

A Random Matrix Perspective on Random Tensors

José Henrique de Morais Goulart

HENRIQUE.GOULART@IRIT.FR

*Université de Toulouse, Toulouse INP, IRIT
31071 Toulouse, France*

Romain Couillet

ROMAIN.COUILLET@UNIV-GRENOBLE-ALPES.FR

*Université Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG
38000 Grenoble, France*

Pierre Comon

PIERRE.COMON@GIPSA-LAB.GRENOBLE-INP.FR

*Université Grenoble Alpes, CNRS, Grenoble INP, GIPSA-lab
38000 Grenoble, France*

Editor: Genevera Allen

Abstract

Several machine learning problems such as latent variable model learning and community detection can be addressed by estimating a low-rank signal from a noisy tensor. Despite recent substantial progress on the fundamental limits of the corresponding estimators in the large-dimensional setting, some of the most significant results are based on spin glass theory, which is not easily accessible to non-experts. We propose a sharply distinct and more elementary approach, relying on tools from random matrix theory. The key idea is to study random matrices arising from contractions of a random tensor, which give access to its spectral properties. In particular, for a symmetric d th-order rank-one model with Gaussian noise, our approach yields a novel characterization of maximum likelihood (ML) estimation performance in terms of a fixed-point equation valid in the regime where weak recovery is possible. For $d = 3$, the solution to this equation matches the existing results. We conjecture that the same holds for any order d , based on numerical evidence for $d \in \{4, 5\}$. Moreover, our analysis illuminates certain properties of the large-dimensional ML landscape. Our approach can be extended to other models, including asymmetric and non-Gaussian ones.

Keywords: Random tensors, random matrix theory, spiked tensor model

1. Introduction

Discovering and exploiting underlying low-dimensional structures in data are key to the success of modern machine learning and signal processing techniques. In many applications, prior knowledge on the process or generative model that produces the observations leads to a low-rank tensor model, which postulates that the sought information takes the form of a tensor decomposition having a few algebraically simple terms—examples can be found in (Anandkumar et al., 2014; Sidiropoulos et al., 2017; Mørup, 2011; Landsberg, 2012). While this approach often yields satisfying results thanks to the the strong uniqueness properties enjoyed by such models (Sidiropoulos et al., 2017), which ensure their identifiability (and therefore their interpretability), the performance of tensor methods is quite difficult to anticipate and very few of them—which are specialized to restrictive models of rather narrow

applicability—are actually accompanied by performance guarantees (Ge et al., 2015; Jain and Oh, 2014; Huang et al., 2015; Anandkumar et al., 2014). This difficulty stems in large part from the notorious and significant conceptual challenges that accompany tensor models, as many results and notions pertaining to matrices do not readily generalize to higher-order tensors.

Recent years have seen some important progress in understanding the properties of low-rank tensor models of the form $\mathbf{Y} = \mathbf{X} + \mathbf{W}$, with \mathbf{W} standing for additive noise and \mathbf{X} having the so-called canonical polyadic decomposition (CPD) form (Hitchcock, 1927), (Landsberg, 2012)

$$\mathbf{x} = \sum_{r=1}^R \lambda_r a_r^{(1)} \otimes \cdots \otimes a_r^{(d)},$$

where the natural number R defines the rank of the model, $a_r^{(i)} \in \mathbb{R}^{N_i}$ and λ_r is a positive number for $i = 1, \dots, d$ and $r = 1, \dots, R$. In particular, the performance of low-rank CPD estimators under various priors for the vectors $a_r^{(i)}$ is now better understood thanks to the efforts of several authors (Montanari and Richard, 2014; Perry et al., 2020; Lesieur et al., 2017; Chen et al., 2021; Jagannath et al., 2020). This progress has been achieved by adopting a large-dimensional inferential statistics viewpoint, which amounts to studying the inference of the random variables $\lambda_r, a_r^{(i)}$ from noisy observations (distributed according to some known law), under the crucial assumption that the dimensions grow large, that is, $N_i \rightarrow \infty$. Just like happens for random matrices, the latter assumption allows exploiting certain concentration-of-measure phenomena and leads to quite precise results, some of which are briefly reviewed below.

Despite these recent results bringing a substantial advancement in the understanding of random tensor models, they are still limited in scope and have in large part been derived by heavily relying on spin glass theory, thus being quite difficult to understand for non-expert readers. Also, their extension to more structured (and more relevant for applications) models is likely difficult with this approach.

We take a completely different path in this work, relying upon years of advances in the field of random matrix theory (RMT) in order to study random tensor models. The key idea lies in the study of *contractions* of a given tensor model with unit norm vectors, which give rise to a set of random matrices. In particular, we show that with a judicious choice of such contractions one gains access to spectral properties of a random tensor in the sense introduced by Lim (2005), which in turn allow a characterization of the performance of the maximum likelihood (ML) estimator of the rank-one symmetric spiked tensor model. For this model, our approach produces precise predictions matching recent results that had been previously obtained only through sophisticated ideas borrowed from spin glass theory by Jagannath et al. (2020). Our findings are obtained by invoking the Gaussian tools developed by Pastur and Shcherbina (2011), and can be extended to other models, notably with non-Gaussian noise. This approach amounts to a significant methodological contribution, as it opens the door to the investigation of other, more general tensor models by means of random matrix tools, on which a vast and mature body of work has been established in recent decades (Pastur and Shcherbina, 2011; Bai and Silverstein, 2010; Tao, 2012).

As such, the present work follows after the latest breakthroughs in the understanding, analysis and improvement of machine learning methods based on spiked large-dimensional random matrix models. Recent examples in data classification include the large-dimensional analysis of support vector machines (Liao and Couillet, 2019; Kammoun and Alouini, 2021), semi-supervised and transfer learning (Mai and Couillet, 2018, 2021; Tiomoko et al., 2020), logistic regression (Mai et al., 2019), as well as neural network dynamics (Advani et al., 2020; Liao and Couillet, 2018). Similar ideas in data and graph clustering are found in (Couillet and Benaych-Georges, 2016; Nadakuditi and Newman, 2012).

1.1 Statistical performance of rank-one spike estimation

Our approach will be leveraged to study ML estimation of a planted signal vector x in a rank-one symmetric model of the form

$$\mathbf{y} = \lambda x^{\otimes d} + \frac{1}{\sqrt{N}} \mathbf{W}, \quad (1)$$

with x uniformly distributed on the unit sphere \mathbb{S}^{N-1} (which is known as a spherical prior) and \mathbf{W} a symmetric Gaussian noise tensor (see Section 2.2 for a precise definition). This is the so-called *spiked rank-one tensor model*, whose “signal” part is referred to as a spike, with the positive number λ effectively playing the role of a signal-to-noise ratio parameter.

Similarly to other high-dimensional statistical inference problems, the fundamental asymptotic performance limits of estimators of x are of utmost importance in this context. Given some estimator $\hat{x}(\mathbf{y})$ taking values on \mathbb{S}^{N-1} , a natural performance measure is the mean *alignment* (or overlap), defined as

$$\bar{\alpha}_{d,N}(\lambda) := \mathbb{E} \{ |\langle x, \hat{x}(\mathbf{y}) \rangle| \} \in [0, 1].$$

Two crucial questions regarding the behavior of this quantity as $N \rightarrow \infty$ are as follows:

1. (*Weak recovery*) For which range of λ is there an estimator whose mean alignment satisfies $\limsup_{N \rightarrow \infty} \bar{\alpha}_{d,N}(\lambda) > 0$?
2. (*Best asymptotic alignment*) What is the largest value of $\limsup_{N \rightarrow \infty} \bar{\alpha}_{d,N}(\lambda)$ attainable by some estimator for a given value of λ ?

A clear understanding of these questions has been achieved in recent years. As it turns out, there exists an $O(1)$ (in N) statistical threshold $\lambda_c(d)$ (which can be numerically computed) such that weak recovery is impossible when $\lambda < \lambda_c(d)$, but is possible when $\lambda > \lambda_c(d)$. Furthermore, for any λ the ML estimator of x , which boils down to computing the best symmetric rank-one approximation of \mathbf{y} or, equivalently, to solving

$$\max_{u \in \mathbb{S}^{N-1}} \langle \mathbf{y}, u^{\otimes d} \rangle = \max_{u \in \mathbb{S}^{N-1}} \sum_{i_1=1}^N \cdots \sum_{i_d=1}^N Y_{i_1 \dots i_d} u_{i_1} \cdots u_{i_d}, \quad (2)$$

attains the highest possible correlation among all measurable estimators of x (Jagannath et al., 2020).

Before summarizing our own contributions, we will briefly review existing studies on this topic, and more generally on the estimation (and detection) of spiked tensor models with other priors.

1.2 Related work

The first study that seems to have looked into the spiked tensor model (1) was carried out by Montanari and Richard (2014), who showed that weak recovery can be performed by ML estimation whenever λ exceeds a number $\mu_0(d)$, defined as the almost sure limit of $\frac{1}{\sqrt{N}} \|\mathbf{W}\|$, where $\|\mathbf{W}\|$ stands for the spectral norm of \mathbf{W} (see Section 2.1 for precise definitions). Then, Montanari et al. (2015) investigated the closely related problem of detecting the presence of a rank-one symmetric spike given a realization of \mathbf{Y} , which amounts to distinguishing the hypothesis $\lambda = 0$ from $\lambda \neq 0$. By using the second-moment method, they proved that the law of \mathbf{Y} is asymptotically indistinguishable¹ from that of $\frac{1}{\sqrt{N}}\mathbf{W}$ when λ is smaller than a certain value $\lambda_{2\text{nd}}(d) < \mu_0(d)$ (for which a variational characterization was given), and thus no hypothesis test can possibly perform this task better than pure random guessing. An extension of this result to a rank- R spiked tensor was later derived by Chevreuil and Loubaton (2018).

Perry et al. (2020) observed that detection is also possible (asymptotically) when $\lambda > \mu_0(d)$, since $\mathbb{P}\{\|\mathbf{Y}\| > \mu_0(d)\} = 1 - o(1)$ holds in this regime. Interestingly, they found that this already holds for λ larger than a certain bound strictly smaller than $\mu_0(d)$, beyond which the spike can thus be detected, and also (weakly) estimated. This effect is reminiscent of the so-called BBP transition of spiked random matrix models, named after the authors that first exposed it (Baik et al., 2005); see also (Benaych-Georges and Nadakuditi, 2011) for a quite general treatment of this phenomenon. By developing a modified second-order method, Perry et al. (2020) also found that detection and weak estimation are impossible when λ is below some threshold $\lambda'_{2\text{nd}}(d) > \lambda_{2\text{nd}}(d)$, thus also improving upon the previously existing lower bound on the statistical threshold. Moreover, they derived similar results for Rademacher and sparse Rademacher priors on x . Borrowing ideas from spin glass theory, Chen (2019) refined the results of Perry et al. (2020) for the Rademacher prior, computing a single critical value that separates the regimes where the rank-one spike is detectable (in the sense that the total variation distance between the laws of \mathbf{Y} and of $\frac{1}{\sqrt{N}}\mathbf{W}$ converges to 1) and can be weakly recovered, and where \mathbf{Y} and $\frac{1}{\sqrt{N}}\mathbf{W}$ are statistically indistinguishable (with the aforementioned distance converging to zero) and the spike cannot be weakly recovered.

Also on the basis of results pertaining to mean-field spin glass models, Chen et al. (2021) derived a variational expression for the threshold that applies to both detection and estimation of a rank-one symmetric spiked tensor model with Gaussian noise. However, their model differs from (1) in that the vector x has i.i.d. elements distributed according to some prior. Specifically, the estimation threshold is defined in that work as the value beyond which the minimum mean-squared error (MMSE) estimator performs strictly better (in the MSE sense) than any so-called “dummy” estimator (that does not depend on the data), and below which it performs as well as a dummy estimator. They went on to extend this result to the case with multiple spikes (that is, higher rank) with possibly different priors: essentially, the signal-to-noise parameter λ_r of at least one spike must exceed the threshold associated with its prior, or else neither detection nor estimation can possibly be performed.

1. More precisely, the total variation distance between these laws converges to zero as $N \rightarrow \infty$.

For the rank-one symmetric setting of (1) with spherical prior and Gaussian noise, Jagannath et al. (2020) finally settled the search for a precise expression for the phase transition threshold $\lambda_c(d)$ that separates the regime where one can detect and weakly estimate x ($\lambda > \lambda_c(d)$) from the one where these tasks are impossible ($\lambda < \lambda_c(d)$), again by relying substantially upon results coming from spin glass theory. They showed that this transition is much more abrupt for tensors of order $d > 2$ than for matrices, by giving also an exact asymptotic expression for the alignment of x with the ML solution $\hat{x}_{\text{ML}}(\mathbf{Y})$, namely $|\langle \hat{x}_{\text{ML}}(\mathbf{Y}), x \rangle|$, which displays a jump discontinuity at $\lambda_c(d)$ (see Section 3.2 for a precise description of that behavior, shown in Figure 3). Their analysis also produced an exact expression for the spectral norm of \mathbf{Y} , which is very much in line with the BPP-like phenomenon reported by Perry et al. (2020), and showed that the ML estimator is information-theoretically optimal, in the sense that its alignment with x matches that of the Bayes-optimal estimator.

Another interesting related work has been carried out by Ben Arous et al. (2019), who studied the complexity of the maximum likelihood landscape by using the Kac-Rice formula to compute the expected number of critical points and of local maxima of (2). Among several other conclusions, their study shows that for $\lambda < \lambda_c(d)$ the objective function values of all local maxima (including the global one) tend to be concentrated (as $N \rightarrow \infty$) on a quite narrow interval, whereas for $\lambda > \lambda_c(d)$ the value achieved by the global maximum “escapes” that interval and grows with λ .

1.3 Summary of contributions

Our main contribution is of a methodological flavor, and is based on studying random matrices obtained by *contractions of a random tensor model*. The fruitfulness of this idea is shown in the case of the rank-one symmetric spiked tensor model \mathbf{Y} given by (1). In this setting, the ML estimator $\hat{x}_{\text{ML}}(\mathbf{Y})$ of the planted vector x is shown to be the dominant eigenvector of the matrix arising from the $(d - 2)$ -fold contraction of \mathbf{Y} with $\hat{x}_{\text{ML}}(\mathbf{Y})$ itself. This is a simple consequence of the tensor eigenvalue equations (see Section 3.1), which by definition characterize the critical points of (2). By pursuing this connection, and under the assumption that there exists a sequence of critical points (μ, u) of (2) such that $|\langle x, u \rangle|$ and μ converge almost surely to deterministic functions of λ , we are able to deploy tools from matrix theory in order to establish a fixed-point characterization of these deterministic quantities. As we will discuss, it turns out that for $d \in \{3, 4, 5\}$ the solution to our fixed-point equation precisely matches the results of Jagannath et al. (2020) describing the behavior of the ML estimator (and thus of global maxima of (2)) beyond the critical value $\lambda_c(d)$. We conjecture that this actually holds for any $d \geq 3$.

