

Estimation in the spiked Wigner model: A short proof of the replica formula

Ahmed El Alaoui* Florent Krzakala†

Abstract

We consider the problem of estimating a rank-one perturbation of a Wigner matrix in a setting of low signal-to-noise ratio. This serves as a simple model for principal component analysis in high dimensions. The mutual information per variable between the spike and the observed matrix, or equivalently, the normalized Kullback-Leibler divergence between the planted and null models are known to converge to the so-called *replica-symmetric* formula, the properties of which determine the fundamental limits of estimation in this model. We provide in this note a short and transparent proof of this formula, based on simple executions of Gaussian interpolations and standard concentration-of-measure arguments. The *Franz-Parisi potential*, that is, the free entropy at a fixed overlap, plays an important role in our proof. Furthermore, our proof can be generalized straightforwardly to spiked tensor models of even order.

1 Introduction

Extracting low-rank information from a data matrix corrupted with noise is a fundamental statistical task. Spiked random matrix models have attracted considerable attention in statistics, probability and machine learning as rich testbeds for theoretical investigation on this problem [Joh01, Péc14, Péc06, BAP05]. A basic such model is the *spiked Wigner model* in which one observes a rank-one deformation of a Wigner matrix \mathbf{W} :

$$\mathbf{Y} = \sqrt{\frac{\lambda}{N}} \mathbf{x}^* \mathbf{x}^{*\top} + \mathbf{W}, \quad (1)$$

where $W_{ij} = W_{ji} \sim \mathcal{N}(0, 1)$ and $W_{ii} \sim \mathcal{N}(0, \sigma^2)$ are independent for all $1 \leq i \leq j \leq N$. The *spike* vector $\mathbf{x}^* \in \mathbb{R}^N$ represents the signal to be recovered, and $\lambda \geq 0$ plays the role of a Signal-to-Noise Ratio (SNR) parameter. The entries of \mathbf{x}^* come i.i.d. from a (Borel) prior $P_{\mathbf{x}}$ on \mathbb{R} with bounded support, so that the scaling in the above model puts the problem in a high-noise regime where only partial recovery of the spike is possible. A basic statistical question about this model is *for what values of the SNR λ is it possible to estimate the spike \mathbf{x}^* with non-trivial accuracy?* Spectral methods, or more precisely, estimation using the top eigenvector of \mathbf{Y} , are known to succeed above a *spectral threshold* and fail below [BGN11]. Since the posterior mean is the estimator with minimal mean squared error, this question boils down to the study of the posterior distribution of \mathbf{x}^* given \mathbf{Y} , which by Bayes' rule, can be written as

$$d\mathbb{P}_{\lambda}(\mathbf{x}|\mathbf{Y}) = \frac{e^{-H(\mathbf{x})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x})}{\int e^{-H(\mathbf{x})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x})}, \quad (2)$$

*Department of EECS, UC Berkeley, CA. Email: elalaoui@berkeley.edu

†Laboratoire de Physique Statistique, CNRS, PSL Universités & Ecole Normale Supérieure, Sorbonne Universités et Université Pierre & Marie Curie, Paris, France.

where H is the (random) Hamiltonian

$$\begin{aligned} -H(\mathbf{x}) &:= \sum_{i < j} \sqrt{\frac{\lambda}{N}} Y_{ij} x_i x_j - \frac{\lambda}{2N} x_i^2 x_j^2 \\ &= \sum_{i < j} \sqrt{\frac{\lambda}{N}} W_{ij} x_i x_j + \frac{\lambda}{N} x_i x_j x_i^* x_j^* - \frac{\lambda}{2N} x_i^2 x_j^2. \end{aligned} \quad (3)$$

Let us define the free entropy¹ of the model as the expected log-partition function (i.e., normalizing constant) of the posterior $\mathbb{P}_\lambda(\cdot|\mathbf{Y})$:

$$F_N = \frac{1}{N} \mathbb{E} \log \int e^{-H(\mathbf{x})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x}), \quad (4)$$

By heuristically analyzing an approximate message passing (AMP) algorithm for this problem, Lesieur et al. [LKZ15] derived an asymptotic—so-called *replica-symmetric* (RS)—formula for the above quantity. This formula is defined as follows: for $r \in \mathbb{R}_+$, let

$$\psi(r) := \mathbb{E}_{x^*, z} \log \int \exp\left(\sqrt{r}zx + rxx^* - \frac{r}{2}x^2\right) dP_{\mathbf{x}}(x),$$

where $z \sim \mathcal{N}(0, 1)$ and $x^* \sim P_{\mathbf{x}}$ are mutually independent. Define the RS potential

$$F(\lambda, q) := \psi(\lambda q) - \frac{\lambda q^2}{4}.$$

The conjectured limit of (4) is the RS formula

$$\phi_{\text{RS}}(\lambda) := \sup_{q \geq 0} F(\lambda, q).$$

This conjecture was then proved shortly after in a series of papers [KXZ16, BDM⁺16, DAM16, LM16] (see also [KM09]):

