



Statistical mechanics of vector Hopfield network near and above saturation

Downloaded from: <https://research.chalmers.se>, 2026-03-05 03:28 UTC

Citation for the original published paper (version of record):

Nicoletti, F., D'Amico, F., Negri, M. (2025). Statistical mechanics of vector Hopfield network near and above saturation. *Journal of Physics A: Mathematical and Theoretical*, 58(50).
<http://dx.doi.org/10.1088/1751-8121/ae2bd0>

N.B. When citing this work, cite the original published paper.

PAPER • OPEN ACCESS

Statistical mechanics of vector Hopfield network near and above saturation

To cite this article: Flavio Nicoletti *et al* 2025 *J. Phys. A: Math. Theor.* **58** 505005

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

You may also like

- [ICRH modelling of DTT in full power and reduced-field plasma scenarios using full wave codes](#)
A Cardinali, C Castaldo, F Napoli et al.
- [Constant maps in equivariant topological strings and geometric modeling of fluxes](#)
Luca Cassia and Kiril Hristov
- [Quantum cellular automata for word statistics facilitated by quantum correlations](#)
Guanhua Chen and Yao Yao

Statistical mechanics of vector Hopfield network near and above saturation

Flavio Nicoletti^{1,2,4} , Francesco D'Amico^{2,3,*} 
and Matteo Negri^{2,3,5} 

¹ Department of Computer Science and Engineering, Chalmers University of Technology and University of Gothenburg, SE-41296 Gothenburg, Sweden

² Dipartimento di Fisica, Sapienza Università di Roma, 00185 Rome, Italy

³ Institute of Nanotechnology, National Research Council of Italy, CNR-NANOTEC, Rome Unit, Italy

E-mail: francesco.damico@uniroma1.it, flavio@chalmers.se
and matteo.negri@uniroma1.it

Received 20 October 2025; revised 2 December 2025

Accepted for publication 11 December 2025

Published 29 December 2025



CrossMark

Abstract

We study analytically and numerically a Hopfield fully-connected network with d -dimensional vector spins. These networks are models of associative memory that generalize the standard Hopfield with Ising spins, where P examples are stored in a network of N units as local minima in an energy landscape. We study the equilibrium and out-of-equilibrium properties of the system, considering the system in its retrieval phase $\alpha < \alpha_c$ and beyond, where $\alpha = P/N$ is the capacity of the system and α_c is its critical value, above which storage fails. We derive the Replica Symmetric solution for the equilibrium thermodynamics of the system, together with its phase diagram: we find that the retrieval phase of the network shrinks with growing spin dimension, having ultimately a vanishing critical capacity $\alpha_c \propto 1/d$ in the large d limit. As a trade-off, we observe that in the same limit vector Hopfield networks are able to denoise corrupted input patterns in the first step of retrieval dynamics, up to very large capacities

⁴ Current address: Department of Computer Science and Engineering Chalmers, University of Technology and University of Gothenburg, SE-41296 Gothenburg, Sweden. The majority of the work was conducted in Dipartimento di Fisica, Sapienza Università di Roma, 00185 Rome, Italy.

⁵ New address: Laboratoire de Physique Théorique et Modélisation, CY Cergy Paris Université, 95302 Cergy-Pontoise, France.

* Author to whom any correspondence should be addressed.



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

$\alpha \propto d$. We also study the static properties of the system at zero temperature, considering the statistical properties of soft modes of the energy Hessian spectrum. We find that local minima of the energy landscape related to memory states have ungapped spectra with rare soft eigenmodes: these excitations are localized, their measure condensating on the noisiest neurons of the memory state.

Keywords: artificial neural networks, disordered systems, Hopfield networks, vector spin glasses

1. Introduction

Associative memories are neural networks that can be thought in physical terms as models where the energy function is designed to have minima in correspondence of given ‘example’ configurations, so that the dynamics of the model can reconstruct the examples if it is initialized close enough to one of those. The presence of an energy function is a major conceptual advantage, as it allows to predict the limits of functionality of these networks with the theoretical tools from the statistical mechanics of disordered systems.

Despite their importance as a theoretical model for biological networks, associative memories played a minor role during the first years of the deep learning revolution. However, many years after the classical studies on systems with higher order interactions [1, 2], developments in machine learning (ML) have renewed interest in models with enhanced capabilities, the so-called modern Hopfield networks, that can store polynomial [3] or even exponential [4, 5] number of examples. Moreover, this class of associative memories appeared in multiple context of modern deep learning, in particular transformers architectures [6] and generative diffusion [7], gathering the interest of the statistical physics community [8].

The connection between associative memories and the attention mechanism at the base of transformers architectures [9] is particularly important: given that the success of transformers is still unexplained, the possibility of using statistical mechanics of associative memories to understand their performance is worth investigating. This connection relies on the interpretation of a single attention forward pass as one step of gradient descent of an energy function [6], which can be thought as the energy of an associative memory where the minima are learned from the data in a complex way rather than just assigned to single examples [10]. This perspective is naturally useful for architectures that repeat the same layer multiple times [11–13], but was also found to be important for a single forward pass [14]. The connection was further strengthened by the finding that associative memories can create additional minima in the energy that correspond to previously unseen examples [15], showing that even the mechanism for generalization may be theoretically accessible.

Still, to successfully apply the physics of associative memories to transformers architectures, we need to address a conceptual node: that transformers operate on sequences of tokens. A token is a vector in a learned embedding space that represents a concept, which can be decoded in different data spaces allowing multi-modal applications such as image generation from a text prompt. Such embedding vectors are often very high dimensional, reaching thousands of components in the largest, most successful architectures [12, 16]. Notably, such a large number of dimensions is an ideal setting for statistical mechanics, since it may be well described by a thermodynamic limit. A first attempt to understand transformers in terms of the statistical physics of vector spins has recently been made in [17].

High-dimensional vector spins also have neurobiological relevance, as they can represent groups of continuous-valued neurons (we will discuss this more at the end of section 5).

Contributions. In this work, we introduce d -dimensional vector spins into a standard Hopfield network [18]. To the best of our knowledge, this is the first time that an associative memory with vector spins of generic dimensionality has been studied. Previous to our work, the authors in [19] use 3-dimensional vector spins, defining the learning rule through vector products in the context of optical neural networks. Moreover, 2-dimensional vector spins have been considered in oscillatory networks [20]: these models, pioneered by so-called phasor networks [21], study the problem of features extraction in composite data, by extending results from resonator networks [22, 23] using complex number neurons. Finally, artificial Kuramoto oscillators neurons were recently proposed as alternatives to threshold units in neural networks layers for ML [24].

We study the statistical mechanics of a vector Hopfield network in great detail, considering both its static and dynamical properties. First, we characterize the equilibrium properties of the system, in terms of temperature T and storage capacity α : we compute the phase diagram for any d , by extending the classic calculations from Amit *et al* [25], identifying spin glass and retrieval transitions; we also identify in the high-dimensional limit $d \rightarrow \infty$ the correct scaling with d of the capacity α for retrieval. We also study the stability of equilibrium states, in terms of $T = 0$ linear low-energy excitations, represented by the lowest modes of the hessian of the energy function. Hessian soft modes provide interesting insights on the physics of continuous-spins disordered systems with structured energy landscape, since they control their linear stability. In the present model, the study of hessian soft modes allows for the characterization of memory or Mattis states and spurious states in terms of their spectral behavior, potentially establishing a connection between the behavior of associative memory neural networks and glassy systems. Finally, we study the dynamics of the model numerically and theoretically, focusing on the ability of the model to de-noise examples. We investigate the dynamics especially in the above saturation regime, where the network cannot store memories of patterns as local minima of the energy landscape.

Summary of results. While the main motivation for our work is to understand basic properties that can be extended to modern Hopfield networks in future research, among the results that we found there are some notable behaviors that are interesting per se. The core findings of our work are the following:

- *Vector Hopfield networks can denoise patterns and mixtures above their storage capacity limits, through first-step retrieval or transient denoising* [26] (sections 3.1 and 4): in particular, the critical storage capacity for patterns vanishes with spin dimension $\alpha_c \propto 1/d$ (section 3.1), whereas their first-step retrieval capacity increases $\tilde{\alpha} \propto d$ (section 4).
- *Memory states of Vector Hopfield networks are stable mean-field glasses* [27, 28] (section 3.2): the spectrum of the Hessian matrix of local minima of the energy landscape representing memory states does not present any gap ($\lambda_{\min} \rightarrow 0$ for $N \rightarrow \infty$). Soft modes are rare and localized, spatially correlated with the noise generated by the storage process. Conversely, spurious states abound with de-localized and featureless soft modes.

Our results are discussed thoroughly in section 5.

Outline of the paper. In section 2 we give a precise definition of the model. In section 3.1 we compute the equilibrium properties of the model. In section 4 we study the out of equilibrium properties. Finally, in section 5 we discuss our results and draw conclusions.

2. Model definition

The vector Hopfield model (VHM) we study in this work is defined by the following energy function of N spins and P examples

$$\mathcal{H}(\underline{\mathbf{S}}; \{\underline{\boldsymbol{\xi}}^{(\mu)}\}_{\mu=1}^P) = -\frac{1}{2} \sum_{i \neq j}^{1,N} \mathbf{S}_i \cdot \mathbb{J}_{ij} \mathbf{S}_j \quad (1)$$

$$\mathbb{J}_{ij} = \frac{1}{N} \sum_{\mu=1}^P \boldsymbol{\xi}_{i;\mu} \boldsymbol{\xi}_{j;\mu}^\top. \quad (1)$$

Neuron spins $\{\mathbf{S}_i\}$ are vectors with d components and fixed norm σ , while example spins $\{\boldsymbol{\xi}_i^{(\mu)}\}$ are independent and identically distributed (iid) random vectors with d components and unit norm, drawn from the uniform distribution on the unit sphere in d dimensions:

$$P_{\boldsymbol{\xi}}(\boldsymbol{\xi}) = \frac{1}{\mathcal{S}_{d-1}(1)} \delta(|\boldsymbol{\xi}| - 1) \quad (2)$$

where $\mathcal{S}_{d-1}(1) = \frac{2\pi^{d/2}}{\Gamma(d/2)}$ is the surface of the unit hypersphere embedded in dimension d . In the following, with the notation $\underline{\mathbf{S}} = (\mathbf{S}_1, \dots, \mathbf{S}_N)^\top$ and $\underline{\boldsymbol{\xi}}^\mu = (\boldsymbol{\xi}_1^\mu, \dots, \boldsymbol{\xi}_N^\mu)^\top$ we identify respectively configuration of spins and patterns.

Our definition in equations (1), (2) is a natural extension of the standard Hopfield Model (SHM) to vector spins: the SHM Hamiltonian in fact is the particular case $d = 1$ of (1), (2). In section 3.1 we study the thermodynamics of the model in (1).

Note that our choice of couplings in (1) implies no global spin symmetries—the invariance under the $O(d)$ group in the case of vector spins. Conversely, the SHM in his original definition in [18, 29, 30] has a global spin-inversion symmetry. In general, the VHM in (1) is not the only possible generalization of the SHM to vector spins: an isotropic VHM with scalar couplings could also be defined. By preliminary studies of such VHM, we found that it bears no practical advantages over our model in (1), but rather unwanted features such as zero modes connected to the $O(d)$ symmetry.

Equation (1) is a model of associative memory as the SHM studied in [18, 29, 30]: configurations of neurons correlated to examples can be stored as local minima of the energy in (1). Depending on how large is P , the energy landscape represented by (1) may include so-called spurious minima, configurations of neurons unrelated to the original set of P examples. The retrieval dynamics equations of the VHM can be obtained from the condition of stationarity of the energy function (1). Stationary points can be identified by introducing Lagrange multipliers $\{\eta_i\}_{i=1}^N$ that enforce spins to have fixed norms:

$$\frac{\partial}{\partial \mathbf{S}_i} \left[\mathcal{H}(\underline{\mathbf{S}}) + \sum_{i=1}^N \eta_i (|\mathbf{S}_i| - \sigma) \right] \Big|_{\underline{\mathbf{S}} = \underline{\mathbf{S}}_*} = 0 \quad (3)$$

so that

$$\begin{cases} \eta_i = |\sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j^*| \\ \mathbf{S}_i^* = \frac{\sigma \sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j^*}{|\sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j^*|}. \end{cases} \quad (4)$$

The Lagrange multipliers in (3) are fixed to equate the norm of the local field vectors acting on the spins. The physical meaning of the $\{\eta_i\}$ is that they represent the strength of the input signal that each spin neuron i receives from all the other spin neurons. In ML, local fields are often called pre-activations or receptive fields, depending on the context. In the following, when we use the term *local field*, we will always refer to the Lagrange multipliers η_i , and not to the corresponding vectors in (3). Equation (3) can be used to build a minimization algorithm for equation (1). Choosing a synchronous update rule for the spins, we introduce the *versor rule* for retrieval dynamics

$$\mathbf{S}_i(t+1) = \frac{\sigma \sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j(t)}{|\sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j(t)|} = \frac{\sigma \boldsymbol{\eta}_i(t)}{|\boldsymbol{\eta}_i(t)|}. \quad (5)$$

According to (5), at each time step each spin is aligned to its local field vector $\boldsymbol{\eta}_i(t) = \sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j(t)$. Equation (5) is a greedy algorithm, since at each time step the spin configuration update follows the steepest descent direction, represented by the negative gradient of the energy. Note that (5) is the generalization to VHM of the *sign rule* introduced by Hopfield in [18]. In practice, the iteration of equation (5) is stopped when the average overlap of the spins with the spins in the previous configuration is larger than a threshold $1 - \epsilon$. In our simulations, we always use $\epsilon = 10^{-12}$. The retrieval dynamics of our VHM is discussed in greater detail in appendix D. In addition, in section 4 we study the retrieval dynamics just after initialization, unveiling the phenomenon of transient denoising.

At variance with the SHM, VHMs are continuous models and thus allow for studying linear excitations around local minima, which are encoded in the $Nd \times Nd$ Hessian matrix of the energy function (1). The Hessian can be written as a block matrix in site indices

$$\mathbb{M}_{ij}(\mathbf{S}) = \mathbb{P}_\perp(\mathbf{S}_i) \left(-\mathbb{J}_{ij} + \frac{\eta_i}{\sigma} \delta_{ij} \mathbb{I}_d \right) \mathbb{P}_\perp(\mathbf{S}_j) \quad (6)$$

where \mathbb{I}_d is the $d \times d$ identity matrix. Orthogonal projectors $\mathbb{P}_\perp(\mathbf{S}) = \mathbb{I}_d - \frac{\mathbf{S}\mathbf{S}^\top}{\sigma^2}$ suppress longitudinal fluctuations and are a consequence of fixed norm constraints on spins. In general, the Hessian (28) depends on the configuration on which it is evaluated, but in mean field systems it is well known that in the thermodynamic limit $N \rightarrow \infty$ physical Hessians can be represented by random matrix ensembles: this simplification allows one to study analytically their properties. Hessian matrices with a block structure similar to that in equation (6) appear frequently in disordered systems; see, for instance, [27, 28, 31–33]. In section 3.2 we will explain how we study (28) through random matrix theory. We are interested in studying the behavior of the spectrum of (28) close to the lower spectral edge: this region is the most relevant for the physics of the system, since Hessian soft modes rule both the long-time close-to-equilibrium relaxation and the linear stability of the system. The fixed points of retrieval, corresponding to either Mattis states or spurious states, can thus be characterized according to the properties of their excitation spectra.

3. The network below saturation

In this section we study the properties of the VHM (1) below saturation, $\alpha < \alpha_c$. In this regime, when initialized close to one of the patterns, the system can retrieve a memory of the pattern through retrieval dynamics (5). Therefore, in this section we only consider the static properties of the system, deriving first the solution at thermodynamic equilibrium for generic

temperatures T and storage capacity $\alpha = P/N$ and then focusing on the linear stability of the system at zero temperature, by studying the statistics of soft modes of the Hessian (28).

