

Statistical Physics of Computation - Exercises

Emanuele Troiani, Vittorio Erba, Yizhou Xu

September 2024

Week 14

Important: The aim of the last exercise session is to give you an idea of what kind of questions we may ask at the exam, and what topics to focus on in your studying. Please note that:

- All the material of the lectures and exercise session is subject to examination, both concerning techniques and the phenomenology of the example problems we studied in class.
- The questions presented in this document cover just a selection of the topics you should know and of the kinds of questions you may be asked at the exam.
- The exam will include also step-by-step exercises of the same style as the various exercise sessions, even though they are not included in this document.

-
1. Consider a Bayesian inference problem with given prior, output channel, and posterior. State and prove the most general version of Nishimori's condition.

Let $w^* \sim P_{\text{prior}}(w)$ be a sample from the prior distribution and $w_1, w_2 \sim P_{\text{posterior}}(w|y)$ independent samples over the posterior distribution on the samples $y \sim P_{\text{out}}(y|w)$, Nishimori's condition states that

$$\mathbb{E}_{y \sim P_{\text{out}}} \mathbb{E}_{w^* \sim P_{\text{prior}}} \mathbb{E}_{w_1 \sim P_{\text{posterior}}} f(w_1, w^*) = \mathbb{E}_{y \sim P_{\text{out}}} \mathbb{E}_{w^* \sim P_{\text{prior}}} \mathbb{E}_{w_1, w_2 \sim P_{\text{posterior}}} f(w_1, w_2) \quad (1)$$

The proof is a simple application of Bayes formula:

$$\mathbb{E}_{y \sim P_{\text{out}}} \mathbb{E}_{w_2 \sim P_{\text{posterior}}} f(w_1, w_2) = \int dy \int dw_2 P_{\text{posterior}}(w_2|y) P_{\text{out}}(y) f(w_1, w_2) = \quad (2)$$

$$\int dy \int dw_2 P_{\text{prior}}(w_2) P_{\text{out}}(y|w_2) f(w_1, w_2) = \mathbb{E}_{y \sim P_{\text{out}}} \mathbb{E}_{w_2 \sim P_{\text{prior}}} f(w_1, w_2) \quad (3)$$

We can then change the name w_2 in w^* to obtain the result.

2. Consider a thermodynamic system with partition function

$$Z(\beta) = \sum_{s_1, \dots, s_N} e^{\beta H(s)} \quad (4)$$

where each variable s_i can assume values in a discrete set. Define the free entropy, and explain how to compute the entropy density as a function of the energy density $s(e)$ in the limit $N \rightarrow \infty$. Feel free to assume that all thermodynamic potentials/energy/entropy are invertible/well behaved if you need to assume that.

The free entropy is defined as $\phi(\beta) = \lim_{N \rightarrow \infty} N^{-1} \log Z(\beta)$. By the series of equalities

$$Z(\beta) = \sum_{s_1, \dots, s_N} e^{\beta H(s)} = \int de e^{N(\beta e + s(e))} = e^{N(\beta e_*(\beta) + s(e_*(\beta)))} \quad (5)$$

where $e_*(\beta)$ maximizes the function $e \rightarrow \beta e + s(e)$ at fixed β and where $s(e)$ is the entropy density at energy density e , we have

$$\phi(\beta) = \beta e_*(\beta) + s(e_*(\beta)). \quad (6)$$

Moreover, we have

$$\partial_\beta \phi(\beta) = e_*(\beta) \quad (7)$$

from which we get

$$s(e_*(\beta)) = \phi(\beta) - \beta \partial_\beta \phi(\beta). \quad (8)$$

Finally, inverting $e_*(\beta)$ (we here assume that this is possible) to express β as a function of e allows to compute $s(e)$ (this is equivalent to plot parametrically in β the curve $(e_*(\beta), s(e_*(\beta)))$).

3. Consider a constraint satisfaction problem with N variables and P constraints, and assume that the correct thermodynamic scaling for the problem is $P = \alpha N$. Define the SAT/UNSAT transition point.

The SAT/UNSAT transition point is the value of α_c such that for $\alpha < \alpha_c$ the CSP has exponentially many solutions (positive entropy) and for $\alpha > \alpha_c$ the CSP has no solution.

4. You are given a generic computational problem that is amenable to replica analysis. Describe (and motivate where appropriate) the main steps of the associated with the replica computation.

- Write down a partition function Z /probability measure relevant to the problem. This allows to set a statistical mechanics treatment of the problem, with the aim of obtaining a low-dimensional characterization of the problem in the thermodynamic limit.
- Use the replica trick to translate the average of the log of Z to the average of the integer moments of Z , as the average of the log is intractable analytically, while the average of the integer moments can be done easily.
- Average over the disorder.
- Introduce the overlap order parameters.
- Impose an appropriate ansatz on the overlap order parameter, as a concrete way to perform the analytic continuation from integer to zero replicas.
- Take the zero replica limit.
- Derive the state equations as the saddle-point condition of the resulting integral.

5. Describe the link between the overlap order parameter and the averaged overlap distribution of a given disordered system.

