}_1, \ldots, \boldsymbol{u}_d) = \sum A_{p_1 \cdots p_d} \prod_{k=1}^{d} [\boldsymbol{u}_k]_{p_k} \in \mathbb{R}$. Partial contractions are denoted similarly but with some arguments replaced by a dot. For instance, $\mathcal{A}(\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{\ell-1}, \cdot, \boldsymbol{u}_{\ell+1}, \ldots, \boldsymbol{u}_d) \in \mathbb{R}^{n_\ell}$.

Tensor norms: The inner product in $\mathbb{R}^{\prod_{\ell=1}^{d} n_\ell}$ is defined as $\langle \mathcal{A}, \mathcal{B} \rangle \equiv \sum_{p_1, \ldots, p_d} A_{p_1 \cdots p_d} B_{p_1 \cdots p_d}$, the ℓ_2-norm or Frobenius norm as $\|\mathcal{A}\|_F \equiv \sqrt{\langle \mathcal{A}, \mathcal{A} \rangle}$, and the spectral norm as $\|\mathcal{A}\| \equiv \sup_{\|\boldsymbol{u}_\ell\|=1} |\mathcal{A}(\boldsymbol{u}_1, \ldots, \boldsymbol{u}_d)|$.

Best rank-one approximation and tensor power iteration: A best rank-one approximation of $\mathcal{T}$ corresponds to a rank-one tensor $\lambda \boldsymbol{u}_1 \otimes \cdots \otimes \boldsymbol{u}_d$, where $\lambda > 0$ and $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_d$ are unitary vectors, that minimizes $\|\mathcal{T} - \lambda \boldsymbol{u}_1 \otimes \cdots \otimes \boldsymbol{u}_d\|_F^2$. In particular, due to the Karush-Kuhn-Tucker conditions, the tuple $(\lambda, \boldsymbol{u}_1, \ldots, \boldsymbol{u}_d)$ satisfies for each $i \in [d]$:

$$\begin{cases} \mathcal{T}(\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{\ell-1}, \cdot, \boldsymbol{u}_{\ell+1}, \ldots, \boldsymbol{u}_d) = \lambda \boldsymbol{u}_\ell, \\ \lambda = \mathcal{T}(\boldsymbol{u}_1, \ldots, \boldsymbol{u}_d). \end{cases} \quad (1)$$

These identities generalize the concept of singular value and vectors to tensors [13] (it is easy to check that (1) yields the classical definition of a matrix singular vector for $d = 2$); the scalar λ coincides with the spectral norm of $\mathcal{T}$.

Numerically, a rank-one approximation can be computed approximately via *tensor power iteration* whose t-th iteration consists in performing for $\ell = 1, \ldots, d$ the update

$$\boldsymbol{u}_\ell^{(t+1)} = \phi \left(\mathcal{T}(\boldsymbol{u}_1^{(t)}, \ldots, \boldsymbol{u}_{\ell-1}^{(t)}, \cdot, \boldsymbol{u}_{\ell+1}^{(t)}, \ldots, \boldsymbol{u}_d^{(t)}) \right),$$

where $\phi(\boldsymbol{x}) = \|\boldsymbol{x}\|^{-1}\boldsymbol{x}$, starting from some initial value [14].

2.2. RANDOM MATRIX THEORY BACKGROUND

In this section, we recall some standard tools from random *matrix* theory (RMT) which are at the heart of our main results. In essence, RMT focuses on describing the distribution of eigenvalues of random matrices. Specifically, we consider the *resolvent* formalism [15] which allows one to characterize the spectral behavior of large symmetric random matrices. Given a symmetric matrix $\boldsymbol{S} \in \mathbb{R}^{n \times n}$, the resolvent of $\boldsymbol{S}$ is defined by $\boldsymbol{R}(z) = (\boldsymbol{S} - z\boldsymbol{I}_n)^{-1}$ for $z \in \mathbb{C} \setminus \text{Sp}(\boldsymbol{S})$, where $\text{Sp}(\boldsymbol{S})$ denotes the set of eigenvalues of $\boldsymbol{S}$.

