

Landscape Complexity for the Empirical Risk of Generalized Linear Models

Antoine Maillard^{*,⊗}, Gérard Ben Arous[†], Giulio Biroli^{*}

Abstract

We present a method to obtain the average and the typical value of the number of critical points of the empirical risk landscape for generalized linear estimation problems and variants. This represents a substantial extension of previous applications of the Kac-Rice method since it allows to analyze the critical points of high dimensional non-Gaussian random functions. We obtain a rigorous explicit variational formula for the *annealed complexity*, which is the logarithm of the average number of critical points at fixed value of the empirical risk. This result is simplified, and extended, using the non-rigorous Kac-Rice replicated method from theoretical physics. In this way we find an explicit variational formula for the *quenched complexity*, which is generally different from its annealed counterpart, and allows to obtain the number of critical points for typical instances up to exponential accuracy.

1 Introduction and main results

1.1 Introduction

Characterizing the landscape of the empirical risk is a key issue in several contexts. Many current machine learning problems are both non-convex and high-dimensional. In these cases, the analysis of optimization algorithms, such as gradient descent and its stochastic variants, represents a very hard challenge. In recent years, there has been a series of works that developed a landscape-based approach to tackle this challenge. The key idea is to study the statistical properties of the empirical risk landscape, and to use these findings to obtain results on the performance of algorithms. Without the aim of being exhaustive this research avenue includes analysis of the landscape of neural networks, matrix completion, tensor factorization and tensor principal component analysis [Fyodorov, 2004, Fyodorov and Nadal, 2012, Kawaguchi, 2016, Soudry and Carmon, 2016, Ge et al., 2016, Freeman and Bruna, 2016, Bhojanapalli et al., 2016, Park et al., 2017, Du et al., 2018, Ge and Ma, 2017, Ge et al., 2017, Lu and Kawaguchi, 2017, Ling et al., 2019, Ben Arous et al., 2019, Ros et al., 2019, Mannelli et al., 2019, Sarao Mannelli et al., 2019, Ben Arous et al., 2018, Biroli et al., 2019]. The majority of these works identifies the region of parameters where the landscape is “easy”, i.e. it focuses on the regime where there shouldn’t be any bad local minima and it proves that indeed there are none. However, gradient descent and other landscape-based algorithms are often observed to work even very far from the region described above where the landscape can be proved mathematically to be “easy”. A possible reason is that the bounds obtained rigorously are not tight enough. Another, more interesting, is that the landscape is “hard”, i.e. spurious minima are present, but their basins of attraction are small and the dynamics is able to avoid them [Mannelli et al., 2019, Sarao Mannelli et al., 2019].

Here we develop a general method that allows to study and focus directly on the “hard” regime, where the empirical risk displays a huge number of bad minima. Our aim is to obtain explicit formulas for the

^{*} Laboratoire de Physique de l’Ecole Normale Supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université Paris-Diderot, Sorbonne Paris Cité, Paris, France.

[†] Courant Institute of Mathematical Sciences, New York University, 251 Mercer Street, New York, NY 10012.

[⊗] To whom correspondence shall be sent: antoine.maillard@ens.fr

number of critical points of the empirical risk landscape, to characterize their indices, and the Hessian associated to them. For a given problem, this will allow to identify the topological transition where the landscape becomes “easy”, and to analyze very precisely the “hard” regime. In recent years, there has been remarkable progress on this subject in the field of spin-glasses and probability theory through the Kac-Rice method [Fyodorov and Williams, 2007, Bray and Dean, 2007, Auffinger and Ben Arous, 2013, Auffinger et al., 2013, Subag, 2017, Ben Arous et al., 2019, Ros et al., 2019]. This line of research has allowed to put on a firm ground results previously obtained in the physics literature [Bray and Moore, 1980, Kurchan, 1991, Crisanti and Sommers, 1995, Cavagna et al., 1999], and it has unveiled important relationships with random matrix theory. Its main domain of application has been the study of the landscapes associated to Gaussian random functions. Its extension to tackle the case of non-Gaussian high-dimensional random functions is an open problem—one that is crucial to address in order to characterize the critical points of the empirical risk.

Here we present an important step forward: an extension of the Kac-Rice method to compute the number of critical points of the empirical risk arising in generalized linear estimation problems [Nelder and Wedderburn, 1972, McCullagh, 1984, Barbier et al., 2019]. Our approach contains both mathematically rigorous analysis and exact results obtained by theoretical physics methods. We work out a rigorous formula for the logarithm of the average of the number of the critical points, called henceforth *annealed complexity*. This quantity already provides interesting information in itself. A more refined quantity, but much more challenging to be analyzed rigorously, is the average of the logarithm of the number of critical points. This so-called *quenched complexity* is truly representative of the typical properties of the landscape for a given instance of the empirical risk and is generically different from its annealed counterpart which is instead dominated by rare instances, see for example the case of tensor PCA [Ros et al., 2019]. In order to obtain the quenched complexity, we develop a non-rigorous but exact approach that combines the Kac-Rice method with the replica theory used by physicists [Ros et al., 2019].

Our main results are two explicit formulas for the quenched and annealed complexities. They open the way towards a full fledged characterization of the landscape of generalized estimation problems and variants, and the analysis of landscape-based algorithms, such as gradient descent. Note that these models can also be viewed as the simplest neural network (single-node) in a teacher-student setting [Engel and Van den Broeck, 2001].

1.2 Main results

We consider two classes of high-dimensional random functions. The first one is a kind of energy that arises in a simple model of neural networks (the perceptron, cf. Engel and Van den Broeck [2001]) and in mean-field glass models [Franz and Parisi, 2016]:

$$L_1(\mathbf{x}) \equiv \frac{1}{m} \sum_{\mu=1}^m \phi(\boldsymbol{\xi}_\mu \cdot \mathbf{x}), \quad (1)$$

where $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is a smooth activation function (the hypotheses on ϕ are precised later), $\mathbf{x} \in \mathbb{S}^{n-1}$, the unit sphere in n dimensions, and $\boldsymbol{\xi}_\mu$ are i.i.d. random variables generated from the standard Gaussian distribution in \mathbb{R}^n . The second class of functions we will consider are related to the loss functions of generalized linear models (GLMs) Nelder and Wedderburn [1972], McCullagh [1984]. In generalized linear estimation an observer has to infer a hidden vector $\mathbf{x}^* \in \mathbb{S}^{n-1}$ from the observation of the m -dimensional output vector $\mathbf{Y} = \{\phi(\boldsymbol{\xi}_\mu \cdot \mathbf{x}^*)\}_{\mu=1}^m$. In this sense, the GLMs generalize the usual linear regression by allowing the output function to be non-linear¹. We consider here *random* GLMs, meaning that the data (or measurement) matrix $\boldsymbol{\xi}$ is taken random, with an i.i.d. standard Gaussian distribution, and we assume that the function ϕ and the data matrix $\boldsymbol{\xi}$ are given to the observer. This naturally leads to the

¹One can also consider GLMs in which the output function is stochastic. Here, we restrict to deterministic outputs.

mean square loss L_2 :

$$L_2(\mathbf{x}) \equiv \frac{1}{2m} \sum_{\mu=1}^m [\phi(\boldsymbol{\xi}_\mu \cdot \mathbf{x}^*) - \phi(\boldsymbol{\xi}_\mu \cdot \mathbf{x})]^2, \quad (2)$$

GLMs (and their random versions) arise in many different areas of statistics, such as e.g. compressed sensing, phase retrieval, logistic regression, or in random artificial neural networks; we refer to [Barbier et al. \[2019\]](#) for a review of its numerous applications.

Here, we are interested in the statistics of the number of critical points, or more precisely the complexity of the associated empirical risk (2). For any open intervals $B \subseteq \mathbb{R}_+$ and $Q \subseteq (-1, 1)$, we consider the (random) number $\text{Crit}_{n,L_2}(B, Q)$ of critical points of the function L_2 with loss value in B and overlap with the signal $q \equiv \mathbf{x} \cdot \mathbf{x}^*$ in Q :

$$\text{Crit}_{n,L_2}(B, Q) \equiv \sum_{\mathbf{x}: \text{grad}L_2(\mathbf{x})=0} \mathbb{1}\{L_2(\mathbf{x}) \in B, \mathbf{x} \cdot \mathbf{x}^* \in Q\}. \quad (3)$$

Here grad is the Riemannian gradient on \mathbb{S}^{n-1} . For L_1 we define the similar quantity $\text{Crit}_{n,L_1}(B)$, dropping the notion of overlap. The average $\mathbb{E} \text{Crit}_{n,L_2}(B, Q)$ is the quantity that can be analyzed rigorously. Its logarithm divided by n is called the *annealed complexity*. However, since the random variable $\text{Crit}_{n,L_2}(B, Q)$ is in general strongly fluctuating and scales exponentially with n , its typical value is different from the mean and can be obtained by taking the exponential of $\mathbb{E} \ln \text{Crit}_{n,L_2}(B, Q)$. This last quantity (divided by n) is called the *quenched complexity*. It is in general different from the annealed one, with very few exceptions [[Subag, 2017](#), [Crisanti and Sommers, 1995](#)].

Our main results consist in explicit formulas for the annealed and quenched complexities for L_1 and L_2 . The formula for the annealed case is obtained by a rigorous Kac-Rice method, whereas the one for the quenched complexity is obtained by theoretical physics methods combining the Kac-Rice method with replica theory. We consider the limit $n, m \rightarrow \infty$ with $m/n \rightarrow \alpha > 1$, a setting called in the statistical physics literature the thermodynamic limit. The condition $\alpha > 1$ is essential, as can be seen for instance in eq. (1): if $m < n$, for each realization of $\{\boldsymbol{\xi}_\mu\}$, the function L_1 has an infinite number of critical points in the set of unit-norm \mathbf{x} orthogonal to all the $\{\boldsymbol{\xi}_\mu\}$, and counting the critical points in this case is meaningless (or one has to quotient the space to take care of the degeneracy). Our results hold for many classical activation functions ϕ , such as e.g. the hyperbolic tangent, the arctangent, the sigmoid, or a smoothed and leaky version of the ReLU activation function¹. Henceforth we shall denote \mathbb{C}_+ the strict upper-half complex plane, and $\mathcal{M}(\mathbb{R}^k)$ the set of probability measures on \mathbb{R}^k . For two probability measures μ and ν we define the relative entropy $H(\mu|\nu) \equiv \int \ln(d\mu/d\nu)d\mu$ if μ is absolutely continuous with respect to ν , and $+\infty$ otherwise. Finally, μ_G is a generic notation for the standard Gaussian measure on any \mathbb{R}^k .

We can now present our main results.

Theorem 1 (The annealed complexity of L_1). *Let $B \subseteq \mathbb{R}$ a non-empty open interval and denote $\mathcal{M}_\phi(B)$ the set of probability measures ν on \mathbb{R} such that $\int \nu(dt)\phi(t) \in B$. Given:*

- $\mathcal{E}_\phi(\nu) \equiv \ln \left[\int \nu(dx) \phi'(x)^2 \right],$
- $t_\phi(\nu) \equiv \int \nu(dx) x \phi'(x),$
- *Let $\mathbf{z} \in \mathbb{R}^{n \times m}$ an i.i.d. standard Gaussian matrix, and $\mathbf{y} \in \mathbb{R}^m$ a vector with components taken i.i.d. from a probability measure ν . Let $D^{(\nu)}$ the diagonal matrix of size m with elements $D_\mu^{(\nu)} = \phi''(y_\mu)$. We define $\mu_{\alpha,\phi}[\nu]$ as the asymptotic spectral measure of $\mathbf{z}D^{(\nu)}\mathbf{z}^\top/m$.*
- $\kappa_{\alpha,\phi}(\nu, C) \equiv \int \mu_{\alpha,\phi}[\nu](dx) \ln |x - C|,$

¹The precise hypotheses on the activation function ϕ are precised in Section 2.

one has¹:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln \mathbb{E} \text{Crit}_{n, L_1}(B) = \frac{1 + \ln \alpha}{2} + \sup_{\nu \in \mathcal{M}_\phi(B)} \left[-\frac{1}{2} \mathcal{E}_\phi(\nu) + \kappa_{\alpha, \phi}(\nu, t_\phi(\nu)) - \alpha H(\nu | \mu_G) \right].$$

A note on free probability Interestingly, the measure $\mu_{\alpha, \phi}[\nu]$ can be interpreted as the free multiplicative convolution of the Marchenko-Pastur law (at ratio α) and the asymptotic spectral distribution of the matrix $D^{(\nu)}$, cf. e.g. Voiculescu [1987], Anderson et al. [2010]². We describe in Section 3 how to explicitly compute the density of $\mu_{\alpha, \phi}[\nu]$, or its linear spectral statistics (as e.g. $\kappa_{\alpha, \phi}(\nu, C)$), via the computation of its Stieljtes transform.

We turn to our second annealed result:

Theorem 2 (The annealed complexity of L_2). *Let $B \subseteq \mathbb{R}_+$ and $Q \subseteq (-1, 1)$ two non-empty open intervals. For $q \in (-1, 1)$ we denote $\mathcal{M}_\phi(B, q)$ the set of probability measures ν on \mathbb{R}^2 such that*

$$\begin{cases} \int \nu(dx, dy) y \phi'(x) \left[\phi(qx + \sqrt{1 - q^2}y) - \phi(x) \right] = 0, \\ \int \nu(dx, dy) \left[\phi(qx + \sqrt{1 - q^2}y) - \phi(x) \right]^2 \in B. \end{cases} \quad (4)$$

Given:

- $\mathcal{E}_\phi(q, \nu) \equiv \ln \left[\int \nu(dx, dy) \phi'(x)^2 \left[\phi(qx + \sqrt{1 - q^2}y) - \phi(x) \right]^2 \right]$,
- $t_\phi(q, \nu) \equiv \int \nu(dx, dy) x \phi'(x) \left[\phi(x) - \phi(qx + \sqrt{1 - q^2}y) \right]$,
- f_q is a function from \mathbb{R}^2 to \mathbb{R} defined by:

$$f_q(x, y) \equiv \phi'(x)^2 - \phi''(x) \left[\phi(qx + \sqrt{1 - q^2}y) - \phi(x) \right], \quad (5)$$

- Let $\mathbf{z} \in \mathbb{R}^{n \times m}$ an i.i.d. standard Gaussian matrix, and $\mathbf{Y} \in \mathbb{R}^{m \times 2}$ with components taken i.i.d. from a probability measure ν on \mathbb{R}^2 . Let $D^{(\nu, q)}$ the diagonal matrix of size m with elements $D_\mu^{(\nu, q)} = f_q(Y_\mu)$. We define $\mu_{\alpha, \phi}[q, \nu]$ as the asymptotic spectral measure of $\mathbf{z} D^{(\nu, q)} \mathbf{z}^\top / m$.
- $\kappa_{\alpha, \phi}(q, \nu) \equiv \int \mu_{\alpha, \phi}[q, \nu](dx) \ln |x - t_\phi(q, \nu)|$,

then one has:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln \mathbb{E} \text{Crit}_{n, L_2}(B, Q) = \frac{1 + \ln \alpha}{2} + \sup_{q \in Q} \sup_{\nu \in \mathcal{M}_\phi(B, q)} \left[\frac{1}{2} \ln(1 - q^2) - \frac{1}{2} \mathcal{E}_\phi(q, \nu) + \kappa_{\alpha, \phi}(q, \nu) - \alpha H(\nu | \mu_G) \right]. \quad (6)$$

The proof of Theorem 1 is presented in Section 2. The proof of Theorem 2 is a straightforward generalization, and sketched in Appendix C.1. The variational problems in Theorems 1 and 2 are challenging, as they imply an optimization on a set of measures, and they involve transforms of this measure that are very hard to access numerically. In Section 3 we present a drastic simplification: a heuristic calculation that allows one to reduce the supremum over the probability measure ν to a much more straightforward optimization over a small number of parameters.

As we have already stressed, the annealed complexity, although interesting in itself, is generically not representative of the landscape corresponding to a given typical instance of the empirical risk. In

¹A fully rigorous statement would imply a lower and an upper bound given by a supremum over the adherence and the interior of $\mathcal{M}_\phi(B)$. For reasons of lightness and clarity of the presentation we write it in the simpler presented form.

²Free multiplication is usually defined for positively-supported measures, however one can generalize it here by explicitly separating the positive and negative parts of ϕ'' (we can show freeness of the resulting two random matrices).

order to obtain the value of the quenched complexity we use the replicated Kac-Rice method, which is an extension to non-Gaussian functions of the one developed in [Ros et al. \[2019\]](#). Although the replica method is non-rigorous, it is considered an exact method in theoretical physics and it has been proven to give correct results for spin-glasses and inference problems [[Talagrand, 2006](#), [Barbier et al., 2019](#)]. We have obtained an explicit formula¹ for the quenched complexity of L_1 and L_2 at fixed values of the empirical risk, and overlap with the solution (in the L_2 case).

For L_1 , using the notations of [Theorem 1](#) we have:

Result 1 (Quenched complexity of L_1). *Let $B \subseteq \mathbb{R}$ an open interval.*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \ln \text{Crit}_{n, L_1}(B) = \frac{\ln \alpha - \alpha \ln 2\pi}{2} + \sup_{\substack{\nu \in \mathcal{M}_\phi(B) \\ q \in (0,1)}} \text{extr} \left\{ \kappa_{\alpha, \phi}(\nu, C) + \frac{1-\alpha}{2} \ln(1-q) \right. \\ \left. + \frac{1-\alpha q}{2(1-q)} - \int \nu(d\lambda) g(\lambda) - \frac{A\hat{A} - a\hat{a}}{2} + C(q\hat{c} - \hat{C}) - \frac{1}{2} \ln[A-a] - \frac{a}{2(A-a)} + \alpha \int_{\mathbb{R}^4} \mathcal{D}\xi \ln I(\xi) \right\}.$$

Here, $\xi = (\xi_q, \xi_a, \xi_c, \xi'_c)$ and $\mathcal{D}\xi$ is the standard Gaussian probability measure on \mathbb{R}^4 . The extr denotes extremization with respect to all variables $(A, \hat{A}, a, \hat{a}, C, \hat{C}, \hat{c}, \{g(\lambda)\})$. We denoted

$$I(\xi) \equiv \int_{\mathbb{R}} d\lambda e^{-\frac{\lambda^2}{2(1-q)} + \frac{g(\lambda)}{\alpha} + \frac{\hat{A}-\hat{a}}{2\alpha} \phi'(\lambda)^2 + \frac{\hat{C}-\hat{c}}{\alpha} \phi'(\lambda)\lambda + \frac{\sqrt{q}}{1-q} \xi_q \lambda + \sqrt{\frac{\alpha}{\alpha}} \xi_a \phi'(\lambda) + \sqrt{\frac{\alpha}{2\alpha}} [\phi'(\lambda)(\xi_c + i\xi'_c) + \lambda(\xi_c - i\xi'_c)]}.$$

In this formula, the notation extr denotes that one should set the partial derivatives with respect to the involved variables to zero. This notation arises from the replica calculation, which mixes saddle-point computations with Lagrange multipliers associated to certain constraints, and the precise meaning of this extremization (as a supremum or infimum) would have to be clarified by a more rigorous method. On a numerical point of view, one would have to solve the associated saddle-point equations, so that this precise meaning is not crucial for applications. We can state a very similar result for L_2 :

Result 2 (Quenched complexity of L_2). *Let $B \subseteq \mathbb{R}$, $Q \subseteq (-1, 1)$ two open intervals and:*

- For $m \in (-1, 1)$, $\mathcal{M}_\phi(B, m)$ is the space of probability measures ν on \mathbb{R}^2 that satisfy:

$$\begin{cases} \frac{1}{2} \int \nu(d\lambda^0, d\lambda) [\phi(\lambda) - \phi(\lambda^0)]^2 \in B, \\ \int \nu(d\lambda^0, d\lambda) \phi'(\lambda) [\phi(\lambda) - \phi(\lambda^0)] (\lambda^0 - m\lambda) = 0. \end{cases} \quad (7)$$

- Let $f(x, y) \equiv \phi''(y)[\phi(y) - \phi(x)] + \phi'(y)^2$. Let $\mathbf{z} \in \mathbb{R}^{n \times m}$ an i.i.d. standard Gaussian matrix, and $\mathbf{Y} \in \mathbb{R}^{m \times 2}$ with components taken i.i.d. from a probability measure ν on \mathbb{R}^2 . Let $D^{(\nu)}$ the diagonal matrix of size m with elements $D_\mu^{(\nu)} = f(Y_\mu)$. We define $\mu_{\alpha, \phi}[\nu]$ as the asymptotic spectral measure of $\mathbf{z} D^{(\nu)} \mathbf{z}^\top / m$.
- $\chi_{\alpha, \phi}(\nu, C) \equiv \int \mu_{\alpha, \phi}[\nu](dx) \ln |x - C|$.

One has:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \ln \text{Crit}_{n, L_2}(B, Q) = \sup_{\substack{m \in Q \\ q \in (0,1)}} \sup_{\nu \in \mathcal{M}_\phi(B, m)} \text{extr} \left[\frac{\ln \alpha - \alpha \ln 2\pi}{2} + \chi_{\alpha, \phi}(\nu, C) \right. \\ \left. + \frac{1-\alpha q - m^2}{2(1-q)} + \frac{1-\alpha}{2} \ln(1-q) - \frac{1}{2} \ln(A-a) - \frac{a}{2(A-a)} - \frac{A\hat{A}}{2} + \frac{a\hat{a}}{2} \right. \\ \left. - C_0 \hat{C}_0 - C \hat{C} + c \hat{c} - \int \nu(d\lambda^0, d\lambda) g(\lambda^0, \lambda) + \alpha \int_{\mathbb{R}^4 \times \mathbb{R}} \mathcal{D}\xi \mathcal{D}\lambda^0 \ln I(\lambda^0, \xi) \right].$$

¹Here we used a replica symmetric structure, which is correct in many cases, and a very good approximation in others where replica symmetry has to be broken [[Mézard et al., 1987](#)].

