

Copycat perceptron: Smashing barriers through collective learning

Giovanni Catania¹, Aurélien Decelle^{1,2} and Beatriz Seoane²

¹*Departamento de Física Teórica, Universidad Complutense de Madrid, 28040 Madrid, Spain*

²*Université Paris-Saclay, CNRS, Inria TAU team, LISN, 91190 Gif-sur-Yvette, France*



(Received 12 December 2023; accepted 3 June 2024; published 28 June 2024)

We characterize the equilibrium properties of a model of γ coupled binary perceptrons in the teacher-student scenario, subject to a suitable cost function, with an explicit ferromagnetic coupling proportional to the Hamming distance between the students' weights. In contrast to recent works, we analyze a more general setting in which thermal noise is present that affects each student's generalization performance. In the nonzero temperature regime, we find that the coupling of replicas leads to a bend of the phase diagram towards smaller values of α : This suggests that the free entropy landscape gets smoother around the solution with perfect generalization (i.e., the teacher) at a fixed fraction of examples, allowing standard thermal updating algorithms such as Simulated Annealing to easily reach the teacher solution and avoid getting trapped in metastable states as happens in the unreplicated case, even in the computationally *easy* regime of the inference phase diagram. These results provide additional analytic and numerical evidence for the recently conjectured Bayes-optimal property of Replicated Simulated Annealing for a sufficient number of replicas. From a learning perspective, these results also suggest that multiple students working together (in this case reviewing the same data) are able to learn the same rule both significantly faster and with fewer examples, a property that could be exploited in the context of cooperative and federated learning.

DOI: [10.1103/PhysRevE.109.065313](https://doi.org/10.1103/PhysRevE.109.065313)

I. INTRODUCTION

Statistical mechanics provides valuable tools for understanding machine learning. The goal is to describe the typical behavior of a neural network with respect to global parameters such as the training set size [1,2] or the gradient descent noise [3,4]. This approach helps us to understand the conditions that favor better learning and when achieving good performance is impossible. An impressive example of this approach's effectiveness is the analysis of the solution space as a function of the proportion of clauses in k -SAT in classical combinatorial optimization problems [5–7].

At zero temperature, accurately classifying labeled data with random labels can be seen as a constraint satisfaction problem (CSP). The examples range from perceptrons [1,8] to more complex architectures such as the Committee Machine [9,10], Support Vector Machines [11], multilayered perceptrons [12,13], or even continuous optimization problems in condensed matter [14]. Surprisingly, despite the vast number of potential solutions with poor generalization, perceptrons and deep neural networks excel in classification tasks. This suggests that the standard training methods do not explore the entire space of quasi-optimal CSP solutions, but use a more efficient approach.

In recent years, researchers have introduced a theoretical framework based on the concept of *local entropy* to better understand the effectiveness of training algorithms and to find solutions that generalize well. This concept, along with the coupled-replicas strategy, has been thoroughly investigated in several studies [15–18], as well as in its connection to quantum annealing protocols [19]. Moreover, a recent work [20] has provided convincing evidence that coupling replicas

can help in identifying favorable local minima and avoiding entrapment in glassy states, as shown in the graph coloring problem.

In this study we focus on the paradigmatic binary perceptron model in the *teacher-student* scenario, a well-established example of a planted inference problem [2,3,21] where a *student* perceptron attempts to learn the classification rule from the examples given by a *teacher* perceptron. In this setting, the ratio between the number of examples and the number of parameters (α) acts as a signal-to-noise ratio.

Previous studies have shown that this model, unlike its continuous counterpart, exhibits a first-order phase transition at zero temperature, corresponding to a sharp decrease in the generalization error as the number of training examples increases [2,22,23]. Recent works have focused on describing this storage performance at zero temperature in the so-called *robust ensemble* [16,18,24], where multiple replicas of the same model interact through a ferromagnetic coupling to favor solutions with high local entropy. This ensemble is closely related to the Franz-Parisi potential in glassy physics [25,26]. Similarly, the idea of coupled-replica models has been used to design efficient cooperative local search strategies in simple neural networks [27] or in deep learning contexts [28].

In this paper, we propose a more general theoretical approach in which we derive the entire (α, T) phase diagram of the binary perceptron in the robust ensemble as a function of the number of coupled replicas γ or the coupling γ . This phase diagram can be used to understand how the structure of the solution space evolves during the training process and how the γ coupling favors broader and flatter landscapes. We show that the phase diagram allows us to propose more effective annealing strategies, especially in non-Bayesian optimal

scenarios. Moreover, we show that the so-called inference-*easy* region is plagued by subdominant 1RSB metastable states that make it practically impossible for annealing strategies to find the teacher in a reasonable time. Using the phase diagram obtained with the replica trick, we show that coupled perceptrons can easily learn the teacher solution in this region by avoiding the 1RSB states.

II. MODEL

The model we consider is defined by a number y of binary perceptron students and a teacher perceptron. Each student $u \in \{1, \dots, y\}$ is parameterized by a weight vector with N components $\mathbf{w}_u = \{w_{i,u}\}_{i=1}^N$ with $w_{i,u} \in \{-1, 1\}$. Students are given the same set of M labeled data points μ , each represented by a pair (ξ^μ, σ_0^μ) : the first is the *training sample* ξ^μ , an N -dimensional vector with binary entries $\xi_i^\mu \in \{-1, 1\}$, and σ_0^μ is the corresponding *label*, assigned by a *teacher* perceptron with weight vector \mathbf{w}_0 . The input-output relation for both the teacher and the students is defined as follows:

$$\sigma_u^\mu = \text{sign}(\omega_u^\mu) \quad \text{and} \quad \omega_u^\mu = \frac{\mathbf{w}_u \cdot \xi^\mu}{\sqrt{N}}, \quad u \in \{0, \dots, y\} \quad (1)$$

for $\mu \in \{1, \dots, M\}$, where the multidot symbol denotes the scalar product and $u = 0$ refers to the teacher. The normalization factor in Eq. (1) ensures that the energy of the model is extensive in the system size N . At fixed planted configuration, the labels generated by the teacher are noiseless according to (1) so that the Bayes-optimal temperature is $T = 0$.

Each student adjusts their weights to minimize an appropriate cost function, which accounts for the number of misclassified examples, i.e., those where the assigned label differs from that of the teacher. Moreover, the students will interact through a pairwise ferromagnetic coupling that favors configurations in which students have a high mutual overlap. In other words: Imitation between students is encouraged.

The learning is defined by specifying the Hamiltonian of the system. Following [29], we first define the *stability parameters* as $\Delta_u^\mu = \sigma_0^\mu \omega_u^\mu$, so that $\Delta_u^\mu > 0$ whenever the input μ is correctly classified by the student u . Second, we define an arbitrary potential $V(\Delta) = (-\Delta)^\nu \Theta(-\Delta)$, where $\nu \in \mathbb{Z}^+$ and $\Theta(x)$ denotes the Heaviside step function. The potential assigns 0 energy to the weight vectors that correctly classify a given example (i.e., those for which $\Delta_u^\mu > 0$) and a positive cost $\propto (-\Delta)^\nu$ to misclassified ones (i.e., when $\sigma_u^\mu \neq \sigma_0^\mu$).

The Hamiltonian of the model can be written as

$$\mathcal{H}(\{\mathbf{w}_u\}_{u=1}^y | \mathbf{w}^0, \mathbb{D}) = \sum_{u,\mu}^{y,M} V(\Delta_u^\mu) - \frac{\gamma}{\beta y} \sum_{u < v} \mathbf{w}_u \cdot \mathbf{w}_v, \quad (2)$$

where $\beta = T^{-1}$ and $\mathbb{D} = \{(\xi^\mu, \sigma_0^\mu)\}_{\mu=1}^M$ denotes the training set of labeled examples. The second term quantifies the interaction between student pairs, and it is proportional to the Hamming distance between their weights. Equation (2) describes a system in which y students try to learn the same teacher's rule, while interacting with a ferromagnetic potential (tuned by γ): The latter favors configurations in which students have a similar \mathbf{w} , i.e., students are actively encouraged to be “inspired” by their peer perceptrons. Different choices of the exponent ν lead to distinct learning dynamics, each with

unique convergence properties and nonzero temperature phase diagrams in the thermodynamic limit [2,30,31]. In this article, we will focus on the case $\nu=1$, the so-called “perceptron rule,” originally introduced by Rosenblatt [32] for continuous weights. The case $\nu=0$ is less interesting as the system remains frozen at all temperatures [30]; the case $\nu \geq 2$ will be the subject of future studies.

We consider a scenario where each student fluctuates with a heat bath at a given β and a fixed interaction strength between students, γ . We seek to characterize whether and how the different operating regimes of the binary perceptron change when considering multiple interacting replicas. In other words, our goal is to understand and model how the collective or cooperative learning of multiple students differs at a fundamental level from that of a single student.