In a celebrated result [16], Marčenko and Pastur showed that under certain technical assumptions on some random matrix $\boldsymbol{S} \in \mathbb{R}^{n \times n}$ with eigenvalues $\lambda_1, \ldots, \lambda_n$, the *empirical spectral measure* of $\boldsymbol{S}$, defined as $\hat{\mu} = \frac{1}{n} \sum_{i=1}^{n} \delta_{\lambda_i}$ with δ_x the Dirac measure at $x \in \mathbb{R}$, converges in the weak sense [17] to some deterministic probability measure μ as $n \to \infty$; RMT aims at describing such μ. To this end, one of the widely considered approaches relies on the *Stieltjes transform* [18]. Given a probability measure μ, let $\mathcal{S}(\mu)$ denote its support. The Stieltjes transform of μ is defined as $g_\mu(z) = \int \frac{d\mu(\lambda)}{\lambda - z}$ with $z \in \mathbb{C} \setminus \sup(\mu)$, and the inverse formula allows one to describe the density of μ (if it exists) as $\mu(dx) = \frac{1}{\pi} \lim_{\varepsilon \to 0} \Im[g_\mu(x + i\varepsilon)]$.

The Stieltjes transform of the empirical spectral measure, denoted by $\hat{\mu}$, is closely related to the resolvent of $\boldsymbol{S}$ through the normalized trace operator. In fact, $g_{\hat{\mu}}(z) = \frac{1}{n} \text{Tr} \boldsymbol{R}(z)$ and the *almost sure* convergence of $g_{\hat{\mu}}(z)$ to some deterministic Stieltjes transform $g(z)$ is equivalent to the weak convergence between the underlying probability measures [18]. Our analysis relies on estimating quantities involving $\frac{1}{n} \text{Tr} \boldsymbol{R}(z)$, making the use of the resolvent approach a natural choice.

3. MODEL AND MAIN RESULTS

We start by formally describing our considered model. We let $r \geq 1$ and $d \geq 3$, and consider the following rank-r order-d spiked tensor model

$$\mathcal{T}_1 = \sum_{i=1}^{r} \beta_i \boldsymbol{x}_{i,1} \otimes \cdots \otimes \boldsymbol{x}_{i,d} + \frac{1}{\sqrt{n}} \mathcal{W}, \quad (2)$$

where $W_{p_1 \cdots p_d} \sim \mathcal{N}(0,1)$ i.i.d., $\boldsymbol{x}_{i,\ell} \in \mathbb{R}^{n_\ell}$ are unit-norm vectors, $n = \sum_{\ell=1}^{d} n_\ell$ and $\beta_i \geq 0$.

Tensor deflation: For tensors, Hotelling deflation is implemented by computing $\mathcal{T}_2, \mathcal{T}_3, \ldots$ sequentially through

$$\mathcal{T}_{i+1} = \mathcal{T}_i - \hat{\lambda}_i \hat{\boldsymbol{u}}_{i,1} \otimes \cdots \otimes \hat{\boldsymbol{u}}_{i,d} \quad \text{for} \quad i \in [r], \qquad (3)$$

where $\hat{\lambda}_i \hat{\boldsymbol{u}}_{i,1} \otimes \cdots \otimes \hat{\boldsymbol{u}}_{i,d}$ is a (in practice, local) minimizer of $\left\| \mathcal{T}_i - \lambda_i \boldsymbol{u}_{i,1} \otimes \cdots \otimes \boldsymbol{u}_{i,d} \right\|_F^2$ under the constraints $\lambda_i > 0$ and $\|\boldsymbol{u}_{i,\ell}\| = 1$ for all $(i, \ell) \in [r] \times [d]$, thus satisfying

$$\mathcal{T}_i(\hat{\boldsymbol{u}}_{i,1}, \ldots, \hat{\boldsymbol{u}}_{i,\ell-1}, \cdot, \hat{\boldsymbol{u}}_{i,\ell+1}, \ldots, \hat{\boldsymbol{u}}_{i,d}) = \hat{\lambda}_i \hat{\boldsymbol{u}}_{i,\ell} \qquad (4)$$

for all $\ell \in [d]$. In the following, we aim at computing the limits of $\hat{\lambda}_i$ and $\langle \boldsymbol{x}_{i,\ell}, \hat{\boldsymbol{u}}_{j,\ell} \rangle$ for $i, j \in [r]$ and $\ell \in [d]$ when the dimensions n_i grow large.

