Statistical Physics of Computation - Exercises

Emanuele Troiani, Vittorio Erba, Yizhou Xu

September 2024

Week 9

9.1 Replica computation for the spiked-Wigner model

We are given as observation a $N \times N$ symmetric matrix Y created as

$$\boldsymbol{Y} = \sqrt{\frac{\lambda}{N}} \underbrace{\boldsymbol{x}^* \boldsymbol{x}^{*\intercal}}_{N\times N \text{ rank-one matrix}} + \underbrace{\boldsymbol{\xi}}_{\text{symmetric iid noise}}$$

where $\boldsymbol{x}^* \in \mathbb{R}^N$ with $x_i^* \overset{\text{i.i.d.}}{\sim} P_X(x)$, $\xi_{ij} = \xi_{ji} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$ for $i \leq j$. We keep the prior P_0 generic, as long as it is factorized over all the components of the vector $\boldsymbol{x}$.

This is called the spiked-Wigner model in statistics. The name "Wigner" refer to the fact that Y is a Wigner matrix (a symmetric random matrix with components sampled randomly from a Gaussian distribution) plus a "spike", that is a rank one matrix $\boldsymbol{x}^* \boldsymbol{x}^{*\intercal}$. We will use it as an example of recovering a signal with low rank structure corrupted by noise which we want to clean up.

Our task shall be to recover the vector $\boldsymbol{x}$ from the knowledge of $\boldsymbol{Y}$, the signal-to-noise ratio λ and the prior P_0. As we saw during the lecture, this can be achieved using the posterior estimation, i.e. by computing the posterior distribution and evaluating some statistics over it.

In this exercise, you will compute the normalization factor of the posterior distribution, i.e. the partition function of the problem, and derive the state equation for the overlap order parameter.

1. Show that the posterior distribution $P(\boldsymbol{x}|\boldsymbol{Y})$ for the problem can be written as

$$P(\boldsymbol{x}|\mathbf{Y}) = \frac{1}{Z(\mathbf{Y})} \left[\prod_{i=1}^{N} P_X(x_i)\right] \left[\prod_{i\leq j} \frac{e^{-\frac{\lambda}{2N} x_i^2 x_j^2 + \sqrt{\frac{\lambda}{N}} x_i x_j y_{ij}}}{\sqrt{2\pi}}\right] \tag{1}$$

   for a specific $Z(\mathbf{Y})$. How is $Z(\mathbf{Y})$ defined for this measure?

   One needs to recognize that the output channel distribution satisfies, for all $i \leq j$, $P_{\text{out}}(y_{ij}|\boldsymbol{x}) = N(y_{ij}, \sqrt{\lambda/N} x_i x_j, 1)$, so that

$$P(\boldsymbol{x}|\mathbf{Y}) = \frac{1}{Z(\mathbf{Y})} \left[\prod_{i=1}^{N} P_X(x_i)\right] \left[\prod_{i\leq j} \frac{e^{-\frac{1}{2}\left(y_{ij} - \sqrt{\frac{\lambda}{N}} x_i x_j\right)^2}}{\sqrt{2\pi}}\right] \tag{2}$$

where we notice in particular that the product in the Gaussian factor runs over $i \le j$ due to the symmetry of the problem. It's convenient to reabsorb all pieces independent of $\boldsymbol{x}$ from the exponential in the normalisation, allowing us to write

$$P(\boldsymbol{x}|\mathbf{Y}) = \frac{1}{Z(\mathbf{Y})} \left[\prod_{i=1}^{N} P_X(x_i) \right] \left[\prod_{i \le j} \frac{e^{-\frac{\lambda}{2N} x_i^2 x_j^2 + \sqrt{\frac{\lambda}{N}} x_i x_j y_{ij}}}{\sqrt{2\pi}} \right] \tag{3}$$

where

$$Z(\mathbf{Y}) = \int \mathrm{d}\boldsymbol{x} \left[\prod_{i=1}^{N} P_X(x_i) \right] \left[\prod_{i \le j} \frac{e^{-\frac{\lambda}{2N} x_i^2 x_j^2 + \sqrt{\frac{\lambda}{N}} x_i x_j y_{ij}}}{\sqrt{2\pi}} \right] \tag{4}$$

We are interested in computing the averaged free entropy associated to the posterior distribution, i.e.

$$\lim_{N \to \infty} \mathbb{E}_{\boldsymbol{Y}} \left[\frac{1}{N} \log Z(\boldsymbol{Y}) \right]$$

which we can compute using the replica method.