3.1. Equilibrium thermodynamics

Thermal averages with respect to the Boltzmann-Gibbs distribution are denoted by the symbol $\langle \dots \rangle$, while averages over examples (2) by $\overline{\dots}$.

3.1.1. Order parameters. The order parameters describing the retrieval phase, as in the SHM, are the so-called Mattis magnetizations

$$m_\mu = \lim_{N \rightarrow \infty} \frac{1}{N\sigma} \sum_{i=1}^N \langle \mathbf{S}_i \rangle \cdot \boldsymbol{\xi}_i^\mu = \frac{1}{\sigma} \overline{\mathbf{S}_1 \cdot \boldsymbol{\xi}_1^\mu}. \quad (7)$$

These quantities describe the correlation of the equilibrium state with the examples. Let us consider the vector $\mathbf{m} = (m_1, \dots, m_P)^\top$ to identify Mattis magnetizations. When the equilibrium state is correlated to a single example, like $\mathbf{m} = (0, \dots, 0, m_\mu, 0, \dots, 0)^\top$, the equilibrium state is a *Mattis* or *memory* state; instead, if the equilibrium state is correlated with a finite number $s \ll P$ of examples, like $\mathbf{m} = (m_1, \dots, m_s, 0, \dots, 0)^\top$, we identify it as a mixture state; finally, states uncorrelated with all examples, $\mathbf{m} = \mathbf{0}$, are spin-glass states. Mixtures and spin glass solutions are termed *spurious states* in the classical Hopfield literature [34]. Recently, it was shown that mixture states can be a pathway to understanding learning and generalization in SHM with random features [15, 35].

When $P \propto N$ ($\alpha > 0$), analogously to the SHM, we find that spurious states are spin glass states with replica symmetry breaking (RSB): the order parameter in this case is the overlap distribution $P(q') = \sum_{\alpha\beta} w_\alpha w_\beta \delta(q' - q_{\alpha\beta})$ between different pure states α, β [36]. In this work, we consider the replica symmetry (RS) ansatz $P(q') \simeq \delta(q' - q)$: in this case, the order parameter is the Edward-Anderson overlap

$$q = \lim_{N \rightarrow \infty} \frac{1}{N\sigma^2} \sum_{i=1}^N |\langle \mathbf{S}_i \rangle|^2 = \frac{1}{\sigma^2} \overline{|\langle \mathbf{S}_1 \rangle|^2}. \quad (8)$$

This order parameter describes the persistence of spins in a disordered magnetic state. We found that for $\alpha > 0$ the RS ansatz for the overlap is exact for Mattis states. For spurious states, the RS ansatz is exact only in the subextensive regime $P \ll N$ ($\alpha = 0$). In this regime, one can show that $q = m^2$, where m is the equilibrium Mattis magnetization.

As in the SHM, when $\alpha > 0$ Mattis states have a magnetization $0 < m < 1$: the network is not able to perfectly retrieve examples, but rather a noisy memory of them. The noise on each Mattis state comes from the feedback of all the other Mattis states and has a typical value $\sqrt{\alpha r}$, where r is the noise parameter

$$r = \frac{1}{P} \sum_{\mu=2}^P \left(\frac{1}{\sqrt{N}} \sum_{i=1}^N \boldsymbol{\xi}_i^\mu \cdot \langle \mathbf{S}_i \rangle \right)^2 = \overline{\left(\frac{1}{\sqrt{N}} \boldsymbol{\xi}^{(2)} \cdot \langle \mathbf{S} \rangle \right)^2}. \quad (9)$$

As we show in appendix A, the noise parameter r is not independent, but rather a function of q , therefore it cannot be strictly considered an order parameter.

3.1.2. Equilibrium free energy. We computed the free energy $f = -\lim_{N \rightarrow \infty} \frac{T}{N} \overline{\log \mathcal{Z}}$ using the Replica Method, a standard technique in the statistical physics of disordered systems [36]: it is based on the renowned Replica trick

$$\overline{\log \mathcal{Z}} = \lim_{n \rightarrow 0} \frac{\overline{\mathcal{Z}^n} - 1}{n}. \quad (10)$$

In practice, with (10), one introduces n uncoupled copies of the original system and then evaluates the partition function of the replicated system. This technique allows in mean field models to decouple the degrees of freedom (in our model, neuron spins) after the disorder average is performed, reducing the computation of the replicated partition function to a saddle-point integral

$$\overline{\mathcal{Z}^n} = \overline{\int D\mathbf{S}_1 \cdots D\mathbf{S}_n e^{-\beta \sum_{a=1}^n \mathcal{H}(\mathbf{S}_a)}} = \int d\mathbb{Q} \prod_{\mu=1}^s \prod_{a=1}^n (dm_{\mu}^a) e^{NA_n(\{m_{\mu}^a\}, \mathbb{Q})}. \quad (11)$$

The replica action A_n depends on the overlap matrix and the replicated Mattis magnetizations

$$Q_{ab} = \frac{1}{N\sigma^2} \mathbf{S}_a \cdot \mathbf{S}_b \quad a \neq b = 1, \dots, n \quad (12)$$

$$m_{\mu}^a = \frac{1}{N\sigma} \mathbf{S}_a \cdot \boldsymbol{\xi}_{\mu} \quad \mu = 1, \dots, s \ll P.$$

In this work, we use the Replica Symmetric (RS) ansatz: $Q_{ab} = q$ for any $a \neq b$ and $m_{\mu}^a = m_{\mu}$ for any a . The free energy of the system is obtained after casting the saddle point evaluation of (11) back into (10)

$$f(\alpha, \beta) = -\frac{1}{\beta} \lim_{n \rightarrow 0} \frac{\max A_n(\{m_{\mu}\}_{\mu=1}^s, q)}{n} = \max \phi(\{m_{\mu}\}_{\mu=1}^s, q), \quad (13)$$

where $\beta \equiv 1/T$ is the inverse temperature⁶. The replica order parameters extremizing the free energy correspond to the physical order parameters (7), (8). The noise parameter (9) corresponds to a conjugated field of q in the replica computation. The complete replica computation of the free energy is presented in appendix A, together with the expression of A_n .

Despite being non-rigorous, Replica method leads to the correct result in several spin glass models [37]. Alternatively, the RS solution of the model can also be derived with the cavity method [36, 38], by generalizing the approach used in [39, 40] to vector spins.

3.1.3. Sub-extensive storage. We found in the $\alpha = 0$ case the free energy

$$f_{\text{eq}}(\beta) = \frac{\sigma^2}{2} \sum_{\mu \leq s} m_{\mu}^2 - \frac{1}{\beta} \ln \sigma^{d-1} \mathcal{K}_d \left(\beta \sigma^2 \left| \sum_{\mu \leq s} m_{\mu} \boldsymbol{\xi}_{\mu} \right| \right) \quad (14)$$

⁶ Note that in last equation we used the known fact that when $n \rightarrow 0$, $A_n = O(n)$ and the saddle point extremizing the action switches from a maximum to a minimum.

where Mattis magnetizations are given by the solution of the saddle point equations

$$m_\mu = \mathfrak{g}_d \left(\beta\sigma^2 \left| \sum_{\mu=1}^s m_\mu \xi_\mu \right| \right) \frac{\xi_\mu \cdot \sum_{\nu=1}^s m_\nu \xi_\nu}{\left| \sum_{\nu=1}^s m_\nu \xi_\nu \right|}. \quad (15)$$

We introduced the functions

$$\mathcal{K}_d(x) = \int d\mathbf{r} \delta(|\mathbf{r}| - 1) e^{r \cdot x} = (2\pi)^{d/2} \frac{I_{d/2-1}(x)}{x^{d/2-1}} \quad (16)$$

$$\mathfrak{g}_d(x) = \frac{1}{\mathcal{K}_d(x)} \frac{d\mathcal{K}_d(x)}{dx} = \frac{I_{d/2}(x)}{I_{d/2-1}(x)}$$

where $I_\nu(x)$ is the modified Bessel function of order ν . The function $\mathcal{K}_d(x)$ is the partition function of a single spin under the action of a field of magnitude x , while the function $0 \leq \mathfrak{g}_d(x) < 1$ represents its related magnetization. For comparisons, in the SHM one has $\mathcal{K}_1(x) = 2 \cosh(x)$ and $\mathfrak{g}_1(x) = \tanh(x)$. Note that for $d = 1$, one recovers the equilibrium free energy and the saddle point equations in [29].

For $\alpha = 0$ the model has a second-order phase transition at $T_c = \sigma^2/d$, from a paramagnetic phase to a ferromagnetic retrieval phase. One can see this from equation (15) in the case of Mattis states ($s = 1$):

$$m = \mathfrak{g}_d(\beta\sigma^2 m). \quad (17)$$

Since $\mathfrak{g}_d(x) \sim x/d$ for small x , a nonzero solution of (17) exists if and only if $\beta > d/\sigma^2 = \beta_c$. Note that in order to fix $T_c = 1$ as in the SHM one needs spins with norm $\sigma = \sqrt{d}$. The retrieval solution of (17) defines Mattis states. The solution of equation (17) behaves as a Ferromagnetic spin magnetization close to the transition, $m \simeq \frac{\beta_c}{\beta} \sqrt{1 + \frac{2}{d}} \sqrt{1 - \frac{\beta_c}{\beta}}$, and converges as a power law to unit in the zero-temperature limit $m \simeq 1 - \frac{d-1}{2\beta\sigma^2}$. Note that this scaling contrasts with that of the SHM, where, because of the discrete nature of the spins, m converges exponentially fast to unit. At $T = 0$, (14) converges to the energy density

$$e = \frac{\sigma^2}{2} \sum_{\mu \leq s} m_\mu^2 - \sigma^2 \left[\sum_{\mu \leq s} m_\mu \xi_\mu \right] \quad (18)$$

with the saddle-point equation

$$m_\mu = \left[\frac{\xi_\mu \cdot \sum_{\nu=1}^s m_\nu \xi_\nu}{\left| \sum_{\nu=1}^s m_\nu \xi_\nu \right|} \right]. \quad (19)$$

The P -degenerate ground state level of the system corresponds to Mattis states ($s = 1$) and its value is $e_1 = -\frac{\sigma^2}{2}$.

We found that the solution of (17) is stable at all temperatures $T < T_c = \sigma^2/d$. In contrast, we found that mixture solutions of (15) are unstable at all temperatures. We focused on the case of symmetric mixtures, $m_\mu = m$ if $\mu \leq s$ and zero otherwise, and we determined that the instability of the solution is related to symmetry-breaking orthogonal fluctuations $m_\mu = m + \delta m_\mu$. We observed numerically that no differences occur in asymmetric mixtures, therefore we conclude that, at variance with the SHM, in VHM mixtures of examples are not stored as spurious states. This different behavior is a consequence of the different nature of the degrees of freedom between VHM (continuous) and SHM (discrete). The instability of mixture states

is related to the presence of transverse fluctuations of $z = \sum_{\mu \leq s} \xi_\mu$, which in the $d = 1$ case are absent. It was already observed by Amit *et al* in [29] that when considering continuous disorder (e.g. examples sampled from the uniform distribution in a symmetric interval) mixture solutions destabilize also in the SHM. In the large d limit we found for any $s > 1$ that mixture solutions become marginally stable. We provide detailed stability computations for Mattis states and mixtures in appendix B.

3.1.4. Extensive storage. We consider the case $\alpha > 0$: since mixture solutions are unstable at $\alpha = 0$, in the following, we will consider only non-retrieval solutions ($s = 0$) and retrieval ones ($s = 1$).

We used the RS ansatz for the replica overlap matrix and Mattis magnetizations

$$Q_{ab}^* = q \quad a \neq b, \quad m_a^* = m \quad (20)$$

where $0 < q < 1$ is the Edward-Anderson order parameter introduced in (8) and m is the Mattis magnetization; with this ansatz, we obtained the free energy density

$$\begin{aligned} f_{\text{eq}}(\alpha, \beta) = & \frac{\sigma^2}{2} m^2 - \frac{1}{\beta} \log \sigma^{d-1} - \frac{1}{\beta} \int_0^\infty dh P_h(h) \log \mathcal{K}_d(\beta \sigma h) \\ & + \frac{\alpha}{2} \left\{ \frac{\sigma^2}{d} + \frac{1}{\beta} \log \left[1 - \frac{\beta \sigma^2}{d} (1-q) \right] + \frac{\beta(1-q)q}{\left[\frac{d}{\sigma^2} - \beta(1-q) \right]^2} - \frac{q}{\frac{d}{\sigma^2} - \beta(1-q)} \right\} \end{aligned} \quad (21)$$

where r, m, q satisfy the saddle point equations

$$\begin{aligned} r &= \frac{q}{\sigma^2 \left[\frac{d}{\sigma^2} - \beta(1-q) \right]^2} \\ q &= \frac{\int \frac{d\mathbf{h}}{[2\pi\alpha r]^{d/2}} \mathfrak{g}_d(\beta\sigma|\mathbf{h} + \sigma m \boldsymbol{\xi}|)^2}{\int \frac{d\mathbf{h}}{[2\pi\alpha r]^{d/2}} \mathfrak{g}_d(\beta\sigma|\mathbf{h} + \sigma m \boldsymbol{\xi}|)^2} = \int_0^\infty dh P_h(h) \mathfrak{g}_d(\beta\sigma h)^2 \\ m &= \frac{\int \frac{d\mathbf{h}}{[2\pi\alpha r]^{d/2}} \frac{(\mathbf{h} + \sigma m \boldsymbol{\xi}) \cdot \boldsymbol{\xi}}{|\mathbf{h} + \sigma m \boldsymbol{\xi}|} \times \mathfrak{g}_d(\beta\sigma|\mathbf{h} + \sigma m \boldsymbol{\xi}|)}{\int \frac{d\mathbf{h}}{[2\pi\alpha r]^{d/2}} \mathfrak{g}_d(\beta\sigma|\mathbf{h} + \sigma m \boldsymbol{\xi}|)} \\ &= \int_0^\infty dh P_h(h) \mathfrak{g}_d\left(\frac{m\sigma h}{\alpha r}\right) \mathfrak{g}_d(\beta\sigma h). \end{aligned} \quad (22)$$

The equilibrium free energy and the saddle point equations reported above for $d = 1$ are the same obtained in [30]. The effective field \mathbf{h} appearing in eqs (21), (22) can be related to the concept of cavity vector: for a given spin S_i , its local fields vectors $\boldsymbol{\eta}_i$ can be decomposed into two contributions, one coming from the interaction with the other spins independently from the value of S_i (the cavity contribution) and one that takes into account the polarizing effect that S_i has on them (Onsager polarization [36]). In mean-field disordered systems, the two contributions are independent. We will come back to Onsager polarization later in section 3.2.

The probability density function (PDF) of $h = |\mathbf{h}|$ can be computed after performing the average over $\boldsymbol{\xi}$ and reads

$$P_h(h) = \begin{cases} \frac{h^{d-1} e^{-\frac{h^2}{2\alpha r}}}{(2\pi\alpha r)^{d/2}} e^{-\frac{m^2\sigma^2}{2\alpha r}} \mathcal{K}_d\left(\frac{m\sigma}{\alpha r} h\right) & m > 0 \\ \frac{h^{d-1} \mathcal{S}_{d-1} e^{-\frac{h^2}{2\alpha r}}}{(2\pi\alpha r)^{d/2}} & m = 0. \end{cases} \quad (23)$$

Note that the cavity pdf of Mattis states is strongly suppressed close to the lower edge $h = 0$: in fact, one finds for small fields $P_h(h; m)/P_h(h; m = 0) \simeq e^{-\frac{m^2\sigma^2}{2\alpha r}} \ll 1$. We will show later in section 3.2 that this property is related to a better stability of Mattis states against perturbations, because the abundance of low fields h is a proxy for that of soft eigenmodes in the Hessian spectrum.

