

# Performance of Rank-One Tensor Approximation on Incomplete Data

Hugo LEBEAU

Inria, ENS de Lyon, CNRS, Université Claude Bernard Lyon 1, LIP, UMR 5668, 69342, Lyon cedex 07, France  
work done while at Université Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, UMR 5217, 38058, Grenoble cedex 09, France

**Résumé** – On s’intéresse à l’estimation d’un signal tensoriel de rang 1 lorsque seulement une portion  $\varepsilon$  de son observation bruitée est disponible. Nous montrons que l’étude de ce problème se ramène à celle d’un modèle de matrice aléatoire dont l’analyse spectrale donne accès aux performances de reconstruction. Ces résultats mettent en lumière et précisent la perte de performance induite par une réduction artificielle du coût en mémoire d’un tenseur via la suppression d’une partie aléatoire de ses entrées.

**Abstract** – We are interested in the estimation of a rank-one tensor signal when only a portion  $\varepsilon$  of its noisy observation is available. We show that the study of this problem can be reduced to that of a random matrix model whose spectral analysis gives access to the reconstruction performance. These results shed light on and specify the loss of performance induced by an artificial reduction of the memory cost of a tensor via the deletion of a random part of its entries.

## 1 Introduction

Processing large amounts of multimodal data is a common challenge of the big data era. Multiway arrays (tensors) have inherently a high volume, scaling exponentially with the number of modes, which can rapidly reach a memory limit. Often, the information sought in such gigantic raw data is a latent and low-dimensional structure. Hence, some techniques have been developed to compute tensor decompositions (such as the Canonical Polyadic Decomposition or the Multilinear SVD [7]) from *incomplete* tensors [8]. At first, they were used to tackle the issue of missing values, which is common in chemometrics data. Nowadays, they can be used to compute tensor decompositions on intentionally punctured data. That is, in order to reduce the volume of a tensor, only a small portion of its entries is actually stored and inference is performed on the incomplete tensor.

A simple yet effective procedure to reduce the volume of a huge tensor  $\mathcal{T} \in \mathbb{R}^{n_1 \times \dots \times n_d}$  is therefore to store a portion  $\varepsilon \in [0, 1]$  of its entries chosen at random. Concretely, this is modeled by the entrywise multiplication  $\mathcal{T} \odot \mathcal{B}$  where  $\mathcal{B}$  has i.i.d. Bernoulli( $\varepsilon$ ) entries (i.e., 1 with probability  $\varepsilon$  and 0 otherwise). It is then natural to ask: what is the loss of performance induced by this procedure? In other words, how different is the inference performed with  $\mathcal{T} \odot \mathcal{B}$  from the one performed with  $\mathcal{T}$ ? A similar procedure on kernel matrices is studied using random matrix tools in [2]. It is shown therein that, when the dimension of the data is not too large, the number of entries of the kernel matrix can be drastically reduced with almost no impairment in the resulting estimation. Yet, in our case, puncturing  $\odot \mathcal{B}$  is directly applied onto the raw data  $\mathcal{T}$ , which may cause a considerable performance loss. Our goal is therefore to quantify the latter in order to help deciding on the best trade-off between performance and memory cost. Our study also provides insights into the performance of algorithms relying on a similar procedure such as the CP-WOPT algorithm [1, 8] which computes a canonical polyadic decomposition from a tensor with missing data.

For the sake of simplicity we consider order-3 tensors  $\mathcal{T} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ , the generalization to order- $d$  tensors  $\mathcal{T} \in$

