Statistical Physics of Computation - Exercises

Emanuele Troiani, Vittorio Erba, Yizhou Xu

September 2024

Week 5

5.1 The p-spin ferromagnet

Let's look at a generalisation Curie–Weiss model where more than two spins are coupled. Consider a system with N binary variables $s_i \in \{-1, 1\}$, $1 \le i \le N$ with energy function

$$\mathcal{H}^p[\mathbf{s}] = -\frac{1}{p! N^\alpha} \sum_{i_1, i_2, \dots, i_p} s_{i_1} s_{i_2} \dots s_{i_p} - \frac{h}{N^\gamma} \sum_i s_i \,. \tag{1}$$

1. (1 pt) For which values of α, γ does the model have a well posed large N limit (as in both pieces are of the same order, and the energy function has the right extensive scaling)?

   We know that $\mathcal{H}^p$ is supposed to be extensive, i.e. proportional to $O(N)$. The energy function has two pieces: the first one is a sum over N^p terms (at leading order in N), the second one over N. This means the right scalings are with $\alpha = p - 1$, $\gamma = 0$.

2. (1 pt) Define the magnetisation as

$$m(\mathbf{s}) = \frac{\sum_{i=1}^N s_i}{N} \tag{2}$$

   Show that:

$$\mathcal{H}^p[\mathbf{s}] = \mathcal{H}^p[m] = -N \left(\frac{m^p}{p!} + hm \right) \tag{3}$$

   We have that:

$$\frac{\sum_{i_1=1}^N \sum_{i_2=1}^N \dots \sum_{i_p=1}^N s_{i_1} s_{i_2} \dots s_{i_p}}{N^p} = \left(\frac{\sum_{i=1}^N s_i}{N} \right)^p = m(\mathbf{s})^p \tag{4}$$

   and similar for the magnetic field term.

3. (2 pt) Compute the partition function $\mathcal{Z}$ and show that

$$\mathcal{Z} = \mathcal{C} \int dm \, d\hat{m} \, e^{N\Phi(m, \hat{m})} \tag{5}$$

   with $\mathcal{C}$ a constant independent on m and $\hat{m}$, and

$$\Phi(m, \hat{m}) = \beta \left(\frac{m^p}{p!} + hm \right) + m\hat{m} + \log\cosh\hat{m} \tag{6}$$

(there is no need to justify a possible Wick's rotation $i\hat{m} \to \hat{m}$).

We write the definition of partition function and use the δ function to introduce the magnetisation as we did in the Curie-Weiss model. You should get

$$\mathcal{Z} = N \int dm e^{\beta N \left(\frac{m^p}{p!} + hm \right)} \sum_{\{s_i\}} \delta \left(Nm - \sum_i s_i \right) \tag{7}$$

Now you can use the Fourier representation for the δ plus a Wick's rotation $i\hat{m} \to \hat{m}$ and obtain

$$\mathcal{Z} = \mathcal{C} \int dm d\hat{m} e^{\beta N \left(\frac{m^p}{p!} + hm + m\hat{m} \right)} \sum_{\{s_i\}} \exp \left(-\hat{m} \sum_i s_i \right) \tag{8}$$

At this point the single spins are decoupled, so you can just sum over them

$$\sum_{\{s_i\}} \exp \left\{ \hat{m} \left(\sum_i s_i \right) \right\} = \mathcal{C} (\cosh \hat{m})^N \tag{9}$$

The result is obtained by plugging this sum in the previous expression and exponentiating again.

4. (1 pt) Derive the state equation for m.

   We do the saddle point first on $\hat{m}$ and get

$$\tanh \hat{m} = -m \tag{10}$$

   Now we do the saddle point on m and get

$$\beta \left(\frac{m^{p-1}}{(p-1)!} + hm \right) = -\hat{m} \tag{11}$$

   The result follows combining the two saddle points into a single equation for m.

$$m = \tanh \beta \left(\frac{m^{p-1}}{(p-1)!} + h \right) \tag{12}$$

From now on consider the case $p = 3$ without an external field $h = 0$, for which the state equation equals

$$m = \tanh \left(\beta \frac{m^2}{2} \right) . \tag{13}$$

We will be studying the solutions of this equation. We recommend to plot both sides of the equation for several values of β to gain some intuition.

5. (1 pt) Argue that there is a paramagnetic solution to this equation for all values of β, and that it is the only solution for large temperature $\beta \to 0$.

   A paramagnetic solution has $m = 0$. We see that this is always a solution to the state equation for all values of β, and that when $\beta = 0$ the right hand side of the equation equals zero irrespective of the value of m at which it is evaluated to as m is bounded in $(-1, 1)$, so in this case $m = 0$ is the only solution.