3.1. SKETCH OF THE ANALYTICAL APPROACH

We follow the approach developed in [9], whereby each tensor $\mathcal{T}_i$ is associated to a structured random matrix given by $\boldsymbol{M} = \boldsymbol{\Phi}_d(\mathcal{T}_i, \hat{\boldsymbol{u}}_{i,1}, \ldots, \hat{\boldsymbol{u}}_{i,d}) \in \mathbb{R}^{n \times n}$, where

$$\boldsymbol{\Phi}_d : \begin{pmatrix} \boldsymbol{X} \\ \boldsymbol{a}_1 \\ \vdots \\ \boldsymbol{a}_d \end{pmatrix} \longmapsto \begin{bmatrix} \boldsymbol{0}_{n_1 \times n_1} & X^{12} & \cdots & X^{1d} \\ (X^{12})^\top & \boldsymbol{0}_{n_2 \times n_2} & \cdots & X^{2d} \\ \vdots & \vdots & \ddots & \vdots \\ (X^{1d})^\top & (X^{2d})^\top & \cdots & \boldsymbol{0}_{n_d \times n_d} \end{bmatrix} \qquad (5)$$

and each block $X^{\ell m}$ is an $n_\ell \times n_m$ matrix resulting from the contraction of X on $d - 2$ vectors, namely

$$X^{\ell m} \equiv X(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_{\ell-1}, \cdot, \boldsymbol{a}_{\ell+1}, \ldots, \boldsymbol{a}_{m-1}, \cdot, \boldsymbol{a}_{m+1}, \ldots, \boldsymbol{a}_d).$$

We further denote $\boldsymbol{R}(z) = (\boldsymbol{M} - z\boldsymbol{I}_n)^{-1}$ the resolvent of $\boldsymbol{M}$. Then, the characterization of the limits of $\hat{\lambda}_i$ and the alignments $\langle \boldsymbol{x}_{i,\ell}, \hat{\boldsymbol{u}}_{j,\ell} \rangle$ when $n_\ell \to \infty$ requires the computation of the Stieltjes transform of the limiting spectral measure of $\boldsymbol{\Phi}_d(\mathcal{T}_i, \hat{\boldsymbol{u}}_{i,1}, \ldots, \hat{\boldsymbol{u}}_{i,d})$. Hence, we need the following definition and technical assumptions.

Assumption 3.1. Let us assume that $n_\ell \to \infty$ for all $\ell \in [d]$, that the limits $c_\ell = \lim \frac{n_\ell}{\sum_{m=1}^d n_m}$ exist, and that $r = O(1)$. For notational convenience, in the sequel the notation $\lim Q$ stands for the limit of the quantity Q in this regime. We further assume that there exists a sequence of critical points such that $\hat{\lambda}_i \xrightarrow{a.s.} \lambda_i$, $|\langle \boldsymbol{x}_{i,\ell}, \hat{\boldsymbol{u}}_{j,\ell} \rangle| \xrightarrow{a.s.} \rho_{ij,\ell}$ and $|\langle \hat{\boldsymbol{u}}_{i,\ell}, \hat{\boldsymbol{u}}_{j,\ell} \rangle| \xrightarrow{a.s.} \eta_{ij,\ell}$ such that $\lambda_i \notin \mathcal{S}(\mu)$ and $\rho_{ij,\ell} > 0$.

Definition 3.2. Let μ be the probability measure with Stieltjes transform $g(z) = \sum_{\ell=1}^d g_\ell(z)$ verifying $\Im[g(z)] > 0$ for $\Im[z] > 0$, where $g_\ell(z)$ satisfies $g_\ell^2(z) - (g(z) + z)g_\ell(z) - c_\ell = 0$, for $z \notin \mathcal{S}(\mu)$ and $\mathcal{S}(\mu)$ stands for the support of μ.

The following result characterizes the limiting spectral measure of $\boldsymbol{\Phi}_d(\mathcal{T}_i, \hat{\boldsymbol{u}}_{i,1}, \ldots, \hat{\boldsymbol{u}}_{i,d})$:

Theorem 3.3. *Under Assumption 3.1, the empirical spectral measure of $\boldsymbol{\Phi}_d(\mathcal{T}_i, \hat{\boldsymbol{u}}_{i,1}, \ldots, \hat{\boldsymbol{u}}_{i,d})$ converges to the deterministic measure μ defined in Definition 3.2.*

In particular, $\frac{1}{n} \operatorname{Tr} \boldsymbol{R}(z) \xrightarrow{a.s.} g(z)$.

