

PAPER

Thermodynamics of bidirectional associative memories

To cite this article: Adriano Barra *et al* 2023 *J. Phys. A: Math. Theor.* **56** 205005

View the [article online](#) for updates and enhancements.

You may also like

- [Free energies of Boltzmann machines: self-averaging, annealed and replica symmetric approximations in the thermodynamic limit](#)
Elena Agliari, Adriano Barra and Brunello Tirozzi
- [The Hopfield model revisited: covariance and quantization](#)
F Belgiorno, S L Cacciatori and F Dalla Piazza
- [Replica symmetry breaking in neural networks: a few steps toward rigorous results](#)
Elena Agliari, Linda Albanese, Adriano Barra *et al.*

Thermodynamics of bidirectional associative memories

Adriano Barra^{1,2} , Giovanni Catania^{3,*} ,
Aurélien Decelle^{3,4}  and Beatriz Seoane^{3,4} 

¹ Dipartimento di Matematica e Fisica, Università del Salento, Campus Ecotekne, 73100 Lecce, Italy

² Istituto Nazionale di Fisica Nucleare, Sezione di Lecce, Campus Ecotekne, 73100 Lecce, Italy

³ Departamento de Física Teórica, Universidad Complutense de Madrid, 28040 Madrid, Spain

⁴ Université Paris-Saclay, CNRS, INRIA Tau team, LISN, 91190 Gif-sur-Yvette, France

E-mail: gcatania@ucm.es

Received 21 November 2022; revised 27 March 2023

Accepted for publication 12 April 2023

Published 2 May 2023



CrossMark

Abstract

In this paper we investigate the equilibrium properties of bidirectional associative memories (BAMs). Introduced by Kosko in 1988 as a generalization of the Hopfield model to a bipartite structure, the simplest architecture is defined by two layers of neurons, with synaptic connections only between units of different layers: even without internal connections within each layer, information storage and retrieval are still possible through the reverberation of neural activities passing from one layer to another. We characterize the computational capabilities of a stochastic extension of this model in the thermodynamic limit, by applying rigorous techniques from statistical physics. A detailed picture of the phase diagram at the replica symmetric level is provided, both at finite temperature and in the noiseless regimes. Also for the latter, the critical load is further investigated up to one step of replica symmetry breaking. An analytical and numerical inspection of the transition curves (namely critical lines splitting the various modes of operation of the machine) is carried out as the control parameters—noise, load and asymmetry between the two layer sizes—are tuned. In particular, with a finite asymmetry between the two layers, it is shown how the BAM can store information more efficiently than the Hopfield model by requiring less parameters to encode a fixed number of patterns. Comparisons are made with numerical simulations of neural dynamics. Finally, a low-load analysis is carried out to explain the retrieval mechanism in the BAM

* Author to whom any correspondence should be addressed.

by analogy with two interacting Hopfield models. A potential equivalence with two coupled Restricted Boltzmann Machines is also discussed.

Keywords: neural networks, associative memory, spin-glasses, restricted Boltzmann machines, Hopfield model

(Some figures may appear in colour only in the online journal)

1. Introduction

Since the original work of Amit, Gutfreund, and Sompolinsky (AGS) [1] who first developed a mean-field theory for the Hopfield model of neural networks in terms of spin glasses, several extensions and generalizations have been made. As for dilution, although synaptic interactions in the Hopfield model are fully connected, it is possible to study artificial networks on sparse topologies, for example, by considering synaptic matrices arranged in random graphs, hierarchical structures, small-world, or scale-free architectures [2–5]. Another possibility is to consider sparsity in the pattern components, leading to equilibrium regimes where the retrieval of composite information is possible [6]. In addition, extensions to different support for the spin variables and/or the patterns (e.g. real-valued spins and/or pattern components) have been deeply analyzed [3, 7–9].

One of the most studied generalizations of the Hopfield model is a bipartite structure, where the network is divided into two layers of neurons and connections are allowed only between neurons of different layers. The most popular application of this type of topology is the so-called restricted Boltzmann machine (RBM) [10–12], which is widely used in computer science [10, 13, 14] (e.g. as fundamental units of modern deep architectures) and can be considered a prototype for machine learning models [14]. The Hopfield model was proposed as a toy model for neurophysiology that accounts for biological learning à la Hebb, and yet strong similarities between the information processing mechanisms of these two types of models have been demonstrated by statistical mechanical studies [11, 12, 15, 16], whose *modus operandi* will therefore be the methodological leitmotif of the present research.

Two main classes of reward emerge from the duality between neural architectures used in machine learning (e.g. RBMs) and biologically inspired neural networks (e.g. Hebbian models). The first is explainability, since by mapping the information processing capabilities of RBMs to those that naturally emerge in the Hopfield network, we can better understand the hidden decision-making mechanisms the machine follows when it performs, for example, denoising or pattern recognition: this can be particularly welcome in the field of Explainable Artificial Intelligence [17]. Moreover, statistical mechanics eventually provides phase diagrams, i.e. it reveals regions in the space of control parameters (e.g. noise, load, etc) where the machine has certain information storage capabilities, and regions where these properties are lost. This, in turn, is particularly welcome in the field of sustainable Artificial Intelligence, since knowledge of the phase diagram allows the machine to be brought into the optimal operational setting, which can lead to significant savings in training costs. In particular, the equivalence between the Hopfield model and certain classes of Boltzmann machines [11, 12, 15] has shown a strong connection between the retrieval properties of biological neural networks and the performance of machine learning algorithms in terms of their ability to correctly describe the empirical distribution of a dataset and/or eventually overfit it [15]; furthermore, the retrieval properties of the Hopfield model can be used to develop new training procedures for RBMs [18]. This means that the properties of biological and artificial information processing networks can be analyzed and exploited using similar techniques.

The bidirectional associative memory (BAM) we study in this paper was introduced by Kosko in [19] as an attempt to overcome the lack of internal organization of information in the original Hopfield model and to account for structured retrieval of patterns. It is another generalization of a neural network based on a bipartite topology: in this machine, as in RBM settings, there are only interactions between units of different layers. However, while RBMs have two layers covering rather different computational roles (the input layer is typically provided with the datasets, e.g. noisy patterns to be recognized and/or classified, and the hidden layer derives correlations in the visible variables from the supplied information, enabling pattern recognition and/or noise reduction) and the weight matrix has no particular biological significance, in the BAM architecture both layers are expected to provide pattern information (one pattern per layer, as an attempt to retrieve pairs of patterns rather than individual ones) and their synaptic matrix strongly resembles the original Hebb proposal. However, the operating principle of the BAM differs from that of the Hopfield model because in the BAM retrieval, information is passed from one layer to another according to appropriate dynamic rules: this mechanism is referred to in the literature as *reverberation* of information [19]. The BAM was analyzed using statistical physics techniques in previous works [20, 21], but only in the noiseless case (i.e. at $T = 0$) since the goal was to determine the storage properties in the saturation regime. Further analysis has been carried out in [22] using the replica trick at the replica symmetric (RS) level. In this paper, we perform an extensive analysis of the equilibrium properties of the BAM: we characterize the phase diagram analytically and numerically, and provide an extension of the Hopfield AGS mean field theory to the BAM. Analytical calculations are obtained independently using the Guerra interpolation scheme [23] and the replica method [24], both at the RS and with 1 step of replica symmetry (RS) breaking. Finally, the BAM has recently become topical again in the context of machine learning [18]. In this context, a detailed study of the phase diagram of the BAM in the presence of an external noise (i.e. at $T \neq 0$) and in terms of the asymmetry between the two layers is necessary to clarify under which circumstances retrieval of information is possible or not, a necessary first step to analyze the pattern extraction process with RBMs.

The paper is organized as follows: section 2 introduces the model and the notation used in the rest of the paper. Section 3 discusses the equilibrium behavior of the model in the thermodynamic limit by analyzing the free energy density computed under the assumption of RS: a detailed discussion of the phase diagram of the model is given both in the presence of noise and in the noiseless limit; furthermore, section 3.3 presents a more peculiar analysis obtained using a 1-step replica symmetry breaking (1-RSB) ansatz. Section 4 reports some numerical simulations used to test theoretical results. Section 5 provides a simple justification of how pattern retrieval is attained in the BAM by interpolating the model with two independent Hopfield models (solely in the low-load regime for the sake of simplicity). Finally, section 6 summarizes our results and presents future directions to be explored. Calculation details are relegated to the appendix, where a formal equivalence between the BAM and two coupled RBMs is further discussed.

2. Model

The BAM model is defined by two sets of neurons, embedded in two vectors σ and $\bar{\sigma}$ for layer 1 and 2, respectively. Each unit in one layer interacts with all the units in the other layer, so that the interaction topology is a bipartite graph; for a schematic representation of the model, see figure 1. Layer 1 (resp. 2) has N (resp. \bar{N}) neurons, called—in the first (resp. second layer) as σ_i (resp. $\bar{\sigma}_j$) with $i \in \{1, \dots, N\}$ (resp. $j \in \{1, \dots, \bar{N}\}$) [25]. Each neuron is hereafter considered




Figure 1. Schematic representation of a BAM with $N = 6$ and $\bar{N} = 4$. The (binary) neurons are represented by black circles; each solid line defines a synaptic weight connecting two neurons of different layers.

as a binary variable (i.e. an Ising spin), namely $\sigma_i, \bar{\sigma}_j \in \{-1, 1\}$, where the state $+1$ is typically associated with a firing state for the neuron, while -1 means the neuron is silent. The BAM Hamiltonian (or cost-function) is given by:

$$H(\sigma, \bar{\sigma}) = - \sum_{i=1}^N \sum_{j=1}^{\bar{N}} w_{ij} \sigma_i \bar{\sigma}_j, \tag{1}$$

where w_{ij} represents the synaptic interaction between node i in the first layer and the node j in the second layer. The interaction matrix $\mathbf{W} = \{w_{ij}\}_{i=1, \dots, N}^{j=1, \dots, \bar{N}}$ is constructed in such a way that the usual Hebb's rule pattern storage is generalized to such a bipartite structure. We therefore define two sets of K patterns (one per layer), denoted with ξ^μ and $\bar{\xi}^\mu$, with $\mu \in \{1, \dots, K\}$. Each pattern has a dimension compatible with the corresponding layer of neurons, so that $\xi^\mu, \bar{\xi}^\mu$ are vectors with N and \bar{N} components, respectively, for each μ . The synaptic weight matrix is constructed by using the following generalized Hebb's rule:

$$w_{ij} = \frac{1}{\sqrt{N\bar{N}}} \sum_{\mu=1}^K \xi_i^\mu \bar{\xi}_j^\mu. \tag{2}$$

The scale factor in the weight matrix is chosen so to have a non-trivial free energy density in the thermodynamic limit where $N, \bar{N}, L := \sqrt{N\bar{N}} \rightarrow \infty$.

To inspect the ability of the BAM to handle pairs of patterns, given the bipartite structure of the network and the absence of intra-synaptic interactions, it is useful to check that retrieval is feasible when both layers have a non-zero overlap with a given pattern. In what follows, we assume that the patterns are drawn independently with i.i.d. components at each layer (see the next section for more details on the distribution from which they are drawn). Under appropriate initial and external conditions (the latter depend on the number of patterns K), a neural dynamics starting with one of the layers close enough to one of the K patterns will eventually reach a fixed point with a certain nonzero overlap on both layers: In the noiseless regime, the attractors of the dynamics—in pairs $(\xi^\mu, \bar{\xi}^\mu)$ —satisfy the following steepest descent dynamical equations:

$$\sigma_i^{t+1} = \text{sign} \left(\sum_j w_{ij} \bar{\sigma}_j^t \right) \quad \longrightarrow \quad \xi_i^\mu = \text{sign} \left(\sum_j w_{ij} \bar{\xi}_j^\mu \right) \quad (3)$$

$$\bar{\sigma}_j^{t+1} = \text{sign} \left(\sum_i w_{ij} \sigma_i^t \right) \quad \longrightarrow \quad \bar{\xi}_j^\mu = \text{sign} \left(\sum_i w_{ij} \xi_i^\mu \right). \quad (4)$$

where the right equations highlight how pattern pairs $(\xi^\mu, \bar{\xi}^\mu)$ are fixed points of such a dynamics. Here we used the orthogonality between patterns, which holds in the thermodynamic limit, namely

$$\frac{1}{N} \sum_i \xi_i^\mu \xi_i^\nu = \delta_{\mu,\nu} + O\left(\frac{1}{\sqrt{N}}\right) \quad \text{and} \quad \frac{1}{N} \sum_j \bar{\xi}_j^\mu \bar{\xi}_j^\nu = \delta_{\mu,\nu} + O\left(\frac{1}{\sqrt{N}}\right). \quad (5)$$

It is easy to check that pattern pairs (i.e. the attractors of equations (3) and (4)) are extremal points of the energy function, by writing it in the following form:

$$H(\sigma, \bar{\sigma}) = - \sum_{i,j} w_{ij} \sigma_i \bar{\sigma}_j = - \frac{1}{L} \sum_{\mu=1}^K \left(\sum_{i=1}^N \xi_i^\mu \sigma_i \right) \left(\sum_{j=1}^{\bar{N}} \bar{\xi}_j^\mu \bar{\sigma}_j \right) = -L \sum_{\mu=1}^K m_\mu(\sigma) \bar{m}_\mu(\bar{\sigma}), \quad (6)$$

where we introduced, as order parameters of the theory, the Mattis overlap $m_\mu(\sigma)$, $\bar{m}_\mu(\bar{\sigma})$ of the configurations $(\sigma, \bar{\sigma})$ with the pattern pair $(\xi^\mu, \bar{\xi}^\mu)$ defined as

$$m_\mu(\sigma) = \frac{1}{N} \sum_i \xi_i^\mu \sigma_i \quad \text{and} \quad \bar{m}_\mu(\bar{\sigma}) = \frac{1}{\bar{N}} \sum_j \bar{\xi}_j^\mu \bar{\sigma}_j. \quad (7)$$

3. Equilibrium analysis in the high-load regime

3.1. Basic setting

In this section, we discuss how to characterize the phase diagram of the BAM using techniques from statistical physics. We begin by defining the partition function at a fixed realization of the pattern components, at a given inverse temperature $\beta \in \mathbb{R}^+$:

$$Z(\beta, \{(\xi^\mu, \bar{\xi}^\mu)\}_{\mu=1}^K) = \sum_{\sigma} \sum_{\bar{\sigma}} \exp \left[\frac{\beta}{L} \sum_{\mu=1}^K \left(\sum_{i=1}^N \xi_i^\mu \sigma_i \right) \left(\sum_{j=1}^{\bar{N}} \bar{\xi}_j^\mu \bar{\sigma}_j \right) \right]. \quad (8)$$

The temperature $T = \beta^{-1}$ plays the role of an external noise affecting the neural dynamics. Patterns are quenched with respect to the timescale of thermal fluctuations in the (physiological) neural dynamics, and are treated, as standard in spin-glass (SG) models, as quenched disorder to be averaged over. Moreover, we assume that synaptic couplings already exist in the form (2), therefore no learning process is considered. The physical quantity containing all the information about the equilibrium behavior of the model is the quenched free energy density

$$f(\beta, \alpha, \gamma) = - \frac{1}{\beta} \lim_{L \rightarrow \infty} \frac{1}{L} \mathbb{E}_{\xi, \bar{\xi}} \log Z(\beta, \{(\xi^\mu, \bar{\xi}^\mu)\}_{\mu=1}^K), \quad (9)$$

where the operator \mathbb{E} denotes the expectation value w.r.t. the patterns' distribution. In the thermodynamic limit, the free energy density will depend on three control parameters, i.e. the thermal noise β , the *network load* α ,

$$\alpha \equiv \frac{K}{L}, \quad (10)$$

that reflects the ratio between the number of patterns embedded in the network and the relevant system size of the model L , and the *asymmetry* between layers (or the *shape* parameter),

$$\gamma \equiv \sqrt{\frac{N}{\bar{N}}}, \quad (11)$$

quantifying the ratio between the two layer sizes [26].

