


Matrix factorization with neural networks

Francesco Camilli^{1,*} and Marc Mézard^{2,†}

¹*Quantitative Life Sciences, International Centre for Theoretical Physics, Trieste 34151, Italy*

²*Department of Computing Sciences, Bocconi University, Milan 20100, Italy*

 (Received 21 November 2022; revised 3 March 2023; accepted 23 May 2023; published 27 June 2023)

Matrix factorization is an important mathematical problem encountered in the context of dictionary learning, recommendation systems, and machine learning. We introduce a decimation scheme that maps it to neural network models of associative memory and provide a detailed theoretical analysis of its performance, showing that decimation is able to factorize extensive-rank matrices and to denoise them efficiently. In the case of binary prior on the signal components, we introduce a decimation algorithm based on a ground-state search of the neural network, which shows performances that match the theoretical prediction.

DOI: [10.1103/PhysRevE.107.064308](https://doi.org/10.1103/PhysRevE.107.064308)

Matrix factorization is a generic inference problem that is simply stated. Consider a matrix \mathbf{Y} which is a noisy observation of the product of two matrices $\mathbf{Y}^* = \mathbf{A}\mathbf{B}$, and some prior knowledge on the distribution of elements of the two factors \mathbf{A} and \mathbf{B} . We would like to answer two main questions: (i) In what regimes of sizes of \mathbf{A} , \mathbf{B} , and noise is it possible to reconstruct the two factors (up to a permutation of the lines of \mathbf{A} and the columns of \mathbf{B})? (ii) Do there exist efficient algorithms that achieve a good performance?

Matrix factorization is important for several concrete tasks. In *dictionary learning* [1–3], one wants to find the two factors \mathbf{A} of size $N \times P$, and \mathbf{B} of size $P \times M$, such that the columns of \mathbf{A} form an overcomplete basis (i.e., $P > N$) and \mathbf{B} is sparse. In other words, one aims for a sparse representation of the data \mathbf{Y} , a common strategy adopted for the denoising of large matrices which aims at reconstructing \mathbf{Y}^* from the measured \mathbf{Y} [4]. The much-studied Boltzmann machines are another instance of this approach [5,6]: Given a set of observations on their visible units, after training they create an internal representation on the hidden units. The core idea of representation learning, and presumably the reason for its effectiveness [7–9], is the extraction of characteristic features, e.g., the overcomplete basis in \mathbf{A} that, if properly recombined, can reconstruct the data or even some of their missing parts, as in recommender systems [10]. Let us stress that the main practical application of matrix factorization is to find a compact representation of the data. One way to do it, as described above, is to force sparsity in \mathbf{B} , though keeping a large rank representation of \mathbf{Y} in general. An alternative is to search for a low rank structure in the data. As we shall recall below, this is a well-studied problem. Another motivation for studying matrix factorization is to find unique ways of training deep networks. Given a desired output, the process of learning can be decomposed layer by layer, in which the task is to find a set of synaptic weights together with the internal representation

of the data in the previous layer. This is a matrix factorization task (complicated by the nonlinearity) that one can hope to turn into a self-consistent solution of deep network training, following what was done for multilayer generalized linear estimation [11].

In the present paper, we propose a way to study matrix factorization, in the difficult regime where the dimensions of the factors go to infinity simultaneously, by mapping it to neural network models used for associative memories. While our analysis can be carried out for the generic problem, for the sake of clarity we present it here in its symmetric version and with Gaussian noise. In this case, one measures the $N \times N$ matrix

$$\mathbf{Y} = \frac{\xi\xi^T}{\sqrt{N}} + \sqrt{\Delta}\mathbf{Z} \quad (1)$$

built from the factors $\xi = (\xi_i^\mu)$, $1 \leq i \leq N$, $1 \leq \mu \leq P$. Here \mathbf{Z} is a symmetric noise matrix in which the $i \leq j$ elements are all independent and Gaussian, with mean 0 and variance $1 + \delta_{ij}$. We shall refer to the vectors $\xi^\mu = (\xi_i^\mu)$, $1 \leq i \leq N$ as *patterns*.

When the rank P is finite, the model (1) goes under the name of spiked Wigner model, introduced in Ref. [12] as a model for principal component analysis and studied in detail in recent years [13–18].

A common setting to tackle (1) is the Bayes-optimal one, in which we suppose to know the noise variance Δ as well as the probabilistic model of the data, including the distribution of the patterns P_ξ . The inference task amounts to reconstructing the patterns stored in ξ from the posterior probability $(1/Z)P_{\text{prior}}(\mathbf{x})P(\mathbf{Y}|\mathbf{x})$. In the presence of Gaussian noise, the latter can be seen as a Boltzmann measure at temperature Δ for a $N \times P$ matrix \mathbf{x} with prior distribution $P_{\text{prior}}(\mathbf{x})$ and energy

$$E(\mathbf{x}|\mathbf{Y}) = \frac{1}{2} \sum_{i,j=1}^N \left(Y_{ij} - \sum_{\mu=1}^P \frac{x_i^\mu x_j^\mu}{\sqrt{N}} \right)^2. \quad (2)$$

*fcamilli@ictp.it

†marc.mezard@unibocconi.it

Bayes optimality means that the prior measure on \mathbf{x} , $P_{\text{prior}}(\mathbf{x})$, is the same measure as P_{ξ} from which the original patterns were drawn. While the Bayes-optimal setting is rarely available in practical inference problem, its relevance resides in its mathematical tractability and in the fact that sampling from the Bayes-optimal Boltzmann measure gives the smallest possible reconstruction error, on average over the observations \mathbf{Y} . This gives an information-theoretical bound for the performance of reconstruction algorithms. For finite P , these benchmark values were computed in the aforementioned literature in the thermodynamic limit $N \rightarrow \infty$, and there exist polynomial-time algorithms known as approximate message passing [19,20] able to saturate the Bayes-optimal performances.

Much less is known in the regime of extensive rank, namely, when P and N go to infinity with fixed $P/N = \alpha$. If one restricts to the simpler question of denoising, a possible strategy, introduced in Ref. [21], is that of using *rotationally invariant estimators* (RIEs) for the hidden matrix $\xi\xi^T/\sqrt{N}$. RIEs produce an estimation of \mathbf{Y}^* that has the same eigenbasis as the observed \mathbf{Y} . The task then reduces to denoising the eigenvalues of \mathbf{Y} . Notice, however, that RIEs are not Bayes optimal if the prior P_{ξ} is not rotational invariant.

Previous attempts at solving matrix factorization in the extensive rank regime [22–24] focused on the Bayes-optimal Boltzmann measure with energy (2). This turns out to be a difficult matrix model with quenched randomness that has defied exact analysis so far. For instance, it could not be established if a solution exists in the extensive rank case.

We propose here an alternative strategy: considering ξ as composed of P patterns ξ^{μ} , $\mu = 1, 2, \dots, P$; we aim for one of them at a time. Assuming that we are able to efficiently get an estimate η^p of a first pattern, say ξ^p , then we can build a rank-one contribution and subtract it from $\mathbf{Y} \equiv \mathbf{Y}_0$, obtaining $\mathbf{Y}_1 = \mathbf{Y}_0 - \eta^p \eta^{pT}/\sqrt{N}$. Then we iterate this process P times. We will refer to this iterative scheme as *decimation*.

Notice that decimation bears some similarity to the unlearning procedure studied in Refs. [25–27]. Unlearning aims at smoothing the energy landscape of attractor neural networks. It consists of iterating the following procedure: One starts from a random configuration, iterates a zero-temperature dynamics until one finds a fixed point η , and subtracts $\mathbf{Y}_1 = \mathbf{Y}_0 - \varepsilon \eta \eta^T/\sqrt{N}$ with a small ε . The main idea is that the zero temperature dynamics will likely be trapped in one of the numerous metastable states, and the subtractions will increase the energy of these metastable states. The process should be run with ε small enough and iterated for an appropriate number of times (too many iterations are disastrous as they flatten the whole energy landscape, including the patterns). Decimation presents two main differences with unlearning. Its aim is to find a η which is close to one of the stored patterns, not a metastable state, and then the subtraction is performed with $\varepsilon = 1$ above. It thus aims at completely removing a pattern from \mathbf{Y} before the next decimation step.