Theorem 1. *For all $\lambda \geq 0$,*

$$\lim_{N \rightarrow \infty} F_N = \phi_{\text{RS}}(\lambda).$$

The above statement contains precious statistical information. It can be written in at least two other equivalent ways, in terms of the mutual information between \mathbf{x}^* and \mathbf{Y} :

$$\lim_{N \rightarrow \infty} \frac{1}{N} I(\mathbf{Y}, \mathbf{x}^*) = \frac{\lambda}{4} (\mathbb{E}_{P_{\mathbf{x}}}[X^2])^2 - \phi_{\text{RS}}(\lambda),$$

or, denoting by \mathbb{P}_λ the probability distribution of the matrix \mathbf{Y} as per (1), in terms of the Kullback-Liebler divergence between \mathbb{P}_λ and \mathbb{P}_0 :

$$\lim_{N \rightarrow \infty} \frac{1}{N} D_{\text{KL}}(\mathbb{P}_\lambda, \mathbb{P}_0) = \phi_{\text{RS}}(\lambda).$$

Furthermore, the point $q^*(\lambda)$ achieving the maximum in the RS formula (which can be shown to be unique and finite for almost every λ) can be interpreted as the best overlap any estimator $\hat{\theta}(\mathbf{Y})$ can have with the spike \mathbf{x}^* . Indeed, the overlap of a draw \mathbf{x} from the posterior $\mathbb{P}_\lambda(\cdot|\mathbf{Y})$ with \mathbf{x}^* concentrates about $q^*(\lambda)$. See [BDM⁺16, LM16, EKJ17] for various forms of this statement.

¹The term “free energy” is also used, although the physics convention requires to put a minus sign in front of the expression in this case.

2 Comment on the existing proofs

The proof of the lower bound $\liminf F_N \geq \phi_{\text{RS}}(\lambda)$ relies on an application of Guerra’s interpolation method [Gue01, GT02], and is fairly short and transparent. (See [KXZ16].) Available proofs of the converse bound $\limsup F_N \leq \phi_{\text{RS}}(\lambda)$ (as well as overlap concentration) are on the other hand highly involved. Barbier et al. [BDM⁺16] and Deshpande et al. [DAM16] adopt an algorithmic approach: they analyze an approximate message passing procedure and show that the produced estimator asymptotically achieves an overlap of $q^*(\lambda)$ with the spike. Thus the posterior mean, being the optimal estimator, must also achieve the same overlap. This allows to prove overlap convergence and thus show the converse bound. A difficulty one has to overcome with this method is that AMP (and supposedly any other algorithm) may fail to achieve the optimal overlap in the presence of first-order phase transitions, which traps the algorithm in a bad local optimum of the RS potential. *Spatial coupling*, an idea from coding theory, is used in [BDM⁺16] to overcome this problem. Lelarge and Miolane [LM16] on the other hand use the Aizenman-Sims-Starr scheme [ASS03], a relative of the cavity method developed within spin-glass theory, to prove the upper bound. Barbier and Macris [BM17] prove the upper bound via a *adaptive* version of the interpolation method that proceeds via a sequence of intermediate interpolation steps. Recently, the optimal rate of convergence and constant order corrections to the RS formula were proved in [EKJ17] using a rigorous incarnation of the cavity method due to Talagrand [Tal11]. However, all the current approaches (perhaps to a lesser extent for [BM17]) require the execution of long and technical arguments.

In this note, we show that the upper bound in Theorem 1 admits a fairly simple proof based on the same interpolation idea that yielded the lower bound, combined with an application of the Laplace method and concentration of measure. The main idea is to consider a version of the free entropy (4) of a subsystem of configurations \mathbf{x} having a *fixed* overlap with the spike \mathbf{x}^* . We then proceed by applying the Guerra bound and optimize over this fixed overlap (which is a free parameter) to obtain an upper bound in the form of a saddle (max-min) formula. A small extra effort is needed to show that this last formula is another representation of the RS formula. The idea of restricting the overlap dates back to Franz and Parisi [FP95, FP98] who introduced it in order to study the relaxation properties of dynamics in spin-glass models. The free entropy at fixed overlap bears the name of the *Franz-Parisi potential*. Our proof thus hinges on a upper bound on this potential, which is may be of independent interest. We first start by presenting the proof of the lower bound, which is a starting point for our argument. We present the proof in the case $\sigma = \infty$, i.e., we omit the diagonal terms of \mathbf{Y} . This is only done to keep the displays concise; recovering the general case is straightforward since the diagonal has vanishing contribution to the overall free entropy. Finally, the method presented here can be easily generalized to all spiked tensor models of even order [RM14], thus recovering the main results of [LML⁺17].

3 Proof of Theorem 1

Let $t \in [0, 1]$ and consider an interpolating Hamiltonian

$$\begin{aligned}
 -H_t(\mathbf{x}) := & \sum_{i < j} \sqrt{\frac{t\lambda}{N}} W_{ij} x_i x_j + \frac{t\lambda}{N} x_i x_i^* x_j x_j^* - \frac{t\lambda}{2N} x_i^2 x_j^2 \\
 & + \sum_{i=1}^N \sqrt{(1-t)r} z_i x_i + (1-t)r x_i x_i^* - \frac{(1-t)r}{2} x_i^2,
 \end{aligned} \tag{5}$$

where the z_i 's are i.i.d. standard Gaussian r.v.'s independent of everything else. For $f : (\mathbb{R}^N)^{n+1} \mapsto \mathbb{R}$, we define the Gibbs average of f as

$$\left\langle f(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}, \mathbf{x}^*) \right\rangle_t := \frac{\int f(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}, \mathbf{x}^*) \prod_{l=1}^n e^{-H_t(\mathbf{x}^{(l)})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x}^{(l)})}{\int \prod_{l=1}^n e^{-H_t(\mathbf{x}^{(l)})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x}^{(l)})}. \quad (6)$$