There are two main computational regimes under which this model can be analyzed: a first one in which the number of patterns is finite with respect to L —or growing up to $\sim \log L$ —so that $\lim_{L \rightarrow \infty} \alpha = 0$, called *low-load* regime (whose treatment is mathematically simpler), and the more challenging *high-load* regime, in which the number of patterns also scales linearly with L , so that $\alpha \in \mathbb{R}^+$ even in the thermodynamic limit (and techniques from statistical mechanics of SGs need to be introduced in the mathematical treatment). In both cases, we let both layer sizes go to infinity, while γ remains finite. The low-load scenario was analyzed in [21], where it was proven that a second-order phase transition splits a paramagnetic phase (with no retrieval) and a ferromagnetic phase, where retrieval spontaneously emerges. In this regime, when considering the retrieval of only one pair of patterns, the BAM shows a qualitatively similar behavior to the (bipartite) Curie–Weiss model [27, 28]. We will come back to the low-load scenario in section 5, where a simple justification of how the BAM achieves pattern-pair retrieval will be discussed. The high-load regime was also explored in [21], but results were provided only at zero-temperature. The calculations in the high-load regime are more delicate because they require a more careful treatment of the extensive noisy term in the Hamiltonian. The expectation value in equation (9) is computed w.r.t. all the pattern pairs $(\xi^\mu, \bar{\xi}^\mu)$. To describe the retrieval properties, we now divide the set of pattern pairs into two subsets: the first consists of a finite (or up to $\sim \log L$) number of binary patterns, indexed by $\mu = 1, \dots, l$; the remaining (extensive in the high-load regime) patterns are drawn from i.i.d. Gaussian for each component for $\mu = l+1, \dots, K$. The first set encodes for pattern pairs that can be retrieved, i.e. those having eventually a non-zero overlap with spin configurations, while the other extensive set will act as a quenched noise w.r.t. the former. As a consequence, the Hamiltonian is split into a *signal* and a *noise* term as well:

$$H(\sigma, \bar{\sigma}) = -\frac{1}{L} \sum_{\mu=1}^l \left(\sum_{i=1}^N \xi_i^\mu \sigma_i \right) \left(\sum_{j=1}^{\bar{N}} \bar{\xi}_j^\mu \bar{\sigma}_j \right) - \frac{1}{L} \sum_{\mu=l+1}^K \left(\sum_{i=1}^N \xi_i^\mu \sigma_i \right) \left(\sum_{j=1}^{\bar{N}} \bar{\xi}_j^\mu \bar{\sigma}_j \right). \quad (12)$$

The reason behind this choice generalizes standard arguments for the Hopfield model with binary neurons, where information retrieval is optimally achieved in the saturation regime (i.e. at finite load $\alpha > 0$) when the patterns to be recalled are binary-valued, see for instance [9, 29, 30]. On the other hand, the extensive part of the pattern set can be treated in different ways, according to standard universality arguments about quenched disorder in SG models [31, 32].

3.1.1. Decoupling. In fully connected spin glasses with quadratic interactions, the Boltzmann weight is usually simplified by introducing a Hubbard–Stratonovich integral transformation. The two-layer structure of the BAM instead requires a different decoupling transformation. Following standard approaches used to analyze bipartite models (such as the bipartite SK model [33] or RBMs [34, 35]), an integral transformation involving a pair of complex conjugate variables (z, z^\dagger) is a more convenient way to decouple the interacting term in the Hamiltonian:

$$\exp[\beta x \bar{x}] = \frac{\beta}{\pi} \int dz dz^\dagger \exp[-\beta z z^\dagger + \beta z x + \beta z^\dagger \bar{x}]. \tag{13}$$

The above identity can be easily verified by writing the two integration variables in terms of their real and imaginary parts, i.e. by setting $z = u + iv, z^\dagger = u - iv$: the resulting object is a simple Gaussian integral over the real plane. Thus, equation (13) is used to decouple the interaction terms in the partition functions equation (8). introducing a pair of complex conjugate variables for each pattern μ . Note that it is necessary to decouple the extensive noisy part of the patterns, i.e. $\mu \in \{l + 1, \dots, K\}$, while the signal term is left untouched. After applying the transformation (13), the partition function (8) reads:

$$\begin{aligned} Z(\beta, \{(\xi^\mu, \bar{\xi}^\mu)\}_{\mu=1}^K) &= \sum_{\sigma, \bar{\sigma}} \exp \left\{ \frac{\beta}{L} \sum_{\mu=1}^l \left(\sum_i \xi_i^\mu \sigma_i \right) \left(\sum_j \bar{\xi}_j^\mu \bar{\sigma}_j \right) \right\} \times \prod_{\mu=l+1}^K \int dz_\mu z_\mu^\dagger \\ &\times \exp \left[-\beta z_\mu z_\mu^\dagger + \frac{\beta}{\sqrt{N}} z_\mu \sum_i \xi_i^\mu \sigma_i + \frac{\beta}{\sqrt{N}} z_\mu^\dagger \sum_j \bar{\xi}_j^\mu \bar{\sigma}_j \right]. \end{aligned} \tag{14}$$

We must now average $\log Z$ over the quenched disorder to calculate the free energy of the system (9), whose extremization w.r.t. the order parameters will provide a set of self-consistent equations that, in turn, determine the phase diagram of the model.

Finally, notice how the integral transform equation (13) can be exploited to derive a structural analogy between the BAM’s partition function and the partition function of two coupled RBMs, by generalizing the well-know equivalence between the Hopfield model and a Binary-Gaussian RBM [15], further discussed in appendix C.

3.2. RS phase diagram

The computation of (9)–(14) is performed using the interpolation technique developed by Guerra [23]. This method, restricting its usage to the RS case analyzed here, allows one to compute the quenched free energy density of a given model whose computation is typically intractable due to the random interactions in the model’s Hamiltonian. The method consists in defining an interpolating free energy function expressed in terms of an additional parameter $t \in [0, 1]$: when the value of this parameter is $t = 0$, the partition function corresponds to a system of independent variables and can be factorized over the spins, which can be computed easily. When $t = 1$, we recover the original free energy to be computed (given here by (9)). For instance, considering just the signal term in equations (12)–(14), this is equivalent to consider the following generalized (t –dependent) Hamiltonian:

$$\mathcal{H}_t^{\text{inter}}(\sigma, \bar{\sigma}) = t \mathcal{H}^{\text{BAM}}(\sigma, \bar{\sigma}) + (1 - t) \beta \sum_{\mu=1}^K \left[\Psi_\mu \sum_i \xi_i^\mu \sigma_i + \bar{\Psi}_\mu \sum_j \bar{\xi}_j^\mu \bar{\sigma}_j \right]. \tag{15}$$

Finally, the method exploits the fundamental theorem of calculus to compute the free energy at $t = 1$ knowing the initial (Cauchy) condition at $t = 0$, plus the integral of the derivative of the interpolating free energy with respect to t : by assuming RS, such an integral becomes analytic. It is worth mentioning that this interpolation method has been shown to be rigorous in some disordered models [36, 37]. This previous qualitative explanation is expanded and adapted to our model's case in appendix A, where all calculation details are explicitly given. The results of this interpolation analysis are described and summarized below.

In the thermodynamic limit, the equilibrium behavior of the BAM is described by a finite number of order parameters, each of which is determined self-consistently by imposing stationary conditions on the resulting free energy. By analogy with the Hopfield model, we can immediately identify the relevant set of order parameters:

$$m_\mu = \frac{1}{N} \sum_i \xi_i^\mu \sigma_i, \quad \bar{m}_\mu = \frac{1}{N} \sum_j \bar{\xi}_j^\mu \bar{\sigma}_j; \quad (16)$$

$$q_{12} = \frac{1}{N} \sum_i \sigma_i^{(1)} \sigma_i^{(2)}, \quad \bar{q}_{12} = \frac{1}{N} \sum_j \bar{\sigma}_j^{(1)} \bar{\sigma}_j^{(2)}; \quad (17)$$

$$p_{12} = \frac{1}{K} \sum_{\mu>l} z_\mu^{(1)} z_\mu^{(2)}, \quad \bar{p}_{12} = \frac{1}{K} \sum_{\mu>l} z_\mu^{\dagger(1)} z_\mu^{\dagger(2)}. \quad (18)$$

The first line reminds the set of *Mattis magnetizations*, representing the projection of a spin configuration (on either one of the two layers) along one of the retrievable patterns, i.e. $\mu \in \{1, \dots, l\}$. By using the short-hand notation $\mathbf{m} = (m_1, \dots, m_l)^T$ and $\bar{\mathbf{m}} = (\bar{m}_1, \dots, \bar{m}_l)^T$, in the retrieval phase, at least one component of each of these vectors (in particular, those with the same index) is expected to be nonzero, so that the free energy minima are characterized by spin configurations with a finite overlap with a particular pattern pair. The second line defines the standard overlap matrix between two replicas (here indexed by 1, 2) for each layer. Finally, the last line defines a set of overlaps between the complex variables introduced by the integral transform (14). Under the RS assumption, each of these order parameters does not fluctuate in the thermodynamic limit and its distribution reaches a delta-function peaked around a unique value determined by imposing stationarity on the free energy, namely

$$\lim_{L \rightarrow \infty} P(\langle \mathcal{O} \rangle) = \delta(\langle \mathcal{O} \rangle - O) \quad (19)$$

where \mathcal{O} denotes any of the order parameters defined in equations (16)–(18), and O denotes the corresponding equilibrium value. To simplify the notation, in the following we highlight the RS value of each of the above order parameters with a capital letter, to be distinguished from the general expression in small. The RS free energy density of the BAM is given by the following expression (see appendix A for calculation details):

$$\begin{aligned} f(\beta, \alpha, \gamma) = & \sum_{\mu=1}^l M_\mu \bar{M}_\mu + \frac{\alpha\beta}{2} P(1 - Q) + \frac{\alpha\beta}{2} \bar{P}(1 - \bar{Q}) + \\ & - \frac{\gamma}{\beta} \mathbb{E}_{\eta, \xi} \log 2 \cosh \left[\beta \sqrt{\gamma} \alpha P \eta + \beta \bar{\gamma} \sum_{\mu} \bar{M}_\mu \xi^\mu \right] - \frac{\bar{\gamma}}{\beta} \mathbb{E}_{\bar{\eta}, \bar{\xi}} \log 2 \cosh \\ & \times \left[\beta \sqrt{\gamma} \alpha \bar{P} \bar{\eta} + \beta \gamma \sum_{\mu} M_\mu \bar{\xi}^\mu \right] + \frac{\alpha}{2\beta} \log \Delta - \frac{\alpha\beta}{2\Delta} [Q(1 - \bar{Q}) + \bar{Q}(1 - Q)], \quad (20) \end{aligned}$$

where $\Delta = 1 - \beta^2 (1 - Q)(1 - \bar{Q})$, $\bar{\gamma} = \gamma^{-1}$ and $\eta, \bar{\eta}$ are two i.i.d. standard Gaussian variables. The above expression requires Δ to be always positive: such a condition is always verified




Figure 2. Phase diagram of the BAM in the plane (α, T) at different values of γ . (a): $\gamma = 1$. (b): $\gamma = 2$; (c): $\gamma = 5$. The labels **P**, **SG**, **R**, **MR** (shown only in (a)) stand for paramagnetic, spin-glass, retrieval and metastable-retrieval, respectively.

at any point of the phase diagram. We stress that the same result can be obtained using the more standard saddle point method through the replica trick [24], as originally carried out in [21]. However, in [21] there clearly was a misprint or a small error in earlier calculations (for instance, the sign of Δ was exchanged, leading to a free-energy being ill-defined at finite β) as results were not consistent with our numerical simulations (discussed in section 4).

The values of the order parameters at fixed control parameters can be obtained by evaluating the saddle points of f of equation (20). The set of self-consistent equations to be fulfilled by these values is shown below:

$$\mathbf{M} = \mathbb{E}_{\eta, \xi} \xi \tanh \left[\beta \sqrt{\gamma \alpha P} \eta + \beta \gamma (\xi \cdot \bar{\mathbf{M}}) \right], \quad \bar{\mathbf{M}} = \mathbb{E}_{\bar{\eta}, \bar{\xi}} \bar{\xi} \tanh \left[\beta \sqrt{\gamma \alpha \bar{P}} \bar{\eta} + \beta \gamma (\bar{\xi} \cdot \mathbf{M}) \right]; \tag{21a}$$

$$Q = \mathbb{E}_{\eta, \xi} \tanh^2 \left[\beta \sqrt{\gamma \alpha P} \eta + \beta \gamma (\xi \cdot \bar{\mathbf{M}}) \right], \quad \bar{Q} = \mathbb{E}_{\bar{\eta}, \bar{\xi}} \tanh^2 \left[\beta \sqrt{\gamma \alpha \bar{P}} \bar{\eta} + \beta \gamma (\bar{\xi} \cdot \mathbf{M}) \right]; \tag{21b}$$

$$P = \Delta^{-2} \left[\bar{Q} + \beta^2 Q (1 - \bar{Q})^2 \right], \quad \bar{P} = \Delta^{-2} \left[Q + \beta^2 \bar{Q} (1 - Q)^2 \right]. \tag{21c}$$

The model phase diagram can be fully characterized by solving the above saddle point equations (21a)–(21c) at any value of the three control parameters (α, β, γ) , and evaluating the corresponding free energy for each of the fixed points. For simplicity, we just consider the retrieval of 1 pattern pair: this is equivalent to assume that only one component of the vectors $\mathbf{M}, \bar{\mathbf{M}}$ is non-zero, namely $M_\nu = M_\mu \delta_{\mu, \nu}$ and $\bar{M}_\nu = \bar{M}_\mu \delta_{\mu, \nu}$ with μ being the pattern index to be retrieved (one could set $\mu = 1$ without loss of generality).

We start from the symmetric case where $\gamma = 1$ (we recall the parameter γ controls the ratio between the size of the two layers, see (11)), and the corresponding phase diagram is shown in figure 2(a) in the plane $(\alpha, T = \beta^{-1})$.

At high temperature, the equilibrium phase is *paramagnetic* (P) with all the order parameters being null. At $T < 1$ and sufficiently small loads α , the model exhibits a ferromagnetic-like, so-called *retrieval* (R) phase, with all the order parameters being non-null. Further increasing the load α over a certain threshold $\alpha_c(T)$ marks the onset of an SG phase with zero Mattis magnetizations over all patterns but non-zero overlaps $Q, \bar{Q} \neq 0$. A more detailed inspection shows how the retrieval phase can be divided in two sub-regimes: a first (red) portion of the

phase diagram where the retrieval fixed point has the lowest free energy, and another one (yellow) where the retrieval fixed point is a local stationary point. In the former, while the retrieval phase is dominant, the SG regime is still locally stable. In the latter, the global minimum of the free energy is the SG fixed point: however, the retrieval solution is locally stable (metastable, MR), so that it is still possible to find it by starting with an initial condition (either in the saddle point equations or in the neural dynamics for finite size systems) that has a sufficiently large overlap with a pattern pair in the retrievable set.