III. MEAN-FIELD THEORY

To characterize the equilibrium properties of the model (2), we need to compute the quenched free entropy in the thermodynamic limit where both $N, M \rightarrow \infty$ keeping the ratio $\alpha = M/N$ finite. The set of patterns $\{\xi^\mu\}_{\mu=1}^M$ and the teacher's weight \mathbf{w}_0 represent the quenched disorder to average over. As each student reviews the same examples, the disorder is the same for all students. The starting point of the derivation is the partition function $Z = \sum_{\mathbf{w}_u} e^{-\beta \mathcal{H}}$. In our definition for the Hamiltonian (2), the thermal noise affects only the first term, but not the coupling between the students: this choice, consistent with previous works [20], allows a finite γ between the replicas even in the $T \rightarrow 0$ limit, where the model is reduced to a CSP in the space of students' weights that satisfies the constraints imposed by the teacher. The physical quantity of interest is the quenched free entropy density

$$\mathcal{G}(\alpha, \beta, \gamma, y) = \lim_{N \rightarrow \infty} \frac{1}{yN} \langle \log Z \rangle_{\mathbb{D}, \mathbf{w}^0}. \quad (3)$$

For simplicity, we restrict ourselves to finite values of y , although in principle the limit $y \rightarrow \infty$ could also be analyzed, since the coupling between the students is rescaled so that (3) is an intensive quantity with respect to y . The computation of (3) can be performed using the usual *replica trick* in spin-glass theory [33,34]. The main difficulty here lies in dealing with the coupled students. Formally, they are “real” replicas: they share the same quenched disorder as the replicas of the replica trick, but they also interact through an explicit pairwise coupling. The mean-field theory requires the introduction of a set of order parameters, to be evaluated with the saddle-point method in the thermodynamic limit. These are the overlap each student has with the teacher vector (in each replica $a \in \{1, \dots, n\}$) and the pairwise overlap between two students u, v in replicas a, b , respectively:

$$R_a^u = \frac{1}{N} \sum_i w_{i,u}^{(a)} w_i^0 \quad \text{and} \quad q_{ab}^{uv} = \frac{1}{N} \sum_i w_{i,u}^{(a)} w_{i,v}^{(b)}. \quad (4)$$

The simplest Replica Symmetric (RS) ansatz imposes a permutation symmetry between the overlaps in the replica space, i.e., $q_{ab}^{uv} = q$, $\forall a \neq b$. The diagonal blocks $a=b$ represent the average correlation between the students, which in general may depend on the topology of the interactions in the replicated space. Assuming a fully connected topology as in




FIG. 1. (α, T) equilibrium phase diagram of the single, non-replicated, binary perceptron (top panel) and the replicated perceptron according to Eq. (2) for a fixed coupling $\gamma = 1$ and y multiple replicas (bottom panel). The color shades refer respectively to the *impossible* (red), *hard* (yellow), and *easy* (green) inference phases as extrapolated from $T = 0$. The (a)–(e) regions are discussed in the text.

(2), it is natural to assume a uniform overlap between students $u \neq v$ in the *same* replica a , so that $q_{aa}^{uv} = \delta_{uv} + p(1 - \delta_{uv}) \forall a$ (the diagonal constraint δ_{uv} comes from the binary nature of the weights). Due to the different nature of the replicas a, b and the students u, v , the two values of the overlap used in this ansatz are expected to be different: in particular, $p \geq q$ due to the ferromagnetic coupling γ in the Hamiltonian (the equality holds at $\gamma = 0$). As far as the *signal* term R_a^u is concerned, the RS ansatz implies $R_a^u = R \forall u, a$. A complete derivation of the explicit form of \mathcal{G} in (3) under this ansatz is given in Appendix A. In the thermodynamic limit, the equilibrium behavior of the model is determined by the maxima of this free entropy, which are found by imposing stationarity w.r.t. the order parameters. In particular, the teacher configuration (i.e., the planted solution) corresponds to a free-entropy maximum with $R \simeq 1$ (the strict equality holding at $T \rightarrow 0$).

IV. RESULTS

It is instructive to first discuss the phase diagram of the single perceptron (i.e., $y = 1$) under the RS ansatz, shown in Fig. 1 (top), and previously derived in [30]. For $0 < T \lesssim 0.2$ five different equilibrium regimes are found as the fraction of training examples α is increased. Above the dashed line [region (a)], the free entropy has a unique maximum with $R < 1$, which corresponds to a solution with imperfect *generalization*. Below the dashed line [region (b)], the teacher's solution ($R \simeq 1$) appears as a metastable (i.e., subdominant in the free entropy) fixed point, and it becomes dominant after crossing the dotted line [region (c)]. Finally, beyond

the spinodal solid line, in region (e), the poor generalization fixed point disappears and only the teacher solution remains. Region (d) (below the dash-dotted line) corresponds instead to the *spin-glass* phase where the solution with poor generalization ($R < 1$) has negative RS entropy. This means that a replica symmetry-breaking ansatz (RSB, with one or more steps) would be required to correctly capture the model's thermodynamics. The colored areas in Fig. 1 stand for the three inference phases present in the $T = 0$ phase diagram [3]: the *impossible* (red) phase for $\alpha < \alpha_{IT} \approx 1.24$, where the teacher solution is surrounded by an exponential number of ground states with 0 training error and therefore it is subdominant in terms of entropy [16]; the *hard* (yellow) phase $\alpha_{IT} < \alpha < \alpha_c \approx 1.49$ where although the teacher is now dominant (after the first-order transition at α_{IT}) there still exists a subextensive (in N) amount of ground states with 0 training error, with a disconnected 1RSB-like structure; finally, the *easy* (green) phase for $\alpha > \alpha_c$ where according to the RS ansatz only the teacher configuration remains as a stable maximum of the free entropy. At $T = 0$, the hard-easy boundary α_c also corresponds to the algorithmic threshold for the Approximate Message Passing (AMP) algorithm (a numerical check on finite-size systems is discussed in Appendix D).

From the phase diagram in Fig. 1 (top), one should be able in principle to infer the performance of a Simulated Annealing (SA) experiment (or, equivalently, a learning process modeled by a slow decrease in the allowable number of errors) for a given value of α . In particular, for $\alpha \leq \alpha_c$, SA (or AMP, or any other known polynomial algorithm) is not able to find the teacher solution, while for $\alpha > \alpha_c$, SA should in principle easily find the planted solution because the equilibrium measure is dominated by configurations close to the teacher vector. However, as pointed out in the next section, SA gets trapped into spurious minima of the energy with nonzero training (and generalization) error even in a finite interval in the easy phase $\alpha > \alpha_c$. At nonzero temperature, the presence of spurious minima can be tracked by computing the dynamic transition temperature T_{dyn} [25,26], shown in red in Fig. 1 (top): this is the temperature at which spurious local minima with a 1-RSB structure appear, which can act as algorithmic traps for SA. Additional details about this quantity and how to compute it are given in Appendix B 2 a, obtained using the procedure discussed in Refs. [25,35]. Following the reason of [20], SA should always be trapped into spurious minima of the energy as long as the dynamic transition (red line) is met before the spinodal of the imperfect generalization solution (solid black line in Fig. 1) upon decreasing the temperature; therefore, the intersection between the two should correspond to the algorithmic threshold for SA, so that $\alpha_c^{SA} \approx 1.56$. In particular, for $\alpha > \alpha_c^{SA}$, SA should always find the solution by melting towards the planted configuration when the spinodal [solid black line in Fig. 1 (top)] is crossed. For the sake of completeness, we recall that another 1RSB-like analysis was performed in [2] in the $T = 0$ model: the authors revealed the presence of 1-RSB frozen metastable states with poor generalization, which cease to exist at $\alpha_c^{RSB} \approx 1.628$. The presence of such 1RSB states at $T = 0$ does not affect the capabilities Bayes-optimal algorithms such as AMP in the regime $\alpha > \alpha_c$ as discussed, e.g., in Ref. [36]: indeed, the algorithmic threshold of AMP is exactly $\alpha_c \approx 1.49$. Conversely, a thermal

algorithm like SA is severely affected by these metastable states, although the interplay between the threshold computed in [2], and the dynamic transition temperature is not yet clear. Apart on this difference between the 1RSB analysis which is not yet clarified, the general conclusion is that the higher $\alpha > \alpha^{\text{SA}}$ is, the easier it should be to find the teacher, as the free entropy landscape becomes smoother as more training examples are available: in particular, on the right side of the intersection between the two spinodals and the transition line (i.e., $\alpha \gtrsim 2.0$), there is no phase transition at all and the generalization error becomes a smooth decreasing function of temperature: in this regime, the energy landscape becomes trivial and the system behaves as an effective paramagnet with a strong external field pointing towards the planted configuration.

Figure 1 (bottom) illustrates how the spinodal lines change when y coupled replicas are considered (in this case for $\gamma = 1$). The most interesting phenomenon is that the right spinodal line, which marks the transition to perfect generalization, is bending towards lower values of α . Therefore, in the range of $\alpha > \alpha_c$ the free entropy landscape becomes smoother just by increasing the number of coupled replicas. In terms of performing a temperature annealing, this means that the student perceptrons should not encounter the poorly generalizing glassy states, as they would melt to the teacher solution before the dynamic transition. In practice, this phase diagram shows us that many coupled students need to go through fewer data examples to perfectly deduce the teacher's rule. The effect of the coupling γ in the phase diagram is discussed in Appendix A 2 (see Fig. 4). A more detailed analysis on the disappearance of metastable glassy states in the coupled system would require a 1-RSB analysis, which is left for future investigations.

