Statistical Physics of Computation 2025 - Exercises

Vittorio Erba, Emanuele Troiani

Week 9

This week we compute the spectrum of random matrices using the replica method. In particular, given a certain symmetric random matrix $M \in \mathbb{R}^{n\times n}$, we are interested in determining the probability density $\rho_M(\lambda)$ of the eigenvalues $\lambda_1, ...\lambda_n$

$$\rho_M(\lambda) = \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \delta(\lambda - \lambda_i) \tag{1}$$

Throughout the exercise we will assume that $\rho_M(\lambda)$ will concentrate to a bounded distribution, and that will be no issue in exchanging the limit $n \to \infty$ with any other one. For convenience we will denote with Tr[M] the trace of M, and with tr[M] = Tr[M]$/n$ the trace divided by the dimension.

9.1 Setting up a replica computation for the spectrum

We start by identifying the Stieltjes transform as the central object we want to compute, and detailing its relation to a replica computation

1. Define the Stieltjes transform $g_M(z)$ of the matrix M as the trace of the resolvent

$$g_M(z) = \mathrm{tr}\frac{1}{z\mathbb{I} - M}\,, \tag{2}$$

where by $1/M$ we mean the matrix inverse of M. We also recall this representation of the delta function given by the Sokhotski–Plemelj theorem

$$\delta(x) = \frac{1}{\pi} \lim_{\epsilon\to 0^+} \mathrm{Im}\,\frac{1}{x - i\epsilon}\,, \tag{3}$$

where Im denotes the imaginary part of the function. Obtain the following relation between $g_M(z)$ and $\rho_M(\lambda)$

$$\rho_M(\lambda) = \frac{1}{\pi} \lim_{\epsilon\to 0^+} \mathrm{Im}\, g_M(\lambda - i\epsilon)\,. \tag{4}$$

2. The previous point shows that knowing $g_M(z)$ is enough to obtain the spectrum, now we show that $g_M(z)$ is related to the determinant of the resolvent. Derive the expression

$$\log\det(z - M) = \sum_{i=1}^{n} \log(z\mathbb{I} - \lambda_i) \tag{5}$$

and use it to write

$$g_M(z) = n^{-1}\partial_z \log\det(z\mathbb{I} - M) \tag{6}$$

3. Turns out that the determinant is a quantity we can compute using the replica method. Argue why we can write

$$\det(z\mathbb{I} - M)^{-\frac{1}{2}} = \int e^{-\frac{z\|w\|^2}{2} + \frac{1}{2}\sum_{i,j=1}^{n} M_{ij} w_i w_j} \, \mathrm{d}w \tag{7}$$

where w is a n-dimensional vector, and w_i are its components. Show that one can write the average Stieltjes transform as

$$\mathbb{E}_M[g_M(z)] = -\frac{2}{n}\partial_z \lim_{r\to 0} \frac{\mathbb{E}_M[\mathcal{Z}^r] - 1}{r} . \tag{8}$$

What is $\mathcal{Z}$? Notice that contrary to Bayes-optimal computations, here there is no hidden signal, hence we have r replicas with $r \to 0$ and not $r+1$.

9.2 The replica method for the Wigner model

In this part we go through the replica method, computing the spectrum for the simplest class of matrices: the Gaussian Orthogonal Ensemble (GOE). Specifically, it's the ensemble of symmetric matrices with $M = \frac{1}{\sqrt{2n}}(G + G^\top)$, with $G_{ij} \sim \mathcal{N}(0,1)$. Apart from an overall scaling, and a difference in the diagonal that does not matter for $n \gg 1$, this is the noise matrix of the spiked-Wigner model.

We want to compute the moments of the partition function $\mathcal{Z}$

$$\mathcal{Z} = \int e^{\frac{z\|w\|^2}{2} - \frac{1}{2}\sum_{i,j=1}^{n} M_{ij} w_i w_j} \, \mathrm{d}w \, . \tag{9}$$

Notice this is slightly different (overall sign in the exponent) from what you got in the previous point. We will clarify this later.