This is the average of f with respect to the posterior distribution of n copies $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}$ of \mathbf{x}^* given the augmented set of observations

$$\begin{cases} Y_{ij} &= \sqrt{\frac{t\lambda}{N}} x_i^* x_j^* + W_{ij}, & 1 \leq i \leq j \leq N, \\ y_i &= \sqrt{(1-t)r} x_i^* + z_i, & 1 \leq i \leq N. \end{cases} \quad (7)$$

The variables $\mathbf{x}^{(l)}, l = 1 \dots, n$ are called *replicas*, and are interpreted as random variables independently drawn from the posterior. When $n = 1$ we simply write $f(\mathbf{x}, \mathbf{x}^*)$ instead of $f(\mathbf{x}^{(1)}, \mathbf{x}^*)$. We shall denote the overlaps between two replicas as follows: for $l, l' = 1, \dots, n, *$, we let

$$R_{l,l'} := \mathbf{x}^{(l)} \cdot \mathbf{x}^{(l')} = \frac{1}{N} \sum_{i=1}^N x_i^{(l)} x_i^{(l')}.$$

A simple consequence of Bayes' rule is that the $n+1$ -tuples $(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n+1)})$ and $(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}, \mathbf{x}^*)$ have the same law under $\mathbb{E}\langle \cdot \rangle_t$ (see Proposition 16 in [LM16]). This bears the name of *the Nishimori property* in the spin glass literature [Nis01].

3.1 The lower bound

Reproducing the argument of [KXZ16], we prove using Guerra's interpolation [Gue01] and the Nishimori property that

$$F_N \geq \phi_{\text{RS}}(\lambda) - \frac{K}{N}.$$

We let $r = \lambda q$ in the definition of H_t and let

$$\varphi(t) := \frac{1}{N} \mathbb{E} \log \int e^{-H_t(\mathbf{x})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x}).$$

A short calculation based on Gaussian integration by parts shows that

$$\begin{aligned} \varphi'(t) &= -\frac{\lambda}{4} \mathbb{E} \langle (R_{1,2} - q)^2 \rangle_t + \frac{\lambda}{4} q^2 + \frac{\lambda}{4N^2} \sum_{i=1}^N \mathbb{E} \langle x_i^{(1)2} x_i^{(2)2} \rangle_t \\ &\quad + \frac{\lambda}{2} \mathbb{E} \langle (R_{1,*} - q)^2 \rangle_t - \frac{\lambda}{2} q^2 - \frac{\lambda}{2N^2} \sum_{i=1}^N \mathbb{E} \langle x_i^2 x_i^{*2} \rangle_t, \end{aligned}$$

By the Nishimori property, the expressions involving the pairs $(\mathbf{x}, \mathbf{x}^*)$ on the one hand and $(\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$ on the other in the brackets are equal. We then obtain

$$\varphi'(t) = \frac{\lambda}{4} \mathbb{E} \langle (R_{1,*} - q)^2 \rangle_t - \frac{\lambda}{4} q^2 - \frac{\lambda}{4N} \mathbb{E} \langle x_N^2 x_N^{*2} \rangle_t.$$

Observe that the last term is $\mathcal{O}(1/N)$ since the variables x_N are bounded. Moreover, the first term is always non-negative so we obtain

$$\varphi'(t) \geq -\frac{\lambda}{4} q^2 - \frac{K}{N}.$$

Since $\varphi(1) = F_N$ and $\varphi(0) = \psi(\lambda q)$, integrating over t , we obtain for all $q \geq 0$, $F_N \geq F(\lambda, q) - \frac{K}{N}$, and this yields the lower bound.

3.2 The upper bound

We prove the converse bound

$$F_N \leq \phi_{\text{RS}}(\lambda) + \mathcal{O}\left(\frac{\log N}{\sqrt{N}}\right).$$

We introduce the Franz-Parisi potential [FP95, FP98]. For $\mathbf{x}^* \in \mathbb{R}^N$ fixed, $m \in \mathbb{R}$ and $\epsilon > 0$ we define

$$\Phi_\epsilon(m, \mathbf{x}^*) := \frac{1}{N} \mathbb{E} \log \int \mathbb{1}\{R_{1,*} \in [m, m + \epsilon]\} e^{-H(\mathbf{x})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x}),$$

where the expectation is over \mathbf{W} . This is the free entropy of a subsystem of configurations having an overlap close to a fixed value m with a planted signal \mathbf{x}^* . It is clear that $\mathbb{E}_{\mathbf{x}^*} \Phi_\epsilon(m, \mathbf{x}^*) \leq F_N$. We will argue via the Laplace method and concentration of measure that $\sup_{m \in \mathbb{R}} \mathbb{E}_{\mathbf{x}^*} \Phi_\epsilon(m, \mathbf{x}^*) \approx F_N$, then use Guerra's interpolation to *upper* bound $\Phi_\epsilon(m, \mathbf{x}^*)$ (notice that this method yielded a *lower* bound on F_N due to the Nishimori property). Let us define a bit of more notation. For $r \in \mathbb{R}_+, s \in \mathbb{R}$, let