After R steps of decimation, and choosing indices such that approximate values of the last R patterns have been found, we must factorize the reduced matrix $\mathbf{Y}_R = \mathbf{Y} - \sum_{\mu=P-R+1}^P \eta^{\mu} \eta^{\mu T}/\sqrt{N}$. We then try to find a pattern which is a single N -component vector $\mathbf{x} = (x_i)$, $1 \leq i \leq N$, with an energy function that can be obtained from (2), replacing \mathbf{Y}

with \mathbf{Y}_R :

$$-E(\mathbf{x}|\mathbf{Y}_R) = \frac{\sqrt{\Delta}}{2\sqrt{N}} \sum_{i,j=1}^N Z_{ij} x_i x_j + \frac{1}{2N} \sum_{\mu=1}^P \left(\sum_{i=1}^N \xi_i^{\mu} x_i \right)^2 - \frac{1}{2N} \sum_{\mu=P-R+1}^P \left(\sum_{i=1}^N \eta_i^{\mu} x_i \right)^2 - \frac{\|\mathbf{x}\|^4}{4N}. \quad (3)$$

We aim at sampling a pattern from the Boltzmann distribution

$$\frac{1}{Z_R} P_{\xi}(\mathbf{x}) e^{-\beta E(\mathbf{x}|\mathbf{Y}_R)} \quad (4)$$

at large β in the case of factorized priors $P_{\xi}(\mathbf{x}) = \prod_{i \leq N} P_{\xi}(x_i)$. Expectations with respect to (4) are denoted by $\langle \cdot \rangle_R$. If we choose $P = 1$ in (3) and decide for finite temperature sampling at $\beta = 1/\Delta$, we recover the Bayes-optimal setting for the rank-1 spiked Wigner model. The assumption of factorized prior is very common in theoretical studies. It can be seen as a starting point for more structured approximations, and is commonly adopted also in empirical risk minimization, used to train artificial neural networks. Let us note that recent developments [28] suggest that under some hypothesis one could also take into account covariances between signal components.

The similarity of (3) with the energy of Hopfield's model [29] for associative memory helps to gain some insight into the physical behavior of decimation. To explain it, one can start from the ideal case where the last R patterns have been reconstructed exactly ($\eta^{\mu} = \xi^{\mu}$ for $\mu > P - R$). Then the energy (3) favors the \mathbf{x} configurations that are most aligned to one of the remaining $P - R$ patterns ξ^{μ} . Decimation may work if the minima of the energy (3) are retrieval states close to the $P - R$ patterns that one wants to reconstruct. In practice, the estimates $(\eta^{\mu})_{\mu \geq P-R+1}$ have an $O(N)$ projection onto the related patterns ξ^{μ} . The R quadratic terms involving the η^{μ} then repel the \mathbf{x} 's from the previously estimated patterns, as desired. The $\|\mathbf{x}\|^4$ term penalizes large norms.

We stress that in the decimation energy (3), there are three sources of noise that might disturb pattern retrieval. The first one is (a) the original Gaussian noise \mathbf{Z} . Then there is (b) a noise due to pattern interferences, the well-known limiting factor for patterns' memorization in the Hopfield model [30,31]. Third, (c) the decimation itself: the retrieval states η^{μ} are blurred versions of the ξ^{μ} .

Along decimation, the interference noise (b) decreases with the number of memorized patterns, but the decimation noise (c) increases. To study the theoretical performance of decimation, we need to identify which of these two mechanisms dominates.

The decimation noise contribution depends on the precision of the retrieval of previous patterns. Using knowledge from Hopfield's model phenomenology, we shall assume that for large N the retrieved states $(\eta^{\mu})_{\mu \geq P-R+1}$ are sampled from a factorized distribution, with an effective local field that is Gaussian centered at $m^{\mu} \xi^{\mu}$ with variance r , as described in Appendix A. The crucial parameter m^{μ} characterizes the retrieval accuracy, i.e., how close the retrieved state η^{μ} is to the stored pattern ξ^{μ} , and it must be determined self-consistently together with the variance r .




FIG. 1. Phase diagrams for the indicated values of inverse temperature β and sparsity ρ . Left panel: $\beta = 1/\Delta$, $\rho = 1$. The R phase is the retrieval phase. In the M phase, the retrieval states exist, but they are metastable. In the SG region, the system is in the spin-glass phase, where there is no retrieval. We do not show here the paramagnetic phase which appears above $\Delta = 1$. Right panel: $\beta \rightarrow \infty$, different values of ρ : red and blue $\rho = 1$, magenta and cyan $\rho = 0.1$, green and yellow $\rho = 0.05$. Dashed and solid lines refer to the same transitions as in the right panel.

The retrieval accuracy m^{P-R} on the pattern retrieved using the measure (4) depends on the whole set of previous retrieval accuracies m^{P-R+q} , $q = 1, \dots, R$. m^{P-R} can be obtained using the replica method [32], which allows us to compute the average of $\Phi = (1/N) \ln \mathcal{Z}_R$ over the distribution of patterns and of retrieved states in the large N limit. We use a *replica symmetric* ansatz, which is known to provide a good approximation to the exact result in neural networks [31]. Defining $t = R/P$, we show in Appendix B that the free entropy Φ can be written in terms of three order parameters: the sought retrieval accuracy m , the overlap $q = \mathbb{E}_{\xi, \eta} \langle x_i \rangle_R^2$, and the variance r . The dependence on m can be written in a form similar to the one of Hopfield's model,

$$\Phi = \text{Extr} \left\{ \Phi_0(\alpha, m^{[0,t]}; q, r, u, v, \hat{v}) - \beta \frac{m^2}{2} + \mathbb{E}_{Z, \xi} \ln \int dP_\xi(x) \exp \left((Z\sqrt{r} + \beta m \xi)x - \frac{u+r}{2} x^2 \right) \right\}, \quad (5)$$

where \hat{v} is an auxiliary variational parameter chosen to impose $v = \mathbb{E}_{\xi, \eta} \langle \|\mathbf{x}\|^2 \rangle_R / N = 1$, and $m^{[0,t]}$ stands for the collection of the accuracies of the previous decimation steps $\{m(\tau), \tau \leq t\}$. The function Φ_0 encapsulates all the noise contributions, including decimation, and is defined in Appendix B. Here, by Extr we intend extremization with respect to the variational parameters m, q, r, u, v, \hat{v} . In case of multiple stationary points, we choose the one with the largest resulting free entropy. The value of m where Φ reaches its maximum is the $R + 1$ th retrieval accuracy m^{P-R} . It depends crucially on the noise strength r which is the sum of three contributions, $r = r_a + r_b + r_c$, corresponding to the three sources of noise discussed above:

$$r_a = \beta^2 \Delta q, \quad r_b = \frac{(1-t)\alpha\beta^2 q}{(1-\beta(1-q))^2}, \quad (6)$$

$$r_c = 2\alpha t \beta^2 q \int_0^t d\tau f(\tau) [1 + 2\beta^2(1-q)^2 f(\tau)]. \quad (7)$$

r_c is the decimation-induced noise level—it depends on $f(\tau) = [1 - m(\tau)^2] / [1 - \beta^2(1-q)^2(1 - m(\tau)^2)]$. The other consistency equations are deferred to Appendix B 1, whereas

in Appendix B 2 we compute their low temperature limit.

We shall study the performance of decimation assuming the existence of an oracle algorithm which could be of two types: (1) Sampling oracle: returns a typical configuration sampled from the Boltzmann distribution (4) at inverse temperature $\beta = 1/\Delta$ and (2) ground-state oracle: returns the lowest energy configuration of energy (3) (equivalent to sampling at $\beta = \infty$). These oracles will return retrieval states correlated with the stored patterns only in some part of the phase diagram controlled by the parameters α, Δ, t .

We start with the first step of decimation, $t = 0$, using for definiteness a measure $P_\xi(x) = (1-\rho)\delta_{x,0} + \rho/2(\delta_{x,1/\sqrt{\rho}} + \delta_{x,-1/\sqrt{\rho}})$, where ρ controls the sparsity. The left panel of Fig. 1 shows the phase diagram for binary patterns with a sampling oracle. The retrieval region R, $\alpha < \alpha_c(\Delta)$, is where the retrieval states dominate the measure. In the M phase, $\alpha_c(\Delta) < \alpha < \alpha_M(\Delta)$, retrieval states are metastable, and in the spin-glass SG phase there are no states correlated with the patterns. This is strongly reminiscent of the phase diagram of the Hopfield model [31] and, in fact, the boundaries of the phase transition lines at $\Delta = 0$ are found at the same values, $\alpha_c(0) = 0.051$ and $\alpha_M(0) = 0.138$. Here Δ plays a role similar to the temperature in Hopfield model. The right panel of Fig. 1 shows that sparsity increases the size of the retrieval phase. It can be actually proved that $\alpha_M(\Delta = 0) \rightarrow \infty$ as $\rho \rightarrow 0$.