The extremum is made over all the variables $(A, a, C_0, C, c, \hat{A}, \hat{a}, \hat{C}, \hat{c}, \hat{C}_0, \{g(\lambda^0, \lambda)\})$. \mathcal{D} denotes the standard Gaussian measure, and the variables C_0, c, C are related by the additional constraint

$$-m(1-q)C_0 - (q-m^2)C + (1-m^2)c = 0.$$

$I(\lambda^0, \boldsymbol{\xi})$ is defined as, with $\boldsymbol{\xi} \equiv (\xi_q, \xi_a, \xi_c, \xi'_c)$:

$$I(\lambda^0, \boldsymbol{\xi}) \equiv \int_{\mathbb{R}} d\lambda e^{\frac{m}{1-q}\lambda^0\lambda - \frac{\lambda^2}{2(1-q)} + \frac{\sqrt{q-m^2}}{1-q}\xi_q\lambda + \frac{g(\lambda^0, \lambda)}{\alpha} + \frac{\hat{C}_0}{\alpha}\phi'(\lambda)[\phi(\lambda) - \phi(\lambda^0)]\lambda^0 + \frac{\hat{C}-\hat{c}}{\alpha}\lambda\phi'(\lambda)[\phi(\lambda) - \phi(\lambda^0)]} e^{\frac{\hat{A}-\hat{a}}{2\alpha}\phi'(\lambda)^2[\phi(\lambda) - \phi(\lambda^0)]^2 + \sqrt{\frac{\hat{a}}{\alpha}}\xi_a\phi'(\lambda)[\phi(\lambda) - \phi(\lambda^0)] + \sqrt{\frac{\hat{c}}{2\alpha}}[\phi'(\lambda)[\phi(\lambda) - \phi(\lambda^0)](\xi_c + i\xi'_c) + \lambda(\xi_c - i\xi'_c)]}.$$

The derivation of Result 1 is given in Section 4. Result 2 can be derived by a straightforward generalization of this computation, see Appendix C.1.

1.3 Conclusion, outcomes and perspectives

We have obtained analytical results for the annealed and quenched complexities of statistical models with non-Gaussian loss functions arising in generalized linear estimation and simple models of glasses and neural networks. Our method is versatile and can be easily extended to other cases. We describe in Appendix C.2 three other inference models to which it applies: binary linear classification, a mixture of two Gaussians, and a simple model of unsupervised learning.

As a sanity check of our results, we have analytically verified by explicit solution that for a linear activation function, the annealed complexities of L_1 is null. It is again a tedious but straightforward computation to check that the annealed complexity of L_1 with a quadratic activation $\phi(x) = x^2$ is also null, as the number of critical points in this case is linear with n . Note that for L_2 , even the case of a linear activation is non trivial, as shown in the very recent analysis of [Fyodorov and Tulin \[2019\]](#).

Our results allow for a complete characterization of the empirical loss landscapes of generalized linear models. The main issue ahead is determining for which class of functions ϕ and in which regimes (e.g. values of α), the annealed and quenched complexities become positive, i.e. when the associated landscape is rough. This will allow to study the connection between landscape properties and dynamics induced by local algorithms. In particular, it will shed light on the relationship between the roughness of the empirical loss landscape and the existence of “hard” phases in the learning of generalized linear models [[Barbier et al., 2019](#)]. It will also provide an interesting benchmark for obtaining the algorithmic thresholds of gradient descent (and variants) only through the knowledge of the landscape properties [[Srao Mannelli et al., 2019](#), [Mannelli et al., 2018](#)]. Based on ongoing works, we can for instance conjecture the existence of a rough landscape for small enough α in phase retrieval [[Lucibello et al., 2019](#)] and retarded learning [[Engel and Van den Broeck, 2001](#)].

Addressing these questions requires additional work, which is beyond the scope of this paper. We have explained in Section 3 how to make tractable the variational problem associated to the annealed and quenched complexity. Its analysis for specific models is an ongoing direction of research and will be presented elsewhere.

Another important extension of our results consists in counting the critical points of a fixed index (i.e. with a fixed number of negative directions in the spectrum of the Hessian). This would provide additional interesting information, in particular it would allow to differentiate local minima from the other critical points of the landscape, as was done for spin glass models in [Auffinger et al. \[2013\]](#). Such a counting would require to understand the large deviations properties of the eigenvalues of the Hessian arising for generalized linear models, a random matrix problem that is hard but hopefully tractable by building on recent developments [[Maida, 2007](#), [Ben Arous et al., 2019](#)]. This is an ongoing work that we are pursuing, with already promising results.

As final note, it is an open problem to generalize our methods to neural network models with many nodes and hidden layers; the random matrix analysis of the Hessian in this case is a particularly exciting challenge.

2 Proof of Theorem 1 for the annealed complexity

In this section we prove Theorem 1. The technique leverages the Kac-Rice formula and Sanov's theorem on the large deviations of the empirical measure of i.i.d. variables. First we precise our hypotheses on ϕ , that we will take in the following set of “well-behaved” activation functions:

Definition 1. $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is “well-behaved” if it is of class \mathcal{C}^3 and if, for $y \sim \mathcal{N}(0, 1)$, the random variable $a = \phi'(y)$ admits a continuous probability density in a neighborhood of $a = 0$.

2.1 The Kac-Rice formula

The first step is to apply the Kac-Rice formula to the random function L_1 :

Lemma 3 (Kac-Rice formula). *For any $\mathbf{x} \in \mathbb{S}^{n-1}$, denote $\text{grad}L_1(\mathbf{x})$ and $\text{Hess}L_1(\mathbf{x})$ the (Riemannian) gradient and Hessian of L_1 at the point \mathbf{x} . Then $\text{grad}L_1(\mathbf{x})$ has a well defined density (on the tangent space $T_{\mathbf{x}}\mathbb{S}^{n-1} \simeq \mathbb{R}^{n-1}$) in a neighborhood of zero, that we denote $\varphi_{\text{grad}L_1(\mathbf{x})}$. Denote $\mu_{\mathbb{S}}$ the usual surface measure on \mathbb{S}^{n-1} . One has:*

$$\mathbb{E} \text{Crit}_{n,L_1}(B) = \int_{\mathbb{S}^{n-1}} \varphi_{\text{grad}L_1(\mathbf{x})}(0) \mathbb{E} \left[\mathbb{1}_{L_1(\mathbf{x}) \in B} |\det \text{Hess}L_1(\mathbf{x})| \mid \text{grad}L_1(\mathbf{x}) = 0 \right] \mu_{\mathbb{S}}(d\mathbf{x}).$$

The proof of this lemma uses necessary conditions for a random function to be a.s. Morse¹, that are stated in [Azaïs and Wschebor \[2009\]](#). The details are given in [Appendix A.1](#).

2.2 The complexity at finite n

In this section we state the result of the Kac-Rice method. For $\mathbf{y} \in \mathbb{R}^m$, let $\Lambda(\mathbf{y}) \in \mathbb{R}^{m \times m}$:

$$\Lambda(\mathbf{y}) \equiv \left(\mathbf{I}_m - \frac{\phi'(\mathbf{y})\phi'(\mathbf{y})^\top}{\|\phi'(\mathbf{y})\|^2} \right) D(\mathbf{y}) \left(\mathbf{I}_m - \frac{\phi'(\mathbf{y})\phi'(\mathbf{y})^\top}{\|\phi'(\mathbf{y})\|^2} \right), \quad (8)$$

in which we (abusively) denote $\phi'(\mathbf{y}) \equiv (\phi'(y_\mu))_{\mu=1}^m$, and $D(\mathbf{y}) \in \mathbb{R}^{m \times m}$ the diagonal matrix with elements $D(\mathbf{y})_\mu = D(y_\mu) = \frac{n}{m} \phi''(y_\mu)$. The main result of this section is:

Lemma 4 (Complexity at finite n).

$$\mathbb{E} \text{Crit}_{n,L_1}(B) = \mathcal{C}_n e^{n \frac{1+\ln \alpha}{2}} \mathbb{E}_{\mathbf{y}} \left[\mathbb{1}_{\frac{1}{m} \sum_{\mu} \phi(y_\mu) \in B} e^{-\frac{n-1}{2} \ln \left(\frac{1}{m} \sum_{\mu} \phi'(y_\mu)^2 \right)} \mathbb{E}_{\mathbf{z}} [|\det H_n^\Lambda(\mathbf{y})|] \right],$$

in which \mathcal{C}_n is exponentially trivial, meaning $\lim_{n \rightarrow \infty} (1/n) \ln \mathcal{C}_n = 0$. The variable $\mathbf{y} \in \mathbb{R}^m$ follows $\mathcal{N}(0, \mathbf{I}_m)$, and $\mathbf{z} \in \mathbb{R}^{(n-1) \times m}$ has i.i.d. standard Gaussian matrix elements, independent of \mathbf{y} . $H_n^\Lambda(\mathbf{y})$ is a square matrix of size $(n-1)$ with the following distribution :

$$H_n^\Lambda(\mathbf{y}) \stackrel{d}{=} \frac{1}{n} \mathbf{z} \Lambda(\mathbf{y}) \mathbf{z}^\top - \left[\frac{1}{m} \sum_{\mu=1}^m y_\mu \phi'(y_\mu) \right] \mathbf{I}_{n-1}. \quad (9)$$

The rest of the section is devoted to the proof of Lemma 4. We start from the result of Lemma 3. The following proposition specifies the joint distribution of $(L_1(\mathbf{x}), \text{grad}L_1(\mathbf{x}), \text{Hess}L_1(\mathbf{x}))$:

¹A Morse function is a function whose critical points are all non-degenerate.

Proposition 5 (Distribution of the gradient and Hessian). *Let $\mathbf{x} \in \mathbb{S}^{n-1}$. Then $(L_1(\mathbf{x}), \text{grad } L_1(\mathbf{x}), \text{Hess } L_1(\mathbf{x}))$ follows the following joint distribution:*

$$\begin{cases} L_1(\mathbf{x}) \stackrel{d}{=} \frac{1}{m} \sum_{\mu=1}^m \phi(y_\mu), & (10a) \\ \text{grad } L_1(\mathbf{x}) \stackrel{d}{=} \frac{1}{m} \sum_{\mu=1}^m \phi'(y_\mu) \mathbf{z}_\mu, & (10b) \\ \text{Hess } L_1(\mathbf{x}) \stackrel{d}{=} \frac{1}{m} \sum_{\mu=1}^m \phi''(y_\mu) \mathbf{z}_\mu \mathbf{z}_\mu^\top - \left[\frac{1}{m} \sum_{\mu=1}^m y_\mu \phi'(y_\mu) \right] \mathbf{I}_{n-1}, & (10c) \end{cases}$$

in which $\mathbf{y} = (y_\mu)_{\mu=1}^m \sim \mathcal{N}(0, \mathbf{I}_m)$, $(\mathbf{z}_\mu)_{\mu=1}^m \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \mathbf{I}_{n-1})$, and all $\{y_\mu, \mathbf{z}_\nu\}$ are independent. We identified in these equations the tangent spaces $T_{\mathbf{x}} \mathbb{S}^{n-1}$ with \mathbb{R}^{n-1} .

Proof. Denote $P_{\mathbf{x}}^\perp$ the orthogonal projection on $\{\mathbf{x}\}^\perp$. For a smooth function $f : \mathbb{S}^{n-1} \rightarrow \mathbb{R}$, ∇f and $\nabla^2 f$ are its Euclidean gradient and Hessian. The Riemannian structure on \mathbb{S}^{n-1} induces the gradient and Hessian of f as $\text{grad } f(\mathbf{x}) = P_{\mathbf{x}}^\perp \nabla f$ and $\text{Hess } f(\mathbf{x}) = P_{\mathbf{x}}^\perp \nabla^2 f P_{\mathbf{x}}^\perp - (\mathbf{x} \cdot \nabla f(\mathbf{x})) P_{\mathbf{x}}^\perp$. Applying these formulas yields:

$$\text{grad } L_1(\mathbf{x}) = \frac{1}{m} \sum_{\mu=1}^m (P_{\mathbf{x}}^\perp \boldsymbol{\xi}_\mu) \phi'(\boldsymbol{\xi}_\mu \cdot \mathbf{x}), \quad (11)$$

$$\text{Hess } L_1(\mathbf{x}) = \frac{1}{m} \sum_{\mu=1}^m \phi''(\boldsymbol{\xi}_\mu \cdot \mathbf{x}) (P_{\mathbf{x}}^\perp \boldsymbol{\xi}_\mu) (P_{\mathbf{x}}^\perp \boldsymbol{\xi}_\mu)^\top - \left[\frac{1}{m} \sum_{\mu=1}^m (\boldsymbol{\xi}_\mu \cdot \mathbf{x}) \phi'(\boldsymbol{\xi}_\mu \cdot \mathbf{x}) \right] P_{\mathbf{x}}^\perp. \quad (12)$$

Letting $y_\mu \equiv \boldsymbol{\xi}_\mu \cdot \mathbf{x}$ and $\mathbf{z}_\mu \equiv P_{\mathbf{x}}^\perp \boldsymbol{\xi}_\mu$ (identified to an element of \mathbb{R}^{n-1}) yields the result. \square

The joint distribution of eq. (10) is invariant with respect to \mathbf{x} , thus we can chose \mathbf{x} to be the North pole $\mathbf{x} = \mathbf{e}_n = (\delta_{i,n})_{i=1}^n$. With $\omega_n \equiv 2\pi^{n/2}/\Gamma(n/2)$ the volume of \mathbb{S}^{n-1} , we obtain from Lemma 3:

$$\mathbb{E} \text{Crit}_{n,L_1}(B) = \omega_n \varphi_{\text{grad } L_1(\mathbf{e}_n)}(0) \mathbb{E} \left[|\det \text{Hess } L_1(\mathbf{e}_n)| \mathbb{1}_{L_1(\mathbf{e}_n) \in B} \mid \text{grad } L_1(\mathbf{e}_n) = 0 \right]. \quad (13)$$

Removing the \mathbf{e}_n indication and conditioning on the distribution of \mathbf{y} in eq. (10), we reach:

$$\mathbb{E} \text{Crit}_{n,L_1}(B) = \omega_n \mathbb{E}_{\mathbf{y}} \left[\mathbb{1}_{\frac{1}{m} \sum_{\mu} \phi(y_\mu) \in B} \varphi_{\text{grad } L_1 | \mathbf{y}}(0) \mathbb{E}_{\mathbf{z}} \left[|\det \text{Hess } L_1| \mid \text{grad } L_1 = 0, \mathbf{y} \right] \right].$$

Once conditioned on \mathbf{y} , eq. (10b) describes a Gaussian density so we can directly compute:

$$\begin{aligned} \omega_n \varphi_{\text{grad } L_1 | \mathbf{y}}(0) &= \frac{2\pi^{n/2}}{\Gamma(n/2)} \exp \left[-\frac{n-1}{2} \ln \left(\frac{2\pi}{m^2} \sum_{\mu=1}^m \phi'(y_\mu)^2 \right) \right], \\ &= C_n \exp \left[\frac{n}{2} + \frac{n}{2} \ln \frac{m}{n} - \frac{n-1}{2} \ln \left(\frac{1}{m} \sum_{\mu=1}^m \phi'(y_\mu)^2 \right) \right], \end{aligned} \quad (14)$$

in which $\ln C_n = \mathcal{O}_n(n)$ (using Stirling's formula). The conditioning of the Hessian by $\text{grad } L_1 = 0$ at fixed \mathbf{y} reduces to a linear conditioning on \mathbf{z} . One thus obtains by classical Gaussian conditioning:

$$\mathbb{E}_{\mathbf{z}} \left[|\det \text{Hess } L_1| \mid \text{grad } L_1 = 0, \mathbf{y} \right] = \mathbb{E}_{\mathbf{z}} \left[|\det H_n^\Lambda(\mathbf{y})| \right], \quad (15)$$

in which $H_n^\Lambda(\mathbf{y})$ is defined by eq. (9). This ends the proof of Lemma 4.

2.3 Large deviations

This section is devoted to the end of the proof of Theorem 1. Denote $\nu_y^m \equiv \frac{1}{m} \sum_{\mu=1}^m \delta_{y_\mu}$ the empirical distribution of \mathbf{y} . We take the notations of Theorem 1 and Lemma 4. We first state an important lemma on the concentration of $\mathbb{E}_z [|\det H_n^\Lambda(\mathbf{y})|]$ ¹:

Lemma 6. *There exists $\eta > 0$ such that for all $t > 0$:*

$$\lim_{n \rightarrow \infty} \frac{1}{n^{1+\eta}} \ln \mathbb{P} \left[\left| \frac{1}{n} \ln \mathbb{E}_z [|\det H_n^\Lambda(\mathbf{y})|] - \kappa_{\alpha, \phi}(\nu_y^m, t_\phi(\nu_y^m)) \right| \geq t \right] = -\infty. \quad (16)$$

The proof of Lemma 6 is detailed in Appendix A.2. Note that we expect this result to actually be valid up to $\eta = 1$, as the large deviations of the spectral distribution of random matrices is typically on the n^2 scale [Ben Arous and Guionnet, 1997, Hiai and Petz, 1998]. The following moment condition, proven in Appendix A.3, will be important:

Lemma 7. *For every $\gamma \in (1, \alpha)$ we have:*

$$\begin{cases} \limsup_{n \rightarrow \infty} \frac{1}{n} \ln \mathbb{E}_y \left[e^{\gamma n \left[-\frac{1}{2} \ln \left(\frac{1}{m} \sum_{\mu} \phi'(y_\mu)^2 \right) + \kappa_{\alpha, \phi}(\nu_y^m, t_\phi(\nu_y^m)) \right]} \right] < +\infty, & (17a) \\ \limsup_{n \rightarrow \infty} \frac{1}{n} \ln \mathbb{E}_y \left[e^{\gamma n \left[-\frac{1}{2} \ln \left(\frac{1}{m} \sum_{\mu} \phi'(y_\mu)^2 \right) + \frac{1}{n} \ln \mathbb{E}_z [|\det H_n^\Lambda(\mathbf{y})|] \right]} \right] < +\infty. & (17b) \end{cases}$$

Then we can conclude using Sanov's large deviation principle [Sanov, 1958, Dembo and Zeitouni, 1998], and Varadhan's lemma. Using our previous lemmas, we obtain the statement of Theorem 1:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln \mathbb{E} \text{Crit}_{n, L_1}(B) = \sup_{\nu \in \mathcal{M}_\phi(B)} \left[\frac{1 + \ln \alpha}{2} - \frac{\mathcal{E}_\phi(\nu)}{2} + \kappa_{\alpha, \phi}(\nu, t_\phi(\nu)) - \alpha H(\nu | \mu_G) \right]. \quad (18)$$

The proof of eq. (18) is detailed in Appendix A.4.

3 Towards a numerical solution to the variational problem

3.1 The logarithmic potential of $\mu_{\alpha, \phi}[\nu]$

Let $\nu \in \mathcal{M}(\mathbb{R})$. The Stieltjes transform $g(z) \equiv \int \mu(dt)(t - z)^{-1}$ of $\mu_{\alpha, \phi}[\nu]$ is given by the unique solution in \mathbb{C}_+ to the implicit equation (as shown for instance in Silverstein and Bai [1995]):

$$\forall z \in \mathbb{C}_+, \quad g(z) = - \left[z - \alpha \int \frac{\phi''(t)}{\alpha + \phi''(t)g(z)} \nu(dt) \right]^{-1}. \quad (19)$$

For any $\mu \in \mathcal{M}(\mathbb{R})$ and $t \in \mathbb{R}$ we define the *logarithmic potential* as $U[\mu](t) \equiv \int \mu(dx) \ln |x - t|$. It is well defined with values in $\mathbb{R} \cup \{\pm\infty\}$, see Faraut [2014] for a review on this subject. Our goal is to numerically compute $U[\mu](t)$ for $\mu = \mu_{\alpha, \phi}[\nu]$ and an arbitrary $t \in \mathbb{R}$, see Theorem 1. For clarity, we will write μ for $\mu_{\alpha, \phi}[\nu]$ for the remainder of this section. Let us define for any $z \in \mathbb{C}_+$, $G(z) \equiv \int \mu(dx) \ln(z - x)$. $G(z)$ is well defined and holomorphic on \mathbb{C}_+ . Moreover, from the Chapter II of Faraut [2014], we know that $U[\mu](t) = \lim_{\epsilon \rightarrow 0^+} \text{Re } G(t + i\epsilon)$.

From this it is clear that in order to get the logarithmic potential in the bulk, we need to be able to evaluate $G(z)$ for $z \in \mathbb{C}_+$. Define, for $z, g \in \mathbb{C}_+$ ²:

$$F(z, g) \equiv -\ln(-g) - zg + \alpha \int \nu(d\lambda) \ln(\alpha + \phi''(\lambda)g) - 1 - \alpha \ln \alpha. \quad (20)$$

¹In the proofs of this section we assume that $x\phi'(x)$ and $\phi''(x)$ are bounded. As one can always smoothly truncate the largest values of ϕ without affecting the complexity, this does not remove any generality to our results.

² $g \in \mathbb{C}_+$, and (since $\alpha > 1$) $\alpha + \phi''(\lambda)g \in \mathbb{C} \setminus (-\infty, 1]$, thus we can use the principal determination of the logarithm.

At any fixed z , $F(z, g)$ is an holomorphic function of g on \mathbb{C}_+ . Its Wirtinger derivative is:

$$\frac{\partial F}{\partial g}(z, g) = -\frac{1}{g} - z + \alpha \int \nu(d\lambda) \frac{\phi''(\lambda)}{\alpha + \phi''(\lambda)g}. \quad (21)$$

Thus $g(z)$ (the Stieltjes transform of μ , cf. eq. (19)) is the *only* $g \in \mathbb{C}_+$ such that $\frac{\partial F}{\partial g}(z, g) = 0$. Moreover, by definition $g(z)$ is an holomorphic function on \mathbb{C}_+ with values in \mathbb{C}_+ . We can thus apply the usual composition of derivatives and obtain:

$$\frac{dF}{dz}(z, g(z)) = -g(z). \quad (22)$$

Furthermore we know $\frac{dG}{dz} = -g(z)$. Computing the remaining constant by investigating the limit $\text{Re}[z] \rightarrow \infty$, we reach that $G(z) = F(z, g(z))$ for every $z \in \mathbb{C}_+$. We thus have the crucial relation:

$$\forall t \in \mathbb{R}, \quad U[\mu](t) = \lim_{\epsilon \rightarrow 0^+} \text{Re} F(t + i\epsilon, g(t + i\epsilon)). \quad (23)$$

This allows for an efficient numerical derivation of the logarithmic potential of $\mu_{\alpha, \phi}[\nu]$, as we will see in more details below.

3.2 Heuristic derivation of the simplified fixed point equations corresponding to Theorem 1

We present here an heuristic derivation of scalar fixed point equations for the numerical resolution of Theorem 1. This technique could be easily extended to Theorem 2 as well as the quenched calculations presented afterwards, but we restrict to this simpler case for the sake of the presentation.