$\mathbb{R}^{n_1 \times \dots \times n_d}$  is straightforward. In order to study the estimation of a planted signal in  $\mathcal{T}$  from  $\mathcal{T} \odot \mathcal{B}$ , we consider the rank-one spiked tensor model  $\mathcal{T} = \beta \mathbf{x} \otimes \mathbf{y} \otimes \mathbf{z} + \frac{1}{\sqrt{N}} \mathcal{N}$  where  $\beta > 0$  is a parameter controlling the signal-to-noise ratio,  $\mathbf{x}, \mathbf{y}, \mathbf{z}$  are fixed unit-norm vectors,  $\otimes$  is the outer product ( $\mathbf{x} \otimes \mathbf{y} = \mathbf{x} \mathbf{y}^\top$ ),  $N = n_1 + n_2 + n_3$  and  $\mathcal{N}$  is an  $n_1 \times n_2 \times n_3$  array with i.i.d.  $\mathcal{N}(0, 1)$  entries. We place ourselves in the large-dimensional regime where  $n_1, n_2, n_3 \rightarrow +\infty$  with the ratios  $(\frac{n_1}{N}, \frac{n_2}{N}, \frac{n_3}{N}) = (c_1, c_2, c_3)$  kept constant (and non-zero), simply denoted “ $N \rightarrow +\infty$ ”. This regime models the fact that all dimensions of the tensor are large, although they remain finite in practice. In this setting, we consider the estimation of  $\beta, \mathbf{x}, \mathbf{y}, \mathbf{z}$  with the best rank-one approximation of  $\mathcal{T} \odot \mathcal{B}$ ,

$$(\sigma^*, \mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*) \in \arg \min_{\sigma, \mathbf{u}, \mathbf{v}, \mathbf{w}} \|\mathcal{T} \odot \mathcal{B} - \sigma \mathbf{u} \otimes \mathbf{v} \otimes \mathbf{w}\|_{\text{F}}^2 \quad (1)$$

where the minimum is taken over  $\sigma > 0$  and  $\mathbf{u}, \mathbf{v}, \mathbf{w}$  of unit norm. The “vanilla” rank-one estimation problem (i.e., without puncturing  $\odot \mathcal{B}$ ) is studied in [6], where the analysis of its critical points is reduced to that of an associated random matrix model  $\Phi$  defined with contractions of the tensor on its singular vectors. This new approach allows the study of the rank-one approximation through that of a random matrix model, thereby allowing the harnessing of random matrix theory. Here, we follow the same approach to study the problem of missing data in the rank-one tensor estimation problem.

Yet, the entrywise multiplication  $\odot \mathcal{B}$  raises some mathematical difficulties and the rigorous analysis of this model is delicate. Therefore, we must rely on the following heuristic.

**Heuristic 1.** A critical point  $(\sigma^*, \mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  of Problem (1) is asymptotically (as  $N \rightarrow +\infty$ ) independent of  $\mathcal{B}$ .

This property is hard to prove, but allows to write  $\mathbb{E}[\sum_{i,j,k} \mathcal{T}_{i,j,k} \mathcal{B}_{i,j,k} u_i^* v_j^* w_k^*] = \varepsilon \mathbb{E}[\sum_{i,j,k} \mathcal{T}_{i,j,k} u_i^* v_j^* w_k^*] + o(1)$  and likewise for other similar “contractions” where the coefficient  $\mathcal{B}_{i,j,k}$  can be replaced by  $\varepsilon$ , up to a vanishing term in the large  $N$  limit. Furthermore, this heuristic leads to precise predictions which are verified by simulations (see Figure 1 and Figure 3). The scope of this paper focuses on these results and their practical implications rather than the technical

mathematical aspects implied by the entrywise multiplication  $\odot \mathcal{B}$ , although they raise many interesting questions about the current state of our tools for tackling this kind of problems.

**Notations.** The set  $\{1, \dots, n\}$  of positive integers smaller or equal to  $n$  is denoted  $[n]$ . The canonical basis in  $\mathbb{R}^n$  is  $(e_i^{(n)})_{i \in [n]}$ . Given a matrix  $M \in \mathbb{R}^{n \times n}$ ,  $\text{Sp } M$  is the set of its eigenvalues. Given a tensor  $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$  and three vectors  $(a, b, c) \in \mathbb{R}^{n_1} \times \mathbb{R}^{n_2} \times \mathbb{R}^{n_3}$ , the *contraction* of  $\mathcal{A}$  onto  $a, b, c$  is  $\mathcal{A}(a, b, c) = \sum_{i=1}^{n_1} \sum_{j=2}^{n_2} \sum_{k=1}^{n_3} \mathcal{A}_{i,j,k} a_i b_j c_k$ .