$$\hat{\psi}(r, s) := \mathbb{E}_z \log \int \exp\left(\sqrt{r}zx + sx - \frac{r}{2}x^2\right) dP_{\mathbf{x}}(x),$$

where $z \sim \mathcal{N}(0, 1)$, and $\bar{\psi}(r, s) = \mathbb{E}_{x^*} \hat{\psi}(r, sx^*)$ where $x^* \sim P_{\mathbf{x}}$. Moreover, let

$$\hat{F}(\lambda, m, q, \mathbf{x}^*) := \frac{1}{N} \sum_{i=1}^N \hat{\psi}(\lambda q, \lambda m x_i^*) - \frac{\lambda m^2}{2} + \frac{\lambda q^2}{4},$$

and similarly define $\bar{F}(\lambda, m, q) = \mathbb{E}_{\mathbf{x}^*} \hat{F}(\lambda, m, q, \mathbf{x}^*) = \bar{\psi}(\lambda q, \lambda m) - \frac{\lambda m^2}{2} + \frac{\lambda q^2}{4}$.

Proposition 2. *There exist $K > 0$ such that for all $\epsilon > 0$, we have*

$$F_N \leq \mathbb{E}_{\mathbf{x}^*} \left[\max_{l \in \mathbb{Z}, |l| \leq K/\epsilon} \Phi_\epsilon(l\epsilon, \mathbf{x}^*) \right] + \frac{\log(K/\epsilon)}{\sqrt{N}}.$$

Now we upper bound Φ_ϵ in terms of \hat{F} :

Proposition 3 (Interpolation upper bound). *There exist $K > 0$ depending on $\lambda \geq 0$ such that for all $m \in \mathbb{R}$ and $\epsilon > 0$ we have*

$$\Phi_\epsilon(m, \mathbf{x}^*) \leq \inf_{q \geq 0} \hat{F}(\lambda, m, q, \mathbf{x}^*) + \frac{\lambda}{2} \epsilon^2 + \frac{K}{N}.$$

Remark: This simple upper bound on the Franz-Parisi potential—which may be of independent interest—can be straightforwardly generalized to spiked tensor models of even order. Indeed, as will be apparent from the proof in the present matrix case, a crucial step in obtaining the inequality is the positivity of a certain hard-to-control remainder term². Tensor models of even order enjoy a convexity property that ensures the positivity of this remainder.

²We note that the adaptive interpolation method of Barbier and Macris [BM17] is able to bypass this issue of positivity of the remainder term along the interpolation path, as long as this interpolation “stays on the Nishimori line”, i.e., corresponds to an inference problem for every t (this is however not true in the case of the FP potential.) They are thus able to compute the free entropy of (asymmetric) spiked tensor models of odd order. See [BMM17, BKM⁺17].

From Propositions 2 and 3, an upper bound on F_N in the form of a saddle formula begins to emerge. For a fixed $m \in \mathbb{R}$ let $\bar{q} = \bar{q}(\lambda, m)$ be any minimizer of $q \mapsto \bar{F}(\lambda, m, q)$ on \mathbb{R}_+ . (By differentiating \bar{F} , we can check that \bar{q} is bounded uniformly in m .) Then we have

$$F_N \leq \mathbb{E}_{\mathbf{x}^*} \left[\max_{\substack{m=l\epsilon \\ |l| \leq K/\epsilon}} \widehat{F}(\lambda, m, \bar{q}(\lambda, m), \mathbf{x}^*) \right] + \frac{\lambda}{2} \epsilon^2 + \frac{\log(K/\epsilon)}{\sqrt{N}}. \quad (8)$$

At this point we need to push the expectation inside the supremum. This will be done using a concentration argument.

Lemma 4. *There exists $K > 0$ such that for all $\lambda \geq 0$, $m \in \mathbb{R}$, $q \geq 0$ and $t \geq 0$,*

$$\Pr_{\mathbf{x}^*} \left(\left| \widehat{F}(\lambda, m, q, \mathbf{x}^*) - \bar{F}(\lambda, m, q) \right| \geq t \right) \leq 2e^{-\frac{Nt^2}{\lambda^2 K(|m|+q)^2}}.$$

It is a routine computation to deduce from Lemma 4 (and boundedness of both m and q) that the expected supremum is bounded by the supremum of the expectation plus a small term (a similar argument is given in the proof of Proposition 2):

$$\mathbb{E} \sup_{\substack{m=l\epsilon \\ |l| \leq K/\epsilon}} \widehat{F}(\lambda, m, \bar{q}(\lambda, m), \mathbf{x}^*) \leq \sup_{\substack{m=l\epsilon \\ |l| \leq K/\epsilon}} \bar{F}(\lambda, m, \bar{q}(\lambda, m)) + \delta,$$

where $\delta = K \log(K/\epsilon)/\sqrt{N}$. Since \bar{q} is a minimizer of \bar{F} , it follows from (8) that

$$F_N \leq \sup_{m \in \mathbb{R}} \inf_{q \geq 0} \bar{F}(\lambda, m, q) + \frac{\lambda}{2} \epsilon^2 + \frac{K \log(K/\epsilon)}{\sqrt{N}}. \quad (9)$$

We now let $\epsilon = N^{-1/4}$, and conclude by noticing that the above saddle formula is another expression for ϕ_{RS} :

Proposition 5. $\phi_{\text{RS}}(\lambda) = \sup_{m \in \mathbb{R}} \inf_{q \geq 0} \bar{F}(\lambda, m, q)$.