3.2.1 Expressing $\kappa_{\alpha, \phi}(\nu, t)$

From eq. (23) we know that for every $t \in \mathbb{R}$:

$$\kappa_{\alpha, \phi}(\nu, t) = \lim_{\epsilon \rightarrow 0^+} \text{Re} \left[-\ln(-g(t + i\epsilon)) - (t + i\epsilon)g(t + i\epsilon) + \alpha \int \nu(d\lambda) \ln [\alpha + \phi''(\lambda)g(t + i\epsilon)] - 1 - \alpha \ln \alpha \right].$$

For every $t \in \mathbb{R}$, $g(t + i\epsilon)$ is the only solution in \mathbb{C}_+ to the partial derivative of the previous equation:

$$-\frac{1}{g} - (t + i\epsilon) + \alpha \int \nu(d\lambda) \frac{\phi''(\lambda)}{\alpha + \phi''(\lambda)g} = 0. \quad (24)$$

So heuristically, we can write that for a small enough ϵ :

$$\kappa_{\alpha, \phi}(\nu, t) = \text{extr}_{g \in \mathbb{C}_+} \left[-\ln |g| - tg_r + \epsilon g_i + \alpha \int \nu(d\lambda) \ln |\alpha + \phi''(\lambda)g| - 1 - \alpha \ln \alpha \right], \quad (25)$$

with $g = g_r + ig_i$ (in practice one iterates over g_r and g_i successively).

3.2.2 Heuristic solution to Theorem 1

We start from the result of Theorem 1. For a function f , we write $\mathbb{E}[f(X)] \equiv \int \nu(dt)f(t)$. We introduce Lagrange multipliers to fix the conditions $\mathbb{E}[\phi(X)] \in B$, and we fix the values of $\mathbb{E}[\phi'(X)^2]$ and $\mathbb{E}[X\phi'(X)]$. We obtain:

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \ln \mathbb{E} \text{ Crit}_{n, L_1}(B) &= \sup_{\substack{l \in B \\ \nu \in \mathcal{M}(\mathbb{R})}} \text{extr}_{\lambda_0, \lambda_1, \lambda_2} \sup_{A, t} \left[\frac{1 + \ln \alpha}{2} - \frac{1}{2} \ln A + \lambda_0 l + \lambda_1 A + \lambda_2 t \right. \\ &\quad \left. + \kappa_{\alpha, \nu}(\nu, t) - \alpha H(\nu | \mu_G) - \lambda_0 \mathbb{E}[\phi(X)] - \lambda_1 \mathbb{E}[\phi'(X)^2] - \lambda_2 \mathbb{E}[X\phi'(X)] \right]. \end{aligned} \quad (26)$$

Note that now the supremum over ν is unconstrained over the set $\mathcal{M}(\mathbb{R})$ of probability distributions. We now make use of eq. (25) to write, with $2K(\alpha) \equiv -1 + \ln \alpha - 2\alpha \ln \alpha$ and a small $\epsilon > 0$:

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \ln \mathbb{E} \text{Crit}_{n, L_1}(B) &= \sup_{\substack{l \in B \\ \nu \in \mathcal{M}(\mathbb{R})}} \text{extr}_{\substack{\{\lambda_i\}, A, t \\ g \in \mathbb{C}_+}} \left[K(\alpha) - \frac{1}{2} \ln A + \lambda_0 l + \lambda_1 A + \lambda_2 t \right. \\ &\quad - \ln |g| - t \text{Re}[g] + \epsilon \text{Im}[g] + \alpha \int \nu(d\lambda) \ln |\alpha + \phi''(\lambda)g| - \alpha H(\nu | \mu_G) - \lambda_0 \mathbb{E}[\phi(X)] \\ &\quad \left. - \lambda_1 \mathbb{E}[\phi'(X)^2] - \lambda_2 \mathbb{E}[X \phi'(X)] \right]. \end{aligned} \quad (27)$$

For any scalar function F , the maximum $\sup_{\nu} [\mathbb{E}[F(X)] - \alpha H(\nu | \mu_G)]$ is attained in ν^* with density proportional to $e^{-x^2/2 + F(x)/\alpha}$, which is called the *Gibbs measure* in statistical physics. This gives (\mathcal{D} is the standard Gaussian measure on \mathbb{R}):

$$\sup_{\nu \in \mathcal{M}(\mathbb{R})} [\mathbb{E}[F(X)] - \alpha H(\nu | \mu_G)] = \alpha \ln \left[\int_{\mathbb{R}} \mathcal{D}x e^{F(x)/\alpha} \right]. \quad (28)$$

Plugging this into our previous equation for the annealed complexity yields:

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \ln \mathbb{E} \text{Crit}_{n, L_1}(B) &= \sup_{l \in B} \text{extr}_{\substack{\{\lambda_i\} \\ g \in \mathbb{C}_+}} \sup_{A, t} \left\{ K(\alpha) + \lambda_0 l + \lambda_1 A + \lambda_2 t \right. \\ &\quad \left. - \frac{1}{2} \ln A - \ln |g| - t \text{Re}[g] + \epsilon \text{Im}[g] \right. \\ &\quad \left. + \alpha \ln \left[\int_{\mathbb{R}} \mathcal{D}x \exp \left\{ -\frac{\lambda_0 \phi(x) + \lambda_1 \phi'(x)^2 + \lambda_2 x \phi'(x)}{\alpha} + \ln |\alpha + \phi''(x)g| \right\} \right] \right\}. \end{aligned} \quad (29)$$

This can be further simplified, as the extrema over A, t are trivially solved and give the value of $\lambda_2 = \text{Re}[g]$ and $\lambda_1 = (2A)^{-1}$. Thus we obtain:

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \ln \mathbb{E} \text{Crit}_{n, L_1}(B) &= \\ &\sup_{\substack{l \in B \\ g \in \mathbb{C}_+}} \text{extr}_{\{\lambda_0, \lambda_1\}} \left\{ K(\alpha) + \lambda_0 l + \frac{1 + \ln 2}{2} + \frac{1}{2} \ln \lambda_1 - \ln |g| + \epsilon \text{Im}[g] \right. \\ &\quad \left. + \alpha \ln \left[\int_{\mathbb{R}} \mathcal{D}x \exp \left\{ -\frac{\lambda_0 \phi(x) + \lambda_1 \phi'(x)^2 + \text{Re}[g] x \phi'(x)}{\alpha} + \ln |\alpha + \phi''(x)g| \right\} \right] \right\}. \end{aligned} \quad (30)$$

Let us now denote

$$\langle \cdots \rangle_{\lambda_0, \lambda_1, g} \equiv \frac{\int_{\mathbb{R}} \mathcal{D}x (\cdots) \exp \left\{ -\alpha^{-1} [\lambda_0 \phi(x) + \lambda_1 \phi'(x)^2 + \text{Re}[g] x \phi'(x)] + \ln |\alpha + \phi''(x)g| \right\}}{\int_{\mathbb{R}} \mathcal{D}x \exp \left\{ -\alpha^{-1} [\lambda_0 \phi(x) + \lambda_1 \phi'(x)^2 + \text{Re}[g] x \phi'(x)] + \ln |\alpha + \phi''(x)g| \right\}},$$

then the fixed point equations of eq. (30) can be written as:

$$l = \langle \phi(x) \rangle_{\lambda_0, \lambda_1, g}, \quad (31a)$$

$$\frac{1}{2\lambda_1} = \langle \phi'(x)^2 \rangle_{\lambda_0, \lambda_1, g}, \quad (31b)$$

$$-\frac{\text{Re}[g]}{|g|^2} = \left\langle x \phi'(x) - \frac{\alpha \phi''(x)(\alpha + \phi''(x)\text{Re}[g])}{|\alpha + \phi''(x)g|^2} \right\rangle_{\lambda_0, \lambda_1, g}, \quad (31c)$$

$$\epsilon - \frac{\text{Im}[g]}{|g|^2} = - \left\langle \frac{\alpha \phi''(x)^2 \text{Im}[g]}{|\alpha + \phi''(x)g|^2} \right\rangle_{\lambda_0, \lambda_1, g}. \quad (31d)$$

These equations are to be iterated over λ_0, λ_1, g , and l (while enforcing the constraint $l \in B$). From experience, the best procedure is to start from the solution of the unconstrained problem (without any constraint on the loss value), before smoothly following the solution while adding the constraint. In the case of $L_2(\mathbf{x})$, one would follow a similar procedure.

4 The quenched complexity and the replica method

In this section we detail the principle of the quenched calculation that gives rise to Results 1-2. For the sake of the presentation we restrict to Result 1, while Result 2 will be discussed in Appendix C.1. We therefore focus on the function L_1 of eq. (1). As the very basis of this calculation is non-rigorous we present this calculation in a fashion closer to theoretical physics standards, differently from Section 2 in which we present rigorous results on the annealed complexity. Some technicalities will be postponed to Appendix B.

4.1 The replica trick and the p -th moment

The replica method is a heuristic tool of theoretical physics that allows to compute the quenched values of observables in the thermodynamic (i.e. $n \rightarrow \infty$) limit from the knowledge of their integer moments, under some assumptions. It is based on the non-rigorous identity (note that it involves an inversion of limits), for a strictly positive function f of a n -dimensional random vector \mathbf{x} :

$$\lim_{n \rightarrow \infty} \mathbb{E} \ln f(\mathbf{x}) = \lim_{p \rightarrow 0^+} \lim_{n \rightarrow \infty} \frac{1}{p} \ln [\mathbb{E} f(\mathbf{x})^p]. \quad (32)$$

Mézard et al. [1987] gives a comprehensive introduction to the replica method and its (many) physical insights and consequences. Let $B \subseteq \mathbb{R}$ an open interval. The Kac-Rice formula can be stated for the p -th moment of the complexity [Azäis and Wschebor, 2008, Adler and Taylor, 2009]:

$$\begin{aligned} \mathbb{E} \text{Crit}_{n, L_1}(B)^p &= \left[\prod_{a=1}^p \int_{\mathbb{S}^{n-1}} \mu_{\mathbb{S}}(d\mathbf{x}^a) \right] \mathbb{1} [\{L_1(\mathbf{x}^a) \in B\}_{a=1}^p] \varphi_{\{\text{grad } L_1(\mathbf{x}^a)\}_{a=1}^p}(0) \\ &\times \mathbb{E} \left[\prod_{a=1}^p |\det \text{Hess } L_1(\mathbf{x}^a)| \left| \{\text{grad } L_1(\mathbf{x}^a)\}_{a=1}^p = 0 \right. \right]. \end{aligned} \quad (33)$$

Here, $\varphi_{\{\text{grad } L_1(\mathbf{x}^a)\}_{a=1}^p}(0)$ represents the joint density of the p gradients, taken at 0. Note that the non-linearity $L_1(\mathbf{x})$ only depends on the parameters $y_\mu^a \equiv \boldsymbol{\xi}_\mu \cdot \mathbf{x}^a$, so we will often write $L_1(\mathbf{y}) \equiv L_1(\mathbf{x})$. Proceeding as in the annealed case, we can rewrite the expectations by conditioning over $\{\mathbf{y}^a\}_{a=1}^p$:

$$\begin{aligned} \mathbb{E} \text{Crit}_{n, L_1}(B)^p &= \left[\prod_{a=1}^p \int \mu_{\mathbb{S}}(d\mathbf{x}^a) \right] \mathbb{E}_{\{\mathbf{y}^a\}} \left[\mathbb{1} [\{L_1(\mathbf{y}^a) \in B\}_{a=1}^p] \varphi_{\{\text{grad } L_1(\mathbf{x}^a)\}_{a=1}^p} \Big|_{\{\mathbf{y}^a\}}(0) \right. \\ &\left. \mathbb{E} \left[\prod_{a=1}^p |\det \text{Hess } L_1(\mathbf{x}^a)| \left| \{\text{grad } L_1(\mathbf{x}^a) = 0, \mathbf{y}^a\}_{a=1}^p \right. \right] \right]. \end{aligned} \quad (34)$$

The gradient and Hessian at \mathbf{x}^a live in the tangent plane to the sphere at \mathbf{x}^a , identified with \mathbb{R}^{n-1} . Note that the $\{y_\mu^a\}$ are Gaussian variables with zero mean and covariance $\mathbb{E}[y_\mu^a y_\nu^b] = \delta_{\mu\nu} q_{ab}$, with $q_{ab} \equiv \mathbf{x}^a \cdot \mathbf{x}^b$ the ‘‘overlap’’ between replicas a and b . We introduce the variables $\{q_{ab}\}$ via delta functions in eq. (34):

$$\begin{aligned} \mathbb{E} \text{Crit}_{n, L_1}(B)^p &= \left[\prod_{a=1}^p \int \mu_{\mathbb{S}}(d\mathbf{x}^a) \right] \left[\prod_{a < b} \int dq_{ab} \delta(q_{ab} - \mathbf{x}^a \cdot \mathbf{x}^b) \right] \mathbb{E}_{\{\mathbf{y}^a\}} \left\{ \mathbb{1} [\{L_1(\mathbf{y}^a) \in B\}_{a=1}^p] \right. \\ &\left. \varphi_{\{\text{grad } L_1(\mathbf{x}^a)\}_{a=1}^p} \Big|_{\{\mathbf{y}^a\}}(0) \mathbb{E} \left[\prod_{a=1}^p |\det \text{Hess } L_1(\mathbf{x}^a)| \left| \{\text{grad } L_1(\mathbf{x}^a) = 0, \mathbf{y}^a\}_{a=1}^p \right. \right] \right\}. \end{aligned} \quad (35)$$

Since we fixed the $\{q_{ab}\}$, the distribution of the $\{\mathbf{y}^a\}$ is fixed, as well as the joint distribution of the loss, gradients and Hessians, as we will explicit in the following. As the number of overlap variables is $p(p-1)/2 = \mathcal{O}_n(1)$, we will perform a saddle-point over the variables $\{q_{ab}\}$ in the thermodynamic limit. The replica-symmetric assumption (see Mézard et al. [1987]) is a crucial hypothesis that can be made

in the framework of the replica method. It amounts to assume that, once the saddle-point is performed, the extremizing $\{q_{ab}\}$ are ‘‘symmetric’’ over the different replicas of the system. Concretely, we assume that the variables $\{q_{ab}\}$ satisfy $q_{aa} = 1$, $q_{ab} = q$ for $a \neq b$. Assuming this structure of the overlap matrix allows to extend the expression of the moments to arbitrary non-integer p , and then to take the $p \rightarrow 0^+$ limit as needed in eq. (32). We used a replica symmetric structure, which is correct in many cases, and a very good approximation in others where replica symmetry has to be broken.

4.1.1 The phase volume factor

Let us first compute the phase space factor in eq. (35). More precisely, the term:

$$\left[\prod_{a=1}^p \int \mu_{\mathbb{S}}(d\mathbf{x}^a) \right] \left[\prod_{a<b} \delta(q_{ab} - \mathbf{x}^a \cdot \mathbf{x}^b) \right] = n^{-\frac{p(p-1)}{2}} \left[\prod_{a=1}^p \int_{\mathbb{R}^n} d\mathbf{x}^a \prod_{a \leq b} \delta(nq_{ab} - n\mathbf{x}^a \cdot \mathbf{x}^b) \right],$$

in which we denoted $q_{aa} = 1$. As we detail in Appendix B.1 we reach, when $p \rightarrow 0^+$ and $n \rightarrow \infty$:

$$\frac{1}{np} \ln \left[\prod_{a=1}^p \int_{\mathbb{R}^n} d\mathbf{x}^a \prod_{a \leq b} \delta(nq_{ab} - n\mathbf{x}^a \cdot \mathbf{x}^b) \right] \simeq \frac{1}{2} \log \frac{2\pi}{n} + \frac{1}{2} \left[\frac{1}{1-q} + \log(1-q) \right]. \quad (36)$$

4.1.2 The joint density of the gradients

We will now compute the joint density of the gradients at $\{\mathbf{x}^a\}$, conditioned on the values of $\{\mathbf{y}^a\}$. The calculation is an extension of Sections V.C and V.E of Ros et al. [2019]. We consider two vectors \mathbf{x}^a and \mathbf{x}^b of overlap $q_{ab} = q$. It is easy to see that $\mathbb{E}[\text{grad } L(\mathbf{x}^a) | \{\mathbf{y}^b\}_{b=1}^p] = 0$ from eq. (10b), so we will focus on the covariance matrix $\mathbb{E}[\text{grad } L(\mathbf{x}^a) \text{grad } L(\mathbf{x}^b)^\top | \{\mathbf{y}^c\}_{c=1}^p]$. After some calculations detailed in Appendix B.2 we get the gradient density at leading exponential order:

$$\begin{aligned} \varphi_{\{\text{grad } L_1(\mathbf{x}^a)\}_{a=1}^p | \{\mathbf{y}^a\}}(0) &\simeq \prod_{a \neq b} \delta \left[\frac{1}{m} \sum_{\mu=1}^m \phi'(y_\mu^a) \left(z_p(q) y_\mu^a + f_p^0(q) y_\mu^b + f_p(q) \sum_{c(\neq a,b)} y_\mu^c \right) \right] \\ &\times \exp \left\{ \frac{np}{2} \log \frac{m}{2\pi} - \frac{n}{2} \ln \det \left[\left(\frac{1}{m} \sum_{\mu=1}^m \phi'(y_\mu^a) \phi'(y_\mu^b) \right)_{1 \leq a, b \leq p} \right] \right\}, \quad (37) \end{aligned}$$

in which the auxiliary functions $(z_p(q), f_p(q), f_p^0(q))$ are defined in eq. (87).

4.1.3 Factorization of the mean product of determinants

The argument of this section is very close to Section V.F of Ros et al. [2019]. We consider the term:

$$\mathbb{E} \left[\prod_{a=1}^p |\det \text{Hess } L_1(\mathbf{x}^a)| \middle| \{\text{grad } L_1(\mathbf{x}^a) = 0, \mathbf{y}^a\}_{a=1}^p \right]. \quad (38)$$

We make two important remarks, which are straightforward transpositions of the arguments of Ros et al. [2019] to our problem, and we refer to this work for more extensive physical justifications.

- The conditioning over the gradients being zero, as in the annealed calculation, only gives a finite-rank change to the Hessians $\text{Hess } L_1(\mathbf{x}^a)$ and thus does not modify the limit at the scale $e^{\Theta(n)}$. At this scale, the statistics of the p matrices $\{\text{Hess } L_1(\mathbf{x}^a)\}_{a=1}^p$ are identical.
- The spectral measure of $\text{Hess } L_1(\mathbf{x}^a)$ concentrates at a rate at least $n^{1+\epsilon}$ for a small enough $\epsilon > 0$ (we expect that the actual rate is n^2), so that at the order $e^{\Theta(n)}$ the expectation value factorizes and we can assume all the Hessians to be independent.

Before stating the consequences of such remarks, we give some definitions:

- $\mu_{G,q}$ is the Gaussian probability measure on \mathbb{R}^p with zero mean and covariance $\mathbb{E}[X_a X_b] = (1-q)\delta_{ab} + q$. Note that $\{\mathbf{y}_\mu\}_{\mu=1}^m$ are i.i.d. variables distributed according to $\mu_{G,q}$.
- We define $\nu_{\mathbf{y}}$ as the empirical measure of $(\mathbf{y}_1, \dots, \mathbf{y}_m)$, that is $\nu_{\mathbf{y}} \equiv \frac{1}{m} \sum_{\mu} \delta_{\mathbf{y}_\mu}$. For every a , we denote $\nu_{\mathbf{y}}^a$ its marginal distribution: $\nu_{\mathbf{y}}^a(d\lambda^a) \equiv \int \prod_{b(\neq a)} \nu_{\mathbf{y}}(d\lambda)$. Then $\nu_{\mathbf{y}}^a$ is also the empirical distribution of $(y_\mu^a)_{\mu=1}^m$.

Our remarks show that we can use the results of the annealed calculation, and we have here by factorization of the expectation of the determinants, at leading exponential order:

$$\mathbb{E} \left[\prod_{a=1}^p |\det \text{Hess } L_1(\mathbf{x}^a)| \left| \left\{ \text{grad } L_1(\mathbf{x}^a) \right\}_{a=1}^p = 0, \{\mathbf{y}^a\} \right. \right] \simeq e^{n \sum_{a=1}^p \kappa_{\alpha, \phi}(\nu_{\mathbf{y}}^a, t_\phi(\nu_{\mathbf{y}}^a))}. \quad (39)$$

4.2 Decoupling the replicas and the $p \rightarrow 0^+$ limit

We can then apply Sanov's theorem to the empirical measure $\nu_{\mathbf{y}} \in \mathcal{M}(\mathbb{R}^p)$. Recall that we have constraints on this measure by the density of the gradient and the fixation of the energy level. More precisely, we denote $\mathcal{M}_\phi^{(p)}(q, B)$ the set of probability measures on \mathbb{R}^p that satisfy the following:

$$\left\{ \begin{array}{l} \forall 1 \leq a \leq p, \int \nu(d\lambda) \phi(\lambda^a) \in B, \\ \forall 1 \leq a \neq b \leq p, \int \nu(d\lambda) \phi'(\lambda^a) \left[z_p(q)\lambda^a + f_p^0(q)\lambda^b + f_p(q) \sum_{c(\neq a,b)} \lambda^c \right] = 0. \end{array} \right. \quad (40a)$$

$$\left[z_p(q)\lambda^a + f_p^0(q)\lambda^b + f_p(q) \sum_{c(\neq a,b)} \lambda^c \right] = 0. \quad (40b)$$

Recall that the functions $(z_p(q), f_p^0(q), f_p(q))$ are defined in eq. (87). Leveraging from the results of eqs. (36), (37) and (39), we obtain from Sanov's theorem and Varadhan's lemma:

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \ln \mathbb{E} [\text{Crit}_{n, L_1}(B)^p] &= \frac{p}{2} \ln \alpha + \sup_{q \in (0,1)} \sup_{\nu \in \mathcal{M}_\phi^{(p)}(q, B)} \left[\frac{p}{2} \left(\frac{1}{1-q} + \ln(1-q) \right) \right. \\ &\quad \left. - \frac{1}{2} \ln \det \left[\left(\int \nu(d\lambda) \phi'(\lambda^a) \phi'(\lambda^b) \right)_{1 \leq a, b \leq p} \right] + \sum_{a=1}^p \kappa_{\alpha, \phi}(\nu^a, t_\phi(\nu^a)) - \alpha H(\nu | \mu_{G,q}) \right]. \end{aligned} \quad (41)$$

Recall that ν^a is the marginal distribution of ν for the variable λ^a . We can then decouple the replicas under an assumption on the measure ν that amounts for replica symmetry. We stress that this replica symmetric assumption in the Kac-Rice calculation actually corresponds to a 1-step replica symmetry breaking (1RSB) structure of the zero-temperature Gibbs measure, that is an exponential number of single-point metastable states that all have the same two-point overlap. While possibly not exact, this assumption should already yield a good approximation to the landscape, and could be analytically checked by studying the stability of the replica-symmetric ansatz within replica theory. This allows to take subsequently the $p \rightarrow 0^+$ limit, and after some simplifications, we reach from eq. (41) the expression of Result 1. These steps are fairly technical, and are postponed to Appendix B.3.