At $T = 0$, the energy density of the system reads

$$e(\alpha) = \frac{\sigma^2}{2} m_0^2 - \sigma \int_0^\infty dh P_h(h; T=0) h + \frac{\alpha}{2} \left[\frac{\sigma^2}{d} + \frac{\beta\chi}{\left(\frac{d}{\sigma^2} - \chi\right)^2} - \frac{1}{\frac{d}{\sigma^2} - \chi} \right]. \quad (24)$$

We introduced the linear susceptibility $\chi = \beta(1 - q)$: this quantity represents the average static linear response of a neuron spin to the action of an external field. In VHMs, a physical solution for the susceptibility is strictly positive and upper bounded, $0 < \chi < d/\sigma^2$. The $T = 0$ saddle point equations read

$$\begin{aligned} r_0 &= \frac{1}{\sigma^2} \frac{1}{\left(\frac{d}{\sigma^2} - \chi\right)^2} \\ \chi &= \frac{(d-1)}{\sigma} \int_0^\infty dh \frac{P_h(h)}{h} \\ m &= \int_0^\infty dh P_h(h) g_d\left(\frac{m\sigma h}{\alpha r}\right) \end{aligned} \quad (25)$$

where it is implicit that in these last equations the parameters of $P_h(h)$, whose expression is given by (23), are the $T = 0$ ones, appearing on the lhs. The susceptibility of the non-retrieval solution is simply $\chi(m = 0) = \frac{cd}{(\sqrt{\alpha+c}\sigma)^2}$, with $c = (d-1) \frac{\mathcal{S}_{d-1}(1)}{\sqrt{2\pi}\mathcal{S}_{d-2}(1)}$. Note that $m_0 < 1$ (compare the third of (25) with the second of equation (16)) for any $\alpha > 0$, as expected.

We numerically solved equations (22), (25) as a function of α . In figure 1, we show the solution we obtained for m , χ and the free energy f as functions of α , for the $d = 3$ VH model and temperature values $T = 0, 0.25, 0.50$. The qualitative physical behavior we find for VH models is the same as that in the SH model: the retrieval solution $m > 0$ exists up to a critical storage α_c , it is a meta-stable solution for $\alpha_m < \alpha < \alpha_c$ and a stable one for $\alpha < \alpha_m$; at $\alpha = \alpha_m$, the system undergoes a first order phase transition.

Stability of the RS solution: the stability of the RS solution is determined by perturbing the Replica Action. The lowest Hessian eigenvalue is known in the spin glass literature as *Replicon* eigenvalue Λ . In our model, we find that this quantity has the following expression

$$\Lambda = 1 - \frac{\alpha\beta^2 r}{q} \int_0^\infty dh P_h(h) \left[(g'_d(\beta\sigma h))^2 + (d-1) \frac{(g_d(\beta\sigma h))^2}{\beta^2 h^2} \right] \quad (26)$$

The Replicon is related to the divergence of the spin glass susceptibility $\chi_{\text{sg}} \propto 1/\Lambda$: only when $\Lambda > 0$, the RS solution is stable, the solution of $\Lambda = 0$ determining the critical line $T_c(\alpha)$ for the RSB spin glass transition.

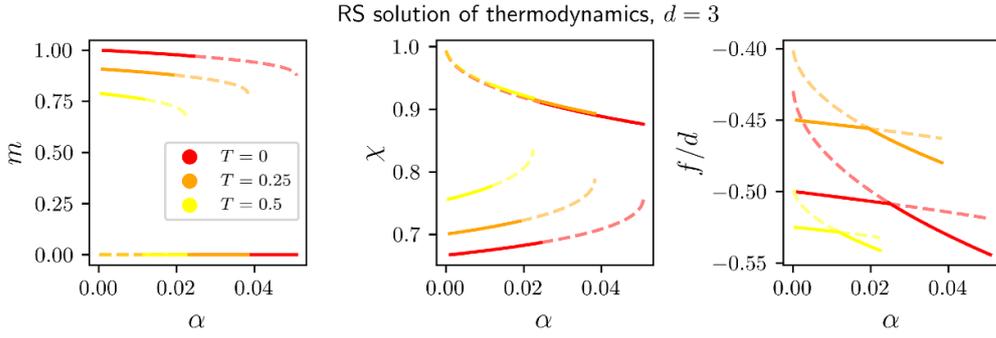


Figure 1. The model is a genuine Hopfield associative memory: Solution of the $d = 3$ VHM (1) in the RS approximation. We show the Mattis magnetization m (left), the linear susceptibility χ (center) and the free energy density (right) as functions of the capacity α , for temperatures $T = 0., 0.25, 0.5$. The behavior is analogous to that of the SHM: for any T , the solution exists up to a critical capacity α_c and undergoes an equilibrium first order phase transition at $0 < \alpha_m < \alpha_c$. We indicate equilibrium solution with continuous lines, while dashed opaque lines identify metastable solutions. Here we show data for spins with norm $\sigma = \sqrt{d}$, thus the free energy density is normalized by a factor d and $0 < \chi < 1$.

We found that the RS solution for spurious states is unstable whenever the equation for q in (22) admits a non-zero solution: this happens at the critical temperature $T_c(\alpha) = \frac{\sigma^2(\sqrt{\alpha} + \sqrt{d})}{d^{3/2}}$, as can be seen by expanding the $m = 0$ equation for q in (22) for small q . As in the SHM, spurious states of VHM are full RSB (fRSB) spin glass states. In particular, at $T = 0$ and for $d > 2$ the Replicon is equal to $-\frac{1}{d-2}$: in the large d limit, the no-retrieval phase becomes a marginally stable RS phase, as in the spherical spin glass [41, 42].

At variance with the SH model ($d = 1$), we observe that the retrieval phase of VHM with $d > 2$ is always RS. Indeed, we found that equation (26) and its $T = 0$ limit

$$\Lambda_0 = 1 - (d-1)\alpha r_0 \int_0^\infty dh \frac{P_h(h)}{h^2} \tag{27}$$

return a positive Replicon in the whole retrieval phase when $d > 2$. This is not the case for $d = 2$, since the integral in (27) for this dimension diverges to $-\infty$, due to $P_h(h) \propto h^{d-1}$ close to $h = 0$. As a consequence, the $d = 2$ VH model must undergo a spin glass transition at finite temperature: using equation (26), we find the transition line $T_c(\alpha) \sim e^{-\frac{1}{2}} e^{\frac{m^2}{\alpha r}}$. In the SH model, the retrieval solution undergoes an RSB transition at $T_c(\alpha) \propto e^{-\frac{m^2}{2\alpha r}}$ [30]. Thus, RSB corrections to the RS solution for $d = 2$ are expected to be even smaller than those observed for $d = 1$, since the transition onsets at very low temperatures: we report $T_c(\alpha_c) \simeq 10^{-6}$ for $d = 2$.

Phase diagram and scaling regimes: In the preceding lines, we discussed the spin glass transition curves for no-retrieval and Mattis states. The retrieval lines of the VHM (1) are analogous to those of the SH: there is a spinodal line $T_c(\alpha)$ that marks the appearance of the retrieval solution as a metastable solution of thermodynamics, with spin glass states dominating the equilibrium measure, and a line $T_m(\alpha) < T_c(\alpha)$ that signals a first order equilibrium phase transition where the retrieval solution becomes the equilibrium one. At variance with the SHM [29, 30], in the VHM there are no retrieval lines for solutions involving mixtures of examples, since we determined that solutions of mixtures of examples are always unstable (see appendix B).

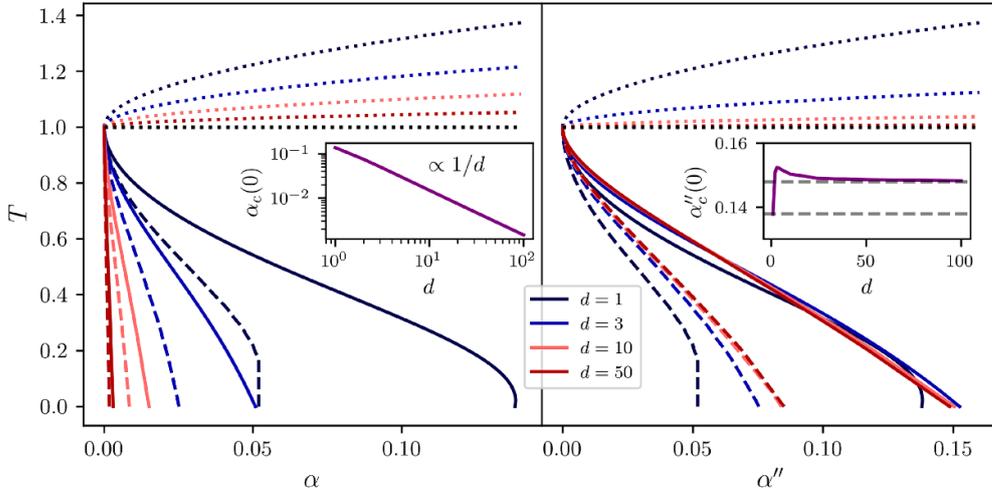


Figure 2. The retrieval phase shrinks with growing spin dimension: Phase diagrams of VHMs with $d = 3, 5, 10, 50$, compared with that of the SHM ($d = 1$). Spins are chosen with norm $\sigma = \sqrt{d}$ in order to fix the no-retrieval spin glass temperature at $\alpha = 0$ to unit. Left: the phase diagram is shown in the T, α plane. Retrieval lines shrink for growing d , while the spin glass line for no-retrieval states converge to unit. In the inset, the zero-temperature critical capacity $\alpha_c(0)$ is shown to decrease as $1/d$. Right: the same phase diagram but with $\alpha'' = \alpha d$ on the x -axis. Retrieval lines in this scaling regime seem to converge to non-trivial curves. In particular, we find that the critical capacity $\alpha_c(0) \rightarrow 4/27$ in the large d limit, as shown in the inset.

In the left panel of figure 2 we show the phase diagram of VHMs with $d = 3, 5, 10, 50$ in the T, α plane, portraying both the spin glass transition curves for no-retrieval states $T_{\text{sg}}^{(\text{nr})} = \frac{\sigma^2(\sqrt{\alpha} + \sqrt{d})}{d^{3/2}}$ (dotted lines) and critical curves for retrieval, respectively the spinodal (full lines) and the first-order transition (dashed lines). We set $\sigma = \sqrt{d}$ for spins norms, in order to have $T_{\text{sg}}^{(\text{nr})}(\alpha = 0) = 1$ independently from d . We derived the retrieval lines by numerically solving eqs (22), (25) for several α, T . More in detail, the line $T_c(\alpha)$ is obtained by finding for a given capacity α the largest temperature such that the retrieval solution $m > 0$ exists, while the line $T_m(\alpha)$ was computed by finding for any α the temperature at which the free energy of memory states is equal to that of spurious spin glass states. In VHM with $d \geq 3$ the spinodal lines $T_c(\alpha)$ we derived are exact, having no reentrance phenomena as in the SHM [30, 43], while first-order transition lines $T_m(\alpha)$ computed by us through the RS ansatz are not exact, since spurious states are RSB spin glass states. We found that for increasing d , the storage capacity of our VHM is vanishing and always lower than that of the SHM, with both the $T = 0$ critical capacities α_c, α_m scaling as $1/d$ for large d : in the inset in the left panel of figure 2 we show the critical capacity $\alpha_c(0)$ as a function of d .

In the right panel of figure 2, we show the phase diagram in the T, α'' plane, with $\alpha'' = \alpha d$; at this scale, we see that retrieval is maximum at $d = 3$ and then apparently saturates to $\alpha''(d \rightarrow \infty) = 4/27$ in the large d limit, as shown in the inset: interestingly, this is the same critical capacity found in [44] for a spherical Hopfield model with both pairwise and four-body interactions. The scaling α'' corresponds to storing $P = \alpha'' \frac{N}{d}$ examples: so for a given system size N , as the number of dimensions of vector spins increases the number of memories that can be stored decreases linearly. We also identified a scaling $\alpha' = \alpha/d$, corresponding to a regime

in which the number of memories increases linearly with the dimension, $P = \alpha'Nd$. With this scaling, the equilibrium free energy (21) in the limit $d \rightarrow \infty$ converges to the free energy of a spherical Hopfield model with pairwise interactions. We will see in section 4 that the scaling $\alpha' = \alpha/d$ is also relevant for the very first steps of the retrieval dynamics in (5), as it sets the limit capacity for transient retrieval.

We discuss the scaling regimes in spin dimension d of the RS thermodynamic solution in appendix C.

3.2. Linear stability

In this section, we study the statistics of the Hessian spectrum of different local minima of the energy landscape, characterizing $T=0$ fluctuations of Mattis and spurious spin glass states. We obtain theoretical equations for the spectrum using random matrix techniques and we test our predictions with data obtained from exact numerical diagonalization. Numerical samples corresponding to local minima of the energy landscapes were obtained through equation (5), using respectively pattern and random initial conditions for memory and spurious states (see also algorithm 1 in appendix D).

The Hessian matrix related to the energy function in (1) reads

$$\mathbb{M}_{ij}(\underline{\mathbf{S}}) = \mathbb{P}_\perp(\mathbf{S}_i) \left(-\mathbb{J}_{ij} + \frac{\eta_i}{\sigma} \delta_{ij} \mathbb{I}_d \right) \mathbb{P}_\perp(\mathbf{S}_j). \quad (28)$$

The projectors $\mathbb{P}_\perp(\mathbf{S}) = \mathbb{I} - \frac{\mathbf{S} \mathbf{S}^\top}{|\mathbf{S}| |\mathbf{S}|}$ ensure that fluctuations keep the spins norms fixed. Therefore, the nontrivial part of the Hessian in (28) is constituted by excitations that are orthogonal to the configuration $\underline{\mathbf{S}}$ on which the Hessian is evaluated. The longitudinal sector of the spectrum is instead associated with the stationarity condition (3), defining the local fields η_i . In the following, we will only consider the orthogonal sector of the spectrum.

We focus our analysis on the statistics of eigenvalues and eigenvectors in the $\lambda \rightarrow 0_+$ limit. This region of the spectrum is the most physically relevant for the static stability of the system against small linear perturbations. In particular, the softest modes rule the long-time relaxation dynamics of the network close to a fixed point.

3.2.1. The spectral equation. The statistical properties of the Hessian spectrum are encoded in the resolvent matrix $\mathbb{G}(z) = (\mathbb{M} - z\mathbb{I})^{-1}$, where z is a complex number that lies outside the spectrum of the Hessian matrix \mathbb{M} (equation (28)).

We derived an equation for the normalized trace of the resolvent $\mathcal{G}(z) = \frac{d-1}{Nd} \text{Tr}(\mathbb{G}(z))$, that is exact in the thermodynamic limit $N \rightarrow \infty$. Setting $z = \lambda - i0_+$, it reads⁷

$$\mathcal{G}(\lambda) = \frac{(d-1)}{\sigma} \int_0^\infty \frac{d\eta P_\eta(\eta)}{\eta - \lambda\sigma - \frac{\alpha\sigma}{d} \frac{\mathcal{G}(\lambda)}{\sigma^2 - \mathcal{G}(\lambda)}} \quad (29)$$

where $P_\eta(\eta)$ is the distribution of the local fields η_i on a stationary state. Equation (29) was derived through a free convolution [45] of the interactions and local field matrices that appear in the Hessian (28), as explained in appendix E.4. Equation (29) is a self-consistent equation for $\mathcal{G}(\lambda)$, in its real and imaginary parts. If the distribution $P_\eta(\eta)$ is known, equation (29) can be solved numerically as explained in appendix E.5, allowing one to compute the distribution of eigenvalues and related eigenvectors, as explained later in the following paragraphs.

⁷ In its derivation, the Green function in equation (29) is implicitly regularized on the spectral line.