## 2 Associated Random Matrix Model

### 2.1 First-Order Optimality Conditions

Problem (1) can be equivalently restated as

$$(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*) \in \arg \max_{\mathbf{u}, \mathbf{v}, \mathbf{w}} |[\mathcal{T} \odot \mathcal{B}](\mathbf{u}, \mathbf{v}, \mathbf{w})| \quad (2)$$

with  $\sigma^* = |[\mathcal{T} \odot \mathcal{B}](\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)|$ . The Lagrangian associated to this problem is

$$\mathcal{L} = [\mathcal{T} \odot \mathcal{B}](\mathbf{u}, \mathbf{v}, \mathbf{w}) - \frac{1}{2} \begin{bmatrix} \lambda_u \\ \lambda_v \\ \lambda_w \end{bmatrix} \cdot \begin{bmatrix} \|\mathbf{u}\|^2 - 1 \\ \|\mathbf{v}\|^2 - 1 \\ \|\mathbf{w}\|^2 - 1 \end{bmatrix}$$

where  $\lambda_u, \lambda_v, \lambda_w$  are Lagrange multipliers associated to the unit-norm constraints and  $\cdot$  denotes the dot product. From the first-order optimality conditions  $\nabla \mathcal{L} = \mathbf{0}$ , we find that  $\sigma^*, \mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*$  must satisfy

$$\begin{cases} [\mathcal{T} \odot \mathcal{B}](\cdot, \mathbf{v}^*, \mathbf{w}^*) = \sigma^* \mathbf{u}^* \\ [\mathcal{T} \odot \mathcal{B}](\mathbf{u}^*, \cdot, \mathbf{w}^*) = \sigma^* \mathbf{v}^* \\ [\mathcal{T} \odot \mathcal{B}](\mathbf{u}^*, \mathbf{v}^*, \cdot) = \sigma^* \mathbf{w}^* \end{cases} \quad (3)$$

### 2.2 Critical Points

As our goal is to reconstruct the signal  $\beta \mathbf{x} \otimes \mathbf{y} \otimes \mathbf{z}$  from the observation  $\mathcal{T} \odot \mathcal{B}$ , we seek the behavior of  $\sigma^*$  (as an estimator of  $\beta$ ) and the alignments of  $\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*$  with  $\mathbf{x}, \mathbf{y}, \mathbf{z}$ . These quantities are expected to *concentrate* around their mean, therefore it is sufficient to study their expectations. To that end, we have the following ‘‘Gaussian integration by parts’’ lemma.

**Lemma 1 (Stein).** *Let  $Z \sim \mathcal{N}(0, 1)$  and  $f : \mathbb{R} \rightarrow \mathbb{C}$  be a polynomially bounded differentiable function such that  $\mathbb{E}[|f'(Z)|] < +\infty$ . Then,  $\mathbb{E}[Zf(Z)] = \mathbb{E}[f'(Z)]$ .*

Based on this lemma and Heuristic 1, we can derive the following relations

$$\begin{aligned} \mathbb{E}[\sigma^*] &\asymp \varepsilon \beta \mathbb{E}[\langle \mathbf{x}, \mathbf{u}^* \rangle \langle \mathbf{y}, \mathbf{v}^* \rangle \langle \mathbf{z}, \mathbf{w}^* \rangle] + \frac{\varepsilon}{\sqrt{N}} \sum_{i,j,k} \mathbb{E} \left[ \frac{\partial u_i^* v_j^* w_k^*}{\partial N_{i,j,k}} \right] \\ \mathbb{E}[\sigma^* \langle \mathbf{x}, \mathbf{u}^* \rangle] &\asymp \varepsilon \beta \mathbb{E}[\langle \mathbf{y}, \mathbf{v}^* \rangle \langle \mathbf{z}, \mathbf{w}^* \rangle] + \frac{\varepsilon}{\sqrt{N}} \sum_{i,j,k} x_i \mathbb{E} \left[ \frac{\partial v_j^* w_k^*}{\partial N_{i,j,k}} \right] \\ \mathbb{E}[\sigma^* \langle \mathbf{y}, \mathbf{v}^* \rangle] &\asymp \varepsilon \beta \mathbb{E}[\langle \mathbf{x}, \mathbf{u}^* \rangle \langle \mathbf{z}, \mathbf{w}^* \rangle] + \frac{\varepsilon}{\sqrt{N}} \sum_{i,j,k} y_j \mathbb{E} \left[ \frac{\partial u_i^* w_k^*}{\partial N_{i,j,k}} \right] \\ \mathbb{E}[\sigma^* \langle \mathbf{z}, \mathbf{w}^* \rangle] &\asymp \varepsilon \beta \mathbb{E}[\langle \mathbf{x}, \mathbf{u}^* \rangle \langle \mathbf{y}, \mathbf{v}^* \rangle] + \frac{\varepsilon}{\sqrt{N}} \sum_{i,j,k} z_k \mathbb{E} \left[ \frac{\partial u_i^* v_j^*}{\partial N_{i,j,k}} \right] \end{aligned}$$