V. NUMERICAL EXPERIMENTS

We now investigate numerically the effects of the shifting of the critical lines associated with the coupling between students when training a binary perceptron model (2) via SA with a fixed cooling rate. Our SA training protocol is constructed as follows: We initialize student weights at high temperature. At each training update, we perform a Monte Carlo sweep (an update of the entire system in random order) and reduce the temperature by η , keeping α fixed. At each α , we repeat the same process 100 times, using different teacher perceptrons to average out the disorder fluctuations (additional implementation details are given in Appendix C). Figure 2(a) shows the results for a very slow cooling rate $\eta = 10^{-5}$ on different number of replicas y . It is striking that the single perceptron (black dots) is not able to find the teacher solution up to $\alpha \approx 1.6$, i.e., deep into the *easy* phase: this is in agreement with the presence of metastable states, as found in the 1-RSB analysis, which trap the annealing dynamics and prevent the system from melting towards the teacher solution. A numerical determination of the precise algorithmic threshold for SA on the single system is, however, out of reach for computational reasons: first, because the system is dense, so that a full MC swipe requires $O(\alpha N^2)$ operations; second, increasing N in order to perform a finite-size scaling for the estimation of α_c^{SA} would require to rescale the annealing rate

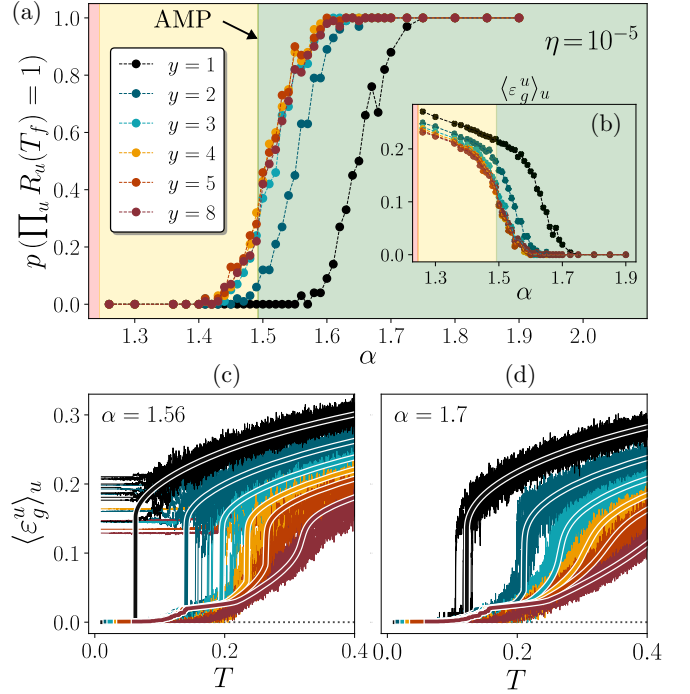


FIG. 2. Numerical performances of replicated SA (RSA) with an annealing rate $\eta = 10^{-5}$. With black lines we show the performance of the single perceptron and with colors that of the coupled perceptron with different y . (a) Empirical probability (from 100 training instances) of finding the teacher configuration at the end of the annealing process, as a function of α , for $N=2001$, $\gamma=1$. The inset (b) shows the corresponding mean generalization error $\varepsilon_g = \pi^{-1} \arccos(R)$. The colored regions are the same as in Fig. 1: in particular, the hard-easy phase boundary corresponds to the algorithmic threshold for AMP at $T = 0$. (c), (d) Examples of typical annealing trajectories for two values of α (10 training trajectories are shown for each y). The settings and color coding are the same as in the top panel. The white outlined lines show the analytical result obtained from solving the RS self-consistent equations, starting from a poorly generalizing solution ($R \ll 1$) at $T = 0.5$, and following the fixed point as the temperature is linearly decreased.

with a certain power law of N [20], further increasing the computational cost. The addition of replicas instead leads to a net shift in the probability of finding the planted solution, which appears to saturate very quickly with y and to approach the theoretical threshold α_c . Such a performance difference between SA and RSA (i.e., SA on the replicated system) is even more remarkable at much faster annealing rates η (see Figs. 5 and 6 in Appendix C), where systems with larger y find the solution faster. We also observe that even when RSA does not find the planted configuration (e.g., in the inference-*hard* phase), it still finds more optimal solutions [in terms of generalization error; see Fig. 2(b)] compared to the nonreplicated system, as previously noted in [16]. The change in the flatterness of the landscape with y is particularly clear in Fig. 2 (bottom panel): the RSA trajectories show increasingly smoother behavior (in terms of generalization error) as the number of students increases, but at the same value of α . Note also the good agreement between simulations and theory with white outlined lines; the disagreement between theory and

simulations for $y = 1$ and $\alpha = 1.56$ [Fig. 2(c)] is expected to be due to the failure of the theory at the RS level.

VI. DISCUSSION

In this paper we have illustrated how the phase diagram of the binary perceptron in the teacher-student scenario changes as more and more students try to learn the same teacher rule while being coupled to each other. From this phase diagram, we can draw conclusions not only about the number of examples required for perfect recovery (as a function of the number of collaborating students), which is much smaller than for a single perceptron, but also about the performance of a thermal learning procedures such as SA or eventually stochastic gradient descent (on models with continuous weights).

The coupled perceptron model discussed here provides a toy model to shed light on the effectiveness of students' collective learning in acquiring a particular set of predetermined rules, as exploited in recent real-life experiences (see, e.g., [37]). Within the framework of this rudimentary model, further generalizations could explore the question of whether diversifying the examples for each student improves the learning experience or whether alternative learning paths could accelerate learning. In particular, our approach could be used to rationalize the impact of federated or collaborative learning in machine learning, a decentralized approach to training models [38]. In both contexts, our phase diagrams facilitate the determination of the ideal number of collaborating teams, the best number of examples, or the appropriate learning pace to ensure optimal learning.

From a theoretical point of view, we believe that this work contributes to a better understanding of the role of the robust ensemble, which has been recently proposed to develop novel algorithmic schemes for solving constraint satisfaction problems. In particular, our results seem to confirm that RSA is a generic and robust inference algorithm that has a close-to-Bayes optimal threshold for a sufficient number of coupled replicas, as recently argued in another setup [20]. Our results also suggest that coupled neural networks can in principle perform better in terms of generalization error for a fixed amount of data. It would be interesting to verify this effect on more complex architectures for a supervised learning task or on other examples of planted models with a hard inference phase. In this regard, a more analytical study of the bending of the phase diagram and the resulting change in the nature of the phase transition in the inference *easy* phase in the limit $y \rightarrow \infty$ will be the subject of a follow-up work.

ACKNOWLEDGMENTS

We thank Carlo Lucibello, Federico Ricci-Tersenghi, Valentina Ros, and Pierfrancesco Urbani for useful discussions and suggestions regarding the article. We further thank Carlo Lucibello for sharing an AMP implementation which we used for comparison. The authors acknowledge financial support by the Comunidad de Madrid and the Complutense University of Madrid (UCM) through the Atracción de Talento program (Refs. 2019-T1/TIC-13298), the Banco Santander and the UCM (Grant No. PR44/21-29937), and to project PID2021-125506NA-I00 financed by the Ministerio de Economía y Competitividad, Agencia Estatal de Investigación (MICIU/AEI/10.13039/501100011033) and the Fondo Europeo de Desarrollo Regional (FEDER, UE).

APPENDIX A: DERIVATION OF QUENCHED FREE ENTROPY

In this Appendix we discuss how to compute the averaged quenched free energy for the model (2). As standard in spin-glass models with quenched disorder, the first step requires the introduction of a number n of replicas, whose limit $n \rightarrow 0$ is taken afterwards. These replicas differ from the ones in the original Hamiltonian because they are independent and no explicit coupling is present between them. In the following derivation, and in order to avoid confusion between the two sets of replicas, we always use indices $a, b \in \{1, \dots, n\}$ to denote “fake” replicas and $u, v \in \{1, \dots, y\}$ to index students (i.e., real replicas in the original Hamiltonian). Replicating all the degrees of freedom $w_{i,u}$ n times (with integer n) we can write the replicated partition function as

$$Z^n = \sum_{\mathbf{w}_u^{(a)}} \exp \left[-\beta \sum_{a,u,\mu} V(\Delta_{u,a}^\mu) + \frac{\gamma}{y} \sum_a \sum_{u < v} \mathbf{w}_u^{(a)} \cdot \mathbf{w}_v^{(a)} \right]. \quad (\text{A1})$$