In the asymmetric case ($\gamma \neq 1$) the phase diagram looks qualitatively similar to the previous one, as clear from figures 2(b) and (c), showing respectively the result at $\gamma = 2$ and $\gamma = 5$. However, the retrieval region of the diagram shrinks as γ departs from 1 (i.e. from the completely symmetric case), so that retrieval loss occurs at lower values of α . Moreover, the critical line separating the paramagnetic from the SG phases moves to higher temperatures. As an overall remark, we can conclude that increasing the asymmetry between the two layers extends the SG phase at the expense of the others. An important property to notice is that the above free energy (and consequently, the phase diagram) is symmetric under the exchange of the two layers, provided that all the order parameters are swapped and $\gamma \leftrightarrow \bar{\gamma} = 1/\gamma$.

A numerical characterization of the transition lines for different values of γ is shown in figure 3: the metastable retrieval-spin glass (MR-SG) in (a), the stable-MR in (b), and the paramagnetic-spin glass (P-SG) in (c). In all cases, we find that the greater the asymmetry of the BAM, the narrower the retrieval region and the wider the spin glass. These three transition lines are obtained as follows: the first-order spinodal of the retrieval state, corresponding to the line between the metastable-retrieval phase and the spin glass one, is found by solving the saddle point equations starting from a large overlap with one pattern, until such a solution ceases to exist upon increasing α . The critical line between the stable and metastable retrieval phase is instead found by looking at the points where the free energies of the retrieval state and the spin glass solution (the latter obtained by setting $M = \bar{M} = 0$ in the saddle point equations) are equal. Finally, the P-SG line allows an analytical description and will therefore be discussed in detail in the next paragraph. We finally notice that in figures 3(a) and (b) the transition lines' reentrance at small values of T is due to an instability of the retrieval solution in the replica space, evidencing the failure of the RS ansatz in this regime, analogously to the Hopfield model. This was noted in [22] in the symmetric case $\gamma = 1$. For completeness, in the next section we investigate the low-temperature behavior of the BAM using a 1-RSB ansatz.

3.2.1. P-SG transition line. The critical line between the paramagnetic and SG phases can be further analyzed analytically because it is of second order. We proceed by setting $M = \bar{M} = 0$ into the saddle-point equations and expand equation (21b) around $Q, \bar{Q} \approx 0$. To the first order, the Jacobian of the linearized systems and its eigenvalues are shown below:

$$\mathcal{J} = \frac{\alpha\beta^2}{(1-\beta^2)^2} \begin{bmatrix} \bar{\gamma}\beta^2 & \bar{\gamma} \\ \gamma & \gamma\beta^2 \end{bmatrix}$$

$$\lambda_{\pm}(\alpha, \beta, \gamma) = \frac{\alpha\beta^2}{2\gamma(1-\beta^2)^2} \left[\beta^2(1+\gamma^2) \pm \sqrt{4\gamma^2 + \beta^4(1-\gamma^2)^2} \right]. \quad (22)$$

At fixed γ , the critical line in the plane (α, T) is given by the first eigenvalue becoming larger than 1 upon lowering the temperature. In particular, in the symmetric case $\gamma = 1$, the equation $\lambda_{\pm}(\alpha, \beta, \gamma = 1) = 1$ can be solved analytically, leading to the following expression for the critical temperature:

$$T_{\text{P-SG}}(\alpha, \gamma = 1) = \sqrt{1 + \frac{\alpha}{2} + \frac{1}{2}\sqrt{\alpha(\alpha + 8)}}. \quad (23)$$

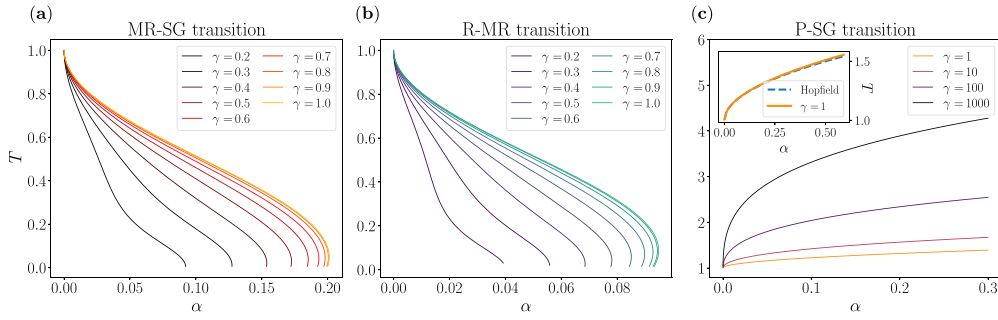


Figure 3. Critical lines separating the different operating regimes of the BAM. In (a), critical lines for the MR-SG. In (b), critical lines for the R-MR transition at different values of γ . In (c), critical lines for the P-SG transition at different values of γ . All the lines are plotted as function of α (please mind the three panels cover different ranges of α), for different values of γ . The inset in panel (c) shows a comparison of the P-SG critical lines between the symmetric BAM, given by (23), and the same critical line computed for the Hopfield model (after a proper rescaling to take into account the different definition of the load in the two models).

A comparison between (23) and the Hopfield critical line as predicted by the AGS theory [1] is shown in the inset of figure 3(c): however, attention must be paid to the different definitions of α in the two models. In the Hopfield case, the network load needs to be defined as $\alpha^{\text{Hopfield}} = K/(N + \bar{N})$ in order to be consistent with the overall dimension of patterns in both models. In the symmetric case where $\gamma = 1$, this implies $\alpha^{\text{Hopfield}} = \alpha^{\text{BAM}}/2$. As a consequence, the Hopfield’s critical line plotted in figure 3 corresponds to $T_c^{\text{Hopfield}} = 1 + \sqrt{\alpha^{\text{BAM}}/2}$, showing a very similar behavior between the two curves (they only depart at very high values of α). In the asymmetric case, critical lines can be found numerically by solving $\lambda_{\pm}(\alpha, \beta, \gamma) = 1$ at different values of γ and α , plotted in figure 3(c). Notice how, at fixed load α , such a phase transition is pushed at higher temperature when the asymmetry between the two layers increases.

3.2.1.1. RS phase diagram at $T=0$. The phase diagram in the noiseless regime can be computed by taking the $T=0$ limit of the free energy (20) and/or the corresponding saddle point equations (21a)–(21c). The calculations follow the same reasoning as discussed in [21], here we just report the final result. Again, we focus on the retrieval of one pattern pair, so that the vectors \mathbf{M} and $\bar{\mathbf{M}}$ defined in equation (16) have only one nonzero component by construction. At $\beta \rightarrow \infty$, the overlaps Q, \bar{Q} tend to 1, in such a way that the SG susceptibilities on the two spin sets $\chi = \beta(1 - Q)$ and $\bar{\chi} = \beta(1 - \bar{Q})$ maintain finite values. The saddle-point equations can therefore be re-arranged in a different form, to determine self-consistently the susceptibilities and the Mattis magnetizations of each layer. After taking the zero-temperature limit, the new saddle-point equations read:

$$\frac{1 + \bar{\chi}^2}{(1 - \chi\bar{\chi})^2} = \frac{\text{erf}^2(\bar{y})}{2\gamma\alpha y^2}, \quad \chi = \frac{2}{\bar{\gamma}\sqrt{\pi}} \frac{y}{\text{erf}(\bar{y})} e^{-y^2}, \quad (24)$$

$$\frac{1 + \chi^2}{(1 - \chi\bar{\chi})^2} = \frac{\text{erf}^2(y)}{2\bar{\gamma}\alpha\bar{y}^2}, \quad \bar{\chi} = \frac{2}{\gamma\sqrt{\pi}} \frac{\bar{y}}{\text{erf}(y)} e^{-\bar{y}^2}, \quad (25)$$

where $\text{erf}(\cdot)$ denotes the error function, and the variables y, \bar{y} are linked to the Mattis magnetizations of each layer through

$$M = \text{erf}(y), \quad \bar{M} = \text{erf}(\bar{y}). \quad (26)$$

The above system can be solved w.r.t. the variables y, \bar{y} at different values of α, γ . The resulting phase diagram is shown in figure 4(a). The gray line separating the MR phase (in yellow) and the SG phase (violet) defines the critical storage capacity $\alpha_c(\gamma)$ of the BAM. Notice how the maximum value is reached at $\gamma = 1$, where $\alpha_c(\gamma = 1) \approx 0.2$, consistently with [20, 22]. Intuitively, this is understood by considering the number of weights encoded in the synaptic matrix (2) of the BAM, which is equal to $N\bar{N} = L^2$. Considering a network with a fixed number of neurons $N + \bar{N}$ (or equivalently, with a fixed pattern dimension), it can be immediately verified that the maximum number of weight entries in the BAM is reached when $\gamma = 1$. The general idea is that the more weights there are in the network, the larger the retrieval phase with respect to the quenched noise induced by the extensive number of patterns stored in the network. We also mention how in the extremely asymmetric limit where e.g. $\gamma \rightarrow 0$ the storage capacity α_c goes to 0 linearly with γ , so that the rescaled capacity $\tilde{\alpha}_c = \alpha_c/\gamma$ tends to a finite value $\tilde{\alpha}_c \rightarrow 0.497$. By symmetry, the same result is achieved when $\gamma \rightarrow \infty$, and defining the rescaled capacity as $\tilde{\alpha}_c = \gamma\alpha_c = \alpha_c/\tilde{\gamma}$. Figures 4(b) and (c) instead show 4 horizontal cuts of the phase diagram at different values of γ (as indicated in the caption). Specifically, in each panel we plot the Mattis magnetization of the fixed point for the two layers (in (b) and (c) respectively) as a function of α , showing a first-order transition between the retrieval and SG phase. Interestingly, the Mattis magnetization of the largest layer, show in figure 4(b), decreases continuously before the critical value α_c .

We finally quantify the difference between the Hopfield model and the BAM in terms of their storage capabilities: in order to have a fair comparison between the two models, it is more convenient to renormalize their critical capacities by the total number of neurons $N + \bar{N}$, which is kept fixed whilst varying the asymmetry between the two layers. In this notation, while the Hopfield's critical capacity is equal to $\alpha_c^{\text{Hopfield}} = K_c/(N + \bar{N}) \approx 0.138$ independently on γ , the BAM's critical capacity is given by

$$\tilde{\alpha}_c^{\text{BAM}} = \frac{K_c^{\text{BAM}}}{N + \bar{N}} = \frac{\alpha_c(\gamma)}{\gamma + \gamma^{-1}}, \quad (27)$$

where at the right-most hand side α_c is the quantity plotted in figure 4(a). These two quantities are shown in figure 5(a): as expected, we conclude that BAM's critical capacity is lower than the Hopfield model even in its most stable configuration (i.e. the symmetric case). On the other hand, the BAM is more efficient from the point of view of the number of weights stored in the network, since for a fixed number of patterns such that the network operates in its retrieval phase, it requires (at most, in the symmetric case) about half of the weights compared to the Hopfield model (see figure 5(b)).

3.3. 1-RSB analysis

A more peculiar investigation of the phase diagram requires analyzing the stability of the RS solution. For the Hopfield model, while the SG solution of the RS fixed point equations is always be unstable below the P-SG critical line (below which the system is not anymore ergodic), the instability of the retrieval state occurs only at very low temperature close to the spinodal curve of the MR-SG [1, 38], below the so-called de Almeida Thouless line [39]. For the BAM, this phenomenon was already observed [22] in the symmetric scenario where

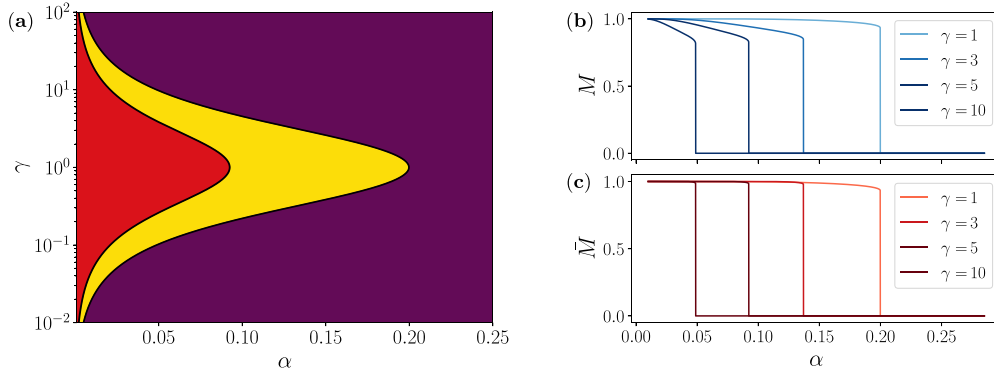


Figure 4. Panel (a): BAM’s phase diagram at $T = 0$ in the plane (α, γ) : the same color-code as in figure 2 is used. The phase boundary separating the MR (yellow) from the SG (violet) phases defines the RS critical capacity α_c as a function of γ . As the vertical axis has a logarithmic scale, we notice again how the phase diagram is symmetric under the transformation $\gamma \rightarrow \gamma^{-1} = \bar{\gamma}$. Panels (b) and (c) show the zero-temperature Mattis magnetizations of the two layers (plotted resp. in panels (b) and (c)) as functions of α for 4 different values of γ .




Figure 5. (a) Critical capacity of BAM and the Hopfield but normalized by the total number of neurons in the network (i.e. $N + \bar{N}$ instead of L as in (10) for the BAM). The green line corresponds to equation (27). (b) Number of weights used to store this critical number of patterns in the BAM normalized by the analogous number of weights in the Hopfield model⁵. In both panels, each quantity is plotted as a function of the asymmetry $\gamma = \sqrt{N/\bar{N}}$.