3.2.2. The distribution of local fields. To have verifiable analytical predictions for the spectrum through (29), it is therefore crucial to have a theoretical expression for the distribution $P_\eta(\eta)$. The information on the type of minimum on which the Hessian is evaluated (memory or spurious spin-glass states) is completely encoded in $P_\eta(\eta)$. In disordered systems on dense networks, local fields η_i are simply related to cavity fields h_i in the thermodynamic limit, through $\eta_i = h_i + p$: the term p is the Onsager reaction, which, as previously explained in 3.1, describes the response of the disordered system to the addition of a novel spin to the network [36]. A nonzero Onsager reaction is a necessary condition for the stability of the system, so the PDF of local fields on local minima of (1) must have a gap: $\eta_{\min} = p > 0$. In mean field spin glasses, the Onsager reaction is exactly the linear susceptibility $p \equiv \chi$, whereas in our VHM (1) we found

$$p = \frac{\alpha\sigma}{d} \frac{\chi}{\frac{d}{\sigma^2} - \chi} \quad (30)$$

where χ is the linear susceptibility defined in equation (25). The expression we obtained in this last equation is an extension to $d > 1$ of that obtained in [40] for the SHM ($d = 1$). Based on these considerations, our theoretical prediction for the distribution of local fields is

$$P_\eta(\eta) = \vartheta(\eta - p) P_h(\eta - p) \quad (31)$$

where $\vartheta(\cdot)$ is Heaviside Theta and $P_h(h)$ is the distribution of cavity fields introduced in (23). We compared the PDFs predicted by equation (31) with the empirical pdfs of local fields of our numerical samples: our results can be found in appendix E.1. We found perfect agreement between theory and numerics for Mattis states minima, but poor agreement in the case of spin glass ones, as predictable from the fact that equation (31) uses the RS solution for $P_h(h)$ and χ .

3.2.3. The spectral density and its pseudogap. We solved the spectral equation (29) obtaining $\mathcal{G}(\lambda)$ and computed the spectral density through Sokhotski—Plemelj formula⁸ as the resolvent imaginary part

$$\rho(\lambda) = \frac{1}{\pi} \Im \mathcal{G}(\lambda - i0_+). \quad (32)$$

In the following, we describe the properties of the spectral density obtained from (32). We briefly discuss the bulk properties and then focus on the lower spectral edge.

The bulk: as explained in appendix F.1, our solution can be approximately decomposed as

$$\rho(\lambda) \simeq \rho_{\text{nt}}(\lambda) \vartheta(\hat{\lambda} - \lambda) + P_\eta \left(\lambda\sigma - \frac{\alpha\sigma}{d} \right) \vartheta(\lambda - \hat{\lambda}). \quad (33)$$

This decomposition results from the properties of the interaction matrix \mathbb{J} in the Hessian in equation (28), which has rank $P < N(d - 1)$ and thus its spectrum for $\alpha < (d - 1)$ has a continuous and a singular part. The 'nontrivial' density ρ_{nt} is given by the free convolution of the continuous part and the local fields distribution, while the trivial term in (33) comes from the free convolution of local fields with the singular part. Therefore, the eigenvalue $\hat{\lambda}$ that marks the crossover between the two bulks corresponds to the rank- P eigenvalue. Since the focus of this section is on low energy excitations, we continue the discussion on the bulk properties in appendix F.1.

⁸ For a real quantity x , it holds $\frac{1}{x \pm i0_+} = \mathcal{P}\left(\frac{1}{x}\right) \pm i\pi\delta(x)$, where $\mathcal{P}(1/x)$ is the principal value integral of $1/x$. Equation (32) for the spectral density can be easily obtained from its definition as the Stieltjes transform of the resolvent function $\mathcal{G}(z) = \int \frac{\rho(\lambda)d\lambda}{\lambda - z}$.

The lower edge: we found that the lower spectral edge of the Hessian spectrum is vanishing in the thermodynamic limit. Indeed, equation (29), that describes the spectrum in the thermodynamic limit, is consistent if and only if the lower spectral edge is $\lambda_{\min} = 0$. We leave a proof of this claim to appendix F.2 for interested readers. We then expanded equation (29) for $\lambda \rightarrow 0$, as reported in appendix F.2, and derived the scaling of the lower tail of the spectral density of Mattis states of $d > 2$ VH models:

$$\rho(\lambda) \approx a_d \lambda^{d-1}, \quad \lambda \ll \lambda_{\text{pg}} \propto \Lambda_0^2. \quad (34)$$

The spectral density follows a power law behavior, up to a scale λ_{pg} controlled by the stability of the RS phase, through the Replicon Λ_0 defined in (27). Slightly above that scale, one has $\rho(\lambda) \simeq \sqrt{\lambda - \lambda_{\text{pg}}}$: the crossover between these two scaling behaviors is an example of a spectral *pseudogap*, qualitatively identical to that observed in vector spin-glass models for glassy excitations in [27, 28]. As a consistency check, note that the global spin glass susceptibility, related to the Hessian spectrum through $\chi_{\text{SG}} = \int d\lambda \rho(\lambda) / \lambda^2$, is finite for any $d > 2$. This is in agreement with our previous observations in section 3.1 that Mattis states are RS and thus should have $\chi_{\text{SG}} \propto 1/\Lambda_0 < \infty$. We did not explicitly derive the scaling for the $d = 2$ VH networks, because they are fRSB and we do not know the correct distribution of local fields necessary to solve the spectral equation. However, by continuity, we claim that $\rho(\lambda) \approx \lambda$ in that case, up to logarithmic correcting factors. In fact, such scaling would return a divergent spin glass susceptibility.

The prefactor a_d (whose expression is reported in appendix F.2) quantifies the abundance of low-energy excitations: indeed, using equation (34) one can see that the fraction of eigenvalues with at most energy λ is $\int_0^\lambda d\lambda' \rho(\lambda') \simeq \frac{a_d}{d-1} \lambda^d$. We found that prefactors a_d have, for the small values of $\alpha = \mathcal{O}(1/d)$ where the retrieval phase exists, exponentially low values, implying that for any reasonable size accessible by numerical simulations, low-energy excitations are quite rare to observe. It follows that soft eigenmodes observed in numerical simulations must be affected by significant finite size effects.

We tested our theoretical predictions by comparing the theoretical spectral density (32) with the empirical one:

$$\rho_{\text{emp}}(\lambda) = \frac{1}{N(d-1)} \sum_{k=1}^{N(d-1)} \overline{\delta(\lambda - \lambda_k)}. \quad (35)$$

In (35) the symbol $\overline{(\cdot)}$ stands for average over numerical samples. We obtained local minima of the energy (1) as fixed points of the retrieval dynamics (5), and obtained their hessian spectra by exact numerical diagonalization. In figure 3 we report plots of the spectral density for the $d = 3$ VH model: the left panel contains the theoretical and numerical spectral densities of Mattis states for some values of capacity $\alpha < \alpha_c(0) \simeq 0.0508$. We found excellent agreement between theory and numerical experiments far from the lower edge, but there are evident finite size effects close to the lower edge. Similar behaviors, which are related to the extreme value statistics of low eigenvalues, are completely analogous to those already observed for vector spin glasses in [27]: therefore, they will not be discussed in this work. In the right panel, we show the spectral densities for the spectrum of spurious spin-glass local minima. We find, as expected, that our solution of the spectral equation (29), which is based on an RS approximation for the theoretical assessment of $P_\eta(\eta)$, produces a wrong result. In particular, close to the lower spectral edge, RS theory predicts a gap that is absent in numerical spectral densities. The correct lower edge behavior for spurious spin glass minima is expected to be $\rho(\lambda) \simeq \sqrt{\lambda}$ (see appendix F.2).

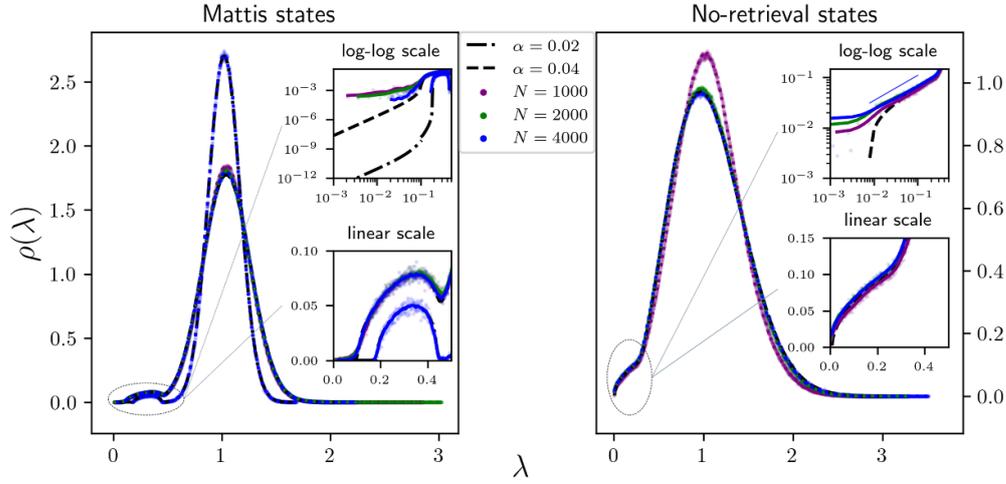


Figure 3. The pseudogap in the spectral density controls the stability of the system. The spectral density obtained from the solution of equation (29) for the $d = 3$ VHM, compared with empirical spectral densities computed from numerical measures. Left: the spectral density for Mattis states. With different dashed lines we show the theoretical predictions for capacities $\alpha = 0.02, 0.04$, while continuous lines of different colors detailed in the legend are numerical measures for sizes $N = 1000, 2000, 4000$. In the inset, we zoom on the non-trivial bulk close to the lower edge of the spectrum. The remainder of the spectral density has a shape that reminds of that of $P_\eta(\eta)$, as one can fully assess by comparing it with figure 12 in appendix E.1. There is a good agreement between theory and measures, with visible finite size effects only at low eigenvalues. In particular, while for capacity $\alpha = 0.02$ the power tail close to $\lambda = 0$ is not captured, for capacity $\alpha = 0.04$ we see that the lower tail of the spectrum progressively approach the asymptotic curve for increasing sizes. Right: same but for spurious spin-glass states. Here we only compare theory and data for $\alpha = 0.04$, to ease graph legibility. One can see systematic errors between the RS theory and numerical measures similar to those found in figure 12 in appendix E.1. In addition, the $\rho(\lambda)$ obtained from (29) using the RS local fields distribution for no-retrieval states has an unphysical gap close to the lower edge: the numerical spectral density instead is clearly un-gapped and seems to approach $\lambda = 0$ as $\rho_{\text{emp}}(\lambda) \sim \sqrt{\lambda}$, consistently with one expects for the spectral density of spin glass states.

Note that the spectral density scales as the PDF of cavity fields (equation (23)) close to the origin: $P_h(h) \approx h^{d-1}$ for small h . This behavior hints at the presence of a strong linear correlation between low-rank eigenvalues and low-rank cavity fields that, as shown in [27] for vector spin glasses and proven rigorously in [46] in the context of the deformed Wigner ensemble of random matrices, is related to the localization of low-rank eigenvectors. In the following paragraph, therefore, we study the localization properties of eigenvectors of our Hessian (28), again singling out their behavior for Mattis states and spin glass states.

3.2.4. Localization of soft eigenvectors. The statistical properties of eigenvectors, in relation to their configuration on the underlying network, carry crucial physical information on how the system responds to small external perturbations. In mean field disordered systems on fully-connected graphs, the dominant local static responses are given by graph nodes where soft modes concentrate their measure. As an example, in terms of hessian eigenmodes, the linear and quadratic local susceptibilities read

$$\chi_{ii} = \sum_{k=1}^N \frac{|\psi_i(\lambda_k)|^2}{\lambda_k} \quad \chi_{ii}^{\text{sg}} = \sum_{k=1}^N \frac{|\psi_i(\lambda_k)|^2}{\lambda_k^2}. \quad (36)$$

Both quantities in this equation are dominated by the softest modes ($\lambda \rightarrow 0$) and the strongest weights ($|\psi_i|^2 = \mathcal{O}(1)$). In equation (36), we identified eigenvectors related to eigenvalues λ with the symbol $\underline{\psi} = (\psi_1^j, \dots, \psi_N^j)^T$, where each ψ_i^j is a vector of dimension d , with the index $i = 1, \dots, N$ identifying network sites and the index $j = 1, \dots, N(d-1)$ identifying the rank of the mode. We consider by definition each eigenvector normalized to unit and adopt this notation for eigenvectors in the following.

In order to measure the localization properties of eigenvectors, we introduce the inverse participation ratio (IPR) $I(\lambda)$ and fixed-rank IPRs (frIPRs) I_j :

$$I(\lambda) = \frac{1}{N(d-1)} \sum_{j=1}^{N(d-1)} \overline{\delta(\lambda - \lambda_j) I_j}$$

$$I_j = \sum_{k=1}^N |\psi_k^j|^4. \quad (37)$$

The frIPRs in (37) are descriptors of the degree of localization of eigenvectors. Because of the normalization constraints $\sum_{k=1}^N |\psi_k^j|^2 = 1$, for any mode the typical value of its eigenvector weights is $|\psi_k^j|^2 = \mathcal{O}(\frac{1}{N})$: thus, when all the weights $|\psi_k^j|^2$ bear such a contribution to normalization one has $I_j = \mathcal{O}(\frac{1}{N})$ and the eigenvector of rank j is said *delocalized* or *extended*. Conversely, if one has $|\psi_k^j|^2 \gg 1/N$ for some sites, then one has either partially delocalized eigenvectors $I_j = \mathcal{O}(\frac{1}{N^a})$ for some $0 < a < 1$ or localized ones $0 < I_j = \mathcal{O}(1) < 1$. The IPR describes localization properties in the thermodynamic limit for fixed eigenvalues: therefore, it is an ideal tool to characterize localization in the bulk of the spectrum. It is convenient to define a rescaled IPR, $\kappa(\lambda) = NI(\lambda)$, where now a finite $\kappa(\lambda)$ identifies delocalized modes and a singular κ , diverging with system size, either partially delocalized ($\kappa = \mathcal{O}(N^a)$, for $0 < a < 1$) or localized modes ($\kappa = \mathcal{O}(N)$).

Equation (37) provide an empirical measure of delocalization. We obtained a theoretical prediction for the IPR in the bulk of the spectrum, in terms of the resolvent function computed from equation (29):

$$\kappa(\lambda) = 3 \left(1 - \frac{1}{d}\right) \left[\frac{(d-1)\alpha}{\sigma^2 \left| \frac{d}{\sigma^2} - \mathcal{G}(\lambda) \right|_{\mathbb{C}}^2} \right]^2 \int \frac{d\eta P_\eta(\eta)}{\left| \eta - \lambda\sigma - \frac{\alpha\sigma}{d} \frac{\mathcal{G}(\lambda)}{\frac{d}{\sigma^2} - \mathcal{G}(\lambda)} \right|_{\mathbb{C}}^4}. \quad (38)$$

Given a solution of equation (29), the IPR can be computed from equation (38) as $I = \kappa/N$. We outline our derivation of equation (38) in appendix F.2.3. We computed the rescaled IPR for several values of α of the $d = 3$ VH model. In what follows, we outline our results for Mattis states and spin glass states.