A schematic representation of the model is given in Fig. 3. We start by introducing the definition of the stabilities

$$\Delta_{u,a}^\mu = \sigma_0^\mu \frac{\mathbf{w}_u^{(a)} \cdot \boldsymbol{\xi}^\mu}{\sqrt{N}}, \quad (\text{A2})$$

with σ_0^μ given by Eq. (1). Enforcing definitions (A2) and (1) by using delta functions and exploiting their Fourier representation, the partition function (A1) can be rewritten as

$$Z^n = \sum_{\mathbf{w}_u^{(a)}} \int \prod_{\mu} d\omega_0^\mu d\hat{\omega}_0^\mu \int \prod_{\mu,a,u} d\Delta_{u,a}^\mu d\hat{\Delta}_{u,a}^\mu \exp \left[-\beta \sum_{a,u,\mu} V(\Delta_{u,a}^\mu) + \frac{\gamma}{y} \sum_a \sum_{u < v} \mathbf{w}_u^{(a)} \cdot \mathbf{w}_v^{(a)} \right] \\ \times \exp \left[i \sum_{\mu} \omega_0^\mu \hat{\omega}_0^\mu - \frac{i}{\sqrt{N}} \sum_{\mu} \hat{\omega}_0^\mu (\mathbf{w}^0 \cdot \boldsymbol{\xi}^\mu) + i \sum_{a,\mu,u} \Delta_{u,a}^\mu \hat{\Delta}_{u,a}^\mu - \frac{i}{\sqrt{N}} \sum_{a,\mu,u} \hat{\Delta}_{u,a}^\mu (\mathbf{w}_u^{(a)} \cdot \boldsymbol{\xi}^\mu) \sigma_0^\mu (\omega_0^\mu) \right], \quad (\text{A3})$$

where $\hat{\omega}_0^\mu$ (resp. $\hat{\Delta}_{u,a}^\mu$) are the conjugate variables of ω_0^μ (resp. $\Delta_{u,a}^\mu$) introduced through a Fourier transform. The dependency of the teacher label on its input simply follows from the perceptron classification rule, i.e., $\sigma_0^\mu (\omega_0^\mu) = \text{sign}(\omega_0^\mu)$, although in the rest of the calculation we will drop this dependency for notation convenience. It is now easy to perform the average over

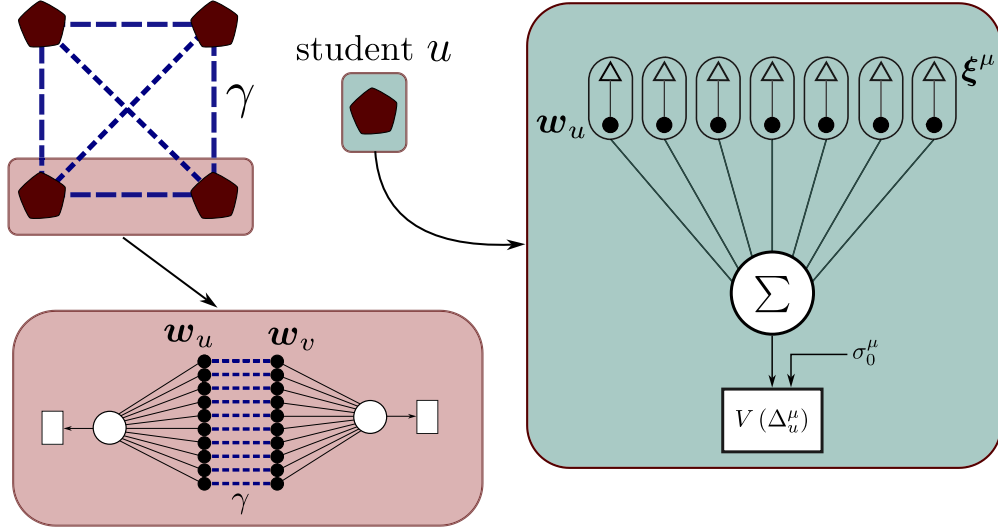


FIG. 3. Schematic representation of the model defined by Eq. (2) for $y = 4$ student perceptrons interacting over a fully connected graph. The dashed blue lines represent the coupling γ between students.

the disorder given by the pattern components. As for now we do not perform the average over the teacher weight vector, but it will become trivial in the final expression. As specified in the main text, we assume the pattern components to be i.i.d. with binary entries, so that $\xi_i^\mu \in \{-1, 1\}$ with equal probability. The average concerns only the second and fourth terms in the second line of (A3):

$$\begin{aligned}
 \langle e^{-\frac{i}{\sqrt{N}} \sum_\mu \hat{\omega}_0^\mu (\mathbf{w}^0 \cdot \boldsymbol{\xi}^\mu) - \frac{i}{\sqrt{N}} \sum_{a,\mu,u} \hat{\Delta}_{u,a}^\mu \sigma_0^\mu (\mathbf{w}_u^{(a)} \cdot \boldsymbol{\xi}^\mu)} \rangle_{\{\boldsymbol{\xi}^\mu\}_{\mu=1}^M} &= \prod_{i,\mu} \left\langle \exp \left[-\frac{i}{\sqrt{N}} \xi_i^\mu \left(\hat{\omega}_0^\mu w_i^0 + \sigma_0^\mu \sum_{a,u} \hat{\Delta}_{u,a}^\mu w_{i,u}^{(a)} \right) \right] \right\rangle_{\xi_i^\mu} \\
 &= \prod_{i,\mu} 2 \cosh \left[\frac{i}{\sqrt{N}} \left(\hat{\omega}_0^\mu w_i^0 + \sigma_0^\mu \sum_{a,u} \hat{\Delta}_{u,a}^\mu w_{i,u}^{(a)} \right) \right] \\
 &\approx \exp \left[-\frac{1}{2N} \sum_{i,\mu} \left(\hat{\omega}_0^\mu w_i^0 + \sigma_0^\mu \sum_{a,u} \hat{\Delta}_{u,a}^\mu w_{i,u}^{(a)} \right)^2 \right], \tag{A4}
 \end{aligned}$$

where in the first line we use the fact that pattern components are i.i.d., and in the last line we expanded for $N \rightarrow \infty$, keeping only the first order, the other ones being subdominant in the thermodynamic limit.

As usual, the disorder average results into an effective coupling between replicas a, b , and another coupling will also be taken into account between students in the same replica a on top of the explicit one in the Hamiltonian. We can now introduce a set of order parameters, namely, the overlap between student u (in replica a) with the teacher and the two-replica overlap between two student vectors u, v , respectively, given by

$$R_a^\mu = \frac{1}{N} \sum_i w_{i,u}^{(a)} w_i^0, \tag{A5}$$

$$q_{ab}^{uv} = \frac{1}{N} \sum_i w_{i,u}^{(a)} w_{i,v}^{(b)}. \tag{A6}$$

It is easy to visualize the overlap matrix (A6) in a block-matrix form. We can write a $n \times n$ block matrix of the type

$$\mathcal{Q} = \begin{bmatrix} \mathcal{Q}_{11} & \cdots & \mathcal{Q}_{1n} \\ \vdots & \ddots & \vdots \\ \mathcal{Q}_{n1} & \cdots & \mathcal{Q}_{nn} \end{bmatrix}, \tag{A7}$$

where each inner matrix \mathcal{Q}_{ab} has dimension $y \times y$. Each of these matrices describes the typical two-point overlap between students with two generic replica indices a and b . Exploiting trivial symmetries under indices permutations, the number of independent overlaps is equal to $n\binom{y}{2} + \binom{n}{2}y^2 = \binom{ny}{2}$. Substituting (A4) and enforcing definitions (A5) and (A6) through delta functions, we can rewrite the averaged partition function as

$$\begin{aligned} \langle Z^n \rangle_{\{\xi^\mu\}_{\mu=1}^M} &= \sum_{\mathbf{w}_u^{(a)}} \int \prod_{a,u} dR_a^u d\hat{R}_a^u \int \prod_{a,b,u,v \in \mathcal{P}} dq_{ab} d\hat{q}_{ab} \int \prod_{\mu} d\omega_0^\mu d\hat{\omega}_0^\mu \int \prod_{\mu,a,u} d\Delta_{u,a}^\mu d\hat{\Delta}_{u,a}^\mu \\ &\times \exp \left[-\beta \sum_{a,u,\mu} V(\Delta_{u,a}^\mu) + i \sum_{\mu} \omega_0^\mu \hat{\omega}_0^\mu + i \sum_{a,\mu,u} \Delta_{u,a}^\mu \hat{\Delta}_{u,a}^\mu + iN \sum_{a,u} R_a^u \hat{R}_a^u + iN \sum_{a,b,u,v \in \mathcal{P}} q_{ab}^{uv} \hat{q}_{ab}^{uv} - \frac{1}{2} \sum_{\mu} (\hat{\omega}_0^\mu)^2 \right. \\ &\quad - \sum_{\mu} \hat{\omega}_0^\mu \sigma_0^\mu \sum_{a,u} \hat{\Delta}_{u,a}^\mu R_a^u - \frac{1}{2} \sum_{\mu,a,u} (\hat{\Delta}_{u,a}^\mu)^2 - \sum_{\mu,a} \sum_{u < v} q_{aa}^{uv} \hat{\Delta}_{u,a}^\mu \hat{\Delta}_{v,a}^\mu - \sum_{\mu} \sum_{a < b} \sum_{u,v} q_{ab}^{uv} \hat{\Delta}_{u,a}^\mu \hat{\Delta}_{v,b}^\mu \\ &\quad \left. + \frac{\gamma}{y} \sum_a \sum_{u < v} \mathbf{w}_u^{(a)} \cdot \mathbf{w}_v^{(a)} - i \sum_{a,u} \hat{R}_a^u \sum_i w_{i,u}^{(a)} w_i^0 - i \sum_{a,b,u,v \in \mathcal{P}} \hat{q}_{ab}^{uv} \sum_i w_{i,u}^{(a)} w_{i,v}^{(b)} \right], \end{aligned} \quad (\text{A8})$$