As decimation proceeds, the retrieval boundary α_c moves to the left because of the increase of retrieval noise r_c , and at the same time the effective α decreases. It turns out that, in all the cases that we have studied, the second effect is more important: If we start the decimation at $t = 0$ in the retrieval phase, then the system remains in the retrieval phase for the whole decimation process. This is the crucial point that makes this approach efficient for matrix factorization. It is illustrated in Fig. 2, where we plot the theoretical value of the retrieval accuracy m at decimation step $R = tP$, computed from (5), versus t . It starts from $m = .9971$ at $t = 0$ and increases to $m = .9996$ at the end of decimation, showing that the precision of pattern retrieval increases during decimation.

We now study the performance of decimation for denoising. The efficiency of a denoiser is measured by the matrix




FIG. 2. The accuracy of pattern retrieval is plotted versus the decimation parameter $t = R/P$. The red curve is the theoretical prediction with $\alpha = 0.03$, $\Delta = 0.08$ at $\beta = \infty$. The blue data points are an average over 300 runs of the ground-state oracle for $N = 1300$, with one standard deviation error bars.

mean square error between the denoiser estimation $\hat{\mathbf{S}}$ and the initial signal $\xi\xi^T/\sqrt{N}$:

$$\text{mMSE} = \frac{1}{2N^2} \mathbb{E} \left\| \hat{\mathbf{S}} - \frac{\xi\xi^T}{\sqrt{N}} \right\|^2 \xrightarrow{N \rightarrow \infty} \alpha - \alpha \int_0^1 d\tau m(\tau)^2. \quad (8)$$

The decimation estimator is $\hat{\mathbf{S}} = \sum_{\mu=1}^P \frac{\eta^\mu \eta^{\mu T}}{\sqrt{N}}$. In Fig. 3, we show that the decimation denoiser, which exploits the structure of the prior, is better than the RIE denoiser.

So far, we have established the theoretical performance of decimation based on an oracle able to find retrieval states close to patterns in an associative memory model. However, this retrieval is far from trivial: the oracle cannot use the property of associative memory, as we have no knowledge of the patterns. In a neurobiological perspective, what we seek is an algorithm able to retrieve the stored information in a neural network, knowing the connectivity and synaptic efficacies. In this context, a message passing algorithm for finite rank matrix factorization was tried recently, but it seems to be limited to the case where the rank scales sublinearly with the number of neurons [33].

We have tested an oracle, whose pseudocode can be found in Appendix C, based on a Monte Carlo method with simulated annealing starting from infinite temperature (a random configuration) and decreasing it to $T = 0$. An extra feature is a restart procedure: when the algorithm finds a configuration x^* , its energy $E(x^*)$ is compared to the ground-state energy E_{GS} at infinite N , which we know analytically from the replica method. When $E(x^*) > \theta E_{\text{GS}}$, with θ a threshold initially fixed to 1, x^* is not accepted as a potential pattern and the procedure is restarted from infinite temperature, with a slightly lower threshold. We have found that the number of restarts needed to find the patterns grows exponentially with N . However, although this algorithm needs an exponential time, the prefactor is small enough that it gives good performance up to N of order 1500 to 2200 (depending on the noise level Δ). It is quite possible that better oracles can be found, without the constraint of satisfying detailed balance.

From our analysis, we are thus able to conclude that (i) symmetric matrix factorization is possible in the extensive rank case and (ii) decimation succeeds at exploiting the structure of the signal in the prior and gets a better denoiser than RIEs. Furthermore, (iii) strong sparsity helps in the reconstruction, making the system more robust against the three noise sources. All these results require having an oracle able to find a retrieval state, given a matrix of coupling in an associative memory network. We saw that simulated annealing with restarts, although it is an exponential algorithm, provides such an oracle for small enough sizes. Finding a polynomial time oracle would be extremely interesting, both for matrix factorization and neurophysiological analysis. Finally, we stress that decimation can be analyzed similarly for the asymmetric matrix factorization problem of \mathbf{Y} into \mathbf{AB} and will be the object of future work.

ACKNOWLEDGMENTS

We thank Florent Krzakala, Lenka Zdeborová, and Jean Barbier for many fruitful discussions. M.M. acknowledges financial support by the PNRR-PE-AI FAIR project funded by the NextGeneration EU program.




FIG. 3. Matrix denoising: Comparison between rotation invariant (red) and denoising based on decimation with a sampling oracle at $\beta = 1/\Delta$ (blue) or with a ground-state oracle at $\beta = \infty$ (green). Performances are measured by the average mean square error over 30 samples for each Δ . Error bars are too small to be seen. Right panel: $N = 2000$, $P = 60$, $\rho = 1$. The magenta points are obtained via our ground-state oracle; error bars are 1 standard deviation of 20 MSEs. The first five points are obtained with $N = 1800$, the remaining three with $N = 1700, 1600, 1500$, respectively. Left panel: $N = 1000$, $P = 400$, $\rho = 0.05$.

APPENDIX A: RETRIEVAL STATES AND RETRIEVAL ERROR

Let us define the Mattis magnetizations:

$$m^\mu(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N \xi_i^\mu x_i, \quad \mu = 1, \dots, P, \quad (\text{A1})$$

$$p^\mu(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N \eta_i^\mu x_i, \quad \mu = P(1-t) + 1, \dots, P. \quad (\text{A2})$$

The energy then takes the useful form

$$-E(\mathbf{x}|\mathbf{Y}_R) = \frac{\sqrt{\Delta}}{2\sqrt{N}} \sum_{i,j=1}^N Z_{ij} x_i x_j + \frac{N}{2} \sum_{\mu=1}^P (m^\mu(\mathbf{x}))^2 - \frac{N}{2} \sum_{\mu=P(1-t)+1}^P (p^\mu(\mathbf{x}))^2 - \frac{\|\mathbf{x}\|^4}{4N} + \frac{\hat{v}}{2}(N - \|\mathbf{x}\|^2). \quad (\text{A3})$$

The Lagrange multiplier \hat{v} has been introduced to fix $\mathbb{E}_{\xi, \eta} \langle \|\mathbf{x}\|^2 \rangle_R / N \xrightarrow{N \rightarrow \infty} 1$. We assume that the oracle will produce η^μ with an asymptotic measure given by

$$\eta_i^\mu \sim \langle \cdot \rangle_{\xi_i^\mu, Z} = \frac{\int dP_\xi(x) e^{(Z\sqrt{r} + \beta m^\mu \xi_i^\mu)x - \frac{r+u}{2}x^2} (\cdot)}{\int dP_\xi(x) e^{(Z\sqrt{r} + \beta m^\mu \xi_i^\mu)x - \frac{r+u}{2}x^2}}, \quad \xi_i^\mu \sim P_\xi, Z \sim \mathcal{N}(0, 1) \text{ independent of other noises}, \quad (\text{A4})$$

where m^μ , i.e., the retrieval accuracy for η^μ and r, u must be determined self-consistently. Define for later convenience the quantities

$$\mathbb{E}_{\eta|\xi}[\eta_i^\mu] = m_i^\mu, \quad \mathbb{E}_{\eta|\xi}[(\eta_i^\mu)^2] = v_i^\mu. \quad (\text{A5})$$

Then, (A4) has the following implications:

$$\mathbb{E}_\xi[\eta_i^\mu] = \mathbb{E}_\xi \mathbb{E}_{\eta|\xi}[\eta_i^\mu] = 0, \quad \mathbb{E}_\xi[\xi_i^\mu m_i^\nu] = m^\mu \delta_{\mu,\nu}, \quad \mathbb{E}_\xi[v_i^\mu] = 1 \quad (\text{A6})$$

that are self-consistent with the fixed point equations for each decimation step.