6. (1 pt) Show that for very large β, $m = 1$ is a solution, and that it is the one dominating the free entropy. In other words, the system will be in a ferromagnetic phase for small temperature.

   At large β we have
$$\tanh\left(\beta\frac{m^2}{2}\right)|_{m=1} \approx 1 \tag{14}$$
   so $m = 1$ is a solution. At large β, the free entropy is just β times the energy, as the Gibbs measure concentrates on its ground states. We see from the expression in point 2 that $m = 1$ minimises the energy, thus it is the dominant solution.

7. (2 pt) It thus seems that the low-temperature ferromagnetic solution $m > 0$ (which becomes $m = 1$ strictly in the $\beta \to \infty$ limit) must appear only for β larger than some critical value. Imagine slowly raising β until a $m \neq 0$ solution first appear. Can this solution be arbitrarily close to $m = 0$? Deduce that the paramagnetic to ferromagnetic transition in this model is of the first order (that is, the order parameter m is discontinuous as a function of β). (Hint: look at the right hand side of the state equation for $m > 0$. It looks like a smooth step function. What is the slope near $m = 0$?)

   Let's define $f(x)$
$$f(x) = \tanh\left(\beta\frac{x^2}{2}\right) \tag{15}$$
   We can think of the solutions of the state equations as the intersection of the graph of $f(x)$ with $y = x$. Notice how $f'(0)$ is zero, which means that the only intersection in the neighbourhood of zero can be $x = 0$. If there is another intersection, it's for sure bounded away from zero. Thus, there is no way for the magnetisation for rise from $m = 0$ to $m \neq 0$ in a continuous way, making the phase transition in this model of the first order, i.e. with discontinuous order parameter.

5.2 Variants of the storage problem

During the lecture we studied in detail the storage problem, in which we are given P labelled points $\xi^\mu \in \mathbb{R}^N$ (and we saw that we can take without loss of generality all labels $\sigma^\mu = +1$), and we want to find a linear classifier (or hyperplane) $J \in \mathbb{R}^N$, with norm constraint $||J||^2 = N$, such that
$$\Delta^\mu = \frac{1}{\sqrt{N}} J^T \xi^\mu \geq \kappa\,, \tag{16}$$
for all μ, and for a given margin $\kappa > 0$. That is, J must classify correctly all points, and must do so respecting a minimal margin κ. We studied the problem as a constraint satisfaction problem, and we were interested in the fundamental question of characterising the SAT/UNSAT transition: given a fixed large N, at which value P_c do we stop finding solutions (because there are too many constraints)? We found that $P_c = \alpha_c(\kappa)N$ in the large N limit, for an explicitly characterised threshold $\alpha_c(\kappa)$ which we computed.

In this exercise, we want to introduce modifications to this problem, to showcase other typically studied variants, and to suggest how to answer questions of different kind. The key take-home message of this exercise is that **the replica computation we did in class generalises easily to many variations of the problem**. This is a recurring theme in the field: there are some standard analytically-tractable "ingredients" that we can combine to perform the replica computation of many different problems.

We recall that, in the lecture, we defined Gardner's volume as

$$\Omega(\{\xi^\mu\}_{\mu=1}^p) = \int d\mu(J) \Pi_{\mu=1}^p \theta\left(\frac{1}{\sqrt{N}} J^T \xi^\mu - \kappa\right) , \tag{17}$$

where

$$d\mu(J) = dJ_1 \dots dJ_N \, \delta\left(||J||^2 - N\right) , \tag{18}$$

and that by replica theory we found this representation for the replicated volume

$$\mathbb{E}_\xi \Omega(\xi)^n = \int \Pi_{a<b} dq^{ab} \exp\left[P s_{\text{energy}}(q^{ab}) + N s_{\text{entropy}}(q^{ab})\right] , \tag{19}$$

where

$$s_{\text{entropy}}(q^{ab}) = \frac{1}{N} \log\left[\int \Pi_a d\mu(J^a) \Pi_{a<b} \delta(N q^{ab} - \sum_{i=1}^N J_i^a J_i^b)\right] , \tag{20}$$

and

$$s_{\text{energy}}(q^{ab}) = \log \int \Pi_a d\Delta^a \mathcal{N}\Big(\{\Delta^a\}_{a=1}^n \Big| 0, q_{ab}\Big) \theta(\Delta^a - k) \, . \tag{21}$$

5.2.1 Computing the minimal number of errors

A first modification we could introduce in our problem is to allow for hyperplanes that do not fit the data, but penalise them based on the amount of errors they do. We introduce a penalty function $v(x)$ such that $v(x) = 0$ for $x \geq 0$, and $v(x) > 0$ for $x < 0$, and to each hyperplane we associate the following energy/penalty function

$$E(J) = \sum_{\mu=1}^P v\left(\frac{1}{\sqrt{N}} J^T \xi^\mu - \kappa\right) . \tag{22}$$

1. (1 pt) What are the ground states (configurations J of minimal energy) of the energy $E(J)$ in the SAT phase, i.e. for $\alpha < \alpha_c(\kappa)$ we computed during the lecture?