Acknowledgments

We thank for interesting discussions Chiara Cammarota, Florent Krzakala, Valentina Ros, Lenka Zdeborová. GB acknowledges support by the Simons Foundation collaboration Cracking the Glass Problem (No. 454935 to G. Biroli). Additional funding is acknowledged by AM from ‘‘Chaire de recherche sur les modèles et sciences des données’’, Fondation CFM pour la Recherche-ENS.

References

- R.J. Adler and J.E. Taylor. *Random Fields and Geometry*. Springer Monographs in Mathematics. Springer New York, 2009.
- Greg W Anderson, Alice Guionnet, and Ofer Zeitouni. *An introduction to random matrices*, volume 118. Cambridge university press, 2010.
- Antonio Auffinger and Gérard Ben Arous. Complexity of random smooth functions on the high-dimensional sphere. *The Annals of Probability*, 41(6):4214–4247, 2013.
- Antonio Auffinger, Gérard Ben Arous, and Jiri Cerny. Random matrices and complexity of spin glasses. *Communications on Pure and Applied Mathematics*, 66(2):165–201, 2013.
- Jean-Marc Azaïs and Mario Wschebor. *Level sets and extrema of random processes and fields*. John Wiley & Sons, 2009.
- Jean-Marc Azäis and Mario Wschebor. *Rice Formulas for Random Fields*, pages 160–183. John Wiley and Sons, Inc., 2008.
- Zhidong Bai and Jack W Silverstein. *Spectral analysis of large dimensional random matrices*, volume 20. Springer, 2010.
- Jean Barbier, Florent Krzakala, Nicolas Macris, Léo Miolane, and Lenka Zdeborová. Optimal errors and phase transitions in high-dimensional generalized linear models. *Proceedings of the National Academy of Sciences*, 116(12):5451–5460, 2019.
- Gérard Ben Arous and Alice Guionnet. Large deviations for wigner’s law and voiculescu’s non-commutative entropy. *Probability theory and related fields*, 108(4):517–542, 1997.
- Gérard Ben Arous, Reza Gheissari, and Aukosh Jagannath. Algorithmic thresholds for tensor pca. *arXiv preprint arXiv:1808.00921*, 2018.
- Gérard Ben Arous, Song Mei, Andrea Montanari, and Mihai Nica. The landscape of the spiked tensor model. *Communications on Pure and Applied Mathematics*, 72(11):2282–2330, 2019.
- Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Global optimality of local search for low rank matrix recovery. In *Advances in Neural Information Processing Systems*, pages 3873–3881, 2016.
- Giulio Biroli, Chiara Cammarota, and Federico Ricci-Tersenghi. How to iron out rough landscapes and get optimal performances: Replicated gradient descent and its application to tensor pca. *arXiv preprint arXiv:1905.12294*, 2019.
- Alan J Bray and David S Dean. Statistics of critical points of gaussian fields on large-dimensional spaces. *Physical review letters*, 98(15):150201, 2007.
- Alan J Bray and MA Moore. Broken replica symmetry and metastable states in spin glasses. *Journal of Physics C: Solid State Physics*, 13(31):L907, 1980.
- Andrea Cavagna, Juan P Garrahan, and Irene Giardina. Quenched complexity of the mean-field p-spin spherical model with external magnetic field. *Journal of Physics A: Mathematical and General*, 32(5): 711, 1999.
- A Crisanti and H-J Sommers. Thouless-anderson-palmer approach to the spherical p-spin spin glass model. *Journal de Physique I*, 5(7):805–813, 1995.
- A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*. Applications of mathematics. Springer, 1998.

- Simon Du, Jason Lee, Yuandong Tian, Aarti Singh, and Barnabas Póczos. Gradient descent learns one-hidden-layer CNN: Don't be afraid of spurious local minima. In *Proceedings of the 35th International Conference on Machine Learning*, pages 1339–1348. PMLR, 2018.
- Andreas Engel and Christian Van den Broeck. *Statistical mechanics of learning*. Cambridge University Press, 2001.
- Jacques Faraut. Logarithmic potential theory, orthogonal polynomials, and random matrices. *Modern methods in multivariate statistics, Lecture Notes of CIMPA-FECYT-UNESCO-ANR. Hermann*, 2014.
- Silvio Franz and Giorgio Parisi. The simplest model of jamming. *Journal of Physics A: Mathematical and Theoretical*, 49(14):145001, 2016.
- C Daniel Freeman and Joan Bruna. Topology and geometry of half-rectified network optimization. *arXiv preprint arXiv:1611.01540*, 2016.
- Yan V Fyodorov. Complexity of random energy landscapes, glass transition, and absolute value of the spectral determinant of random matrices. *Physical review letters*, 92(24):240601, 2004.
- Yan V Fyodorov and Celine Nadal. Critical behavior of the number of minima of a random landscape at the glass transition point and the tracy-widom distribution. *Physical review letters*, 109(16):167203, 2012.
- Yan V Fyodorov and Rashel Tublin. Counting stationary points of the loss function in the simplest constrained least-square optimization. *arXiv preprint arXiv:1911.12452*, 2019.
- Yan V Fyodorov and Ian Williams. Replica symmetry breaking condition exposed by random matrix calculation of landscape complexity. *Journal of Statistical Physics*, 129(5-6):1081–1116, 2007.
- Rong Ge and Tengyu Ma. On the optimization landscape of tensor decompositions. In *Advances in Neural Information Processing Systems*, pages 3656–3666, 2017.
- Rong Ge, Jason D Lee, and Tengyu Ma. Matrix completion has no spurious local minimum. In *Advances in Neural Information Processing Systems*, pages 2973–2981, 2016.
- Rong Ge, Chi Jin, and Yi Zheng. No spurious local minima in nonconvex low rank problems: A unified geometric analysis. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, page 1233–1242, 2017.
- Fumio Hiai and Dénes Petz. Eigenvalue density of the wishart matrix and large deviations. *Infinite Dimensional Analysis, Quantum Probability and Related Topics*, 1(04):633–646, 1998.
- Alan J Hoffman and Helmut W Wielandt. The variation of the spectrum of a normal matrix. In *Selected Papers Of Alan J Hoffman: With Commentary*, pages 118–120. World Scientific, 2003.
- Kenji Kawaguchi. Deep learning without poor local minima. In *Advances in Neural Information Processing Systems*, pages 586–594, 2016.
- J Kurchan. Replica trick to calculate means of absolute values: applications to stochastic equations. *Journal of Physics A: Mathematical and General*, 24(21):4969, 1991.
- Shuyang Ling, Ruitu Xu, and Afonso S Bandeira. On the landscape of synchronization networks: A perspective from nonconvex optimization. *SIAM Journal on Optimization*, 29(3):1879–1907, 2019.
- Haihao Lu and Kenji Kawaguchi. Depth creates no bad local minima. *arXiv preprint arXiv:1702.08580*, 2017.

- Carlo Lucibello, Luca Saglietti, and Yue Lu. Generalized approximate survey propagation for high-dimensional estimation. In *International Conference on Machine Learning*, pages 4173–4182, 2019.
- Mylène Maida. Large deviations for the largest eigenvalue of rank one deformations of gaussian ensembles. *Electronic Journal of Probability*, 12:1131–1150, 2007.
- Stefano Sarao Mannelli, Giulio Biroli, Chiara Cammarota, Florent Krzakala, Pierfrancesco Urbani, and Lenka Zdeborová. Marvels and pitfalls of the langevin algorithm in noisy high-dimensional inference. *arXiv preprint arXiv:1812.09066*, 2018.
- Stefano Sarao Mannelli, Florent Krzakala, Pierfrancesco Urbani, and Lenka Zdeborova. Passed & spurious: Descent algorithms and local minima in spiked matrix-tensor models. In *International Conference on Machine Learning*, pages 4333–4342, 2019.
- Peter McCullagh. Generalized linear models. *European Journal of Operational Research*, 16(3):285–292, 1984.
- Song Mei, Yu Bai, Andrea Montanari, et al. The landscape of empirical risk for nonconvex losses. *The Annals of Statistics*, 46(6A):2747–2774, 2018.
- Marc Mézard, Giorgio Parisi, and Miguel Virasoro. *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, volume 9. World Scientific Publishing Company, 1987.
- John Ashworth Nelder and Robert WM Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)*, 135(3):370–384, 1972.
- Dohyung Park, Anastasios Kyrillidis, Constantine Carmanis, and Sujay Sanghavi. Non-square matrix sensing without spurious local minima via the Burer-Monteiro approach. In *Artificial Intelligence and Statistics*, pages 65–74, 2017.
- Valentina Ros, Gérard Ben Arous, Giulio Biroli, and Chiara Cammarota. Complex energy landscapes in spiked-tensor and simple glassy models: Ruggedness, arrangements of local minima, and phase transitions. *Physical Review X*, 9(1):011003, 2019.
- Ivan N Sanov. On the probability of large deviations of random variables. Technical report, North Carolina State University. Dept. of Statistics, 1958.
- Stefano Sarao Mannelli, Giulio Biroli, Chiara Cammarota, Florent Krzakala, and Lenka Zdeborová. Who is afraid of big bad minima? analysis of gradient-flow in spiked matrix-tensor models. In *Advances in Neural Information Processing Systems 32*, pages 8679–8689. 2019.
- Jack W Silverstein. Strong convergence of the empirical distribution of eigenvalues of large dimensional random matrices. *Journal of Multivariate Analysis*, 55(2):331–339, 1995.
- Jack W Silverstein and ZD Bai. On the empirical distribution of eigenvalues of a class of large dimensional random matrices. *Journal of Multivariate analysis*, 54(2):175–192, 1995.
- Jack W Silverstein and Sang-II Choi. Analysis of the limiting spectral distribution of large dimensional random matrices. *Journal of Multivariate Analysis*, 54(2):295–309, 1995.
- Daniel Soudry and Yair Carmon. No bad local minima: Data independent training error guarantees for multilayer neural networks. *arXiv preprint arXiv:1605.08361*, 2016.
- Eliran Subag. The complexity of spherical p -spin models—a second moment approach. *The Annals of Probability*, 45(5):3385–3450, 2017.
- Michel Talagrand. Free energy of the spherical mean field model. *Probability theory and related fields*, 134(3):339–382, 2006.

Dan Voiculescu. Multiplication of certain non-commuting random variables. *Journal of Operator Theory*, pages 223–235, 1987.

Hermann Weyl. Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen (mit einer anwendung auf die theorie der hohlraumstrahlung). *Mathematische Annalen*, 71(4): 441–479, 1912.

A Technical steps of the proof of Theorem 1

A.1 Proof of Lemma 3

We will apply the Kac-Rice machinery in the form of the remark made in Paragraph 6.1.4 of [Azäis and Wschebor \[2008\]](#). We recall it as a theorem:

Theorem 8 (Azais-Wschebor). *Let $k, d \in \mathbb{N}^*$. Let $Z : U \rightarrow \mathbb{R}^d$ be a random field, in which U is an open subset of \mathbb{R}^d . Assume that for every $t \in U$, we can write $Z(t) = H[Y(t)]$, such that:*

- (i) $\{Y(t), t \in U\}$ is a Gaussian random field with values in \mathbb{R}^k , C^1 paths, and such that for every $t \in U$, the distribution of $Y(t)$ is non-degenerate.
- (ii) $H : \mathbb{R}^k \rightarrow \mathbb{R}^d$ is a C^1 function.
- (iii) For all $t \in U$, $Z(t)$ has a density $\varphi_{Z(t)}(x)$, which is a continuous function of $(t, x) \in U \times \mathbb{R}^d$.
- (iv) $\mathbb{P}[\exists t \in U \text{ s.t. } Z(t) = 0 \text{ and } \det[\nabla Z(t)] = 0] = 0$.

Define, for every compact set $B \subseteq U$, $N(Z, B)$ to be the (finite) number of zeros of Z in B . Then:

$$\mathbb{E}[N(Z, B)] = \int_B \mathbb{E} \left[|\det \nabla Z(t)| \mathbb{1}_{Z(t) = 0} \right] \varphi_{Z(t)}(0) dt. \quad (42)$$

We wish to apply this theorem to the gradient $\text{grad} L_1(\mathbf{x})$. Verifying its hypotheses will end the proof of Lemma 3. We denote $\boldsymbol{\xi} \in \mathbb{R}^{n \times m}$ the matrix $\{\xi_{i\mu}\} = \{(\boldsymbol{\xi}_\mu)_i\}$, ∇L_1 the Euclidean gradient of L_1 , and $P_{\mathbf{x}}^\perp$ the orthogonal projection on $T_{\mathbf{x}}S^{n-1}$. Since $\text{grad} L_1(\mathbf{x}) = P_{\mathbf{x}}^\perp \nabla L_1(\mathbf{x})$ we have:

$$\text{grad} L_1(\mathbf{x}) = \frac{1}{m} \sum_{\mu=1}^m \left(P_{\mathbf{x}}^\perp \boldsymbol{\xi}_\mu \right) \phi'(\boldsymbol{\xi}_\mu \cdot \mathbf{x}). \quad (43)$$

We will apply Theorem 8 with $d = n - 1$ and $k = m \times n$. The Gaussian random field $Y(\mathbf{x}) \in \mathbb{R}^{n \times m}$ is defined as $Y(\mathbf{x}) \equiv \begin{pmatrix} P_{\mathbf{x}}^\perp \boldsymbol{\xi}_1 & \cdots & P_{\mathbf{x}}^\perp \boldsymbol{\xi}_m \\ \boldsymbol{\xi}_1 \cdot \mathbf{x} & \cdots & \boldsymbol{\xi}_m \cdot \mathbf{x} \end{pmatrix}$. Since $Y(\mathbf{x})$ is just $\boldsymbol{\xi}$ written in an orthonormal basis of \mathbb{R}^n whose first vector is \mathbf{x} , its distributions is non-degenerate. $H : \mathbb{R}^{n \times m} \rightarrow \mathbb{R}^{n-1}$ is defined as:

$$\forall 1 \leq i < n, \quad H(Y)_i \equiv \frac{1}{m} \sum_{\mu=1}^m Y_{i,\mu} \phi'(Y_{n,\mu}), \quad (Y \in \mathbb{R}^{n \times m}). \quad (44)$$

Since ϕ is C^2 , H is C^1 . This verifies (i) and (ii). We turn our attention to verifying (iii). One can write the distribution of the gradient of eq. (43) as $\text{grad} L_1(\mathbf{x}) \stackrel{d}{=} (1/m) \sum_{\mu=1}^m \phi'(y_\mu) \mathbf{z}_\mu$, in which $y_\mu \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$, $\mathbf{z}_\mu \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_{n-1})$, and all $\{y_\mu, \mathbf{z}_\nu\}$ are independent. Since the distribution of $\text{grad} L_1(\mathbf{x})$ does not depend on \mathbf{x} , it is enough to check that its density exists and is a continuous function. To do so, we will show that its characteristic function $\hat{\varphi}_{\text{grad} L_1(\mathbf{x})} \in L^1(\mathbb{R}^{n-1})$. We denote $\hat{\varphi}_a$ the characteristic function of the random variable $a \equiv \phi'(y)$, and one obtains:

$$\begin{aligned} \|\hat{\varphi}_{\text{grad} L_1(\mathbf{x})}\|_1 &= \int_{\mathbb{R}^{n-1}} dt \left| \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(0, I_{n-1})} \hat{\varphi}_a \left(\frac{\mathbf{t} \cdot \mathbf{z}}{m} \right) \right|^m = \int_{\mathbb{R}^{n-1}} dt \left| \mathbb{E}_{z \sim \mathcal{N}(0,1)} \hat{\varphi}_a \left(\frac{\|\mathbf{t}\| z}{m} \right) \right|^m, \\ &= \frac{2\pi^{\frac{n-1}{2}} m^{n-1}}{\Gamma(\frac{n-1}{2})} \int_0^\infty dq q^{n-2} |\mathbb{E}_z \hat{\varphi}_a(qz)|^m. \end{aligned}$$

Since $\alpha > 1$, if $[q \mathbb{E}_z \hat{\varphi}_a(qz)] = \mathcal{O}_{q \rightarrow \infty}(1)$ we can conclude that $\|\hat{\varphi}_{\text{grad} L_1(\mathbf{x})}\|_1 < \infty$. And:

$$q \mathbb{E}_z \hat{\varphi}_a(qz) = \int_{\mathbb{R}} \frac{dz}{\sqrt{2\pi}} e^{-\frac{z^2}{2q^2}} \hat{\varphi}_a(z) = \int_{\mathbb{R}} \frac{dz}{\sqrt{2\pi}} \mathbb{E} \left[e^{-\frac{z^2}{2q^2}} e^{iaz} \right] = \frac{1}{q} \mathbb{E} \left[e^{-\frac{q^2 a^2}{2}} \right],$$

by Fubini's theorem. Therefore $q \mathbb{E}_z \hat{\varphi}_a(qz) \rightarrow_{q \rightarrow \infty} \varphi_a(0)$ by continuity of φ_a around $a = 0$ (Definition 1), so $\|\hat{\varphi}_{\text{grad}L_1(\mathbf{x})}\|_1 < \infty$. Thus, $\text{grad}L_1(\mathbf{x})$ admits the following probability density:

$$\varphi_{\text{grad}L_1(\mathbf{x})}(\mathbf{u}) = \frac{1}{(2\pi)^{n-1}} \int_{\mathbb{R}^{n-1}} d\mathbf{t} e^{i\mathbf{u}\cdot\mathbf{t}} \left[\mathbb{E}_z \left\{ \hat{\varphi}_a \left(\frac{\|\mathbf{t}\|z}{m} \right) \right\} \right]^m, \quad (45)$$

which is a continuous function of \mathbf{u} , since $\hat{\varphi}_{\text{grad}L_1(\mathbf{x})} \in L^1(\mathbb{R}^{n-1})$. This shows (iii). In order to show (iv), we will use Proposition 6.5 of [Azais and Wschebor \[2008\]](#), that we recall here:

Lemma 9 (Azais-Wschebor). *Let $d \in \mathbb{N}^*$, and U a compact subset of \mathbb{R}^d . Consider $Z : U \rightarrow \mathbb{R}^d$ a random field, such that (a): The paths of Z are of class \mathcal{C}^2 , and (b): There exists $C > 0$ such that for all $t \in U$ and all u in a neighborhood of 0, the density $\varphi_{Z(t)}$ of Z verifies $\varphi_{Z(t)}(u) \leq C$. Then $\mathbb{P}[\exists t \in U \text{ s.t. } Z(t) = 0 \text{ and } \det Z'(t) = 0] = 0$.*

Since ϕ is assumed to be of class \mathcal{C}^3 , hypothesis (a) is verified for $Z = \text{grad}L_1$. Notice then that we can fix $C > 0$ such that $|\mathbb{E}_{z \sim \mathcal{N}(0,1)} \hat{\varphi}_a(qz)| \leq \frac{C}{1+q}$ for all $q \geq 0$. Starting from eq. (45):

$$|\varphi_{\text{grad}L_1(\mathbf{x})}(\mathbf{u})| \leq C_n \int_0^\infty dq \frac{q^{n-2}}{(1+q)^m} \leq D_n,$$

with C_n, D_n constants depending only on n , using that $m \geq n$ ($\alpha > 1$). This shows (b), so by Lemma 9, hypothesis (iv) of Theorem 8 follows. This ends the proof of Lemma 3.

A.2 Proof of Lemma 6

The proof is done in several parts, of which some are inspired by arguments of [Silverstein \[1995\]](#), [Silverstein and Bai \[1995\]](#), [Silverstein and Choi \[1995\]](#), [Bai and Silverstein \[2010\]](#).

A.2.1 Technicalities on the Hessian

We begin by a quick lemma on $\Lambda(\mathbf{y})$, defined in eq. (8).

Lemma 10 (Low-rank perturbation). *Since the distributions of \mathbf{z} and \mathbf{y} are independent, by rotation invariance we can assume that $\Lambda(\mathbf{y})$ is a diagonal matrix with elements $\Lambda_\mu(\mathbf{y})$. There exists a constant, denoted $\|D\|_\infty$, such that for all n, y , $|D(y)| \leq \|D\|_\infty$. Then we have:*

$$(i) \sup_{\mathbf{y} \in \mathbb{R}^m} \sup_{1 \leq \mu \leq m} |\Lambda_\mu(\mathbf{y})| \leq 4\|D\|_\infty.$$

(ii) *Let $\mathbf{Z} \in \mathbb{R}^{(n-1) \times m}$ be i.i.d. variables with zero mean and unit variance. We denote $\mu_D^{(n)}$ and $\mu_\Lambda^{(n)}$ the empirical eigenvalue distributions of $\frac{1}{n}\mathbf{Z}D(\mathbf{y})\mathbf{Z}^\top$ and $\frac{1}{n}\mathbf{Z}\Lambda(\mathbf{y})\mathbf{Z}^\top$ respectively. Then for all $\eta \in (0, 1)$, $\left\{ n^\eta \mathbb{E}_z \left[\mu_D^{(n)} - \mu_\Lambda^{(n)} \right] \right\} \rightarrow_{n \rightarrow \infty} 0$ weakly and uniformly in $\mathbf{y} \in \mathbb{R}^m$.*

We have some control of the boundedness of the Hessian:

Lemma 11. *Denote $\rho_n(\mathbf{y})$ the spectral radius of $H_n^\Lambda(\mathbf{y})$. There exists $C > 0$ such that:*

$$(i) \text{ With probability 1, } \limsup_{n \rightarrow \infty} \sup_{\mathbf{y} \in \mathbb{R}^m} \rho_n(\mathbf{y}) < C.$$

(ii) *The support of $\mu_{\alpha, \phi}[v_y^m]$ is included in $(-C, C)$ uniformly over \mathbf{y} and n .*

(iii) *For all $\mathbf{y} \in \mathbb{R}^m$, $\mu_{\alpha, \phi}[v_y^m]$ has a well-defined and continuous density outside $x = 0$.*

Points (ii) and (iii) of Lemma 11 are consequences of Theorem 1.1 of [Silverstein and Choi \[1995\]](#), while item (i) follows from the boundedness of $\Lambda(\mathbf{y})$ by Lemma 10, and the one of $x\phi'(x)$. We now prove Lemma 10.