APPENDIX B: REPLICAS SYMMETRIC FREE ENTROPY

In this Appendix, we compute the large N limit of the free entropy

$$\Phi_N = \frac{1}{N} \mathbb{E} \ln \int dP_\xi(\mathbf{x}) \exp[-\beta E(\mathbf{x}|\mathbf{Y}_R)], \quad (\text{B1})$$

where \mathbb{E} is taken with respect to all the disorder: \mathbf{Z}, ξ, η . This is done using the *replica method* [32]. We thus introduce

$$\mathbb{E}_{\mathbf{Z}_N^n} := \mathbb{E}_{\mathbf{Z}} \mathbb{E}_{\xi, \eta} \int \prod_{a=1}^n dP_\xi(\mathbf{x}_a) \exp \left[-\beta \sum_{a=1}^n E(\mathbf{x}_a|\mathbf{Y}_R) \right]. \quad (\text{B2})$$

We decompose this computation and start with the first noise terms in (A3) and the related $\mathbb{E}_{\mathbf{Z}}$ average

$$\begin{aligned} \mathbb{E}_{\mathbf{Z}} \exp \left(\frac{\beta\sqrt{\Delta}}{2\sqrt{N}} \sum_{i,j=1}^N Z_{ij} \sum_{a=1}^n x_{a,i} x_{a,j} \right) &= \exp \left(\frac{\beta^2 \Delta}{4N} \sum_{i,j=1}^N \sum_{a,b=1}^n x_{a,i} x_{a,j} x_{b,i} x_{b,j} \right) \\ &= \exp \left(\frac{N\beta^2 \Delta}{4} \sum_{a \neq b} Q^2(\mathbf{x}_a, \mathbf{x}_b) + \beta^2 \Delta \frac{\|\mathbf{x}_a\|^4}{4N} \right). \end{aligned} \quad (\text{B3})$$

Now we take care of the penalizing p terms in (A3). After replicating, their contribution to the partition function is

$$A := \prod_{\mu=P(1-t)+1}^P \prod_{a=1}^n e^{-\frac{N\beta}{2} (p^\mu(\mathbf{x}_a))^2} = \prod_{\mu=P(1-t)+1}^P \prod_{a=1}^n \int \frac{ds_a^\mu}{\sqrt{2\pi}} e^{-\frac{(s_a^\mu)^2}{2} + i\sqrt{\frac{\beta}{N}} s_a^\mu \sum_{j=1}^N \eta_j^\mu x_{a,j}}. \quad (\text{B4})$$

Notice that, thanks to the introduction of the auxiliary Gaussian variables $(s_a^\mu)_{a \leq n, P(1-t) < \mu \leq P}$, the exponential is now decoupled over the particle indices j . Consider then the expectation of A with respect to $\boldsymbol{\eta}$, given $\boldsymbol{\xi}$ with the assumptions (A6):

$$\begin{aligned} \mathbb{E}_{\boldsymbol{\eta}|\boldsymbol{\xi}}[A] &= \prod_{\mu=P(1-t)+1}^P \prod_{a=1}^n \int \frac{ds_a^\mu}{\sqrt{2\pi}} e^{-\frac{(s_a^\mu)^2}{2}} \prod_{i=1}^N \mathbb{E}_{\boldsymbol{\eta}_i^\mu | \boldsymbol{\xi}_i^\mu} \exp \left(i \sqrt{\frac{\beta}{N}} \eta_i^\mu \sum_{a=1}^n s_a^\mu x_{a,i} \right) \\ &= \prod_{\mu=P(1-t)+1}^P \prod_{a=1}^n \int \frac{ds_a^\mu}{\sqrt{2\pi}} \exp \left(-\frac{(s_a^\mu)^2}{2} + \sum_{i=1}^N \ln \mathbb{E}_{\boldsymbol{\eta}_i^\mu | \boldsymbol{\xi}_i^\mu} e^{i \sqrt{\frac{\beta}{N}} \eta_i^\mu \sum_{a=1}^n s_a^\mu x_{a,i}} \right). \end{aligned} \quad (\text{B5})$$

Now we can expand the exponential inside the \ln up to second order; the remaining terms will be of subleading order and thus neglected in the following:

$$\begin{aligned} \mathbb{E}_{\boldsymbol{\eta}|\boldsymbol{\xi}}[A] &= \prod_{\mu=P(1-t)+1}^P \prod_{a=1}^n \int \frac{ds_a^\mu}{\sqrt{2\pi}} \exp \left(-\frac{(s_a^\mu)^2}{2} + \sum_{a=1}^n i s_a^\mu \sqrt{\frac{\beta}{N}} \sum_{i=1}^N m_i^\mu x_{a,i} - \frac{\beta}{2} \sum_{a,b=1}^n s_a^\mu s_b^\mu \sum_{i=1}^N \frac{(v_i^\mu - (m_i^\mu)^2)}{N} x_{a,i} x_{b,i} \right) \\ &= \prod_{\mu=P(1-t)+1}^P \prod_{a=1}^n \int \frac{ds_a^\mu}{\sqrt{2\pi}} \exp \left[-\frac{1}{2} \sum_{a,b=1}^n s_a^\mu s_b^\mu \left(\delta_{ab} + \beta \sum_{i=1}^N \frac{(v_i^\mu - (m_i^\mu)^2)}{N} x_{a,i} x_{b,i} \right) + \sum_{a=1}^n i s_a^\mu \sqrt{\frac{\beta}{N}} \sum_{i=1}^N m_i^\mu x_{a,i} \right]. \end{aligned} \quad (\text{B6})$$

To continue, we assume condensation on a finite number of patterns, say the first k . We focus now on the remaining ones, namely, for $\mu > k$:

$$B := \exp \left[\frac{\beta N}{2} \sum_{a=1}^n \sum_{\mu=k+1}^P (m^\mu(\mathbf{x}_a))^2 \right] = \int \prod_{\mu=k+1}^P \prod_{a=1}^n \frac{dz_a^\mu}{\sqrt{2\pi}} \exp \left[-\sum_{a=1}^n \sum_{\mu=k+1}^P \frac{(z_a^\mu)^2}{2} + \sqrt{\frac{\beta}{N}} \sum_{a=1}^n \sum_{\mu=k+1}^P z_a^\mu \sum_{i=1}^N x_{a,i} \xi_i^\mu \right]. \quad (\text{B7})$$

Putting A and B together, their overall average over $(\boldsymbol{\xi}^\mu)_{\mu > k}$ takes the form

$$\begin{aligned} \mathbb{E}_{(\boldsymbol{\xi}^\mu)_{\mu > k}}[AB] &= \int \prod_{\mu=P(1-t)+1}^P \prod_{a=1}^n \frac{ds_a^\mu}{\sqrt{2\pi}} \int \prod_{\mu=k+1}^P \prod_{a=1}^n \frac{dz_a^\mu}{\sqrt{2\pi}} e^{-\frac{1}{2} \sum_{a=1}^n \left(\sum_{\mu=P(1-t)+1}^P \frac{(s_a^\mu)^2}{2} + \sum_{\mu=k+1}^P \frac{(z_a^\mu)^2}{2} \right)} \\ &\quad \times \exp \left[\sum_{i=1}^N \sum_{\mu=k+1}^P \ln \mathbb{E}_{\boldsymbol{\xi}_i^\mu} e^{\sqrt{\frac{\beta}{N}} \sum_{a=1}^n x_{a,i} \left(\xi_i^\mu z_a^\mu + i \chi(\mu > P-R) m_i^\mu s_a^\mu \right) - \chi(\mu > P-R) \sum_{a,b=1}^n s_a^\mu s_b^\mu \sum_{i=1}^N \frac{\beta(v_i^\mu - (m_i^\mu)^2) x_{a,i} x_{b,i}}{2N}} \right]. \end{aligned} \quad (\text{B8})$$