In the process of deriving these results, we demonstrate that the spectral measure of contractions of \mathbf{Y} with sequences of independent vectors or of local maximizers of the ML problem (satisfying a mild technical condition) converges weakly almost surely to a semi-circle law supported on $[-\beta_d, \beta_d]$, where $\beta_d = 2/\sqrt{d(d - 1)}$. This sheds light on one feature of the maximum likelihood landscape that had been exposed by Ben Arous et al. (2019), namely, the fact that the value of the objective function of (2) at every local maxima asymptotically lies on the interval $[2\beta_3, \|\mathbf{Y}\|]$ when $d = 3$. It turns out that every such a maximum is the largest eigenvalue of a contraction of \mathbf{Y} , and the second-order necessary




Figure 1: Tensor power iteration method applied to find a local maximum of the ML problem (2), with $d = 3$ and $N = 500$. This algorithm produces a sequence $v^{(k)}$ which (under certain conditions) converges to a tensor eigenvector u associated with a local maximum of (2) (Kolda and Mayo, 2011).

optimality conditions combined with the derived spectral measure of such contractions entail the lower bound $(d - 1)\beta_d = 2\beta_3$, while the upper bound follows by definition of $\|\mathbf{Y}\|$.

Our analysis also offers another perspective into the behavior of the tensor power iteration method (Kolda and Mayo, 2011), which can be seen to produce a sequence of *matrices*² $\mathbf{Y} \cdot (v^{(k)})^{d-2}$ converging (under appropriate conditions) to a contraction of \mathbf{Y} with one of its eigenvectors u (in the sense introduced by Lim (2005)), as $v^{(k)} \rightarrow u$. In this interpretation, each iteration is tantamount to a single step of a *matrix* power iteration method applied to $\mathbf{Y} \cdot (v^{(k)})^{d-2}$, which then changes in the next iteration. In particular, when this algorithm is applied to a spiked tensor model of the form (1), the spectra of such matrices $\mathbf{Y} \cdot (v^{(k)})^{d-2}$ exhibit a “bulk” obeying a semi-circle law and a spike that moves away from that bulk until it reaches the position of an eigenvalue associated with a local maximum (see Figure 1 for an illustration), which asymptotically must be at least as large as $(d - 1)\beta_d$. The presence of this spike in $\mathbf{Y} \cdot (v^{(k)})^{d-2}$ is observed regardless of the value of λ , and is in fact a consequence of the dependence between \mathbf{Y} and $v^{(k)}$.

Finally, the high degree of versatility of our approach should be emphasized, as it can be extended to other, less stringent models with other noise distributions and asymmetric structure. We also note that our approach based on contractions with suitably chosen vectors fundamentally differs from previous spectral techniques which relied on matrix unfoldings of the given tensor model, such as those of Montanari and Richard (2014) and Ben Arous et al. (2021). See Section 2.3 for a discussion.

We point out that this work was first presented at the seminar in Random Tensors of Texas A&M University (Goulart et al., 2021a) and at the 2021 SIAM Conference on Applied Linear Algebra (Goulart et al., 2021b).

2. Here, the notation $\mathbf{Y} \cdot (v^{(k)})^{d-2}$ stands for the $(d - 2)$ -fold contraction of \mathbf{Y} with the vector $v^{(k)}$ (see Section 2.1 for our notational conventions).

2. Preliminaries and main results

We will start by introducing some notation and defining the objects of interest in our study, before describing our main results. For simplicity, in the following we shall stick to symmetric random tensors, but the main ideas generalize to the asymmetric case.

2.1 Tensors and contractions

The entries of a d -th order tensor $\mathbf{Y} \in \mathbb{R}^{N \times \dots \times N}$ will be denoted by uppercase letters, as in $Y_{i_1 \dots i_d}$. The same convention applies to matrices ($d = 2$) but not to vectors, whose components are written in lowercase letters, as in u_i . A tensor is said to be symmetric if and only if its entries are invariant with respect to any permutation of the indices, that is, $Y_{i_1 \dots i_d} = Y_{j_1 \dots j_d}$, where (j_1, \dots, j_d) is an arbitrary permutation of (i_1, \dots, i_d) . The space of d -th order real-valued symmetric tensors of dimensions $N \times \dots \times N$ will be denoted $\mathcal{S}^d(N)$.

Given a tensor $\mathbf{Y} \in \mathcal{S}^d(N)$ and vectors $u^{(1)}, \dots, u^{(p)} \in \mathbb{R}^N$ with $p \leq d$, we define the contraction $\mathbf{Y} \cdot (u^{(1)}, \dots, u^{(p)})$ as the $(d - p)$ th-order tensor with components

$$\left(\mathbf{Y} \cdot (u^{(1)}, \dots, u^{(p)}) \right)_{i_{p+1}, \dots, i_d} = \sum_{i_1=1}^N \dots \sum_{i_p=1}^N Y_{i_1 \dots i_d} u_{i_1}^{(1)} \dots u_{i_p}^{(p)}.$$

The order of the vectors in $(u^{(1)}, \dots, u^{(p)})$ is immaterial, due to the symmetry of \mathbf{Y} . Moreover, the outcome is clearly symmetric as well, by symmetry of \mathbf{Y} . In particular, when $p = d - 2$, $p = d - 1$ and $p = d$ we get a symmetric $N \times N$ matrix, an N -dimensional vector and a scalar, respectively. A special case of interest to us is when $u^{(j)} = u$ for all $j \in \{1, \dots, p\}$, for which we introduce the notation

$$\mathbf{Y} \cdot u^p := \mathbf{Y} \cdot \underbrace{(u, \dots, u)}_{p \text{ times}}.$$

In a similar vein, the multilinear transformation of $\mathbf{Y} \in \mathcal{S}^d(N)$ by a square matrix $U \in \mathbb{R}^{N \times N}$ is defined as

$$\mathbf{Y}' = \mathbf{Y} \cdot U^d \Leftrightarrow Y'_{i_1 \dots i_d} = \sum_{n_1=1}^N \dots \sum_{n_d=1}^N Y_{n_1 \dots n_d} U_{i_1 n_1} \dots U_{i_d n_d}.$$

When U is nonsingular, this operation amounts to a change of basis in $\mathcal{S}^d(N)$.

A d -th order tensor $\mathbf{X} \in \mathbb{R}^{N \times \dots \times N}$ is said to be of rank one if and only there exist vectors $u^{(1)}, \dots, u^{(d)} \in \mathbb{R}^N$ such that the entries of \mathbf{X} decompose as $X_{i_1 \dots i_d} = u_{i_1}^{(1)} \dots u_{i_d}^{(d)}$. This relation can be expressed in terms of the tensor product \otimes as follows:

$$\mathbf{X} = u^{(1)} \otimes \dots \otimes u^{(d)}.$$

In particular, a symmetric rank-one tensor can be written in this form with $u^{(i)} = u$ for all i , in which case we denote it as $u^{\otimes d}$. A tensor is said to have rank R if it can be written as a sum of R rank-one terms, but not fewer. Similarly, a symmetric tensor has symmetric rank S if and only if it can be written as a sum of S symmetric rank-one terms, but not fewer.

By equipping $\mathbb{R}^{N \times \dots \times N}$ with the Euclidean scalar product (Hackbusch, 2012)

$$\langle \mathbf{y}, \mathbf{x} \rangle := \sum_{i_1=1}^N \cdots \sum_{i_d=1}^N Y_{i_1 \dots i_d} X_{i_1 \dots i_d},$$

we have the identity $\langle \mathbf{y}, u^{\otimes d} \rangle = \mathbf{y} \cdot u^d$ for any symmetric rank-one d th-order tensor $u^{\otimes d}$, and also $\langle u^{\otimes d}, v^{\otimes d} \rangle = \langle u, v \rangle^d$ for any two such tensors. Naturally, the Euclidean or Frobenius norm on $\mathbb{R}^{N \times \dots \times N}$ is defined as $\|\mathbf{y}\|_{\mathbb{F}} := \sqrt{\langle \mathbf{y}, \mathbf{y} \rangle}$.

2.2 Random tensors and the Gaussian orthogonal tensor ensemble

A random tensor \mathbf{Y} is simply a tensor-valued random variable with an associated probability measure $\rho_{\mathbf{y}}$.

In particular, the Gaussian orthogonal matrix ensemble can be readily generalized to tensors in $\mathcal{S}^d(N)$ by defining the density

$$f(\mathbf{W}) = \frac{1}{Z_d(N)} e^{-\frac{1}{2}\|\mathbf{W}\|_{\mathbb{F}}^2}, \quad (3)$$

where $Z_d(N)$ is a normalization constant. Just as in the matrix case, because f depends on \mathbf{W} only through its Euclidean norm, it is invariant with respect to an orthogonal change of basis of $\mathcal{S}^d(N)$, as

$$\forall U \in \text{O}(N), \quad \|\mathbf{W} \cdot U^d\|_{\mathbb{F}} = \|\mathbf{W}\|_{\mathbb{F}},$$

where $\text{O}(N)$ denotes the orthogonal group. Another consequence of the definition (3) is that, by symmetry of \mathbf{W} , the variance of an entry $W_{i_1 \dots i_d}$ equals the reciprocal of the number of distinct permutations of its indices (i_1, \dots, i_d) . For instance, in the $d = 3$ case one can expand $\|\mathbf{W}\|_{\mathbb{F}}^2$ as

$$\|\mathbf{W}\|_{\mathbb{F}}^2 = \sum_i W_{iii}^2 + 3 \sum_{i \neq j} W_{ijj}^2 + 6 \sum_{i < j < k} W_{ijk}^2,$$

which shows that on-diagonal entries of \mathbf{W} have variance 1, entries with two distinct indices have variance 1/3 and entries with three distinct indices have variance 1/6:

$$\sigma_{W_{ijk}}^2 = \frac{1}{6} + \delta_{ij} \frac{1}{6} + \delta_{ik} \frac{1}{6} + \delta_{jk} \frac{1}{6} + \delta_{ij} \delta_{jk} \frac{1}{3},$$

where δ_{ij} equals 1 if $i = j$ and 0 otherwise.

2.3 The contraction ensemble of a random tensor

A commonly used expedient when dealing with tensors is to consider matrix slices or matrix unfoldings. A slice is obtained by fixing all but two indices of a tensor, which yields a matrix. For instance, one can define the map $S_{\mathbf{y}} : \{1, \dots, N\}^{d-2} \rightarrow \mathbb{R}^{N \times N}$ by requiring that the components of the matrix $S_{\mathbf{y}}(n_1, \dots, n_{d-2})$ satisfy

$$(S_{\mathbf{y}}(n_1, \dots, n_{d-2}))_{i,j} = Y_{n_1, \dots, n_{d-2}, i, j}.$$

Note that the order of the indices on the right-hand side (and of the arguments as well) is immaterial, as \mathbf{Y} is symmetric. Hence, this map produces an object of dimension lower than that of \mathbf{Y} . By contrast, an unfolding (also known as flattening or matricization) of a tensor accommodates all its entries inside a matrix by partitioning its set of indices into two subsets, which are then bijectively mapped into row and column indices. These two definitions are very useful as they enable the application of matrix-analytic concepts and tools to tensors, and have long been studied in algebraic geometry (Landsberg, 2012) and utilized in various scientific domains for data analysis purposes (Tucker, 1966; Bro et al., 1997). In particular, a spectral method for the estimation of the spike in the rank-one model (1) based on unfoldings has been proposed by Montanari and Richard (2014), and later studied by Ben Arous et al. (2021).

When it comes to the study of random tensors, one is tempted to resort to these definitions in order to borrow tools from the well-developed RMT. However, merely taking slices or unfoldings of a tensor turns out to be insufficient for our purposes, as these matrices cannot describe the whole picture and in particular do not provide access to the *spectral* properties of the tensor (for us here its ℓ_2 -eigenpairs, which are connected to the ML estimation problem as we will shortly discuss). Nevertheless, realizing that the above defined map $S_{\mathbf{Y}}$ amounts to the contraction of \mathbf{Y} with unit norm vectors

$$S_{\mathbf{Y}}(n_1, \dots, n_{d-2}) = \mathbf{Y} \cdot (e^{(n_1)}, \dots, e^{(n_{d-2})}), \quad \text{where } e_i^{(n)} = \delta_{in},$$

one can more generally consider arbitrary contractions of a random tensor with unit-norm vectors. For our purposes, though, it will be sufficient to restrict these contractions to the case where all vectors are the same. This motivates the following definition.

Definition 1 (Contraction ensemble) *Let $\mathbf{Y} \in \mathcal{S}^d(N)$ be a random tensor. We define the contraction ensemble of \mathbf{Y} as the set of random matrices*

$$\mathcal{M}(\mathbf{Y}) := \left\{ \mathbf{Y} \cdot v^{d-2} : v \in \mathbb{S}^{N-1} \right\}.$$

As will be shown ahead in Section 3, by studying appropriately chosen random matrices from this set one gains access to spectral properties of \mathbf{Y} .

2.4 Main results

We are now ready to describe our main results.

2.4.1 SPECTRAL MEASURE OF CONTRACTIONS OF THE GAUSSIAN ORTHOGONAL TENSOR ENSEMBLE

Our first result concerns the spectral measure of random matrices from the contraction ensemble $\mathcal{M}\left(\frac{1}{\sqrt{N}}\mathbf{W}\right)$ of a random tensor \mathbf{W} distributed according to (3). This result is of its own interest, as taking contractions of a Gaussian symmetric tensor with independent entries (apart from the symmetries) is a natural way of introducing dependencies in a

random matrix model.³ Furthermore, it plays a major role in our analysis of the ML problem (2). Its statement is as follows.

Theorem 2 *Let $\mathcal{W} \in \mathcal{S}^d(N)$ be a sequence of random tensors with density given by (3), with $d \geq 3$. For any deterministic sequence⁴ of vectors $v \in \mathbb{S}^{N-1}$, the empirical spectral measure of $\frac{1}{\sqrt{N}}\mathcal{W} \cdot v^{d-2}$ converges weakly almost surely to the semi-circle distribution whose Stieltjes transform reads*

$$m_d(z) = \frac{2}{\beta_d^2} \left(-z + z \sqrt{1 - \frac{\beta_d^2}{z^2}} \right),$$

with $\beta_d := 2/\sqrt{d(d-1)}$. Its density is thus given by

$$\rho(dx) = \frac{2}{\pi \beta_d^2} \sqrt{(\beta_d^2 - x^2)^+} dx, \tag{4}$$

and is supported on $[-\beta_d, \beta_d]$.