1. Write the r-th moment of $\mathcal{Z}$ as

$$\mathbb{E}_M[\mathcal{Z}^r] = \int e^{\frac{z}{2}\sum_{a=1}^{r}\|w^a\|^2} \prod_{i,j}^{n} \mathbb{E}_{g\sim N(0,1)}\left[e^{-\frac{1}{\sqrt{2n}}\sum_{a=1}^{r} g w_i^a w_j^a}\right] \mathrm{d}w^a \tag{10}$$

where by $\mathrm{d}w^a$ we mean that we integrate over all the r replicas.

2. Show that

$$\prod_{i,j=1}^{n} \mathbb{E}_{g\sim N(0,1)}\left[e^{-\frac{1}{\sqrt{2n}}\sum_{a=1}^{r} g w_i^a w_j^a}\right] = e^{\frac{n}{4}\sum_{a,b=1}^{r}\left[\sum_{i=1}^{n}\frac{w_i^a w_i^b}{n}\right]^2} \tag{11}$$

3. Let's recap for a moment. At this point of the computation you should agree that we have this expression for the moment of the partition function

$$\mathbb{E}_M[\mathcal{Z}^r] = \int e^{\frac{zn}{2}\sum_{a=1}^{r}\frac{\|w^a\|^2}{n} + \frac{n}{4}\sum_{a,b=1}^{r}\left[\sum_{i=1}^{n}\frac{w_i^a w_i^b}{n}\right]^2} \mathrm{d}w^a \, . \tag{12}$$

It's the moment to introduce our order parameter q_{ab}

$$q_{ab} = \frac{w_i^a w_i^b}{n} \, . \tag{13}$$

Rewrite the moment as

$$\mathbb{E}_M[\mathcal{Z}^r] = \int e^{\frac{zn}{2}\sum_{a=1}^r q_{aa} + \frac{n}{4}\sum_{a,b=1}^r q_{ab}^2 + I_{\text{entropy}}(q)}\, \mathrm{d}q_{ab}\,. \tag{14}$$

Here we use the convention $\mathrm{d}q_{ab}$ to indicate the integral over $\{dq_{ab}\}_{1\leq a\leq b\leq r}$. How is $I_{\text{entropy}}(q)$ defined?

4. For the following, you can assume that

$$I_{\text{entropy}}(q) = \frac{n}{2}\log\det q\,. \tag{15}$$

You will derive this in a later part of the exercise. Make sure you agree that

$$\mathbb{E}_M[\mathcal{Z}^r] = \int e^{\frac{zn}{2}\operatorname{Tr} q + \frac{n}{4}\operatorname{Tr} q^2 + \frac{n}{2}\log\det q}\, \mathrm{d}q_{ab}\,. \tag{16}$$

Since q is symmetric we can decompose it into a diagonal matrix of eigenvalues Λ and a rotation matrix O as $q = O^\top \Lambda O$. Show that the argument of the integral only depends on Λ. Then, show that

$$\mathbb{E}_M[\mathcal{Z}^r] = \int e^{\frac{n}{2}\sum_{a=1}^r \left[z\Lambda_a + \frac{\Lambda_a^2}{2} + \log\Lambda_a\right] + \sum_{a<b}\log|\Lambda_a - \Lambda_b|}\, \mathrm{d}\Lambda\,, \tag{17}$$

using (and not deriving) Weyl's integration formula

$$\mathrm{d}q = \mu(O)\prod_{a<b}|\Lambda_a - \Lambda_b|\,\mathrm{d}\Lambda\,\mathrm{d}O\,, \tag{18}$$

where $\mathrm{d}O$ means integration over $n\times n$ rotation matrices O, and $\mu(O)$ is the uniform measure over rotation matrices.