   The energy is non-negative, and equals zero for all J s.t.

   $$\frac{1}{\sqrt{N}} J^T \xi^\mu \geq \kappa \, . \tag{23}$$

   Thus, the ground states are all hyperplanes that classify correctly all points if they exists., and they do because we are in the SAT phase.

2. (1 pt) Find the expression of $v(x)$ such that the energy function $E(J)$ counts the number of points the hyperplane J *incorrectly* classifies.

   It is sufficient to choose $v(x) = +1$ for all $x < 0$ and zero otherwise, i.e. each configuration gets exactly one unit of penalty for each of the misclassified points.

We now introduce the canonical partition function associated to the energy $E(J)$ as

$$\begin{aligned} Z(\{\xi^\mu\}_{\mu=1}^p, \beta) &= \int d\mu(J) \exp\left(-\beta \sum_{\mu=1}^P v\left(\frac{1}{\sqrt{N}} J^T \xi^\mu - \kappa\right)\right) \\ &= \int d\mu(J) \prod_{\mu=1}^P \exp\left(-\beta v\left(\frac{1}{\sqrt{N}} J^T \xi^\mu - \kappa\right)\right) . \end{aligned} \tag{24}$$

For the choice $v(x) = +\infty$ for all $x < 0$, this reduces to the Gardner's volume defined in Eq. (17). Thus, you should not be too surprised that the replica computation we did at lecture for the SAT/UNSAT transition can be generalised to include the energy term $E(J)$, allowing us to get information about how bad our hyperplanes perform in the UNSAT phase. Indeed, the partition function for $\beta \to +\infty$ concentrates on the ground states. In the SAT phase $\alpha < \alpha_c(\kappa)$ the ground states are the zero-energy correctly-classifying hyperplanes, so the $\beta \to +\infty$ partition function there behaves as the Gardner's volume we already studied. In the UNSAT phase $\alpha > \alpha_c(\kappa)$ the ground states are the hyperplanes which achieve the minimal value of the energy $E(J)$. If we choose $v(x)$ as in point 2, i.e. each configuration gets energy exactly equal to the number of misclassified points, then the average energy in the UNSAT phase for $\beta \to \infty$ will count the minimal number of errors that an hyperplane can do.

The full replica computation for the partition function $Z(\{\xi^\mu\}_{\mu=1}^p, \beta)$, or better for the associated averaged free entropy $N^{-1}\mathbb{E}_\xi \log Z(\{\xi^\mu\}_{\mu=1}^p, \beta)$, is outside the scope of this exercise. Yet, it is a useful exercise to understand how the replica computation we did at lecture generalises to this setting. You can convince yourself that Eq. (17) and Eq. (24) are structurally identical, with just the θ function being substituted by a different function of the same argument, $\frac{1}{\sqrt{N}} J^T \xi^\mu - \kappa$. Thus, while the small details of the computation will change, and the final analysis of the state equations will be different, the structure itself of the computation will not change drastically.

Consider the energetic-entropic decomposition of Eq (19) for Gardner's volume.

3. (1 pt) Argue that the entropic term s'_{entropy} associated to the replica computation for the partition function (24) is the same as the entropic term s_{entropy} associated to the replica computation for the Gardner's volume (17).

   The entropic term depends only on the ambient space of the variables J. In both cases, they are N dimensional vectors with norm $||J||^2 = N$, so the entropic term does not change.

4. (1pt) The energetic term instead needs a small update. Argue that the energetic term $s'_{\text{energetic}}$ associated to the replica computation for the partition function (24) equals

$$s'_{\text{energy}}(q^{ab}) = \log \int \Pi_a d\Delta^a \mathcal{N}\Big(\{\Delta^a\}_{a=1}^n \Big| 0, q_{ab} \Big) \exp(-\beta v(\Delta^a - k)) \,. \tag{25}$$

   Eq. (24) is identical to Eq. (17) under the substitution $\theta(x) \leftrightarrow \exp(-\beta v(x))$. It suffices to make the same substitution in the energetic term, as the replica computation does not use any specific property of the function θ.