The proof of this result can be essentially obtained as a byproduct of that of Theorem 3. Concretely, it will turn out that the limiting spectral measure of the contractions $\mathcal{W} \cdot v^{d-2}$ with deterministic $v \in \mathbb{S}^{N-1}$ is exactly the same as that of $\mathcal{Y} \cdot u^{d-2}$, where the vectors u form a sequence of critical points u of the ML problem satisfying certain conditions explained ahead. We comment on how the proof of Theorem 2 can be derived from that of Theorem 3 in Section 4.4.4.

Though the statement of Theorem 2 might at first glance seem trivial since $\mathcal{W} \cdot v^{d-2}$ is a linear combination of Gaussian matrices, these matrices are *not* independent due to the symmetries in \mathcal{W} . Nevertheless, at the end we show that, just as for standard Wigner matrices, the limiting Stieltjes transform also satisfies a quadratic equation, as the additional terms due to the dependence structure all vanish asymptotically in the derivation. An illustration of the derived spectral measure is shown on Figure 2.

One heuristic argument to see why the size of the support should scale like $O(1/d)$ is as follows. Take every vector v in the considered sequence to be the first canonical vector of the standard basis, denoted by $e^{(1)}$. Then, the contraction $\frac{1}{\sqrt{N}}\mathcal{W} \cdot v^{d-2} = \frac{1}{\sqrt{N}}\mathcal{W} \cdot (e^{(1)})^{d-2}$ yields a matrix C with entries

$$C_{ij} = \frac{1}{\sqrt{N}} W_{1\dots 1ij}.$$

It follows from the definition of \mathcal{W} that C is by construction a Gaussian Wigner matrix whose entries have variance equal to the reciprocal of the number of distinct permutations

3. Indeed, according to the authors of (Au and Garza-Vargas, 2021), the first draft of our work motivated them to study the spectral asymptotics of contractions of tensors of a Wigner type. They argued that existing methods for the study of random matrices with dependent entries do not apply in this case, and then proposed a graph formalism for dealing with it on the basis of combinatorial arguments. In its turn, our proof is based on analytic tools and focused on Gaussian tensors, but can be extended to other distributions by means of the tools from Pastur and Shcherbina (2011).

4. To keep the notation simple, we adopt the usual practice of omitting the index of the sequences of vectors v and of random tensors \mathcal{W} . This convention is adopted throughout the paper.




Figure 2: Density of the spectral measure of Theorem 2 (in red) for $d = 3, 4$ and histogram of eigenvalues for one realization of $\frac{1}{\sqrt{N}} \mathbf{W} \cdot v$, with a random vector v drawn uniformly on \mathbb{S}^{N-1} and independently of \mathbf{W} .

of $(1, \dots, 1, i, j)$. With the exception of the first row and first column of C , the variances are therefore $\frac{1}{d(d-1)}$ off-diagonal and $\frac{2}{d(d-1)}$ on the main diagonal. Hence, apart from the first row and column (whose contribution vanishes as $N \rightarrow \infty$), we see that C equals a standard GOE matrix (with variance 1 off-diagonal and 2 on the diagonal) multiplied by $\frac{1}{\sqrt{d(d-1)}}$, and hence its limiting spectrum should be a semicircle with support⁵ $[-\beta_d, \beta_d]$.

2.4.2 FIXED-POINT CHARACTERIZATION OF THE ML ESTIMATOR

We now state our main result concerning the performance of the ML estimator of model (1). In order to derive it, we study a sequence of contractions $\mathbf{y} \cdot u^{d-2}$, where the vectors u are critical points of the ML problem satisfying some conditions discussed below. The result is as follows.

Theorem 3 *Let $d \geq 3$, fix $\lambda > 0$ and suppose that there exists a sequence of critical points u of the maximum likelihood problem (2) such that*

$$\langle x, u \rangle \xrightarrow{a.s.} \alpha_{d,\infty}(\lambda), \quad (5)$$

$$\mathbf{y} \cdot u^d \xrightarrow{a.s.} \mu_{d,\infty}(\lambda) \quad (6)$$

for some deterministic functions $\alpha_{d,\infty}(\lambda)$ and $\mu_{d,\infty}(\lambda)$ defined on λ . Furthermore, assume that the following inequalities hold: $\alpha_{d,\infty}(\lambda) > 0$ and $\mu_{d,\infty}(\lambda) > (d-1)\beta_d$. Then, $\mu_{d,\infty}(\lambda)$ satisfies the fixed-point equation $\mu_{d,\infty}(\lambda) = \phi_d(\mu_{d,\infty}(\lambda), \lambda)$, where

$$\phi_d(z, \lambda) := \lambda \omega_d^d(z, \lambda) - \frac{1}{d-1} m_d \left(\frac{z}{d-1} \right), \quad (7)$$

$$\omega_d(z, \lambda) := \left[\frac{1}{\lambda} \left(z + \frac{1}{d} m_d \left(\frac{z}{d-1} \right) \right) \right]^{\frac{1}{d-2}}. \quad (8)$$

5. This is the same result found by Au and Garza-Vargas (2021) for Wigner-type tensors when all $d-2$ contractions are with the same vector. By contrast, when the contractions are allowed to involve $d-2$ distinct vectors, the support size can shrink down to $O(1/\sqrt{d!})$ —this can be understood along the lines of our heuristic argument by considering contractions with the canonical basis vectors $e^{(1)}, \dots, e^{(d-2)}$.

Furthermore, $\alpha_{d,\infty}(\lambda) = \omega_d(\mu_{d,\infty}(\lambda), \lambda)$.

For $d = 3$, we are able to verify that this fixed-point equation leads to the same exact expressions that follow from the study of Jagannath et al. (2020) for the asymptotic spectral norm of \mathbf{Y} (given by $\mu_{d,\infty}$) and the alignment $|\langle x, \hat{x}_{\text{ML}}(\lambda) \rangle|$ of the ML estimator (given by $\alpha_{d,\infty}(\lambda)$). These expressions are as follows.

Lemma 4 *For $d = 3$, the only positive solution to the fixed-point equation of Theorem 3 for positive values of λ is given by*

$$\mu_{3,\infty}(\lambda) = \frac{3\lambda^2 + \lambda\sqrt{9\lambda^2 - 12} + 4}{\sqrt{18\lambda^2 + 6\lambda\sqrt{9\lambda^2 - 12}}},$$

which holds for $\lambda \geq 2/\sqrt{3}$. Furthermore,

$$\alpha_{3,\infty}(\lambda) = \sqrt{\frac{1}{2} + \sqrt{\frac{3\lambda^2 - 4}{12\lambda^2}}}.$$

Unfortunately, showing that this holds true for any $d \geq 3$ is likely to be difficult, and we do not pursue this path here. In particular, for $d \geq 6$ the variational characterization of the alignment given by Jagannath et al. (2020) does not admit an explicit expression, and thus another strategy would be needed for such a proof. Nonetheless, for $d = 4$ and $d = 5$, explicit expressions for both quantities (spectral norm and alignment) do exist, and we verify numerically that they satisfy the equations (7)–(8). We thus conjecture that the same happens for every $d \geq 3$.

Before moving on, let us comment on the assumptions of Theorem 3. The condition $\mu_{d,\infty}(\lambda) > (d - 1)\beta_d$ is essentially imposed for a technical reason, related to the use of the implicit function theorem when computing derivatives of u and μ with respect to \mathcal{W} . As we will argue ahead, it is not restrictive for our purposes, since every local maximum u of the ML problem (including global ones) is such that $\mathbf{Y} \cdot u^d \geq (d - 1)\beta_d$ with probability $1 - o(1)$. The condition $\alpha_{d,\infty}(\lambda) > 0$, in its turn, is required at the end of our asymptotic analysis for deriving equations (7)–(8). Since $\alpha_{d,\infty}(\lambda)$ is the limiting alignment $\langle x, u \rangle$, this condition means that the considered sequence cannot converge to any critical point (asymptotically) lying on the orthogonal complement of x . Therefore, it prevents us from “seeing” the ML solution \hat{x}_{ML} in the regime where it is asymptotically uncorrelated with x , that is, below the critical value $\lambda_c(d)$.

Finally, the assumption that there exists a sequence of critical points such that both quantities in (5)–(6) converge almost surely is of course a strong one, and is crucial to our analysis. At present, it is not clear how one can rigorously show that such a sequence exists in the framework of our approach: it seems hard to unambiguously pin down a sequence of critical points for which standard arguments (namely, a moment control combined with Markov’s inequality and Borel-Cantelli) can be applied to show almost sure convergence. Our result indicates that, for $\lambda > \lambda_c(d)$, such a sequence asymptotically “behaves” like a global maximum, in the sense that its alignment and cost function value match (for $d = 3, 4, 5$) what is known for the ML estimator in that regime. However, the only characterization that we have at our disposal corresponds to the tensor eigenvalue equations, which

hold by definition for all critical points, combined with the condition $\mu_{d,\infty}(\lambda) > (d-1)\beta_d$, due to which one could probably restrict the sequence to contain local maxima only.

As mentioned before, the limiting spectral measure of the contractions $\mathbf{y} \cdot u^{d-2}$ with the sequence of critical points u considered in Theorem 3 is the same as that described by Theorem 2—indeed, the Stieltjes transform $m_d(z)$ is one of the main ingredients in the proof of Theorem 3. As we will see, it turns out that the additional terms introduced by the dependence of u on \mathcal{W} (which are computed using the implicit function theorem and Gaussian integration by parts) all vanish asymptotically. Hence, the limiting spectrum of $\mathbf{y} \cdot u^{d-2}$ is still a semicircle law on $[-\beta_d, \beta_d]$. (Moreover, if u is a local maximum, then a spike can be seen in the spectrum of $\mathbf{y} \cdot u^{d-2}$ regardless of the value of λ , because u is an eigenvector of that matrix with an eigenvalue larger than $(d-1)\beta_d$ with probability $1 - o(1)$ —see Figure 1 for an illustration).

3. Maximum likelihood estimation of the symmetric rank-one spiked tensor model

In the following, we introduce the key ideas that underlie our analysis of the symmetric rank-one spiked tensor model (1). Since the spherical prior imposed on x does not bring useful information for estimation purposes, for simplicity we will take the point of view in which x is an arbitrary deterministic vector on the unit sphere.

3.1 Maximum likelihood estimation, tensor and matrix eigenpairs

By writing the density of \mathbf{y} as

$$f(\mathbf{y}) = \frac{1}{Z_d(N)} e^{-\frac{N}{2} \|\mathbf{y} - \lambda x^{\otimes d}\|_F^2},$$

one can readily see that ML estimation of x boils down to solving the best rank-one approximation problem

$$\min_{u \in \mathbb{S}^{N-1}} \|\mathbf{y} - \lambda u^{\otimes d}\|_F^2.$$

Since $\lambda > 0$, it is straightforward to show that the above problem is equivalent to (2). For convenience, we will add a $\frac{1}{d}$ factor in the cost function of that problem, rewriting it as follows:

$$\max_{u \in \mathbb{S}^{N-1}} f(u), \quad \text{where} \quad f(u) := \frac{1}{d} \langle \mathbf{y}, u^{\otimes d} \rangle = \frac{1}{d} \mathbf{y} \cdot u^d. \quad (9)$$

Now, writing the Lagrangian as $\mathcal{L}(u, \mu) := f(u) - \frac{\mu}{2}(\langle u, u \rangle - 1)$, it follows that stationary points (μ, u) must satisfy the Karush-Kuhn-Tucker conditions

$$\frac{\partial \mathcal{L}}{\partial u} = \nabla f(u) - \mu u = \mathbf{y} \cdot u^{d-1} - \mu u = 0 \quad \text{and} \quad \langle u, u \rangle = 1,$$

implying that they are (real-valued) solutions to the equations

$$\mathbf{y} \cdot u^{d-1} = \mu u, \quad \langle u, u \rangle = 1. \quad (10)$$

These are called tensor ℓ_2 -eigenvalue⁶ equations, following the variational definition of tensor eigenpairs introduced by Lim (2005). For simplicity, we will refer to (μ, u) satisfying (10) as a tensor eigenpair, and to μ and u as an eigenvalue and an eigenvector of \mathcal{Y} , respectively. It should be borne in mind that, unlike the matrix case, the normalization $\|u\| = 1$ is essential, since $\mathcal{Y} \cdot (\alpha u)^{d-1} = \mu \alpha^{d-1} u \neq \mu \alpha u$ for $\alpha \notin \{0, 1\}$ and $d > 2$.

The above discussion shows that the spiked rank-one tensor model resembles in many respects its matrix counterpart, in that its ML estimator also corresponds to the eigenvector associated with the largest eigenvalue. However, there are several striking differences. One of them is computational, and has strong implications in practice: whereas this estimator can be computed in polynomial time in the matrix case, solving (2) for $d > 2$ is worst-case NP-hard (Hillar and Lim, 2013), and thus in general one can at best hope to find a local maximum of the nonconvex problem (9). Another difference concerns the sharp phase transition behavior that is observed in the tensor case, as we will discuss ahead.

Now, a simple but crucial observation that brings into focus the relevance of Definition 1 is the following, whose proof follows trivially from the identity $\mathcal{Y} \cdot u^{d-1} = (\mathcal{Y} \cdot u^{d-2}) u$:

Proposition 5 *Let $\mathcal{Y} \in \mathcal{S}^d(N)$ and $u \in \mathbb{S}^{N-1}$. The pair (μ, u) is an eigenpair of the tensor \mathcal{Y} if and only if it is an eigenpair of the matrix $\mathcal{Y} \cdot u^{d-2} \in \mathcal{M}(\mathcal{Y})$.*

An immediate consequence is that, by a suitable choice of matrices from the contraction ensemble $\mathcal{M}(\mathcal{Y})$, we can gain access to spectral properties of \mathcal{Y} . We will thoroughly rely on this link to derive our results on the spiked symmetric rank-one tensor model.

Remark 1 *Before moving on, we make another observation that will later prove useful. First note that, if (μ, u) is an eigenpair of \mathcal{Y} , then it is also an eigenpair of the Hessian $\nabla_{\mathcal{L}}^2 := \frac{\partial^2 \mathcal{L}}{\partial u^2}$ evaluated at u , since:*

$$\nabla_{\mathcal{L}}^2(u) = (d-1) \mathcal{Y} \cdot u^{d-2} - \mu I \quad \implies \quad \nabla_{\mathcal{L}}^2(u)u = (d-1) \mathcal{Y} \cdot u^{d-1} - \mu u = (d-2)\mu u. \quad (11)$$

A necessary condition for u to be a local maximum of problem (9) is that $\nabla_{\mathcal{L}}^2(u)$ be negative semidefinite on the tangent space at u (see, e.g., Theorem 12.5 of Nocedal and Wright (2006)), that is,

$$\langle \nabla_{\mathcal{L}}^2(u) w, w \rangle \leq 0, \quad \forall w \in u^\perp := \{w \in \mathbb{R}^N \mid \langle u, w \rangle = 0\}.$$

Under the above condition, it follows from (11) that

$$\max_{w \in \mathbb{S}^{N-1} \cap u^\perp} \left\langle (\mathcal{Y} \cdot u^{d-2}) w, w \right\rangle \leq \frac{\mu}{d-1}.$$

This means that μ must be the largest eigenvalue of (the matrix) $\mathcal{Y} \cdot u^{d-2}$, with all others bounded by $\mu/(d-1)$, which is true in particular for the largest eigenvalue of \mathcal{Y} , as discussed above.