Proof. One inequality follows from (9) and the lower bound $F_N \geq \phi_{\text{RS}}(\lambda) - o_N(1)$. For the converse inequality, we notice that for all $m \in \mathbb{R}$

$$\inf_{q \geq 0} \bar{F}(\lambda, m, q) \leq \bar{F}(\lambda, m, |m|) = \bar{\psi}(\lambda|m|, \lambda m) - \frac{\lambda}{4}|m|^2.$$

Now we use the fact that the function $\bar{\psi}$ is largest when its second argument is positive:

Lemma 6. *For all $r \geq 0$ we have $\bar{\psi}(r, -r) \leq \bar{\psi}(r, r)$.*

This implies $\inf_{q \geq 0} \bar{F}(\lambda, m, q) \leq \bar{F}(\lambda, |m|) - \frac{\lambda}{4}|m|^2$. Taking the supremum over m yields the converse bound. \blacksquare

Proof of Proposition 2. Let $\epsilon > 0$. Since the prior $P_{\mathbf{x}}$ has bounded support, we can grid the set of the overlap values $R_{1,*}$ by $2K/\epsilon$ many intervals of size ϵ for some $K > 0$. This allows the following discretization, where l runs over the finite range $\{-K/\epsilon, \dots, K/\epsilon\}$:

$$\begin{aligned} F_N &= \frac{1}{N} \mathbb{E} \log \sum_l \int \mathbf{1}\{R_{1,*} \in [l\epsilon, (l+1)\epsilon)\} e^{-H(\mathbf{x})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x}) \\ &\leq \frac{1}{N} \mathbb{E} \log \frac{2K}{\epsilon} \max_l \int \mathbf{1}\{R_{1,*} \in [l\epsilon, (l+1)\epsilon)\} e^{-H(\mathbf{x})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x}) \\ &= \frac{1}{N} \mathbb{E} \max_l \log \int \mathbf{1}\{R_{1,*} \in [l\epsilon, (l+1)\epsilon)\} e^{-H(\mathbf{x})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x}) + \frac{\log(2K/\epsilon)}{N}. \end{aligned} \quad (10)$$

In the above, \mathbb{E} is w.r.t. both \mathbf{W} and \mathbf{x}^* . We use concentration of measure to push the expectation over \mathbf{W} to the left of the maximum. Let

$$Z_l := \int \mathbf{1}\{R_{1,*} \in [l\epsilon, (l+1)\epsilon)\} e^{-H(\mathbf{x})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x}).$$

We show that each term $X_l = \frac{1}{N} \log Z_l$ concentrates about its expectation (in the randomness of \mathbf{W}). Let \mathbb{E}' denote the expectation w.r.t. \mathbf{W} .

Lemma 7. *There exists a constant $K > 0$ such that for all $\gamma \geq 0$ and all l ,*

$$\mathbb{E}' e^{\gamma(X_l - \mathbb{E}'[X_l])} \leq \frac{K\gamma}{\sqrt{N}} e^{K\gamma^2/N}.$$

Therefore, the expectation of the maximum concentrates as well:

$$\begin{aligned} \mathbb{E}' \max_l (X_l - \mathbb{E}'[X_l]) &\leq \frac{1}{\gamma} \log \mathbb{E}' \exp \left(\gamma \max_l (X_l - \mathbb{E}'[X_l]) \right) \\ &= \frac{1}{\gamma} \log \mathbb{E}' \max_l e^{\gamma(X_l - \mathbb{E}'[X_l])} \\ &\leq \frac{1}{\gamma} \log \mathbb{E}' \sum_l e^{\gamma(X_l - \mathbb{E}'[X_l])} \\ &\leq \frac{1}{\gamma} \log \left(\frac{2K}{\epsilon} \frac{\gamma K}{\sqrt{N}} e^{\gamma^2 K/N} \right) \\ &= \frac{\log(2K/\epsilon)}{\gamma} + \frac{1}{\gamma} \log \frac{\gamma K}{\sqrt{N}} + \frac{\gamma K}{N}. \end{aligned}$$

We set $\gamma = \sqrt{N}$ and obtain

$$\mathbb{E}' \max_l (X_l - \mathbb{E}'[X_l]) \leq \frac{\log(K/\epsilon)}{\sqrt{N}}.$$

Therefore, plugging the above estimates into (10), we obtain

$$\begin{aligned} F_N &\leq \mathbb{E}_{\mathbf{x}^*} \max_l \mathbb{E}' X_l + \frac{\log(K/\epsilon)}{\sqrt{N}} + \frac{\log(K/\epsilon)}{N} \\ &\leq \mathbb{E}_{\mathbf{x}^*} \max_l \Phi_\epsilon(l\epsilon, \mathbf{x}^*) + 2 \frac{\log(K/\epsilon)}{\sqrt{N}}. \end{aligned}$$