The proof follows similar arguments as in [9] and requires some additional arguments for controlling the statistical dependencies between the $\hat{\boldsymbol{u}}_{i,\ell}$'s and the noise $\mathcal{W}$. In fact, the result is identical to that of [9] because $r = O(1)$, since the contribution of the rank-one terms vanish asymptotically. As shown in [9], in the case $c_\ell = \frac{1}{d}$ for all $\ell \in [d]$, the measure μ describes a semi-circle or Wigner-type law of compact support $\mathcal{S}(\mu) = [-2\sqrt{\frac{d-1}{d}}, 2\sqrt{\frac{d-1}{d}}]$, the Stieltjes transform of which writes explicitly as

$$g(z) = \frac{-zd + d\sqrt{z^2 - \frac{4(d-1)}{d}}}{2(d-1)}, \quad z \notin \mathcal{S}(\mu). \qquad (6)$$

Limiting spectral norms and alignments: We introduce the quantities $\alpha_{ij,\ell} = \lim |\langle \boldsymbol{x}_{i,\ell}, \boldsymbol{x}_{j,\ell} \rangle|$, $f(z) = z + g(z)$, $h_\ell(z) = -\frac{c_\ell}{g_\ell(z)}$ that shall be used subsequently. The main result brought by this paper describes the asymptotic singular values and alignments obtained after each tensor deflation step as stated by the following theorem.

Theorem 3.4. *Assume that Assumption 3.1 holds. Then, λ_i, $\rho_{ij,\ell}$ and $\eta_{ij,\ell}$ satisfy the following system of equations*

$$\begin{cases} f(\lambda_i) + \sum_{k=1}^{i-1} \lambda_k \prod_{\ell=1}^d \eta_{ki,\ell} - \sum_{k=1}^r \beta_k \prod_{\ell=1}^d \rho_{ki,\ell} = 0 \,, \, 1 \leq i \leq r \\[2ex] h_\ell(\lambda_i)\rho_{ji,\ell} + \sum_{k=1}^{i-1} \lambda_k \rho_{jk,\ell} \prod_{m \neq \ell}^d \eta_{ki,m} - \sum_{k=1}^r \beta_k \alpha_{kj,\ell} \prod_{m \neq \ell}^d \rho_{ki,m} = 0 \\[2ex] \quad 1 \leq \ell \leq d, 1 \leq i, j \leq r \\[2ex] h_\ell(\lambda_i)\eta_{ji,\ell} + g_\ell(\lambda_j) \prod_{m \neq \ell}^d \eta_{ji,m} + \sum_{k=1}^{i-1} \lambda_k \eta_{kj,\ell} \prod_{m \neq \ell}^d \eta_{ki,m} + \cdots \\[2ex] \quad - \sum_{k=1}^r \beta_k \rho_{kj,\ell} \prod_{m \neq \ell}^d \rho_{ki,m} = 0, \, 1 \leq \ell \leq d, 1 \leq i < j \leq r \end{cases}$$

Sketch of the proof. We use similar arguments as in [9]: we evaluate the expectation of the scalar product between (4) and $\boldsymbol{x}_{i,\ell}$ or $\hat{\boldsymbol{u}}_{j,\ell}$ using Stein's Lemma.[1] Along the proof, the variances of $\hat{\lambda}_i$ and of the alignments $\rho_{ij,\ell}$, $\eta_{ij,\ell}$ are shown to decay as $O(n^{-1})$. $\qquad \square$

3.2. PARTICULAR CASE OF A RANK-2, ORDER-3 TENSOR

For the sake of clarity, let us consider the example of a rank-2 order-3 spiked tensor with $n_1 = n_2 = n_3$, for which we can get more insightful analytical results. The model in (2) becomes

$$\mathcal{T}_1 = \sum_{i=1}^2 \beta_i \boldsymbol{x}_{i,1} \otimes \boldsymbol{x}_{i,2} \otimes \boldsymbol{x}_{i,3} + \frac{1}{\sqrt{n}} \mathcal{W}. \qquad (7)$$