2. Show that the averaged replicated partition function equals

$$\mathbb{E}_{\mathbf{Y}}[Z^n] = \int \mathrm{d}\boldsymbol{Y} e^{-\frac{1}{2}\sum_{i \le j} y_{ij}^2} \prod_{\alpha=0}^{n} \int \mathrm{d}\boldsymbol{x}^{(\alpha)} \left(\prod_{i=1}^{N} P_X\left(x_i^{(\alpha)}\right) \right) \left(\prod_{i \le j} \frac{e^{-\frac{\lambda}{2N}\left(x_i^{(\alpha)}\right)^2 \left(x_j^{(\alpha)}\right)^2 + \sqrt{\frac{\lambda}{N}} x_i^{(\alpha)} x_j^{(\alpha)} y_{ij}}}{\sqrt{2\pi}} \right) \tag{5}$$

where you should notice that we are taking the product $n+1$ replicas.

We use the identity

$$\mathbb{E}_{\boldsymbol{Y}}[F[\mathbf{Y}]] = \int \mathrm{d}\boldsymbol{Y} P(\mathbf{Y}) F[\mathbf{Y}] = \int \mathrm{d}\boldsymbol{Y} \int \mathrm{d}\boldsymbol{x}^* P_{\text{out}}(\mathbf{Y}|\boldsymbol{x}^*) P_X(\boldsymbol{x}^*) F[\mathbf{Y}] = \int \mathrm{d}\boldsymbol{Y} Z[\mathbf{Y}] e^{-\frac{1}{2}\sum_{i \le j} y_{ij}^2} F[\mathbf{Y}] \tag{6}$$

where the last step crucially depends on the fact that we reabsorbed all x-independent terms of the posterior in the definition of the partition function. Then, one has

$$\mathbb{E}_{\mathbf{Y}}[Z[\mathbf{Y}]^n] = \int \mathrm{d}\boldsymbol{Y} e^{-\frac{1}{2}\sum_{i \le j} y_{ij}^2} Z[\mathbf{Y}]^{n+1} \tag{7}$$

and one gets the result by plugging in the definition of the partition function.

3. Integrate over the disorder, i.e. the observation Y, to get at leading order in N

$$\mathbb{E}_{\mathbf{Y}}[Z^n] = \int \prod_{\alpha, i} P_X\left(x_i^{(\alpha)}\right) \mathrm{d}x_i^{(\alpha)} \ \exp\left(\frac{\lambda N}{2} \sum_{\alpha < \beta} \left(\sum_i \frac{x_i^{(\alpha)} x_i^{(\beta)}}{N} \right)^2 \right) \tag{8}$$

$$
\begin{aligned}
\mathbb{E}_{\mathbf{Y}}\left[Z^n\right] &= \int \mathrm{d}\mathbf{Y} \prod_{\alpha=0}^{n} \int \mathrm{d}\boldsymbol{x}^{(\alpha)} \prod_i P_X\left(x_i^{(\alpha)}\right) (2\pi)^{-\frac{N(N+1)}{2}} \exp\left(\sum_{\alpha=0}^{n} \sum_{i\le j}\left[-\frac{\lambda}{2N} {x_i^{(\alpha)}}^2 {x_j^{(\alpha)}}^2 + \sqrt{\frac{\lambda}{N}} y_{ij} x_i^{(\alpha)} x_j^{(\alpha)}\right]\right) \\
&= \int \prod_{\alpha,i} P_X\left(x_i^{(\alpha)}\right) \mathrm{d}x_i^{(\alpha)} \ \exp\left(\sum_{i\le j}\left[-\frac{\lambda}{2N}\sum_\alpha {x_i^{(\alpha)}}^2 {x_j^{(\alpha)}}^2\right]\right) \underbrace{\prod_{i\le j}\int \mathrm{d}y_{ij}\, \frac{e^{-\frac{y_{ij}^2}{2} + y_{ij}\left(\sqrt{\frac{\lambda}{N}}\sum_\alpha x_i^{(\alpha)} x_j^{(\alpha)}\right)}}{\sqrt{2\pi}}}_{\overset{(a)}{=} \exp\left(\frac{\lambda}{2N}\sum_{i\le j}\sum_{\alpha,\beta} x_i^{(\alpha)} x_j^{(\alpha)} x_i^{(\beta)} x_j^{(\beta)}\right)} \\
&\overset{(b)}{=} \int \prod_{\alpha,i} P_X\left(x_i^{(\alpha)}\right) \mathrm{d}x_i^{(\alpha)} \ \exp\left(-\frac{\lambda N}{4}\sum_\alpha\left(\sum_i \frac{{x_i^{(\alpha)}}^2}{N}\right)^2 + \frac{\lambda N}{4}\sum_{\alpha,\beta}\left(\sum_i \frac{x_i^{(\alpha)} x_i^{(\beta)}}{N}\right)^2\right) \\
&= \int \prod_{\alpha,i} P_X\left(x_i^{(\alpha)}\right) \mathrm{d}x_i^{(\alpha)} \ \exp\left(\frac{\lambda N}{2}\sum_{\alpha<\beta}\left(\sum_i \frac{x_i^{(\alpha)} x_i^{(\beta)}}{N}\right)^2\right)
\end{aligned}
$$