Mattis states: in the top panel of figure 4 we show on the left (figure 4(a)) the rescaled IPR $\kappa(\lambda)$ as computed from equation (38) and its empirical counterpart $\kappa_{\text{emp}}(\lambda) = NI(\lambda)$, where $I(\lambda)$ is the IPR defined in (37), for capacities and sizes detailed in the legend. There is excellent agreement between our theoretical prediction and numerical data in the bulk of the spectrum, with finite size effects becoming increasingly smaller for increasing sizes N . However, we observe that, close to the edges of the spectrum, numerical data do not agree with the theoretical prediction: instead, $\kappa_{\text{emp}}(\lambda)$ seems to increase roughly linearly with N . This observation is a

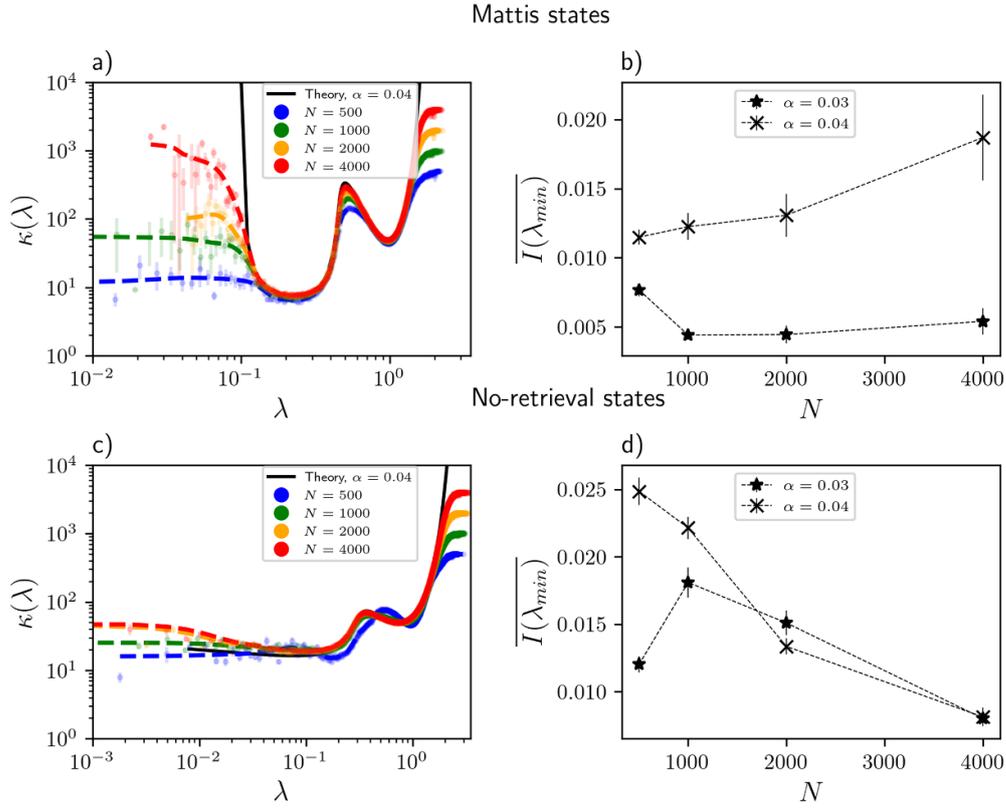


Figure 4. Soft modes of memory states are localized. The rescaled IPR $\kappa = NI$ of the $d=3$ VHM, for Mattis states (top panel) and Spin glass states (bottom panel). Top: figure (a) compares the theoretical prediction (38) with empirical measures based on equation (37). In the bulk there is an excellent agreement between theory and numerics, while at the spectral edges deviations consistent with eigenvectors localization appear. Figure (b) shows the sample-averaged IPR of the softest mode for several values of α detailed in the legend. For increasing sizes, the IPR is not vanishing, validating the localization hypothesis. Bottom: same but for no-retrieval states. Figure (c) shows the comparison between theoretical rescaled IPR and numerics: again, the RS theory is not correct since it predicts a non-physical spectral gap. The empirical rescaled IPR seems to be regular for $\lambda \rightarrow 0$. In figure (d) we show the sample-averaged IPR of the smallest mode: this quantity is decreasing for increasing sizes in a range $N = \mathcal{O}(10^2) : \mathcal{O}(10^3)$.

clear sign of localization for $\lambda \rightarrow 0$. To corroborate our claim, in figure 4(b) on the right we show the sample-averaged IPR of the first eigenvector versus N for capacities $\alpha = 0.03, 0.04$. In all cases, the IPR is not vanishing for increasing N , possibly saturating to a constant value for $N \rightarrow \infty$.

The type of localization occurring in the soft modes of Mattis states is a random matrix condensation phenomenon, similar to that observed in [27, 28, 47] for mean field disordered systems or [46, 48] for more general deformed Wigner matrices: for each low-rank (with rank k) eigenvector, a single component bears a finite contribution to normalization, while each of its remaining components are vanishing in the thermodynamic limit. In formulas, we obtain for the IPR at the lower edge, in the thermodynamic limit

$$I(0) = \Lambda_0^2 + O\left(N^{-1+\frac{2}{d}}\right) \quad (39)$$

so that the condensation phenomenon is tied to the stability of the RS phase ($\Lambda_0 > 0$), as already shown in [27]. The finite-size corrections to the asymptotic value in (39) in the case $d = 3$ analyzed by us are not negligible, being $O(N^{-\frac{1}{3}})$, thus hindering the condensation phenomenon, as previously described, up to very large sizes, inaccessible by numerical simulations. We further discuss the condensation phenomenon in appendix F.2.

Spin glass no-retrieval states: we show our results in the bottom panel of figure 4. Here, as expected, our RS theoretical prediction of the rescaled IPR leads to incorrect results, especially at low eigenvalues: numerical data clearly show that the spectrum is not gapped. The rescaled IPR appears to be regular in the limit $\lambda \rightarrow 0$, as is apparent from the progressive collapse of data curves as the size N of the system increases. It follows that soft modes are delocalized, as it should be in a spin glass phase.

In this paragraph, we showed that Mattis states representing memories stored in the network feature localized soft modes. In the next paragraph, we unveil the physical meaning of such a phenomenon in this model, discussing the relation between eigenvectors, local fields, and noise in Mattis states.

3.2.5. Relation between noise and low energy modes. The relation between the rescaled square components $u_{i,k} \equiv N|\psi_{i,k}^2|$ of rank- k eigenvectors and local fields η_i is the following: the average $u_{i,k}$ for fixed η_i reads

$$\langle u_{i,k} \rangle = \frac{\alpha/\sigma^2}{\left|\frac{d}{\sigma^2} - \mathcal{G}(\lambda_k)\right|_{\mathbb{C}}^2} \frac{d-1}{\left|\eta_i - \lambda_k\sigma - \frac{\alpha\sigma}{d} \frac{\mathcal{G}(\lambda_k)}{\frac{d}{\sigma^2} - \mathcal{G}(\lambda_k)}\right|_{\mathbb{C}}^2}. \quad (40)$$

Indeed, the components $\psi_{i,k}^a$ of any eigenvectors are in the thermodynamic limit Gaussian variables with variance given by (40). Equation (40) suggests that the components of the eigenvectors and local fields are anticorrelated. In order to characterize the distribution of the noise that the network experience in the retrieval of a pattern ξ^1 given its corresponding Mattis state \mathbf{S} , we introduce the local noises

$$\ell_i = \frac{1}{2} - \frac{\xi_i^1 \cdot \mathbf{S}_i}{2\sigma}. \quad (41)$$

Local noises are bounded in the unit interval $0 \leq \ell_i \leq 1$: spins highly correlated with pattern spins ($\frac{1}{\sigma}\mathbf{S}_i \cdot \xi_i^1 \simeq 1$) have very low local noise, while those weakly correlated ($\frac{1}{\sigma}\mathbf{S}_i \cdot \xi_i^1 \simeq 0$) have local noise close to $1/2$; finally, spins negatively correlated with pattern spins have $1/2 < \ell_i \leq 1$. The average of local noises is simply connected to Mattis overlap in equation (25): $\bar{\ell} = \frac{1-m}{2}$. The local fields η_i can be simply related to local noises: by using the stationarity condition in equation (3), one has

$$\begin{aligned} \eta_i &= \frac{1}{\sigma} \sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j \cdot \mathbf{S}_i \simeq (1 - 2\ell_i) \sigma m + \mathbf{n}_i \cdot \mathbf{S}_i \\ \mathbf{n}_i &= \frac{1}{\sqrt{N}} \sum_{\mu=2}^P \left(\frac{1}{\sqrt{N}} \sum_{j:j \neq i} \xi_j^\mu \cdot \mathbf{S}_j \right) \xi_i^\mu. \end{aligned} \quad (42)$$

The term $\mathbf{n}_i \cdot \mathbf{S}_i$ represents the local field in absence of pattern $\mu = 1$: the vector \mathbf{n}_i is in fact an isotropic Gaussian vector with variance αr (compare with equation (9)). Equation (42)

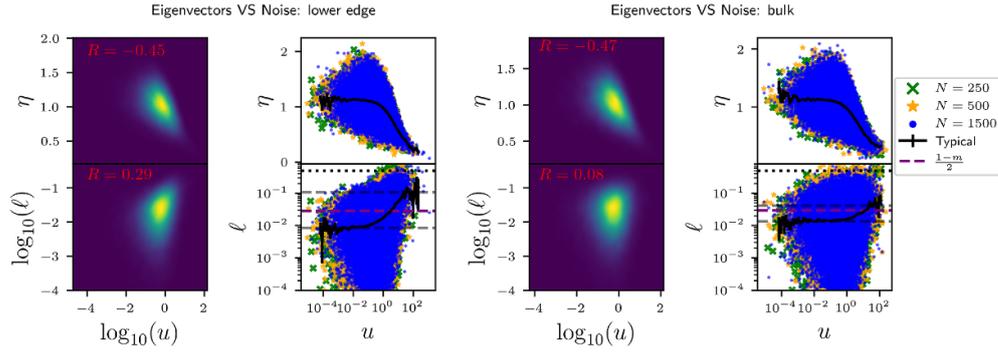


Figure 5. Soft modes of Mattis states localize on noisy spins. In this figure we show our measures of correlation for respectively local fields η_i and rescaled weights $u_i = N|\psi_i|^2$ and local noises ℓ_i and rescaled weights, for the $d = 3$ VHM. We show heatmaps in panels (a) and (c) and scatter plots in figures (b) and (d). Superimposed to scatter plots we show with dark continuous lines typical curves of η and ℓ at fixed u , computed respectively as $\eta_{typ}(u) = \exp \frac{1}{N} \sum_i \log(\eta_i) \delta(u - u_i)$ and $\ell_{typ}(u) = \exp \frac{1}{N} \sum_i \log(\ell_i) \delta(u - u_i)$. The dashed purple line marks the position of the sample-averaged local noise $\bar{\ell} = \frac{1-m}{2}$, while black dashed lines mark respectively the empirical lower and upper bounds of $\tilde{\ell}_{typ}$. Sizes are $N = 250, 500, 1500$, capacity is $\alpha = 0.04$. We considered a number of samples per size such that $N_s = 1500 \times \frac{100}{N}$.

suggests that local fields η_i and local noises ℓ_i should be negatively correlated, and consequently thanks to equation (40) that local noises and rescaled weights u_i are positively correlated. The vectors \mathbf{n}_i are correlated to local noises in a non-trivial way, thus making the relation between local fields, weights and local noises hardly accessible through analytical techniques.

We claim that the correlation of local noises ℓ_i with rescaled weights u_i , or equivalently, the anticorrelation of local noises ℓ_i with local fields η_i , is stronger in lower edge modes than in bulk ones. We expect nodes of localization of soft modes to pinpoint spins of the Mattis state with high levels of noise ($\ell_i \approx 1/2$) with respect to the respective pattern spins. The reasoning behind our conjecture is the following: in mean-field disordered systems, localization is induced by extremal values of local fields, as it happens for Anderson localization and related models [49, 50] or for soft modes of mean field models of glassy systems, such as the models studied in [27, 28, 47] or the present model. In the latter case, spins with very low local fields are then on nodes where soft modes tend to localize in the thermodynamic limit. Since spins of the Mattis states well-aligned with their respective patterns spins will have on average local fields $\eta \approx 1$ (compare with equation (5) with a spin $S_i \simeq \sigma \xi_i^1$), it follows that local fields $\eta_i \ll 1$ may be related to noisy spins in the Mattis state.

We tested this conjecture numerically, by measuring the Spearman correlation of local fields and local noises with the weights of the first eigenvector and the correlation of local fields and local noises with the weights of the eigenvector of rank $k = P/2 \gg 1$. We considered several samples of the $d = 3$ VHM at capacity $\alpha = 0.04$, for sizes $N = 250, 500, 1500$. We report our findings in figure 5. Figures in the top panel show that local fields are anti-correlated to eigenvector weights, with almost the same correlation for lower edge and bulk modes. The bottom panel discusses the correlation between local noises ℓ_i and rescaled weights u_i : it appears that the correlation between local noises and weights is stronger for lower edge modes ($R = 0.29$) than for bulk ones ($R = 0.08$).

Spins with high noise in the memory states can also have consequences on the computational costs of retrieval, as we discuss in appendix D. Specifically, we study the correlation between the number of iterations needed for convergence of (5) and the Mattis magnetization of the final state. We study this dependence also locally, measuring the correlation between local noises and local convergence times, i.e. the different number of iterations necessary for the different spins to align to the desired accuracy. We found that, as expected, noisy spins are also slow spins. Computationally, the number of iterations needed to align the slowest spins seems to asymptotically have a very weak dependence on system size.

4. The network above saturation

In this section we study the network in the saturated regime $\alpha > \alpha_c$. No retrieval states exist as fixed point of equation (5). We thus concentrate on the out-of-equilibrium properties of the system, focusing on the transient behavior at the very first steps of the numerical retrieval dynamics. We study numerically the first-step, zero-temperature dynamics of the system and compare our results with theoretical predictions from a signal-to-noise (STN) analysis of equation (5): here, our analysis spans a wide range of capacities below and above saturation, considering both Mattis and mixture states as initial conditions. For numerical simulations of dynamics we use always systems with $N = 1000$ because finite size effects are already negligible at this system size for the phenomenology studied, so that using values of $N = \{250, 500, 1000, 2000\}$ results in overlapping curves.

4.1. Transient denoising phenomenon

4.1.1. Mattis states above saturation. We studied Mattis states above saturation ($\alpha > \alpha_c$), out of the storage phase: here the dynamics initialized at a given pattern converges to a spin glass state with vanishing correlation with the same pattern.

We observed that, starting from a magnetization $0 < m_0 < 1$ and small enough, m_t has a peak at $t = 1$, with $m_1 > m_0$. So, at the very first step the model has a denoising capability, even though the dynamics have not fixed points strongly correlated with examples. We define this phenomenon as *transient* or *first-step* denoising: it appears that this phenomenon is the stronger the larger d . In figure 6(a) we report an example of dynamics at capacity $\alpha = 0.5 > \alpha_c(d = 1) \simeq 0.138$. In section 4.2 we study this phenomenon more in depth.

4.1.2. Mixtures states. In section 3.1 we stated that there is no stable RS solution involving mixture of examples and in appendix B we thoroughly discuss the origin of such instability. Our theoretical prediction is restricted to the case $\alpha = 0$ and to symmetric mixtures: here we show numerically that mixture solutions are unstable also for $\alpha > 0$.

In figure 6(b) we show our results for dynamical runs starting from the symmetric mixture $\mathbf{S}_{\text{ini}}(0) = \frac{(\xi_i^1 + \xi_i^2 + \xi_i^3)\sigma}{|\xi_i^1 + \xi_i^2 + \xi_i^3|}$ and from an initial condition with overlap $m_{\text{ini}} = 0.5$ with the same mixture. We chose $\alpha < \alpha_c(d = 16) \approx 0.009$ in such a way there exist fixed points correlated to examples. In figure 6(b) we show the magnetization $m_t = \frac{1}{N\sigma} \mathbf{S}(0) \cdot \mathbf{S}(t)$ during the dynamics. The dynamics always converges to a Mattis state related to one of the examples composing the mixture, therefore no mixtures stable solutions exist in VHMs. However a transient retrieval ability of the network happens also for mixtures states: at the very first step the dynamics is attracted by the mixture, but at later times it ends in one of the states composing the mixture.