$\gamma = 1$ (generalization to the asymmetric case being straightforward). Therefore, we now perform a 1-RSB analysis of the BAM to investigate how it impacts the critical capacity of the model. In general, this ansatz implies that the overlap distribution acquires a bi-modal structure with two equilibrium values, and this property holds for any of the four overlaps defined in equations (17) and (18). Specifically, let us denote with \mathcal{O} either one of the four overlaps $q_{12}, \bar{q}_{12}, p_{12}, \bar{p}_{12}$, its probability distribution will be given in the thermodynamic limit by

$$\lim_{L \rightarrow \infty} P(\langle \mathcal{O} \rangle) = \theta \delta(\langle \mathcal{O} \rangle - O_1) + (1 - \theta) \delta(\langle \mathcal{O} \rangle - O_2) \quad (28)$$

⁵ Note that this quotient is equal to $N\bar{N} / \binom{N+\bar{N}}{2}$.

where O_1, O_2 are the two equilibrium values weighted by $\theta \in [0, 1]$. On the other hand, magnetizations still are supposed to self-average as in the RS case (cf equation (19)) The computation of the 1-RSB free energy can be easily carried out either by using the replica trick at the 1-RSB level or by extending the technique presented in [40] where an interpolation method at the 1-RSB level is presented for the Hopfield model. Regarding the BAM, one just needs to adapt the definition of the interpolating pressure as discussed in appendix A for the RS case. In either cases, and assuming and $O_2 > O_1$ for any of the four overlaps in equations (17) and (18) the final result reads

$$\begin{aligned}
 f(\beta, \alpha, \gamma) = & \sum_{\mu=1}^l M_{\mu} \bar{M}_{\mu} + \frac{\alpha\beta}{2} P_2 [1 + (\theta - 1) Q_2] + \frac{\alpha\beta}{2} \bar{P}_2 [1 + (\theta - 1) \bar{Q}_2] \\
 & - \frac{\alpha\beta}{2} \theta [Q_1 P_1 + \bar{Q}_1 \bar{P}_1] + \frac{\alpha}{2\beta} \log \Delta + \frac{\alpha}{2\beta\theta} \log \frac{\Delta_{\theta}}{\Delta} - \frac{\alpha}{2\Delta_{\theta}} (\bar{Q}_1 \kappa + Q_1 \bar{\kappa}) \\
 & - \frac{\gamma}{\beta\theta} \mathbb{E}_{\xi, \tau} \log \mathbb{E}_z 2^{\theta} \cosh^{\theta} \Xi - \frac{\bar{\gamma}}{\beta\theta} \mathbb{E}_{\bar{\xi}, \tau} \log \mathbb{E}_z 2^{\theta} \cosh^{\theta} \bar{\Xi}
 \end{aligned} \tag{29}$$

where we defined

$$\begin{aligned}
 \Xi &= \beta \left[\bar{\gamma} (\mathbf{M} \cdot \xi) + \sqrt{\bar{\gamma}\alpha P_1} \tau + \sqrt{\bar{\gamma}\alpha (P_2 - P_1)} z \right], \\
 \bar{\Xi} &= \beta \left[\gamma (\mathbf{M} \cdot \bar{\xi}) + \sqrt{\gamma\alpha \bar{P}_1} \tau + \sqrt{\gamma\alpha (\bar{P}_2 - \bar{P}_1)} z \right],
 \end{aligned} \tag{30}$$

and

$$\kappa = \beta (1 - Q_2) + \beta\theta\delta Q \tag{31a}$$

$$\bar{\kappa} = \beta (1 - \bar{Q}_2) + \beta\theta\delta \bar{Q} \tag{31b}$$

$$\Delta_{\theta} = 1 - \kappa\bar{\kappa}, \tag{31c}$$

where $\delta Q = Q_2 - Q_1$, $\delta \bar{Q} = \bar{Q}_2 - \bar{Q}_1$. Finally, Δ has the same expression as in the previous section (see equation (20)). Note that the RS free-energy is recovered in the limit $O_2 - O_1 \rightarrow 0$ with $O \in \{Q, \bar{Q}, P, \bar{P}\}$, or alternatively in one of the two extreme values of $\theta = 0, 1$ where the dependency on one of the two values assumed by each overlap disappears. Once again, the free energy solely depends on the three control parameters (α, β, γ) , the value of all the order parameters being determined self-consistently through the following equations:

$$\mathbf{M} = \mathbb{E}_{\xi, \tau} \xi \frac{\mathbb{E}_z \tanh \Xi \cosh^{\theta} \Xi}{\mathbb{E}_z \cosh^{\theta} \Xi}, \quad \bar{\mathbf{M}} = \mathbb{E}_{\bar{\xi}, \tau} \bar{\xi} \frac{\mathbb{E}_z \tanh \bar{\Xi} \cosh^{\theta} \bar{\Xi}}{\mathbb{E}_z \cosh^{\theta} \bar{\Xi}}, \tag{32a}$$

$$Q_1 = \mathbb{E}_{\xi, \tau} \left[\frac{\mathbb{E}_z \tanh \Xi \cosh^{\theta} \Xi}{\mathbb{E}_z \cosh^{\theta} \Xi} \right]^2, \quad \bar{Q}_1 = \mathbb{E}_{\bar{\xi}, \tau} \left[\frac{\mathbb{E}_z \tanh \bar{\Xi} \cosh^{\theta} \bar{\Xi}}{\mathbb{E}_z \cosh^{\theta} \bar{\Xi}} \right]^2, \tag{32b}$$

$$Q_2 = \mathbb{E}_{\xi, \tau} \frac{\mathbb{E}_z \tanh^2 \Xi \cosh^{\theta} \Xi}{\mathbb{E}_z \cosh^{\theta} \Xi}, \quad \bar{Q}_2 = \mathbb{E}_{\bar{\xi}, \tau} \frac{\mathbb{E}_z \tanh^2 \bar{\Xi} \cosh^{\theta} \bar{\Xi}}{\mathbb{E}_z \cosh^{\theta} \bar{\Xi}}, \tag{32c}$$

$$P_1 = \Delta_{\theta}^{-2} (\bar{Q}_1 + \bar{\kappa}^2 Q_1), \quad \bar{P}_1 = \Delta_{\theta}^{-2} (Q_1 + \kappa^2 \bar{Q}_1), \tag{32d}$$

$$\begin{aligned}
 P_2 &= P_1 + (\Delta \Delta_{\theta})^{-1} [\delta \bar{Q} + \beta (1 - \bar{Q}_2) \bar{\kappa} \delta Q], \\
 \bar{P}_2 &= \bar{P}_1 + (\Delta \Delta_{\theta})^{-1} [\delta Q + \beta (1 - Q_2) \kappa \delta \bar{Q}].
 \end{aligned} \tag{32e}$$

Notice that also θ plays the role of an order parameter, whose equilibrium value is determined by another stationary condition obtained by taking $\partial f/\partial\theta = 0$.

3.3.1. 1-RSB phase diagram at $T=0$. Instead of solving numerically equation (32) at finite temperature we just focus on the $T \rightarrow 0$ limit, with the main goal of determining if the critical capacity increases with respect to the RS estimation, thus reducing the re-entrant behavior of the phase diagram at low temperatures. The analytic limit of the free energy can be easily computed by following the same kind of calculations as in [41, 42]. As discussed in the previous section, the two largest spin overlaps Q_2, \bar{Q}_2 tend to 1 in such a way that the quantities $\chi = \beta(1 - Q_2), \bar{\chi} = \beta(1 - \bar{Q}_2)$ have a finite limit when $\beta \rightarrow \infty$. In addition, the parameter θ is rescaled by β so that $\Theta = \beta\theta$, as pointed out in [41, 43]. With these rescalings, and defining for convenience the overlap differences $\delta Q = 1 - Q_1$ and $\delta\bar{Q} = 1 - \bar{Q}_1$, the inner expectations in equations (29)–(32) (denoted with \mathbb{E}_τ), can be performed analytically; further simplifying by considering just one pattern pair to be retrieved (i.e. $l = 1$ as in section 3.2.1), the 1-RSB free energy in the $T \rightarrow 0$ limit reads

$$\begin{aligned}
 f(\alpha, \gamma) = & M\bar{M} + \frac{\alpha}{2}(P_2\chi + \bar{P}_2\bar{\chi}) + \frac{\alpha}{2}\Theta(P_1\delta Q + \bar{P}_1\delta\bar{Q}) + \frac{\alpha}{2\Theta}\log\frac{\Delta_\Theta}{\Delta} \\
 & - \frac{\alpha}{2\Delta_\Theta}[(1 - \delta\bar{Q})\kappa + (1 - \delta Q)\bar{\kappa}] - \frac{\gamma}{\Theta}\mathbb{E}_\tau\log\mathcal{G}_+(\tau) \\
 & - \frac{\bar{\gamma}}{\Theta}\mathbb{E}_\tau\log\bar{\mathcal{G}}_+(\tau) + \frac{\gamma + \bar{\gamma}}{\Theta}\log 2
 \end{aligned} \tag{33}$$

where $\Delta = 1 - \chi\bar{\chi}$, Δ_Θ is the analogous limit at $T = 0$ of equation (31c) and we defined the following quantities:

$$\begin{aligned}
 \mathcal{G}_\pm(\tau) &= e^{b\Theta}(1 + \operatorname{erf}\omega_+) \pm e^{-b\Theta}(1 + \operatorname{erf}\omega_-), \\
 \bar{\mathcal{G}}_\pm(\tau) &= e^{\bar{b}\Theta}(1 + \operatorname{erf}\bar{\omega}_+) \pm e^{-\bar{b}\Theta}(1 + \operatorname{erf}\bar{\omega}_-),
 \end{aligned} \tag{34a}$$

$$a = \sqrt{\bar{\gamma}\alpha(P_2 - P_1)}, \quad \bar{a} = \sqrt{\gamma\alpha(\bar{P}_2 - \bar{P}_1)}, \tag{34b}$$

$$b(\tau) = \bar{\gamma}M + \sqrt{\bar{\gamma}\alpha P_1}\tau, \quad \bar{b}(\tau) = \gamma M + \sqrt{\gamma\alpha \bar{P}_1}\tau, \tag{34c}$$

$$\omega_\pm(\tau) = \frac{\Theta a}{\sqrt{2}} \pm \frac{b}{\sqrt{2a}}, \quad \bar{\omega}_\pm(\tau) = \frac{\Theta \bar{a}}{\sqrt{2}} \pm \frac{\bar{b}}{\sqrt{2\bar{a}}}. \tag{34d}$$

Note that the functions $\mathcal{G}, \bar{\mathcal{G}}$ explicitly depend on τ through ω_\pm, b and $\bar{\omega}_\pm, \bar{b}$ respectively (equations (34c) and (34d)), but for simplicity we dropped all the dependencies at the r.h.s. of equation (34). The self-consistent equations can be either obtained by imposing stationarity of equation (33) w.r.t. the $M, \bar{M}, \chi, \bar{\chi}, \delta Q, \delta\bar{Q}$ or by directly taking the $T \rightarrow 0$ limit of equation (35). In either case, the final result is given by

$$M = \mathbb{E}_\tau \frac{\mathcal{G}_-(\tau)}{\mathcal{G}_+(\tau)}, \quad \bar{M} = \mathbb{E}_\tau \frac{\bar{\mathcal{G}}_-(\tau)}{\bar{\mathcal{G}}_+(\tau)}, \tag{35a}$$

$$\delta Q = 4\mathbb{E}_\tau \frac{(1 + \operatorname{erf}\omega_+)(1 + \operatorname{erf}\omega_-)}{\mathcal{G}_+^2(\tau)}, \quad \delta\bar{Q} = 4\mathbb{E}_\tau \frac{(1 + \operatorname{erf}\bar{\omega}_+)(1 + \operatorname{erf}\bar{\omega}_-)}{\bar{\mathcal{G}}_+^2(\tau)}, \tag{35b}$$

$$\chi = \sqrt{\frac{8}{\pi}} \frac{1}{a} e^{-\frac{1}{2}a^2\Theta^2} \mathbb{E}_\tau \frac{e^{-\frac{b^2}{2a^2}}}{\mathcal{G}_+(\tau)}, \quad \bar{\chi} = \sqrt{\frac{8}{\pi}} \frac{1}{\bar{a}} e^{-\frac{1}{2}\bar{a}^2\Theta^2} \mathbb{E}_\tau \frac{e^{-\frac{\bar{b}^2}{2\bar{a}^2}}}{\bar{\mathcal{G}}_+(\tau)}, \tag{35c}$$

$$P_1 = \Delta_\Theta^{-2} [1 - \delta\bar{Q} + \bar{\kappa}^2(1 - \delta Q)], \quad \bar{P}_1 = \Delta_\Theta^{-2} [1 - \delta Q + \kappa^2(1 - \delta\bar{Q})], \tag{35d}$$

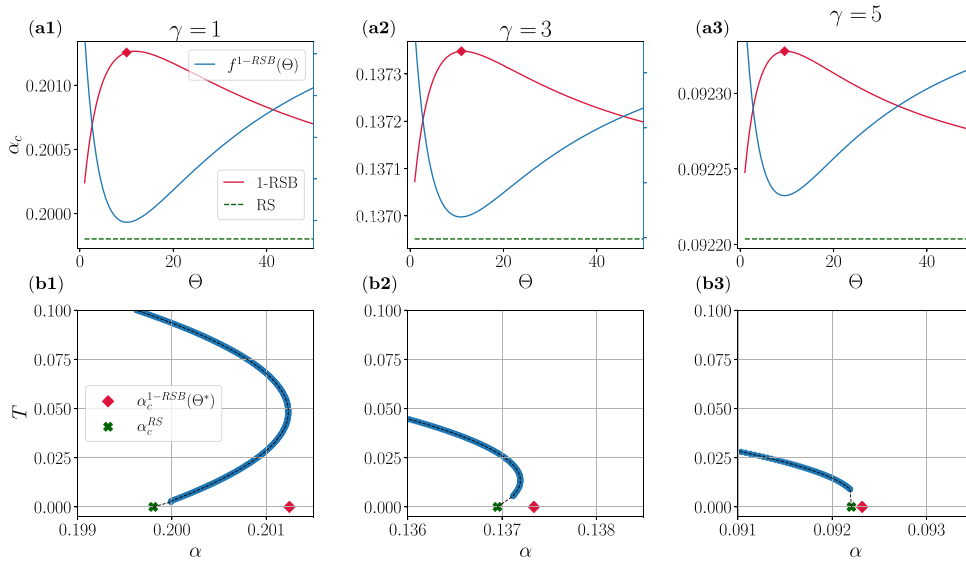


Figure 6. The upper panels show the critical load α_c estimated through the 1-RSB self-consistent equation (35) at fixed Θ and plotted as a function of the latter; we superimpose the free energy of the retrieval state at the critical point (blue lines); for comparison, dashed green lines report the RS critical capacity computed in the previous section (see figure 4.). The bottom panels show a zoom of the BAM phase diagram at low temperature close to the MR-SG transition, where the re-entrance of the critical line is manifest. Scatter points identify the RS critical capacity (green) obtained as a continuation of the RS spinodal line (blue points); red dots identify the critical capacity obtained at the optimal value Θ^* minimizing the free energy of the corresponding upper panel. Each column corresponds to one value of γ : (a1)–(b1): $\gamma = 1$; (a2)–(b2): $\gamma = 3$; (a3)–(b3): $\gamma = 5$.

$$P_2 = P_1 + (\Delta\Delta_\Theta)^{-1} (\delta\bar{Q} + \bar{\kappa}\bar{\chi}\delta Q), \quad \bar{P}_2 = \bar{P}_1 + (\Delta\Delta_\Theta)^{-1} (\delta Q + \kappa\chi\delta\bar{Q}). \quad (35e)$$

The numerical solution of the above set of equations allows us to determine the location of the critical load for a specific value of γ . In principle, an additional equation must be solved together with equation (35) regarding the stationarity of f w.r.t. the rescaled Parisi parameter Θ . However, numerical solutions involving also the latter are typically more involved from the point of view of numerical stability. We instead proceed by fixing a certain value Θ and solving equation (35); then, the optimal value of Θ is determined *a-posteriori* by looking at the minimum of the free energy (33) w.r.t. Θ . A summary set of results is shown in figure 6: the upper panels display the 1-RSB critical load computed as a function of Θ (one column correspond to a fixed value of γ), and the corresponding free energy of the Retrieval state at the critical point; the lower panels show the BAM’s phase diagram (as in figure 2) in the region close to the MR-SG spinodal, where we highlight the new position of the critical capacity as estimated through the 1-RSB equations. We first notice that 1-RSB critical capacity is always higher than the corresponding RS estimation (green dashed lines in the upper panels of figure 6), and the latter is recovered in the two limits $\Theta \rightarrow \{0, \infty\}$. However, at a difference w.r.t. what discussed in [43], the optimal (w.r.t. Θ) critical load [44] slight exceeds the maximum value attained at finite T . In any case, the difference between the 1-RSB and the RS critical loads are of order $O(10^{-3} - 10^{-4})$ so the difference between the two estimations are in any case in agreement with the simulation results discussed in the next section. As a final remark, note

that the re-entrant behavior of the phase diagram reduces as the asymmetry between the two layers increases: this is also confirmed by the optimal value Θ^* at the critical point, which numerically turns out to be a decreasing function of $\gamma \leq 1$. In other words, the temperature at which the RS instability occurs for the retrieval solution shifts to lower temperatures.