We now compute the new energetic term in the replica symmetric (RS) ansatz, in which $q^{ab} = \delta_{ab} + (1 - \delta_{ab})q$.

5. (1 pt) Review the replica computation we did in class, and argue that the energetic term at leading order for $n \to 0$ satisfies

$$s'_{\text{energy}}(q^{ab}) \approx n \int Dz \log \left[\int Dt \exp\left\{ -\beta v(\sqrt{1-q}\,t + \sqrt{q}\,z - \kappa) \right\} \right] \tag{26}$$

   where $Dz = \exp(-z^2/2)/\sqrt{2\pi} dz$ and similar for Dt. Hint: no need to redo the full computation...

   Eq. (24) is identical to Eq. (17) under the substitution $\theta(x) \leftrightarrow \exp(-\beta v(x))$. It suffices to make the same substitution in the energetic term, as the replica computation does not use any specific property of the function θ. Notice that while the inner integral was "solvable" in the θ case, it is now not "solvable" in general, as the function v is generic.

5.2.2 Changing the hyperplane space

Let's now forget about the energy function $E(J)$, and let us go back to the Gardner's volume Eq. (17) and to the problem of computing SAT/UNSAT transitions.

Another classic modification to the problem is to add additional constraints on the hyperplane's parameter J. For example, instead of letting the elements of J be continuous, we could consider a discrete setting in which each coordinate of the hyperplane vector $J_i \in \{-1, 1\}$, i.e. it's binary.

Contrary to the previous modifications of the problem (the introduction of and energy function, and the sign constraints we looked at in Hw 4), this modification is much more disrupting. Indeed, we are completely changing the nature of the constraint satisfaction problem, which moves from having continuous variables to having discrete variables. Moreover, the solution space of the problem is not convex anymore. The solution set is discrete, so it makes no sense to continuously interpolate between solutions as we did to check for the convexity of the solution space in the spherical problem (the one we saw in the lecture). Even computing the SAT/UNSAT transition as the α for which $q \to 1$ makes little sense, as the solution space is not a nice convex subset, and may disappear suddenly without shrinking to a single point.

All these considerations will affect the interpretation of the replica computation, but from the technical point of view the computation remains very similar to the spherical case.

1. (1 pt) Write the Gardner's volume for the binary storage problem.

$$\Omega_{\text{binary}}(\{\xi^\mu\}_{\mu=1}^p) = \sum_{J_1,\dots,J_N=\pm 1} \Pi_{\mu=1}^p \theta\left(\frac{1}{\sqrt{N}} J^T \xi^\mu - k\right) . \tag{27}$$

As discussed in class, normalising by the total number of configurations 2^N is not needed, but not wrong. It just amounts to a shift of the free entropy by $\log 2$.

2. (1 pt) Argue that in the entropic/energetic decomposition (19) for the replica computation in this new binary problem, the energetic term is the same.

The energetic term depends only on the classification constraints, and not on the ambient space of the J. Thus it remains unchanged.

3. (1 pt) Argue that in the entropic/energetic decomposition (19) for this new binary problem, the entropic term becomes

$$s''_{\text{entropy}}(q^{ab}) = \frac{1}{N} \log\left[\sum_{J_1^a,\dots,J_N^a=\pm 1} \Pi_{a<b} \delta\left(N q^{ab} - \sum_{i=1}^N J_i^a J_i^b\right)\right] , \tag{28}$$

where the sum runs over $J_i^a = \pm 1$ for all replicas $a = 1, \dots, n$ and all coordinates $i = 1, \dots, N$.

The entropic term depends on the ambient space of the J only in the integration measure. Before we integrated on the N-dim sphere with radius $\sqrt{N}$, now we sum over all the binary configurations $J \in \{-1, +1\}^N$.