where  $\asymp$  means equality up to a vanishing term as  $N \rightarrow +\infty$ .

**Remark 1.** In order to avoid any differentiability issues, we consider, from now on,  $(\sigma^*, \mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  as a fixed *critical point* of Problem (1) rather than a global minimizer.

**Remark 2.** The concentration of relevant quantities stems from standard arguments such as the Poincaré–Nash inequality for Gaussian vectors (see Lemma 2.14 in [3]) and the subgaussianity of the entries of  $\mathcal{N} \odot \mathcal{B}$ .

From the previous relations, we see that it is necessary to find an expression for the derivatives of  $\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*$  with respect to the entries of  $\mathcal{N}$ . Despite their intricate dependence on the noise, differentiating conditions given in Equation (3) leads to the following result.

**Proposition 1.** *Define the  $N \times N$  matrix*

$$\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*) = \begin{bmatrix} \mathbf{0}_{n_1 \times n_1} & \Phi^{(1,2)} & \Phi^{(1,3)} \\ \Phi^{(1,2)\top} & \mathbf{0}_{n_2 \times n_2} & \Phi^{(2,3)} \\ \Phi^{(1,3)\top} & \Phi^{(2,3)\top} & \mathbf{0}_{n_3 \times n_3} \end{bmatrix}$$

with  $\Phi^{(1,2)} = [\mathcal{T} \odot \mathcal{B}](\cdot, \cdot, \mathbf{w}^*)$ ,  $\Phi^{(1,3)} = [\mathcal{T} \odot \mathcal{B}](\cdot, \mathbf{v}^*, \cdot)$  and  $\Phi^{(2,3)} = [\mathcal{T} \odot \mathcal{B}](\mathbf{u}^*, \cdot, \cdot)$ . If  $\sigma^*$  is not an eigenvalue of  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$ , then, for all  $(i, j, k) \in [n_1] \times [n_2] \times [n_3]$ ,

$$\frac{\partial}{\partial N_{i,j,k}} \begin{bmatrix} \mathbf{u}^* \\ \mathbf{v}^* \\ \mathbf{w}^* \end{bmatrix} = -\frac{\mathcal{B}_{i,j,k}}{\sqrt{N}} \mathbf{Q}(\sigma^*) \begin{bmatrix} v_j^* w_k^* (e_i^{(n_1)} - u_i^* \mathbf{u}^*) \\ u_i^* w_k^* (e_j^{(n_2)} - v_j^* \mathbf{v}^*) \\ u_i^* v_j^* (e_k^{(n_3)} - w_k^* \mathbf{w}^*) \end{bmatrix}$$

where  $\mathbf{Q}(\sigma^*) = [\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*) - \sigma^* \mathbf{I}_N]^{-1}$ .

This Proposition shows the link between the derivatives of  $\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*$  and the associated matrix model  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  through a matrix  $\mathbf{Q}(\sigma^*)$ , known as the *resolvent* — a common tool in random matrix theory which is used to study the spectral properties of random matrices [3]. The reconstruction performance of  $(\sigma^*, \mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  can therefore be revealed by a spectral analysis of  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$ .