where the symbol \mathcal{P} is a short-hand notation to indicate all the possible independent overlaps. In particular, $\sum_{a,b,u,v \in \mathcal{P}} = \sum_a \sum_{u < v} + \sum_{a < b} \sum_{u,v}$. We can notice now how the integrals over μ -dependent quantities can be factorized, as well as the sum over weight components i . We can therefore rewrite Eq. (A8) in a saddle point form. Using the property that the teacher vector components are i.i.d, we can write

$$\langle Z^n \rangle_{\{\xi^\mu\}_{\mu=1}^M, \mathbf{w}_0} = \int \prod_{a,u} dR_a^u d\hat{R}_a^u \int \prod_{a,b,u,v \in \mathcal{P}} dq_{ab} d\hat{q}_{ab} \exp [Ny \mathcal{G}(\{R_a^u, q_{ab}^{uv}, \hat{R}_a^u, \hat{q}_{ab}^{uv}\})], \quad (\text{A9})$$

$$\mathcal{G}(\{R_a^u, q_{ab}^{uv}, \hat{R}_a^u, \hat{q}_{ab}^{uv}\}) = \frac{i}{y} \sum_{a,u} R_a^u \hat{R}_a^u + \frac{i}{y} \sum_{a,b,u,v \in \mathcal{P}} q_{ab}^{uv} \hat{q}_{ab}^{uv} + \frac{\alpha}{y} G_E(\{R_a^u, q_{ab}^{uv}\}) + \frac{1}{y} G_I(\{\hat{R}_a^u, \hat{q}_{ab}^{uv}\}), \quad (\text{A10})$$

where \mathcal{G} plays the role of a free entropy with opposite sign. The equilibrium behavior is thus determined by the maximum of Eq. (A10) w.r.t. the order parameters. The quantities G_E and G_I represent the usual entropic and energetic terms, respectively, given by

$$G_I = \mathbb{E}_{w_0^a} \log \sum_{w_u^a \in \{-1,1\}} \exp \left[\frac{\gamma}{y} \sum_a \sum_{u < v} w_u^a w_v^a - i w_0^a \sum_{a,u} \hat{R}_a^u w_u^a - i \sum_{a,b,u,v \in \mathcal{P}} \hat{q}_{ab}^{uv} w_u^a w_v^b \right], \quad (\text{A11})$$

$$\begin{aligned} G_E &= \log \int d\omega_0 d\hat{\omega}_0 d\Delta_u^a d\hat{\Delta}_u^a \exp \left[-\beta \sum_{a,u} V(\Delta_u^a) + i\omega_0 \hat{\omega}_0 + i \sum_{a,u} \Delta_u^a \hat{\Delta}_u^a - \frac{1}{2} \hat{\omega}_0^2 - \frac{1}{2} \sum_{a,u} (\hat{\Delta}_{u,a}^\mu)^2 \right. \\ &\quad \left. - \hat{\omega}_0 \sigma_0 \sum_{a,u} \hat{\Delta}_u^a R_a^u - \frac{1}{2} \sum_{a,u \neq v} \hat{\Delta}_u^a \hat{\Delta}_v^a q_{aa}^{uv} - \frac{1}{2} \sum_{a \neq b} \sum_{u,v} \hat{\Delta}_u^a \hat{\Delta}_v^b q_{ab}^{uv} \right], \end{aligned} \quad (\text{A12})$$

where in the last line we exploited the fact that the teacher vector components are i.i.d. by assumption. Notice that the expectation over one representative component w_0 becomes dummy by means of a gauge transformation $w_u^a \rightarrow w_u^0 w_u^a$ for all the weight components, so we will drop it from now on. Before going on, note that the integral over $\hat{\omega}_0$ in Eq. (A12) can be carried out explicitly, leading to

$$\begin{aligned} G_E &= \log \int \mathcal{D}\omega_0 d\Delta_u^a d\hat{\Delta}_u^a \exp \left\{ -\beta \sum_{a,u} V(\Delta_u^a) + i \sum_{a,u} \hat{\Delta}_u^a (\Delta_u^a - \omega_0 \sigma_0 R_a^u) - \frac{1}{2} \sum_{a,u} [1 - (R_a^u)^2] (\hat{\Delta}_{u,a}^\mu)^2 \right. \\ &\quad \left. - \frac{1}{2} \sum_{a,u \neq v} \hat{\Delta}_u^a \hat{\Delta}_v^a (q_{aa}^{uv} - R_a^u R_a^v) - \frac{1}{2} \sum_{a \neq b} \sum_{u,v} \hat{\Delta}_u^a \hat{\Delta}_v^b (q_{ab}^{uv} - R_a^u R_b^v) \right\}, \end{aligned} \quad (\text{A13})$$

where $\mathcal{D}x = e^{-x^2/2} dx / \sqrt{2\pi}$ denotes the standard Gaussian probability measure.

1. Replica symmetric ansatz on both spaces

The simplest possible ansatz corresponds to assume a permutation symmetry over *both* replica spaces. For what concerns the “fake”-replicated space, this is the simplest choice in any spin-glass model [33]. On the other hand, a symmetry ansatz between the pairwise correlations between students (i.e., the overlaps q_{aa}^{uv}) comes naturally from the fully connected topology of interactions between the students’ weight vectors as assumed in Eq. (2). Choosing different topologies would imply to parametrize the matrices \mathcal{Q}_{aa} in such a way to reflect the behavior of nonconnected correlations in that specific graph, which is itself a nontrivial problem unless in very specific architectures (e.g., trees or planar graphs). However, restricting to the fully connected topology as in the main text, within this extended RS assumption we have only three order parameters (and their conjugates) left, namely,

$$R_a^u = R \quad \forall u, a, \quad q_{aa}^{uv} = p \quad \forall u \neq v, a, \quad q_{ab}^{uv} = q \quad \forall u, v, a \neq b, \quad (\text{A14a})$$

$$\hat{R}_a^u = i\hat{R} \quad \forall u, a, \quad \hat{q}_{aa}^{uv} = i\hat{p} \quad \forall u \neq v, a, \quad \hat{q}_{ab}^{uv} = i\hat{q} \quad \forall u, v, a \neq b. \quad (\text{A14b})$$

By inserting the above ansatz into Eqs. (A10), (A13), and (A11) and taking the $n \rightarrow 0$ limit, after some calculations we can rewrite the RS quenched free entropy as

$$\mathcal{G} = -R\hat{R} - \frac{(y-1)}{2}p\hat{p} + \frac{\hat{q}}{2}(yq-1) + \frac{2\alpha}{y} \int \mathcal{D}t \Phi(at) \log \int \mathcal{D}\tau \Xi^y + \frac{1}{y} \int \mathcal{D}z \log \mathcal{Z}_{(y)}, \quad (\text{A15})$$

where

$$\Xi = \int \frac{d\Delta}{\sqrt{1-p}} \exp \left[-\beta V(\Delta) - \frac{1}{2(1-p)} (\Delta + \sqrt{qt} + \sqrt{p-q}\tau)^2 \right], \quad (\text{A16})$$

$$\mathcal{Z}_{(y)} = \sum_{w_u \in \pm 1} \exp \left[\left(\frac{y}{y} + \hat{p} - \hat{q} \right) \sum_{u < v} w_u w_v + (w^0 \hat{R} + \sqrt{\hat{q}}z) \sum_u w_u \right], \quad (\text{A17})$$

with $\Phi(x) = \int_x^\infty \mathcal{D}z = \frac{1}{2} \text{erfc}(x/\sqrt{2})$ and for notational brevity we defined the following three quantities:

$$a = \frac{R}{\sqrt{q-R^2}}, \quad b = \sqrt{\frac{q}{1-p}}, \quad c = \sqrt{\frac{p-q}{1-p}}. \quad (\text{A18})$$

Concerning the energetic term in (A15), explicit formulas for Ξ (A16) depend on the specific choice for the potential. For a potential of the type $V(\Delta) = (-\Delta)^v \Theta(-\Delta)$, with $v = 1$, it reads

$$\begin{aligned} V(\Delta) &= -\Delta \Theta(-\Delta) \longrightarrow \\ \Xi &= \Phi(bt + c\tau) + e^{\frac{\beta^2(1-p)}{2} - \beta\sqrt{1-p}(bt+c\tau)} \\ &\quad \times \Phi(\beta\sqrt{1-p} - bt - c\tau). \end{aligned} \quad (\text{A19})$$

Interestingly, the entropic term can be seen as an averaged free entropy (apart on a sign) of a reduced system of y degrees of freedom (in this case, binary spins due to the binary nature of the starting weights): specifically, at fixed z Eq. (A17) defines the partition function of a Curie-Weiss model of y spins with an effective ferromagnetic coupling and a global external field given by the sum of two terms: a signal one proportional to \hat{R} , and a *global* Gaussian field z to be further averaged over.