(a) uses the fact that $\int \mathcal{D}z\ e^{az} = e^{a^2/2}$

(b) uses the fact that

$$
\sum_{i\le j} \frac{a_i}{N}\frac{a_j}{N} = \frac{1}{2}\left(\sum_i \frac{a_i}{N}\right)^2 + \frac{1}{2}\sum_i \frac{a_i^2}{N^2}
$$

and neglect the second term since it scales like $\mathrm{O}(N^{-1})$.

4. Introduce the appropriate order parameters and obtain

$$
\mathbb{E}_{\mathbf{Y}}[Z^n] = \int \prod_{\alpha<\beta} \mathrm{d}\hat{q}_{\alpha\beta}\, \mathrm{d}q_{\alpha\beta}\ \exp\left(N I_{\text{energy}}(q_{\alpha\beta}, \hat{q}_{\alpha\beta}) + N I_{\text{entropy}}(\hat{q}_{\alpha\beta})\right) \tag{9}
$$

where we defined

$$
I_{\text{energy}}(q_{\alpha\beta}, \hat{q}_{\alpha\beta}) = \frac{\lambda}{2}\sum_{\alpha<\beta} q_{\alpha\beta}^2 - \sum_{\alpha<\beta} q_{\alpha\beta}\hat{q}_{\alpha\beta} \tag{10}
$$

and

$$
I_{\text{entropy}}(\hat{q}_{\alpha\beta}) = \log\left(\int \prod_\alpha P_X\left(x_\alpha\right) \mathrm{d}x_\alpha\ \exp\left\{\sum_{\alpha<\beta} \hat{q}_{\alpha\beta} x_\alpha x_\beta\right\}\right) \tag{11}
$$

where we stress that here the integral over dx_α runs over the real numbers for all $\alpha = 0, \ldots, n$.

$$\mathbb{E}_{\mathbf{Y}}\left[Z^n\right] \overset{(a)}{=} \int \prod_{\alpha,i} P_X\left(x_i^{(\alpha)}\right) \mathrm{d}x_i^{(\alpha)} \int \prod_{\alpha<\beta} \delta\left(N q_{\alpha\beta} - \sum_i x_i^{(\alpha)} x_i^{(\beta)}\right) \mathrm{d}q_{\alpha\beta} \ \exp\left(\frac{\lambda N}{2} \sum_{\alpha<\beta} q_{\alpha\beta}^2\right)$$

$$\overset{(b)}{=} \int \prod_{\alpha,i} P_X\left(x_i^{(\alpha)}\right) \mathrm{d}x_i^{(\alpha)} \int \prod_{\alpha<\beta} e^{-\hat{q}_{\alpha\beta} N q_{\alpha\beta} + \hat{q}_{\alpha\beta} \sum_i x_i^{(\alpha)} x_i^{(\beta)}} \mathrm{d}\hat{q}_{\alpha\beta}\, \mathrm{d}q_{\alpha\beta} \ \exp\left(\frac{\lambda N}{2} \sum_{\alpha<\beta} q_{\alpha\beta}^2\right)$$

$$\overset{(c)}{=} \int \prod_{\alpha<\beta} \mathrm{d}\hat{q}_{\alpha\beta}\, \mathrm{d}q_{\alpha\beta} \ \exp\left(\frac{\lambda N}{2} \sum_{\alpha<\beta} q_{\alpha\beta}^2 - N \sum_{\alpha<\beta} q_{\alpha\beta} \hat{q}_{\alpha\beta}\right)$$

$$\left\{\int \prod_{\alpha} P_X\left(x_\alpha\right) \mathrm{d}x_\alpha \ \exp\left(\sum_{\alpha<\beta} \hat{q}_{\alpha\beta} x_\alpha x_\beta\right)\right\}^N$$