4. Numerical simulations

In this section we provide a numerical characterization of the BAM's dynamics, with the aim of numerically determining the retrieval-spin glass spinodal point. At $T = 0$, the dynamic rules for updating a single neuron belonging to one of the two layers are given by:

$$\sigma_i^{t+1} = \text{sign} \left(\sum_j w_{ij} \bar{\sigma}_j^t \right), \quad \bar{\sigma}_j^{t+1} = \text{sign} \left(\sum_i w_{ij} \sigma_i^t \right), \quad (36)$$

respectively if the chosen neuron belongs to layer 1 or layer 2. Equation (36) represents a noiseless version of the Glauber's sequential dynamics for the BAM. At finite β , the state of each neuron is sampled from a single-variable distribution, uniquely defined by a local field. In particular, for spins belonging to either layer, these probabilities are given by:

$$\sigma_i^{t+1} \sim \frac{e^{\beta h_i(\bar{\sigma}^t) \sigma_i^{t+1}}}{2 \cosh \beta h_i(\bar{\sigma}^t)}, \quad h_i(\bar{\sigma}) = \sum_j w_{ij} \bar{\sigma}_j, \quad (37)$$

$$\bar{\sigma}_j^{t+1} \sim \frac{e^{\beta h_j(\sigma^t) \bar{\sigma}_j^{t+1}}}{2 \cosh \beta h_j(\sigma^t)}, \quad h_j(\sigma) = \sum_i w_{ij} \sigma_i. \quad (38)$$

It is easy to show how the above update rules satisfy detailed balance at finite β and the corresponding equilibrium distribution is given by the Boltzmann probability with the BAM's Hamiltonian (1): indeed, it is sufficient to rewrite the original model as a generalized Hopfield with $N + \bar{N}$ neurons, having the following synaptic matrix

$$\mathbf{J} = \begin{pmatrix} \mathbf{0} & \mathbf{W} \\ \mathbf{W}^T & \mathbf{0} \end{pmatrix}, \quad (39)$$

where \mathbf{W} is the BAM's weight matrix defined in equation (2). In the following, we will focus on the retrieval of pure states (i.e. of a single pattern pair) by storing an increasing number of patterns in the network to characterize the first-order spinodal between the retrieval and the SG phases. To study the retrieval properties of the model, the initial condition on at least one of the two layers must have a high overlap with the pattern to be retrieved: otherwise, the dynamics gets stuck in spurious attractors reminiscent of the SG fixed point (which, as mentioned earlier, also exists in the retrieval phase) having zero overlap with any pattern. Therefore, the dynamics starts by choosing an initial configuration that is close enough to a pattern pair μ (in what follows we assume without loss of generality $\mu = 1$), by random flipping each neuron with probability $\varepsilon \in (0, 1/2)$ w.r.t the pattern component. As a consequence, the initial Mattis magnetization over a layer is equal (on average) to $M^{(t=0)} = 1 - 2\varepsilon$. In principle, since there are two layers in the BAM, we use ε_1 (or ε_2) to indicate the initial noise on the first (or second) layer: as will become clear below, choosing different values for the initial noise on the largest/smallest layer have strongly different outcomes if the BAM is asymmetric.

Rather than using a random sequential update scheme at each iteration of the dynamics, the bipartite structure of the BAM naturally suggests to simultaneously update the spins in one layer given the entire state of the other layer: such a 'parallel' update exploits the reverberation mechanism in which information on one layer is transferred to the other at once, as originally

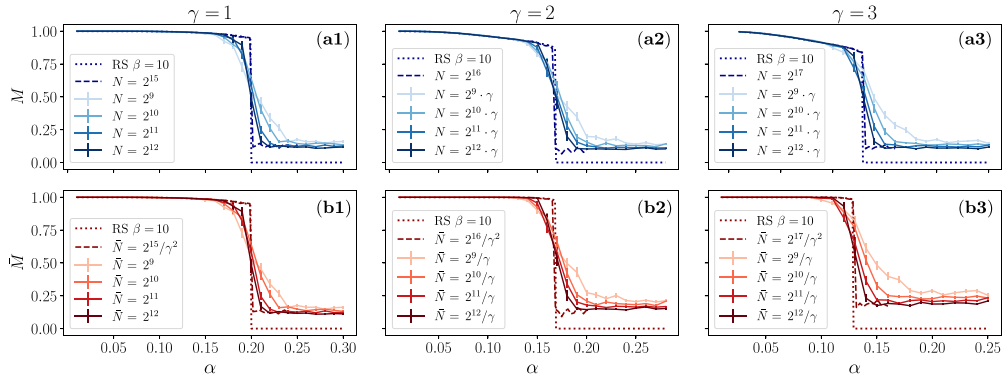


Figure 7. Numerical estimation of the retrieval/spin-glass phase transition. Simulations are performed using a parallel dynamics at $\beta = 10$ with $N_{\text{steps}} = 5 \times 10^4$ for 3 values of γ . (a1)–(b1): $\gamma = 1$; (a2)–(b2): $\gamma = 2$; (a3)–(b3): $\gamma = 3$; Each panel shows the final Mattis magnetization over each layer w.r.t. to the pattern to be retrieved. In all cases, the dynamics starts from the pattern to be retrieved with a small noise $\varepsilon_1 = 0.1$ on layer 1. Results are plotted as a function of α for four different system sizes (shown in the caption) and averaged over 100 different realizations of the disorder. Dashed lines in each panel show the results for a single instance with a very large system size. Dotted lines in each panel denote the RS saddle point solution at $\beta = 10$ (i.e. a horizontal cut in the phase diagrams of figure 2).

described in [19]: for example, note how the local fields in equations (37) and (38) acting on each neuron in one layer depend exclusively on the neuron states of the other layer, again by virtue the bipartite structure of the model. This choice also allows for a significant speedup of simulations (especially on GPU architectures), as shown by the typical sizes that can be simulated in a reasonable amount of time. Therefore, all the results presented in the current section are obtained using a simultaneous update of equation (37) for all the neurons in layer 1 and equation (38) for all the neurons in layer 2, in alternate sequence, for a total of N_{steps} . Without loss of generality, we always start the dynamics by updating layer 2, so that the only relevant initial condition is the starting configuration on layer 1 (eventually, with an initial noise ε_1). For completeness, in appendix B we also report an equivalent set of results using a random sequential update of equations (37) and (38) (using smaller system sizes), which display an identical behavior.

A summary of numerical results is shown in figure 7, considering three values of the asymmetry, i.e. $\gamma = 1, 2, 3$ (each value corresponds to one column). Each panel shows the final overlap of spin configurations calculated at the end of the simulation with respect to the original pattern (i.e. the Mattis magnetizations (7)), on layer 1 (top panels) and 2 (bottom panels), plotted as a function of α . We find an overall good agreement between the numerical simulations in predicting the location of the transition between retrieval/spin glass, compared to the RS theory (plotted in dotted lines). Moreover, the empirical transition curves become sharper as the total system size L is increased, as expected for finite sizes. Figure 7 also reports the results for a single large system size, confirming the self-averaging property of the free energy.

In figure 7 we further notice how, even beyond the predicted critical capacity, the magnetization does not fall to 0: a first explanation is that finite-size effect can result into the dynamics being stuck in spurious attractors characterized by a partial overlap with respect to more than

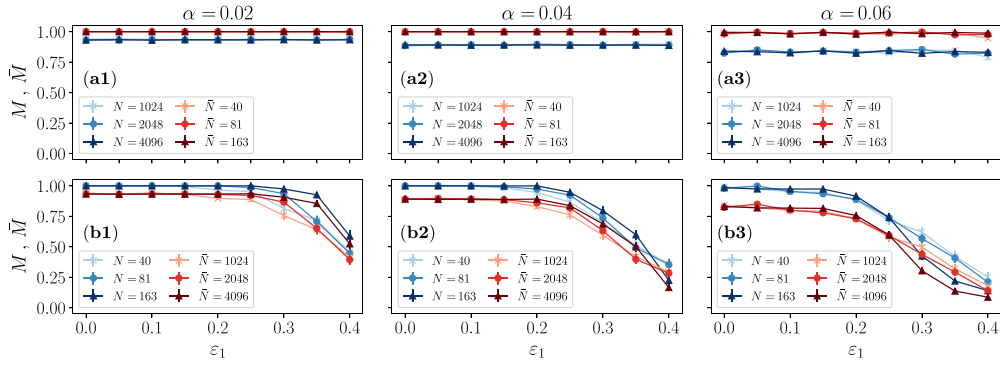


Figure 8. Basin of attraction in the asymmetric BAM. Simulations performed using a parallel dynamics at $\beta = 10$ for $N_{\text{steps}} = 10^4$. The upper row (panels (a1)–(a3)) shows results at $\gamma = 5$, so that the first layer is the largest one (i.e. $N > \bar{N}$). The bottom row (panels (b1)–(b3)) refers to the opposite scenario where the two layers are swapped, namely $\gamma = 1/5$ and $N < \bar{N}$. Each panel shows the final Mattis magnetization of the two layers M, \bar{M} (in blue/red shades respectively), plotted as functions of the initial noise on layer 1 (ε_1), for three different system sizes $L \in \{2^{10}/\gamma, 2^{11}/\gamma, 2^{12}/\gamma\}$. Each datapoint represents the average over 100 realizations of the disorder. Each column refers instead to a different value of the network load $\alpha < \alpha_c(\gamma = 5)$.

one pattern (mixed states), that are not present in the MF theory in this regime, a phenomenon already observed in [8] for the Hopfield model. Secondly, this phenomenon can be understood also in terms of the lack of patterns' orthogonality at finite size, and it would explain why in both figures 7 (a2)–(b2)–(a3)–(b3) (i.e. in the asymmetric scenarios with $\gamma = 2, 3$) the final magnetizations on the smallest layer are slightly higher than the largest layer's ones even at $\alpha > \alpha_c$.

4.1. Basin of attraction

Finally, we report a numerical experiment to quantify the different basins of attraction of the two layers in the asymmetric BAM (i.e. $\gamma \neq 1$). A summary set of results is shown in figure 8, where we performed simulations on a BAM with increasing size L while keeping fixed the asymmetry between the two layers (in all the results of figure 8, γ is equal either to 5 or 1/5). The network load α is chosen in such a way that the network is expected to be inside the retrieval phase within the MF theory, so that $\alpha < \alpha_c(\gamma = 5) \approx 0.092$. The initial noise (varying on the horizontal scale in all the panels in figure 8) is always set on the first layer (ε_1). What changes between the upper and bottom panels is that, in the former, the first layer is the largest one ($\gamma = 5$), while in the latter the situation is swapped ($\gamma = 1/5$). Therefore, when the first layer is the largest one, even a strong initial noise does not affect the retrieval capabilities of the machine: in particular, the values plotted in all the upper panels of the final Mattis magnetizations (on both layers) are consistent with the MF prediction. On the other hand, in the bottom panels the two layers are inverted, so that the dynamics is initialized with an increasing noise on the smallest layer. As a consequence, by increasing such noise the BAM is not anymore able to retrieve the pattern pairs on the two layers, as evidenced by the decreasing Mattis magnetizations in figures 8 (b1)–(b3). This is consistent with the original claim in [21] that the smallest layer has a smaller basin of attraction.

5. Understanding BAM's retrieval mechanism: analogy with coupled Hopfield models in the low-load regime

In this section, you will find a heuristic discussion to better understand how retrieval is attained in the BAM. The main differences between the Hopfield model and the BAM is that in the latter, the patterns retrieved in both layers correspond to two different sets of variables: ξ and $\bar{\xi}$. This is clear from the analytical formulation of the two models: In the Hopfield model, the Hamiltonian can be expressed as the square of a Mattis magnetization, whereas in the BAM, it is the product over the Mattis magnetization over each layer. In fact, we can observe that the Hopfield potential is a quadratic function of the Mattis magnetization: $m^2(\sigma)$, hence being stable. While, the BAM Hamiltonian resembles a saddle given by the product of the Mattis magnetization between the two different layers: $m(\sigma)\bar{m}(\bar{\sigma})$: the main difference is that, in the BAM, we cannot retrieve a certain pattern on both layers if the two magnetizations have opposite signs. This structure naturally suggests that we can interpolate between the two situations to construct some sort of generalization of the Hopfield model with tunable intra/inter synaptic connections, or to put it differently, to move from a quadratic potential to a one with unstable directions.

Consequently, the BAM can be viewed as a generalized Hopfield model over the union set of the neurons of the two layers $\sigma \cup \bar{\sigma}$, where only the intersynaptic connections between the two subsets are maintained, while all intrasynaptic connections are set to 0 (see equation (39)). This mechanism allows for the BAM to retrieve only associated information on each layer: namely, if the pattern ξ^μ is recalled on layer 1, layer 2 will recall its corresponding 'dual' pattern $\bar{\xi}^\mu$. In other words, the BAM is able to retrieve a pattern represented by $\xi^\mu \cup \bar{\xi}^\mu$ (for the set $\sigma \cup \bar{\sigma}$ of neurons), with a reduced number of interaction weights, $N \times \bar{N}$: a number to be compared with the $\sim(N+\bar{N})^2$ weights that would be needed in the equivalent Hopfield case.

For the following analysis, we consider a generalization of the starting Hamiltonian (1), introducing an additional parameter τ :

$$H_\tau(\sigma, \bar{\sigma}) = -\sum_{\mu=1}^K \left[\frac{(1-\tau)}{2} \left(\frac{1}{\sqrt{N}} \sum_i \xi_i^\mu \sigma_i \right)^2 + \frac{(1-\tau)}{2} \left(\frac{1}{\sqrt{\bar{N}}} \sum_j \bar{\xi}_j^\mu \bar{\sigma}_j \right)^2 + \tau \left(\frac{1}{\sqrt{N}} \sum_i \xi_i^\mu \sigma_i \right) \left(\frac{1}{\sqrt{\bar{N}}} \sum_j \bar{\xi}_j^\mu \bar{\sigma}_j \right) \right]. \tag{40}$$

Note that a similar kind of interpolation has been employed to overcome the obstacle of the non convexity of the interaction for the free energy of layered SG models [45, 46]. Clearly, equation (40) recovers the original BAM Hamiltonian when $\tau = 1$. On the other hand, when $\tau = 0$, the generalized Hamiltonian describes two independent Hopfield models with N and \bar{N} units, respectively, each of which has its own set of stored patterns in its synaptic matrix (which is fully connected within each layer), i.e.

$$H_{\tau=0}(\sigma, \bar{\sigma}) = -\frac{1}{2} \sum_{\mu=1}^K \left(\frac{1}{\sqrt{N}} \sum_i \xi_i^\mu \sigma_i \right)^2 - \frac{1}{2} \sum_{\mu=1}^K \left(\frac{1}{\sqrt{\bar{N}}} \sum_j \bar{\xi}_j^\mu \bar{\sigma}_j \right)^2 = -\frac{N}{2} \sum_{\mu=1}^K (m^\mu(\sigma))^2 - \frac{\bar{N}}{2} \sum_{\mu=1}^K (\bar{m}^\mu(\bar{\sigma}))^2. \tag{41}$$