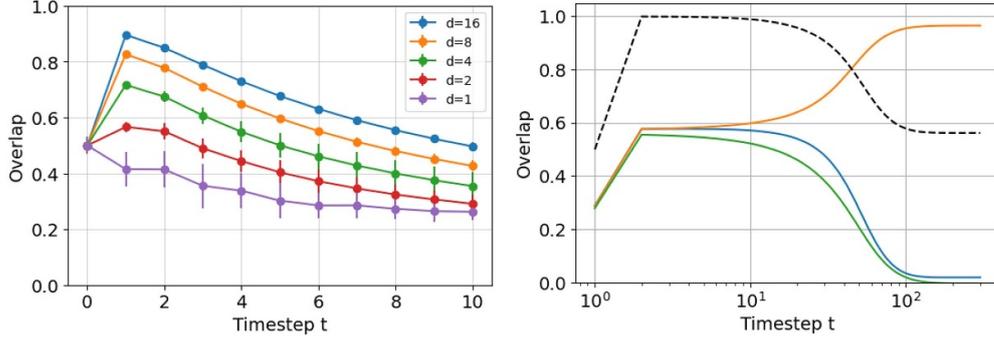


Figure 6. Transient denoising: the network can denoise examples and mixtures even above the storage capacity. Panel (a) shows the overlap with an examples as the function of (synchronous) time steps, for various dimensions of the spins d . Panel (b) shows the overlap with a mixture of examples (black dashed line) as the function of (synchronous) time steps at $d = 16$. The solid lines correspond to the overlaps with examples within the mixture.

4.2. STN analysis at first step

We study the first-step, zero-temperature dynamics of the system using the STN analysis of equation (5). We consider the above saturation regime and noisy examples as initial condition. However the same computation holds also in the case of near saturation regime, and in appendix G we discuss the case of mixtures as initial conditions.

4.2.1. Mattis states. We perform the STN analysis at the first step of the zero temperature dynamics of the model. The relevant parameters are spins dimension d , relative number of examples α and initial correlation m_0 with one of them. We choose as initial condition for the dynamics $\frac{1}{\sigma} S_i(0) = m_0 \xi_i^1 + \sqrt{1 - m_0^2} \pi_i$ where π_i are unit vectors orthogonal to examples spins, $\pi_i \cdot \xi_i^1 = 0$. The initial overlap m_0 is a control parameter to quantify the initial amount of noise with respect to pattern ξ^1 . The local fields vectors at the first step of the dynamics read

$$\frac{1}{\sigma} \sum_{j:j \neq i} J_{ij} S_j(0) = m_0 \xi_i^1 + \mathbf{n}_i. \tag{43}$$

The noise vectors \mathbf{n}_i is an isotropic Gaussian vector with zero mean and variance $\frac{\alpha}{d^2}$.

Thus, the first step Mattis magnetization is

$$\begin{aligned} m_1 &= \frac{1}{N\sigma} \sum_{i=1}^N \xi_i^1 \cdot S_i(1) = \int \frac{d\mathbf{n}}{[(2\frac{\pi}{d})(\frac{\alpha}{d})]^{d/2}} e^{-\frac{n^2 d}{2(\alpha/d)}} \frac{m_0 + \mathbf{n} \cdot \xi}{|m_0 + \mathbf{n} \cdot \xi|} \\ &= \int_0^\infty \frac{dn n^{d-1} e^{-\frac{n^2}{2} - \frac{m_0^2 d}{2(\alpha/d)}}}{(2\pi)^{d/2}} \times \mathcal{K}_d \left(\frac{m_0 n \sqrt{d}}{\sqrt{\alpha/d}} \right) \mathfrak{g}_d \left(\frac{m_0 n \sqrt{d}}{\sqrt{\alpha/d}} \right) \\ &= \Phi_d \left(\frac{m_0}{\sqrt{\alpha/d}} \right). \end{aligned} \tag{44}$$

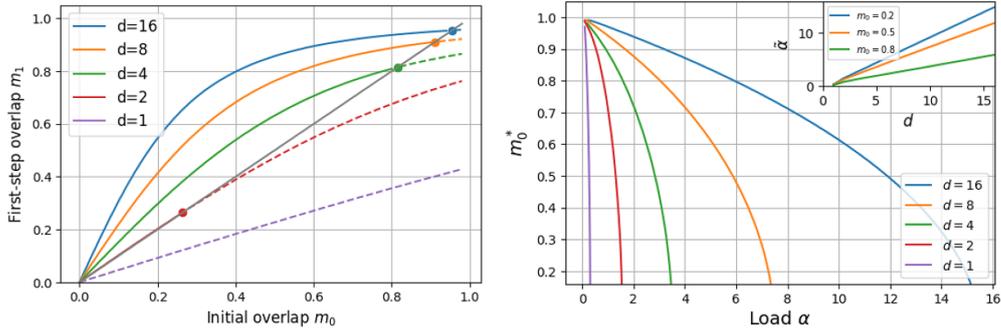


Figure 7. The denoising performance improves with the dimension of the spins. Panel (a) We show the retrieval maps of examples for various d . We also show the line $m_1 = m_0$ (gray), to highlight where the model denoises effectively (solid lines, $m_1 > m_0$) and where it does not (dashed lines, $m_1 < m_0$). Panel (b) the curves show the value of m_0 up to which $m_1 > m_0$ in the retrieval maps, as a function of the load $\alpha = P/N$. Higher values of m_0 mean better denoising.

In figure 13(a) the analytical curves are obtained substituting in equation (44) the values of interest, and they correctly predict the numerical simulations. In appendix G we show more details on STN method and we explain how to deal with first step denoising of mixture states. We introduced the function $\Phi_d(\cdot)$ to express in a more compact way the dependence of m_1 on m_0, α, d . We note that m_1 depends only on $\frac{m_0}{\sqrt{\alpha/d}}$. In the special case $d = 1$ of the SHM one can show that $\Phi_1(\cdot) = \text{Erf}\left(\frac{\cdot}{\sqrt{2}}\right)$. It is also worth to consider the large d limit of (44): we found that with the scaling $\alpha' = \alpha/d$ the first-step magnetization m_1 converges to a non-trivial function

$$\lim_{d \rightarrow \infty} \Phi_d\left(\frac{m_0}{\sqrt{\alpha'}}\right) = \frac{m_0/\sqrt{\alpha'}}{\sqrt{1 + (m_0/\sqrt{\alpha'})^2}}. \quad (45)$$

We define $x = m_0/\sqrt{\alpha'}$ and in figure 13(b) we plot curves $y = \Phi_d(x)$ for different values of d , showing that the curves collapse rapidly to the $d \rightarrow \infty$ case. It means that the dependence on d of equation (44) is mainly in its argument rather than in Φ_d itself. Consequently, to compute first step value m_1 all parameters of the system m_0, α, d are relevant only through the combination x .

In figure 7(a) we show retrieval maps at different d . The transient denoising, when present, occurs for values of $m_0 \in (0, m_0^*)$. This is the opposite of standard retrieval where the system is capable of magnetize a pattern if $m_0 > \tilde{m}_0$. In figure 7(b) we show $m_0^*(\alpha)$ for various d . We obtain that transient denoising at a given d can happen up to a certain α^* , with such α^* increasing with d . The same is true not only for $m_0^* = 0$, but for all values k of $m_0^* = k$. In particular α^* increases linearly in d , with a coefficient depending on k .

5. Discussion

Motivated by understanding the phenomenology of neural networks with vector variables, in this work we studied a classical Hopfield network with continuous vector spins. We first characterized the static properties of the system below saturation, solving the model in an RS

ansatz and computing its phase diagram for varying storage capacity and temperature, identifying retrieval and spin glass lines. Secondly, we focused on the physics at $T = 0$ and studied the low-energy excitations of the system, characterized by the soft modes of the Hessian matrix associated with the energy function. Finally, we studied the dynamical properties of the system, for patterns and mixture states, focusing on the behavior of the network in the early stages of the dynamics.

In the following, we discuss separately our two main results: i) Vector Hopfield networks trade off the classic retrieval of patterns, where memory states highly correlated with patterns are stored as local minima of the energy (static retrieval), for the ability of denoising patterns transiently [26], in one dynamical step (first-step retrieval). The critical capacity for static retrieval shrinks with spin dimension, $\alpha_c \propto 1/d$, but the critical capacity associated with first-step retrieval grows, $\tilde{\alpha} \propto d$. ii) Continuous associative memories with symmetric interactions, such as the VHM of this work, are equivalent to mean-field models of low-temperature glasses, in the behavior of the low-energy excitations of their local minima [27, 28]: minima related to memory states can be mapped to stable glasses, with soft modes localized on the noisiest spins of the memory state, while minima of spurious states are linked to marginal glasses, having delocalized and featureless soft modes.

Below we also discuss the possible future developments for each of these two results, which should be considered separately from the main development suggested in this work, that is to introduce vector spins in modern Hopfield networks.

(a) First-step retrieval: The phenomenon of transient denoising (in our model first-step retrieval) has recently been studied in great detail in [26] for very general continuous associative memory networks, through a dynamical mean field theory approach. The author interprets the phenomenon of transient retrieval as a geometric memory effect of the energy landscape in the vicinity of patterns, with attractive features for the dynamics that survive well beyond capacity saturation. We believe that such an explanation holds also for our model, even though at present we do not know why the first-step retrieval phase grows with spin dimension d : we reserve the test of the geometric memory hypothesis and the characterization of its dependence on spin dimension to future research. We point out that the landscape of our model in the vicinity of any memory state does not change sharply at saturation: indeed, since the Replicon eigenvalue in equation (27) remains finite at the critical load $\alpha_c = \mathcal{O}(1/d)$ for any d , the pseudogap width of the spectral density in equation (34) remains finite, and therefore there are no changes in the statistics of low energy modes at saturation. We expect that stable memory states are destabilized smoothly, with the appearance at saturation of a single negative outlier eigenvalue, together with a non-zero gradient direction⁹.

The author in [26] suggests that biological memory circuits may take advantage of dynamical transients to perform memory retrieval. In fact, transients may be fundamental for neural computations [51]. Dynamical denoising and its trade-off with static storage is also relevant for current applications in ML, considering the relation observed recently between the self-attention mechanism in Transformers [9] and modern Hopfield networks [5, 6, 13]. In such ML architectures, the computation of neuron activations through Self-Attention can be seen as a single-step update for the retrieval dynamics of an exponential Hopfield network [6]. In [13] transient denoising is observed for Self-Attention layers of Visual Transformers, in a simplified setting where the image denoising task is performed by repeating a single layer

⁹ Since there is no stationary (but unstable) retrieval solution of the thermodynamics beyond saturation, we do not expect that memory states become saddle points of the landscape.

of Attention. Moreover we highlight the relation between denoising and most common ML tasks such as autoregression or classification in Transformer models: special mask vectors in the input have to be mapped to correct labels, in a denoising task where all of the noise is in that mask vectors while all other input vectors have not been corrupted. Interestingly, a simple scaling argument shows that state-of-the-art large language models (LLMs) have typical parameters compatible with that phenomenon in our model. In a typical pre-trained Transformer for language modeling [52, 53], the order of magnitudes are $P \sim 10^8$ training sequences $N \sim 10^3$ tokens of $d \sim 10^4$ dimensions. With those values, the order parameter for denoising assumes a reasonable value $\alpha' = \frac{P}{Nd} \sim 10$, while the memorization order parameter assumes an unreasonable scale $\alpha'' = \frac{Pd}{N} \sim 10^9$, which would rule out memorization phenomena. While LLMs have many differences with respect to a single update of a vector Hopfield network, it would be interesting to observe if transient denoising still plays a role in the performances of those models (as suggested in [13]), especially because the task of masked token prediction is similar to aligning a chosen vector spin to the local field produced by the other spins. If this is the case, we note that the high dimensionality of variables itself could produce one-step retrieval of a large number of training data (proportional to the vector dimension as $P \sim Nd$), without the need for the higher-order interactions of modern Hopfield networks.

There are several interesting follow-ups that may be worth exploring to better understand first-step retrieval in Vector Hopfield networks.

- First, a necessary development of our research would be to go beyond random patterns and consider hidden-manifold models for their generation process, such as those studied in [15, 35] for classical Hopfield networks with Ising spins. Since we showed that Vector Hopfield networks can perform first-step retrieval also for mixtures of patterns, transient denoising in VHMs with structured data could be possible also for mixtures of hidden features, allowing for retrieval and generation of patterns up to very large capacities.
- It would be interesting to extend the pseudo-likelihood approach to associative memories of [54] to vector variables. Couplings tensor \mathbb{J} would be obtained from minimization of pseudo-likelihood cost instead of being fixed with Hebb's rule. A interesting question is whether or not scalings with d of equilibrium fixed points and transient denoising would be the same.
- Another perspective could be to check if first-step retrieval of VHMs is enhanced or weakened by algorithms that increase the capacity of Hopfield networks, such as daydreaming algorithms [55] or Hebbian unlearning [56].
- Finally, there is also the possibility of studying first-step retrieval in the case of hetero-associative Hopfield networks [57, 58], generalized to vector spins. These networks can store sequences of patterns, allowing for more complex computations than those performed by ordinary Hopfield networks.

(b) Vector Hopfield memories as stable glasses: The nature of low-energy excitations of glasses is still debated: recent years have seen many numerical efforts to understand the density of states (DoS) of amorphous through computer glass models, being still unclear whether the DoS exhibits quasi-localized modes with a universal quartic scaling at vanishing frequencies, beyond the lowest phonon excitation [59, 60], or rather if the universal quartic DoS scaling law is an artifact of numerical simulations [61]. Mean-field analytical approaches on phenomenological [47] and first-principle [27, 28] glassy models have shown that key features of glassy soft modes can be reproduced by means of random matrix theory: these models have Hessians of their respective energy functions that, in the thermodynamic limit, are instances of a deformed Wigner ensemble, where a Gaussian matrix with vanishing $\mathcal{O}(1/\sqrt{N})$ entries (interaction matrix or second derivatives of the energy) is perturbed by a diagonal matrix with

heterogeneous $\mathcal{O}(1)$ random entries (local fields or stiffness parameters). From the knowledge of the distribution of diagonal entries, by using free convolution of random matrices [45] it is possible to evaluate the spectrum of the deformed Wigner matrix. Depending on the relative strength of the diagonal matrix with respect to the off-diagonal one, rigorous results [46] predict the existence of localization for eigenmodes of deformed Wigner ensembles close to the spectral edges, provided that the distribution of the diagonal elements satisfies specific scaling conditions close to the edges of their support. The type of localization occurring in these models, which are defined on densely connected networks, is a condensation phenomenon, where edge eigenvectors concentrate their measure on single nodes, following a mathematical mechanism reminiscent of Bose–Einstein condensation [48]. Note that this type of localization is different from Anderson localization [49] and related models [50, 62], where bulk modes at the center of the spectrum undergo a localization transition driven by strong spatial disorder, and also from Lifshitz tails in the spectrum of the adjacency matrices of sparse locally tree-like heterogeneous networks [63, 64], which are induced by anomalous fluctuations in the degrees of the nodes.

The Hessian matrix (28) studied in this work in the thermodynamic limit is a random matrix of a deformed Wishart ensemble, where a Wishart-like matrix (the interaction matrix in (28)) is added to a diagonal matrix (the matrix of local fields in (28)). To our knowledge, this is the first work where such a random matrix ensemble is studied in relation to a physical model. The observation that memory states behave in their low-energy excitations as stable mean-field glasses corroborates the link between continuous associative memories and glassy models. It is important to stress that localization and its physical meaning in the present model are an effect of the heterogeneity of the Lagrange multipliers in (3), i.e. of the local fields: in the present model, this property is a consequence of the local constraints on the norm of neuron spins, but for models with unconstrained soft neuron spins, such as that studied in [65], heterogeneity can be achieved through local site potentials, thus our results should be valid for several continuous models of associative memories. Note that in spherical models, such as that studied in [44], the Hessian diagonal is homogeneous, therefore the rich phenomenology unveiled in this work is absent. Moreover, our results in appendix F.2.4 show that the nontrivial properties observed for the lower edge spectrum of the Hessian are lost in the high-dimensional $d \rightarrow \infty$ limit, therefore they are not relevant for applications in Transformers, where the dimension of token vectors is very large, but may be relevant for activations in neural network layers, as shown in [24], where the authors use Kuramoto Oscillators (two-dimensional vector spins) to perform computations, instead of customary threshold units. There are several possible developments for our research, relatively to the properties of low-energy excitations of finite- d VHMs.