---

[1] $\mathbb{E}[W f(W)] = \mathbb{E}[f'(W)]$ for $W \sim \mathcal{N}(0,1)$ and a continuously differentiable f having at most polynomial growth.
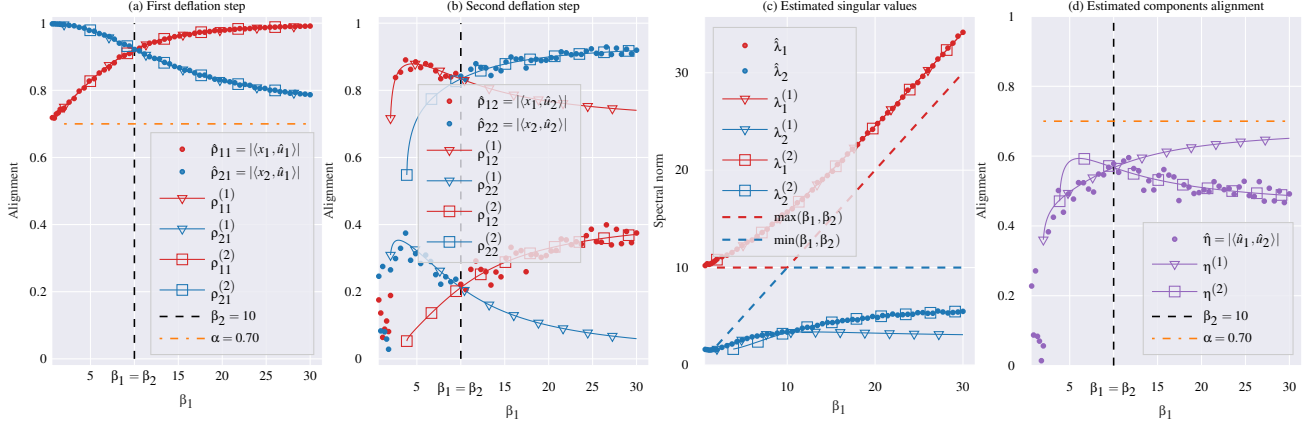
Fig. 1: Illustration of $\langle \boldsymbol{x}_i, \hat{\boldsymbol{u}}_j \rangle$ for $(i,j) \in \{1,2\}$, $\langle \hat{\boldsymbol{u}}_1, \hat{\boldsymbol{u}}_2 \rangle$, $\hat{\lambda}_1$ and $\hat{\lambda}_2$ vs. their limits for $\beta_2 = 10$ and $\alpha = \langle \boldsymbol{x}_1, \boldsymbol{x}_2 \rangle = 0.7$ of the two spikes tensor model in (7). (a) shows the alignments between the signal components and the first singular vectors corresponding to the best rank-one approximation of $\mathcal{T}_1$. (b) shows the alignments with the second singular vectors computed after deflation. (c) depicts the singular values. (d) shows the alignments between the singular vectors computed at each step of the deflation procedure. Simulations were performed on a tensor of dimensions $(50, 50, 50)$.

Furthermore, we assume that for all $i \neq j$ and each $\ell \in [3]$, $\lim |\langle \boldsymbol{x}_{i,\ell}, \boldsymbol{x}_{j,\ell} \rangle| = \alpha \in [0,1]$. In this case, since all the dimensions n_i are equal, the limits of $|\langle \boldsymbol{x}_{i,\ell}, \hat{\boldsymbol{u}}_{j,\ell} \rangle|$ and of $|\langle \hat{\boldsymbol{u}}_{1,\ell}, \hat{\boldsymbol{u}}_{2,\ell} \rangle|$ are both independent of ℓ by symmetry. Therefore, we drop their dependence on ℓ in our notations. The system of equations in Theorem 3.4 reduces to seven equations, as detailed in the following corollary:

Corollary 3.5. *Denote $\rho_{ij} = \lim |\langle \boldsymbol{x}_{i,\ell}, \hat{\boldsymbol{u}}_{j,\ell} \rangle|$ for $i,j \in [2]$ and $\eta = \lim |\langle \hat{\boldsymbol{u}}_{1,\ell}, \hat{\boldsymbol{u}}_{2,\ell} \rangle|$ and suppose that Assumption 3.1 holds, then λ_i, ρ_{ij} and η satisfy $\psi(\boldsymbol{\lambda}, \boldsymbol{\beta}, \boldsymbol{\rho}) = \boldsymbol{0}$ with $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \eta)$, $\boldsymbol{\rho} = (\rho_{11}, \rho_{12}, \rho_{21}, \rho_{22})$, $\boldsymbol{\beta} = (\beta_1, \beta_2, \alpha)$ and*