Proof. Recall that $|D(y)| = (n/m)|\phi''(y)|$. Since $m/n \rightarrow \alpha > 1$ and ϕ'' is bounded, $|D(y)|$ is bounded (uniformly over n, \mathbf{y}) by a constant that we denote $\|D\|_\infty$. Note that $\sup_{1 \leq \mu \leq m} |\Lambda_\mu(\mathbf{y})| = \sup_{\|\mathbf{u}\|=1} \mathbf{u}^\top \Lambda(\mathbf{y}) \mathbf{u}$. Using eq. (8) and denoting $\mathbf{v}(\mathbf{y}) \equiv \phi'(\mathbf{y})/|\phi'(\mathbf{y})|$, we reach

$$\begin{aligned} \sup_{\|\mathbf{u}\|=1} \mathbf{u}^\top \Lambda(\mathbf{y}) \mathbf{u} &\leq \|D\|_\infty + \sup_{\|\mathbf{u}\|=1} [|\mathbf{v}^\top D \mathbf{v}|(\mathbf{u}^\top \mathbf{v})^2 + 2(\mathbf{u}^\top \mathbf{v})|\mathbf{v}^\top D \mathbf{u}|], \\ &\leq 2\|D\|_\infty + 2 \sup_{\|\mathbf{u}\|=1} [(\mathbf{u}^\top \mathbf{v})|\mathbf{v}^\top D \mathbf{u}|] \leq 4\|D\|_\infty, \end{aligned}$$

in which we used the uniform boundedness of $|D(y_\mu)|$, and the Cauchy-Schwarz inequality. This proves (i). We note that $\frac{1}{n} \mathbf{z} \Lambda(\mathbf{y}) \mathbf{z}^\top$ and $\frac{1}{n} \mathbf{z} D(\mathbf{y}) \mathbf{z}^\top$ differ by a rank-2 matrix. (ii) is thus an immediate application of the following lemma (from a course of C. Bordenave):

Lemma 12. *Let $n \geq 1$, and A, B two symmetric matrices of size n , such that the rank of $A - B$ is r . Denote F_A (resp. F_B) the c.d.f. of the empirical spectral distribution of A (resp. B). Then*

$$\sup_{t \in \mathbb{R}} |F_A(t) - F_B(t)| \leq \frac{r}{n}.$$

This ends the proof of Lemma 10. □

Lemma 12. We note $\lambda_1(A) \geq \dots \geq \lambda_n(A)$ the eigenvalues of A (and similarly for B). Recall weak Weyl's interlacing inequalities [Weyl, 1912]: for every $1 \leq i \leq n$, $\lambda_{i+r}(A) \leq \lambda_i(B) \leq \lambda_{i-r}(A)$ (we use the convention $\lambda_{1-i} = +\infty$ and $\lambda_{n+i} = -\infty$ for $i \geq 1$). Let $t \in \mathbb{R}$, and i, j be the smallest indices such that $\lambda_i(A) \leq t$ and $\lambda_j(B) < t$. By the interlacing inequalities, $|i - j| \leq r$. And $n|F_A(t) - F_B(t)| = |(n+1-i) - (n+1-j)| \leq r$. □

A.2.2 The cut-off and the logarithmic potential

For any $\epsilon > 0$, define $\ln_\epsilon : x \in \mathbb{R}_+^* \mapsto \ln(\max(x, \epsilon))$, then $x \mapsto \ln_\epsilon |x|$ is a ϵ^{-1} -Lipschitz function on \mathbb{R} . Let $\delta \in (0, 1)$. In this section, we show that a cut-off $\epsilon_n = n^{-\delta}$ on the smallest eigenvalues does not perturb the logarithmic potential at the thermodynamical scale. We rely on the following result, proven in an ongoing work of Ben Arous, Bourgade and McKenna. For any $\delta \in (0, 1)$, there exists $\eta > 0$ such that for all $t > 0$:

$$\left\{ \lim_{n \rightarrow \infty} \frac{1}{n^{1+\eta}} \ln \mathbb{P} \left[\left| \frac{1}{n} \sum_{\lambda \in \text{Sp}(H_n^\Lambda(\mathbf{y}))} \ln |\lambda| \mathbb{1} \{ |\lambda| \leq n^{-\delta} \} \right| \geq t \right] = -\infty, \quad (46a)$$

$$\left\{ \lim_{n \rightarrow \infty} \frac{1}{n^{1+\eta}} \ln \mathbb{P} \left[\int_{|x - t_\phi(\nu_{\mathbf{y}}^m)| \leq n^{-\delta}} \mu_{\alpha, \phi}[\nu_{\mathbf{y}}^m](dx) \ln |x - t_\phi(\nu_{\mathbf{y}}^m)| \leq -t \right] = -\infty. \quad (46b)$$

Physically, this makes explicit that, with large probability, there should not be enough eigenvalues of $H_n^\Lambda(\mathbf{y})$ around zero so that they contribute macroscopically to the logarithmic potential. This is a consequence the natural fluctuations and repulsion of the eigenvalues of $H_n^N(\mathbf{y})$. Denote $\{\lambda_i\}_{i=1}^{n-1}$ the (sorted) eigenvalues of $H_n^\Lambda(\mathbf{y})$. We can now state:

Lemma 13. *There exists $\eta > 0$ such that for all $K > 0$:*

$$\lim_{n \rightarrow \infty} \frac{1}{n^{1+\eta}} \ln \mathbb{P} \left[\left| \frac{1}{n} \ln \mathbb{E} |\det H_n^\Lambda(\mathbf{y})| - \frac{1}{n} \ln \mathbb{E} e^{\sum_{i=1}^{n-1} \ln_{\epsilon_n} |\lambda_i|} \right| \geq K \right] = -\infty. \quad (47)$$

Proof. Since this lemma is not used in the proofs of the following Lemma 17 and Proposition 18, we will refer to them. We consider η given by eq. (46a). Let $t > 0$. We denote $A_t^{(n)}$ the event

$$A_t^{(n)} \equiv \left\{ \left| \frac{1}{n} \sum_{i=1}^{n-1} \ln |\lambda_i| \mathbb{1} \{ |\lambda_i| \leq n^{-\delta} \} \right| \geq t \right\}. \quad (48)$$

We have for all \mathbf{y} and $t > 0$ ($\bar{A}_t^{(n)}$ being the complementary event to $A_t^{(n)}$):

$$\frac{1}{n} \ln \mathbb{E}_{\mathbf{z}} e^{\sum_{i=1}^{n-1} \ln |\lambda_i|} \geq \frac{1}{n} \ln \mathbb{E}_{\mathbf{z}} \left[e^{\sum_{i=1}^{n-1} \ln |\lambda_i|} \mathbb{1} \left[\bar{A}_t^{(n)} \right] \right] \geq -t + \frac{1}{n} \ln \mathbb{E}_{\mathbf{z}} \left[e^{\sum_{i=1}^{n-1} \ln_{\epsilon_n} |\lambda_i|} \mathbb{1} \left[\bar{A}_t^{(n)} \right] \right].$$

So that (using $\ln_{\epsilon_n}(x) \geq \ln(x)$ for all $x > 0$):

$$0 \leq \frac{1}{n} \ln \mathbb{E} e^{\sum_{i=1}^{n-1} \ln_{\epsilon_n} |\lambda_i|} - \frac{1}{n} \ln \mathbb{E} |\det H_n^\Lambda(\mathbf{y})| \leq t - \frac{1}{n} \ln \left[1 - \frac{\mathbb{E}_{\mathbf{z}} \left[e^{\sum_{i=1}^{n-1} \ln_{\epsilon_n} |\lambda_i|} \mathbb{1} \left[A_t^{(n)} \right] \right]}{\mathbb{E}_{\mathbf{z}} \left[e^{\sum_{i=1}^{n-1} \ln_{\epsilon_n} |\lambda_i|} \right]} \right].$$

We know $\ln_{\epsilon_n} |x| \geq -\delta \ln(n)$ and moreover, by Lemma 17 and Proposition 18, for every $\gamma > 0$:

$$\limsup_{n \rightarrow \infty} \sup_{\mathbf{y} \in \mathbb{R}^m} \frac{1}{n} \ln \mathbb{E}_{\mathbf{z}} \left[e^{\gamma \sum_{i=1}^{n-1} \ln_{\epsilon_n} |\lambda_i|} \right] = \limsup_{n \rightarrow \infty} \sup_{\mathbf{y} \in \mathbb{R}^m} \frac{\gamma}{n} \mathbb{E}_{\mathbf{z}} \sum_{i=1}^{n-1} \ln_{\epsilon_n} |\lambda_i| < +\infty,$$

in which the last inequality is a consequence of (i) of Lemma 11. Fixing $\gamma > 1$ and using Hölder's inequality, there exists therefore $C > 0$ such that for all $K > 0$ and $t \in (0, K)$:

$$\begin{aligned} & \limsup_{n \rightarrow \infty} \frac{1}{n^{1+\eta}} \ln \mathbb{P}_{\mathbf{y}} \left[\left| \frac{1}{n} \ln \mathbb{E}_{\mathbf{z}} |\det H_n^\Lambda(\mathbf{y})| - \frac{1}{n} \ln \mathbb{E}_{\mathbf{z}} e^{\sum_{i=1}^{n-1} \ln_{\epsilon_n} |\lambda_i|} \right| \geq K \right] \\ & \leq \limsup_{n \rightarrow \infty} \frac{1}{n^{1+\eta}} \ln \mathbb{P}_{\mathbf{y}} \left[\mathbb{P}_{\mathbf{z}} \left[A_t^{(n)} \right]^{1/\gamma} \geq e^{-n(\delta \ln(n)+C)} \left[1 - e^{n(t-K)} \right] \right], \\ & \stackrel{(a)}{\leq} \limsup_{n \rightarrow \infty} \frac{1}{n^{1+\eta}} \ln \left\{ \frac{\mathbb{P}[A_t^{(n)}]}{e^{-\gamma n(\delta \ln(n)+C)} \left[1 - e^{n(t-K)} \right]^\gamma} \right\} \stackrel{(b)}{=} -\infty, \end{aligned}$$

in which we used the Markov inequality in (a) and eq. (46a) in (b). \square

In the following, for simplicity we will often abusively denote $\ln_{\epsilon_n} |\det H_n^\Lambda(\mathbf{y})| \equiv e^{\sum_{i=1}^{n-1} \ln_{\epsilon_n} |\lambda_i|}$ and $\ln_{\epsilon_n} \mathbb{E} |\det H_n^\Lambda(\mathbf{y})| \equiv \ln \mathbb{E} e^{\sum_{i=1}^{n-1} \ln_{\epsilon_n} |\lambda_i|}$.

A.2.3 Truncation and renormalization of the variables $\{z_{i\mu}\}$

Truncation of the distribution of the matrix elements of \mathbf{z} will be needed in the following section A.2.4, to circumvent some difficulties arising in the concentration. The idea is to truncate the variables $\{z_{i\mu}\}$, and to show that this does not perturb the result. Let $\gamma \in (0, 1/2)$ such that $\gamma + 2\delta < 1$ and define a truncated and normalized version of the \mathbf{z} variables:

$$\tilde{z}_{i\mu} \equiv \left(z_{i\mu} \mathbb{1} \left[|z_{i\mu}| \leq n^{\frac{\gamma}{2}} \right] \right) / \left(\mathbb{E} \left\{ z_{i\mu}^2 \mathbb{1} \left[|z_{i\mu}| \leq n^{\frac{\gamma}{2}} \right] \right\} \right)^{1/2}. \quad (49)$$

Note that $\{\tilde{z}_{i\mu}\}$ are still i.i.d. variables, with zero mean and unit variance. We denote the corresponding version of the Hessian:

$$\tilde{H}_n^\Lambda(\mathbf{y}) \equiv \frac{1}{n} \tilde{\mathbf{z}} \Lambda(\mathbf{y}) \tilde{\mathbf{z}}^\top - \left[\frac{1}{m} \sum_{\mu=1}^m y_\mu \phi'(y_\mu) \right] \mathbf{I}_{n-1}. \quad (50)$$

The following shows that truncation does not impact the logarithmic potential at exponential scale:

Proposition 14. *There exists $C_1, C_2 > 0$ such that for all Lipschitz function $f : \mathbb{R} \rightarrow \mathbb{R}$ with Lipschitz constant $\|f\|_{\mathcal{L}}$, and uniformly over \mathbf{y} :*

$$\mathbb{E}_{\mathbf{z}} \left| \text{Tr} f(H_n^\Lambda(\mathbf{y})) - \text{Tr} f(\tilde{H}_n^\Lambda(\mathbf{y})) \right| \leq C_1 \|f\|_{\mathcal{L}} n^{3/2} e^{-C_2 n^\gamma}.$$

Proof. Denote $\lambda_1 \leq \dots \leq \lambda_{n-1}$ the (sorted) eigenvalues of $H_n^\Lambda(\mathbf{y})$ and $\tilde{\lambda}_1 \leq \dots \leq \tilde{\lambda}_{n-1}$ the ones of $\tilde{H}_n^\Lambda(\mathbf{y})$. One has:

$$\begin{aligned} \mathbb{E}_{\mathbf{z}} \left| \text{Tr} f(H_n^\Lambda(\mathbf{y})) - \text{Tr} f(\tilde{H}_n^\Lambda(\mathbf{y})) \right| &\stackrel{(a)}{\leq} \|f\|_{\mathcal{L}} \sum_{i=1}^{n-1} \mathbb{E} \left| \lambda_i - \tilde{\lambda}_i \right|, \\ &\stackrel{(b)}{\leq} \|f\|_{\mathcal{L}} n^{1/2} \mathbb{E} \left[\sum_{i=1}^{n-1} (\lambda_i - \tilde{\lambda}_i)^2 \right]^{1/2}, \\ &\stackrel{(c)}{\leq} \|f\|_{\mathcal{L}} n^{-1/2} \mathbb{E} \|\mathbf{z}\Lambda(\mathbf{y})\mathbf{z}^\top - \tilde{\mathbf{z}}\Lambda(\mathbf{y})\tilde{\mathbf{z}}^\top\|_2, \end{aligned}$$

in which (a) uses the Lipschitzity of f , (b) Cauchy-Schwarz inequality and (c) uses Hoffman-Wielandt inequality [Hoffman and Wielandt, 2003] that we recall here:

Lemma 15 (Hoffman-Wielandt inequality for L_2 norm). *Let $k \in \mathbb{N}^*$, and $A, B \in \mathbb{R}^{k \times k}$ be two symmetric matrices with respective eigenvalues $\lambda_1(A) \leq \dots \leq \lambda_k(A)$ and $\lambda_1(B) \leq \dots \leq \lambda_k(B)$. Then $\sum_{i=1}^k (\lambda_i(A) - \lambda_i(B))^2 \leq \|A - B\|_2^2$.*

Coming back to our original problem:

$$\begin{aligned} \mathbb{E}_{\mathbf{z}} \left| \text{Tr} f(H_n^\Lambda(\mathbf{y})) - \text{Tr} f(\tilde{H}_n^\Lambda(\mathbf{y})) \right| &\leq \frac{\|f\|_{\mathcal{L}}}{\sqrt{n}} \mathbb{E} \left[\sum_{1 \leq i, j < n} \left(\sum_{\mu=1}^m \Lambda_\mu(\mathbf{y}) (z_{i\mu} z_{j\mu} - \tilde{z}_{i\mu} \tilde{z}_{j\mu}) \right)^2 \right]^{1/2}, \\ &\stackrel{(a)}{\leq} 4 \|f\|_{\mathcal{L}} \|D\|_\infty \left(\frac{m}{n} \right)^{1/2} \mathbb{E} \left[\sum_{1 \leq i, j < n} \sum_{\mu=1}^m (z_{i\mu} z_{j\mu} - \tilde{z}_{i\mu} \tilde{z}_{j\mu})^2 \right]^{1/2}, \\ &\stackrel{(b)}{\leq} 4 \|f\|_{\mathcal{L}} \|D\|_\infty \left(\frac{m^2}{n} \right)^{1/2} \left[\mathbb{E} \sum_{1 \leq i, j < n} (z_{i1} z_{j1} - \tilde{z}_{i1} \tilde{z}_{j1})^2 \right]^{1/2}, \\ &\stackrel{(c)}{\leq} 4 \|f\|_{\mathcal{L}} \|D\|_\infty (m^2 n)^{1/2} \left[n \mathbb{E} (z_{11} z_{21} - \tilde{z}_{11} \tilde{z}_{21})^2 + n \mathbb{E} (z_{11}^2 - \tilde{z}_{11}^2)^2 \right]^{1/2}, \end{aligned} \quad (51)$$

in which we used Cauchy-Schwarz in (a), the concavity of $x \mapsto x^{1/2}$ in (b), along with the fact that all the variables $z_{i\mu}$ are independent and identically distributed in (c). We can now bound each one of the two remaining terms. To do so, we use a simple technical lemma that we will use for each of these terms, its proof is deferred to Section A.6.

Lemma 16. *Let $z_{11} \sim \mathcal{N}(0, 1)$. There exists strictly positive constants $(K_i)_{i=1}^4$ such that:*

$$\begin{cases} \mathbb{E} \left[\mathbb{E} \left(z_{11}^2 \mathbf{1}_{|z_{11}| \leq n^{\gamma/2}} \right) - \mathbf{1}_{|z_{11}| \leq n^{\gamma/2}} \right]^4 \leq K_1 e^{-K_2 n^\gamma}, \\ \mathbb{E} \left[\left(\mathbb{E} \left(z_{11}^2 \mathbf{1}_{|z_{11}| \leq n^{\gamma/2}} \right) \right)^{1/2} - \mathbf{1}_{|z_{11}| \leq n^{\gamma/2}} \right]^8 \leq K_3 e^{-K_4 n^\gamma}. \end{cases}$$

Let us bound the two terms of eq. (51). We start with the second one:

$$\mathbb{E} (z_{11}^2 - \tilde{z}_{11}^2)^2 = \mathbb{E} \left[z_{11}^4 \left(\mathbb{E} \left(z_{11}^2 \mathbf{1}_{|z_{11}| \leq n^{\gamma/2}} \right) - \mathbf{1}_{|z_{11}| \leq n^{\gamma/2}} \right)^2 \right] / \left(\mathbb{E} z_{11}^2 \mathbf{1}_{|z_{11}| \leq n^{\gamma/2}} \right)^2.$$

This last equation, combined with Cauchy-Schwarz inequality, the first part of Lemma 16, and the fact that $\mathbb{E} \left[z_{11}^2 \mathbf{1}_{|z_{11}| \leq n^{\gamma/2}} \right] \geq 1/2$ for large enough n yields constants $C_1, C_2 > 0$ such that

$$\mathbb{E} (z_{11}^2 - \tilde{z}_{11}^2)^2 \leq C_1 e^{-C_2 n^\gamma}. \quad (53)$$

Similarly, one can show:

$$\begin{aligned} & \mathbb{E} (z_{11}z_{21} - \tilde{z}_{11}\tilde{z}_{21})^2 \\ & \leq 4\mathbb{E} \left[z_{11}^2 z_{21}^2 \left(\sqrt{\mathbb{E} [z_{11}^2 \mathbb{1}_{|z_{11}| \geq n^{\gamma/2}}]} \sqrt{\mathbb{E} [z_{12}^2 \mathbb{1}_{|z_{12}| \geq n^{\gamma/2}}]} - \mathbb{1}_{\substack{|z_{11}| \geq n^{\gamma/2} \\ |z_{12}| \geq n^{\gamma/2}}} \right)^2 \right], \\ & \leq 4 [\mathbb{E} z_{11}^4 z_{21}^4]^{1/2} \left[\mathbb{E} \left(\sqrt{\mathbb{E} [z_{11}^2 \mathbb{1}_{|z_{11}| \geq n^{\gamma/2}}]} \sqrt{\mathbb{E} [z_{12}^2 \mathbb{1}_{|z_{12}| \geq n^{\gamma/2}}]} - \mathbb{1}_{\substack{|z_{11}| \geq n^{\gamma/2} \\ |z_{12}| \geq n^{\gamma/2}}} \right)^4 \right]^{1/2}. \end{aligned}$$

For all positive x, y, x', y' , one has $(xy - x'y')^2 \leq 2y^2(x - x')^2 + 2x'^2(y - y')^2$. Combined with Cauchy-Schwarz inequality, and since all $\{z_{i\mu}\}$ are identically distributed, this gives a constant $K > 0$ such that :

$$\mathbb{E} (z_{11}z_{21} - \tilde{z}_{11}\tilde{z}_{21})^2 \leq K \mathbb{E} \left[\left(\mathbb{E} (z_{11}^2 \mathbb{1}_{|z_{11}| \leq n^{\gamma/2}}) \right)^{1/2} - \mathbb{1}_{|z_{11}| \leq n^{\gamma/2}} \right]^8.$$

The second part of Lemma 16 finally yields $C_3, C_4 > 0$ such that:

$$\mathbb{E} (z_{11}z_{21} - \tilde{z}_{11}\tilde{z}_{21})^2 \leq C_3 e^{-C_4 n^\gamma}. \quad (54)$$

Recall that $\lim_{n \rightarrow \infty} \frac{m}{n} = \alpha > 1$. Combining eqs. (51),(53),(54) finishes the proof of Proposition 14. \square

Recall that $\epsilon_n = n^{-\delta}$ with $\delta \in (0, 1)$. We can apply Proposition 14 to $f(x) = \ln_{\epsilon_n} |x|$, which has a Lipschitz constant $\|f\|_{\mathcal{L}} = n^\delta$. In particular, this implies the following lemma:

Lemma 17 (Truncation and renormalization).

$$\lim_{n \rightarrow \infty} \left\{ \sup_{\mathbf{y} \in \mathbb{R}^m} \mathbb{E}_{\mathbf{z}} \left| \frac{1}{n} \ln_{\epsilon_n} |\det H_n^\Lambda(\mathbf{y})| - \frac{1}{n} \ln_{\epsilon_n} |\det \tilde{H}_n^\Lambda(\mathbf{y})| \right| \right\} = 0.$$

A.2.4 Concentration of the logarithmic potential with a cut-off

We show here that discarding the eigenvalues of the Hessian¹ that are close to 0 using a cutoff $\epsilon_n \equiv n^{-\delta}$ (recall that $\gamma \in (0, 1/2)$ and $\gamma + 2\delta < 1$), we have concentration of the logarithmic potential. Let us thus fix a couple (γ, δ) verifying the above conditions. Once we are in this setting, it is actually an easy consequence of the classical Lipschitz concentration for random variables:

Proposition 18 (Concentration of the logarithmic potential). *There exists constants $K_1, K_2 > 0$ (independent of \mathbf{y}) such that almost surely over \mathbf{y} and for all n :*

$$\forall t > 0, \mathbb{P}_{\mathbf{z}} \left[\frac{1}{n} \left| \ln_{\epsilon_n} |\det \tilde{H}_n^\Lambda(\mathbf{y})| - \mathbb{E}_{\mathbf{z}} \ln_{\epsilon_n} |\det \tilde{H}_n^\Lambda(\mathbf{y})| \right| \geq t \right] \leq K_1 \exp \left\{ -K_2 n^{2-\gamma-2\delta} t^2 \right\}.$$