If we call $\mathbb{E}_{\boldsymbol{\xi}} m_i^{\mu 2} =: \bar{M}^{\mu 2}$, a further expansion of the exponential yields

$$\begin{aligned} \mathbb{E}_{(\boldsymbol{\xi}^\mu)_{\mu > k}}[AB] &= \int \prod_{\mu=P(1-t)+1}^P \prod_{a=1}^n \frac{ds_a^\mu}{\sqrt{2\pi}} \exp \left[-\frac{1}{2} \sum_{\mu=P(1-t)+1}^P \mathbf{s}^\mu \cdot (\mathbb{1} + \beta(1 - \bar{M}^{\mu 2})\mathcal{Q})\mathbf{s}^\mu \right] \\ &\quad \times \int \prod_{\mu=k+1}^P \prod_{a=1}^n \frac{dz_a^\mu}{\sqrt{2\pi}} \exp \left\{ -\sum_{\mu=k+1}^P \sum_{a=1}^n \frac{(z_a^\mu)^2}{2} + \frac{\beta}{2} \sum_{\mu=k+1}^P \sum_{a,b=1}^n z_a^\mu z_b^\mu \mathcal{Q}(\mathbf{x}_a, \mathbf{x}_b) \right. \\ &\quad \left. + i\beta \sum_{\mu=P(1-t)+1}^P \mathbb{E}_{\boldsymbol{\xi}} [\xi_1^\mu m_1^\mu] \sum_{a,b=1}^n z_a^\mu s_b^\mu \mathcal{Q}(\mathbf{x}_a, \mathbf{x}_b) - \frac{\beta}{\Delta} \sum_{\mu=P(1-t)+1}^P \sum_{a,b=1}^n (\bar{M}^\mu)^2 s_a^\mu s_b^\mu \mathcal{Q}(\mathbf{x}_a, \mathbf{x}_b) \right\}. \end{aligned} \quad (\text{B9})$$

We can now perform a Gaussian integration over the variables $\mathbf{z}^\mu = (z_a^\mu)_{a \leq n}$:

$$\begin{aligned} \mathbb{E}_{(\boldsymbol{\xi}^\mu)_{\mu > k}}[AB] &= \int \prod_{\mu=P(1-t)+1}^P \prod_{a=1}^n \frac{ds_a^\mu}{\sqrt{2\pi}} \exp \left[-\frac{1}{2} \sum_{\mu=P(1-t)+1}^P \mathbf{s}^\mu \cdot \left(\mathbb{1} + \beta\mathcal{Q} + \beta^2 \mathcal{Q} \frac{\mathbb{E}_{\boldsymbol{\xi}}^2 [\xi_1^\mu m_1^\mu]}{\mathbb{1} - \beta\mathcal{Q}} \right) \mathbf{s}^\mu \right] \\ &\quad \times \exp \left[-\frac{\alpha N}{2} \ln \det (\mathbb{1} - \beta\mathcal{Q}) \right]. \end{aligned} \quad (\text{B10})$$

Finally, after an integration over the remaining Gaussian variables \mathbf{s}^μ , and using (A6), we get

$$\mathbb{E}_{(\boldsymbol{\xi}^\mu)_{\mu > k}}[AB] = \exp \left[-\frac{\alpha(1-t)N}{2} \ln \det (\mathbb{1} - \beta\mathcal{Q}) - \frac{1}{2} \sum_{\mu=P(1-t)+1}^P \ln \det (\mathbb{1} - (1 - m^2(\tau^\mu))\beta^2 \mathcal{Q}^2) \right], \quad (\text{B11})$$

where $\tau^\mu = (1 - (\mu - 1)/P)$, and $m(\tau^\mu) = m^\mu$ are the previous retrieval accuracies. It remains to analyze the contribution given by $(\xi^\mu)_{\mu \leq k}$:

$$C := \exp \left[\frac{\beta N}{2} \sum_{a=1}^n \sum_{\mu=1}^k (m^\mu(\mathbf{x}_a))^2 \right] = \int \prod_{a=1}^n \prod_{\mu=1}^k dm_a^\mu \sqrt{\frac{\beta N}{2\pi}} \exp \left[\sum_{a=1}^n \sum_{\mu=1}^k \left(-N\beta \frac{(m_a^\mu)^2}{2} + \beta m_a^\mu \sum_{i=1}^N \xi_i^\mu x_{a,i} \right) \right]. \quad (\text{B12})$$

Before plugging the contributions coming from A , B , and C into $\mathbb{E} Z_N^n$, we need to introduce a collection of Dirac deltas to fix the desired order parameters that are organized in the overlap matrix $(Q(\mathbf{x}_a, \mathbf{x}_b))_{a,b=1}^n$:

$$1 = \int \prod_{a \leq b \leq n} dq_{ab} \delta(Q(\mathbf{x}_a, \mathbf{x}_b) - q_{ab}) = \int \prod_{a \leq b \leq n} \frac{N dr_{ab} dq_{ab}}{4\pi i} \exp \left[-\frac{1}{2} \sum_{a,b=1}^n r_{ab} \left(Nq_{ab} - \sum_i x_{a,i} x_{b,i} \right) \right]. \quad (\text{B13})$$

Hence, the averaged replicated partition function, at leading exponential order in N , takes the form

$$\begin{aligned} \mathbb{E} Z_N^n &= \int \prod_{a \leq b \leq n} \frac{N dr_{ab} dq_{ab}}{4\pi i} \int \prod_{a=1}^n \prod_{\mu=1}^k dm_a^\mu \sqrt{\frac{N\beta}{2\pi}} \exp \left[-\frac{N}{2} \sum_{a,b} r_{ab} q_{ab} - \frac{\beta N}{2} \sum_{a=1}^n \sum_{\mu=1}^k (m_a^\mu)^2 \right] \\ &\times \exp \left[-\frac{1}{2} \sum_{\mu=P(1-t)+1}^P \ln \det(\mathbb{1} - (1 - m^2(\tau^\mu))\beta^2 Q^2) \right] \\ &\times \exp \left[-\frac{\alpha(1-t)N}{2} \ln \det(\mathbb{1} - \beta Q) + N\beta^2 \Delta \sum_{a \neq b, 1}^n \frac{q_{ab}^2}{4} + N\beta \sum_{a=1}^n \left(\frac{\beta \hat{v}}{2} (1 - q_{aa}) + \frac{\beta \Delta - 1}{4} q_{aa}^2 \right) \right] \\ &\times \left(\int \prod_{\mu=1}^k dP_\xi(\xi^\mu) \prod_{a=1}^n dP_\xi(x_a) \exp \left[\frac{1}{2} \sum_{a,b=1}^n r_{ab} x_a x_b + \beta \sum_{\mu=1}^k \sum_{a=1}^n m_a^\mu \xi^\mu x_a \right] \right)^N, \end{aligned} \quad (\text{B14})$$

where we denote $Q = (q_{ab})_{a,b=1}^n$. We can finally express the replicated free entropy with a variational principle coming from a saddle-point argument applied to the formula above:

$$\begin{aligned} \Phi_n &:= \lim_{N \rightarrow \infty} \Phi_{N,n} = \frac{1}{n} \text{Extr} \left\{ -\frac{1}{2} \sum_{a,b} r_{ab} q_{ab} - \frac{\beta}{2} \sum_{a=1}^n \sum_{\mu=1}^k (m_a^\mu)^2 - \frac{\alpha(1-t)N}{2} \ln \det(\mathbb{1} - \beta Q) \right. \\ &+ \beta \sum_{a=1}^n \left(\frac{\hat{v}(1 - q_{aa})}{2} + \frac{\beta \Delta - 1}{4} q_{aa}^2 \right) - \frac{\alpha t}{2R} \sum_{\mu=P(1-t)+1}^P \ln \det[\mathbb{1} - (1 - m^2(\tau^\mu))\beta^2 Q] \\ &\left. + \beta^2 \Delta \sum_{a \neq b, 1}^n \frac{q_{ab}^2}{4} + \ln \int \prod_{\mu=1}^k \mathbb{E}_{\xi^\mu} \int \prod_{a=1}^n dP_\xi(x_a) \exp \left[\frac{1}{2} \sum_{a,b=1}^n r_{ab} x_a x_b + \beta \sum_{\mu=1}^k \sum_{a=1}^n m_a^\mu \xi^\mu x_a \right] \right\}. \end{aligned} \quad (\text{B15})$$

The normalized sum over $\mu = P(1-t) + 1, \dots, P$ on the second line can be turned into an integral $\int_0^t d\tau \dots$ in the large N limit.