6. Other definitions of eigenpairs have been proposed for tensors. In the book on tensor spectral theory by Qi and Luo (2017), solutions to (10) are termed E-eigenpairs, and real-valued solutions are called Z-eigenpairs.

Remark 1 implies that, at every local maximum u of the ML problem, the second largest eigenvalue of $\mathbf{Y} \cdot u^{d-2}$ cannot be larger than $\mu/(d-1)$ and, consequently, the largest eigenvalue μ is isolated from the rest of the spectrum of $\mathbf{Y} \cdot u^{d-2}$. At first sight this seems to be a particularly unsettling conclusion since, unlike the spiked Wigner random *matrix* model, this implies the systematic presence of a “spike” in the spectrum of $\mathbf{Y} \cdot u^{d-2}$ (positioned at least $d-1$ times as far as the bulk spectrum of $\mathbf{Y} \cdot u^{d-2}$) irrespective of the (possibly null) signal-to-noise ratio λ . Yet, as previously discussed, note that this spike can merely be induced by the contraction of \mathbf{W} with a dependent vector u , and thus may bear no relation whatsoever with the “signal” term $\lambda x^{\otimes d}$.

3.2 Asymptotic performance of ML estimation

Before stating our results, we briefly recall the contributions recently made by Jagannath et al. (2020). Let

$$\varphi_d(\lambda, t) := \lambda^2 t^d + \log(1-t) + t,$$

and define $\lambda_c(d) := \sup\{\lambda \geq 0 : \sup_{t \in [0,1[} \varphi_d(\lambda, t) \leq 0\}$. It was shown in (Jagannath et al., 2020) that the critical value λ_c defines the phase transition point, below which ML estimation of the planted vector x of the model (1) suddenly breaks down, in the sense that its solution (asymptotically) ceases to exhibit any correlation with x as measured by the alignment $|\langle x, \hat{x}_{\text{ML}}(\lambda) \rangle|$. More precisely, their result states that for $\lambda \neq \lambda_c(d)$,

$$|\langle x, \hat{x}_{\text{ML}}(\lambda) \rangle| \xrightarrow{\text{a.s.}} \sqrt{q_d(\lambda)}$$

as $N \rightarrow \infty$, where

$$q_d(\lambda) := \begin{cases} 0, & \text{if } \lambda < \lambda_c(d), \\ \arg \max_{t \in [0,1[} \varphi_d(\lambda, t), & \text{if } \lambda > \lambda_c(d). \end{cases}$$

Furthermore, the quantity $\mu_d^*(\lambda) = \mathbf{Y} \cdot (\hat{x}_{\text{ML}}(\lambda))^d$ also converges almost surely to a deterministic expression given by

$$\mu_d^*(\lambda) \xrightarrow{\text{a.s.}} \begin{cases} \mu_0(d), & \lambda \leq \lambda_c(d), \\ \sqrt{d} \frac{1 + \lambda^2 q_d(\lambda)^{d-1}}{\sqrt{1 + \lambda^2 d q_d(\lambda)^{d-1}}} & \lambda > \lambda_c(d), \end{cases} \quad (12)$$

where $\mu_0(d) = (d + z^*)/\sqrt{d + dz^*}$, with z^* denoting the unique positive solution of

$$\frac{1+z}{z^2} \log(1+z) - \frac{1}{z} - \frac{1}{d} = 0.$$

As discussed in the Introduction, $\mu_0(d)$ is in fact the (almost surely) asymptotic value of $\frac{1}{\sqrt{N}} \|\mathbf{W}\|$, and it can be shown that $\lambda_c(d) < \mu_0(d)$ for all d —hence the claim of a BBP-like phase transition taking place for all d . Similarly, $\mu_d^*(\lambda)$ is in effect the asymptotic spectral norm of \mathbf{Y} .

For $d = 3, 4, 5$, one can analytically solve $\frac{\partial \varphi_d(\lambda, t)}{\partial t} = d\lambda^2 t^{d-1} + t/(t-1) = 0$ to find $q_d(\lambda)$ for $\lambda_c(d)$, which gives

$$q_3(\lambda) = \frac{1}{2} + \sqrt{\frac{3\lambda^2 - 4}{12\lambda^2}}, \quad \lambda > \lambda_c(3),$$

$$\begin{aligned}
 q_4(\lambda) &= \frac{1}{3} + \frac{g_4(\lambda)}{6} + \frac{2}{3g_4(\lambda)}, \quad \lambda > \lambda_c(4), \\
 q_5(\lambda) &= \frac{1}{4} + \frac{\sqrt{15} h_5(\lambda)}{60} + \frac{\sqrt{450}}{60} \sqrt{\frac{\left(\lambda g_5(\lambda)^{\frac{1}{3}} - \frac{g_5(\lambda)^{\frac{2}{3}}}{15} - 16\right) h_5(\lambda) + \sqrt{15} \lambda g_5(\lambda)^{\frac{1}{3}}}{g_5(\lambda)^{\frac{1}{3}} h_5(\lambda) \lambda}}, \\
 &\quad \lambda > \lambda_c(5),
 \end{aligned}$$

with

$$\begin{aligned}
 g_4(\lambda) &= \left(\frac{8\lambda^2 + 3\sqrt{81 - 48\lambda^2} - 27}{\lambda^2} \right)^{\frac{1}{3}}, \\
 g_5(\lambda) &= 60\sqrt{15} \sqrt{135\lambda^2 - 256} + 2700\lambda, \\
 h_5(\lambda) &= \sqrt{\frac{2g_5(\lambda)^{\frac{2}{3}} + 15\lambda g_5(\lambda)^{\frac{1}{3}} + 480}{\lambda g_5(\lambda)^{\frac{1}{3}}}},
 \end{aligned}$$

and hence

$$\begin{aligned}
 \mu_3^*(\lambda) &= \frac{3\lambda^2 + \lambda\sqrt{9\lambda^2 - 12} + 4}{\sqrt{18\lambda^2 + 6\lambda\sqrt{9\lambda^2 - 12}}}, \quad \lambda > \lambda_c(3), \\
 \mu_4^*(\lambda) &= \frac{2 + 2\lambda^2 \left(\frac{1}{3} + \frac{h_4(\lambda)}{6} + \frac{2}{3h_4(\lambda)} \right)^3}{\sqrt{1 + 4\lambda^2 \left(\frac{1}{3} + \frac{h_4(\lambda)}{6} + \frac{2}{3h_4(\lambda)} \right)^3}}, \quad \lambda > \lambda_c(4).
 \end{aligned}$$

These quantities are plotted in Figure 3. (As the expression of μ_5^* is too long, we choose to omit it here.) In fact, one can easily check that the above explicit formulas are still algebraically valid slightly below $\lambda_c(d)$, namely, for $\lambda_s(d) \leq \lambda < \lambda_c(d)$ with $\lambda_s(d) := \sqrt{(d-1)^{d-1}/d(d-2)^{d-2}}$. As discussed by Jagannath et al. (2020) and Ben Arous et al. (2019), a local maximizer of the likelihood having alignment given (asymptotically) by $\sqrt{q_d(\lambda)}$ with x exists in this regime as N grows large. However, the above results are relevant not only to ML estimation, but in fact to any estimator: it turns out that the ML estimator is information-theoretically optimal, in the sense that no other estimator can attain a higher correlation $|\langle x, \hat{x}_{\text{ML}}(\lambda) \rangle|$ for any $\lambda \geq 0$, as also shown by Jagannath et al. (2020). This implies, in particular, that no estimator $\hat{x}(\mathcal{Y})$ of x can (asymptotically) achieve nonzero alignment with x in the regime $\lambda < \lambda_c(d)$ without additional information on x . Therefore, the above mentioned local maximizer that is positively correlated with x for $\lambda < \lambda_c(d)$ cannot be a valid estimator of x —that is, it cannot be expressed a measurable function of \mathcal{Y} .

The reader will notice that the above discussion makes no reference to the spectral viewpoint discussed in the previous subsection. The omission is intentional, as the statement and derivation of the results of (Jagannath et al., 2020) do not rely on or mention in any way the connection between ML estimation and tensor eigenpairs. By bringing this connection into the picture and combining it with our proposed approach described in the previous section, we are able to reach some of the same predictions but using standard tools from RMT rather than from spin glass theory.




Figure 3: Phase transition for the ML estimator of the symmetric rank-one spiked tensor model (1) with $d = 3, 4, 5$, according to Jagannath et al. (2020): (a) asymptotic alignment between x and ML solution $\hat{x}_{\text{ML}}(\lambda)$; (b) asymptotic maximum likelihood (or spectral norm of \mathcal{Y}). The argument of $\lambda_c(d)$ has been omitted for simplicity.

3.3 Analysis under the light of random matrix theory

In this section, we explain how we leverage our approach to obtain the result described by Theorem 3, concerning ML estimation of the planted vector x of the symmetric spiked rank-one tensor model (1). Then, we will discuss connections with other works, limitations and possible extensions of this result. The proof of Theorem 3 itself is postponed to Section 4.

3.3.1 ANALYZING CONTRACTIONS OF CRITICAL POINTS

Let us now turn back to the goal of studying $\mathcal{Y} \cdot (\hat{x}_{\text{ML}}(\lambda))^{d-2}$. As there seems to be no simple way of taking into account the fact that $\hat{x}_{\text{ML}}(\lambda)$ is a global maximum, we consider more generally random contractions $\mathcal{Y} \cdot u^{d-2}$ with critical points u . This can be done by relying on the tensor eigenvalue equations as a characterization of critical points. Nevertheless,

these critical points must asymptotically behave as local maxima, since we also impose the condition that the largest eigenvalue of $\mathbf{Y} \cdot u^{d-2}$ must (asymptotically) be strictly greater than $(d-1)\beta_d$ (cf. Remark 1). The tensor eigenvalue equations give

$$\mu u = \mathbf{Y} \cdot u^{d-1} = \lambda \langle x, u \rangle^{d-1} x + \frac{1}{\sqrt{N}} \mathbf{W} \cdot u^{d-1},$$

from which follows by taking scalar products with u and x (and recalling that $\|u\| = \|x\| = 1$) that

$$\mu = \lambda \langle x, u \rangle^d + \frac{1}{\sqrt{N}} \mathbf{W} \cdot u^d, \quad (13)$$

$$\langle x, u \rangle = \frac{\lambda}{\mu} \langle x, u \rangle^{d-1} + \frac{1}{\mu \sqrt{N}} \langle x, \mathbf{W} \cdot u^{d-1} \rangle, \quad (14)$$

respectively. Taking the expectation of these expressions involves considerable difficulty, as u depends on \mathbf{W} in an *a priori* intricate manner, and thus cannot be pulled out from expectations. The technical tools that allow managing this dependency are the implicit function theorem and Stein's lemma (also known simply as Gaussian integration by parts; see Appendix C for a statement). The former gives us the derivatives of u and μ with respect to components of \mathbf{W} , which read

$$\begin{pmatrix} \frac{\partial u}{\partial W_{i_1 \dots i_d}} \\ \frac{\partial \mu}{\partial W_{i_1 \dots i_d}} \end{pmatrix} = \begin{pmatrix} -\frac{1}{(d-1)\sqrt{N}} R\left(\frac{\mu}{d-1}\right) \phi + \frac{1}{(d-2)\mu} \frac{\partial \mu}{\partial W_{i_1 \dots i_d}} u \\ \frac{1}{\sigma_{W_{i_1 \dots i_d}}^2 \sqrt{N}} \prod_{j=1}^d u_{i_j} \end{pmatrix}, \quad (15)$$

where $R(z)$ is the resolvent of $\mathbf{Y} \cdot u^{d-2}$ (a definition is given in Appendix C) and ϕ is an N -dimensional vector that satisfies

$$\sigma_{W_{i_1 \dots i_d}}^2 \phi = \frac{1}{d} \sum_{j=1}^d u_1 \dots u_{i_{j-1}} u_{i_{j+1}} \dots u_d e^{(i_j)},$$

with $e^{(i_j)}$ denoting the i_j canonical basis vector. The expectations on the right-hand sides of (13) and (14) can then be computed by resorting to (15), thanks to Stein's lemma.

Now, one must be careful when taking the large- N limit of these expectations, because of the term involving the resolvent $R(z)$ evaluated at $z = \mu/(d-1)$ in (15). In fact, along the proof we show that

$$\forall z \in \mathbb{C} \text{ such that } \Im\{z\} > 0, \quad \frac{1}{N} \operatorname{tr} R(z) \xrightarrow{\text{a.s.}} m_d(z),$$

where $m_d(z)$ is as given by Theorem 2, and thus by the Stieltjes continuity theorem (see Appendix C) the spectral measure of $\mathbf{Y} \cdot u^{d-2}$ converges weakly almost surely to the same semicircle law of (4), which is supported on $[-\beta_d, \beta_d]$. Consequently, the sequence of critical points u considered in Theorem 3 must be such that the almost sure limit $\mu_{d,\infty}(\lambda)$ of

$\mu = \mathfrak{Y} \cdot u^d$ (assumed to exist) must be strictly greater than $(d-1)\beta_d$, so that $(d-1)^{-1} \mu_{d,\infty}(\lambda)$ does not fall on the limiting spectrum of $\mathfrak{Y} \cdot u^{d-2}$.

Once in possession of the asymptotic expressions of these expectations as $N \rightarrow \infty$, we invoke the assumption that the random quantities μ and $\langle x, u \rangle$ converge almost surely to a deterministic expression depending only on λ and on d (as assumed in the statement of Theorem 3). This step (along with the fact that the normalized trace of the resolvent of $\mathfrak{Y} \cdot u^{d-2}$ also converges almost surely, as in Theorem 2) allows breaking the resulting expectations of products into products of expectations, leading to the final fixed-point equation described in the statement of Theorem 3. Moreover, it turns out that this equation effectively describes the performance of ML in the regime $\lambda > \lambda_c(d)$ for $d = 3, 4, 5$, as we already mentioned and will further discuss next.

3.3.2 DISCUSSION

One can readily verify that the expressions of q_3 and μ_3^* given in Section 3.2 precisely match those given by Lemma 4, meaning that our fixed point equation describes the asymptotic performance of the ML solution in the regime $\lambda > \lambda_c(3)$. Even though we did not compute the solution of our fixed-point equation explicitly for other values of d , it can be numerically verified for $d = 4, 5$ that they are again consistent with the characterization of Jagannath et al. (2020) described above. We therefore conjecture the following:

Conjecture 6 *For any $d \geq 3$ and $\lambda > \lambda_c(d)$, the quantities characterized by Theorem 3 match those described in Section 3.2, that is,*

$$\alpha_{d,\infty}(\lambda) = \sqrt{q_d(\lambda)}, \quad \mu_{d,\infty}(\lambda) = \mu_d^*(\lambda).$$

Furthermore, the first above equation also holds in the regime $\lambda_s(d) \leq \lambda \leq \lambda_c(d)$.