Indeed, on each layer one can define a synaptic matrix constructed via the usual Hebb's rule, namely,

$$J_{i_1, i_2} = \frac{1}{N} \sum_{\mu=1}^K \xi_{i_1}^{\mu} \xi_{i_2}^{\mu}, \quad \bar{J}_{j_1 j_2} = \frac{1}{N} \sum_{\mu=1}^K \bar{\xi}_{j_1}^{\mu} \bar{\xi}_{j_2}^{\mu}, \quad (42)$$

as it can be easily verified by expanding the squares in equation (41). For a generic value of $\tau \in (0, 1)$, equation (40) describes a mixture of two Hopfield models with both intra-synaptic and inter-synaptic connections whose relative strength is tuned by the value of τ . The main difference between the two limiting cases is that when $\tau = 0$, the partition function of the model defined in equation (40), factorizes completely over the two spin sets, which means that each Hopfield model by itself can retrieve any pattern, regardless on the other. On the other hand, when $\tau = 1$ only retrieval of corresponding couples of patterns with the same index μ is possible. For this reason, the question naturally arises under what conditions the retrieval of individual information is possible as a function of τ . The following analysis is performed in the regime of low load, i.e. when the number of patterns is K is finite relative to the number of spins on each layer, using the standard procedure used for the Hopfield model (e.g. as discussed in detail in [38]). For simplicity, we consider all patterns to be binary vectors with i.i.d. components drawn with equal probability between $\{-1, 1\}$. Note, however, that retrieval in the low-load scenario should be robust w.r.t. the distribution chosen for the pattern components by generalizing recent arguments about the Hopfield model (see, e.g. [9]) In the low-load regime, the noise induced by the patterns is subextensive relative to the system size, so it is not necessary to introduce replicas or implement Guerra's interpolation technique to address this. The average quenched free energy can be easily calculated by a saddle point method after introducing the usual set of Mattis magnetizations on each layer. After some calculations, the density of the free energy corresponding to the equation (40) model can be written as

$$\begin{aligned} f(\beta, \gamma, \tau) = & + \frac{1}{2} \gamma (1 - \tau) \sum_{\mu} m_{\mu}^2 - \frac{\gamma}{\beta} \mathbb{E}_{\xi} \log \cosh \left\{ \beta \sum_{\mu} [(1 - \tau) m_{\mu} + \tau \bar{\gamma} \bar{m}_{\mu}] \xi^{\mu} \right\} \\ & + \frac{1}{2} \bar{\gamma} (1 - \tau) \sum_{\mu} \bar{m}_{\mu}^2 - \frac{\bar{\gamma}}{\beta} \mathbb{E}_{\bar{\xi}} \log \cosh \left\{ \beta \sum_{\mu} [(1 - \tau) \bar{m}_{\mu} + \tau \gamma m_{\mu}] \bar{\xi}^{\mu} \right\} \\ & + \tau \sum_{\mu} m_{\mu} \bar{m}_{\mu}, \end{aligned} \quad (43)$$

where m_{μ}, \bar{m}_{μ} are defined as in equation (16). It is easy to check how the above expression recovers the BAM's free energy in the low load limit—i.e. (20) at $\alpha = 0$ —when $\tau = 1$, and the sum of two Hopfield free energies (again, in the low-load limit) - each rescaled by the shape parameter (resp. γ for layer 1 and $\bar{\gamma}$ for layer 2) - when $\tau = 0$. The values of $\{m_{\mu}, \bar{m}_{\mu}\}_{\mu=1}^K$ can be found by imposing stationarity of (43), which determines the following self-consistent equations:

$$m_{\mu} = \mathbb{E}_{\xi} \xi^{\mu} \tanh \left\{ \beta \sum_{\nu} [(1 - \tau) m_{\nu} + \tau \bar{\gamma} \bar{m}_{\nu}] \xi^{\nu} \right\}, \quad (44)$$

$$\bar{m}_{\mu} = \mathbb{E}_{\bar{\xi}} \bar{\xi}^{\mu} \tanh \left\{ \beta \sum_{\nu} [(1 - \tau) \bar{m}_{\nu} + \tau \gamma m_{\nu}] \bar{\xi}^{\nu} \right\}. \quad (45)$$

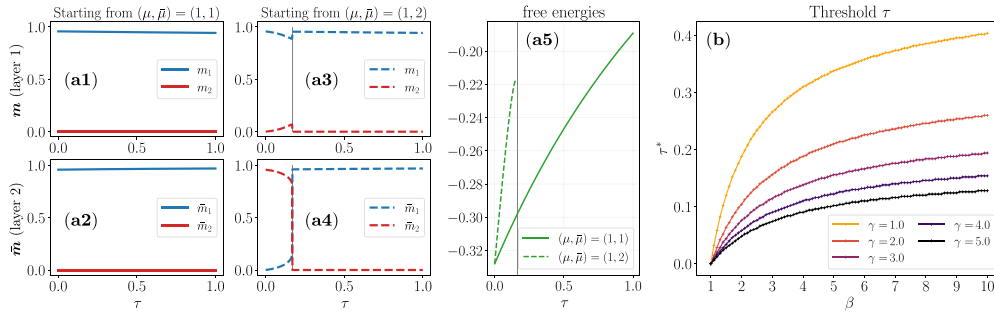


Figure 9. Panels labeled with letter (a) show results on the saddle point equations (44) and (45) obtained varying τ for the retrieval of the same/different pattern index on the two layers. (a1) and (a2): Mattis magnetizations on 2 pattern components (shown resp. in blue and red solid lines and labeled with subscripts 1, 2) over layer 1 (a1) and layer 2 (a2), obtained by initializing the saddle point equations with an initial high magnetization over pattern 1 on both layers, i.e. $\mu, \bar{\mu} = 1$. (a3) and (a4) (again showing the Mattis magnetizations on layer 1 and 2 respectively) represent the same results but starting from a high magnetization on different patterns for the two layers: in particular, starting from an initial high overlap with pattern 1 on the first layer and pattern 2 on the second. Panel (a5) shows the free energies of both solutions: the solid green line refers to the saddle points obtained in (a1) and (a2), while the dashed line refers to (a3) and (a4). All the results of (a)-panels are obtained at $\beta = 2$ and $\gamma = 1.1$, and plotted as a function of τ . (b): Threshold value of τ after which retrieval of different pattern indexes in each of the layers is not possible, plotted as a function of $\beta > 1$ for different values of γ (shown in the legend).

The equilibrium behavior of the model can again be studied by numerically solving the above equations. The simplest nontrivial scenario arises with $K = 2$, so that the magnetization vectors have at most two nonzero components for each layer, parameterized as follows:

$$\mathbf{m} = (m_1, m_2)^T \quad \text{and} \quad \bar{\mathbf{m}} = (\bar{m}_1, \bar{m}_2)^T. \quad (46)$$

It is easy to check, both numerically and analytically, that independently on the value of τ the only solution at $\beta < 1$ is the paramagnetic one, where all the magnetizations are null. At $\beta = 1$, a ferromagnetic solution appears, independently on the value of τ . It is perhaps more interesting to look at the numerical solutions by focusing on a certain value of $\beta > 1$ (and γ) and analyze its behavior as a function of τ : an example set of results is shown in figure 9, panels (a1)→(a5), obtained at $\beta = 2$ and $\gamma = 1.1$ (although the qualitative behavior remains unchanged by lowering β or changing γ). The four left panels (a1)→(a4) show the two components of the Mattis magnetization for the two layers (respectively in the upper and lower panel), obtained by solving equations (44) and (45) at different values of τ (varying on the horizontal axis), and with respect to the two different patterns (with indexes $\mu, \bar{\mu} = 1, 2$). The left panels (a1) and (a2) refer to fixed points obtained by starting from a high magnetizations on the *same* pattern index on both layers, i.e. $\mathbf{m}^0 = (1 - \varepsilon, \varepsilon)^T$ and $\bar{\mathbf{m}}^0 = (1 - \bar{\varepsilon}, \bar{\varepsilon})^T$, with $\varepsilon, \bar{\varepsilon}$ being arbitrary small. Such a ferromagnetic solution is always found at any value of τ , and the corresponding free energy is plotted in the panel (a5) (full green line). We conclude that the retrieval of the same pattern is a stable fixed point in the full range $\tau \in [0, 1]$. Panels (a3) and (a4) show the same quantities, but this time the initial magnetization vector has a high value on a *different* pattern over each layer, e.g. $\mathbf{m}^0 = (1 - \varepsilon, \varepsilon)^T$ and $\bar{\mathbf{m}}^0 = (\bar{\varepsilon}, 1 - \bar{\varepsilon})^T$. As clear from the figure, the retrieval of different patterns (in each of the layers) can be attained only at small values of τ , i.e. when the two Hopfield models are almost independent and the strength

of their inter-synaptic connections falls below a certain (temperature-dependent) threshold. Above such threshold—called τ^* in the following—one of the two layers (in particular, the smallest one) switches the components of its magnetization vector so to be aligned with the other layer. The free energy corresponding to these solutions is shown again in panel (a5) (dashed green line): it should be noticed how, at $\tau < \tau^*$, these solutions are thermodynamically unfavored, i.e. their free energy is higher than the solution corresponding to the same pattern index (except in the trivial case at $\tau = 0$ where the two Hopfield models are independent), so in this sense they are metastable. To summarize, while varying τ there is first a (trivial) phase co-existence at $\tau = 0$; then, the recall of two distinct patterns is locally stable, but unfavored, until the state completely disappear at τ^* . Such threshold values are computed numerically and plotted figure 9(b) as functions of $\beta > 1$ for different values of γ . The value of τ^* turns out to increase upon lowering the temperature for a fixed γ and viceversa at fixed β when γ departs from 1.

6. Conclusions

In this paper, we have presented a comprehensive and exhaustive statistical mechanics treatment of the BAM. We have characterized the phase diagram of the BAM, and in particular its operating regimes, i.e. the existence and extent of its retrieval phase, by both an analytic treatment in the asymptotic limit and numerical simulations, also by overcoming limitations of previous works. We have paid particular attention to the effects of asymmetry between the sizes of the layers of the bipartite network and concluded that the asymmetry tends to damage the retrieval capacities. We have also shown that the BAM can be used as an efficient alternative to the Hopfield model from the point of view of the storage and retrieval properties of recurrent neural networks, in terms of the number of weights required to store a fixed set of patterns, and the efficiency of the reverberation mechanism explained in section 4 using parallel dynamics, the latter allowing significant speedup of sampling mechanisms (particularly efficient on GPUs), compared to the standard Hopfield model. However, the price to be paid is, on the one hand, a lower network capacity compared to the (fully connected) Hopfield model and, on the other hand, a smaller basin of attraction for the smallest layer in the asymmetric BAM. In principle, this phenomenon could be further amplified by adding more and more layers: the resulting network might retrieve an overall smaller number of patterns than in the Hopfield case, but the efficiency of storage would be improved by the structure of the network. We leave this topic for future investigations.

Additionally, this work lays the foundation for the use of BAM as an energy-based generative model in unsupervised learning. An earlier attempt to use the retrieval properties of the Hopfield model in training RBMs was recently proposed [18], but such a training process is likely to introduce correlations between patterns when using real datasets. This is a situation we have not considered in this work and should be addressed in order to properly characterize the model's phase diagram during learning. Although, for instance, the Hopfield model with correlated patterns has already been analyzed in the simple low-load regime [47], we still lack a general theory that can be used in a context of learning.

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgment

GC, AD and BS acknowledge financial support by the Comunidad de Madrid and the Complutense University of Madrid (Spain) through the Atracción de Talento programs (references 2019-T1/TIC-13298 for AD and GC and 2019-T1/TIC-12776 for BS), the Banco Santander and the UCM (Grant PR44/21-29937), and Ministerio de Economía y Competitividad, Agencia Estatal de Investigación and Fondo Europeo de Desarrollo Regional (FEDER) (Spain and European Union) through the Grant PID2021-125506NA-I00. BS acknowledges the support provided by the Nationale de la Recherche (France) under grant number ANR CPJ-ESE 2022. AB acknowledges financial support by Ministero degli Affari Esteri e della Cooperazione Internazionale (Italy-Israel collaboration, BULBUL Grant, CUP Project No. F85F21006230001). GC acknowledges useful discussions with Linda Albanese and Alberto Fachechi. All the authors acknowledge hospitality during the Alan Turing Institute’s Theory and Methods Challenge Fortnights event ‘Physics-informed machine learning’.

Appendix A. Interpolation method

We now illustrate how to compute the BAM’s quenched free energy using the interpolation technique developed by Guerra [23, 48]. We proceed by analogy with previous related works on SG models (see e.g. [49]), by defining the following quenched intensive pressure at fixed size L

$$\mathcal{A}_L(\beta, \alpha, \gamma) = \frac{1}{L} \mathbb{E}_{\xi, \bar{\xi}} \log \sum_{\sigma, \bar{\sigma}} \exp \left[\frac{\beta}{L} \sum_{\mu=1}^K \left(\sum_{i=1}^N \xi_i^\mu \sigma_i \right) \left(\sum_{j=1}^{\bar{N}} \bar{\xi}_j^\mu \bar{\sigma}_j \right) \right], \quad (\text{A1})$$

and we are interested in evaluating its thermodynamic limit, namely

$$\mathcal{A}(\beta, \alpha, \gamma) = \lim_{L \rightarrow \infty} \mathcal{A}_L(\beta, \alpha, \gamma). \quad (\text{A2})$$

The above quantity is related to the usual Helmholtz free energy density through $f = -\mathcal{A}/\beta$. Considering the choice of pattern components for the signal and noise term discussed in section 3.1 of the main text, the operator $\mathbb{E}_{\xi, \bar{\xi}}$ reads

$$\mathbb{E}_{\xi, \bar{\xi}} = \prod_{i=1}^N \prod_{j=1}^{\bar{N}} \left[\prod_{\mu=1}^l \mathbb{E}_{\xi_i^\mu \in \{-1,1\}} \mathbb{E}_{\bar{\xi}_j^\mu \in \{-1,1\}} \right] \left[\prod_{\mu=l+1}^K \mathbb{E}_{\xi_i^\mu \sim \mathcal{N}(0,1)} \mathbb{E}_{\bar{\xi}_j^\mu \sim \mathcal{N}(0,1)} \right]. \quad (\text{A3})$$

The interpolation method developed by Guerra allows to compute equation (A1) in the thermodynamic limit $L \rightarrow \infty$ when the extensive set of disordered quantities are Gaussian i.i.d. random variables. The interpolating structure of the quenched free energy is constructed in terms of an interpolating real parameter $t \in [0, 1]$, whose expression is shown below [50]:

$$\begin{aligned}
 \mathcal{A}_L(\beta, \alpha, \gamma, t) = & \frac{1}{L} \mathbb{E} \log \sum_{\sigma, \bar{\sigma}} \exp \left\{ t \frac{\beta}{L} \sum_{\mu=1}^l \left(\sum_i \xi_i^\mu \sigma_i \right) \left(\sum_j \bar{\xi}_j^\mu \bar{\sigma}_j \right) \right. \\
 & \left. + (1-t) \beta \sum_{\mu=1}^l \left[\Psi_\mu \sum_i \xi_i^\mu \sigma_i + \bar{\Psi}_\mu \sum_j \bar{\xi}_j^\mu \bar{\sigma}_j \right] \right\} \\
 & \times \int \prod_{\mu} dz_{\mu} z_{\mu}^{\dagger} \exp \left\{ -\beta \sum_{\mu} z_{\mu} z_{\mu}^{\dagger} + \beta \sqrt{\frac{t}{N}} \sum_{i, \mu} z_{\mu} \xi_i^\mu \sigma_i + \beta \sqrt{\frac{t}{N}} \sum_{j, \mu} z_{\mu}^{\dagger} \bar{\xi}_j^\mu \bar{\sigma}_j \right. \\
 & \left. + \sqrt{1-t} \left(c \sum_i \eta_i \sigma_i + \bar{c} \sum_j \bar{\eta}_j \bar{\sigma}_j \right) + \sqrt{1-t} \left(d \sum_{\mu} y_{\mu} z_{\mu} + \bar{d} \sum_{\mu} \bar{y}_{\mu} z_{\mu}^{\dagger} \right) \right. \\
 & \left. + \frac{(1-t)}{2} \sum_{\mu} (z_{\mu}, z_{\mu}^{\dagger}) \mathbf{E} (z_{\mu}, z_{\mu}^{\dagger})^T \right\}, \tag{A4}
 \end{aligned}$$

where we have already split the signal and noise term in the starting Hamiltonian, as discussed in the main text (see section 3.1). In the above expression, an additional set of i.i.d. standard Gaussian variables has been introduced (whose average is included in the expectation \mathbb{E}), namely $\{\eta_i\}_{i=1}^N, \{\bar{\eta}_j\}_{j=1}^{\bar{N}}, \{y_{\mu}, \bar{y}_{\mu}\}_{\mu=1}^K$. The scalars $\{c, \bar{c}, d, \bar{d}\}$ and the entries of the 2×2 matrix \mathbf{E} are constant real values yet to be fixed. It is trivial to check that the original intensive pressure equation (A1) is recovered when $t = 1$. On the other hand, at $t = 0$ the interacting term in the original Hamiltonian disappears, and the corresponding intensive pressure is that of a factorized model over the spin components (on both layers): the only difference here with the standard Hopfield model is that the $t = 0$ case contains a series of 2-body problems in the pairs $(z_{\mu}, z_{\mu}^{\dagger})$, which is still easy to handle. Notice that the two terms on the first line of equation (A4) - accounting for the signal (retrieval) part of the free energy - are interpolated linearly with t , following the standard procedure for the Curie-Weiss model [51]. All the others are interpolated using the square root function as usually done for SGs. The key point behind such a method is that the quenched pressure of the BAM, obtained by setting $t = 1$, can be obtained exploiting the calculus theorem:

$$\mathcal{A}_L(\beta, \alpha, \gamma, t = 1) = \mathcal{A}_L(\beta, \alpha, \gamma, t = 0) + \int_0^1 dt' \frac{d\mathcal{A}_L(\beta, \alpha, \gamma, t')}{dt'}. \tag{A5}$$

Therefore, we just need to compute the initial (Cauchy) condition at $t = 0$ and the t -derivative of the interpolating quenched pressure.