### 2.3 First Remarks on the Associated Matrix

Notice that  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  can be decomposed as

$$\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*) = \varepsilon \beta \mathbf{V} \mathbf{S} \mathbf{V}^\top + \Phi_0(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*) + \mathbf{E} \quad (4)$$

where  $\mathbf{V} = \begin{bmatrix} \mathbf{x} & \mathbf{0}_{n_1} & \mathbf{0}_{n_1} \\ \mathbf{0}_{n_2} & \mathbf{y} & \mathbf{0}_{n_2} \\ \mathbf{0}_{n_3} & \mathbf{0}_{n_3} & \mathbf{z} \end{bmatrix}$ ,  $\mathbf{S} = \begin{bmatrix} 0 & \langle \mathbf{z}, \mathbf{w}^* \rangle & \langle \mathbf{y}, \mathbf{v}^* \rangle \\ \langle \mathbf{z}, \mathbf{w}^* \rangle & 0 & \langle \mathbf{x}, \mathbf{u}^* \rangle \\ \langle \mathbf{y}, \mathbf{v}^* \rangle & \langle \mathbf{x}, \mathbf{u}^* \rangle & 0 \end{bmatrix}$ ,  $\Phi_0$  is defined as  $\Phi$  but with  $\mathcal{T}$  replaced by  $\frac{1}{\sqrt{N}} \mathcal{N}$  (i.e.,  $\beta = 0$ ) and  $\|\mathbf{E}\| \rightarrow 0$  almost surely as  $N \rightarrow +\infty$  (Heuristic 1). Hence,  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  has the structure of a *spiked model* — it is a rank-3 perturbation of a random matrix. Moreover, based on the definition of  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  and a second-order optimality condition of Problem (2), we can state the following properties on its spectrum.

**Proposition 2.** *1.  $2\sigma^*$  is an eigenvalue of  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  with multiplicity 1 and eigenvector  $\begin{bmatrix} \mathbf{u}^* \\ \mathbf{v}^* \\ \mathbf{w}^* \end{bmatrix}$ .*  
*2.  $-\sigma^*$  is an eigenvalue of  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  with multiplicity (at least) 2 and eigenspace spanned by  $\begin{bmatrix} \mathbf{u}^* & \mathbf{u}^* \\ -\mathbf{v}^* & \mathbf{0}_{n_2} \\ \mathbf{0}_{n_3} & -\mathbf{w}^* \end{bmatrix}$ .*  
*3. For all eigenvalue  $\lambda$  of  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$ , either  $\lambda = 2\sigma^*$  or  $|\lambda| \leq \sigma^*$ .*

These properties are illustrated by the histograms of eigenvalues depicted in Figure 1.

### 3 Random Matrix Analysis

#### 3.1 Limiting Spectral Distribution

In order to study the eigenvalue distribution of the associated random matrix  $\Phi$ , we must introduce an important tool, namely, the Stieltjes transform of a probability distribution.

**Definition 1.** Given a real probability measure  $\mu$ , its *Stieltjes transform* is  $m_\mu : z \in \mathbb{C} \setminus \mathbb{R} \mapsto \int_{\mathbb{R}} \frac{d\mu(t)}{t-z}$ .

The Stieltjes transform satisfies many interesting properties (see [3]) and the underlying probability measure can be reconstructed with the inverse Stieltjes transform formula  $\mu([a, b]) = \frac{1}{\pi} \lim_{\eta \downarrow 0} \int_a^b \Im[m_\mu(x + i\eta)] dx$  provided that  $a < b$  are not atoms of  $\mu$ .

Firstly, we study the limiting behavior of the *empirical spectral distribution*  $\mu_N = \frac{1}{N} \sum_{\lambda \in \text{Sp } \Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)} \delta_\lambda$ . From the decomposition (4), we can already expect that it has the same behavior as that of  $\Phi_0(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  since the rank-3 perturbation will have a negligible weight. Yet, the entries of  $\Phi_0(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  show non-trivial dependencies because of the contractions on  $\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*$ . Nevertheless, these dependencies vanish asymptotically as precised in the following theorem.