(a) partitions the huge integral according to overlap between two distinct replicas $q_{\alpha\beta}$ with definitions

$$q_{\alpha\beta} = \frac{1}{N} \sum_i x_i^{(\alpha)} x_i^{(\beta)}, \quad \forall\, \alpha < \beta$$

(b) introduces Fourier representation of the Dirac's delta

(c) change the order of integral and expectation. Moreover, $x_i^{(\alpha)}$ are iid for each i, the tuple $x_i^{(0)}, x_i^{(1)}, \cdots, x_i^{(n)}$ are identical distributed, so we switch to subscript notation $x_{(0)}, x_{(1)}, \cdots, x_{(n)}$ to get rid of component index i but keep the replica index α. This leads to the power N around the curly brackets.

5. Which replica ansatz should you impose in this computation?

It makes sense to use the Replica Symmetric one, as we are considering the Bayes optimal setting. We pick $q_{\alpha\beta} = q$, $\hat{q}_{\alpha\beta} = \hat{q}$.

6. Show that in the ansatz you discussed in the previous point, the energetic term satisfies

$$\lim_{n\to 0} \frac{1}{n} I_{\text{energy}}(q_{\alpha\beta}, \hat{q}_{\alpha\beta}) = \frac{\lambda}{4} q^2 - \frac{q\hat{q}}{2} \tag{12}$$

for q and $\hat{q}$ real numbers.

We need to expand I_{energy} for small n. Recall that $q_{\alpha\beta}$ is a square matrix of size $n+1$.

$$\begin{aligned} I_{\text{energy}}(q_{\alpha\beta}, \hat{q}_{\alpha\beta}) =& \frac{\lambda}{2} \sum_{\alpha<\beta} q_{\alpha\beta}^2 + \sum_{\alpha<\beta} q_{\alpha\beta} \hat{q}_{\alpha\beta} \\ =& \frac{\lambda}{2} \frac{(n+1)n}{2} q^2 - \frac{(n+1)n}{2} q\hat{q} \\ \approx& \frac{\lambda n}{4} q^2 - \frac{nq\hat{q}}{2} \end{aligned}$$

7. Show that in the ansatz you discussed in the previous point, the entropic term satisfies

$$I_{\text{entropy}}(\hat{q}_{\alpha\beta}) = \log\left(\int Dz \left[\left(\int P_X\left(x\right) \mathrm{d}x \ \exp\left\{-\frac{\hat{q}}{2} x^2 + \sqrt{\hat{q}} x z\right\}\right)^{n+1}\right]\right) \tag{13}$$

for $\hat{q}$ real numbers.

We have

$$
\begin{aligned}
I_{\text{entropy}}(\hat{q}_{\alpha\beta}) &= \log\left(\int \prod_\alpha P_X\left(x_\alpha\right) \mathrm{d}x_\alpha \ \exp\left\{-\sum_{\alpha<\beta} \hat{q}_{\alpha\beta} x_\alpha x_\beta\right\}\right) \\
&= \log\left(\int \prod_\alpha P_X\left(x_\alpha\right) \mathrm{d}x_\alpha \ \exp\left\{-\hat{q}\sum_{\alpha<\beta} x_\alpha x_\beta\right\}\right) \\
&= \log\left(\int \prod_\alpha P_X\left(x_\alpha\right) \mathrm{d}x_\alpha \ \exp\left\{-\frac{\hat{q}}{2}\sum_{\alpha\neq\beta} x_\alpha x_\beta\right\}\right) \\
&= \log\left(\int \prod_\alpha P_X\left(x_\alpha\right) \mathrm{d}x_\alpha \ \exp\left\{-\frac{\hat{q}}{2}\sum_{\alpha} x_\alpha^2 + \frac{\hat{q}}{2}\left(\sum_\alpha x_\alpha\right)^2\right\}\right) \\
&= \log\left(\int \mathcal{D}z \int \prod_\alpha P_X\left(x_\alpha\right) \mathrm{d}x_\alpha \ \exp\left\{\sum_{\alpha}\left(-\frac{\hat{q}}{2} x_\alpha^2 + \sqrt{\hat{q}}x_\alpha z\right)\right\}\right) \\
&= \log\left(\int \mathcal{D}z \left(\int P_X\left(x\right) \mathrm{d}x \ \exp\left\{-\frac{\hat{q}}{2} x^2 + \sqrt{\hat{q}}xz\right\}\right)^{n+1}\right)
\end{aligned}
$$

where in the fifth passage we used the Hubbard stratonovish transformation (see previous replica computation).