It can seem somewhat surprising that our analysis is able to predict the performance of ML estimation in the regime $\lambda > \lambda_c(d)$, despite relying on the tensor eigenvalue equations satisfied not only by the ML solution but also by all other critical points of the ML problem—which are on average exponentially numerous, according to the results of Ben Arous et al. (2019). The key to this “selectivity” lies of course in the additional assumptions of Theorem 3, which seem to restrict the considered sequence of critical points u to asymptotically behave as global maxima of the ML problem in the regime $\lambda > \lambda_c(d)$. However, a precise understanding of this effect remains yet to be achieved.

Another problem that remains open is whether our RMT-based approach can also predict the phase transition point $\lambda_c(d)$. This being said, from a practical standpoint, it is also known that the ML solution cannot be numerically estimated in polynomial time for values of λ close to $\lambda_c(d)$ (it is even conjectured that λ must scale like $O(N^\ell)$ with $\ell > 0$ for polynomial-time algorithms to recover the solution, as discussed by Jagannath et al., 2020; Hopkins et al., 2015); this, in a way, makes the very notion of a phase transition for the ML solution unessential.

4. Proof of the main results

In this section, we will provide the proof of Theorem 3, and explain how one part of it can be adapted to show also Theorem 2. For clarity, we have organized the proof into separate subsections which deal with each required ingredient.

Recall that we consider a sequence of critical points (μ, u) of the ML problem, and hence satisfy the tensor eigenvalue equations $\mathbf{Y} \cdot u^{d-2} = \mu u$ and $\|u\| = 1$. Let us start by computing the expectation of μ as given by (13).

4.1 Notational conventions

Since we manipulate explicit summations over indices i_1, \dots, i_d arising from the contraction of \mathbf{Y} with u and x , it is convenient to introduce some notation for avoiding too lengthy expressions. Namely, we will abbreviate the indices i_1, \dots, i_d by a boldface notation \mathbf{i} (and similarly for \mathbf{j} and ℓ). Under this convention, we have for instance $W_{i_1, \dots, i_d} = W_{\mathbf{i}}$ and $\sum_{i_1, \dots, i_d} = \sum_{\mathbf{i}}$, where each i_j runs from 1 to N . We will also simplify the symbol used for the variance of $W_{\mathbf{i}}$, writing it more compactly as $\sigma_{\mathbf{i}}^2$ instead of $\sigma_{W_{\mathbf{i}}}^2$. Finally, the shorthand $\prod_{k \neq j}^d u_{i_k}$ will be used to denote $\prod_{k \in \{1, \dots, d\} \setminus \{j\}} u_{i_k}$. Similarly, $\prod_{k \notin S}^d u_{i_k}$ stands for $\prod_{k \in \{1, \dots, d\} \setminus S} u_{i_k}$.

4.2 Expectation of μ

The starting point is the expression

$$\mathbb{E}\{\mu\} = \lambda \mathbb{E}\{\langle x, u \rangle^d\} + \frac{1}{\sqrt{N}} \mathbb{E}\{\mathbf{W} \cdot u^d\}.$$

To develop it, we apply Stein's lemma (which is simply Gaussian integration by parts, see Appendix C) to the second term on the right-hand side, by relying on the derivative of u with respect to components of \mathbf{W} , as computed in Appendix A. This step requires the following result, which may not be immediately clear. Its proof will be postponed to Appendix B.

Lemma 7 *There exists a function $G : \mathcal{S}^d(N) \rightarrow \mathbb{R}^{N+1}$ which is almost everywhere continuously differentiable and such that $G(\mathbf{W}) = (\mu(\mathbf{W}), u(\mathbf{W}))$ is an eigenpair of \mathbf{W} (for almost every \mathbf{W}).*

The components of $u(\mathbf{W})$ are bounded and $\mu(\mathbf{W})$ has polynomial growth, and hence Stein's lemma applies. We also note that all subsequent applications of this tool in the following sections can be justified in the same manner.

Now, according to our computation in Appendix A, the partial derivative $\frac{\partial u_{i_j}}{\partial W_{\mathbf{i}}}$ of a component of such a function u is expressed in terms of the resolvent of $\mathbf{Y} \cdot u^{d-2}$, that is, of $R(z) = (\mathbf{Y} \cdot u^{d-2} - zI)^{-1}$, and also of a certain vector ϕ defined in (28), which depends on u and on the variance of $W_{\mathbf{i}}$. For simplicity, we will use the notation $\bar{R} := R(\mu/(d-1))$, as in Appendix A.

With these definitions, we can write

$$\mathbb{E}\{\mathbf{W} \cdot u^d\} = \sum_{\mathbf{i}} \mathbb{E}\{W_{\mathbf{i}} u_{i_1} \dots u_{i_d}\}$$

$$\begin{aligned}
 &= \sum_i \sigma_i^2 \sum_{j=1}^d \mathbb{E} \left\{ \frac{\partial u_{i_j}}{\partial W_i} \prod_{k \neq j}^d u_{i_k} \right\} \\
 &= \sum_i \sigma_i^2 \sum_{j=1}^d \mathbb{E} \left\{ -\frac{1}{(d-1)\sqrt{N}} [\bar{R}\phi]_{i_j} \prod_{k \neq j}^d u_{i_k} + \frac{1}{\sigma_i^2} \frac{1}{(d-2)\mu\sqrt{N}} \prod_{k=1}^d u_{i_k}^2 \right\}.
 \end{aligned}$$

Using now (28) and $\|u\| = 1$, we get

$$\begin{aligned}
 \mathbb{E} \left\{ \mathbf{W} \cdot u^d \right\} &= -\frac{1}{(d-1)\sqrt{N}} \sum_i \sum_{j=1}^d \mathbb{E} \left\{ \left(\prod_{k \neq j}^d u_{i_k} \right) \frac{1}{d} \sum_{\ell=1}^d \bar{R}_{i_j, i_\ell} \prod_{m \neq \ell}^d u_{i_m} \right\} \\
 &\quad + \frac{d}{(d-2)\sqrt{N}} \mathbb{E} \left\{ \frac{1}{\mu} \right\}.
 \end{aligned}$$

Now, for each j , the summation in ℓ in the first expectation contains exactly one term of the form $\bar{R}_{i_j, i_j} \prod_{k \neq j}^d u_{i_k}^2$, which gives rise to $\text{tr } \bar{R}$ when summing over all indices i_1, \dots, i_d . All the other $d-1$ terms are of the form $(\prod_{k \notin \{j, \ell\}}^d u_{i_k}^2) u_{i_j} \bar{R}_{i_j, i_\ell} u_{i_\ell}$. Because $((\mu - \mu/(d-1))^{-1}, u)$ is an eigenpair of \bar{R} , summing each such term over the indices i_1, \dots, i_d yields

$$\sum_i \left(\prod_{k \notin \{j, \ell\}}^d u_{i_k}^2 \right) u_{i_j} \bar{R}_{i_j, i_\ell} u_{i_\ell} = \langle u, \bar{R}u \rangle = \frac{d-1}{(d-2)\mu}.$$

Hence, using these results we arrive at

$$\begin{aligned}
 \mathbb{E} \left\{ \mathbf{W} \cdot u^d \right\} &= -\frac{1}{d(d-1)\sqrt{N}} \mathbb{E} \left\{ d \text{tr } \bar{R} + \frac{d(d-1)^2}{(d-2)\mu} \right\} + \frac{d}{(d-2)\sqrt{N}} \mathbb{E} \left\{ \frac{1}{\mu} \right\} \\
 &= \frac{1}{(d-2)\sqrt{N}} \mathbb{E} \left\{ \frac{1}{\mu} \right\} - \frac{1}{(d-1)\sqrt{N}} \mathbb{E} \left\{ \text{tr } \bar{R} \right\}.
 \end{aligned}$$

Therefore, we have shown that

$$\mathbb{E} \{ \mu \} = \lambda \mathbb{E} \left\{ \langle x, u \rangle^d \right\} + \frac{1}{(d-2)N} \mathbb{E} \left\{ \frac{1}{\mu} \right\} - \frac{1}{(d-1)N} \mathbb{E} \left\{ \text{tr } \bar{R} \right\}. \quad (16)$$

4.3 Expectation of $\langle x, u \rangle$

Now let us carry out a similar computation for the expected alignment:

$$\mathbb{E} \left\{ \langle x, u \rangle \right\} = \lambda \mathbb{E} \left\{ \frac{1}{\mu} \langle x, u \rangle^{d-1} \right\} + \frac{1}{\sqrt{N}} \mathbb{E} \left\{ \frac{1}{\mu} \langle x, \mathbf{W} \cdot u^{d-1} \rangle \right\}.$$

By explicitly developing the scalar product of the last term, we write

$$\mathbb{E} \left\{ \frac{1}{\mu} \langle x, \mathbf{W} \cdot u^{d-1} \rangle \right\} = \sum_i \mathbb{E} \left\{ \frac{1}{\mu} W_i u_{i_1} \dots u_{i_{d-1}} x_{i_d} \right\}.$$

Treating now μ and each component of u as functions of the Gaussian tensor \mathbf{W} , by Stein's lemma we obtain

$$\mathbb{E} \left\{ \frac{1}{\mu} \langle x, \mathbf{W} \cdot u^{d-1} \rangle \right\} = \sum_{\mathbf{i}} \sigma_{\mathbf{i}}^2 \mathbb{E} \left\{ -\frac{1}{\mu^2} \frac{\partial \mu}{\partial W_{\mathbf{i}}} \left(\prod_{k=1}^{d-1} u_{i_k} \right) x_{i_d} + \frac{1}{\mu} x_{i_d} \sum_{j=1}^{d-1} \frac{\partial u_{i_j}}{\partial W_{\mathbf{i}}} \prod_{k \neq j}^{d-1} u_{i_k} \right\}.$$

Since $\sigma_{\mathbf{i}}^2 \sqrt{N} \frac{\partial \mu}{\partial W_{\mathbf{i}}} = \prod_j u_{i_j}$, after multiplying by $\sigma_{\mathbf{i}}^2$ and summing over \mathbf{i} the first term in the expectation becomes $-N^{-1/2} \langle x, u \rangle / \mu^2$, as all factors u_{i_j} appear twice except for u_{i_d} . Using now the expression for the derivative of u_{i_j} , we get

$$\begin{aligned} \mathbb{E} \left\{ \frac{1}{\mu} \langle x, \mathbf{W} \cdot u^{d-1} \rangle \right\} &= -\frac{1}{\sqrt{N}} \mathbb{E} \left\{ \frac{\langle x, u \rangle}{\mu^2} \right\} \\ &+ \sum_{\mathbf{i}} \mathbb{E} \left\{ x_{i_d} \frac{\sigma_{\mathbf{i}}^2}{\mu} \sum_{j=1}^{d-1} \left(\prod_{k \neq j}^{d-1} u_{i_k} \right) \left[-\frac{[\bar{R}\phi]_{i_j}}{(d-1)\sqrt{N}} + \frac{u_{i_j}^2}{\mu(d-2)\sqrt{N}\sigma_{\mathbf{i}}^2} \prod_{m \neq j}^d u_{i_m} \right] \right\}. \end{aligned}$$

The second term inside the inner brackets is easy to compute, since the products of u_{i_k} and of u_{i_m} combine to yield $u_{i_d} \prod_{k \neq j}^{d-1} u_{i_k}^2$, and thus the summation over \mathbf{i} leaves only $\langle x, u \rangle$ divided by $\mu^2(d-2)\sqrt{N}$ for each j . As for the first term, we can expand it as $[\bar{R}\phi]_{i_j} = (d\sigma_{\mathbf{i}}^2)^{-1} \sum_{\ell=1}^d \bar{R}_{i_j, i_\ell} \prod_{m \neq \ell}^d u_{i_m}$ and note that for $j = \ell$ (there are $d-1$ such terms), we recover after summing over \mathbf{i} a term of the form $-(d\sqrt{N}\mu)^{-1} \langle x, u \rangle \text{tr} \bar{R}$. The terms for which $j \neq \ell$ (there are $(d-1)^2$) can be obtained by noting that $\bar{R}u = \frac{d-1}{(d-2)\mu} u$ and by matching the indices of the components of u (with only u_{i_d} remaining free to contract with x_{i_d}). This computation gives

$$\begin{aligned} \mathbb{E} \left\{ \frac{1}{\mu} \langle x, \mathbf{W} \cdot u^{d-1} \rangle \right\} &= -\frac{1}{\sqrt{N}} \mathbb{E} \left\{ \frac{\langle x, u \rangle}{\mu^2} + \frac{(d-1)\langle x, u \rangle}{(d-2)\mu^2} - \frac{\langle x, u \rangle}{d\mu} \text{tr} \bar{R} - \frac{(d-1)^2 \langle x, u \rangle}{d(d-2)\mu^2} \right\} \\ &= \frac{d-(d-1)^2}{d(d-2)\sqrt{N}} \mathbb{E} \left\{ \frac{\langle x, u \rangle}{\mu^2} \right\} - \frac{1}{d\sqrt{N}} \mathbb{E} \left\{ \frac{\langle x, u \rangle}{\mu} \text{tr} \bar{R} \right\}. \end{aligned}$$

We have thus derived the formula

$$\mathbb{E} \{ \langle x, u \rangle \} = \lambda \mathbb{E} \left\{ \frac{1}{\mu} \langle x, u \rangle^{d-1} \right\} + \frac{d-(d-1)^2}{d(d-2)N} \mathbb{E} \left\{ \frac{\langle x, u \rangle}{\mu^2} \right\} - \frac{1}{dN} \mathbb{E} \left\{ \frac{\langle x, u \rangle}{\mu} \text{tr} \bar{R} \right\}. \quad (17)$$

4.4 Concentration of $N^{-1} \text{tr} R(z)$ around $m_d(z)$

Since both expressions (16) and (17) involve the normalized trace of the resolvent R of $\mathbf{Y} \cdot u^{d-2}$ evaluated at $\mu/(d-1)$, we will need the asymptotic expression of that quantity. It turns out that this normalized trace converges almost surely to $m_d(z)$, that is, the same Stieltjes transform that was computed for contractions of \mathbf{W} with independent unit norm vectors v in Theorem 2. We will now show this fact.