**Theorem 1.** Assume Heuristic 1 is true. Let  $\mathbf{a}, \mathbf{b}, \mathbf{c}$  be unit-norm vectors in dimension  $n_1, n_2, n_3$  respectively. The empirical spectral distributions of  $\Phi_0(\mathbf{a}, \mathbf{b}, \mathbf{c})$  and  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  both converge weakly almost surely to the same probability distribution  $\bar{\mu}$  whose Stieltjes transform is  $\bar{m} = m_1 + m_2 + m_3$  with  $m_\ell$  ( $\ell \in \{1, 2, 3\}$ ) such that, for all  $z \in \mathbb{C} \setminus \mathbb{R}$ ,<sup>1</sup>

$$\varepsilon m_\ell(z)(\bar{m}(z) - m_\ell(z)) + z m_\ell(z) + c_\ell = 0. \quad (5)$$

The Stieltjes transform  $\bar{m}$  of the limiting spectral distribution  $\bar{\mu}$  is the sum of 3 Stieltjes transforms  $m_\ell$  which satisfy the fixed-point equation (5). In two different settings, Figure 1 shows good agreement between the empirical spectral distribution of  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  and the density of  $\bar{\mu}$  predicted by Theorem 1. If  $n_1 = n_2 = n_3$ , we see that  $\bar{\mu}$  corresponds to a semicircle distribution. This observation can be justified by the fact that, in this case,  $m_1 = m_2 = m_3$  therefore  $\bar{m}$  satisfies

$$\frac{2\varepsilon}{3} \bar{m}^2(z) + z \bar{m}(z) + 1 = 0, \quad z \in \mathbb{C} \setminus \mathbb{R},$$

which is the equation characterizing the Stieltjes transform of the semicircle distribution on  $[-2\sqrt{2\varepsilon/3}, +2\sqrt{2\varepsilon/3}]$ .

**Remark 3** (Universality). With the change of variables  $(z, m(z)) \rightsquigarrow (\sqrt{\varepsilon} \tilde{z}, \frac{1}{\sqrt{\varepsilon}} \tilde{m}(\tilde{z}))$  applied to Equation (5), we recover the relation given in [6] for the model without puncturing ( $\varepsilon = 1$ ). Hence, the system of equations satisfied by  $m_\ell(z)$  is the same as that of the model *without* mask  $\mathcal{B}$  but with an  $\mathcal{N}(0, \varepsilon)$ -noise instead of  $\mathcal{N}(0, 1)$ . This suggests that the limiting spectral distribution of  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  only depends on the first and second moments of the entries of  $\mathcal{T}$ .

#### 3.2 Asymptotic Spike Recovery

Relying on our previous spectral analysis of  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  and Proposition 1, we can now express the derivatives

1. Recall that  $c_\ell = n_\ell/N$ .

of  $\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*$  with respect to the entries of  $\mathcal{N}$  and describe the asymptotic behavior of  $\sigma^*$  and the alignments  $\langle \mathbf{x}, \mathbf{u}^* \rangle, \langle \mathbf{y}, \mathbf{v}^* \rangle, \langle \mathbf{z}, \mathbf{w}^* \rangle$  thanks to the relations found with Stein's lemma in Section 2.2.

**Theorem 2.** Assume Heuristic 1 is true. There exists a threshold  $\beta_s > 0$  such that, as  $N \rightarrow +\infty$ , if  $\beta > \beta_s$  then  $\sigma^* \rightarrow \sigma_\infty$  almost surely and

$$|\langle \mathbf{x}, \mathbf{u}^* \rangle| \xrightarrow{a.s.} q_1, \quad |\langle \mathbf{y}, \mathbf{v}^* \rangle| \xrightarrow{a.s.} q_2, \quad |\langle \mathbf{z}, \mathbf{w}^* \rangle| \xrightarrow{a.s.} q_3,$$

where, for  $\ell \in \{1, 2, 3\}$ ,  $q_\ell^2 = 1 - \frac{\varepsilon m_\ell^2(\sigma_\infty)}{c_\ell}$  and  $\sigma_\infty$  satisfies

$$\sigma_\infty + \varepsilon \bar{m}(\sigma_\infty) - \varepsilon \beta q_1 q_2 q_3 = 0. \quad (6)$$

Figure 2 depicts the values of  $\sigma_\infty$  and  $q_1, q_2, q_3$  against  $\beta$  in the same settings as Figure 1 (where  $\beta$  was fixed to 4). We observe a phase transition phenomenon: it is only above the threshold value  $\beta_s$  that  $\sigma_\infty$  is well-defined and  $q_1, q_2, q_3 > 0$ . As  $\beta$  crosses the threshold, the alignments suddenly jump from 0 to a high value, revealing the appearance of an informative critical point of Problem (1) as described in [4]. Then, as  $\beta$  increases, the alignments converge to 1 while  $\sigma_\infty$  approaches its asymptote  $\varepsilon\beta$  from above.