5. This is almost an expression where all eigenvalues Λ are independent, except for the term that came from the change of variables. It's reasonable to drop it in a first approximation for $n \gg 1$. Let's compute the integral by the saddle-point method. Show that the saddle point equation for each Λ_a equals

$$-z = \Lambda_a + \frac{1}{\Lambda_a} \tag{19}$$

6. The equation above is not a bad approximation. It is in fact exact in the large n limit. Show that without dropping the logarithm term the saddle point equation is for each a

$$z + \Lambda_a + \frac{1}{\Lambda_a} + \frac{1}{n}\sum_{b\neq a}\frac{1}{\Lambda_a - \Lambda_b} = 0\,. \tag{20}$$

Now let's stare at this equation for a moment. If the last piece is negligible (as we want to claim) it means that all the Λ_a will be identical. This will however make the last fraction diverge, so the answer can't be that simple! In fact, what we should look at is an ansatz of this form

$$\Lambda_a = \Lambda^* + \frac{\epsilon_a}{\sqrt{n}} \tag{21}$$

where Λ^* is a solution of the saddle point equation without the extra piece. What we are saying here is that at finite n the Λ_a will be different, and in the limit $n\to\infty$ they will have a typical distance of the order $1/\sqrt{n}$. Show that the ansatz of the scaling $1/\sqrt{n}$ is consistent, and derive (but do not solve) the equation satisfied by the corrections ϵ_a. Finally, argue that the correction term is truly subleading.

7. We now want to connect this replica computation with (8). Argue that the $\mathcal{Z}$ in (9) can for our purposes replace the $\det(z-M)^{-\frac{1}{2}}$ defined in (7).

8. Show that $\mathbb{E}_M[g_M(z)]$ (for convenience just g from now onward) is a solution of the equation

$$z = g + \frac{1}{g} \tag{22}$$

9. We are at the end of the computation. Solve $g(z)$ in terms of z. To choose the right solution look at the definition of $g(z)$ for very large $|z|$.

10. Obtain finally the spectral distribution $\rho_M(\lambda)$

$$\rho_M(\lambda) = \mathbf{1}_{|\lambda|<2} \frac{\sqrt{4-\lambda^2}}{2\pi} \tag{23}$$

where $\mathbf{1}_{|\lambda|<2}$ is the indicator function. Notice that the distribution is a semicircle of radius 2. This should be familiar to you from exercise 4.

9.3 Computation of entropic part

Here we compute the entropic piece

$$I_{\text{entropy}}(q) = \log\left[\int \prod_{a\leq b}^{r} \delta\left(nq_{ab} - \sum_{i=1}^{n} w_i^a w_i^b\right) \mathrm{d}w^a\right]. \tag{24}$$

We will disregard all the 2π factors, as they will be completely irrelevant in the result.

1. Introduce the Lagrange multiplier matrix $\hat{q}$ to rewrite the integral in the log as

$$I_{\text{entropy}}(q) = \int \mathrm{d}\hat{q}\, \exp\left\{n\sum_{a\leq b}^{r} \hat{q}_{ab}q_{ab}\right\} \prod_{i=1}^{n} \int \exp\left\{-\sum_{a\leq b}^{r} \hat{q}_{ab} w_i^a w_i^b\right\} \mathrm{d}w^a \tag{25}$$

2. What we have above is almost a Gaussian integral, except that in the exponent the sums run over $a \leq b$. Show that one can write

$$I_{\text{entropy}}(q) = \int \mathrm{d}\hat{q}\, \exp\left\{\frac{n}{2}\sum_{a,b}^{r} \bar{q}_{ab}q_{ab}\right\} \prod_{i=1}^{n} \int \exp\left\{-\frac{1}{2}\sum_{a,b}^{r} \bar{q}_{ab} w_i^a w_i^b\right\} \mathrm{d}w^a \tag{26}$$

where $\bar{q}_{ab} = (1+\delta_{ab})\hat{q}_{ab}$.

3. Perform the integral over the weights to have

$$I_{\text{entropy}}(q) = \int \mathrm{d}\hat{q}\, \exp\left\{\frac{n}{2}\sum_{a,b}^{r} \bar{q}_{ab}q_{ab} - \frac{n}{2}\log\det\bar{q}\right\} \tag{27}$$

4. Solve the remaining integral using a saddle point method, obtaining the result

$$I_{\text{entropy}}(q) = \frac{n}{2}\log\det q \tag{28}$$

up to constant additive terms. It is useful to know that for any matrix M, then $\partial_{M_{ij}}[\log\det M] = [M^{-1}]_{ij}$.