From now on, for simplicity of notation we shall omit the argument of R whenever it is simply z . As in the previous sections, the notation \bar{R} will stand for $R(\mu/(d-1))$, where μ

is the eigenvalue associated with the vector u that defines the contraction $\mathfrak{Y} \cdot u^{d-2}$. Recall that R satisfies the recurrence relation

$$R = -\frac{1}{z}I + \frac{1}{z} \left(\mathfrak{Y} \cdot u^{d-2} \right) R. \quad (18)$$

Replacing \mathfrak{Y} by its definition and taking the expectation of the normalized trace,

$$\mathbb{E} \left\{ \frac{1}{N} \operatorname{tr} R \right\} = -\frac{1}{z} + \frac{\lambda}{Nz} \mathbb{E} \left\{ \langle x, u \rangle^{d-2} \langle x, Rx \rangle \right\} + \frac{1}{zN\sqrt{N}} \sum_i \mathbb{E} \left\{ W_i \left(\prod_{k=2}^{d-1} u_{i_k} \right) R_{i_1 i_d} \right\}. \quad (19)$$

As usual, Stein's lemma will provide us an expression for the second expectation of the right-hand side:

$$\frac{1}{N\sqrt{N}} \sum_i \mathbb{E} \left\{ W_i \left(\prod_{k=2}^{d-1} u_{i_k} \right) R_{i_1 i_d} \right\} = A + B, \quad (20)$$

where

$$A := \frac{1}{N\sqrt{N}} \sum_i \sigma_i^2 \mathbb{E} \left\{ \frac{\partial R_{i_1 i_d}}{\partial W_i} \left(\prod_{k=2}^{d-1} u_{i_k} \right) \right\},$$

$$B := \frac{1}{N\sqrt{N}} \sum_i \sigma_i^2 \mathbb{E} \left\{ R_{i_1 i_d} \sum_{\ell=2}^{d-1} \frac{\partial u_{i_\ell}}{\partial W_i} \left(\prod_{k=2, k \neq \ell}^{d-1} u_{i_k} \right) \right\}.$$

In the following, we will compute asymptotic expressions for A and B .

4.4.1 ASYMPTOTIC EXPRESSION OF A IN (20)

To compute A , we will need the derivative of R with respect to each component \mathfrak{W}_i . Starting again from (18), it is easy to derive

$$\frac{\partial R}{\partial W_i} = -R \frac{\partial}{\partial W_i} \left(\mathfrak{Y} \cdot u^{d-2} \right) R,$$

which component-wise reads

$$\frac{\partial R_{pq}}{\partial W_i} = -\frac{1}{\sqrt{N}} \sum_j R_{pj_1} \left(\frac{\partial W_j}{\partial W_i} \right) \left(\prod_{k=2}^{d-1} u_{j_k} \right) R_{j_d q} - \sum_j \sum_{\ell=2}^{d-1} R_{pj_1} \mathfrak{Y}_j \left(\frac{\partial u_{j_\ell}}{\partial W_i} \right) \left(\prod_{k=2, k \neq \ell}^{d-1} u_{j_k} \right) R_{j_d q}.$$

Hence, we can write $A = A_1 + A_2$, where

$$A_1 := -\frac{1}{N^2} \sum_i \sum_j \sigma_i^2 \mathbb{E} \left\{ R_{i_1 j_1} \left(\frac{\partial W_j}{\partial W_i} \right) \left(\prod_{k=2}^{d-1} u_{i_k} u_{j_k} \right) R_{j_d i_d} \right\}$$

and

$$A_2 := -\frac{1}{N\sqrt{N}} \sum_i \sum_j \sigma_i^2 \mathbb{E} \left\{ \sum_{\ell=2}^{d-1} R_{i_1 j_1} \mathfrak{Y}_j \left(\frac{\partial u_{j_\ell}}{\partial W_i} \right) \left(\prod_{k=2, k \neq \ell}^{d-1} u_{i_k} u_{j_k} \right) R_{i_d j_d} u_{i_\ell} \right\}. \quad (21)$$

We focus first on A_1 . Since $\frac{\partial W_j}{\partial W_i} = 1$ if \mathbf{j} is a permutation of \mathbf{i} and $\frac{\partial W_j}{\partial W_i} = 0$ otherwise, and since σ_i^2 is given by reciprocal of the number of distinct permutations of (the particular value of) \mathbf{i} , the value of A_1 is proportional to the sum of all $d!$ possible d -fold contractions obtained by matching the indices of \mathbf{j} with one of the possible permutations of \mathbf{i} . This conclusion is reached by noting that σ_i^2 is precisely equal to $1/d!$ times the number of “repetitions” of the term with indices \mathbf{i} that are required in order to “complete” all such contractions. Therefore, by taking into account the result of all such possible contractions, we conclude that A_1 has the form

$$A_1 = -\frac{1}{N^2 d!} \left(K_1 \mathbb{E} \{(\operatorname{tr} R)^2\} + K_2 \mathbb{E} \{\operatorname{tr} R^2\} + K_3 \mathbb{E} \left\{ \frac{1}{(\mu - z)^2} \right\} + K_4 \mathbb{E} \left\{ \operatorname{tr} R \frac{1}{\mu - z} \right\} \right),$$

where the integers K_i count the number of contractions yielding terms of each one of the forms in this expression. In particular, it is easy to see that $K_1 = (d-2)!$, since $(\operatorname{tr} R)^2$ results from each contraction in which (i_1, j_1) and (i_d, j_d) are matched, while all others $d-2$ indices j_2, \dots, j_{d-1} can be arbitrarily permuted to match i_2, \dots, i_{d-1} . At this point, we will simplify the subsequent calculations by already dropping terms which will vanish at the end in the large- N limit. We have

$$\lim_{N \rightarrow \infty} A_1 = -\frac{1}{d(d-1)} \lim_{N \rightarrow \infty} \mathbb{E} \left\{ \left(\frac{1}{N} \operatorname{tr} R \right)^2 \right\}, \quad (22)$$

since the terms multiplied by K_2 , K_3 and K_4 are all $o(N^2)$. We will see in a moment that $(\frac{1}{N} \operatorname{tr} R)^2$ will in fact converge almost surely to $(m_d(z))^2$.

The key in the above argument is of course the fact that, in expectation, $\langle u, Ru \rangle = 1/(\mu - z)$ and $\langle u, u \rangle$ are $O(1)$, while $\operatorname{tr} R$ is $O(N)$. By the same token, the term A_2 defined in (21) can also be neglected, since the derivative of u_{j_ℓ} contains a $1/\sqrt{N}$ factor and the contractions in (21) only produce $o(N^2)$ terms: all resolvents appearing in the expression are ultimately contracted with one “copy” of u (hence no second-order moment $(\frac{1}{N} \operatorname{tr} R)^2$ arises), and contractions of u with \mathbf{y} are also $O(1)$.

Hence, in summary, A is asymptotically given by the expression derived in (22).

4.4.2 THE TERM B IN (20) VANISHES ASYMPTOTICALLY

Let us proceed by inserting the derivative of u into B to get $B = B_1 + B_2$, where

$$\begin{aligned} B_1 &:= \frac{1}{N\sqrt{N}} \sum_{\mathbf{i}} \sigma_{\mathbf{i}}^2 \mathbb{E} \left\{ R_{i_1 i_d} \left(\prod_{k=2, k \neq \ell}^{d-1} u_{i_k} \right) \sum_{\ell=2}^{d-1} -\frac{1}{(d-1)\sqrt{N}} \frac{1}{d\sigma_{\mathbf{i}}^2} \sum_{m=1}^d \bar{R}_{i_\ell, i_m} \left(\prod_{n \neq m}^d u_{i_n} \right) \right\} \\ &= -\frac{1}{N^2 d(d-1)} \sum_{\mathbf{i}} \mathbb{E} \left\{ \sum_{\ell=2}^{d-1} \sum_{m=1}^d R_{i_1 i_d} \left(\prod_{k=2, k \neq \ell}^{d-1} u_{i_k} \right) \bar{R}_{i_\ell, i_m} \left(\prod_{n \neq m}^d u_{i_n} \right) \right\} \end{aligned}$$

and

$$B_2 := \frac{1}{N\sqrt{N}} \sum_{\mathbf{i}} \sigma_{\mathbf{i}}^2 \mathbb{E} \left\{ R_{i_1 i_d} \left(\prod_{k=2, k \neq \ell}^{d-1} u_{i_k} \right) \frac{1}{(d-2)\mu \sigma_{\mathbf{i}}^2 \sqrt{N}} \left(\prod_{j=1}^d u_{i_j} \right) u_{i_\ell} \right\}$$

$$\begin{aligned}
&= \frac{1}{N^2(d-2)} \sum_i \mathbb{E} \left\{ \frac{1}{\mu} R_{i_1 i_d} \left(\prod_{k=2}^{d-1} u_{i_k} \right) \left(\prod_{j=1}^d u_{i_j} \right) \right\} \\
&= \frac{1}{N^2(d-2)} \mathbb{E} \left\{ \frac{1}{\mu} \frac{1}{\mu - z} \right\} \\
&= O(1/N^2).
\end{aligned}$$

It is not hard to see that $B_1 = O(1/N)$, since only one trace (that of \bar{R} when $m = \ell$) arises from the contractions. It thus follows that $B = O(1/N)$.

4.4.3 CONCENTRATION AND ASYMPTOTIC EXPRESSION FOR THE NORMALIZED TRACE

Relying on the same arguments used in the analytic proof of Wigner's semicircle law, it is not difficult to show that, for any z outside the real line, the (analytic) expected normalized trace $\mathbb{E} \left\{ \frac{1}{N} \text{tr} R(z) \right\}$ converges to a function of z

$$m_d(z) := \lim_{N \rightarrow \infty} \mathbb{E} \left\{ \frac{1}{N} \text{tr} R \right\}, \quad \forall z \in \mathbb{C} \setminus \mathbb{R},$$

which is also analytic by Vitali's theorem. The first ingredient is the inequality $\|R^p(z)\| \leq 1/|\Im(z)|^p$, valid for all such z and all integers p , which entails that all moments $\frac{1}{N} \text{tr} R^p$ are bounded on $\mathbb{C} \setminus \mathbb{R}$ with probability one, thus ensuring the existence of convergent subsequences. Then, it can be shown that every such subsequence must actually converge to a same limit $m_d(z)$, which by (19) must then satisfy

$$m_d(z) = -\frac{1}{z} - \frac{1}{z d(d-1)} \lim_{N \rightarrow \infty} \mathbb{E} \left\{ \left(\frac{1}{N} \text{tr} R \right)^2 \right\}. \quad (23)$$

It is now a standard exercise to show that $\frac{1}{N} \text{tr} R$ concentrates around $m_d(z)$, and hence the above moment can be replaced by $(m_d(z))^2$. This can be done by resorting to the Nash-Poincaré inequality (see Pastur and Shcherbina, 2011) to bound the variance of the normalized trace as in

$$\text{Var} \left(\frac{1}{N} \text{tr} R \right) \leq \frac{1}{N^2} \mathbb{E} \left\{ \sum_i \sigma_i^2 \left(\sum_p \frac{\partial R_{pp}}{\partial W_i} \right)^2 \right\}. \quad (24)$$

Now, denoting the (i, j) component of R^2 by R_{ij}^2 , we can write

$$\sum_p \frac{\partial R_{pp}}{\partial W_i} = -\frac{1}{\sqrt{N}} \sum_j R_{j_1 j_d}^2 \left(\frac{\partial W_j}{\partial W_i} \right) \left(\prod_{k=2}^{d-1} u_{j_k} \right) - \sum_j \sum_{m=2}^{d-1} R_{j_1 j_d}^2 \mathbf{y}_j \left(\frac{\partial u_{j_m}}{W_i} \right) \left(\prod_{k=2, k \neq m}^{d-1} u_{j_k} \right),$$

Using this expression, it is tedious but not hard to verify that the expectation in (24) is $O(1)$. This essentially comes from the fact that each term in that expectation comprises a factor $1/N$ and at most a single trace (of R^4). Because all moments $\frac{1}{N} \text{tr} R^p$ are bounded on $\mathbb{C} \setminus \mathbb{R}$, it follows that $\text{Var} \left(\frac{1}{N} \text{tr} R(z) \right) = O(1/N^2)$ for $z \in \bar{\mathbb{C}}$, and thus by a standard

combination of Markov's inequality and Borel-Cantelli we have that $\frac{1}{N} \operatorname{tr}(R(z)) \xrightarrow{\text{a.s.}} m_d(z)$ on that set. Hence, from (23) we deduce that, as $N \rightarrow \infty$,

$$\frac{1}{z d(d-1)} (m_d(z))^2 + z m_d(z) + 1 = 0, \quad (25)$$

thus yielding the solutions

$$m_d(z) = \frac{-d(d-1) \pm \sqrt{d^2(d-1)^2 z^2 - 4d(d-1)}}{2}.$$

At last, using our definition $\beta_d = 2/\sqrt{d(d-1)}$ and keeping the solution corresponding to a Stieltjes transform (behaving as $O(z^{-1})$ at infinity), we arrive at

$$m_d(z) = \frac{2}{\beta_d^2} \left(-z + z \sqrt{1 - \frac{\beta_d^2}{z^2}} \right),$$

which is the same expression of Theorem 2. Hence, by the Stieltjes continuity theorem (see Appendix C), the spectrum of $\mathbf{Y} \cdot u^{d-2}$ converges weakly almost surely to a semi-circle law supported on $[-\beta_d, \beta_d]$.

Let us pause for a moment to reflect on the implications of this intermediate result. Since the almost sure limit $\mu_{d,\infty}(\lambda)$ of μ obtained by our calculations must be such that $\mu_{d,\infty}(\lambda)/(d-1)$ is not in the limiting spectrum of $\mathbf{Y} \cdot u^{d-2}$ (because of the factor $R(\mu/(d-1))$ in the formulas for the derivatives of μ and u), and since the limiting spectral measure of this matrix is a semicircle law supported on $[-\beta_d, \beta_d]$, it follows that we must have $\mu_{d,\infty}(\lambda) > (d-1)\beta_d$ for our formulas to make sense—just as we assume in the statement of Theorem 3.

Incidentally, the observation of Remark 1, combined with the limiting spectrum of $\mathbf{Y} \cdot u^{d-2}$, implies that the existence of local maxima associated with negative eigenvalues μ becomes increasingly unlikely, since the symmetric support of the limiting spectral measure $[-\beta_d, \beta_d]$ is incompatible with the condition that all other eigenvalues must lie in $]-\infty, \mu/(d-1)[$. This observation is coherent with the results of Ben Arous et al. (2019), which show that the expected number of local maxima associated to negative eigenvalues decays exponentially as $N \rightarrow 0$.