4. (1pt) Show that the entropic term can be rewritten as

$$\exp(N s''_{\text{entropy}}(q^{ab})) = \int \mathrm{d}\hat{q}\, e^{N \sum_{a<b} \hat{q}_{ab} q_{ab} + NI} \tag{29}$$

where

$$I = \log\left(\sum_{J^1,\dots,J^n=\pm 1} \exp\left\{-\sum_{a<b}\hat{q}^{ab} J^a J^b\right\}\right) \tag{30}$$

Notice that here the sum runs only over $J^a = \pm 1$ for all replicas $a = 1,\dots,n$: we lost the coordinate index. (There is no need to justify any Wick's rotation of the kind $\hat{q} \to \pm i\hat{q}$).

$$\sum_{J_1^a,\dots,J_N^a=\pm 1} \prod_{a<b} \delta\left(Nq^{ab} - \sum_{i=1}^N J_i^a J_i^b\right) = \tag{31}$$

$$\int \mathrm{d}\hat{q} \sum_{J_1^a,\dots,J_N^a=\pm 1} \prod_{a<b} \exp\left\{N\hat{q}^{ab}q^{ab} - \hat{q}^{ab}\sum_{i=1}^N J_i^a J_i^b\right\} = \tag{32}$$

$$\int \mathrm{d}\hat{q} \sum_{J_1^a,\dots,J_N^a=\pm 1} \exp\left\{N\sum_{a<b}\hat{q}^{ab}q^{ab} - \sum_{a<b}\hat{q}^{ab}\sum_{i=1}^N J_i^a J_i^b\right\} = \tag{33}$$

$$\int \mathrm{d}\hat{q}\, e^{N\sum_{a<b}\hat{q}_{ab}q_{ab}} \sum_{J_1^a,\dots,J_N^a=\pm 1} \exp\left\{-\sum_{a<b}\hat{q}^{ab}\sum_{i=1}^N J_i^a J_i^b\right\} = \tag{34}$$

$$\int \mathrm{d}\hat{q}\, e^{N\sum_{a<b}\hat{q}_{ab}q_{ab}} \left[\sum_{J^1,\dots,J^n=\pm 1} \exp\left\{-\sum_{a<b}\hat{q}^{ab}J^a J^b\right\}\right]^N \tag{35}$$

5. (3 pt) Use the RS ansatz $\hat{q}^{ab} = -\hat{q}$ (for $a < b$, the only values for which $\hat{q}^{ab}$ is defined) and sum over the spins to get

$$I = -n\frac{\hat{q}}{2} + \log\left[\int Dz \left[2\cosh z\sqrt{\hat{q}}\right]^n\right] \tag{36}$$

where $Dz = \exp(-z^2/2)/\sqrt{2\pi}\, dz$. This computation is very similar in spirit to the computation of the energetic term that we saw in class.

Since J has binary entries, we have $(J^a)^2 = 1$.

$$\sum_{J^1,\dots,J^n=\pm 1} \exp\left\{-\sum_{a<b} \hat{q}^{ab} J^a J^b\right\} \tag{37}$$

$$\sum_{J^1,\dots,J^n=\pm 1} \exp\left\{\frac{\hat{q}}{2}\sum_{a,b} J^a J^b - \sum_a \frac{\hat{q}}{2}\right\} = \tag{38}$$

$$e^{-n\hat{q}/2} \sum_{J^1,\dots,J^n=\pm 1} \exp\left\{\left[\sqrt{\hat{q}}\sum_a J^a\right]^2\right\} = \tag{39}$$

$$e^{-n\hat{q}/2} \int Dz \sum_{J^1,\dots,J^n=\pm 1} \exp\left\{z\sqrt{\hat{q}}\sum_a J^a\right\} = \tag{40}$$

$$e^{-n\hat{q}/2} \int Dz \left[\sum_{J=\pm 1} \exp\left\{z\sqrt{\hat{q}}J\right\}\right]^n = \tag{41}$$

$$e^{-n\hat{q}/2} \int Dz \left[2\cosh z\sqrt{\hat{q}}\right]^n \tag{42}$$

At this point, one could take the small n limit as we did for the energetic term in the lecture, concluding the computation of this different entropic term. The final expression for the entropic term is given by

$$\exp(N s''_{\text{entropy}}(q)) = \int d\hat{q} \exp\left[nN\left(-\frac{1}{2}(1-q)\hat{q} + \int Dz \log(2\cosh(z\sqrt{\hat{q}}))\right)\right] \tag{43}$$

where we are not able to integrate over the $\hat{q}$ variable exactly (as we were instead able to do in the spherical case). This is a generic feature of replica computations, in which this conjugate variable $\hat{q}$ for the overlap q remains non-integrated, and needs to be dealt with at the level of the saddle-point equations (much like in the Curie-Weiss model both m and $\hat{m}$ appeared in the final saddle-point problem).

A final remark: in this binary problem, the RS ansatz is not anymore strictly correct. One can see that by considering the RS ansatz, some inconsistencies between the result of the computation and some bounds that one can derive for the SAT/UNSAT threshold arise, suggesting that a more refined ansatz is necessary. The crucial ingredient that goes missing is the convexity of the constraint satisfaction problem.