The overlap order parameter is linked to the averaged overlap distribution by the condition

$$\mathbb{E}_J \text{Prob}(q(\sigma, \tau) = q) = \lim_{n \rightarrow 0} \frac{2}{n(n-1)} \sum_{a < b} \delta(q - Q_{ab}) \quad (9)$$

where σ and τ are independent draws from the Gibbs distribution associated to the system studied with disorder J , and Q_{ab} is the $n \times n$ overlap matrix order parameter arising in replica computations.

6. Consider a regression problem in which you generated data/label pairs $(x, y) \sim P_{\text{data}}$ with $x \in \mathbb{R}^d$ and $y \in \mathbb{R}$. A predictor/student function $f_w(x)$ parametrized by some weights w has been trained on a dataset. Write down the definition of the generalization error of the student. Use the mean-square-error in label space to assess the performance of the student.

$$E_{\text{gen}}(f_w) = \mathbb{E}_{(x_{\text{new}}, y_{\text{new}})} \|f_w(x_{\text{new}}) - y_{\text{new}}\|^2 \quad (10)$$

7. Consider a teacher-student ERM problem amenable to replica computations in the RS ansatz. What are the fundamental order parameters describing the physics of the problem?

They are the replica-replica overlap, telling us how close independent replicas of the system are, and the replica-teacher overlap, telling us how similar a sample from the Gibbs measure is to the teacher.

8. Consider a probability measure $p(w)$ for a student vector w , where we aim at recovering some hidden teacher vector w^* . Show that the MSE

$$\|\hat{w} - w^*\|^2$$

of the average estimator $\hat{w}_{\text{mean}} = \mathbb{E}_{w \sim p(w)}$ is 1/2 times the average MSE of the Gibbs estimator $\hat{w}_{\text{Gibbs}} \sim p(w)$.

The MSE of the Gibbs estimator is

$$\mathbb{E}_{w \sim p(w)} \|w - w^*\|^2 \quad (11)$$

We can expand the square to have

$$\mathbb{E}_{w \sim p(w)} \|w - w^*\|^2 = \mathbb{E}_{w \sim p(w)} \|w\|^2 + \|w^*\|^2 - 2\mathbb{E}_{w \sim p(w)} w \cdot w^* = \|w^*\|^2 - \hat{w} \cdot w^* \quad (12)$$

where in the last step we used Nishimori. Doing the same on the MSE of the average estimator we get

$$\|\hat{w} - w^*\|^2 = \|\hat{w}\|^2 + \|w^*\|^2 - 2\hat{w} \cdot w^* = 2(\|w^*\| - \hat{w} \cdot w^*) \quad (13)$$

which is exactly the relation we want.

9. Give the definition of computational hard phase in a Bayes optimal inference problem. You can suppose that the phase diagram of the problem is given as a function of some signal-to-noise ratio (SNR) parameter, and you can assume that AMP is the best efficient algorithm for the problem at hand.

An hard phase is an interval on the SNR axis where the performance of the BO estimator is strictly larger than the performance of AMP.

10. Is it possible for a Bayes optimal inference problem to always be in an impossible/hard phase? If so, provide an example.

There is no reason preventing the absence of easy phases, where the BO estimator has non-zero performance but AMP is stuck to either zero or strictly lower performance than the BO estimator. We saw during the course that the spiked-tensor model has this property.

11. Consider the factorized probability distribution

$$p(x, y, z, w) = f_1(x, y)f_2(z, w). \quad (14)$$

Is the corresponding factor graph connected? Justify your answer.

The factor graph is not connected as there is no interaction term between the variables $\{x, y\}$ and the variables $\{z, w\}$.

12. Under what condition on the factor graph is BP exact? Under what condition on the factor graph we can conjecture that BP will be exact in the thermodynamic limit?

BP is exact on tree (loop-free) factor graphs. In the thermodynamic limit, if the factor graph is locally tree-like, meaning that typical loops are of length that diverges with the system size, we can conjecture that BP will apply.

13. Suppose that a given disordered system is replica symmetric, and has replica symmetric overlap q_{RS} . Compute the averaged overlap distribution

$$\mathbb{E}_J \text{Prob}(q(\sigma, \tau) = q) \quad (15)$$

where σ and τ are independent draws from the Gibbs distribution associated to the system.

The overlap order parameter is linked to the averaged overlap distribution by the condition

$$\mathbb{E}_J \text{Prob}(q(\sigma, \tau) = q) = \lim_{n \rightarrow 0} \frac{2}{n(n-1)} \sum_{a < b} \delta(q - Q_{ab}). \quad (16)$$

In the RS case, $Q_{ab} = q_{\text{RS}}$ for all $a < b$ giving

$$\mathbb{E}_J \text{Prob}(q(\sigma, \tau) = q) = \delta(q - q_{\text{RS}}). \quad (17)$$

14. What is the assumption that we implicitly use whenever we average the free entropy over the disorder, and then claim that the averaged free entropy describes also the properties of specific disorder realizations?

We assume that in the thermodynamic limit the free entropy is self-averaging, i.e. that it is not a random variable anymore (where the randomness comes from the variability of the disorder). This means that with probability going to one over the disorder realizations, the free entropy of a specific disorder realization is the same as the averaged free entropy.