A.1. Derivative

We now evaluate the t -derivative of equation (A4). For simplicity, let us call the interpolated exponent of (A4) as $-H(t)$: in the following, we will indicate with the letter $\omega(\mathcal{O})$ the expectation value of a certain observable \mathcal{O} w.r.t. the Boltzmann measure for a fixed realization of the quenched disorder:

$$\omega(\mathcal{O}) = \frac{\sum_{\sigma, \bar{\sigma}} \int \prod_{\mu} dz_{\mu} z_{\mu}^{\dagger} \mathcal{O} \exp[-H(t)]}{\sum_{\sigma, \bar{\sigma}} \int \prod_{\mu} dz_{\mu} z_{\mu}^{\dagger} \exp[-H(t)]}, \tag{A6}$$

and with $\langle \mathcal{O} \rangle = \mathbb{E}\omega(\mathcal{O})$: from now on, the operator \mathbb{E} denotes the expectation value over all the possible disorder components appearing in the quenched pressure. The total t -derivative of (A4) reads:

$$\begin{aligned} \frac{d\mathcal{A}}{dt} = & + \frac{\beta}{L^2} \sum_{\mu=1}^l \mathbb{E}\omega \left[\left(\sum_i \xi_i^\mu \sigma_i \right) \left(\sum_j \bar{\xi}_j^\mu \bar{\sigma}_j \right) \right] \\ & - \frac{\beta}{L} \sum_{\mu=1}^l \mathbb{E}\omega \left[\Psi_\mu \sum_i \xi_i^\mu \sigma_i + \bar{\Psi}_\mu \sum_j \bar{\xi}_j^\mu \bar{\sigma}_j \right] \\ & + \frac{\beta}{2L\sqrt{N}\sqrt{t}} \sum_{i,\mu} \mathbb{E}\omega(z_\mu \xi_i^\mu \sigma_i) + \frac{\beta}{2L\sqrt{N}\sqrt{t}} \sum_{j,\mu} \mathbb{E}\omega(z_\mu^\dagger \bar{\xi}_j^\mu \bar{\sigma}_j) \\ & - \frac{c}{2L\sqrt{1-t}} \sum_i \mathbb{E}\omega(\eta_i \sigma_i) - \frac{\bar{c}}{2L\sqrt{1-t}} \sum_j \mathbb{E}\omega(\bar{\eta}_j \bar{\sigma}_j) \\ & - \frac{d}{2L\sqrt{1-t}} \sum_\mu \mathbb{E}\omega(y_\mu z_\mu) - \frac{\bar{d}}{2L\sqrt{1-t}} \sum_\mu^K \mathbb{E}\omega(\bar{y}_\mu z_\mu^\dagger) \\ & - \frac{1}{2L} \sum_\mu^K \mathbb{E}\omega \left[(z_\mu, z_\mu^\dagger) \mathbf{E} (z_\mu, z_\mu^\dagger)^T \right]. \end{aligned} \tag{A7}$$

All the terms explicitly depending on Gaussian quenched disordered can be further simplified by exploiting Wick's Theorem:

$$\mathbb{E}[xf(x)] = \mathbb{E}[\partial_x f(x)], \tag{A8}$$

where $x \sim \mathcal{N}(0, 1)$ is a standard normal random variable. After applying equation (A8) to the above derivative, we get:

$$\begin{aligned} \frac{d\mathcal{A}}{dt} = & \frac{\beta}{L^2} \sum_{\mu=1}^l \mathbb{E}\omega \left[\left(\sum_i \xi_i^\mu \sigma_i \right) \left(\sum_j \bar{\xi}_j^\mu \bar{\sigma}_j \right) \right] \\ & - \frac{\beta}{L} \sum_{\mu=1}^l \mathbb{E}\omega \left[\Psi_\mu \sum_i \xi_i^\mu \sigma_i + \bar{\Psi}_\mu \sum_j \bar{\xi}_j^\mu \bar{\sigma}_j \right] \\ & + \frac{\beta^2}{2LN} \sum_{i,\mu} \mathbb{E}[\omega(z_\mu^2) - \omega^2(z_\mu \sigma_i)] + \frac{\beta^2}{2LN} \sum_{j,\mu} \mathbb{E}[\omega(z_\mu^{\dagger 2}) - \omega^2(z_\mu^\dagger \bar{\sigma}_j)] \\ & - \frac{c^2}{2L} \sum_i \mathbb{E}[1 - \omega^2(\sigma_i)] - \frac{\bar{c}^2}{2L} \sum_j \mathbb{E}[1 - \omega^2(\bar{\sigma}_j)] \\ & - \frac{d^2}{2L} \sum_\mu \mathbb{E}[\omega(z_\mu^2) - \omega^2(z_\mu)] - \frac{\bar{d}^2}{2L} \sum_\mu \mathbb{E}[\omega(z_\mu^{\dagger 2}) - \omega^2(z_\mu^\dagger)] \\ & - \frac{1}{2L} \sum_\mu^K \mathbb{E}\omega \left[(z_\mu, z_\mu^\dagger) \mathbf{E} (z_\mu, z_\mu^\dagger)^T \right]. \end{aligned} \tag{A9}$$

Now, introducing the order parameters using the definitions given in the main text, namely equations (16)–(18) we can further simplify the above expression as follows:

$$\frac{dA}{dt} = \beta \sum_{\mu=1}^l [\langle m_\mu \bar{m}_\mu \rangle - \gamma \Psi_\mu \langle m_\mu \rangle - \bar{\gamma} \bar{\Psi}_\mu \langle \bar{m}_\mu \rangle] + \tag{A10}$$

$$- \frac{\alpha \beta^2}{2} \langle q_{12} p_{12} \rangle - \frac{\alpha \beta^2}{2} \langle \bar{q}_{12} \bar{p}_{12} \rangle - \frac{\gamma c^2}{2} (1 - \langle q_{12} \rangle) - \frac{\bar{\gamma} \bar{c}^2}{2} (1 - \langle \bar{q}_{12} \rangle) \tag{A11}$$

$$- \frac{\alpha d^2}{2} \langle p_{12} \rangle - \frac{\alpha \bar{d}^2}{2} \langle \bar{p}_{12} \rangle - \frac{1}{2L} \sum_{\mu=l+1}^K \langle (z_\mu, z_\mu^\dagger) \tilde{E} (z_\mu, z_\mu^\dagger)^T \rangle, \tag{A12}$$

where

$$\tilde{E} = E - \begin{bmatrix} \beta^2 - d^2 & 0 \\ 0 & \beta^2 - \bar{d}^2 \end{bmatrix}. \tag{A13}$$

The main idea now is to fix the scalar parameters $c, \bar{c}, d, \bar{d}, \{E_{ij}\}$ in such a way to explicitly show the fluctuations of the order parameters. First notice that the above expression can be rewritten as

$$\begin{aligned} \frac{dA}{dt} = & +\beta \sum_{\mu=1}^l [\langle m_\mu \bar{m}_\mu \rangle - \gamma \Psi_\mu \langle m_\mu \rangle - \bar{\gamma} \bar{\Psi}_\mu \langle \bar{m}_\mu \rangle] \\ & - \frac{\alpha \beta^2}{2} \left[\langle q_{12} p_{12} \rangle - \frac{\gamma c^2}{\alpha \beta^2} \langle q_{12} \rangle - \frac{d^2}{\beta^2} \langle p_{12} \rangle \right] - \frac{\gamma c^2}{2} \\ & - \frac{\alpha \beta^2}{2} \left[\langle \bar{q}_{12} \bar{p}_{12} \rangle - \frac{\bar{\gamma} \bar{c}^2}{\alpha \beta^2} \langle \bar{q}_{12} \rangle - \frac{\bar{d}^2}{\beta^2} \langle \bar{p}_{12} \rangle \right] - \frac{\bar{\gamma} \bar{c}^2}{2} \\ & - \frac{1}{2L} \sum_{\mu} \langle (z_\mu, z_\mu^\dagger) \tilde{E} (z_\mu, z_\mu^\dagger)^T \rangle. \end{aligned} \tag{A14}$$

Finally, by fixing the scalar parameters as follows

$$\begin{aligned} \Psi_\mu &= \bar{\gamma} \langle \bar{m}_\mu \rangle & \bar{\Psi}_\mu &= \gamma \langle m_\mu \rangle \\ c &= \beta \sqrt{\bar{\gamma} \alpha \langle p_{12} \rangle} & \bar{c} &= \beta \sqrt{\gamma \alpha \langle \bar{p}_{12} \rangle} \\ d &= \beta \sqrt{\langle q_{12} \rangle} & \bar{d} &= \beta \sqrt{\langle \bar{q}_{12} \rangle} \\ E_{11} &= \beta^2 - d^2 & E_{22} &= \beta^2 - \bar{d}^2 \end{aligned} \tag{A15}$$

and $E_{12} = E_{21} = 0$ we can simplify the total stream as:

$$\begin{aligned} \frac{dA}{dt} = & \beta \sum_{\mu=1}^l \langle (m_\mu - \langle m_\mu \rangle) (\bar{m}_\mu - \langle \bar{m}_\mu \rangle) \rangle - \beta \sum_{\mu=1}^l \langle m_\mu \bar{m}_\mu \rangle \\ & - \frac{\alpha \beta^2}{2} \langle (q_{12} - \langle q_{12} \rangle) (p_{12} - \langle p_{12} \rangle) \rangle - \frac{\alpha \beta^2}{2} \langle p_{12} \rangle (1 - \langle q_{12} \rangle) \\ & - \frac{\alpha \beta^2}{2} \langle (\bar{q}_{12} - \langle \bar{q}_{12} \rangle) (\bar{p}_{12} - \langle \bar{p}_{12} \rangle) \rangle - \frac{\alpha \beta^2}{2} \langle \bar{p}_{12} \rangle (1 - \langle \bar{q}_{12} \rangle). \end{aligned} \tag{A16}$$

A.1.1. RS solution. Under the RS approximation, all the order parameters concentrate around a fixed value (to be determined self-consistently at any α, β, γ) in the thermodynamic limit and do not fluctuate, namely

$$\begin{aligned} m_\mu &\xrightarrow{\text{RS}} M_\mu & \langle q_{12} \rangle &\xrightarrow{\text{RS}} Q & \langle p_{12} \rangle &\xrightarrow{\text{RS}} P \\ \bar{m}_\mu &\xrightarrow{\text{RS}} \bar{M}_\mu & \langle \bar{q}_{12} \rangle &\xrightarrow{\text{RS}} \bar{Q} & \langle \bar{p}_{12} \rangle &\xrightarrow{\text{RS}} \bar{P} \end{aligned}$$

With this assumption, all the fluctuation terms in (A16) disappear, and the final expression of the time derivative for the interpolating quenched pressure simplifies to

$$\frac{d\mathcal{A}}{dt} = -\beta \sum_{\mu=1}^l M_\mu \bar{M}_\mu - \frac{\alpha\beta^2}{2} P(1-Q) - \frac{\alpha\beta^2}{2} \bar{P}(1-\bar{Q}). \tag{A17}$$

A.2. Initial condition

The last missing piece to evaluate (A16) is the initial (Cauchy) condition at $t=0$. As already discussed, such a computation is trivial as the interacting part disappears and we are left with a series of one-body terms w.r.t. to all the spin variables, plus set of 2-body problems in the pairs (z_μ, z_μ^\dagger) . The quenched pressure evaluated at $t=0$ reads:

$$\begin{aligned} \mathcal{A}_L(\beta, \alpha, \gamma, t=0) &= \frac{1}{L} \mathbb{E} \log \sum_{\sigma, \bar{\sigma}} \exp \left[\beta \sum_{\mu=1}^l \left(\Psi_\mu \sum_i \xi_i^\mu \sigma_i + \bar{\Psi}_\mu \sum_j \bar{\xi}_j^\mu \bar{\sigma}_j \right) \right. \\ &\quad \left. + c \sum_i \eta_i \sigma_i + \bar{c} \sum_j \bar{\eta}_j \bar{\sigma}_j \right] \\ &\quad \times \prod_\mu \int dz_\mu z_\mu^\dagger \exp \left[-\beta z_\mu z_\mu^\dagger + dy_\mu z_\mu + \bar{d}\bar{y}_\mu z_\mu^\dagger + \frac{1}{2} (z_\mu, z_\mu^\dagger) \mathbf{E} (z_\mu, z_\mu^\dagger)^T \right]. \end{aligned} \tag{A18}$$

The traces over $\sigma, \bar{\sigma}$ can be easily computed since everything is factorized. Moreover, each integral in the second line (they are all equal) can be transformed into a Gaussian integral over a pair of real variables u_μ, v_μ by exploiting $z_\mu = u_\mu + iv_\mu$ and $z_\mu^\dagger = u_\mu - iv_\mu$. After some calculations, we get:

$$\begin{aligned} \mathcal{A}_L(\beta, \alpha, \gamma, t=0) &= \gamma \mathbb{E}_{\eta, \xi} \log 2 \cosh \left[c\eta + \beta \sum_{\mu=1}^l \Psi_\mu \xi^\mu \right] \\ &\quad + \bar{\gamma} \mathbb{E}_{\bar{\eta}, \bar{\xi}} \log 2 \cosh \left[\bar{c}\bar{\eta} + \beta \sum_{\mu=1}^l \bar{\Psi}_\mu \bar{\xi}^\mu \right] + -\frac{\alpha}{2} \log \det \mathbf{K} \\ &\quad + \frac{\alpha}{2} \mathbb{E}_{y, \bar{y}} \left[(dy + \bar{d}\bar{y}, idy - i\bar{d}\bar{y}) \mathbf{K}^{-1} (dy + \bar{d}\bar{y}, idy - i\bar{d}\bar{y})^T \right], \end{aligned} \tag{A19}$$

where the matrix \mathbf{K} is given by

$$\mathbf{K} = \begin{bmatrix} 2\beta - (E_{11} + E_{22}) & i(E_{22} - E_{11}) \\ i(E_{22} - E_{11}) & 2\beta + (E_{11} + E_{22}) \end{bmatrix}. \tag{A20}$$