The extremization is taken with respect to the collection of parameters $(r_{ab}, q_{ab})_{a,b=1}^n, (m_a^\mu)_{a=1, \mu=1}^{n,k}$. Within the replica symmetric ansatz,

$$\begin{cases} r_{ab} = r, & a \neq b \\ r_{aa} = -u \end{cases} \quad \begin{cases} q_{ab} = q, & a \neq b \\ q_{aa} = v \end{cases} \quad m_a^\mu = m^\mu, \quad Q = \begin{pmatrix} v & q & q & \dots & q \\ q & v & q & \dots & q \\ q & q & v & \dots & q \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ q & q & q & \dots & v \end{pmatrix} \in \mathbb{R}^{n \times n}. \quad (\text{B16})$$

The determinants of $\mathbb{1} - \beta Q$ and $\mathbb{1} - \beta^2(1 - m^2(\tau))Q$ are then easily computed:

$$\det(\mathbb{1} - \beta Q) = (1 - \beta(v - q))^n \left[1 - n \frac{\beta q}{1 - \beta(v - q)} \right], \quad (\text{B17})$$

$$\det(\mathbb{1} - (1 - m^2(\tau))\beta^2 Q^2) = [1 - (1 - m^2(\tau))\beta^2(v - q)^2]^{n-1} \times [1 - (1 - m^2(\tau))\beta^2(v + (n-1)q)^2]. \quad (\text{B18})$$

Further simplifications occur for the other terms in the replicated free entropy. In particular, the remaining \ln integral has to be treated as follows:

$$\begin{aligned}
& \int \prod_{\mu=1}^k \mathbb{E}_{\xi^\mu} \int \prod_{a=1}^n dP_\xi(x_a) \exp \left[\frac{r}{2} \sum_{a \neq b, 1}^n x_a x_b - \frac{u}{2} \sum_{a=1}^n x_a^2 + \beta \sum_{\mu=1}^k m^\mu \xi^\mu \sum_{a=1}^n x_a \right] \\
&= \int \prod_{\mu=1}^k \mathbb{E}_{\xi^\mu} \int \prod_{a=1}^n dP_\xi(x_a) \exp \left[\frac{r}{2} \left(\sum_{a=1}^n x_a \right)^2 - \frac{u+r}{2} \sum_{a=1}^n x_a^2 + \beta \sum_{\mu=1}^k m^\mu \xi^\mu \sum_{a=1}^n x_a \right] \\
&= \mathbb{E}_Z \int \prod_{\mu=1}^k \mathbb{E}_{\xi^\mu} \int \prod_{a=1}^n dP_\xi(x_a) \exp \left[\sqrt{r} Z x_a - \frac{u+r}{2} x_a^2 + \beta \sum_{\mu=1}^k m^\mu \xi^\mu x_a \right] \\
&= \mathbb{E}_Z \mathbb{E}_\xi \left[\int dP_\xi(x) \exp \left((Z\sqrt{r} + \beta \mathbf{m} \cdot \boldsymbol{\xi}) x - \frac{u+r}{2} x^2 \right) \right]^n, \tag{B19}
\end{aligned}$$

where $Z \sim \mathcal{N}(0, 1)$, $\boldsymbol{\xi} = (\xi^1, \dots, \xi^k)$, $\mathbf{m} = (m^1, \dots, m^k)$. Finally, expanding at first order in n , one has

$$\begin{aligned}
\Phi_n := \text{Extr} & \left\{ \frac{rq + uv}{2} - \beta \sum_{\mu=1}^k \frac{(m^\mu)^2}{2} - \frac{\beta^2 \Delta q^2}{4} - \frac{\alpha(1-t)}{2} \left[\ln(1 - \beta(v-q)) - \frac{\beta q}{1 - \beta(v-q)} \right] \right. \\
& - \frac{\alpha t}{2} \int_0^t d\tau \left[\ln(1 - (1 - m^2(\tau))\beta^2(v-q)^2) - \frac{2\beta^2 q(v-q)(1 - m^2(\tau))}{1 - \beta^2(1 - m^2(\tau))(v-q)^2} \right] \\
& \left. + \beta \left(\frac{\hat{v}(1-v)}{2} + \frac{\beta \Delta - 1}{4} v^2 \right) + \mathbb{E}_{Z, \boldsymbol{\xi}} \ln \int dP_\xi(x) \exp \left((Z\sqrt{r} + \beta \mathbf{m} \cdot \boldsymbol{\xi}) x - \frac{u+r}{2} x^2 \right) \right\} + O(n). \tag{B20}
\end{aligned}$$

It suffices then to let $n \rightarrow 0$. From the previous, we also deduce the expression of Φ_0 :

$$\begin{aligned}
\Phi_0(\alpha, m^{[0,t]}; q, r, u, v, \hat{v}) &= \frac{rq + uv}{2} - \frac{\beta^2 \Delta q^2}{4} - \frac{\alpha(1-t)}{2} \left[\ln(1 - \beta(v-q)) - \frac{\beta q}{1 - \beta(v-q)} \right] \\
& - \frac{\alpha t}{2} \int_0^t d\tau \left[\ln(1 - (1 - m^2(\tau))\beta^2(v-q)^2) - \frac{2\beta^2 q(v-q)(1 - m^2(\tau))}{1 - \beta^2(1 - m^2(\tau))(v-q)^2} \right] \\
& + \beta \left(\frac{\hat{v}(1-v)}{2} + \frac{\beta \Delta - 1}{4} v^2 \right). \tag{B21}
\end{aligned}$$

1. Fixed point equations

The stationarity conditions coming from (B20) are

$$v = 1, \tag{B22}$$

$$v = \mathbb{E}_{\xi, \eta} \langle X^2 \rangle_R \tag{B23}$$

$$m^\mu = \mathbb{E}_{\xi, \eta} \xi^\mu \langle X \rangle_R, \quad \mu = 1, \dots, k, \tag{B24}$$

$$q = \mathbb{E}_{\xi, \eta} \langle X \rangle_R^2, \tag{B25}$$

$$r = \frac{\alpha(1-t)\beta^2 q}{(1 - \beta(v-q))^2} + \beta^2 \Delta q + \alpha t \int_0^t d\tau \frac{2q\beta^2(1 - m^2(\tau))}{1 - \beta^2(1 - m^2(\tau))(v-q)^2} \left[1 + \frac{2\beta^2(v-q)^2(1 - m^2(\tau))}{1 - \beta^2(1 - m^2(\tau))(v-q)^2} \right], \tag{B26}$$

$$\begin{aligned}
u &= \beta \hat{v} + \beta(1 - \beta \Delta)v - \alpha(1-t)\beta \frac{1 - \beta(v-2q)}{(1 - \beta^2(v-q))^2} \\
& - \alpha t \int_0^t d\tau \left[\frac{2v\beta^2(1 - m^2(\tau))}{1 - \beta^2(1 - m^2(\tau))(v-q)^2} + q \frac{4\beta^4(v-q)^2(1 - m^2(\tau))^2}{(1 - \beta^2(1 - m^2(\tau))(v-q)^2)^2} \right]. \tag{B27}
\end{aligned}$$

The first equation, corresponding to \hat{v} , can be directly eliminated. Notice that the effect of decimation is visible only in the variables u and r that affect the local measure (A4). For all practical purposes, we will make finite size simulations and use the discretized form present in (B15) of the integral accounting for decimation contributions, starting from step 0, when no pattern has been retrieved yet. Finally, notice that mixed states solutions are possible, with the estimates aligning to more than 1 pattern,

i.e., different $m^{\mu'}$ s in (B24) are nonvanishing. This is not desirable in inference, since one wants to estimate one pattern at a time with the best possible performance.