In the ‘‘cubic setting’’  $c_1 = c_2 = c_3$ , we have the closed-form expression  $\beta_s = \sqrt{\frac{d-1}{\varepsilon d} \left( \frac{d-1}{d-2} \right)^{\frac{d-2}{2}}}$  which we state for a generic tensor order  $d \geq 3$  to show that it coincides with the value found in [4] for a symmetric spike without puncturing. Moreover, we can also precise the value of the alignments right after the phase transition,  $\lim_{\beta \downarrow \beta_s} q_\ell = \sqrt{\frac{d-2}{d-1}}$ , which surprisingly does not depend  $\varepsilon$ . This shows that keeping less entries of  $\mathcal{T}$  (reducing  $\varepsilon$ ) only causes the phase transition to recede ( $\beta_s \propto \frac{1}{\sqrt{\varepsilon}}$ ) but does not change the alignments after  $\beta_s$ .

#### 3.3 Performance vs. Cost Trade-Off

Given a punctured tensor  $\mathcal{T} \odot \mathcal{B}$  with  $\mathcal{T} = \beta \mathbf{x} \otimes \mathbf{y} \otimes \mathbf{z} + \frac{1}{\sqrt{N}} \mathcal{N}$ , the reconstruction performance of  $\sigma^*, \mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*$  hinges upon the strength of the signal  $\beta$  and the portion of non-missing values  $\varepsilon$ . Although  $\beta$  is fixed by the problem,  $\varepsilon$  may be chosen by the user in a context where the data is intentionally punctured to store the tensor at a lower memory cost.

From a general perspective, applying the change of variables  $(z, m(z)) \rightsquigarrow (\sqrt{\varepsilon} \tilde{z}, \frac{1}{\sqrt{\varepsilon}} \tilde{m}(\tilde{z}))$  to Equation (6) shows that the reconstruction performance of the best rank-one approximation of  $\mathcal{T} \odot \mathcal{B}$  is the same as the performance achieved *without* puncturing ( $\varepsilon = 1$ ) but with a signal strength  $\tilde{\beta} = \sqrt{\varepsilon} \beta$  instead. This dilation by a factor  $\sqrt{\varepsilon}$  summarizes the impact of puncturing on the reconstruction performance. For small values of  $\varepsilon$  this may represent a significant loss of performance.

Figure 3 displays the empirical alignments computed on a  $50 \times 100 \times 350$  tensor  $\mathcal{T}$  with  $\beta = 2.5$  as a function of the puncturing level  $\varepsilon$ . Theoretical curves also reveal a phase transition at a threshold value of  $\varepsilon$  (here, approximately 0.17). Yet, it is only for values of  $\varepsilon \gtrsim 0.53$  that empirical alignments match the predictions of Theorem 2. This reveals a *computational-to-statistical gap*:  $0.17 \lesssim \varepsilon \lesssim 0.53$  corresponds to a *hard phase* where recovery is statistically possible but computationally too difficult. The problem becomes computationally feasible above the *algorithmic threshold* which can also be characterized with a random matrix analysis, as in [5].




Figure 1 – Empirical spectral distribution (ESD) and limiting spectral distribution (LSD) of  $\Phi(\mathbf{u}^*, \mathbf{v}^*, \mathbf{w}^*)$  with two different tensor shapes. **Parameters:**  $\beta = 4, \varepsilon = 0.25$ . The histogram on the **left** only counts the 599 (out of 1 000) *non-zero* eigenvalues.