$$\psi(\boldsymbol{\lambda}, \boldsymbol{\beta}, \boldsymbol{\rho}) = \begin{pmatrix} f(\lambda_1) - \beta_1 \rho_{11}^3 - \beta_2 \rho_{21}^3 \\ h(\lambda_1)\rho_{11} - \beta_1 \rho_{11}^2 - \beta_2 \alpha \rho_{21}^2 \\ h(\lambda_1)\rho_{21} - \beta_1 \alpha \rho_{11}^2 - \beta_2 \rho_{21}^2 \\ f(\lambda_2) + \lambda_1 \eta^3 - \beta_1 \rho_{12}^3 - \beta_2 \rho_{22}^3 \\ h(\lambda_2)\rho_{12} + \lambda_1 \rho_{11}\eta^2 - \beta_1 \rho_{12}^2 - \beta_2 \alpha \rho_{22}^2 \\ h(\lambda_2)\rho_{22} + \lambda_1 \rho_{21}\eta^2 - \beta_1 \alpha \rho_{12}^2 - \beta_2 \rho_{22}^2 \\ h(\lambda_2)\eta + q(\lambda_1)\eta^2 - \beta_1 \rho_{11}\rho_{12}^2 - \beta_2 \rho_{21}\rho_{22}^2 \end{pmatrix} \quad (8)$$

where $h(z) = \frac{-1}{g(z)}$ and $q(z) = z + \frac{g(z)}{3}$ with $g(z)$ given by (6) for $d = 3$ and we recall that $f(z) = z + g(z)$.

Fixing $\boldsymbol{\beta} = (\beta_1, \beta_2, \alpha)$, one can solve $\psi(\boldsymbol{\lambda}, \boldsymbol{\beta}, \boldsymbol{\rho}) = \boldsymbol{0}$ in $(\boldsymbol{\lambda}, \boldsymbol{\rho})$ using a Newton-Raphson method with 100 random initializations. We keep the solutions ensuring that $0 \leq \eta, \rho_{ij} \leq 1$ and $\lambda_1, \lambda_2 > 2\sqrt{\frac{2}{3}}$. Note that the equation $\psi(\cdot, \boldsymbol{\beta}, \cdot) = \boldsymbol{0}$ can have several distinct solutions. These different solutions correspond to different sequences of critical points.

Indeed, we illustrate in Figure 1 the matching between $\hat{\boldsymbol{\lambda}} = (\hat{\lambda}_1, \hat{\lambda}_2, \hat{\eta})$, $\hat{\boldsymbol{\rho}} = (\hat{\rho}_{11}, \hat{\rho}_{12}, \hat{\rho}_{21}, \hat{\rho}_{22})$ (where the rank-one approximations are performed using tensor power iteration initialized by tensor SVD [19]) and their asymptotic limits $\boldsymbol{\lambda}, \boldsymbol{\rho}$. We observe that, when $\alpha = 0.7$, the equation $\psi(\cdot, \boldsymbol{\beta}, \cdot) =$

$\boldsymbol{0}$ has two solutions (we note $(\boldsymbol{\rho}^{(1)}, \boldsymbol{\lambda}^{(1)})$ the solution with the highest ρ_{12} and $(\boldsymbol{\rho}^{(2)}, \boldsymbol{\lambda}^{(2)})$ the other solution) and our initialization of the tensor power method seems to favor the global minimizer. Moreover, note that for β_1 large enough (e.g. $\beta_1 > 2\sqrt{2/3}$, as per Figure 1), signal detection is always possible and the singular vector $\hat{\boldsymbol{u}}_1$ presents a higher alignment with the signal components having the highest β_i. However, when β_1 is not large enough, there exists a region (see in particular Figure 1(b)) where it is information-theoretically impossible to detect a signal, while outside this region estimation becomes possible with the MLE. We refer the reader to [6, 9] for detailed discussions about the detectability threshold in the rank-one case.

4. CONCLUSION AND DISCUSSION

We have provided an analysis of a tensor deflation method in the high-dimensional regime and assumed a low-rank spiked tensor model with correlated signal components. Our analysis allows for a precise description of the asymptotic behavior of such models. Moreover, our results can be exploited to design an estimator of the model parameters by solving the system $\psi(\hat{\boldsymbol{\lambda}}, \cdot, \cdot) = \boldsymbol{0}$, therefore, possibly resulting in improved deflation. This paves a new way for the analysis of more sophisticated tensor decomposition methods and the understanding of more general tensor models through random tensor theory.