Proof. Using traditional Lipschitz concentration bound (cf for instance Anderson et al. [2010]), we will end the proof if we can show that, \mathbf{y} being fixed, the function $G(\mathbf{z}) \equiv (1/n) \ln_{\epsilon_n} |\det \tilde{H}_n^\Lambda(\mathbf{y})|$ is a Lipschitz function of $\mathbf{z} \in \mathbb{R}^{(n-1) \times m}$, with Lipschitz norm bounded as $\|G\|_{\mathcal{L}} \leq C n^{\frac{\gamma+2\delta}{2}}$, for a constant $C > 0$. We will do it by bounding $\|\nabla_{\mathbf{z}} G\|_{\infty}$. Let $f_n(x) \equiv \ln_{\epsilon_n} |x|$ for $x \in \mathbb{R}$. We have:

$$\sum_{i=1}^{n-1} \sum_{\mu=1}^m \left(\frac{\partial G(\mathbf{z})}{\partial z_{i\mu}} \right)^2 = \frac{1}{n^4} \sum_{i=1}^{n-1} \sum_{\mu=1}^m \left[\text{Tr} \left\{ f'_n \left(\frac{1}{n} \mathbf{z} \Lambda(\mathbf{y}) \mathbf{z}^\top \right) \Delta_{i\mu} \right\} \right]^2, \quad (55)$$

¹With truncated $\{z_{i\mu}\}$, see section A.2.3.

in which $\Delta_{i\mu} \in \mathbb{R}^{(n-1) \times (n-1)}$ with $(\Delta_{i\mu})_{jk} \equiv \Lambda_\mu(\mathbf{y}) (\delta_{ij} z_{k\mu} + \delta_{ik} z_{j\mu})$. So one shows easily:

$$\sum_{i=1}^{n-1} \sum_{\mu=1}^m \left(\frac{\partial G(\mathbf{z})}{\partial z_{i\mu}} \right)^2 = \frac{4}{n^3} \text{Tr} \left[\left(f'_n \left(\frac{1}{n} \mathbf{z} \Lambda(\mathbf{y}) \mathbf{z}^\top \right) \right)^2 \left(\frac{1}{n} \mathbf{z} \Lambda(\mathbf{y})^2 \mathbf{z}^\top \right) \right]. \quad (56)$$

By Lemma 15, if A and B are positive matrices with sorted eigenvalues $\{\lambda_i(A), \lambda_i(B)\}$, one has $\text{Tr}[AB] \leq \sum_i \lambda_i(A) \lambda_i(B)$. Using this identity in eq. (56) and the n^δ -Lipschitzity of f_n yields:

$$\sum_{i=1}^{n-1} \sum_{\mu=1}^m \left(\frac{\partial G(\mathbf{z})}{\partial z_{i\mu}} \right)^2 \leq \frac{4n^{2\delta}}{n^4} \text{Tr} [\mathbf{z} \Lambda(\mathbf{y})^2 \mathbf{z}^\top] \leq \frac{4^3 n^{2\delta} \|D\|_\infty^2}{n^4} \sum_{\mu=1}^m \sum_{i=1}^{n-1} z_{i\mu}^2,$$

in which we used Lemma 10. Recall (cf section A.2.3) that we truncated the \mathbf{z} variables so that we impose $|z_{i\mu}| \leq n^{\frac{\gamma}{2}}$, which yields:

$$\sum_{i=1}^{n-1} \sum_{\mu=1}^m \left(\frac{\partial G(\mathbf{z})}{\partial z_{i\mu}} \right)^2 \leq \frac{4^3 n^{\gamma+2\delta} \|D\|_\infty^2}{n^2}. \quad (57)$$

Recall the Lipschitz concentration of independent variables with laws satisfying the logarithmic Sobolev inequality with a uniform constant c (see for instance Anderson et al. [2010]):

Lemma 19 (Herbst). *Let $n \in \mathbb{N}^*$ and P be a probability distribution on \mathbb{R}^n satisfying the Logarithmic Sobolev Inequality with constant $c > 0$. Let G be a Lipschitz function on \mathbb{R}^n with Lipschitz constant $\|G\|_{\mathcal{L}}$. Then for all $t > 0$, $\mathbb{P}[|G - \mathbb{E}G| \geq t] \leq 2 \exp[-t^2/(2c \|G\|_{\mathcal{L}}^2)]$.*

It is easy to check that the truncated law of the $\{z_{i\mu}\}$ satisfies the Logarithmic Sobolev Inequality with constant $c = 1$. Applying Lemma 19 alongside eq. (57) finishes the proof. \square

A.2.5 The logarithmic potential of the asymptotic measure

In this part, we relate the expected logarithmic potential to the logarithmic potential of the measure $\mu_{\alpha,\phi}[\nu_{\mathbf{y}}^m]$, cf Theorem 1.

Proposition 20 (Concentration on $\kappa_{\alpha,\phi}$). *There exists $\eta > 0$ such that for all $t > 0$:*

$$\lim_{n \rightarrow \infty} \frac{1}{n^{1+\eta}} \ln \mathbb{P} \left[\left| \mathbb{E}_{\mathbf{z}} \frac{1}{n} \ln_{\epsilon_n} \left| \det \tilde{H}_n^\Lambda(\mathbf{y}) \right| - \kappa_{\alpha,\phi}(\nu_{\mathbf{y}}^m, t_\phi(\nu_{\mathbf{y}}^m)) \right| \geq t \right] = -\infty. \quad (58)$$

Proof. The proof goes in two parts. First, we show that there exists $\eta_1 > 0$ such that¹:

$$\lim_{n \rightarrow \infty} \left[n^{\eta_1} \sup_{\mathbf{y} \in \mathbb{R}^m} \left| \mathbb{E}_{\mathbf{z}} \frac{1}{n} \ln_{\epsilon_n} \left| \det \tilde{H}_n^\Lambda(\mathbf{y}) \right| - \int_{\mathbb{R}} \ln_{\epsilon_n} |x - t_\phi(\nu_{\mathbf{y}}^m)| \mu_{\alpha,\phi}[\nu_{\mathbf{y}}^m](dx) \right| \right] = 0. \quad (59)$$

We will then conclude by showing that there exists $\eta_2 > 0$ such that for all $t > 0$:

$$\lim_{n \rightarrow \infty} \frac{1}{n^{1+\eta_2}} \mathbb{P} \left[\left| \int_{\mathbb{R}} \ln_{\epsilon_n} |x - t_\phi(\nu_{\mathbf{y}}^m)| \mu_{\alpha,\phi}[\nu_{\mathbf{y}}^m](dx) - \kappa_{\alpha,\phi}(\nu_{\mathbf{y}}^m, t_\phi(\nu_{\mathbf{y}}^m)) \right| \geq t \right] = -\infty. \quad (60)$$

We begin by eq. (60). We take η_2 given by eq. (46b). We have

$$\begin{aligned} & \int_{\mathbb{R}} \ln_{\epsilon_n} |x - t_\phi(\nu_{\mathbf{y}}^m)| \mu_{\alpha,\phi}[\nu_{\mathbf{y}}^m](dx) - \kappa_{\alpha,\phi}(\nu_{\mathbf{y}}^m, t_\phi(\nu_{\mathbf{y}}^m)) \\ &= -\delta \ln(n) \mu_{\alpha,\phi}[\nu_{\mathbf{y}}^m](t_\phi(\nu_{\mathbf{y}}^m) - \epsilon_n, t_\phi(\nu_{\mathbf{y}}^m) + \epsilon_n) - \int_{t_\phi - \epsilon_n}^{t_\phi + \epsilon_n} \ln |x - t_\phi(\nu_{\mathbf{y}}^m)| \mu_{\alpha,\phi}[\nu_{\mathbf{y}}^m](dx). \end{aligned}$$

¹Note that this result is uniform over \mathbf{y} , and thus stronger than what is needed to show Proposition 20.

Since $\ln_\epsilon(x) \geq \ln(x)$, we reach from this:

$$\begin{aligned} & \limsup_{n \rightarrow \infty} \frac{1}{n^{1+\eta_2}} \mathbb{P} \left[\left| \int_{\mathbb{R}} \ln_{\epsilon_n} |x - t_\phi(\nu_{\mathbf{y}}^m)| \mu_{\alpha, \phi}[\nu_{\mathbf{y}}^m](dx) - \kappa_{\alpha, \phi}(\nu_{\mathbf{y}}^m, t_\phi(\nu_{\mathbf{y}}^m)) \right| \geq t \right], \\ & \leq \limsup_{n \rightarrow \infty} \frac{1}{n^{1+\eta_2}} \mathbb{P} \left[\int_{t_\phi - \epsilon_n}^{t_\phi + \epsilon_n} \ln |x - t_\phi(\nu_{\mathbf{y}}^m)| \mu_{\alpha, \phi}[\nu_{\mathbf{y}}^m](dx) \leq -t \right], \end{aligned}$$

and using eq. (46b), we reach eq. (60). Let us show eq. (59). Its proof is based on the following lemma, a consequence of the analysis of [Silverstein and Bai \[1995\]](#), [Bai and Silverstein \[2010\]](#):

Lemma 21. *Denote $g_n(z)$ the Stieltjes transform of $\frac{1}{n} \tilde{\mathbf{z}} \Lambda(\mathbf{y}) \tilde{\mathbf{z}}^\top$, and $g_{\alpha, \phi}[\nu_{\mathbf{y}}^m](z)$ the one of $\mu_{\alpha, \phi}[\nu_{\mathbf{y}}^m]$, for $z \in \mathbb{C}_+$. Then there exists $\eta \in (0, 1)$ such that for all $z \in \mathbb{C}_+$:*

$$\lim_{n \rightarrow \infty} \left\{ \sup_{\mathbf{y} \in \mathbb{R}^m} n^\eta |\mathbb{E}_{\mathbf{z}}(g_n(z)) - g_{\alpha, \phi}[\nu_{\mathbf{y}}^m](z)| \right\} = 0. \quad (61)$$

The proof of this lemma is postponed to Appendix A.5. Let us fix η given by this lemma. As stated for instance in Theorem 2.4.4 of [Anderson et al. \[2010\]](#), a consequence of the Stieltjes inversion theorem is that for every Borel set $E \subseteq \mathbb{R}$: $\sup_{\mathbf{y} \in \mathbb{R}^m} [n^\eta |\mathbb{E} \mu_n(E) - \mu_{\alpha, \phi}[\nu_{\mathbf{y}}^m](E)|] \rightarrow 0$, in which μ_n is the empirical spectral distribution of $\frac{1}{n} \tilde{\mathbf{z}} \Lambda(\mathbf{y}) \tilde{\mathbf{z}}^\top$. Fix $\eta_1 < \eta$. We have, uniformly over \mathbf{y} :

$$\begin{aligned} & n^{\eta_1} \left| \mathbb{E}_{\mathbf{z}} \frac{1}{n} \ln_{\epsilon_n} \left| \det \tilde{H}_n^\Lambda(\mathbf{y}) \right| - \int_{\mathbb{R}} \ln_{\epsilon_n} |x - t_\phi(\nu_{\mathbf{y}}^m)| \mu_{\alpha, \phi}[\nu_{\mathbf{y}}^m](dx) \right| \\ & \leq n^{\eta_1} \int_{|x - t_\phi(\nu_{\mathbf{y}}^m)| > 1} \ln |x - t_\phi(\nu_{\mathbf{y}}^m)| [\mathbb{E} \mu_n - \mu_{\alpha, \phi}[\nu_{\mathbf{y}}^m]](dx) \\ & \quad + \delta \ln(n) n^{\eta_1} \int_{|x - t_\phi(\nu_{\mathbf{y}}^m)| < 1} [\mathbb{E} \mu_n - \mu_{\alpha, \phi}[\nu_{\mathbf{y}}^m]](dx). \end{aligned} \quad (62)$$

Let us fix $C > 0$ given by item (i) of Lemma 11. Moreover, we can also bound $t_\phi(\nu_{\mathbf{y}}^m)$ by $\|x\phi'\|_\infty$. This gives that for n large enough the quantity of eq. (62) is bounded (uniformly over \mathbf{y}) by:

$$n^{\eta_1} [\ln(C + \|x\phi'\|_\infty) + \delta \ln(n)] \{ [\mathbb{E} \mu_n - \mu_{\alpha, \phi}[\nu_{\mathbf{y}}^m]](-C, C) \}.$$

Using the remark below Lemma 21 and since $\eta_1 < \eta$, this shows eq. (59). \square

A.2.6 Conclusion of the proof

Let us conclude from the rest of Section A.2. We fix $\delta, \gamma > 0$ such that $\gamma + 2\delta < 1$ and $\gamma < 1/2$. First, note that Lemma 17 and Proposition 18, as they are uniform results on \mathbf{y} , are much stronger than what we required. In particular, they imply that there exists $\eta > 0$ such that for all $t > 0$:

$$\left\{ \lim_{n \rightarrow \infty} \frac{1}{n^{1+\eta}} \ln \mathbb{P} \left[\left| \frac{1}{n} \ln_{\epsilon_n} \mathbb{E}_{\mathbf{z}} \left| \det H_n^\Lambda(\mathbf{y}) \right| - \frac{1}{n} \ln_{\epsilon_n} \mathbb{E}_{\mathbf{z}} \left| \det \tilde{H}_n^\Lambda(\mathbf{y}) \right| \right| \geq t \right] = -\infty, \quad (63a) \right.$$

$$\left. \left\{ \lim_{n \rightarrow \infty} \frac{1}{n^{1+\eta}} \ln \mathbb{P} \left[\left| \frac{1}{n} \ln_{\epsilon_n} \mathbb{E}_{\mathbf{z}} \left| \det \tilde{H}_n^\Lambda(\mathbf{y}) \right| - \frac{1}{n} \mathbb{E}_{\mathbf{z}} \ln_{\epsilon_n} \left| \det \tilde{H}_n^\Lambda(\mathbf{y}) \right| \right| \geq t \right] = -\infty. \quad (63b) \right.$$

From this and Lemma 13 and Proposition 20, we reach the conclusion of Lemma 6.

A.3 Proof of Lemma 7

Proof. Let $\gamma \in (1, \alpha)$. We fix $C > 0$ given by Lemma 11. Then for all $y \in \mathbb{R}^m$:

$$\begin{aligned} \kappa_{\alpha, \phi}(\nu_{\mathbf{y}}^m, t_\phi(\nu_{\mathbf{y}}^m)) & \leq \int \mu_{\alpha, \phi}[\nu_{\mathbf{y}}^m](dx) \ln(1 + |x|) + \ln(1 + |t_\phi(\nu_{\mathbf{y}}^m)|), \\ & \leq \ln(1 + \|x\phi'(x)\|_\infty) + \ln(1 + C). \end{aligned}$$

Furthermore, using $\ln_{\epsilon_n}(x) \geq \ln(x)^1$, along with Lemma 17 and Proposition 18, we have:

$$\limsup_{n \rightarrow \infty} \sup_{\mathbf{y} \in \mathbb{R}^m} \left[\frac{1}{n} \ln \mathbb{E} |\det H_n^\Lambda(\mathbf{y})| \right] \leq \limsup_{n \rightarrow \infty} \sup_{\mathbf{y} \in \mathbb{R}^m} \mathbb{E}_{\mathbf{z}} \ln \rho_n(\mathbf{y}),$$

in which $\rho_n(\mathbf{y})$ is the spectral radius of $\tilde{H}_n^\Lambda(\mathbf{y})$ (see Section A.2.3 for the definition of $\tilde{H}_n^\Lambda(\mathbf{y})$). Using item (i) of Lemma 11, this implies that $\limsup_{n \rightarrow \infty} \sup_{\mathbf{y} \in \mathbb{R}^m} \left[\frac{1}{n} \ln \mathbb{E} |\det H_n^\Lambda(\mathbf{y})| \right] < \infty$. Therefore, in order to prove Lemma 7, it only remains to show that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \ln \mathbb{E}_{\mathbf{y}} \left[e^{-\frac{\gamma n}{2} \ln(\frac{1}{m} \sum_{\mu} \phi'(y_{\mu})^2)} \right] < \infty. \quad (64)$$

Let us now prove eq. (64). We denote $m \equiv \mathbb{E}_{y \sim \mathcal{N}(0,1)}[\phi'(y)^2]$ and $A \equiv \|\phi'\|_{\infty}^2$. Since $A < \infty$, we can apply Cramer's theorem to $S \equiv \frac{1}{m} \sum_{\mu} \phi'(y_{\mu})^2$, so that we have:

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \ln \mathbb{E}_{\mathbf{y}} \left[e^{-\frac{\gamma n}{2} \ln(\frac{1}{m} \sum_{\mu} \phi'(y_{\mu})^2)} \right] \leq \sup_{S \in (0,A)} \left[-\frac{\gamma}{2} \ln S - \alpha \Lambda^*(S) \right], \quad (65)$$

in which $\Lambda^*(S)$ is defined as the Legendre transform of the moment generating function of $\phi'(y)^2$:

$$\Lambda^*(S) \equiv \begin{cases} \sup_{\theta \geq 0} \left\{ \theta S - \ln \mathbb{E}_{y \sim \mathcal{N}(0,1)} \left[e^{\theta \phi'(y)^2} \right] \right\} & \text{if } S \geq m, \\ \sup_{\theta \geq 0} \left\{ -\theta S - \ln \mathbb{E}_{y \sim \mathcal{N}(0,1)} \left[e^{-\theta \phi'(y)^2} \right] \right\} & \text{if } S \leq m. \end{cases}$$

By continuity of the involved functions, in order to conclude from eq. (65) we just need to be able to show that (i) : $\limsup_{S \rightarrow A} (-\Lambda^*(S)) < \infty$ and (ii) : $\limsup_{S \rightarrow 0} \left[-\frac{\gamma}{2} \ln S - \alpha \Lambda^*(S) \right] < \infty$. Point (i) is trivial since $\Lambda^*(S) \geq 0$ for all $S \in (0, A)$ (it is a rate function). To show (ii), we use the fact that for all $S \in (0, m)$ and $\theta \geq 0$ we have $\Lambda^*(S) \geq -\theta S - \ln \mathbb{E}[e^{-\theta \phi'(y)^2}]$. In particular, for $\theta = S^{-1}$ we have $\Lambda^*(S) \geq -1 - \ln \mathbb{E}[e^{-S^{-1} \phi'(y)^2}]$. Since $a = \phi'(y)$ has by Definition 1 a density φ_a continuous around 0, we fix $a_0 > 0$ such that φ_a is continuous in $[-a_0, a_0]$. For every $\theta > 0$:

$$\ln \mathbb{E}[e^{-\theta \phi'(y)^2}] \leq \ln \left[\mathbb{E} \left(e^{-\theta a^2} \mathbf{1}_{|a| \leq a_0} \right) + e^{-\theta a_0^2} \right] \leq \ln \left[\left(\sup_{|a| \leq a_0} |f(a)| \right) \frac{\sqrt{\pi}}{\sqrt{\theta}} + e^{-\theta a_0^2} \right],$$

and thus $\ln \mathbb{E}[e^{-\theta \phi'(y)^2}] \leq C - (1/2) \ln \theta$ with a constant $C > 0$. Using this bound and the remark before we reach

$$-\frac{\gamma}{2} \ln S - \Lambda^*(S) \leq \frac{\alpha - \gamma}{2} \ln S + \alpha(1 + C).$$

Since $\alpha - \gamma > 0$, we have $\lim_{S \rightarrow 0} \left[-\frac{\gamma}{2} \ln S - \alpha \Lambda^*(S) \right] = -\infty$, which obviously implies point (ii), which in turn shows eq. (64). \square

A.4 Proof of eq. (18)

Let $t > 0$, and fix $\eta > 0$ given by Lemma 6. We define $E_n^{(t)}$, A_n , B_n :

$$\left\{ \begin{aligned} E_n^{(t)} &\equiv \left\{ \left| \frac{1}{n} \ln \mathbb{E}_{\mathbf{z}} [|\det H_n^\Lambda(\mathbf{y})|] - \kappa_{\alpha, \phi}(\nu_{\mathbf{y}}^m, t_{\phi}(\nu_{\mathbf{y}}^m)) \right| \geq t \right\}, & (66a) \\ A_n &\equiv \frac{1}{n} \ln \mathbb{E} \left[\mathbf{1}[L_1(\mathbf{y}) \in B] e^{-\frac{\eta}{2} \ln(\frac{1}{m} \sum_{\mu} \phi'(y_{\mu})^2)} \mathbb{E} |\det H_n^\Lambda(\mathbf{y})| \right], & (66b) \\ B_n &\equiv \frac{1}{n} \ln \mathbb{E} \left[\mathbf{1}[L_1(\mathbf{y}) \in B] e^{-\frac{\eta}{2} \ln(\frac{1}{m} \sum_{\mu} \phi'(y_{\mu})^2) + n \kappa_{\alpha, \phi}(\nu_{\mathbf{y}}^m, t_{\phi}(\nu_{\mathbf{y}}^m))} \right]. & (66c) \end{aligned} \right.$$

¹Recall that $\epsilon_n = n^{-\delta}$ and $\delta \in (0, 1)$.