4.4.4 A MOMENTARY DIGRESSION: EXTENDING THE ABOVE PROOF TO THEOREM 2

The proof of Theorem 3 can be obtained by adapting the above derivation. The idea is to set $\lambda = 0$ and replace u by a deterministic vector $v \in \mathbb{S}^{N-1}$ in (18):

$$Q = -\frac{1}{z} I + \frac{1}{z \sqrt{N}} (\mathbf{W} \cdot v^{d-2}) Q,$$

where we use Q to denote this new resolvent to avoid any ambiguities. We are as before interested in the expected normalized trace of Q ,

$$\mathbb{E} \left\{ \frac{1}{N} \operatorname{tr} Q \right\} = -\frac{1}{z} + \frac{1}{z N \sqrt{N}} \sum_{\mathbf{i}} \mathbb{E} \left\{ W_{\mathbf{i}} \left(\prod_{k=2}^{d-1} v_{i_k} \right) Q_{i_1 i_d} \right\}. \quad (26)$$

Since now v does not depend on \mathbf{W} , the calculations are greatly simplified: the derivative of Q is

$$\frac{\partial Q_{pq}}{\partial W_i} = -\frac{1}{\sqrt{N}} \sum_j Q_{pj1} \left(\frac{\partial W_j}{\partial W_i} \right) \left(\prod_{k=2}^{d-1} v_{j_k} \right) Q_{j_d q},$$

and therefore the application of Stein's lemma to compute the expectation in the right-hand side of (26) gives

$$\sum_i \mathbb{E} \left\{ W_i \left(\prod_{k=2}^{d-1} v_{i_k} \right) Q_{i_1 i_d} \right\} = -\frac{1}{\sqrt{N}} \sum_i \sum_j \sigma_i^2 \mathbb{E} \left\{ Q_{i_1 j_1} \left(\frac{\partial W_j}{\partial W_i} \right) \left(\prod_{k=2}^{d-1} v_{i_k} v_{j_k} \right) Q_{j_d i_d} \right\}.$$

Here again, by considering all possible matchings of the indices \mathbf{i} and \mathbf{j} , we conclude that

$$\lim_{N \rightarrow \infty} \frac{1}{N\sqrt{N}} \sum_i \mathbb{E} \left\{ W_i \left(\prod_{k=2}^{d-1} v_{i_k} \right) Q_{i_1 i_d} \right\} = -\frac{1}{d(d-1)} \lim_{N \rightarrow \infty} \mathbb{E} \left\{ \left(\frac{1}{N} \operatorname{tr} Q \right)^2 \right\},$$

since all other terms are $o(1)$. In particular, contractions of Q with v are $O(1)$ in expectation, since by Cauchy-Schwarz $|\langle Qv, v \rangle| \leq \|Qv\| \|v\| \leq \|Q\| \leq 1/|\Im(z)|$ for all $z \in \mathbb{C} \setminus \mathbb{R}$. We thus get the same expression as in the study of R .

Finally, the same reasoning followed in the previous subsection to show concentration of $\frac{1}{N} \operatorname{tr} R$ can be mimicked to show concentration of $\frac{1}{N} \operatorname{tr} Q$, thus leading to the same asymptotic equation as before, namely, to (25). This shows that the spectrum of $\frac{1}{\sqrt{N}} \mathbf{W} \cdot v^{d-2}$ also converges weakly almost surely to a semicircle law supported on $[-\beta_d, \beta_d]$, as claimed. In retrospect, we can see that the additional terms that are present in the computation of $\frac{1}{N} \operatorname{tr} R$ (originating from the dependence of u on \mathbf{W}) are all $o(1)$, and thus the resulting spectral measures are the same in both cases.

4.5 Asymptotic fixed-point equation

We turn back now to the proof of Theorem 3, by invoking the key assumption that⁷

$$\begin{aligned} \mu &\xrightarrow{\text{a.s.}} \mu_{d,\infty}(\lambda) > (d-1)\beta_d, \\ \langle x, u \rangle &\xrightarrow{\text{a.s.}} \alpha_{d,\infty}(\lambda) > 0 \end{aligned}$$

as $N \rightarrow \infty$. In particular, the first above statement guarantees that the derived derivative formulas hold as $N \rightarrow \infty$, since we know from the previous subsection that the spectrum of $\mathbf{Y} \cdot u^{d-2}$ converges to a semi-circle law on $[-\beta_d, \beta_d]$, and hence $\mu_{d,\infty}(\lambda)/(d-1)$ does not fall on this support. Hence, by inspection of (16)–(17) we have

$$\alpha_{d,\infty}(\lambda) = \frac{\lambda}{\mu_{d,\infty}(\lambda)} (\alpha_{d,\infty}(\lambda))^{d-1} - \frac{1}{d} \frac{\alpha_{d,\infty}(\lambda)}{\mu_{d,\infty}(\lambda)} m_d \left(\frac{\mu_{d,\infty}(\lambda)}{d-1} \right) \quad (27)$$

7. The reader will note that this assumption implicitly imposes that $\lim_{N \rightarrow \infty} \frac{1}{N} \operatorname{tr} R = m_d(z)$ not only outside the real line, but also outside the support $[-\beta_d, \beta_d]$, since the formulas (16)–(17) both involve m_d evaluated on the *real* (random) quantity $\mu/(d-1)$, which by hypothesis converges to a number $\mu_{d,\infty}$ outside the support.

$$\mu_{d,\infty}(\lambda) = \lambda (\alpha_{d,\infty}(\lambda))^d - \frac{1}{d-1} m_d \left(\frac{\mu_{d,\infty}(\lambda)}{d-1} \right).$$

Since $\alpha_{d,\infty}(\lambda) > 0$, we get from (27)

$$\mu_{d,\infty}(\lambda) = \lambda (\alpha_{d,\infty}(\lambda))^{d-2} - \frac{1}{d} m_d \left(\frac{\mu_{d,\infty}(\lambda)}{d-1} \right)$$

Solving now for $\alpha_{d,\infty}(\lambda)$, we get

$$\alpha_{d,\infty}(\lambda) = \left[\frac{1}{\lambda} \left(\mu_{d,\infty}(\lambda) + \frac{1}{d} m_d \left(\frac{\mu_{d,\infty}(\lambda)}{d-1} \right) \right) \right]^{\frac{1}{d-2}}.$$

The equations (7)–(8) are thus established, and the proof is complete.

Acknowledgments

This work was supported by the MIAI LargeDATA Chair at University Grenoble Alpes led by R. Couillet, and also by the Labex CIMI project of the first author.

We would like to thank Gérard Ben-Arous, Mohamed Seddik, Yang Qi and Arthur Marmin for helpful discussions on the subject. Finally, the anonymous Referees are also thanked for their constructive suggestions which led to significant improvements in the manuscript.

Appendix A. Derivative of tensor eigenpairs

Given an eigenpair (μ, u) , the derivatives of μ and u with respect to any entry $W_{i_1 \dots i_d}$ can be computed by resorting to the implicit value theorem (see, e.g., Krantz and Parks, 2013). To this end, we start from the system of equations that characterizes the eigenpair (μ, u) :

$$F(\mu, u, \mathbf{W}) = \begin{pmatrix} \mathbf{Y} u^{d-1} - \mu u \\ 1 - \langle u, u \rangle \end{pmatrix} = 0.$$

The Jacobian of F with respect to (μ, u) is

$$J(\mu, u, \mathbf{W}) = \begin{pmatrix} (d-1)\mathbf{Y} u^{d-2} - \mu I & -u \\ -2u^\top & 0. \end{pmatrix}$$

Now, if $\frac{\mu}{d-1}$ is not in the spectrum of $\mathbf{Y} u^{d-2}$, then the first block of the above Jacobian is invertible, and its inverse is simply $(d-1)R\left(\frac{\mu}{d-1}\right)$, where $R(z) = (\mathbf{Y} u^{d-2} - zI)^{-1}$. For simplicity of notation, we will write $\bar{R} := R\left(\frac{\mu}{d-1}\right)$ in the following. In particular, this implies

$$\bar{R}u = \frac{1}{\mu - \frac{\mu}{d-1}} u = \frac{d-1}{(d-2)\mu} u.$$

Defining also

$$A := (d-1)\bar{R}^{-1}, \quad B := -u, \quad \text{and} \quad C := -2u^\top$$

and assuming that $\frac{\mu}{(d-1)}$ is not in the spectrum of $\mathbf{Y} u^{d-2}$, the Schur complement formula yields

$$J^{-1}(\mu, u, \mathbf{W}) = \begin{pmatrix} A^{-1} + A^{-1}BS^{-1}CA^{-1} & -A^{-1}BS^{-1} \\ -S^{-1}CA^{-1} & S^{-1}, \end{pmatrix}$$

where

$$S = -CA^{-1}B = -\frac{2}{d-1}u^T \bar{R}u = -\frac{2}{(d-2)\mu}.$$

Hence,

$$J^{-1}(\mu, u, \mathbf{W}) = \begin{pmatrix} \frac{1}{d-1}\bar{R} - \frac{1}{(d-2)\mu}uu^T & -\frac{1}{2}u \\ -u^T & -\frac{d-2}{2}\mu \end{pmatrix}.$$

Now, by the implicit function theorem we have

$$\begin{pmatrix} \frac{\partial u}{\partial W_{i_1 \dots i_d}} \\ \frac{\partial \mu}{\partial W_{i_1 \dots i_d}} \end{pmatrix} = -J^{-1}(\mu, u, \mathbf{W}) \frac{\partial F}{\partial W_{i_1 \dots i_d}}(\mu, u, \mathbf{W}).$$

The required derivative $\frac{\partial F}{\partial W_{i_1 \dots i_d}}$ turns out to be

$$\frac{\partial F}{\partial W_{i_1 \dots i_d}}(\mu, u, \mathbf{W}) = \begin{pmatrix} \frac{1}{\sqrt{N}}\phi \\ 0 \end{pmatrix},$$

where ϕ is an N -dimensional vector with at most d nonzero entries given by

$$\sigma_{W_{i_1 \dots i_d}}^2 \phi = \frac{1}{d} \sum_{j=1}^d u_1 \dots u_{i_{j-1}} u_{i_{j+1}} \dots u_d e^{(i_j)}. \quad (28)$$

Hence, provided that $\frac{\mu}{d-1}$ is not an eigenvalue of $\mathbf{Y} u^{d-2}$, we finally get the derivatives

$$\begin{pmatrix} \frac{\partial u}{\partial W_{i_1 \dots i_d}} \\ \frac{\partial \mu}{\partial W_{i_1 \dots i_d}} \end{pmatrix} = \begin{pmatrix} -\frac{1}{(d-1)\sqrt{N}}\bar{R}\phi + \frac{1}{(d-2)\mu} \frac{\partial \mu}{\partial W_{i_1 \dots i_d}} u \\ \frac{1}{\sigma_{w_{i_1 \dots i_d}}^2 \sqrt{N}} \prod_{j=1}^d u_{i_j} \end{pmatrix}. \quad (29)$$

For instance, for $d = 3$ we have

$$\begin{pmatrix} \frac{\partial u}{\partial W_{iml}} \\ \frac{\partial \mu}{\partial W_{iml}} \end{pmatrix} = \begin{pmatrix} -\frac{1}{2\sqrt{N}}\bar{R}\phi + \frac{1}{\mu} \frac{\partial \mu}{\partial W_{iml}} u \\ \frac{1}{\sigma_{W_{iml}}^2 \sqrt{N}} u_i u_m u_l \end{pmatrix}.$$

Appendix B. Proof of Lemma 7

For every triple $T = (\mu, u, \mathbf{W})$ such that the Jacobian $J(\mu, u, \mathbf{W})$ of F (see Appendix A) with respect to (μ, u) is nonsingular, by the implicit function theorem there exists a neighborhood of T such that (μ, u) is given by a (unique) continuously differentiable function of \mathbf{W} whose derivative satisfies (29). Hence, provided that the set

$$B = \{\mathbf{W} \in \mathcal{S}^d(N) : J(\mu, u, \mathbf{W}) \text{ is singular } \forall (\mu, u) \text{ satisfying } F(\mu, u, \mathbf{W}) = 0\}$$

has measure zero, then for almost every $\mathbf{W} \in \mathcal{S}^d(N)$ there exists at least one couple (μ, u) such that the theorem can be applied to (μ, u, \mathbf{W}) . It follows that one can construct (at least) one function $G : \mathcal{S}^d(N) \rightarrow \mathbb{R}^{N+1}$ defined and continuously differentiable almost everywhere on $\mathcal{S}^d(N)$ such that $G(\mathbf{W}) = (\mu, u)$ is an eigenpair of \mathbf{W} , with $u \in \mathbb{S}^{N-1}$.

Hence, we are left with the proof that the above set B indeed has measure zero. In particular, every tensor $\mathbf{W} \in \mathcal{S}^d(N)$ having at least one eigenpair (μ, u) such that $\mu/(d-1)$ is not an eigenvalue of the contraction $\mathbf{W} \cdot u^{d-2}$ cannot belong to B , since the resolvent of such contraction is well defined on $\mu/(d-1)$, and hence we can apply the explicit expression (in terms of that resolvent) for the inverse of $J(\mu, u, \mathbf{W})$ given in Appendix A. But it turns out that the set

$$B' = \left\{ \mathbf{W} \in \mathcal{S}^d(N) : \mathbf{W} \text{ has an eigenpair } (\mu, u) \text{ such that } \frac{\mu}{d-1} \text{ is an eigenvalue of } \mathbf{W} \cdot u^{d-2} \right\},$$

which contains B , has measure zero. To show this, it suffices to construct an example of a tensor \mathbf{W}_0 not in B' , since this implies the existence of a non-empty open set of tensors U (a neighborhood of \mathbf{W}_0) such that $U \cap B' = \emptyset$. In other words, B' would then be proper Zariski closed in $\mathcal{S}^d(N)$, and hence of zero measure. Such an example can be easily constructed. For instance, take any rank-one tensor $v^{\otimes d}$ with $v \in \mathbb{S}^{N-1}$: $(1, v)$ is clearly an eigenpair of \mathbf{W} , but $1/(d-1)$ is not an eigenvalue of $v \otimes v$.

Appendix C. Main tools in Pastur's Stein approach

We describe in the following the main tools used in the so-called Pastur's Stein approach for the study of random matrices, which mainly revolve around the resolvent formalism. In this description, we will denote the N eigenvalues of a symmetric matrix $Y \in \mathbb{R}^{N \times N}$ by $\mu_1(Y) \leq \dots \leq \mu_N(Y)$, and its spectral decomposition by

$$Y = \sum_{i=1}^N \mu_i(Y) u_i(Y) u_i^\top(Y),$$

where the dependence of the eigenvalues and eigenvectors on Y will be omitted whenever no ambiguity arises, for the sake of simplicity.

C.1 The resolvent matrix and its properties

Let us start by defining the resolvent of a symmetric matrix.

Definition 8 (Resolvent) *Given a symmetric matrix $Y \in \mathbb{R}^{N \times N}$, we define its resolvent matrix as*

$$R_Y(z) := (Y - zI)^{-1}, \quad z \in \mathbb{C} \setminus \sigma(Y),$$

where $\sigma(Y)$ stands for the spectrum of Y .