- It would be interesting to consider how diluting the interaction matrix would affect the stability properties of the system. Neurons with few connections may be more sensitive to noise, possibly affecting the features of soft modes. There could be phases with concentrated but delocalized soft modes, as reported in [66].
- VHMs could be studied for dynamical models with non-symmetric interactions, as attractor associative memories in recurrent neural networks. It would be insightful to see how our findings on linear stability extend to attractor models.
- Higher-order interactions of modern Hopfield networks [3] could enrich the phenomenology of low energy modes.

Finally, vector spins may be relevant for neuroscience applications. Potts spins, i.e. discrete degrees of freedom with multiple states available [67], have been used to represent the state of groups of binary neurons [68, 69]. Vector spins can be seen as a continuous version of Potts

spins (notice the tensor coupling common between the two), and therefore they may be useful to represent the state of groups of continuous-valued neurons—for instance, graded response neurons [65].

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: https://github.com/Francill99/vector_hopfield.git.

Acknowledgments

The authors thank Federico Ricci-Tersenghi for insightful discussions; Elena Agliari, Dario Bocchi and Alberto Fachechi for feedbacks. MN acknowledges the financial support of PNRR MUR Project PE0000013-FAIR. This study was conducted using the DARIAH HPC-AI cluster at CNR-NANOTEC in Lecce, funded by the “MUR PON Ricerca e Innovazione 2014-2020” project, code PIR01_00022 and H2IOSC Project - Humanities and cultural Heritage Italian Open Science Cloud funded by the European Union – NextGenerationEU – NRRP M4C2 - Project code IR0000029.

Author contributions

Flavio Nicoletti  0000-0003-3962-418X

Conceptualization (equal), Data curation (equal), Formal analysis (lead), Investigation (lead), Methodology (lead), Project administration (equal), Software (equal), Supervision (equal), Validation (equal), Visualization (equal), Writing – original draft (lead), Writing – review & editing (equal)

Francesco D’Amico  0009-0006-0777-2160

Conceptualization (equal), Data curation (equal), Formal analysis (supporting), Investigation (equal), Methodology (supporting), Project administration (equal), Software (lead), Writing – original draft (equal)

Matteo Negri  0000-0002-8941-0705

Conceptualization (equal), Funding acquisition (lead), Project administration (lead), Supervision (lead), Validation (equal), Writing – original draft (equal), Writing – review & editing (equal)

Appendix A. Replica computation of the free energy

The equilibrium free energy is evaluated using the replica trick

$$f(\alpha, \beta) = - \lim_{N \rightarrow \infty} \frac{1}{N\beta} \overline{\log \mathcal{Z}(\alpha, \beta)} = - \lim_{N \rightarrow \infty} \frac{1}{N\beta} \lim_{n \rightarrow 0} \frac{\overline{\mathcal{Z}^n} - 1}{n} \quad (\text{A1})$$

which allows one to exchange the disorder average of the log with that of the n th moments of \mathcal{Z} . If one assumes n to be an integer, we can introduce n fictitious replica of the original system and perform the average over the disorder on this system. We have

$$\begin{aligned}
 \overline{Z}^n &= \int D\mathbf{S}_1 \dots \int d\mathbf{S}_n e^{\frac{N\beta}{2} \sum_{\mu=1}^s \left(\frac{1}{N} \sum_{j=1}^N \mathbf{S}_j^a \cdot \boldsymbol{\xi}_j^\mu \right)^2 + \frac{\beta}{2} \sum_{\mu=s+1}^p \left(\frac{1}{\sqrt{N}} \sum_{j=1}^N \mathbf{S}_j^a \cdot \boldsymbol{\xi}_j^\mu \right)^2 - \frac{\beta}{2N} \sum_{\mu=1}^p \sum_{j=1}^N (\mathbf{S}_j^a \cdot \boldsymbol{\xi}_j^\mu)^2} \\
 &\propto e^{-\frac{\beta p n \sigma^2}{2d}} \int D\mathbf{m} D\hat{\mathbf{m}} \exp \left(-iN \sum_{\mu \leq s} \sum_a m_\mu^a \hat{m}_\mu^a + \frac{\beta N \sigma^2}{2} \sum_{\mu \leq s} \sum_a (m_\mu^a)^2 \right. \\
 &\quad \left. -i \sum_{\mu > s} \sum_a m_\mu^a \hat{m}_\mu^a + \frac{\beta}{2} \sum_{\mu > s} \sum_a (m_\mu^a)^2 \right) \\
 &\quad \times \int \prod_a D\mathbf{S}_a \int \prod_{\mu} \prod_{k=1}^N [d\boldsymbol{\xi}_k^\mu P(\boldsymbol{\xi}_k^\mu)] \exp \left(\frac{i}{\sigma} \sum_{\mu \leq s} \sum_a \hat{m}_\mu^a \sum_k \boldsymbol{\xi}_k^a \cdot \mathbf{S}_k^a \right. \\
 &\quad \left. + \frac{i}{\sqrt{N}} \sum_{\mu > s} \sum_a \hat{m}_\mu^a \sum_k \boldsymbol{\xi}_k^a \cdot \mathbf{S}_k^a \right) \tag{A2}
 \end{aligned}$$

In the last equation we do not show prefactors that are not relevant for the final saddle point evaluation. We introduced magnetizations order parameters through delta functions:

$$m_{\mu,a} = \begin{cases} \frac{1}{N\sigma} \sum_{k=1}^N \boldsymbol{\xi}_k^\mu \cdot \mathbf{S}_k^a & \mu \leq s \\ \frac{1}{\sqrt{N}} \sum_{k=1}^N \boldsymbol{\xi}_k^\mu \cdot \mathbf{S}_k^a & \mu > s. \end{cases}$$

The disorder average over noncondensated magnetizations reads

$$\begin{aligned}
 &\int \prod_{\mu > s} \prod_{k=1}^N [d\boldsymbol{\xi}_k^\mu P(\boldsymbol{\xi}_k^\mu)] \exp \left(\frac{i}{\sqrt{N}} \sum_{\mu > s} \sum_a \hat{m}_\mu^a \sum_k \boldsymbol{\xi}_k^a \cdot \mathbf{S}_k^a \right) \\
 &\simeq \exp \left(-\frac{1}{2dN} \sum_{\mu > s} \sum_{ab} \hat{m}_{\mu,a} \hat{m}_{\mu,b} \sum_k \mathbf{S}_k^a \cdot \mathbf{S}_k^b \right). \tag{A3}
 \end{aligned}$$

Next, we introduce the spin glass overlap matrix

$$Q_{ab} = \frac{1}{N\sigma^2} \sum_{k=1}^N \mathbf{S}_k^a \cdot \mathbf{S}_k^b \tag{A4}$$

through delta functions. After the last passages, the replicated partition function reads

$$\begin{aligned}
 \overline{Z}^n &\propto e^{-\frac{\beta p n \sigma^2}{2d}} \int D\mathbf{m}_{\mu \leq s} D\hat{\mathbf{m}}_{\mu \leq s} D\mathbf{Q} D\hat{\mathbf{Q}} \\
 &\quad \times \exp \left(-iN \sum_{\mu \leq s} \sum_a m_\mu^a \hat{m}_\mu^a + \frac{\beta N \sigma^2}{2} \sum_{\mu \leq s} \sum_a (m_\mu^a)^2 + iN \sum_{ab} \hat{Q}_{ab} Q_{ab} \right) \int \prod_a D\mathbf{S}_a \\
 &\quad \times \int \prod_{\mu \leq s} \prod_{k=1}^N [d\boldsymbol{\xi}_k^\mu P(\boldsymbol{\xi}_k^\mu)] \exp \left(\frac{i}{\sigma} \sum_{\mu \leq s} \sum_a \hat{m}_\mu^a \sum_k \boldsymbol{\xi}_k^a \cdot \mathbf{S}_k^a - \frac{i}{\sigma^2} \sum_{ab} \hat{Q}_{ab} \sum_{k=1}^N \mathbf{S}_k^a \cdot \mathbf{S}_k^b \right) \\
 &\quad \times \int D\mathbf{m}_{\mu > s} \frac{D\hat{\mathbf{m}}_{\mu > s}}{(2\pi)^{P-s}} \exp \left(-\frac{\sigma^2}{2d} \sum_{\mu > s} \sum_{ab} Q_{ab} \hat{m}_\mu^a \hat{m}_\mu^b - i \sum_{\mu > l} \sum_a m_\mu^a \hat{m}_\mu^a + \frac{\beta}{2} \sum_{\mu > s} \sum_a (m_\mu^a)^2 \right). \tag{A5}
 \end{aligned}$$

The integrals over noncondensated magnetizations $\{m_{\mu,a}\}_{\mu>s}$ and their conjugated fields $\{\hat{m}_{\mu,a}\}_{\mu>s}$, after the integration over disorder in (A3), are Gaussian integrals and thus can be performed straightforwardly

$$\begin{aligned}
 & \int D\mathbf{m}_{\mu>s} \exp\left(\frac{\beta}{2} \sum_{\mu>s} \sum_a (m_{\mu}^a)^2\right) \int \frac{D\hat{\mathbf{m}}_{\mu>s}}{(2\pi)^{P-s}} \\
 & \times \exp\left(-\frac{\sigma^2}{2d} \sum_{\mu>s} \sum_{ab} Q_{ab} \hat{m}_{\mu}^a \hat{m}_{\mu}^b - i \sum_{\mu>s} \sum_a m_{\mu}^a \hat{m}_{\mu}^a\right) \\
 & = \left(\det \frac{\sigma^2}{d} \mathbb{Q}\right)^{-\frac{(P-s)}{2}} \int \frac{D\mathbf{m}_{\mu>s}}{(2\pi)^{\frac{P-s}{2}}} \exp\left(-\frac{1}{2} \sum_{\mu>s} \sum_{ab} \left(\frac{d(Q^{-1})_{ab}}{\sigma^2} - \beta\delta_{ab}\right) m_{\mu}^a m_{\mu}^b\right) \\
 & = \left[\det\left(\mathbb{I}_n - \frac{\beta\sigma^2}{d} \mathbb{Q}\right)\right]^{-\frac{(P-s)}{2}}. \tag{A6}
 \end{aligned}$$

We can easily eliminate $\{\hat{m}_{\mu,a}\}_{\mu\leq s}$, Q_{aa} and $\hat{\mathbb{Q}}$ through saddle point evaluations

$$\begin{aligned}
 Q_{aa} &= 1 & i\hat{m}_{\mu}^a &= \beta\sigma^2 m_{\mu}^a & i\hat{Q}_{ab} &= \frac{\alpha}{2} \frac{\partial}{\partial Q_{ab}} \log \det\left(\mathbb{I}_n - \frac{\beta\sigma^2}{d} \mathbb{Q}\right) \\
 & & & & &= -\frac{\alpha\beta\sigma^2}{2d} \left(\mathbb{I}_n - \frac{\beta\sigma^2}{d} \mathbb{Q}\right)_{ab}^{-1} \equiv -\frac{\alpha\beta^2\sigma^2}{2} R_{ab}
 \end{aligned}$$

The replicated partition function can be finally written as a saddle point integral

$$\begin{aligned}
 \bar{Z}^n &\propto \int D\mathbf{m}_{\mu\leq s} D\mathbb{Q} \exp(N A_n(\mathbf{m}_{\mu\leq s}, \mathbb{Q})) \\
 A_n(\mathbf{m}_{\mu\leq s}, \mathbb{Q}) &= -\frac{\alpha\beta\sigma^2 n}{2d} - \frac{\beta\sigma^2}{2} \sum_{\mu\leq s} \sum_a (m_{\mu}^a)^2 - \frac{\alpha\beta^2\sigma^2}{2} \sum_{ab} R_{ab} Q_{ab} \\
 &\quad - \frac{\alpha}{2} \log \det\left(\mathbb{I}_n - \frac{\beta\sigma^2}{d} \mathbb{Q}\right) + \log W(\mathbf{m}_{\mu\leq s}, \mathbb{R}) \\
 W(\mathbf{m}_{\mu\leq s, a}, \mathbb{R}) &= \int \prod_{a=1}^n (dS_a \delta(S_a - \sigma)) \\
 &\quad \times \int \prod_{\mu\leq s} \left(\frac{d\xi^{\mu} \delta(\xi^{\mu} - 1)}{\mathcal{S}_{d-1}(1)}\right) e^{\beta\sigma \sum_a \sum_{\mu\leq s} m_{\mu}^a \xi_{\mu} \cdot S_a + \frac{\alpha\beta^2}{2} \sum_{ab} R_{ab} S_a \cdot S_b}. \tag{A7}
 \end{aligned}$$

The original problem has been reduced to a problem of n interacting spins $\{S_a\}_{a=1}^n$: once the saddle point integral is evaluated, the free energy can be finally evaluated after continuing analytically to continuous n and performing the $n \rightarrow 0$ limit. The saddle point equations related to (A7) read

$$m_{\mu}^a = \frac{1}{\sigma} \langle \xi_{\mu} \cdot S_a \rangle_{\text{eff}} \tag{A8}$$

$$Q_{ab} = \frac{1}{\sigma^2} \langle S_a \cdot S_b \rangle_{\text{eff}} \tag{A9}$$

$$\begin{aligned} \langle \cdot \rangle_{\text{eff}} &= \frac{1}{W} \int \prod_{a=1}^n (dS_a \delta(S_a - \sigma)) \\ &\times \int \prod_{\mu \leq s} \left(\frac{d\xi^\mu \delta(\xi^\mu - 1)}{\mathcal{S}_{d-1}(1)} \right) e^{\beta\sigma \sum_a \sum_{\mu \leq s} m_\mu^a \xi_\mu \cdot S_a + \frac{\alpha\beta^2}{2} \sum_{ab} R_{ab} S_a \cdot S_b} (\cdot). \end{aligned} \quad (\text{A10})$$

Unfortunately, saddle point equations (A8), (A9) cannot be solved for generic $\underline{m}_{\mu \leq s} \in \mathbb{Q}$, but an ansatz is needed. The natural choice is a RS ansatz, as the original problem is invariant under permutation of replica indices:

$$m_\mu^a = m_\mu \quad Q_{ab} = \delta_{ab} + (1 - \delta_{ab})q. \quad (\text{A11})$$

With this ansatz the action in (A7) can be further simplified, allowing us to finally evaluate the free energy of the system. First, the noise parameter \mathbb{R} becomes

$$R_{ab} \equiv \begin{cases} \frac{1}{\beta\sigma^2} \frac{1}{\frac{d}{\sigma^2} - \beta(1-q)} + \frac{q}{\sigma^2 [\frac{d}{\sigma^2} - \beta(1-q)]^2} + O(n) \equiv r_d & a = b \\ \frac{q}{\sigma^2 [\frac{d}{\sigma^2} - \beta(1-q)]^2} + O(n) \equiv r & a \neq b. \end{cases} \quad (\text{A12})$$

The action becomes

$$\begin{aligned} A_n &= -\frac{\alpha\beta\sigma^2 n}{2d} - \frac{\beta\sigma^2 n}{2} \sum_{\mu \leq s} m_\mu^2 - \frac{\alpha\beta^2\sigma^2 n(n-1)}{2} r q \\ &\quad - \frac{\alpha\beta^2\sigma^2 n}{2} r - \frac{\alpha}{2} \left[\log \left(1 - \frac{\beta\sigma^2}{d} (1 + (n-1)q) \right) \right. \\ &\quad \left. + (n-1) \log \left(1 - \frac{\beta\sigma^2}{d} (1-q) \right) \right] \\ &\quad + \log \int_{(\mathbb{S}_{d-1}(1))^n} \prod_a (ds_a) \sigma^{n-1} e^{\beta\sigma^2 \sum_{\mu \leq s} m_\mu \xi_\mu \cdot \sum_a s_a + \frac{\alpha\beta^2\sigma^2 r}{2} (\sum_a s_a)^2} (\{\xi_\mu\}_{\mu \leq s}). \end{aligned} \quad (\text{A13})$$

The first term with the log can be easily simplified

$$\begin{aligned} &\frac{\alpha}{2} \left[\log \left(1 - \frac{\beta\sigma^2}{d} (1 + (n-1)q) \right) + (n-1) \log \left(1 - \frac{\beta\sigma^2}{d} (1-q) \right) \right] \\ &= \frac{\alpha n}{2} \log \left(1 - \frac{\beta\sigma^2}{d} (1-q) \right) + \frac{\alpha\beta\sigma^2}{2d} \frac{nq}{1 - \frac{\beta\sigma^2}{d} (1-q)} + O(n^2). \end{aligned}$$

The integral in the last entropic term can be further simplified through an Hubbard-Stratonovich transform

$$e^{\frac{\alpha\beta^2\sigma^2 r}{2} (\sum_a s_a)^2} = \int \frac{d\mathbf{h}}{(2\pi\alpha r)^{d/2}} e^{-\frac{\mathbf{h}^2}{2\alpha r} + \beta\sigma\mathbf{h} \cdot \sum_a s_a}.$$

With this last passage, we have

$$\begin{aligned} & \overline{\int_{(\mathbb{S}_{d-1}(1))^n} \prod_a (ds_a) \sigma^{n-1} e^{\beta\sigma^2 \sum_{\mu \leq s} m_\mu \boldsymbol{\xi}_\mu \cdot \sum_a s_a + \frac{\alpha\beta^2\sigma^2 r}{2} (\sum_a s_a)^2} (\{\boldsymbol{\xi}_\mu\}_{\mu \leq s})} \\ &= n \int \frac{d\mathbf{h}}{(2\pi\alpha r)^{d/2}} e^{-\frac{h^2}{2\alpha r}} \log \mathcal{K}_d \left(\beta\sigma \left| \mathbf{h} + \sigma \sum_{\mu \leq s} m_\mu \boldsymbol{\xi}_\mu \right. \right) (\{\boldsymbol{\xi}_\mu\}_{\mu \leq s}) \end{aligned}$$

where we used $\mathcal{K}_d(x) = \int_{\mathbb{S}_{d-1}(1)} d\mathbf{s} e^{s \cdot x}$ defined in (16). Combining all these passages, using (13) one can straightforwardly obtain (21) and related saddle point equations (22) by extremizing (21).