Our current results still have some limitations. In particular, we consider Assumption 3.1 which is standard in the analysis of large spiked random tensors [6] but the almost sure convergences stated therein need to be proved through concentration results. Also, the existence and uniqueness of the solutions of the involved systems of equations are not discussed in this paper and are left for future investigation.

5. REFERENCES

[1] Emile Richard and Andrea Montanari, "A statistical model for tensor PCA," *Advances in neural information processing systems*, vol. 27, 2014.

[2] Amelia Perry, Alexander S Wein, and Afonso S Bandeira, "Statistical limits of spiked tensor models," in *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*. Institut Henri Poincaré, 2020, vol. 56, pp. 230–264.

[3] Thibault Lesieur, Léo Miolane, Marc Lelarge, Florent Krzakala, and Lenka Zdeborová, "Statistical and computational phase transitions in spiked tensor estimation," in *2017 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2017, pp. 511–515.

[4] Aukosh Jagannath, Patrick Lopatto, and Leo Miolane, "Statistical thresholds for tensor PCA," *The Annals of Applied Probability*, vol. 30, no. 4, pp. 1910–1933, 2020.

[5] Wei-Kuo Chen, Madeline Handschy, and Gilad Lerman, "Phase transition in random tensors with multiple independent spikes," *The Annals of Applied Probability*, vol. 31, no. 4, pp. 1868–1913, 2021.

[6] José Henrique de Morais Goulart, Romain Couillet, and Pierre Comon, "A random matrix perspective on random tensors," *Journal of Machine Learning Research*, vol. 23, no. 264, pp. 1–36, 2022.

[7] Christopher J Hillar and Lek-Heng Lim, "Most tensor problems are NP-hard," *Journal of the ACM (JACM)*, vol. 60, no. 6, pp. 1–39, 2013.

[8] Gérard Ben Arous, Daniel Zhengyu Huang, and Jiaoyang Huang, "Long random matrices and tensor unfolding," *arXiv preprint arXiv:2110.10210*, 2021.

[9] Mohamed El Amine Seddik, Maxime Guillaud, and Romain Couillet, "When random tensors meet random matrices," *Annals of Applied Probability*, 2023, in press, preprint arXiv:2112.12348.

[10] Harold Hotelling, "Analysis of a complex of statistical variables into principal components," *Journal of Educational Psychology*, vol. 24, no. 6, pp. 417–441, Sept. 1933.

[11] Alhussein Fawzi, Matej Balog, Aja Huang, Thomas Hubert, Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Francisco J R Ruiz, Julian Schrittwieser, Grzegorz Swirszcz, et al., "Discovering faster matrix multiplication algorithms with reinforcement learning," *Nature*, vol. 610, no. 7930, pp. 47–53, 2022.

[12] Alwin Stegeman and Pierre Comon, "Subtracting a best rank-1 approximation may increase tensor rank," *Linear Algebra and its Applications*, vol. 433, no. 7, pp. 1276–1300, 2010.

[13] Lek-Heng Lim, "Singular values and eigenvalues of tensors: a variational approach," in *Proc. IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 2005, pp. 129–132.

[14] Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky, "Tensor decompositions for learning latent variable models," *Journal of machine learning research*, vol. 15, pp. 2773–2832, 2014.

[15] Walid Hachem, Philippe Loubaton, and Jamal Najim, "Deterministic equivalents for certain functionals of large random matrices," *The Annals of Applied Probability*, vol. 17, no. 3, pp. 875–930, 2007.

[16] V. A. Marčenko and L. A. Pastur, "Distribution of eigenvalues for some sets of random matrices," *Mathematics of the USSR-Sbornik*, vol. 1, no. 4, pp. 457–483, Apr. 1967.

[17] Aad W Van Der Vaart and Jon A Wellner, "Weak convergence," in *Weak convergence and empirical processes*, pp. 16–28. Springer, 1996.

[18] Terence Tao, *Topics in random matrix theory*, vol. 132, American Mathematical Soc., 2012.

[19] Arnab Auddy and Ming Yuan, "On estimating rank-one spiked tensors in the presence of heavy tailed errors," *IEEE Transactions on Information Theory*, 2022.