Saddle point equations

The equilibrium behavior of the system at fixed control parameters is given by the maximum of \mathcal{G} w.r.t. all the order parameters, whose values are determined by imposing stationarity of the free entropy. After some calculations, we can write the self-consistent equations for the six-order parameters

($R, q, \hat{R}, \hat{q}, \hat{p}$) as

$$\begin{aligned} \hat{R} &= \sqrt{\frac{2}{\pi}} \alpha \beta \sqrt{\frac{q}{q-R^2}} \int \mathcal{D}t e^{-\frac{a^2 t^2}{2}} \\ &\quad \times \frac{\int \mathcal{D}\tau \mathcal{T} \Phi(\beta\sqrt{1-p} - bt - c\tau) \Xi^{y-1}}{\int \mathcal{D}\tau \Xi^y}, \end{aligned} \quad (\text{A20a})$$

$$\begin{aligned} \hat{q} &= 2\alpha\beta^2 \int \mathcal{D}t \Phi(at) \\ &\quad \times \frac{[\int \mathcal{D}\tau \Xi^{y-1} \mathcal{T} \Phi(\beta\sqrt{1-p} - bt - c\tau)]^2}{[\int \mathcal{D}\tau \Xi^y]^2}, \end{aligned} \quad (\text{A20b})$$

$$\begin{aligned} \hat{p} &= 2\alpha\beta^2 \int \mathcal{D}t \frac{\Phi(at)}{\int \mathcal{D}\tau \Xi^y} \int \mathcal{D}\tau \mathcal{T}^2 \Xi^{y-2} \Phi^2 \\ &\quad \times (\beta\sqrt{1-p} - bt - c\tau), \end{aligned} \quad (\text{A20c})$$

for the conjugate order parameters—where we defined for convenience $\mathcal{T} = e^{\frac{\beta^2(1-p)}{2} - \beta\sqrt{1-p}(bt+c\tau)}$ —and

$$R = \int \mathcal{D}z \Omega(w_1), \quad (\text{A21a})$$

$$q = \int \mathcal{D}z \Omega(w_1) \Omega(w_2) = \int \mathcal{D}z \Omega^2(w_1), \quad (\text{A21b})$$

$$p = \int \mathcal{D}z \Omega(w_1 w_2), \quad (\text{A21c})$$




FIG. 4. Phase diagram of $y = 2$ coupled binary perceptrons in the $(T - \alpha)$ plane at different values of γ (values given in the legend). Comparison with the single perceptron in the RS approach (black lines).

with $\Omega(O)$ denoting the expectation of a generic observable O for the partition function (A17), namely,

$$\Omega(O) = \frac{1}{\mathcal{Z}_{(y)}(\hat{R}, \hat{p}, \hat{q}, \gamma)} \sum_{w_u \in \pm 1} O \exp \left[\left(\frac{\gamma}{y} + \hat{p} - \hat{q} \right) \times \sum_{u < v} w_u w_v + (\hat{R} + \sqrt{\hat{q}}z) \sum_u w_u \right]. \quad (\text{A22})$$

In the limit $\gamma \rightarrow 0$ the different students are uncoupled and the free entropy becomes identical to that of the nonreplicated model: in other terms, the only possible solution of the saddle point equations at $\gamma = 0$ is $p = q$ and $\hat{p} = \hat{q}$.

2. Effect of γ

In Fig. 4 we show the effect of changing γ at a fixed number of coupled students y . The spinodal lines continuously approach the result for the single system $y = 1$ when $\gamma \rightarrow 0$, confirming the validity of the theory. This holds independently on the number y of students. As long as γ is increased, the phase diagram seems also to shift at higher temperatures: this would imply that the larger γ , the higher the starting temperature for RSA should be, in such a way to start the annealing in a regime where the interaction between replicas has a lower effect w.r.t. the cost function of each student.

APPENDIX B: MF THEORY OF SINGLE PERCEPTRON

For the sake of completeness, in this Appendix we report the expressions of the free entropy—and the corresponding self-consistent equations for the order parameters—for the single perceptron (i.e., at $y = 1$), obtained both through a replica-symmetric (RS) and one-step replica-symmetry-breaking (1-RSB) ansatzs. We omit the derivation of all the following expressions as they have been extensively computed in several works (see e.g., [2,23,29] and references therein). In both cases, the two control parameters are simply the fraction α of examples provided to the student and the temperature T .

The following expressions are reported considering the same learning rule as in the main text, i.e., $V(\Delta) = -\Delta\Theta(-\Delta)$, where Δ is the stability parameter.

1. RS free entropy

In this case, there are two order parameters R, q (and their conjugates \hat{R}, \hat{q}) where R is the typical overlap between student and teacher's weights, and q is the (unique by assumption) two-replica overlap. The free entropy as a function of these order parameters (R, q, \hat{R}, \hat{q}) and the corresponding self-consistent equations read

$$\mathcal{G} = -R\hat{R} + \frac{1}{2}(q-1)\hat{q} + 2\alpha \int \mathcal{D}t \Phi(at) \log \Xi + \int \mathcal{D}z \log 2 \cosh(\hat{R} + \sqrt{\hat{q}}z), \quad (\text{B1})$$

$$\Xi = \Phi(bt) + e^{\frac{\beta^2(1-q)}{2} - \beta\sqrt{\hat{q}}t} \Phi(\beta\sqrt{1-q} - bt), \quad (\text{B2})$$

and

$$R = \int \mathcal{D}z \tanh(\hat{R} + \sqrt{\hat{q}}z), \quad \hat{R} = \frac{2\alpha\beta}{\sqrt{2\pi}} \int \mathcal{D}z \left[1 + \frac{\Phi(vz)}{\Phi(\beta\sqrt{1-q} - vz)} e^{\beta\sqrt{q-R^2}z - \frac{\beta^2(1-q)}{2}} \right]^{-1}, \quad (\text{B3a})$$

$$q = \int \mathcal{D}z \tanh^2(\hat{R} + \sqrt{\hat{q}}z) \quad \hat{q} = 2\alpha\beta^2 \int \mathcal{D}t \Phi(at) \left[1 + \frac{\Phi(bt)}{\Phi(\beta\sqrt{1-q} - bt)} e^{\beta\sqrt{\hat{q}}t - \frac{\beta^2(1-q)}{2}} \right]^{-2}, \quad (\text{B3b})$$

where

$$a = \frac{R}{\sqrt{q-R^2}}, \quad b = \sqrt{\frac{q}{1-q}}, \quad v = \sqrt{\frac{q-R^2}{1-q}}. \quad (\text{B4})$$

2. 1-RSB free entropy

In this case there are three order parameters R, q_1, q_0 (and their conjugates $\hat{R}, \hat{q}_1, \hat{q}_0$) with $q_0 \geq q_1$ (resp. $\hat{q}_0 \geq \hat{q}_1$): the two values of the overlap arise from the usual 1-RSB structure of the overlap matrix $\mathbf{Q} = \{q_{ab}\}$ [33]. There is an additional order parameter θ that tunes the relative weight of the two overlaps, so that the overlap distribution follows $P(\langle q_{ab} \rangle) = \theta \delta(\langle q_{ab} \rangle - q_0) + (1-\theta) \delta(\langle q_{ab} \rangle - q_1)$. The free entropy as a function of these order parameters ($R, q_1, q_0, \hat{R}, \hat{q}_1, \hat{q}_0$) and the corresponding self-consistent equations read

$$\mathcal{G} = -R\hat{R} + \frac{1}{2}\hat{q}_0(q_0-1) - \frac{\theta}{2}(q_0\hat{q}_0 - q_1\hat{q}_1) + \frac{2\alpha}{\theta} \int \mathcal{D}t \Phi(at) \log \int \mathcal{D}\tau \Xi^\theta + \frac{1}{\theta} \int \mathcal{D}\tau \log \int \mathcal{D}z 2^\theta \cosh^\theta \chi, \quad (\text{B5})$$