The resolvent is a rich object that encodes spectral properties of Y in its complex-analytic structure. Among its various properties, we can mention the following:

- (i) it satisfies the recurrence relations

$$R_Y(z) = -\frac{1}{z}I + \frac{1}{z}YR_Y(z) = -\frac{1}{z}I + \frac{1}{z}R_Y(z)Y;$$

- (ii) its spectral decomposition $R_Y(z) = \sum_{i=1}^N \frac{1}{\mu_i - z} u_i u_i^\top$ is a meromorphic function with poles on the eigenvalues of Y ;
- (iii) its trace $\text{tr} R(z) = \sum_{i=1}^N \frac{1}{\mu_i - z}$ is a complex analytic function on $\mathbb{C} \setminus \sigma(Y)$;
- (iv) as a consequence of (ii), the contour integral

$$P_{\mu_i} = -\frac{1}{2\pi i} \oint_{C_{\mu_i}} R_N(z) dz$$

yields a projector onto the eigenspace associated with μ_i , where C_{μ_i} denotes a sufficiently small positively oriented contour around eigenvalue μ_i alone.

Property (iv) is often utilized to get the alignment of an arbitrary vector x with the eigenvector u of an eigenvalue μ of multiplicity one, via:

$$\langle u, x \rangle^2 = x^\top P_\mu x = -\frac{1}{2\pi i} \oint_{C_\mu} x^\top R_Y(z) x dz.$$

In particular, when studying a sequence of random $N \times N$ matrices of growing dimension N , the above formula is useful for computing the asymptotic alignment of a given sequence of vectors x with the sequence of eigenvectors u whose eigenvalues converge to a limit μ which is bounded away from the support of the (rest of the) limiting spectrum.

C.2 The Stieltjes transform

One typical question that arises in the study of random matrices is whether the empirical spectral distribution

$$\rho_{Y_N} := \frac{1}{N} \sum_{i=1}^N \delta_{\mu_i(Y_N)}, \tag{30}$$

where δ_μ is a Dirac mass on μ , of a sequence $\{Y_N\}$ of random $N \times N$ matrices converges to some deterministic probability measure ρ . One widely used method that allows answering this question is based upon an analytic tool known as the Stieltjes transform, defined next.

Definition 9 (Stieltjes transform) *Given a probability measure ρ , its Stieltjes transform is defined by*

$$s_\rho(z) := \int \frac{1}{t-z} \rho(dt), \quad z \in \mathbb{C} \setminus \text{supp}(\rho).$$

The usefulness of this tool in RMT is explained by the following result, discussed e.g. by Tao (2012) (see Exercise 1.1.25 of that book).

Theorem 10 (Stieltjes' continuity theorem) *The sequence of random probability measures ρ_k supported on \mathbb{R} converges almost surely weakly⁸ to the deterministic measure ρ if and only if $s_{\rho_k}(z) \xrightarrow{a.s.} s_\rho(z)$ for all complex z in the upper-half plane $\overline{\mathbb{C}} := \mathbb{R} + i\mathbb{R}_+$.*

It should be noted that the same result holds also for other forms of convergence apart from almost sure (for instance, convergence in probability). In any case, the above posed question of convergence of $\rho_{Y_N} \xrightarrow{?} \rho$ can be translated into a question of pointwise convergence of a sequence of Stieltjes transforms. Furthermore, if ρ has a density f , then having determined the limiting transform $s_\rho(z)$ this density can be recovered by the following identity (often termed Sokhotski–Plemelj formula):

$$f(x) = \frac{1}{\pi} \lim_{\epsilon \rightarrow 0^+} \Im\{s_\rho(x + i\epsilon)\};$$

see, e.g., (Tao, 2012).

In the particular case of a sequence of probability measures (30), it can be seen that

$$s_{\rho_{Y_N}}(z) = \frac{1}{N} \sum_{i=1}^N \frac{1}{\mu_i - z} = \frac{1}{N} \text{tr} R_{Y_N}(z),$$

which reinforces the central role played by the resolvent formalism.

C.3 Gaussian integration by parts (Stein's lemma)

The last technical tool that we will review is Stein's lemma, also known as Stein's identity or Gaussian integration by parts, which allows replacing the expected product of a Gaussian variable with a differentiable function f (of polynomially bounded growth) by the variance of that variable times the expectation of f' .

Lemma 11 (Stein, 1981) *If $x \sim \mathcal{N}(0, \sigma^2)$ and $f : \mathbb{R} \rightarrow \mathbb{R}$ is continuously differentiable almost everywhere and has at most polynomial growth, then*

$$\mathbb{E}\{x f(x)\} = \sigma^2 \mathbb{E}\{f'(x)\}.$$

8. While this result is often stated in terms of convergence in the vague topology rather than the weak topology, these are equivalent when the limiting measure ρ is assumed to be a probability measure, since no mass is lost.

Appendix D. Maple solution of fixed-point equation for $d = 3$

By simply defining on Maple:

$$\begin{aligned} h(z) &:= \text{sqrt}(z^2 - 2/3) \\ \omega(z) &:= (1/2*z + h(1/2*z))/\lambda \\ \phi(z) &:= \lambda*\omega(z)^3 + (3/4)*z - (3/2)*h(z/2) \end{aligned}$$

(note that $m_3(z) = -3z + 3h(z)$ with $h(z) = \sqrt{z^2 - 2/3}$), one can find the solutions of $z = \phi(z, \lambda)$ for the fixed-point equation in the case $d = 3$ by using the command:

$$\text{solve}(z = \phi(z), z).$$

This produced four solutions:

$$z_{1,2}^*(\lambda) = \frac{\sqrt{18\lambda^4 + 72\lambda^2 \pm 2\sqrt{3} \sqrt{\lambda^2(3\lambda^2 - 4)^3}}}{6\lambda}, \quad z_3^*(\lambda) = -z_1^*(\lambda) \quad \text{and} \quad z_4^*(\lambda) = -z_2^*(\lambda).$$

We now show that the solution $z_1^*(\lambda)$ corresponds precisely to the formula (12) for the case $\lambda > \lambda_c(d)$. As for the other solutions, z_3^* and z_4^* are negative and thus not of our interest, while z_2^* holds only for negative values of λ . Letting $\theta = \lambda\sqrt{3}$, we get

$$\mu_3^*(\lambda) = \frac{1}{\sqrt{6}} \frac{\theta^2 + \theta\sqrt{\theta^2 - 4} + 4}{\sqrt{\theta^2 + \theta\sqrt{\theta^2 - 4}}} \quad \text{and} \quad z_1^*(\lambda) = \frac{1}{\sqrt{6}} \sqrt{\theta^2 + 12 + \frac{1}{\theta}(\theta^2 - 4)^{3/2}}.$$

Defining also $\omega = \sqrt{\theta^2 - 4} = \sqrt{3\lambda^2 - 4}$, we can write the squares of these functions as

$$[\mu_3^*(\lambda)]^2 = \frac{1}{24\theta} (\theta^2 + \theta\omega + 4)^2 (\theta - \omega)$$

and

$$[z_1^*(\lambda)]^2 = \frac{1}{6\theta} (\theta^3 + \theta^2\omega + 12\theta - 4\omega).$$

Finally, by developing the above expressions and replacing ω^2 by $(\theta^2 - 4)$ to obtain a polynomial in (θ, ω) of partial degree 1 in ω , one can easily show that $[\mu_3^*(\lambda)]^2 = [z_1^*(\lambda)]^2$. Since both formulas yield positive values, this implies $\mu_3^*(\lambda) = z_1^*(\lambda)$, as claimed.

References

- M. S. Advani, A. M. Saxe, and H. Sompolinsky. High-dimensional dynamics of generalization error in neural networks. *Neural Networks*, 132:428–446, 2020.
- A. Anandkumar, R. Ge, D. Hsu, S. M. Kakade, and M. Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15:2773–2832, 2014.
- B. Au and J. Garza-Vargas. Spectral asymptotics for contracted tensor ensembles. *arXiv preprint arXiv:2110.01652*, 2021.

- Z. Bai and J. W. Silverstein. *Spectral analysis of large dimensional random matrices*. Springer New York, NY, 2nd edition, 2010.
- J. Baik, G. Ben Arous, and S. Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Annals of Probability*, 33(5):1643–1697, 2005.
- G. Ben Arous, S. Mei, A. Montanari, and M. Nica. The landscape of the spiked tensor model. *Communications on Pure and Applied Mathematics*, 72(11):2282–2330, 2019.
- G. Ben Arous, D. Z. Huang, and J. Huang. Long random matrices and tensor unfolding. *arXiv preprint arXiv:2110.10210*, 2021.
- F. Benaych-Georges and R. R. Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, 2011.
- R. Bro, J. J. Workman Jr., P. R. Mobley, and B. R. Kowalski. Review of chemometrics applied to spectroscopy: 1985-95, Part 3 – Multi-way analysis. *Applied Spectroscopy Reviews*, 32(3):237–261, 1997.
- W.-K. Chen. Phase transition in the spiked random tensor with Rademacher prior. *The Annals of Statistics*, 47(5):2734–2756, 2019.
- W.-K. Chen, M. Handschy, and G. Lerman. Phase transition in random tensors with multiple independent spikes. *The Annals of Applied Probability*, 31(4):1868 – 1913, 2021.
- A. Chevreuil and P. Loubaton. On the non-detectability of spiked large random tensors. In *Proceedings of the IEEE Workshop on Statistical Signal Processing (SSP)*, pages 443–447, Freiburg, Germany, June 2018.
- R. Couillet and F. Benaych-Georges. Kernel spectral clustering of large dimensional data. *Electronic Journal of Statistics*, 10(1):1393–1454, 2016.
- R. Ge, F. Huang, C. Jin, and Y. Yuan. Escaping from saddle points — online stochastic gradient for tensor decomposition. In *Proceedings of the 28th Conference on Learning Theory (COLT)*, volume 40 of *Proceedings of Machine Learning Research*, pages 797–842, Paris, France, July 2015.
- J. H. de M. Goulart, R. Couillet, and P. Comon. A random matrix perspective on random tensor models. Seminar in Random Tensors, Texas A&M University, February 2021a.
- J. H. de M. Goulart, R. Couillet, and P. Comon. A random matrix perspective on random spiked tensors. SIAM Conference on Applied Linear Algebra, May 2021b.
- W. Hackbusch. *Tensor Spaces and Numerical Tensor Calculus*. Series in Computational Mathematics. Springer Berlin, Heidelberg, 2012.
- C. J. Hillar and L.-H. Lim. Most tensor problems are NP-hard. *Journal of the ACM*, 60(6):45:1–45:39, 2013.

- F. L. Hitchcock. The expression of a tensor or a polyadic as a sum of products. *Journal of Mathematics and Physics*, 6(1-4):164–189, 1927.
- S. Hopkins, J. Shi, and D. Steurer. Tensor principal component analysis via sum-of-square proofs. In *Proceedings of the 28th Conference on Learning Theory (COLT)*, volume 40 of *Proceedings of Machine Learning Research*, pages 956–1006, Paris, France, July 2015.
- B. Huang, C. Mu, D. Goldfarb, and J. Wright. Provable models for robust low-rank tensor completion. *Pacific Journal of Optimization*, 11(2):339–364, 2015.
- A. Jagannath, P. Lopatto, and L. Miolane. Statistical thresholds for Tensor PCA. *The Annals of Applied Probability*, 30(4):1910–1933, 2020.
- P. Jain and S. Oh. Provable tensor factorization with missing data. In *Advances in Neural Information Processing Systems (NIPS)*, volume 27, pages 1431–1439, Montréal, Canada, December 2014.
- A. Kammoun and M.-S. Alouini. On the precise error analysis of support vector machines. *IEEE Open Journal of Signal Processing*, 2:99–118, 2021.
- T. G. Kolda and J. R. Mayo. Shifted power method for computing tensor eigenpairs. *SIAM Journal on Matrix Analysis and Applications*, 32(4):1095–1124, 2011.
- S. G. Krantz and H. R. Parks. *The implicit function theorem: History, theory, and applications*. Birkhäuser New York, NY, 2013.
- J. M. Landsberg. *Tensors: Geometry and Applications*, volume 128 of *Graduate Studies in Mathematics*. American Mathematical Society, 2012.
- T. Lesieur, L. Miolane, M. Lelarge, F. Krzakala, and L. Zdeborova. Statistical and computational phase transitions in spiked tensor estimation. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, pages 511–515, Aachen, Germany, June 2017.
- Z. Liao and R. Couillet. The dynamics of learning: A random matrix approach. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, pages 3072–3081, July 2018.
- Z. Liao and R. Couillet. A large dimensional analysis of least squares support vector machines. *IEEE Transactions on Signal Processing*, 67(4):1065–1074, 2019.
- L.-H. Lim. 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)*, pages 129–132, Puerto Vallarta, Mexico, December 2005.
- X. Mai and R. Couillet. A random matrix analysis and improvement of semi-supervised learning for large dimensional data. *Journal of Machine Learning Research*, 19(1):3074–3100, 2018.
- X. Mai and R. Couillet. Consistent semi-supervised graph regularization for high dimensional data. *Journal of Machine Learning Research*, 22(94):1–48, 2021.

- X. Mai, Z. Liao, and R. Couillet. A large scale analysis of logistic regression: Asymptotic performance and new insights. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3357–3361, Brighton, UK, May 2019.
- A. Montanari and E. Richard. A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems (NIPS)*, volume 27, pages 2897–2905, Montréal, Canada, December 2014.
- A. Montanari, D. Reichman, and O. Zeitouni. On the limitation of spectral methods: From the Gaussian hidden clique problem to rank-one perturbations of Gaussian tensors. In *Advances in Neural Information Processing Systems*, volume 28, pages 227–225, Montréal, Canada, December 2015.
- M. Mørup. Applications of tensor (multiway array) factorizations and decompositions in data mining. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 1(1):24–40, 2011.
- R. R. Nadakuditi and M. E. J. Newman. Graph spectra and the detectability of community structure in networks. *Physical Review Letters*, 108(18):188701, 2012.
- J. Nocedal and S. J. Wright. *Numerical optimization*. Springer New York, NY, 2nd edition, 2006.
- L. A. Pastur and M. Shcherbina. *Eigenvalue distribution of large random matrices*, volume 171 of *Mathematical surveys and monographs*. American Mathematical Society, 2011.
- A. Perry, A. S. Wein, and A. S. Bandeira. Statistical limits of spiked tensor models. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 56(1):230–264, 2020.
- L. Qi and Z. Luo. *Tensor analysis: Spectral theory and special tensors*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2017.
- 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, 2017.
- C. M. Stein. Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics*, 9(6):1135–1151, 1981.
- T. Tao. *Topics in random matrix theory*, volume 132 of *Graduate Studies in Mathematics*. American Mathematical Society, 2012.
- M. Tiomoko, H. T. Ali, and R. Couillet. Deciphering and optimizing multi-task learning: a random matrix approach. In *Proceedings of the International Conference on Learning Representations (ICLR)*, Addis Ababa, Ethiopia, April 2020.
- L. R. Tucker. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31(3):279–311, 1966.