2. The zero temperature limit

Let us express the $\beta \rightarrow \infty$ limit of the free entropy with a prior of the form

$$P_{\xi} = (1 - \rho)\delta_0 + \frac{\rho}{2}[\delta_{-1/\sqrt{\rho}} + \delta_{1/\sqrt{\rho}}], \quad \rho \in (0, 1]. \tag{B28}$$

For future convenience, we introduce the notations

$$C := \beta(1 - q) \in [0, 1), \quad \bar{r} := r/\beta^2, \quad U := \frac{u + r}{\beta}, \tag{B29}$$

where q is intended as the stationary value of the overlap solving the fixed point equations. Denote $\mathbf{m} = (m^{\mu})_{\mu=1}^k$, where k is the maximum number of condensed patterns. In the low temperature limit, the free entropy, rescaled by β and evaluated at the stationary values of the parameters involved, has the form

$$\frac{1}{\beta} \Phi = -\frac{\bar{r}C}{2} + \frac{U}{2} + \frac{\alpha(1-t)}{2(1-C)} - \frac{1}{4} - \frac{\mathbf{m}^2}{2} + \frac{\Delta C}{2} + \psi + \alpha t \int_0^t d\tau \frac{C(1 - m^2(\tau))}{1 - (1 - m^2(\tau))C^2}, \tag{B30}$$

where

$$\psi = \frac{1}{\beta} \mathbb{E}_{\xi, Z} \ln \left[1 - \rho + \rho \cosh \frac{\beta}{\sqrt{\rho}} (Z\sqrt{\bar{r}} + \mathbf{m} \cdot \xi) \exp \left(-\frac{\beta U}{2\rho} \right) \right]. \tag{B31}$$

When $\beta \rightarrow \infty$, we have to distinguish two cases in the Z average:

$$\psi = O\left(\frac{1}{\beta}\right) + \frac{1}{\beta} \mathbb{E}_{\xi} \left(\int_{-\mathbf{m} \cdot \xi / \sqrt{\bar{r}} + U/2\sqrt{\bar{r}\rho}}^{\infty} + \int_{-\infty}^{-\mathbf{m} \cdot \xi / \sqrt{\bar{r}} - U/2\sqrt{\bar{r}\rho}} \right) \frac{dz e^{-\frac{z^2}{2}}}{\sqrt{2\pi}} \ln \left[1 - \rho + \rho \cosh \frac{\beta}{\sqrt{\rho}} (z\sqrt{\bar{r}} + \mathbf{m} \cdot \xi) e^{-\frac{\beta U}{2\rho}} \right]. \tag{B32}$$

The $O(\beta^{-1})$ instead comes from integration on the interval $[-\mathbf{m} \cdot \xi / \sqrt{\bar{r}} - U/2\sqrt{\bar{r}\rho}, -\mathbf{m} \cdot \xi / \sqrt{\bar{r}} + U/2\sqrt{\bar{r}\rho}]$ of the same integrand, which can be easily bounded.

Let us now focus on the first integral in (B32). The hyperbolic cosine and the exponential in U dominate on the other terms in the ln. Taking into account the exponential growth in the selected range of z values, the first integral can be approximated with

$$\begin{aligned} \mathbb{E}_{\xi} \int_{-\mathbf{m} \cdot \xi / \sqrt{\bar{r}} + U/2\sqrt{\bar{r}\rho}}^{\infty} \frac{dz}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} \left(\frac{Z\sqrt{\bar{r}} + \mathbf{m} \cdot \xi}{\sqrt{\rho}} - \frac{U}{2\rho} \right) &= \sqrt{\frac{\bar{r}}{2\pi\rho}} \mathbb{E}_{\xi} e^{-\frac{1}{2\bar{r}} \left(\frac{U}{2\sqrt{\rho}} - \mathbf{m} \cdot \xi \right)^2} \\ &+ \mathbb{E}_{\xi} \left(\frac{\mathbf{m} \cdot \xi}{\sqrt{\rho}} - \frac{U}{2\rho} \right) \int_{-\mathbf{m} \cdot \xi / \sqrt{\bar{r}} + U/2\sqrt{\bar{r}\rho}}^{\infty} \frac{dz}{\sqrt{2\pi}} e^{-\frac{z^2}{2}}. \end{aligned} \tag{B33}$$

The second integral in (B32) can be treated similarly. Putting all the terms together one gets

$$\begin{aligned} \frac{1}{\beta} \Phi &= -\frac{\bar{r}C}{2} + \frac{\Delta C}{2} + \frac{U}{2} + \frac{\alpha(1-t)}{2(1-C)} - \frac{1}{4} - \frac{\mathbf{m}^2}{2} + \sqrt{\frac{\bar{r}}{2\pi\rho}} \mathbb{E}_{\xi} \left[e^{-\frac{1}{2\bar{r}} \left(\frac{U}{2\sqrt{\rho}} - \mathbf{m} \cdot \xi \right)^2} + e^{-\frac{1}{2\bar{r}} \left(\frac{U}{2\sqrt{\rho}} + \mathbf{m} \cdot \xi \right)^2} \right] \\ &+ \mathbb{E}_{\xi} \frac{\mathbf{m} \cdot \xi}{2\sqrt{\rho}} \left[\operatorname{erf} \left(\frac{\mathbf{m} \cdot \xi + \frac{U}{2\sqrt{\rho}}}{\sqrt{2\bar{r}}} \right) + \operatorname{erf} \left(\frac{\mathbf{m} \cdot \xi - \frac{U}{2\sqrt{\rho}}}{\sqrt{2\bar{r}}} \right) \right] \\ &- \frac{U}{4\rho} \mathbb{E}_{\xi} \left[2 + \operatorname{erf} \left(\frac{\mathbf{m} \cdot \xi - \frac{U}{2\sqrt{\rho}}}{\sqrt{2\bar{r}}} \right) - \operatorname{erf} \left(\frac{\mathbf{m} \cdot \xi + \frac{U}{2\sqrt{\rho}}}{\sqrt{2\bar{r}}} \right) \right] + \alpha t \int_0^t d\tau \frac{C(1 - m^2(\tau))}{1 - (1 - m^2(\tau))C^2}. \end{aligned} \tag{B34}$$

Using the fact that all the parameters are evaluated at their stationary values, the previous formula can be further simplified by looking at the limiting version of the fixed point equations. In particular, we have that

$$C = \frac{1}{\sqrt{2\pi\rho\bar{r}}} \mathbb{E}_{\xi} \left[\exp \left(-\left(\frac{U/2\sqrt{\rho} - \mathbf{m} \cdot \xi}{\sqrt{2\bar{r}}} \right)^2 \right) + \exp \left(-\left(\frac{U/2\sqrt{\rho} + \mathbf{m} \cdot \xi}{\sqrt{2\bar{r}}} \right)^2 \right) \right]. \tag{B35}$$

The value of \bar{r} can be found directly from (B26) by multiplying it by β^{-2} :

$$\bar{r} = \frac{\alpha(1-t)}{(1-C)^2} + \Delta + \alpha t \int_0^t d\tau \frac{2(1 - m^2(\tau))}{1 - (1 - m^2(\tau))C^2} \left[1 + \frac{2C^2(1 - m^2(\tau))}{1 - (1 - m^2(\tau))C^2} \right]. \tag{B36}$$




FIG. 4. Logarithm of the number of restarts in the whole decimation procedure versus number of signal components for $\Delta = 0.05$, $\alpha = 0.03$. Recall that $\beta \rightarrow \infty$. Each data point is the average of 30 runs, the error bars are one standard deviation.

Furthermore, from (B23) with $v = 1$, one has

$$\mathbb{E}_{\xi} \left[\operatorname{erf} \left(\frac{U/2\sqrt{\rho} - \mathbf{m} \cdot \xi}{\sqrt{2\bar{r}}} \right) + \operatorname{erf} \left(\frac{U/2\sqrt{\rho} + \mathbf{m} \cdot \xi}{\sqrt{2\bar{r}}} \right) \right] = 2(1 - \rho). \quad (\text{B37})$$

We stress that the left-hand side of the previous equation is monotonic in U , and thus (B27) has a unique solution for fixed sparsity parameter ρ . Finally, from (B24) and (B31):

$$\mathbf{m} = \mathbb{E} \xi(X)_{Z,\xi} = \frac{\partial \psi}{\partial \mathbf{m}} = \mathbb{E}_{\xi} \frac{\xi}{2\sqrt{\rho}} \left[\operatorname{erf} \left(\frac{\mathbf{m} \cdot \xi - U/2\sqrt{\rho}}{\sqrt{2\bar{r}}} \right) + \operatorname{erf} \left(\frac{U/2\sqrt{\rho} + \mathbf{m} \cdot \xi}{\sqrt{2\bar{r}}} \right) \right]. \quad (\text{B38})$$

If we insert these conditions in (B34), we get

$$\frac{\Phi}{\beta} = \frac{\alpha(1-t)}{2(1-C)^2} + \Delta C - \frac{1}{4} + \frac{\mathbf{m}^2}{2} + 2\alpha t \int_0^t d\tau \frac{C(1-m^2(\tau))}{[1 - (1-m^2(\tau))C^2]^2}. \quad (\text{B39})$$

APPENDIX C: GROUND-STATE ORACLE FOR BINARY SPINS

Here we provide the pseudocode for the ground-state oracle for binary spins, based on an adaptation of simulated annealing. Algorithm 1 is itself run multiple times, typically five times or more, and we chose the outcome that required the least number of restarts.