■

Proof of Proposition 3. Let $t \in [0, 1]$ and consider a slightly modified interpolating Hamiltonian that has two parameters $r = \lambda q \geq 0$ and $s = \lambda m \in \mathbb{R}$:

$$\begin{aligned} -H_t(\mathbf{x}) &:= \sum_{i < j} \sqrt{\frac{t\lambda}{N}} W_{ij} x_i x_j + \frac{t\lambda}{N} x_i x_i^* x_j x_j^* - \frac{t\lambda}{2N} x_i^2 x_j^2 \\ &\quad + \sum_{i=1}^N \sqrt{(1-t)r} z_i x_i + (1-t)s x_i x_i^* - \frac{(1-t)r}{2} x_i^2, \end{aligned} \tag{11}$$

where the z_i 's are i.i.d. standard Gaussian r.v.'s independent of everything else. Let

$$\varphi(t) := \frac{1}{N} \mathbb{E} \log \int \mathbf{1}\{R_{1,*} \in [m, m + \epsilon)\} e^{-H_t(\mathbf{x})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x}),$$

where \mathbb{E} is over the Gaussian disorder \mathbf{W} and \mathbf{z} (\mathbf{x}^* is fixed). Let $\langle \cdot \rangle_t$ be the corresponding Gibbs average, similarly to (6). By differentiation and Gaussian integration by parts,

$$\begin{aligned}\varphi'(t) &= -\frac{\lambda}{4} \mathbb{E} \langle (R_{1,2} - q)^2 \rangle_t + \frac{\lambda}{4} q^2 + \frac{\lambda}{4N^2} \sum_{i=1}^N \mathbb{E} \langle (x_i^{(1)} x_i^{(2)})^2 \rangle_t \\ &\quad + \frac{\lambda}{2} \mathbb{E} \langle (R_{1,*} - m)^2 \rangle_t - \frac{\lambda}{2} m^2 - \frac{\lambda}{2N^2} \sum_{i=1}^N \mathbb{E} \langle (x_i x_i^*)^2 \rangle_t,\end{aligned}$$

Notice that by the overlap restriction, $\mathbb{E} \langle (R_{1,*} - m)^2 \rangle_t \leq \epsilon^2$. Moreover, the last terms in the first and second lines in the above are of order $1/N$ since the variables x_i are bounded. Next, since $\mathbb{E} \langle (R_{1,2} - q)^2 \rangle_t$ has non-negative sign (this is a crucial fact), we can ignore it and obtain an upper bound:

$$\varphi'(t) \leq -\frac{\lambda}{2} m^2 + \frac{\lambda}{4} q^2 + \frac{\lambda}{2} \epsilon^2 + \frac{K}{N}.$$

Integrating over t , we obtain

$$\Phi_\epsilon(m, \mathbf{x}^*) \leq -\frac{\lambda}{2} m^2 + \frac{\lambda}{4} q^2 + \frac{\lambda}{2} \epsilon^2 + \varphi(0) + \frac{K}{N}.$$

Now we use a trivial upper bound on $\varphi(0)$:

$$\begin{aligned}\varphi(0) &= \frac{1}{N} \mathbb{E} \log \int \mathbb{1}\{R_{1,*} \in [m, m + \epsilon]\} e^{-H_0(\mathbf{x})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x}) \\ &\leq \frac{1}{N} \mathbb{E} \log \int e^{-H_0(\mathbf{x})} dP_{\mathbf{x}}^{\otimes N}(\mathbf{x}) \\ &= \frac{1}{N} \sum_{i=1}^N \widehat{\psi}(\lambda q, \lambda m x_i^*).\end{aligned}$$

Hence,

$$\Phi_\epsilon(m, \mathbf{x}^*) \leq \widehat{F}(\lambda, m, q, \mathbf{x}^*) + \frac{\lambda}{2} \epsilon^2 + \frac{K}{N}. \quad \blacksquare$$

Proof of lemma 4. The random part of $\widehat{F}(\lambda, m, \bar{q}, \mathbf{x}^*)$ is the average of i.i.d. terms $\widehat{\psi}(\lambda q, \lambda m x_i^*)$. Since $|\partial_s \widehat{\psi}(r, s x^*)| \leq K^2$, $|\partial_r \widehat{\psi}(r, s x^*)| \leq K^2/2$ and $\widehat{\psi}(0, 0) = 0$, where K is a bound on the support of $P_{\mathbf{x}}$, we have $|\widehat{\psi}(r, s x^*)| \leq K^2(r/2 + |s|)$. For bounded r and s , the claim follows from concentration of the average of i.i.d. bounded r.v.'s. \blacksquare

Proof of Lemma 7. We notice that X_l seen as a function of \mathbf{W} is Lipschitz with constant $K\sqrt{\frac{\lambda}{N}}$. By Gaussian concentration of Lipschitz functions (the Tsirelson-Ibragimov-Sudakov inequality [BLM13]), there exist a constant K depending only on λ such that for all $t \geq 0$,

$$\Pr(X_l - \mathbb{E}' X_l \geq t) \leq e^{-Nt^2/K}.$$

Then we conclude by means of the identity

$$\mathbb{E}' e^{\gamma(X_l - \mathbb{E}'[X_l])} = \gamma \int_{-\infty}^{+\infty} \Pr(X_l - \mathbb{E}'[X_l] \geq t) e^{\gamma t} dt,$$