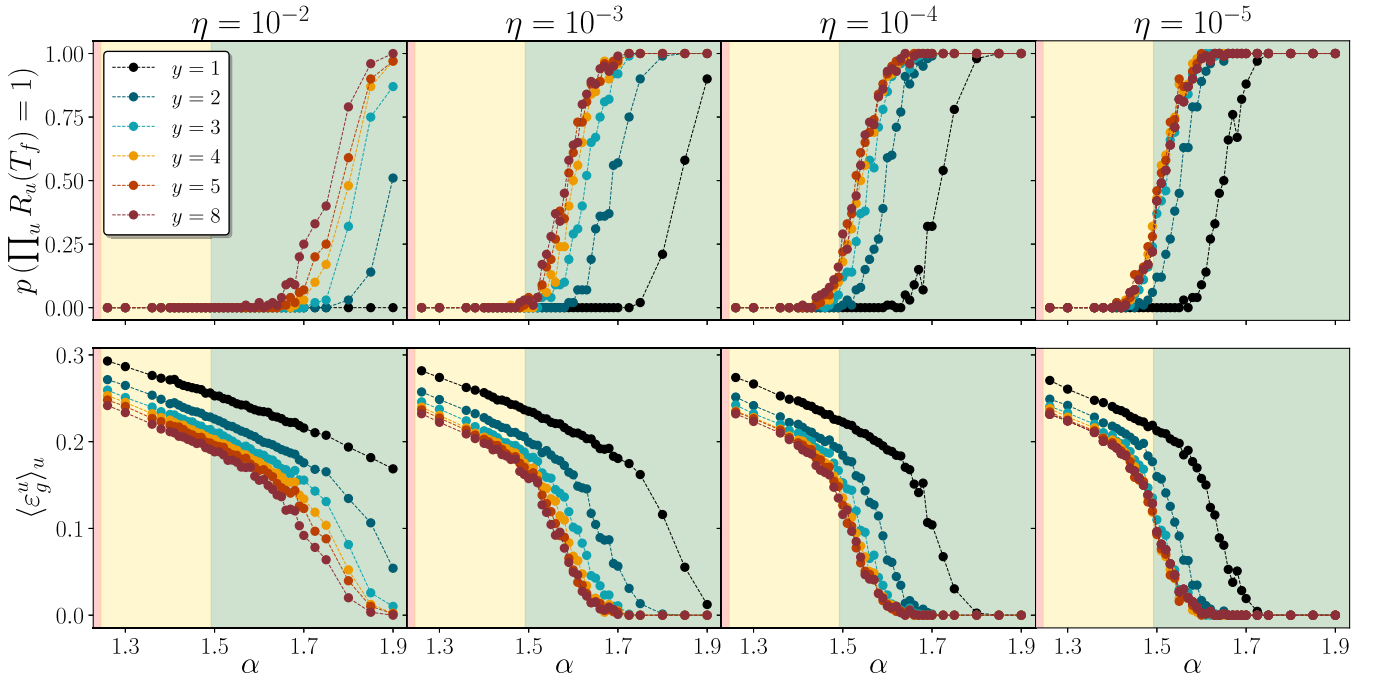


FIG. 5. Numerical performances of SA/RSA: the top row shows the empirical probability that the teacher configuration is found at the end of the SA, the bottom row shows the mean generalization error; both quantities are plotted vs the fraction of samples α , for a system with $N = 2001$ weights, for different values of the annealing rate η (shown at the top of each column). Comparison between the single nonreplicated perceptron (black lines) and a system of y coupled perceptrons with different values of y , with fixed coupling $\gamma = 1$. For $y \geq 2$, the top row shows the probability that *all* the y students find the teacher configuration at the end of the SA, and the bottom row displays the mean generalization error averaged also on the y students. However, we numerically observe that, except at extremely fast annealing rates (e.g., the first panel with $\eta = 10^{-2}$), the students display a practically identical behavior in temperature. The rightmost panels are the same shown in Fig. 2.

$$\Xi = \Phi(bt + c\tau) + e^{\frac{\beta^2(1-p)}{2} - \beta\sqrt{1-q_0}(bt+c\tau)} \times \Phi(\beta\sqrt{1-q_0} - bt - c\tau), \quad (\text{B6})$$

$$\chi = \hat{R} + \sqrt{\hat{q}_0 - \hat{q}_1}z + \sqrt{\hat{q}_1}\tau, \quad (\text{B7})$$

$$\hat{q}_0 = 2\alpha\beta^2 \int \mathcal{D}t \frac{\Phi(at)}{\int \mathcal{D}\tau \Xi^\theta} \int \mathcal{D}\tau \mathcal{T}^2 \Xi^{\theta-2} \Phi^2 \times (\beta\sqrt{1-q_0} - bt - c\tau), \quad (\text{B8c})$$

where $\mathcal{T} = e^{\frac{\beta^2(1-q_0)}{2} - \beta\sqrt{1-q_0}(bt+c\tau)}$ and

$$a = \frac{R}{\sqrt{q_1 - R^2}}, \quad b = \sqrt{\frac{q_1}{1 - q_0}}, \quad c = \sqrt{\frac{q_0 - q_1}{1 - q_0}}. \quad (\text{B9})$$

and

$$R = \int \mathcal{D}\tau \frac{\int \mathcal{D}z \tanh \chi \cosh^\theta \chi}{\int \mathcal{D}z \cosh^\theta \chi}, \quad \hat{R} = \sqrt{\frac{2}{\pi}} \alpha \beta \sqrt{\frac{q_0}{q_1 - R^2}} \int \mathcal{D}t \times e^{-\frac{a^2 \tau^2}{2}} \frac{\int \mathcal{D}\tau \mathcal{T} \Phi(\beta\sqrt{1-q_0} - bt - c\tau) \Xi^{\theta-1}}{\int \mathcal{D}\tau \Xi^\theta}, \quad (\text{B8a})$$

$$q_1 = \int \mathcal{D}\tau \frac{(\int \mathcal{D}z \tanh \chi \cosh^\theta \chi)^2}{(\int \mathcal{D}z \cosh^\theta \chi)^2} \quad \hat{q}_1 = 2\alpha\beta^2 \int \mathcal{D}t \Phi(at) \times \frac{[\int \mathcal{D}\tau \Xi^{\theta-1} \mathcal{T} \Phi(\beta\sqrt{1-q_0} - bt - c\tau)]^2}{[\int \mathcal{D}\tau \Xi^\theta]^2}, \quad (\text{B8b})$$

$$q_0 = \int \mathcal{D}\tau \frac{\int \mathcal{D}z \tanh^2 \chi \cosh^\theta \chi}{\int \mathcal{D}z \cosh^\theta \chi}$$

Note that the structure of both the free entropy and the self-consistent equations (in particular the block related to the energetic term) are almost equivalent to the ones shown in Appendix A for the coupled system with y students [see in particular Eqs. (A15) and (A20)], provided the mapping $p \leftrightarrow q_0$, $q \leftrightarrow q_1$ (and the same for their conjugates) and $y \leftrightarrow \theta$, although the interpretation of the two parameters $y \in \mathbb{Z}_+$ and $\theta \in [0, 1]$ in the two models is completely different and an additional ferromagnetic interaction γ is present in the first case. Given this mapping, also the left equations in (B8) are equivalent to (A21a), (A21b), and (A21c) upon a Hubbard-Stratonovich transform.

Dynamic transition line

The dynamic 1-RSB transition temperature $T_{\text{dyn}}(\alpha)$ can be found as the largest temperature at which the 1-RSB free entropy admits a local maximum at $\theta \rightarrow 1$ with $q_0 > q_1$. We recall the procedure discussed in [35] that we used also in

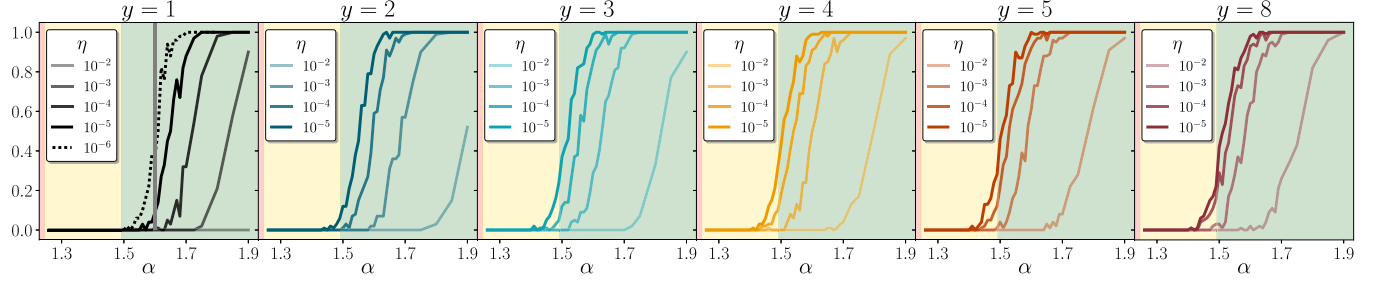


FIG. 6. Same results as in Fig. 5, but this time each panel refers to one value of y : each panel shows the probability of finding the teacher solution at the end of the SA for different values of η : we can observe how the cooling rate seems to converge to the same curves in the case $y = 5$ and 8 . The leftmost panel ($y = 1$) shows also the result obtained with an annealing rate $\eta = 10^{-6}$ over 50 seeds (dotted line). The vertical gray line in the leftmost panel shows the critical value of α at which 1-RSB glassy states disappear, $\alpha_c^{\text{RSB}} \approx 1.6$. The settings are identical to those of Fig. 5.

this case. We start from the 1-RSB free entropy (B5), and we expand it around $\theta = 1$ to the first order:

$$\begin{aligned} \mathcal{G}^{\text{1RSB}}(R, \hat{R}q_1, q_0, \hat{q}_1, \theta) \\ = \mathcal{G}^{\text{RS}}(R, q_1, \hat{R}, \hat{q}_1) \\ + (\theta - 1)\partial_\theta \mathcal{G}^{\text{1RSB}}|_{\theta=1} + O[(\theta - 1)^2]. \end{aligned} \quad (\text{B10})$$

In particular, the 1-RSB free entropy computed at $\theta = 1$ gives the RS expression by construction with $q_0 = q_1$. On the other hand, the first-order contribution $\delta\mathcal{G} = \partial_\theta \mathcal{G}^{\text{1RSB}}|_{\theta=1}$ is given by

$$\begin{aligned} \delta\mathcal{G} = \frac{\partial \mathcal{G}^{\text{1RSB}}}{\partial \theta}|_{\theta=1} = -\frac{1}{2}(q_0\hat{q}_0 - q_1\hat{q}_1) \\ - \mathbb{E}_{w^0} \int D\tau \log \int Dz \cosh \chi \\ + \mathbb{E}_{w^0} \int D\tau \frac{\int Dz \cosh \chi \log \cosh \chi}{\int Dz \cosh \chi} \\ - 2\alpha \int Dt \Phi(at) \log \int Dz \Xi \\ + 2\alpha \int Dt \Phi(at) \frac{\int Dz \Xi \log \Xi}{\int Dz \Xi}. \end{aligned} \quad (\text{B11})$$