Algorithm 1. Ground-state oracle with $P_{\xi}(x) = \frac{1}{2}(\delta_{x,-1} + \delta_{x,1})$.

Require: N, P (or α), \mathbf{Y} , threshold ($\in \mathbb{R}_+$), maxr ($\in \mathbb{N}$)

restarts $\leftarrow 0$

$\mathbf{S} \leftarrow N \times N$ zeros matrix

while $i \leq P$ **do**

 s, trials \leftarrow SA(N, \mathbf{Y} , threshold, maxr, restarts) \triangleright See the code of the SA routine below

 restarts \leftarrow restarts + trials

if (trials > 299 **or** restarts > maxr) **then**

break and start over

end if

$\mathbf{Y} \leftarrow \mathbf{Y} - \frac{\mathbf{ss}^T}{\sqrt{N}}$

$\mathbf{S}_i \leftarrow \mathbf{s}$ $\triangleright \mathbf{S}_i = i$ th column vector in \mathbf{S}

 threshold \leftarrow threshold $\cdot 0.9975$

$i \leftarrow i + 1$

end while

Algorithm 1 implements the decimation procedure, whereas Algorithm 2 is the simulated annealing. Notice that in the latter we have introduced a restarting criterion based on a moving energy threshold that lowers any time we find a candidate pattern that is good enough, in terms of energy, or we do not find anything acceptable in 20 trials. The number of restarts needed increases exponentially with the size of the system, as in Fig. 4, but the algorithm proved to be efficient up to sizes $N \simeq 1500$ for any noise level Δ we have tested, and above $N = 2000$ when the noise is not too close to its critical value.

Algorithm 2. Simulated annealing (SA).

Require: N , \mathbf{Y} , threshold ($\in \mathbb{R}_+$), maxr ($\in \mathbb{N}$), restarts ($\in \mathbb{N}$)

```

itry  $\leftarrow$  0
found=False
while itry < 300 do
  stop  $\leftarrow$  0
   $\beta \leftarrow$  0
   $\mathbf{s} \leftarrow$  random sample from  $\prod_{i=1}^N P_{\xi}$ 
  itry  $\leftarrow$  itry + 1
  if itry + restarts > maxr then
    returns, itry
  end if
  if itry % 20 = 0 then
    threshold  $\leftarrow$  threshold  $\cdot$  0.9975
  end if
  while  $k < 200$  do
     $k \leftarrow k + 1$ 
     $\beta \leftarrow 1 + \frac{k}{200} \cdot 99$ 
     $\mathbf{h} \leftarrow \frac{\mathbf{Y}}{\sqrt{N}} \mathbf{s}$ 
    define  $ss_i \leftarrow -s_i$  with probability  $\frac{1}{1+e^{-2\beta h_i}} \forall i$ 
    if  $\mathbf{s} - \mathbf{ss} = \mathbf{0}$  then
      stop  $\leftarrow$  stop + 1
      if stop > 5 then
        if  $\frac{1}{2} \mathbf{s}^T \frac{\mathbf{Y}}{\sqrt{N}} \mathbf{s} >$  threshold then
          return  $\mathbf{s}$ , itry
        else
          break
        end if
      end if
    else
      stop  $\leftarrow$  0
       $\mathbf{s} \leftarrow \mathbf{ss}$ 
    end if
  end while
end while

```

\triangleright When the search does not move for five consecutive times.

\triangleright wrong energy, try again

-
- [1] B. A. Olshausen and D. J. Field, *Nature (London)* **381**, 607 (1996).
- [2] B. A. Olshausen and D. J. Field, *Vision Res.* **37**, 3311 (1997).
- [3] K. Kreuz-Delgado, J. F. Murray, B. D. Rao, K. Engan, T.-W. Lee, and T. J. Sejnowski, *Neural Comput.* **15**, 349 (2003).
- [4] Y. Bengio, A. Courville, and P. Vincent, *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 1798 (2013).
- [5] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, *Cognit. Sci.* **9**, 147 (1985).
- [6] R. Salakhutdinov, A. Mnih, and G. Hinton, *Proceedings of the 24th International Conference on Machine Learning, ICML'07* (Association for Computing Machinery, New York, 2007), pp. 791–798.
- [7] J. Mairal, M. Elad, and G. Sapiro, *IEEE Trans. Image Process.* **17**, 53 (2008).
- [8] J. Wright, Y. Ma, J. Mairal, G. Sapiro, T. S. Huang, and S. Yan, *Proc. IEEE* **98**, 1031 (2010).
- [9] M. Elad and M. Aharon, *Trans. Img. Proc.* **15**, 3736 (2006).
- [10] A. Mnih and R. R. Salakhutdinov, in *Advances in Neural Information Processing Systems*, edited by J. Platt, D. Koller, Y. Singer, and S. Roweis (Curran Associates, Inc., New York, USA, 2007), Vol. 20.
- [11] A. Manoel, F. Krzakala, M. Mézard, and L. Zdeborová, *2017 IEEE International Symposium on Information Theory (ISIT)* (IEEE Press, New York City, USA, 2017), pp. 2098–2102.
- [12] I. M. Johnstone, *Ann. Stat.* **29**, 295 (2001).
- [13] T. Lesieur, F. Krzakala, and L. Zdeborová, *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)* (IEEE, New York City, USA, 2015), pp. 680–687.
- [14] F. Krzakala, J. Xu, and L. Zdeborová, *2016 IEEE Information Theory Workshop (ITW)* (IEEE, New York City, USA, 2016), pp. 71–75.
- [15] J. Barbier and N. Macris, *J. Phys. A: Math. Theor.* **52**, 294002 (2019).
- [16] A. E. Alaoui, F. Krzakala, and M. Jordan, *Ann. Stat.* **48**, 863 (2020).

- [17] J. Baik, G. Ben-Arous, and S. Péché, *Ann. Probab.* **33**, 1643 (2005).
- [18] F. Benaych-Georges and R. R. Nadakuditi, *Adv. Math.* **227**, 494 (2011).
- [19] Y. Kabashima, *J. Phys. A: Math. Gen.* **36**, 11111 (2003).
- [20] D. L. Donoho, A. Maleki, and A. Montanari, *Proc. Natl. Acad. Sci. USA* **106**, 18914 (2009).
- [21] J. Bun, R. Allez, J.-P. Bouchaud, and M. Potters, *IEEE Trans. Inf. Theory* **62**, 7475 (2016).
- [22] Y. Kabashima, F. Krzakala, M. Mézard, A. Sakata, and L. Zdeborová, *IEEE Trans. Inf. Theory* **62**, 4228 (2016).
- [23] A. Maillard, F. Krzakala, M. Mézard, and L. Zdeborová, *J. Stat. Mech.: Theory Exp.* (2022) 083301.
- [24] J. Barbier and N. Macris, *Phys. Rev. E* **106**, 024136 (2022).
- [25] J. J. Hopfield, D. I. Feinstein, and R. G. Palmer, *Nature (London)* **304**, 158 (1983).
- [26] J. Van Hemmen, L. Ioffe, R. Kühn, and M. Vaas, *Physica A* **163**, 386 (1990).
- [27] M. Benedetti, E. Ventura, E. Marinari, G. Ruocco, and F. Zamponi, *J. Chem. Phys.* **156**, 104107 (2022).
- [28] Y. Dandi, L. Stephan, F. Krzakala, B. Loureiro, and L. Zdeborová, [arXiv:2302.08933](https://arxiv.org/abs/2302.08933).
- [29] J. J. Hopfield, *Proc. Natl. Acad. Sci. USA* **79**, 2554 (1982).
- [30] D. J. Amit, H. Gutfreund, and H. Sompolinsky, *Phys. Rev. A* **32**, 1007 (1985).
- [31] D. J. Amit, H. Gutfreund, and H. Sompolinsky, *Phys. Rev. Lett.* **55**, 1530 (1985).
- [32] M. Mézard, G. Parisi, and M. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1986), <https://www.worldscientific.com/doi/pdf/10.1142/0271>.
- [33] S. Goldt, F. Krzakala, N. Brunel, and L. Zdeborová, *PLOS Comput. Biol.* **19**, e1010813 (2023).