Figure 2 – Asymptotic values of  $\sigma^*$  (**top**) and alignments  $|\langle \mathbf{x}, \mathbf{u}^* \rangle|, |\langle \mathbf{y}, \mathbf{v}^* \rangle|, |\langle \mathbf{z}, \mathbf{w}^* \rangle|$  (**bottom**) predicted by Theorem 2 against  $\beta$  with  $\varepsilon = 0.25$  and tensor shapes  $(c_1, c_2, c_3)$  corresponding to the two settings of Figure 1.




Figure 3 – Empirical alignments and their asymptotic values against  $\varepsilon$  with  $\beta = 2.5$  and tensor shape  $(\frac{1}{10}, \frac{2}{10}, \frac{7}{10})$ . Dots correspond to mean alignments averaged over 20 runs and the corresponding filled areas show the standard deviation.

## 4 Conclusion and Perspectives

Studying the performances of inference on a punctured tensor  $\mathcal{T} \odot \mathcal{B}$  is of interest in a context of missing data or memory limitations. Here, we study its best rank-one approximation and show that the behavior of the critical points of Problem (1) is linked to a random matrix  $\Phi$ . The spectral analysis of the latter gives access to the performance of reconstruction of a planted rank-one signal  $\beta \mathbf{x} \otimes \mathbf{y} \otimes \mathbf{z}$  in  $\mathcal{T}$ . In particular, the reconstruction performance from the punctured tensor  $\mathcal{T} \odot \mathcal{B}$  is the same as the performance achieved from  $\mathcal{T}$  directly but with a smaller signal strength  $\tilde{\beta} = \sqrt{\varepsilon} \beta$  instead. Although Theorem 2 describes the alignment with the signal of an informative optimum to Problem (1), in practice, recovering this optimum is hard and becomes feasible only if the signal strength is above an *algorithmic threshold*, whose theoretical behavior is also accessible with random matrix tools, see [5].

The analysis of punctured tensors is rich in exciting challenges. On the theoretical side, the rigorous justification of Heuristic 1 requires tools to untangle the dependence of solutions to Problem (1) on the entries of  $\mathcal{B}$ . Furthermore, in order to face strong memory limitations, more clever data-dependent puncturing approaches should be considered. Although their theoretical study is expected to be much more delicate than that of random puncturing (considered here), relevant policies such as *sparsification* via hard thresholding or even *quantization* of the entries may prove to be much more efficient.

## References

- [1] E. Acar, D. M. Dunlavy, T. G. Kolda, and M. Mørup. Scalable tensor factorizations for incomplete data. *Chemometrics and Intelligent Laboratory Systems*, 106(1):41–56, March 2011.
- [2] R. Couillet, F. Chatelain, and N. Le Bihan. Two-way kernel matrix puncturing: towards resource-efficient PCA and spectral clustering. In *Proceedings of the 38th International Conference on Machine Learning*, pages 2156–2165. PMLR, July 2021.
- [3] R. Couillet and Z. Liao. *Random Matrix Methods for Machine Learning*. Cambridge University Press, Cambridge, 2022.
- [4] A. Jagannath, P. Lopatto, and L. Miolane. Statistical thresholds for tensor PCA. *The Annals of Applied Probability*, 30(4):1910–1933, August 2020.
- [5] H. Lebeau, F. Chatelain, and R. Couillet. A Random Matrix Approach to Low-Multilinear-Rank Tensor Approximation. *Journal of Machine Learning Research*, 26(7):1–64, 2025.
- [6] M. E. A. Seddik, M. Guillaud, and R. Couillet. When Random Tensors meet Random Matrices, November 2022. arXiv:2112.12348 [math, stat].
- [7] N. D. Sidiropoulos, L. De Lathauwer, X. Fu, K. Huang, E. E. Papalexakis, and C. Faloutsos. Tensor Decomposition for Signal Processing and Machine Learning. *IEEE Transactions on Signal Processing*, 65(13):3551–3582, July 2017.
- [8] N. Vervliet, O. Debals, L. Sorber, and L. De Lathauwer. Breaking the Curse of Dimensionality Using Decompositions of Incomplete Tensors: Tensor-based scientific computing in big data analysis. *IEEE Signal Processing Magazine*, 31(5):71–79, September 2014.