and integrate the tail. ■

Proof of Lemma 6. Let $\nu = P_x$, and let μ be the symmetric part of P_x , i.e., $\mu(A) = (P_x(A) + P_x(-A))/2$ for all Borel $A \subseteq \mathbb{R}$. Observe that ν is absolutely continuous with respect to μ . The argument relies on a linearly interpolating between the two measures μ and ν . Let $t \in [0, 1]$ and let $\rho_t = (1-t)\mu + t\nu$. Further, let $r > 0$ be fixed, and

$$\varphi^\pm(t) := \mathbb{E}_z \int \left(\log \int \exp \left(\sqrt{r}zx \pm rxx^* - \frac{r}{2}x^2 \right) d\rho_t(x) \right) d\rho_t(x^*),$$

where $z \sim \mathcal{N}(0, 1)$. Now let $\phi(t) = \varphi^+(t) - \varphi^-(t)$. We have $\phi(1) = \bar{\psi}(r, r) - \bar{\psi}(r, -r)$ on the one hand, and since μ is a symmetric distribution, $\phi(0) = 0$ on the other. We will show that ϕ is a convex increasing function on the interval $[0, 1]$. Then we deduce that $\phi(1) \geq 0$. First, we have

$$\begin{aligned} \frac{d}{dt}\varphi^+(t) &= \mathbb{E}_z \int \log \int e^{\sqrt{r}zx + rxx^* - \frac{r}{2}x^2} d\rho_t(x) d(\nu - \mu)(x^*) \\ &\quad + \mathbb{E}_z \int \frac{\int e^{\sqrt{r}zx + rxx^* - \frac{r}{2}x^2} d(\nu - \mu)(x)}{\int e^{\sqrt{r}zx + rxx^* - \frac{r}{2}x^2} d\rho_t(x)} d\rho_t(x^*), \end{aligned}$$

and

$$\begin{aligned} \frac{d^2}{dt^2}\varphi^+(t) &= 2\mathbb{E}_z \int \frac{\int e^{\sqrt{r}zx + rxx^* - \frac{r}{2}x^2} d(\nu - \mu)(x)}{\int e^{\sqrt{r}zx + rxx^* - \frac{r}{2}x^2} d\rho_t(x)} d(\nu - \mu)(x^*) \\ &\quad - 2\mathbb{E}_z \int \left(\frac{\int e^{\sqrt{r}zx + rxx^* - \frac{r}{2}x^2} d(\nu - \mu)(x)}{\int e^{\sqrt{r}zx + rxx^* - \frac{r}{2}x^2} d\rho_t(x)} \right)^2 d\rho_t(x^*). \end{aligned}$$

Similar expressions holds for φ^- where x^* is replaced by $-x^*$ inside the exponentials. We see from the expression of the first derivative at $t = 0$ that $\varphi^{+'}(0) = \varphi^{-'}(0)$. This is because $\rho_0 = \mu$ is symmetric about the origin, so a sign change (of x for the first term, and x^* for the second term) does not affect the value of the integrals. Hence $\phi'(0) = 0$. Now, we focus on the second derivative. Observe that since μ is the symmetric part of ν , $\nu - \mu$ is anti-symmetric. This implies that the first term in the expression of the second derivative changes sign under a sign change in x^* and keeps the same modulus. As for the second term, a sign change in x^* induces integration against $d\rho_t(-x^*)$. Hence we can write the difference $(\varphi^+ - \varphi^-)''$ as

$$\begin{aligned} \phi''(t) &= 4\mathbb{E}_z \int \frac{\int e^{\sqrt{r}zx + rxx^* - \frac{r}{2}x^2} d(\nu - \mu)(x)}{\int e^{\sqrt{r}zx + rxx^* - \frac{r}{2}x^2} d\rho_t(x)} d(\nu - \mu)(x^*) \\ &\quad - 2\mathbb{E}_z \int \left(\frac{\int e^{\sqrt{r}zx + rxx^* - \frac{r}{2}x^2} d(\nu - \mu)(x)}{\int e^{\sqrt{r}zx + rxx^* - \frac{r}{2}x^2} d\rho_t(x)} \right)^2 (d\rho_t(x^*) - d\rho_t(-x^*)). \end{aligned}$$

For any Borel A , we have $\rho_t(A) - \rho_t(-A) = (1-t)(\mu(A) - \mu(-A)) + t(\nu(A) - \nu(-A)) = 2t(\nu - \mu)(A)$. Therefore the second term in the above expression becomes

$$-4t\mathbb{E}_z \int \left(\frac{\int e^{\sqrt{r}zx + rxx^* - \frac{r}{2}x^2} d(\nu - \mu)(x)}{\int e^{\sqrt{r}zx + rxx^* - \frac{r}{2}x^2} d\rho_t(x)} \right)^2 d(\nu - \mu)(x^*).$$

Since both μ and ν are absolutely continuous with respect to ρ_t for all $0 \leq t < 1$ we can write

$$\phi''(t) = 4 \mathbb{E}_{z, x^*} \left\langle \frac{d(\nu - \mu)}{d\rho_t}(x) \frac{d(\nu - \mu)}{d\rho_t}(x^*) \right\rangle - 4t \mathbb{E}_{z, x^*} \left\langle \frac{d(\nu - \mu)}{d\rho_t}(x) \right\rangle^2,$$

where the Gibbs average is with respect to the posterior of x given z, x^* under the Gaussian channel $y = \sqrt{r}x^* + z$, and the expectation is under $x^* \sim \rho_t$ and $z \sim \mathcal{N}(0, 1)$. By the Nishimori property, we simplify the above expression to

$$\phi''(t) = 4(1 - t) \mathbb{E} \left[\left\langle \frac{d(\nu - \mu)}{d\rho_t}(x) \right\rangle^2 \right] \geq 0,$$

where the expression is valid for all $0 \leq t < 1$. From here we see that the function ϕ is convex on $[0, 1]$. Since $\phi(0) = \phi'(0) = 0$, ϕ is also increasing on $[0, 1]$. ■

Acknowledgments. Florent Krzakala acknowledges funding from the ERC under the European Union 7th Framework Programme Grant Agreement 307087-SPARCS.