A_n is related to the complexity by Lemma 4, and by Varadhan's lemma (and Lemma 7), we have:

$$\lim_{n \rightarrow \infty} B_n = \sup_{\nu \in \mathcal{M}_\phi(B)} \left[-\frac{1}{2} \mathcal{E}_\phi(\nu) + \kappa_{\alpha, \phi}(\nu, t_\phi(\nu)) - \alpha H(\nu | \mu_G) \right] \in [-\infty, +\infty). \quad (67)$$

The factor α in front of the relative entropy arises as we consider the empirical distribution of $m = \alpha n$ i.i.d. variables. For all $t > 0$, we have by definition of A_n, B_n :

$$\left\{ \begin{array}{l} A_n - B_n \geq -t + \frac{1}{n} \ln \left[1 - \frac{\mathbb{E} \left[\mathbb{1}_{L_1(\mathbf{y}) \in B; E_n^{(t)}} e^{-\frac{n}{2} \ln(\frac{1}{m} \sum_\mu \phi'(y_\mu)^2) + n\kappa_{\alpha, \phi}(\nu_y^m, t_\phi(\nu_y^m))} \right]}{\mathbb{E} \left[\mathbb{1}_{L_1(\mathbf{y}) \in B} e^{-\frac{n}{2} \ln(\frac{1}{m} \sum_\mu \phi'(y_\mu)^2) + n\kappa_{\alpha, \phi}(\nu_y^m, t_\phi(\nu_y^m))} \right]} \right], \end{array} \right. \quad (68a)$$

$$\left\{ \begin{array}{l} A_n - B_n \leq t + \frac{1}{n} \ln \left[1 + e^{-nt} \frac{\mathbb{E} \left[\mathbb{1}_{L_1(\mathbf{y}) \in B; E_n^{(t)}} e^{-\frac{n}{2} \ln(\frac{1}{m} \sum_\mu \phi'(y_\mu)^2)} \mathbb{E} |\det H_n^\Lambda(\mathbf{y})| \right]}{\mathbb{E} \left[\mathbb{1}_{L_1(\mathbf{y}) \in B} e^{-\frac{n}{2} \ln(\frac{1}{m} \sum_\mu \phi'(y_\mu)^2)} \mathbb{E} |\det H_n^\Lambda(\mathbf{y})| \right]} \right]. \end{array} \right. \quad (68b)$$

Using Hölder's inequality and Lemma 7, there exists $\gamma > 1$ and a constant $C > 0$ such that:

$$-t + \frac{1}{n} \ln \left[1 - \frac{\mathbb{P}[E_n^{(t)}]^\frac{1}{\gamma}}{e^{nB_n - nC}} \right] \leq A_n - B_n \leq t + \frac{1}{n} \ln \left[1 + \frac{\mathbb{P}[E_n^{(t)}]^\frac{1}{\gamma}}{e^{nt + nA_n - nC}} \right]. \quad (69)$$

Assume that $\lim B_n = -\infty$ and $\limsup A_n > -\infty$. Let us fix a lower-bounded sub-sequence $A_{\varphi(n)}$ of A_n , so that $\lim[A_{\varphi(n)} - B_{\varphi(n)}] = +\infty$. However, by eq. (69) and Lemma 6, we have $\limsup[A_{\varphi(n)} - B_{\varphi(n)}] \leq t$, as $(1/n) \ln \mathbb{P}[E_n^{(t)}] \rightarrow -\infty$. So we showed that $\lim B_n = -\infty \Rightarrow \lim A_n = -\infty$, which shows eq. (18) in this case. Let us now assume that $\lim B_n > -\infty$. Using the left inequality of eq. (69) and Lemma 6, we reach in the same way that $\liminf[A_n - B_n] \geq -t$, which implies that $\liminf A_n > -\infty$. Thus we can use the right inequality of eq. (69) to show similarly that $\limsup[A_n - B_n] \leq t$. Taking the $t \rightarrow 0$ limit finishes the proof of eq. (18).

A.5 Proof of Lemma 21

Proof. This proof is a generalization of the arguments of Silverstein and Bai [1995]. To fix the notations here, $(\mu_n, g_n(z))$ are the ESD¹ and Stieltjes transform of $\frac{1}{n} \tilde{\mathbf{z}} \Lambda(\mathbf{y}) \tilde{\mathbf{z}}^\top$, $(\mu_n^D, g_n^D(z))$ the ones of $\frac{1}{n} \tilde{\mathbf{z}} D(\mathbf{y}) \tilde{\mathbf{z}}^\top$, and $g_{\alpha, \phi}[\nu_y^m](z)$ is the Stieltjes transform of $\mu_{\alpha, \phi}[\nu_y^m]$. We show:

(i) There exists $\eta_1 \in (0, 1)$ such that for all $z \in \mathbb{C}_+$:

$$\lim_{n \rightarrow \infty} \left\{ n^{\eta_1} \sup_{\mathbf{y} \in \mathbb{R}^m} |\mathbb{E}_{\mathbf{z}} g_n(z) - \mathbb{E}_{\mathbf{z}} g_n^D(z)| \right\} = 0.$$

(ii) There exists $\eta_2 \in (0, 1)$ and $K > 0$ such that for all $z \in \mathbb{C}_+$:

$$\limsup_{n \rightarrow \infty} \left\{ n^{\eta_2} \sup_{\mathbf{y} \in \mathbb{R}^m} |\mathbb{E}_{\mathbf{z}} g_n^D(z) - g_{\alpha, \phi}[\nu_y^m](z)| \right\} < K.$$

Points (i) – (ii) obviously imply Lemma 21, so we now prove them.

Proof of (i) This is a direct consequence of the second part of Lemma 10, which implies that we can fix $\eta_1 \in (0, 1)$ such that uniformly in \mathbf{y} , $\{n^{\eta_1} \mathbb{E}_{\mathbf{z}} [\mu_n^D - \mu_n]\} \rightarrow_{n \rightarrow \infty} 0$ weakly. By a classical characterization of the Stieltjes transform (cf for instance Theorem 2.4.4 of Anderson et al. [2010]), this implies (i).

¹Empirical Spectral Distribution.

Proof of (ii) This part is more involved. Let us give some definitions and conventions. We define $M_n \equiv \frac{1}{n} \tilde{\mathbf{z}} D(\mathbf{y}) \tilde{\mathbf{z}}^\top$. If \mathbf{z}_μ is the column of $\tilde{\mathbf{z}}$ indexed by μ , we denote $M_n^{(\mu)} \equiv \frac{1}{n} \sum_{\nu \neq \mu} D(y_\nu) \mathbf{z}_\nu \mathbf{z}_\nu^\top$, which is independent of \mathbf{z}_μ , and $g_n^{(\mu)}(z)$ its Stieltjes transform¹. Finally, we define x and $x^{(\mu)}$ as:

$$x \equiv \frac{\alpha}{m} \sum_{\mu=1}^m \frac{\phi''(y_\mu)}{\alpha + \phi''(y_\mu) g_n(z)}, \quad x^{(\mu)} \equiv \frac{\alpha}{m} \sum_{\nu \neq \mu} \frac{\phi''(y_\nu)}{\alpha + \phi''(y_\nu) g_n(z)}. \quad (70)$$

We fix $z = z_1 + iz_2 \in \mathbb{C}_+$ and we start from the trivial identity:

$$-\frac{1}{z-x} \mathbf{I}_n = (M_n - z \mathbf{I}_n)^{-1} - \frac{1}{z-x} (M_n - x \mathbf{I}_n) (M_n - z \mathbf{I}_n)^{-1}. \quad (71)$$

For every invertible matrix B , vector \mathbf{q} , and $\tau \in \mathbb{R}$, we have by the Sherman–Morrison formula:

$$\mathbf{q}^\top (B + \tau \mathbf{q} \mathbf{q}^\top)^{-1} = \frac{1}{q + \tau \mathbf{q}^\top B^{-1} \tau} \mathbf{q}^\top B^{-1}.$$

Plugging it into the last equation and taking the averaged trace yields

$$\begin{aligned} -\frac{1}{z-x} - g_n(z) &= -\frac{1}{z-x} \frac{1}{n} \text{Tr} \left[\left(\frac{1}{n} \sum_{\mu} \phi''(y_\mu) \mathbf{z}_\mu \mathbf{z}_\mu^\top - x \mathbf{I}_n \right) (M_n - z \mathbf{I}_n)^{-1} \right], \\ &= \frac{1}{m} \sum_{\mu=1}^m \frac{\alpha \phi''(y_\mu)}{\alpha + \phi''(y_\mu) g_n(z)} d_\mu, \end{aligned} \quad (72)$$

in which we defined

$$d_\mu \equiv \frac{g_n(z)}{z-x} - \frac{1}{z-x} \frac{\alpha + \phi''(y_\mu) g_n(z)}{\alpha + \phi''(y_\mu) \frac{1}{n} \mathbf{z}_\mu^\top (M_n^{(\mu)} - z \mathbf{I}_n)^{-1} \mathbf{z}_\mu} \frac{1}{n} \mathbf{z}_\mu^\top (M_n^{(\mu)} - z \mathbf{I}_n)^{-1} \mathbf{z}_\mu. \quad (73)$$

Let us denote $L_n(z, g) \equiv -g - \left[z - \frac{\alpha}{m} \sum_{\mu=1}^m \frac{\phi''(y_\mu)}{\alpha + \phi''(y_\mu) g} \right]^{-1}$. We know by [Silverstein and Bai \[1995\]](#) that $g_{\alpha, \phi}[\nu_{\mathbf{y}}^m](z)$ is defined as the only solution in \mathbb{C}_+ to $L_n(z, g) = 0$. Let us first show, using eq. (72), that there exists $\eta \in (0, 1)$ such that:

$$\lim_{n \rightarrow \infty} n^\eta \sup_{\mathbf{y} \in \mathbb{R}^m} |L_n(z, \mathbb{E} g_n(z))| = 0. \quad (74)$$

Proof of eq. (74). We use Lemma 3.1 of [Silverstein and Bai \[1995\]](#), which gives a $K > 0$ (independent of \mathbf{y} and z) such that uniformly over \mathbf{y} :

$$\mathbb{E}_{\mathbf{z}} \left| \frac{1}{z-x^{(\mu)}} \frac{1}{n} \mathbf{z}_\mu^\top (M_n^{(\mu)} - z \mathbf{I}_n)^{-1} \mathbf{z}_\mu - \frac{g_n^{(\mu)}(z)}{z-x^{(\mu)}} \right|^6 \leq K z_2^{-12} n^{6\gamma-3}. \quad (75)$$

Moreover, let us denote $\Delta \equiv \frac{z_2}{2(C^2 + z_1^2) + z_2^2}$, in which C is the constant of Lemma 11. Thanks to this lemma, we have that (uniformly in \mathbf{y} and a.s.) $\liminf_{n \rightarrow \infty} \text{Im}[g_n(z)] \geq \Delta$. In particular, using this bound, we can easily show (cf Part 4 of [Silverstein and Bai \[1995\]](#) for a completely similar argument) that, uniformly in \mathbf{y} and almost surely:

$$\left\{ \max_{1 \leq \mu \leq m} \left| \frac{\alpha \phi''(y_\mu)}{\alpha + \phi''(y_\mu) g_n(z)} \right| \leq \frac{\alpha}{\Delta}, \right. \quad (76a)$$

$$\left. \left[\lim_{n \rightarrow \infty} \left[\max_{1 \leq \mu \leq m} \max \left\{ \left| \frac{\alpha + \phi''(y_\mu) g_n(z)}{\alpha + \phi''(y_\mu) \frac{1}{n} \mathbf{z}_\mu^\top (M_n^{(\mu)} - z \mathbf{I}_n)^{-1} \mathbf{z}_\mu} \right|, |x - x^{(\mu)}| \right\} \right] \right] = 0. \right. \quad (76b)$$

¹To lighten the results, we state the results as if these matrices were of size n , even though they are of size $(n-1)$, as it does not change anything to the argument.

Since $\gamma < 1/2$, combining eqs. (75),(76) into eqs. (72),(73) yields that there exists $\eta \in (0, 1)$ such that eq. (74) is satisfied. \blacksquare

By eq. (74) there exists $K > 0$ and a function $K(\mathbf{y}) \in \mathbb{C}$, such that for all \mathbf{y} : $|K(\mathbf{y})| \leq K$, and

$$\mathbb{E}g_n(z) = - \left[z - \frac{\alpha}{m} \sum_{\mu=1}^m \frac{\phi''(y_\mu)}{\alpha + \phi''(y_\mu)\mathbb{E}g_n(z)} \right]^{-1} + \frac{K(\mathbf{y})}{n^\eta}. \quad (77)$$

We write $g_{\alpha,\phi}[\nu_{\mathbf{y}}^m](z) = m_1(z) + im_2(z)$ and $\mathbb{E}g_n(z) = m_{n,1}(z) + im_{n,2}(z)$. Note that all imaginary parts are strictly positive. Taking the imaginary part of eq. (77):

$$m_{n,2}(z) = \frac{\text{Im}(K(\mathbf{y}))}{n^\eta} + \frac{z_2 + \frac{\alpha}{m} \sum_{\mu=1}^m \frac{\phi''(y_\mu)^2 m_{n,2}(z)}{|\alpha + \phi''(y_\mu)\mathbb{E}g_n(z)|^2}}{\left| z - \frac{\alpha}{m} \sum_{\mu=1}^m \frac{\phi''(y_\mu)}{\alpha + \phi''(y_\mu)\mathbb{E}g_n(z)} \right|^2}. \quad (78)$$

Using eq. (77) and its counterpart for $g_{\alpha,\phi}[\nu_{\mathbf{y}}^m](z)$ (which does not have a second term), we reach

$$g_{\alpha,\phi}[\nu_{\mathbf{y}}^m](z) - \mathbb{E}g_n(z) = \frac{K(\mathbf{y})}{n^\eta} + [g_{\alpha,\phi}[\nu_{\mathbf{y}}^m](z) - \mathbb{E}g_n(z)] A_n(z), \quad (79)$$

with

$$A_n(z) \equiv \frac{\frac{\alpha}{m} \sum_{\mu} \frac{\phi''(y_\mu)^2}{(\alpha + \phi''(y_\mu)\mathbb{E}g_n(z))(\alpha + \phi''(y_\mu)g_{\alpha,\phi}[\nu_{\mathbf{y}}^m](z))}}{\left[z - \frac{\alpha}{m} \sum_{\mu=1}^m \frac{\phi''(y_\mu)}{\alpha + \phi''(y_\mu)\mathbb{E}g_n(z)} \right] \left[z - \frac{\alpha}{m} \sum_{\mu=1}^m \frac{\phi''(y_\mu)}{\alpha + \phi''(y_\mu)g_{\alpha,\phi}[\nu_{\mathbf{y}}^m](z)} \right]}. \quad (80)$$

By the Cauchy-Schwarz inequality, $|A_n(z)| \leq \sqrt{A_1(z)A_2(z)}$, with

$$\left\{ \begin{array}{l} A_1(z) \equiv \frac{\alpha}{m} \sum_{\mu=1}^m \frac{\phi''(y_\mu)^2}{|\alpha + \phi''(y_\mu)\mathbb{E}g_n(z)|^2} \left| z - \frac{\alpha}{m} \sum_{\mu=1}^m \frac{\phi''(y_\mu)}{\alpha + \phi''(y_\mu)\mathbb{E}g_n(z)} \right|^{-2}, \end{array} \right. \quad (81a)$$

$$\left\{ \begin{array}{l} A_2(z) \equiv \frac{\alpha}{m} \sum_{\mu=1}^m \frac{\phi''(y_\mu)^2}{|\alpha + \phi''(y_\mu)g_{\alpha,\phi}[\nu_{\mathbf{y}}^m](z)|^2} \left| z - \frac{\alpha}{m} \sum_{\mu=1}^m \frac{\phi''(y_\mu)}{\alpha + \phi''(y_\mu)g_{\alpha,\phi}[\nu_{\mathbf{y}}^m](z)} \right|^{-2}. \end{array} \right. \quad (81b)$$

In particular, using the counterpart of eq. (78) for $m_2(z)$, we have:

$$A_1(z) = \frac{\alpha}{m} \sum_{\mu=1}^m \frac{\phi''(y_\mu)^2 m_2(z)}{|\alpha + \phi''(y_\mu)\mathbb{E}g_n(z)|^2} \frac{1}{z_2 + \frac{\alpha}{m} \sum_{\mu=1}^m \frac{\phi''(y_\mu)^2 m_{n,2}(z)}{|\alpha + \phi''(y_\mu)\mathbb{E}g_n(z)|^2}}. \quad (82)$$

Using items (i) and (ii) of Lemma 11, we have the inequalities: $\frac{z_2}{2(C^2+z_1^2)+z_2^2} \leq m_2(z) \leq \frac{1}{z_2}$ and $\frac{z_2}{2(C^2+z_1^2)+z_2^2} \leq m_{n,2}(z) \leq \frac{1}{z_2}$. This implies from eq. (82):

$$A_1(z) \leq \frac{z_2}{2(C^2+z_1^2)+z_2^2} \frac{1}{z_2 + \frac{z_2}{2(C^2+z_1^2)+z_2^2}}. \quad (83)$$

In particular, there exists a constant $\Gamma \in (0, 1)$ such that $A_1(z) \leq \Gamma$ uniformly in n, \mathbf{y} . Similarly, we find that $A_2(z) \leq \Gamma(1 + \frac{K}{z_2 n^\eta})$ uniformly in \mathbf{y} . So we have from eq. (79):

$$\limsup_{n \rightarrow \infty} \left\{ n^\eta \sup_{\mathbf{y} \in \mathbb{R}^m} |\mathbb{E}g_n^D(z) - g_{\alpha,\phi}[\nu_{\mathbf{y}}^m](z)| \right\} \leq \frac{K}{1-\Gamma} < \infty, \quad (84)$$

which proves item (ii). \square

A.6 Proof of Lemma 16

For simplicity, we denote here z_{11} by z . Since $\mathbb{E}z^2 = 1$, one has:

$$\begin{aligned} \mathbb{E} \left[\mathbb{E} \left(z^2 \mathbb{1}_{|z| \leq n^{\gamma/2}} \right) - \mathbb{1}_{|z| \leq n^{\gamma/2}} \right]^4 &= \mathbb{E} \left[\mathbb{E} \left(z^2 \mathbb{1}_{|z| > n^{\gamma/2}} \right) - \mathbb{1}_{|z| > n^{\gamma/2}} \right]^4, \\ &\leq 8 \left[\mathbb{E} \left(z^2 \mathbb{1}_{|z| > n^{\gamma/2}} \right)^4 + \mathbb{P}_{|z| > n^{\gamma/2}} \right] \leq 80 \mathbb{P}_{|z| > n^{\gamma/2}}. \end{aligned}$$

This immediately yields the first part of Lemma. 16. We now tackle with the second part:

$$\mathbb{E} \left[\left(\mathbb{E} \left(z^2 \mathbb{1}_{|z| \leq n^{\gamma/2}} \right) \right)^{1/2} - \mathbb{1}_{|z| \leq n^{\gamma/2}} \right]^8 \leq 2^7 \left[\left(\mathbb{E} \left(z^2 \mathbb{1}_{|z| \leq n^{\gamma/2}} \right) \right)^4 + \mathbb{P}_{|z| > n^{\gamma/2}} \right] \leq 2^8 5 \mathbb{P}_{|z| > n^{\gamma/2}},$$

which finishes the proof.

B Technical aspects of the quenched calculation

B.1 The phase volume factor

Introducing the Fourier transform of the delta functions, we reach at leading exponential order in n :

$$\frac{1}{n} \ln \left[\prod_{a=1}^p \int_{\mathbb{R}^n} d\mathbf{x}^a \prod_{a \leq b} \delta \left(nq_{ab} - n\mathbf{x}^a \cdot \mathbf{x}^b \right) \right] \simeq \frac{p}{2} \ln \frac{2\pi}{n} + \frac{1}{2} \sup_{\{\hat{q}_{ab}\}} \left[\sum_{a,b} q_{ab} \hat{q}_{ab} - \ln \det \hat{\mathbf{q}} \right].$$

The replica symmetric assumption can be made on the variables $\hat{\mathbf{q}}$ that achieve this supremum : $\hat{q}_{aa} = \hat{q}_0$ and $\hat{q}_{ab} = -\hat{q}$ for $a \neq b$. This leads to $\det \hat{\mathbf{q}} = (\hat{q}_0 + \hat{q})^{p-1} (\hat{q}_0 - (p-1)\hat{q})$, and after taking the $p \rightarrow 0^+$ limit, we reach:

$$\frac{1}{np} \ln \left[\prod_{a=1}^p \int_{\mathbb{R}^n} d\mathbf{x}^a \prod_{a \leq b} \delta \left(nq_{ab} - n\mathbf{x}^a \cdot \mathbf{x}^b \right) \right] \simeq \frac{1}{2} \log \frac{2\pi}{n} + \frac{1}{2} \sup_{\hat{q}_0, \hat{q}} \left[\hat{q}_0 + q\hat{q} - \log(\hat{q}_0 + \hat{q}) + \frac{\hat{q}}{\hat{q}_0 + \hat{q}} \right].$$

The diverging term $-\frac{1}{2} \log n$ will be canceled out by the joint density of the gradients as we will see later. The solution of the supremum is easy to carry out, and we finally reach eq. (36).

B.2 The joint density of the gradients

We denote $S = \text{Span}(\{\mathbf{x}^a\}_{a=1}^p)$. Following Ros et al. [2019], for every $1 \leq a \leq p$ we can construct an orthonormal basis of S , denoted $(\mathbf{e}_b^a)_{1 \leq b \leq p}$ for which \mathbf{x}^a is the first vector, that is $\mathbf{e}_a^a = \mathbf{x}^a$. This basis is convenient, since $\{\mathbf{x}^a\}^\perp \cap S = \text{Span}(\{\mathbf{e}_b^a\}_{b(\neq a)})$. We can also chose an arbitrary orthonormal basis $(\mathbf{e}_{p+1}, \dots, \mathbf{e}_n)$ of S^\perp . With this choice of basis, we can see that the gradient $\text{grad } L(\mathbf{x}^a)$ is identified with the vector in \mathbb{R}^{n-1} with components:

$$\text{grad } L(\mathbf{x}^a) = (\{\nabla L(\mathbf{x}^a) \cdot \mathbf{e}_i^a\}_{i=1}^{a-1}, \{\nabla L(\mathbf{x}^a) \cdot \mathbf{e}_i^a\}_{i=a+1}^p, \{\nabla L(\mathbf{x}^a) \cdot \mathbf{e}_i^a\}_{i=p+1}^n). \quad (85)$$

Recall that $\nabla L(\mathbf{x}^a) = \frac{1}{m} \sum_{\mu} \xi_{\mu} \phi'(y_{\mu}^a)$. Let us make a few remarks:

- For every a , the basis $(\mathbf{e}_b^a)_{b=1}^p$ is only a function of the values of the overlaps $\{q_{ab}\}$.
- We consider the joint density of the gradients *conditioned* by the value of $\{\mathbf{y}^a\}$. In particular, this means that for every $a \neq b$, $\nabla L(\mathbf{x}^a) \cdot \mathbf{e}_b^a$ is fixed by the values of $\{\mathbf{y}^c\}_{c=1}^p$ and the overlaps q_{ab} . In particular, the first $(p-1)$ components of eq. (85) are deterministic, thus their density will yield delta functions that are constraints on $\{\mathbf{y}^a\}$ and $\{q_{ab}\}$.