The last term in the second line of equation (A20) can be further simplified by explicitly computing the average over y, \bar{y} , being two i.i.d standard Gaussian variables. After substituting all the scalar parameters as fixed by (A15) within the RS ansatz, the final expression for the initial condition on the quenched intensive pressure reads:

$$\begin{aligned} \mathcal{A}_L(\beta, \alpha, \gamma, t=0) &= \gamma \mathbb{E}_{\eta, \xi} \log 2 \cosh \left[\beta \sqrt{\gamma \alpha P} \eta + \beta \gamma \sum_{\mu} \bar{M}_{\mu} \xi^{\mu} \right] \\ &+ \bar{\gamma} \mathbb{E}_{\bar{\eta}, \bar{\xi}} \log 2 \cosh \left[\beta \sqrt{\gamma \alpha \bar{P}} \bar{\eta} + \beta \gamma \sum_{\mu} M_{\mu} \bar{\xi}^{\mu} \right] + \\ &- \frac{\alpha}{2} \log [1 - \beta^2 (1 - Q)(1 - \bar{Q})] + \frac{\alpha \beta^2}{2} \frac{[Q(1 - \bar{Q}) + \bar{Q}(1 - Q)]}{1 - \beta^2 (1 - Q)(1 - \bar{Q})}. \end{aligned} \tag{A21}$$

Finally, putting together equations (A22) and (A17), and using the calculus theorem (A5) we get the expression for the RS quenched intensive pressure:

$$\begin{aligned} \mathcal{A}(\beta, \alpha, \gamma) &= -\beta \sum_{\mu=1}^l M_{\mu} \bar{M}_{\mu} - \frac{\alpha \beta^2}{2} P(1 - Q) - \frac{\alpha \beta^2}{2} \bar{P}(1 - \bar{Q}) \\ &+ \gamma \mathbb{E}_{\eta, \xi} \log 2 \cosh \left[\beta \sqrt{\gamma \alpha P} \eta + \beta \gamma \sum_{\mu} \bar{M}_{\mu} \xi^{\mu} \right] \\ &+ \bar{\gamma} \mathbb{E}_{\bar{\eta}, \bar{\xi}} \log 2 \cosh \left[\beta \sqrt{\gamma \alpha \bar{P}} \bar{\eta} + \beta \gamma \sum_{\mu} M_{\mu} \bar{\xi}^{\mu} \right] + \\ &- \frac{\alpha}{2} \log [1 - \beta^2 (1 - Q)(1 - \bar{Q})] + \frac{\alpha \beta^2}{2} \frac{[Q(1 - \bar{Q}) + \bar{Q}(1 - Q)]}{1 - \beta^2 (1 - Q)(1 - \bar{Q})} \\ &\equiv -\beta f(\beta, \alpha, \gamma) \end{aligned} \tag{A22}$$

with $f(\beta, \alpha, \gamma)$ being the RS free energy expression in the main text (equation (20)).

Appendix B. Additional numerical results

We report additional results on the numerical estimation of the retrieval/ SG spinodal point at low temperatures, by using a random sequential dynamics (whose update rules are given by equations (37) and (38) in section 4). In all the simulations shown in figure 10, the dynamics starts with an initial configuration having noises $\varepsilon_1 = \varepsilon_2 = 0.05$ on both layers. At each step, the dynamics is carried out by selecting one neuron among the union set of the two layers and running equations (37) or (38), depending on which layer the selected neuron belongs to; this procedure is repeated for $(N + \bar{N}) \times N_{\text{steps}}$ single neuron updates, so that -on average- each neuron is updated N_{steps} times. In figure 10 the system sizes are smaller than the ones shown in figure 7 of the main text: the latter would require an unreasonable amount of time using a random sequential dynamics. In general, despite the curve are less sharp than in figure 7 of the main text (again, because of the smaller sizes), we still find an overall good agreement between the numerical prediction of the transition w.r.t. the MF theory. We have numerically verified that these results remain qualitatively unchanged when the noise on both layers is

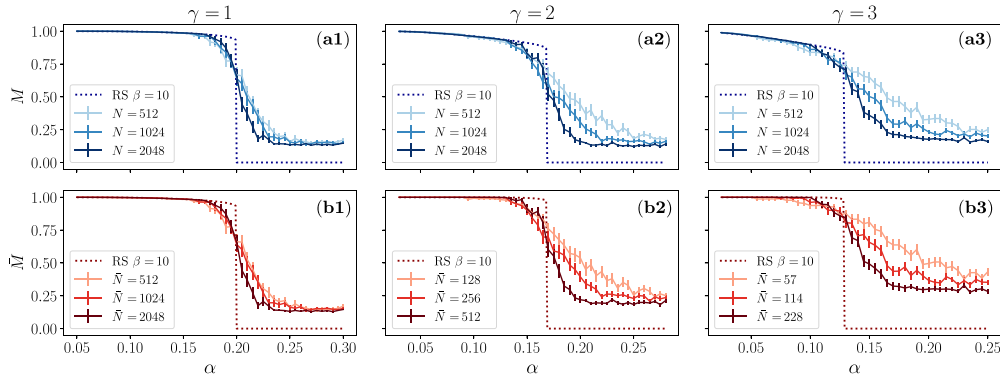


Figure 10. Numerical estimation of the retrieval/spin-glass phase transition. Simulations are performed using a random sequential dynamics at $\beta = 10$ with $N_{\text{steps}} = 10^4$ for 3 values of γ . (a1)–(b1): $\gamma = 1$; (a2)–(b2): $\gamma = 2$; (a3)–(b3): $\gamma = 3$; Each panel shows the final Mattis magnetization over each layer w.r.t. to the pattern pair to be retrieved. In all cases, the dynamics start from the pattern pair to be retrieved with a small noise $\varepsilon_1 = \varepsilon_2 = 0.05$ on both layers. Results are plotted as a function of α for three different system sizes (shown in the caption) and averaged over 50 different realizations of the disorder. Dashed lines in each panel denote the RS saddle point solution at $\beta = 10$ (i.e. a horizontal cut in the phase diagrams of figure 2).

increased up to $\varepsilon_1 = \varepsilon_2 \approx 0.3$: in the strongly asymmetric case, however, using this type of dynamics might affect the basin of attraction of the two layers in different ways w.r.t. what discussed in section 4 of the main text. As a final remark, note that the initial configuration on both layers must be constructed in such a way that the starting Mattis magnetizations have the same sign: otherwise, retrieval is lost and both Mattis magnetizations will tend to 0 (even at $\alpha \ll \alpha_c(\gamma)$), which is a clear consequence of the saddle type of the BAM Hamiltonian (1).

Appendix C. Equivalence between BAM and coupled RBMs

Recent works [15] have shown how the Hopfield model can be mapped into a binary-Gaussian RBM whose weights are directly linked to the pattern components stored in the Hopfield counterpart. Such an analogy provides an efficient way to determine the operating regimes of this specific type of artificial network for unsupervised learning: specifically, it allows on one hand to simulate more efficiently the retrieval properties of the Hopfield model [15] and, on the other hand, to describe the functioning regimes of the artificial (RBM) counterpart in terms of the ratio between the number of hidden and visible variables, which further allow for efficient and computationally cheap pre-training procedures [16]. In principle, such an analogy can be extended to the BAM, at least from a structural point of view: indeed, considering the integral transformation (13) used to decouple the interacting term in the BAM Hamiltonian (1), it is possible to interpret the right hand side of equation (14) as the partition function of a 4-partite system similar to two coupled RBMs. This is shown in figure 11: the two RBMs on (b) have each one a binary visible layer (with sizes N and \bar{N} respectively); the two hidden layers are encoded in the vectors $\mathbf{z}, \mathbf{z}^\dagger$: each of them has a size equal to the number of patterns in the BAM (K), and for each $\mu \in \{1, \dots, K\}$ the hidden nodes z_μ, z_μ^\dagger are complex conjugates, interacting through a fixed potential.




Figure 11. Schematic representation of the equivalence between the BAM (a) and two coupled RBMs (b). In (a), each link is a synaptic connection in the BAM (constructed using the Hebb rule through equation (2)). In (b), the left-most and right-most interactions correspond to the patterns components ξ_i^μ and $\bar{\xi}_j^\mu$ respectively. Red dots represent hidden nodes acting on each layer, whose number is equal to the number of patterns K in the BAM.

However, even if the analogy seems to be consistent from a structural point of view, it is not clear how to exploit it to perform learning in the resulting artificial network depicted in figure 11(b): for instance, gradient ascent on the pattern components would require sampling from a probability distribution defined on complex variables. We leave this issue for future investigation.

ORCID iDs

Adriano Barra  <https://orcid.org/0000-0003-4255-7678>
 Giovanni Catania  <https://orcid.org/0000-0002-9532-7879>
 Aurélien Decelle  <https://orcid.org/0000-0002-3017-0858>
 Beatriz Seoane  <https://orcid.org/0000-0003-4007-9406>

References

- [1] Amit D J, Gutfreund H and Sompolinsky H 1985 *Phys. Rev. Lett.* **55** 1530
- [2] Castillo I P and Skantzos N 2004 *J. Phys. A: Math. Gen.* **37** 9087
- [3] Agliari E, Asti L, Barra A, Burioni R and Uguzzoni G 2012 *J. Phys. A: Math. Theor.* **45** 365001
- [4] Castillo I P, Wemmenhove B, Hatchett J P L, Coolen A C C, Skantzos N S and Nikolettopoulos T 2004 *J. Phys. A: Math. Gen.* **37** 8789
- [5] Agliari E, Barra A, Galluzzi A, Guerra F, Tantari D and Tavani F 2015 *Phys. Rev. Lett.* **114** 028103
- [6] Agliari E, Barra A, Galluzzi A, Guerra F and Moauro F 2012 *Phys. Rev. Lett.* **109** 268101
- [7] Bollé D, Nieuwenhuizen T M, Castillo I P and Verbeiren T 2003 *J. Phys. A: Math. Gen.* **36** 10269
- [8] Leuzzi L, Patti A and Ricci-Tersenghi F 2022 *J. Stat. Mech.* **073301**
- [9] Barra A, Genovese G, Sollich P and Tantari D 2018 *Phys. Rev. E* **97** 022310
- [10] Smolensky P 1986 Information processing in dynamical systems: foundations of harmony theory *Technical Report* (Institution Colorado Univ at Boulder Dept of Computer Science)
- [11] Shimagaki K and Weigt M 2019 *Phys. Rev. E* **100** 032128
- [12] Decelle A and Furtlehner C 2021 *Chin. Phys. B* **30** 40202
- [13] Hinton G E 2002 *Neural Comput.* **14** 1771

- [14] Hinton G E and Salakhutdinov R R 2006 *Science* **313** 504
- [15] Barra A, Bernacchia A, Santucci E and Contucci P 2012 *Neural Netw.* **34** 1–9
- [16] Leonelli F E, Agliari E, Albanese L and Barra A 2021 *Neural Netw.* **143** 314
- [17] Alemanno F, Aquaro M, Kanter I, Barra A and Agliari E 2023 *Europhys. Lett.* **141** 11001
- [18] Pozas-Kerstjens A, Muñoz-Gil G, Piñol E, García-March M A, Acín A, Lewenstein M and Grzybowski P R 2021 *Mach. Learn.: Sci. Technol.* **2** 025026
- [19] Kosko B 1988 *IEEE Trans. Syst. Man Cybern.* **18** 49
- [20] Englisch H, Mastropietro V and Tirozzi B 1995 *J. Physique I* **5** 85
- [21] Kurchan J, Peliti L and Saber M 1994 *J. Physique I* **4** 1627
- [22] Tanaka T, Kakiya S and Kabashima Y 2000 *Proc. 7th Int. Conf. Neural Information Processing, Taejeon, Korea* vol 2 (Citeseer) pp 779–84
- [23] Guerra F 2006 *Mathematical Statistical Physics (Les Houches)* vol 83, ed A Bovier, F Dunlop, A van Enter, F den Hollander and J Dalibard (Amsterdam: Elsevier) pp 243–71
- [24] Mézard M, Parisi G and Virasoro M A 1987 *Spin Glass Theory and Beyond: An Introduction to the Replica Method and its Applications* (World Scientific) (<https://doi.org/10.1142/0271>)
- [25] In the whole discussion, we will always indicate with an overbar $\bar{}$ quantities referring to layer 2
- [26] The square root in equation (11) is chosen for notation convenience
- [27] Gallo I and Contucci P 2008 Bipartite mean field spin systems. Existence and solution *Math. Phys. Electron. J.* **14** 1-21
- [28] Collet F 2014 *J. Stat. Phys.* **157** 1301
- [29] Agliari E, Barra A, Longo C and Tantari D 2017 *J. Stat. Phys.* **168** 1085
- [30] Barra A, Genovese G, Sollich P and Tantari D 2017 *Phys. Rev. E* **96** 042156
- [31] Carmona P and Hu Y 2006 *Ann. Inst. Henri Poincaré B* **42** 215
- [32] Genovese G 2012 *J. Math. Phys.* **53** 123304
- [33] Hartnett G S, Parker E and Geist E 2018 *Phys. Rev. E* **98** 022116
- [34] Decelle A, Fissore G and Furtlehner C 2017 *Europhys. Lett.* **119** 60001
- [35] Decelle A, Fissore G and Furtlehner C 2018 *J. Stat. Phys.* **172** 1576
- [36] Barbier J and Macris N 2019 *Probab. Theory Relat. Fields* **174** 1133
- [37] Barbier J and Macris N 2019 *J. Phys. A: Math. Theor.* **52** 294002
- [38] Coolen A C C 2000 Statistical mechanics of recurrent neural networks I. Statics (arXiv:cond-mat/0006010)
- [39] de Almeida J R L and Thouless D J 1978 *J. Phys. A: Math. Gen.* **11** 983
- [40] Agliari E, Albanese L, Barra A and Ottaviani G 2020 *J. Phys. A: Math. Theor.* **53** 415005
- [41] Crisanti A, Amit D J and Gutfreund H 1986 *Europhys. Lett.* **2** 337
- [42] Albanese L, Alemanno F, Alessandrelli A and Barra A 2022 *J. Stat. Phys.* **189** 24
- [43] Steffan H and Kühn R 1994 *Z. Phys. B* **95** 249
- [44] Note that the optimal value of the critical load w.r.t. Θ does not necessarily coincide with its maximum value, i.e. $\alpha_c(\arg \min_{\Theta} f(\Theta)) \neq \max_{\Theta} \alpha_c(\Theta)$
- [45] Alberici D, Barra A, Contucci P and Mingione E 2020 *J. Stat. Phys.* **180** 665
- [46] Alberici D, Contucci P and Mingione E 2021 *Ann. Henri Poincaré* **22** 2619
- [47] Agliari E, Barra A, De Antoni A and Galluzzi A 2013 *Neural Netw.* **38** 52
- [48] Guerra F 2001 *Fields Inst. Commun.* **30** 161–70
- [49] Barra A, Genovese G, Guerra F and Tantari D 2012 *J. Stat. Mech.* 07009
- [50] Unless otherwise specified, the symbol \sum_{μ} represents a summation over the noisy part of the pattern, i.e. $\sum_{\mu} \equiv \sum_{\mu=l+1}^K$
- [51] Barra A 2008 *J. Stat. Phys.* **132** 787