References

- [ASS03] Michael Aizenman, Robert Sims, and Shannon L Starr. Extended variational principle for the Sherrington-Kirkpatrick spin-glass model. *Physical Review B*, 68(21):214403, 2003.
- [BAP05] Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Annals of Probability*, 33(5):1643–1697, 2005.
- [BDM⁺16] Jean Barbier, Mohamad Dia, Nicolas Macris, Florent Krzakala, Thibault Lesieur, and Lenka Zdeborová. Mutual information for symmetric rank-one matrix estimation: A proof of the replica formula. In *Advances in Neural Information Processing Systems (NIPS)*, pages 424–432, 2016.
- [BGN11] Florent Benaych-Georges and Raj Rao Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, 2011.
- [BKM⁺17] Jean Barbier, Florent Krzakala, Nicolas Macris, Léo Miolane, and Lenka Zdeborová. Phase transitions, optimal errors and optimality of message-passing in generalized linear models. *arXiv preprint arXiv:1708.03395*, 2017.
- [BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- [BM17] Jean Barbier and Nicolas Macris. The stochastic interpolation method: A simple scheme to prove replica formulas in bayesian inference. *arXiv preprint arXiv:1705.02780*, 2017.
- [BMM17] Jean Barbier, Nicolas Macris, and Léo Miolane. The layered structure of tensor estimation and its mutual information. *arXiv preprint arXiv:1709.10368*, 2017.
- [DAM16] Yash Deshpande, Emmanuel Abbé, and Andrea Montanari. Asymptotic mutual information for the binary stochastic block model. In *IEEE International Symposium on Information Theory (ISIT)*, pages 185–189. IEEE, 2016.
- [EKJ17] Ahmed El Alaoui, Florent Krzakala, and Michael I Jordan. Finite size corrections and likelihood ratio fluctuations in the spiked Wigner model. *arXiv preprint arXiv:1710.02903*, 2017.

- [FP95] Silvio Franz and Giorgio Parisi. Recipes for metastable states in spin glasses. *Journal de Physique I*, 5(11):1401–1415, 1995.
- [FP98] Silvio Franz and Giorgio Parisi. Effective potential in glassy systems: theory and simulations. *Physica A: Statistical Mechanics and its Applications*, 261(3-4):317–339, 1998.
- [GT02] Francesco Guerra and Fabio Lucio Toninelli. The thermodynamic limit in mean field spin glass models. *Communications in Mathematical Physics*, 230(1):71–79, 2002.
- [Gue01] Francesco Guerra. Sum rules for the free energy in the mean field spin glass model. *Fields Institute Communications*, 30:161–170, 2001.
- [Joh01] Iain M Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *Annals of Statistics*, pages 295–327, 2001.
- [KM09] Satish Babu Korada and Nicolas Macris. Exact solution of the gauge symmetric p-spin glass model on a complete graph. *Journal of Statistical Physics*, 136(2):205–230, 2009.
- [KXZ16] Florent Krzakala, Jiaming Xu, and Lenka Zdeborová. Mutual information in rank-one matrix estimation. In *Information Theory Workshop (ITW)*, pages 71–75. IEEE, 2016.
- [LKZ15] Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. Phase transitions in sparse PCA. In *IEEE International Symposium on Information Theory (ISIT)*, pages 1635–1639. IEEE, 2015.
- [LM16] Marc Lelarge and Léo Miolane. Fundamental limits of symmetric low-rank matrix estimation. *arXiv preprint arXiv:1611.03888*, 2016.
- [LML⁺17] Thibault Lesieur, Léo Miolane, Marc Lelarge, Florent Krzakala, and Lenka Zdeborová. Statistical and computational phase transitions in spiked tensor estimation. In *IEEE International Symposium on Information Theory (ISIT)*, pages 511–515, June 2017.
- [Nis01] Hidetoshi Nishimori. *Statistical physics of spin glasses and information processing: an introduction*, volume 111. Clarendon Press, 2001.
- [Péc06] Sandrine Péché. The largest eigenvalue of small rank perturbations of Hermitian random matrices. *Probability Theory and Related Fields*, 134(1):127–173, 2006.
- [Péc14] Sandrine Péché. Deformed ensembles of random matrices. In *Proceedings of the International Congress of Mathematicians, Seoul*, pages 1059–1174. ICM, 2014.
- [RM14] Emile Richard and Andrea Montanari. A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2897–2905, 2014.
- [Tal11] Michel Talagrand. *Mean field models for spin glasses. Volume I: Basic examples*, volume 54. Springer Science & Business Media, 2011.