8. Show that in the small n limit, the entropic term satisfies

$$\lim_{n\to 0}\frac{1}{n} I_{\text{entropy}}(\hat{q}_{\alpha\beta}) = \int Dz\, I(\hat{q},z)\log\left(I(\hat{q},z)\right) \tag{14}$$

where we defined $I(\hat{q},z)$ as

$$I(\hat{q},z) = \int P_X\left(x\right)\mathrm{d}x \ \exp\left\{-\frac{\hat{q}}{2}x^2 + \sqrt{\hat{q}}xz\right\} \tag{15}$$

and that it can be equivalently rewritten as

$$\lim_{n\to 0}\frac{1}{n} I_{\text{entropy}}(\hat{q}_{\alpha\beta}) = \int Dz \int P_X\left(x_0\right)\mathrm{d}x_0 \log\left(\int P_X\left(x\right)\mathrm{d}x \ \exp\left\{-\frac{\hat{q}}{2}x^2 + \sqrt{\hat{q}}xz + \hat{q}xx_0\right\}\right) \tag{16}$$

We will need the following identity: for a random quantity X with mean one, in the small n limit we have

$$\mathbb{E}[X^{n+1}] \sim \exp(n\mathbb{E}[X\log(X)]) \tag{17}$$

$$
\begin{aligned}
\mathbb{E}[X^{n+1}] &= \mathbb{E}\left[Xe^{n\log(X)}\right] \sim \mathbb{E}[X + nX\log(X)] = \mathbb{E}[X] + n\mathbb{E}[X\log(X)] = 1 + n\mathbb{E}[X\log(X)] \\
&= \exp(\log(1+n\mathbb{E}[X\log(X)])) \sim \exp(n\mathbb{E}[X\log(X)])
\end{aligned}
$$

Now we notice that

$$\int \mathcal{D}z I(\hat{q}, z) = \int \mathcal{D}z \int P_X(x)\, \mathrm{d}x \;\exp\left\{-\frac{\hat{q}}{2}x^2 + \sqrt{\hat{q}}xz\right\} = \int P_X(x)\, \mathrm{d}x = 1$$

which gives us the first result. To get the final result one needs to do a bit of manipulations. First, we notice that

$$\exp\left\{-\frac{1}{2}z^2\right\} I(\hat{q}, z) = \int P_X(x)\, \mathrm{d}x \exp\left\{-\frac{1}{2}\left(\sqrt{\hat{q}}x - z\right)^2\right\} \tag{18}$$

Finally, after a change of variable we have

$$z \to t = z - \sqrt{\hat{q}}x_0 \tag{19}$$

where we indicate with x_0 the integration variable inside the first $I(\hat{q}, z)$ (the one outside of the logarithm). This gives us

$$\int Dz\, I(\hat{q}, z) \log\left(I(\hat{q}, z)\right) = \int Dt \int P_X(x_0)\, \mathrm{d}x_0\, I(\hat{q}, z) \log\left(I(\hat{q}, t + \sqrt{\hat{q}}x_0)\right) \tag{20}$$

which is exactly the result we want.

9. Argue finally that the free entropy

$$\phi = \lim_{N\to\infty} \mathbb{E}_{\boldsymbol{Y}}\left[\frac{1}{N}\log Z(\boldsymbol{Y})\right] \tag{21}$$

can be expressed as

$$\phi = \mathrm{extr}_{q,\hat{q}}\left[\frac{\lambda}{4}q^2 - \frac{q\hat{q}}{2} + \int Dz \int P_X(x_0)\, \mathrm{d}x_0 \log\left(\int P_X(x)\, \mathrm{d}x \;\exp\left\{-\frac{\hat{q}}{2}x^2 + \sqrt{\hat{q}}xz + \hat{q}xx_0\right\}\right)\right] \tag{22}$$

or equivalently as

$$\phi = \mathrm{extr}_{q}\left[-\frac{\lambda}{4}q^2 + \int Dz P_X(x_0)\, \mathrm{d}x_0 \log\left(\int P_X(x)\, \mathrm{d}x \;\exp\left\{-\frac{\lambda q}{2}x^2 + \sqrt{\lambda q}xz + \lambda q x x_0\right\}\right)\right] \tag{23}$$

The first expression just combines the previous results and uses the replica trick

$$\phi \approx \frac{\mathbb{E}Z(Y)^n - 1}{Nn}\,. \tag{24}$$

The second expression comes from taking the fixed point on q, which gives $\hat{q} = \lambda q$, and substituting it back in.