- The last $n - p$ components of eq. (85) are (at fixed $\{\mathbf{y}^a\}$) zero mean Gaussian random variables with covariance given by $\mathbb{E} [\text{grad}L(\mathbf{x}^a)_i \text{grad}L(\mathbf{x}^b)_j] = \frac{\delta_{ij}}{m^2} \sum_{\mu} \phi'(y_{\mu}^a) \phi'(y_{\mu}^b)$. Their joint density taken at 0 is thus at leading exponential order in the $n \rightarrow \infty$ limit:

$$\exp \left\{ \frac{np}{2} \log \frac{m}{2\pi} - \frac{n}{2} \log \det \left[\left(\frac{1}{m} \sum_{\mu=1}^m \phi'(y_{\mu}^a) \phi'(y_{\mu}^b) \right)_{1 \leq a, b \leq p} \right] \right\}. \quad (86)$$

Given these remarks and eq. (86), in order to complete the calculation of the joint gradient density we need to compute the quantities $(\nabla L(\mathbf{x}^a) \cdot \mathbf{e}_b^a)$ for every $a \neq b$ as a function of $\{y_{\mu}^a\}$ and $\{q_{ab}\}$. In order to simplify the calculation, we will already make use of the replica-symmetric assumption on q , that is we assume $q_{aa} = 1$ and $q_{ab} = q$ for $a \neq b$. Let us now describe a possible construction for the basis $(\mathbf{e}_b^a)_{b=1}^p$. We introduce three auxiliary quantities that are functions of q and p :

$$f_p^0(q) \equiv \frac{1}{p-1} \left[\frac{p-2}{\sqrt{1-q}} + \frac{1}{\sqrt{1+(p-2)q-(p-1)q^2}} \right], \quad (87a)$$

$$f_p(q) \equiv \frac{1}{p-1} \left[-\frac{1}{\sqrt{1-q}} + \frac{1}{\sqrt{1+(p-2)q-(p-1)q^2}} \right], \quad (87b)$$

$$z_p(q) \equiv -\frac{q}{\sqrt{1+(p-2)q-(p-1)q^2}}. \quad (87c)$$

With these definitions, we can consider:

$$\begin{cases} \mathbf{e}_a^a & \equiv \mathbf{x}^a, \\ \mathbf{e}_b^a & \equiv z_p(q)\mathbf{x}^a + f_p^0(q)\mathbf{x}^b + f_p(q) \sum_{c(\neq a,b)} \mathbf{x}^c, \quad (b \neq a). \end{cases} \quad (88)$$

It is straightforward to check from eq. (88) that we have for all a, b, c that $\mathbf{e}_b^a \cdot \mathbf{e}_c^a = \delta_{bc}$. We can now see that the delta term of the joint density of the gradients taken at 0 is:

$$\prod_{a \neq b} \delta [\nabla L_1(\mathbf{x}^a) \cdot \mathbf{e}_b^a] = \prod_{a \neq b} \delta \left[\frac{1}{m} \sum_{\mu=1}^m \phi'(y_{\mu}^a) \left(z_p(q)y_{\mu}^a + f_p^0(q)y_{\mu}^b + f_p(q) \sum_{c(\neq a,b)} y_{\mu}^c \right) \right]. \quad (89)$$

The product of eq. (86) and eq. (89) gives eq. (37).

B.3 Decoupling the replicas and taking the $p \rightarrow 0^+$ limit

B.3.1 Replica symmetry and decoupling

In order to apply the replica method, we need to be able to take the $p \rightarrow 0$ limit, by analytically extending eq. (41) to all $p > 0$. The main idea is that we expect replica symmetry to influence the measure ν that appears in eq. (41). More precisely, we expect that for all permutation $\pi \in \mathcal{S}_p$, we have $\nu(d\lambda^1, \dots, d\lambda^p) = \nu(d\lambda^{\pi(1)}, \dots, d\lambda^{\pi(p)})$. We make in eq. (41) the substitution:

$$\sup_{\nu \in \mathcal{M}(p,q)} \rightarrow \sup_{\{\mu_a\}_{a=1}^p \in \mathcal{M}(\mathbb{R})} \sup_{\substack{\nu \in \mathcal{M}(p,q) \\ \text{s.t. } \{\nu^a = \mu_a\}}} \quad (90)$$

In this last expression, the replica symmetric assumption leads us in particular to assume that $\mu_a = \mu$ for all a . In order to make the remaining calculation tractable we will also need to fix some linear statistics of ν via Lagrange multipliers:

- For every $a \leq b$, we will fix the linear statistics $\int \nu(d\lambda) \phi'(\lambda^a) \phi'(\lambda^b) = A_{ab}$, with Lagrange multipliers \hat{A}_{ab} . Note that by replica symmetry, we actually assume that $A_{ab} = a$ for $a \neq b$ and $A_{aa} = A$ (and samely for the Lagrange multipliers).

- For all a, b we will fix the linear statistics $\int \nu(d\lambda) \phi'(\lambda^a) \lambda^b = B_{ab}$, with Lagrange multipliers \hat{B}_{ab} . By replica symmetry, we assume that $B_{aa} = B$ and $B_{ab} = b$.

Combining these remarks, we reach that the ν -dependent term of eq. (41) is equal to:

$$\begin{aligned} & \sup_{\mu \in \mathcal{M}_\phi(B)} \sup_{\substack{A,a \\ B,b}} \text{extr}_{\substack{\hat{A},\hat{a} \\ \hat{B},\hat{b}}} \sup_{\nu \in \mathcal{M}(\mathbb{R}^n)} \left[p\kappa_{\alpha,\phi}[\mu, t_\phi(\mu)] - \frac{1}{2} \ln \det [\{A_{ab}\}] - \sum_{a,b} \left[\frac{1}{2} A_{ab} \hat{A}_{ab} + B_{ab} \hat{B}_{ab} \right] \right. \\ & \left. + \sum_{a,b} \left[\frac{1}{2} \hat{A}_{ab} \int \nu(d\lambda) \phi'(\lambda^a) \phi'(\lambda^b) + \hat{B}_{ab} \int \nu(d\lambda) \phi'(\lambda^a) \lambda^b \right] - \alpha H(\nu | \mu_{G,q}) \right]. \end{aligned} \quad (91)$$

Note that here we did not always explicit the replica-symmetry assumption on all the variables to obtain more compact expressions. The supremum over B, b is moreover constrained by the following condition of eq. (40b): $\forall a \neq b, z_p(q)B_{aa} + f_p^0(q)B_{ab} + f_p(q) \sum_{c(\neq a,b)} B_{ac} = 0$. Under the replica symmetric assumption, this becomes:

$$z_p(q)B + f_p^0(q)b + f_p(q)(p-2)b = 0. \quad (92)$$

Again, we introduce Lagrange multipliers C_{ab} to fix these conditions, that reduce to $C_{ab} = C$ because of replica symmetry. Finally, in order to fix the marginal distributions of ν , we will have to introduce 'functional' Lagrange multipliers $g^a(\lambda^a)$. Again, by replica symmetry, we expect all of them to be equal to $g(\lambda^a)$. In the end, we reach the simplification of eq. (91):

$$\begin{aligned} & \sup_{\substack{\mu \in \mathcal{M}_\phi(B) \\ \nu \in \mathcal{M}(\mathbb{R}^n)}} \sup_{\substack{A,a \\ B,b}} \text{extr}_{\substack{C,\hat{A},\hat{a} \\ \hat{B},\hat{b},\{g(\lambda)\}}} \left\{ p\kappa_{\alpha,\phi}[\mu, t_\phi(\mu)] - \frac{1}{2} \ln \det [\{A_{ab}\}] - \sum_{a,b} \left[\frac{1}{2} \hat{A}_{ab} A_{ab} + \hat{B}_{ab} B_{ab} \right] \right. \\ & - p \int \mu(d\lambda) g(\lambda) + p(p-1)C [Bz_p(q) + b(f_p^0(q) + (p-2)f_p(q))] - \alpha H(\nu | \mu_{G,q}) \\ & \left. + \sum_{a,b} \left[\frac{\hat{A}_{ab}}{2} \int \nu(d\lambda) \phi'(\lambda^a) \phi'(\lambda^b) + \hat{B}_{ab} \int \nu(d\lambda) \phi'(\lambda^a) \lambda^b \right] + \sum_a \int \nu(d\lambda) g(\lambda^a) \right\}. \end{aligned} \quad (93)$$

We can now solve exactly the supremum over ν . By a classical Gibbs measure calculation that we already detailed in section 2, we obtain (recall that $Q \in \mathbb{R}^{p \times p}$ is the overlap matrix):

$$\begin{aligned} & \sup_{\nu \in \mathcal{M}(\mathbb{R}^n)} \left\{ \int \nu(d\lambda) \left[\sum_{a,b} \left(\frac{\hat{A}_{ab}}{2} \phi'(\lambda^a) \phi'(\lambda^b) + \hat{B}_{ab} \phi'(\lambda^a) \lambda^b \right) + \sum_a g(\lambda^a) \right] - \alpha H(\nu | \mu_{G,q}) \right\} \\ & = \alpha \ln \left[\int_{\mathbb{R}^p} \frac{d\lambda}{\sqrt{2\pi^p} \sqrt{\det Q}} e^{\sum_{a,b} \left(-\frac{1}{2} (Q^{-1})_{ab} \lambda^a \lambda^b + \frac{\hat{A}_{ab}}{2\alpha} \phi'(\lambda^a) \phi'(\lambda^b) + \frac{\hat{B}_{ab}}{\alpha} \phi'(\lambda^a) \lambda^b \right) + \sum_a \frac{g(\lambda^a)}{\alpha}} \right]. \end{aligned} \quad (94)$$

To completely decouple the replicas, we will make use of two classical identities, for any x, y :

$$e^{\frac{x^2}{2}} = \int \mathcal{D}\xi e^{\xi x}, \quad e^{xy} = \int \mathcal{D}\xi \mathcal{D}\xi' e^{\frac{x}{\sqrt{2}}(\xi + i\xi') + \frac{y}{\sqrt{2}}(\xi - i\xi')}.$$

Thanks to replica symmetry, we can compute Q^{-1} and $\det Q$ as:

$$\begin{cases} \det Q = (1-q)^{p-1} [1 + (p-1)q], & (95a) \\ Q_{ab}^{-1} = \frac{1 + (p-1)q}{1 + (p-2)q - (p-1)q^2} \delta_{ab} - \frac{q}{(1-q)(1 + (p-1)q)}. & (95b) \end{cases}$$

We define

$$\begin{cases} d_{0,p}(q) \equiv \frac{1 + (p-1)q}{1 + (p-2)q - (p-1)q^2}, & (96a) \\ d_p(q) \equiv \frac{q}{(1-q)(1 + (p-1)q)}. & (96b) \end{cases}$$

Using all the above, we can now simplify eq. (94):

$$\begin{aligned} & \alpha \ln \left[\int_{\mathbb{R}^p} \frac{d\lambda}{\sqrt{2\pi^p} \sqrt{\det Q}} e^{\sum_{a,b} \left(-\frac{1}{2}(Q^{-1})_{ab} \lambda^a \lambda^b + \frac{\hat{A}_{ab}}{2\alpha} \phi'(\lambda^a) \phi'(\lambda^b) + \frac{\hat{B}_{ab}}{\alpha} \phi'(\lambda^a) \lambda^b \right) + \sum_a \frac{g(\lambda^a)}{\alpha}} \right] \\ &= -\frac{\alpha p}{2} \ln 2\pi - \frac{\alpha(p-1)}{2} \ln(1-q) - \frac{\alpha}{2} \ln[1 + (p-1)q] + \alpha \ln \left[\int_{\mathbb{R}^4} \mathcal{D}\xi I_p(\xi)^p \right], \end{aligned} \quad (97)$$

in which we defined $\xi \equiv (\xi_q, \xi_a, \xi_b, \xi'_b)$ and

$$I_p(\xi) \equiv \int d\lambda e^{\frac{g(\lambda)}{\alpha} - \frac{d_{0,p}(q)\lambda^2}{2} + \frac{\hat{A}-\hat{a}}{2\alpha} \phi'(\lambda)^2 + \frac{\hat{B}-\hat{b}}{\alpha} \phi'(\lambda)\lambda + \sqrt{d_p(q)}\xi_q\lambda + \sqrt{\frac{\hat{a}}{\alpha}}\xi_a\phi'(\lambda) + \sqrt{\frac{\hat{b}}{2\alpha}}[\phi'(\lambda)(\xi_b + i\xi'_b) + \lambda(\xi_b - i\xi'_b)]}.$$

While the involved expressions are very cumbersome, we have successfully decoupled the replicas.

B.3.2 The $p \rightarrow 0$ limit and final result

We begin by a remark on eq. (97). Note that $\lim_{p \rightarrow 0} (1/p) \ln \left[\int \mathcal{D}\xi I_p(\xi)^p \right] = \int \mathcal{D}\xi \ln [I_0(\xi)]$. Thus, after multiplication by $(1/p)$, the $p \rightarrow 0^+$ limit of eq. (97) will yield:

$$-\frac{\alpha}{2} \ln 2\pi - \frac{\alpha}{2} \ln(1-q) - \frac{\alpha q}{2(1-q)} + \alpha \int \mathcal{D}\xi \ln I(\xi), \quad (98)$$

in which $I(\xi)$ is defined in Result 1. We can wrap up the calculation. We make two remarks. Firstly the condition eq. (92) reduces, in the $p \rightarrow 0$ limit, to $b = qB$, so that we will be able to simplify the terms involving the Lagrange multiplier C . Secondly, the variable B is equal to $t_\phi(\mu)$, defined in Theorem 1. We combine now eqs. (40a),(41),(93) and (98) with the two remarks above. Changing notations from μ to ν and B to C , we obtain finally the conclusion of Result 1.

C Generalization to more involved models

C.1 Annealed and quenched calculations for L_2

We give here a sketch of the generalization of our annealed and quenched calculations to L_2 , thus yielding Theorem 2 and Result 2. A more detailed derivation of these results will be included in a future extended version of this work. We restrict here to the annealed calculation (the quenched calculation is completely similar). The majority of the arguments being identical to the L_1 case, we will only highlight the main differences and give the important intermediary results.

In the Kac-Rice formula, one has now to integrate over the overlap $q \equiv \mathbf{x} \cdot \mathbf{x}^*$ as well. Moreover, we condition over the joint values of $a_\mu \equiv \xi_\mu \cdot \mathbf{x}$ and $b_\mu \equiv (1 - q^2)^{-1/2} [(\xi_\mu \cdot \mathbf{x}^*) - qa_\mu]$, rather than just $\xi_\mu \cdot \mathbf{x}$ (as we did for L_1). Note that (a_μ, b_μ) follows a joint standard Gaussian distribution. Using these definitions, we can obtain the counterpart of Lemma 4 for L_2 :

$$\mathbb{E} \text{Crit}_{n,L_2}(B, Q) = \mathcal{C}_n \int_Q dq e^{\frac{n(1+\ln \alpha + \ln(1-q^2))}{2}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} \left[\delta(P_n(\mathbf{a}, \mathbf{b})) \mathbb{1}_{L_2(\mathbf{a}, \mathbf{b}) \in B} e^{-n\mathcal{E}_n(\mathbf{a}, \mathbf{b})} \mathbb{E}_{\mathbf{z}} |\det H_n(\mathbf{a}, \mathbf{b})| \right],$$

in which \mathcal{C}_n is exponentially trivial, and we defined:

$$\begin{cases} P_n(\mathbf{a}, \mathbf{b}) \equiv \frac{1}{m} \sum_{\mu=1}^m b_\mu \phi'(a_\mu) \left[\phi \left(qa_\mu + \sqrt{1-q^2} b_\mu \right) - \phi(a_\mu) \right], & (99a) \\ \mathcal{E}_n(\mathbf{a}, \mathbf{b}) \equiv \frac{1}{2} \ln \left\{ \sum_{\mu=1}^m \phi'(a_\mu)^2 \left[\phi \left(qa_\mu + \sqrt{1-q^2} b_\mu \right) - \phi(a_\mu) \right]^2 \right\}, & (99b) \\ H_n(\mathbf{a}, \mathbf{b}) \equiv \frac{1}{m} \sum_{\mu=1}^m \left[\phi'(a_\mu)^2 - \theta''(a_\mu) \left[\phi \left(\sqrt{1-q^2} b_\mu + qa_\mu \right) - \phi(a_\mu) \right] \right] \mathbf{z}_\mu \mathbf{z}_\mu^T \\ \quad - \left[\frac{1}{m} \sum_{\mu=1}^m a_\mu \phi'(a_\mu) \left[\phi(a_\mu) - \phi \left(qa_\mu + \sqrt{1-q^2} b_\mu \right) \right] \right] \mathbb{I}_{n-2}. & (99c) \end{cases}$$

Here $\mathbf{z} \in \mathbb{R}^{(n-2) \times m}$ is an i.i.d. standard Gaussian matrix. The condition $P_q(\mathbf{a}, \mathbf{b}) = 0$ arises from the conditioning on the nullity of the gradient in the linear subspace spanned by \mathbf{x}^* , and $\mathcal{E}_n(\mathbf{a}, \mathbf{b})$ from the density of the gradient in the subspace orthogonal to $\{\mathbf{x}, \mathbf{x}^*\}$. A crucial feature of this equation is that the joint distribution of $(L_2(\mathbf{x}), \text{grad}L_2(\mathbf{x}), \text{Hess}L_2(\mathbf{x}))$ only depends on \mathbf{x} via the overlap $q = \mathbf{x} \cdot \mathbf{x}^*$ with the “true” solution. Once conditioned over the values of q , it thus becomes clear why the calculations made for L_1 will generalize here.

As in Section 2.3, one can then show the concentration of the empirical logarithmic potential on the functional $\kappa_{\alpha, \phi}(q, \nu_{\mathbf{a}, \mathbf{b}}^m)$, in which $\nu_{\mathbf{a}, \mathbf{b}}^m \in \mathcal{M}(\mathbb{R}^2)$ is now the empirical measure of $\{a_\mu, b_\mu\}_{\mu=1}^m$. We obtain the counterpart of Lemma 6: there exists $\eta > 0$ such that for all $t > 0$:

$$\lim_{n \rightarrow \infty} \frac{1}{n^{1+\eta}} \ln \mathbb{P} \left[\left| \frac{1}{n} \ln \mathbb{E}_{\mathbf{z}} |\det H_n(\mathbf{a}, \mathbf{b})| - \kappa_{\alpha, \phi}(q, \nu_{\mathbf{a}, \mathbf{b}}^m) \right| \geq t \right] = -\infty. \quad (100)$$

Thanks to this result, we perform then a saddle point on the overlap q and the empirical measure $\nu \in \mathcal{M}(\mathbb{R}^2)$, using Sanov’s theorem and Varadhan’s lemma. This yields the result of Theorem 2.

As a final note, there exists similar results to the one presented in Section 3 that allow to compute the density (and the logarithmic potential) of $\mu_{\alpha, \phi}[q, \nu]$, via the computation of its Stieltjes transform.

C.2 Generalizations to other models

Our calculations, both annealed and quenched, can be generalized straightforwardly to many other loss functions and models. As is clear for instance in the annealed computation of Section 2, the key features that must be present are a Gaussian distribution for the data, and a loss function $L(\mathbf{x})$ that only depends on the data samples ξ_μ via their projection over a few vectors (as \mathbf{x} for $L_1(\mathbf{x})$ and \mathbf{x}, \mathbf{x}^* for $L_2(\mathbf{x})$). We give thereafter three examples of models, that can be found in Engel and Van den Broeck [2001], Mei et al. [2018], and for which the calculations can be performed.

Model 1 (Binary linear classification). Consider $n, m \geq 1$ such that $m/n \rightarrow \alpha > 1$. Let $\sigma : \mathbb{R} \rightarrow [0, 1]$ a smooth threshold function. We are given m samples $(y_\mu, \mathbf{x}_\mu)_{\mu=1}^m$ with $y_\mu \in \{0, 1\}$ and $\mathbf{x}_\mu \in \mathbb{R}^n$. The elements of $(y_\mu)_{\mu=1}^m$ are generated according to $\mathbb{P}(Y_\mu = 1 | \mathbf{X}_\mu = \mathbf{x}) = \sigma(\boldsymbol{\theta}_0 \cdot \mathbf{x})$, and the \mathbf{x}_μ are i.i.d. standard Gaussian random variables in \mathbb{R}^n . We want to learn the vector $\boldsymbol{\theta}_0 \in \mathbb{S}^{n-1}$ by minimizing the loss function:

$$L(\boldsymbol{\theta}) \equiv \frac{1}{2m} \sum_{\mu=1}^m [y_\mu - \sigma(\boldsymbol{\theta} \cdot \mathbf{x}_\mu)]^2, \quad \boldsymbol{\theta} \in \mathbb{S}^{n-1}. \quad (101)$$

Model 2 (Mixture of two Gaussians). Consider $n, m \geq 1$ such that $m/n \rightarrow \alpha > 1$. We are given m samples $\mathbf{y}_\mu \in \mathbb{R}^n$, generated as $\mathbf{y}_\mu \stackrel{\text{i.i.d.}}{\sim} \sum_{a=1,2} p_a \mathcal{N}(\boldsymbol{\theta}_a^0, \mathbb{I}_n)$. The proportions p_1, p_2 are known, and we wish to recover $\boldsymbol{\theta}_1^0$ and $\boldsymbol{\theta}_2^0$ by minimizing the maximum-likelihood estimator:

$$L(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \equiv -\frac{1}{m} \sum_{\mu=1}^m \ln \left[\frac{1}{2} \sum_{a=1,2} \frac{1}{\sqrt{2\pi}^n} \exp \left\{ -\frac{1}{2} \|\mathbf{y}_\mu - \boldsymbol{\theta}_a\|^2 \right\} \right]. \quad (102)$$

Model 3 (Simple unsupervised learning model). Consider $n, m \geq 1$ such that $m/n \rightarrow \alpha > 1$. Let $\phi : \mathbb{R} \rightarrow \mathbb{R}$ a smooth activation function, $V : \mathbb{R} \rightarrow \mathbb{R}_+$ a “potential”, and $\mathbf{x}^0 \in \mathbb{S}^{n-1}$ a fixed vector. We assume that we are given i.i.d. data samples $\{\xi_\mu\}_{\mu=1}^m \in \mathbb{R}^n$ distributed such that their projection on \mathbf{x}^0 has a probability density $P(\xi_\mu \cdot \mathbf{x}^0 = h) \propto e^{-\frac{1}{2}h^2 - V(h)}$, and the other coordinates of ξ_μ are i.i.d. standard Gaussian variables. We wish to recover the vector \mathbf{x}^0 by minimizing:

$$L(\mathbf{x}) \equiv \frac{1}{m} \sum_{\mu=1}^m \phi(\xi_\mu \cdot \mathbf{x}), \quad \mathbf{x} \in \mathbb{S}^{n-1}. \quad (103)$$

For each of these three models one can replicate the annealed and quenched calculations of Sections 2 and 4, under suitable technical hypotheses.

A note on non-spherical priors Finally, it is clear from the calculation of Appendix B (particularly Section B.1) that we can also generalize these techniques (at least heuristically) to non-spherical prior distributions on the vectors \mathbf{x} . The most natural hypothesis that allows the computation to be generalized is that the prior distribution takes the decoupled form $P(d\mathbf{x}) = \prod_i P(dx_i)$.