The procedure is to find the global maximum of the 0th-order term (i.e., the RS part); then, at fixed (R, q_1) we compute the equilibrium value of the first-order correction (B11) through another saddle point w.r.t. q_0 . The latter order parameter and its conjugate \hat{q}_0 are fixed by the following saddle point equations:

$$q_0 = \int D\tau \frac{\int Dz \tanh \chi \sinh \chi}{\int Dz \cosh \chi}, \quad (\text{B12})$$

$$\begin{aligned} \hat{q}_0 = 4\alpha \int Dt \Phi(at) \left[\frac{\int Dz (\partial_{q_0} \Xi) \log \Xi}{\int Dz \Xi} \right. \\ \left. - \left(\frac{\int Dz \Xi \log \Xi}{\int Dz \Xi} \right) \left(\frac{\int Dz \partial_{q_0} \Xi}{\int Dz \Xi} \right) \right], \end{aligned} \quad (\text{B13})$$

where χ, Ξ are given by Eqs. (B7) and (B6), respectively.

APPENDIX C: SIMULATED ANNEALING IMPLEMENTATION DETAILS AND ADDITIONAL NUMERICAL RESULTS

In this Appendix we discuss the implementation parameters of the simulated annealing (SA). The model (2) is initialized at a temperature $T_0 = 0.4$, and each student's weight vector \mathbf{w}_u is drawn at random from $\{-1, 1\}^N$, independently on the others. At each temperature we perform one Monte Carlo sampling sweep, where a move is proposed for every degree of freedom $w_{i,u}$, and it is accepted/rejected according to the Metropolis choice. Then the temperature is linearly decreased so that $T_l = T_{l-1} - \eta$ with η a suitable annealing rate, and the sampling is repeated starting from the last configuration at the previous (higher) temperature. The annealing process continues up to a final temperature $T_f = 10^{-2}$. In all the simulations shown performed in this work, we used $\eta \in 10^{\{-2, -3, -4, -5, -6\}}$ (for computational time reasons, the latter value has been used only for $y = 1$; see Fig. 6). The energy shifts can be efficiently computed using the same procedure discussed in [17] (Supporting Information).




FIG. 7. AMP's empirical probability of finding the teacher configuration vs the fraction of examples α . Each curve corresponds to a different system size N (shown in the caption) and shows the amount of instances for which AMP converges to the teacher, i.e., to a set of marginals \mathbf{m} such that $1 - \sum_i m_i w_i^0 / N < \epsilon$ where \mathbf{w}^0 is the teacher configuration and ϵ is the convergence threshold on the AMP marginal probabilities used for these experiments (here $\epsilon = 10^{-7}$). The simulations are done at the Bayes optimal temperature $T = 0$.

APPENDIX D: APPROXIMATE MESSAGE PASSING

In this section we discuss some numerical experiments on the Approximate Message Passing (AMP) algorithm that we used to check the spinodal point at $T = 0$, i.e., the hard-easy boundary as found by the RS theory. We do not discuss the details of the method itself, as has been extensively analyzed in several previous works: in particular, the AMP implementation we used here is taken from the general derivation for generalized linear estimation problems (GLMs, the binary perceptron in the teacher-student scenario being an example) presented in the review article [3] (Sec. VI C). Similarly to the SA experiments, we run the AMP algorithm by fixing the

system size N and varying the amount of examples $\alpha = M/N$. For each instance, AMP is run until convergence is reached on the marginal probabilities for each weight (with a finite tolerance $\varepsilon = 10^{-7}$). We then compute the empirical probability (averaging over all the instances) that the final overlap between the teacher configuration and the AMP's marginal is 1. Results are shown in Fig. 7 for different system sizes N . We can clearly see how the curves get sharper and sharper by increasing N , and the transition to a perfect recovery approaches the boundary between the hard and the easy phases as computed from the RS phase diagram at $T = 0$, so that $\alpha_{\text{AMP}} = \alpha_c$.

-
- [1] E. Gardner and B. Derrida, *J. Phys. A: Math. Gen.* **21**, 271 (1988).
 - [2] H. S. Seung, H. Sompolinsky, and N. Tishby, *Phys. Rev. A* **45**, 6056 (1992).
 - [3] L. Zdeborová and F. Krzakala, *Adv. Phys.* **65**, 453 (2016).
 - [4] S. Sarao Mannelli, G. Biroli, C. Cammarota, F. Krzakala, P. Urbani, and L. Zdeborová, *Phys. Rev. X* **10**, 011057 (2020).
 - [5] M. Mézard, G. Parisi, and R. Zecchina, *Science* **297**, 812 (2002).
 - [6] L. Zdeborová and F. Krzakala, *Phys. Rev. E* **76**, 031131 (2007).
 - [7] F. Krzakala and L. Zdeborová, *Phys. Rev. Lett.* **102**, 238701 (2009).
 - [8] W. Krauth, M. Mézard, and J.-P. Nadal, *Complex Syst.* **2**, 387 (1988).
 - [9] R. Monasson and R. Zecchina, *Phys. Rev. Lett.* **75**, 2432 (1995).
 - [10] R. Monasson and R. Zecchina, *Mod. Phys. Lett. B* **09**, 1887 (1995).
 - [11] R. Dietrich, M. Oppen, and H. Sompolinsky, *Phys. Rev. Lett.* **82**, 2975 (1999).
 - [12] E. Barkai, D. Hansel, and H. Sompolinsky, *Phys. Rev. A* **45**, 4146 (1992).
 - [13] E. Cornacchia, F. Mignacco, R. Veiga, C. Gerbelot, B. Loureiro, and L. Zdeborová, *Mach. Learn.: Sci. Technol.* **4**, 015019 (2023).
 - [14] S. Franz and G. Parisi, *J. Phys. A: Math. Theor.* **49**, 145001 (2016).
 - [15] C. Baldassi, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina, *J. Stat. Mech.* (2016) 023301.
 - [16] C. Baldassi, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina, *Phys. Rev. Lett.* **115**, 128101 (2015).
 - [17] C. Baldassi, C. Borgs, J. T. Chayes, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina, *Proc. Natl. Acad. Sci. USA* **113**, E7655 (2016).
 - [18] H. Huang, K. Y. M. Wong, and Y. Kabashima, *J. Phys. A: Math. Theor.* **46**, 375002 (2013).
 - [19] C. Baldassi and R. Zecchina, *Proc. Natl. Acad. Sci. USA* **115**, 1457 (2018).
 - [20] M. C. Angelini and F. Ricci-Tersenghi, *Phys. Rev. X* **13**, 021011 (2023).
 - [21] W. Barthel, A. K. Hartmann, M. Leone, F. Ricci-Tersenghi, M. Weigt, and R. Zecchina, *Phys. Rev. Lett.* **88**, 188701 (2002).
 - [22] G. Györgyi, *Phys. Rev. A* **41**, 7097 (1990).
 - [23] T. L. H. Watkin, A. Rau, and M. Biehl, *Rev. Mod. Phys.* **65**, 499 (1993).
 - [24] H. Huang and Y. Kabashima, *Phys. Rev. E* **90**, 052813 (2014).
 - [25] R. Monasson, *Phys. Rev. Lett.* **75**, 2847 (1995).
 - [26] S. Franz and G. Parisi, *J. Phys. I* **5**, 1401 (1995).
 - [27] H. Huang and H. Zhou, *Europhys. Lett.* **96**, 58003 (2011).
 - [28] S. Zhang, A. E. Choromanska, and Y. LeCun, in *Advances in Neural Information Processing Systems*, edited by C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett (Curran Associates, Red Hook, New York, 2015), Vol. 28.
 - [29] A. Engel and C. Van den Broeck, *Statistical Mechanics of Learning* (Cambridge University Press, Cambridge, 2001).
 - [30] H. Horner, *Z. Phys. B* **87**, 371 (1992).
 - [31] H. Horner, *Z. Phys. B* **86**, 291 (1992).
 - [32] F. Rosenblatt, *Psychol. Rev.* **65**, 386 (1958).
 - [33] M. Mézard, G. Parisi, and M. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1986).
 - [34] P. Charbonneau, E. Marinari, M. Mézard, G. Parisi, F. Ricci-Tersenghi, G. Sicuro, and F. Zamponi, *Spin Glass Theory and Far Beyond* (World Scientific, Singapore, 2023).
 - [35] S. Franz, G. Parisi, M. Sevelev, P. Urbani, and F. Zamponi, *SciPost Phys.* **2**, 019 (2017).
 - [36] F. Antenucci, S. Franz, P. Urbani, and L. Zdeborová, *Phys. Rev. X* **9**, 011020 (2019).
 - [37] ic.kampal.com.
 - [38] Q. Yang, Y. Liu, T. Chen, and Y. Tong, *ACM Trans. Intell. Syst. Technol. (TIST)* **10**, 1 (2019).