Appendix B. The instability of mixtures

In this appendix, we study the stability of mixture solutions ($s > 1$) of equation (15). We have numerical evidence that mixtures are unstable for any α and generic types of mixtures (see for instance figure 6(b)). Here we focus on the sublinear case $\alpha = 0$ and on symmetric mixtures and show analytically that these states are unstable at all temperatures.

First, for readers convenience, we write again the expression of the free energy density in the $\alpha = 0$ case (equation (14)), but for symmetric mixtures

$$f_{\text{eq}}(\beta) = \frac{\sigma^2 s}{2} m^2 - \frac{1}{\beta} \ln \sigma^{d-1} \mathcal{K}_d \left(\beta\sigma^2 m \left| \sum_{\mu \leq s} \boldsymbol{\xi}_\mu \right. \right). \quad (\text{B1})$$

Here m denotes the Mattis magnetization of the symmetric s -mixture. Second, we also rewrite the saddle point equation for m :

$$m = g_d \left(\beta\sigma^2 m \left| \sum_{\mu=1}^s \boldsymbol{\xi}_\mu \right. \right) \frac{\boldsymbol{\xi}_\mu \cdot \sum_{\nu=1}^s \boldsymbol{\xi}_\nu}{\left| \sum_{\nu=1}^s \boldsymbol{\xi}_\nu \right|}. \quad (\text{B2})$$

The Hessian of (14), when evaluated on a symmetric mixture, reads ¹⁰

$$\begin{aligned} & \frac{1}{\sigma^2} \frac{\partial^2 f}{\partial m_\mu \partial m_\nu} \Big|_{\{m_\mu = m\}} \\ &= \delta_{\mu\nu} - \beta\sigma^2 \boldsymbol{\xi}_\mu^T \left[g'_d(\beta\sigma^2 m |z|) \mathbb{P}_\parallel \left(\frac{\mathbf{z}}{|z|} \right) + \frac{g_d(\beta\sigma^2 m |z|)}{\beta\sigma^2 m |z|} \mathbb{P}_\perp \left(\frac{\mathbf{z}}{|z|} \right) \right] \boldsymbol{\xi}_\nu \end{aligned} \quad (\text{B3})$$

where $\mathbf{z} \equiv \sum_{\mu=1}^s \boldsymbol{\xi}_\mu$. We recall that $\mathcal{K}_d(\cdot)$ and $g_d(\cdot)$ are defined in equation (16), while \mathbb{P}_\parallel and \mathbb{P}_\perp are respectively longitudinal and orthogonal projectors. In all equations (B1), (B2) and (B3), the average $\bar{\cdot}$ is understood to be over the patterns of the mixture $\{\boldsymbol{\xi}_\mu\}_{\mu=1}^s$.

¹⁰ To correctly evaluate the Hessian, one should rewrite the free energy in equation (14) as a function of the full configuration of Mattis magnetizations $\mathbf{m}_\mu = (m_1, \dots, m_p)$, do the derivatives and only after that restrict to symmetric mixtures of order s .

The Hessian is a $P \times P$ matrix with only three distinct non-zero entries, respectively for the three set of indices $\mu \neq \nu \leq s$, $\mu = \nu \leq s$ and $P > \mu = \nu > s$. In the following, we explain how to carry on the disorder average over the patterns of the mixture.

B.1. Average over mixtures

We decomposed the mixture vector into two contributions, $\mathbf{z} = \sum_{\mu \leq s} \boldsymbol{\xi}_\mu = \sum_{\mu \leq s-1} \boldsymbol{\xi}_\mu + \boldsymbol{\xi}_s \equiv \mathbf{w} + \boldsymbol{\xi}_s$. One has $|\mathbf{z}|^2 = |\mathbf{w}|^2 + 1 + 2|\mathbf{w}|u$ with $u = \frac{\mathbf{w}}{|\mathbf{w}|} \cdot \boldsymbol{\xi}_s$. The average over the disorder is thus expressed as an average over the random variables \mathbf{w} and u . They represent respectively the mixture vector of an order $s-1$ symmetric mixture and the scalar product of two random unit vectors. The pdfs of u and $w = |\mathbf{w}|$ are the following:

$$P_u(u) = \frac{(1-u^2)^{\frac{d-3}{2}}}{\int_0^\pi (\sin \vartheta)^{d-2} d\vartheta} \quad -1 \leq u \leq 1 \quad (\text{B4})$$

$$P_w^{(s)}(w) = \mathcal{S}_{d-1}(1) w^{d-1} \rho_{s-1}(w) \quad (\text{B5})$$

where $\rho_s(\mathbf{z}) \equiv P_z(\mathbf{z})$, the pdf of the mixture vector \mathbf{z} , satisfies the recursive equation

$$\begin{aligned} \rho_1(\mathbf{z}) &= \frac{\delta(\mathbf{z}-1)}{\mathcal{S}_{d-1}(1)} \\ \rho_s(\mathbf{z}) &= \frac{\int_0^\pi d\vartheta \sin^{d-2}(\vartheta) \rho_{s-1}(\sqrt{z^2+1-2z\cos\vartheta})}{\int_0^\pi d\vartheta \sin^{d-2}(\vartheta)} \quad s > 1. \end{aligned} \quad (\text{B6})$$

Note that $\rho_s(\mathbf{z})$ is the distribution of the position vector of a discrete-time random walk in d dimensions at time s , where at each time step the walker makes a unit-length move in a random direction.

For simplicity, we consider only the cases $s = 2, 3$: for these values, the distributions $\rho_2(\mathbf{z})$ and $\rho_3(\mathbf{z})$ read

$$\rho_2(\mathbf{z}) = \frac{1}{\mathcal{S}_{d-1}(1)} \frac{(4-z^2)^{\frac{d-3}{2}} \vartheta(\mathbf{z}) \vartheta(2-z)}{2^{d-3} z \int_0^\pi d\vartheta \sin^{d-2}(\vartheta)} \quad (\text{B7})$$

$$\begin{aligned} \rho_3(\mathbf{z}) &= \frac{\int_0^{\Theta(\mathbf{z})} d\vartheta \sin^{d-2}(\vartheta) \frac{[4-(z^2+1-2z\cos\vartheta)]^{\frac{d-3}{2}}}{2^{d-3} \sqrt{z^2+1-2z\cos\vartheta}}}{\mathcal{S}_{d-1}(1) \left[\int_0^\pi d\vartheta \sin^{d-2}(\vartheta) \right]^2}, \\ \Theta(\mathbf{z}) &= \pi \vartheta(1-z) + \arccos\left(\frac{z}{2} - \frac{3}{2z}\right) \vartheta(z-1) \vartheta(3-z). \end{aligned} \quad (\text{B8})$$

We could not find explicit expressions for $s > 3$ and generic d , but only for odd values of the spin dimension d . For instance, for $d = 3$ one has

$$\rho_3(\mathbf{z}; d=3) = \frac{1}{8\pi} \theta(1-z) + \frac{3-z}{16\pi z} \theta(z-1) \theta(3-z) \quad (\text{B9})$$

$$\rho_4(\mathbf{z}; d=3) = \frac{1}{16\pi} \left(1 - \frac{z}{2}\right) \theta(2-z) + \frac{1}{16\pi z} \theta(z-2) \theta(4-z). \quad (\text{B10})$$

B.2. Eigenvalues of the Hessian

With the notation introduced in the last paragraph, the saddle point equation for Mattis magnetization m , equation (15), and the elements of the hessian in equation (B3) can be rewritten as

$$m = \int_0^{s-1} dw P_w^{(s)}(w) \int_{-1}^1 du P_u(u) g_d \left(\beta \sigma^2 m \sqrt{1+w^2+2wu} \right) \frac{1+uw}{\sqrt{1+w^2+2wu}} \quad (\text{B11})$$

- For $\mu = \nu \leq s$

$$\begin{aligned} \frac{1}{\rho_1^2 \rho_2^2} \frac{\partial^2 f}{\partial m_\mu \partial m_\mu} &= 1 - \beta \sigma^2 I_1 \equiv a \\ I_1 &= \int_0^{s-1} dw P_w^{(s)}(w) \int_{-1}^1 du P_u(u) \left[\frac{(1+uw)^2}{1+w^2+2wu} g'_d \left(\beta \sigma^2 m \sqrt{1+w^2+2wu} \right) \right. \\ &\quad \left. + \left(1 - \frac{(1+uw)^2}{1+w^2+2wu} \right) \frac{g_d \left(\beta \sigma^2 m \sqrt{1+w^2+2wu} \right)}{\beta \sigma^2 m \sqrt{1+w^2+2wu}} \right]. \end{aligned} \quad (\text{B12})$$

- For $\mu = \nu > s$

$$\frac{1}{\sigma^2} \frac{\partial^2 f}{\partial m_\mu \partial m_\mu} = 1 - \beta \sigma^2 I_2 \equiv c \quad (\text{B13})$$

$$I_2 = \frac{1}{d} \int_0^s dw P_w^{(s+1)}(w) \left[g'_d(\beta \sigma^2 mw) + (d-1) \frac{g_d(\beta \sigma^2 mw)}{\beta \sigma^2 mw} \right].$$

- For $\mu \neq \nu \leq s$ by computing $\sum_{\mu\nu}^{1,s} \frac{\partial^2 f}{\partial m_\mu \partial m_\nu}$ using (B3) and knowing that $\frac{\partial^2 f}{\partial m_\mu \partial m_\nu}$ are expected to be all equal, one can show that

$$\frac{1}{\sigma^2} \frac{\partial^2 f}{\partial m_\mu \partial m_\nu} = \frac{\beta \sigma^2}{s-1} \left(I_1 - \frac{I_3}{s} \right) \equiv b \quad (\text{B14})$$

$$I_3 = \int_0^s dw P_w^{(s+1)}(w) g'_d(\beta \sigma^2 mw) w^2.$$

- For $\mu \neq \nu > s$, the Hessian entries are zero.

We now proceed to evaluate the Hessian spectrum. The Hessian (B3) has a simple block structure (compare with equations (B12), (B14), (B13))

$$\left(\begin{array}{cc|cc} a & \mathbf{b} & & \\ & \ddots & & \mathbf{0} \\ \mathbf{b} & a & & \\ \hline & & c & \mathbf{0} \\ & \mathbf{0} & & \ddots \\ & & \mathbf{0} & c \end{array} \right).$$

Due to this simple structure, the Hessian has three distinct eigenvalues

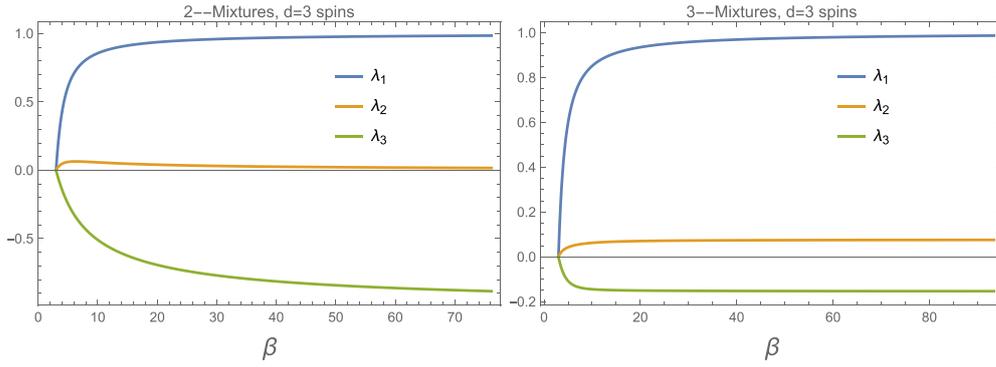


Figure 8. The three distinct eigenvalues of the $\alpha = 0$ free energy Hessian in (B3) evaluated on symmetric mixtures, as functions of inverse temperature β . The left panel is for $s = 2$ mixtures, the right one for $s = 3$ ones. Mixtures of patterns are never stable states: the instability is due to the eigenvalue λ_3 , related to fluctuations that break the symmetry of the mixture.

- The nondegenerate eigenvalue

$$\lambda_1 = a + (s - 1)b = 1 - \frac{\beta\sigma^2}{s}I_3. \tag{B15}$$

This eigenvalue is related to longitudinal fluctuations of the symmetric mixture, i.e. fluctuations that preserve the degree and symmetry of the mixture.

- The eigenvalue

$$\lambda_2 = c = 1 - \beta\sigma^2I_2 \tag{B16}$$

with degeneracy $M - s$. This eigenvalue is related to fluctuations that increase the degree of the mixture, without breaking its symmetry.

- The eigenvalue

$$\lambda_3 = a - b = 1 - \frac{s}{s - 1}\beta\sigma^2I_1 + \frac{\beta\sigma^2}{s(s - 1)}I_3 \tag{B17}$$

with degeneracy $s - 1$. This eigenvalue represents fluctuations that break the symmetry of the mixtures.

In figure 8 we show for the $d = 3$ VHM that symmetric mixtures with $s = 2, 3$ are unstable for any $T \geq 0$: the instability is due to the eigenvalue λ_3 , therefore, to fluctuations that break the symmetry of the mixture. For the SHM ($d = 1$) one has instead that odd¹¹ symmetric mixtures stabilize, the eigenvalue λ_3 becoming positive for sufficiently low temperatures [29]. The cause of this different behavior is the different nature of the degrees of freedom: continuous for VHM and discrete for SHM. Mathematically, in the SHM model ($d = 1$) and for odd-order mixtures the integrals I_1, I_3 appearing in (B17) vanish exponentially fast in the zero temperature limit, resulting in $\lambda_3 \rightarrow 1$. Conversely, for $d > 1$ the integral I_1 decays only as a power-law, $I_1 \simeq \frac{\tilde{I}_1}{\beta\sigma^{2m}}$.

¹¹ Symmetric mixtures of even orders are always unstable in the SHM, because of the possibility of having null mixtures with finite probability.

B.3. Large s and large d limits

We now consider the instability of mixtures in the limits $s \rightarrow \infty, d = \mathcal{O}(1)$ and $d \rightarrow \infty, s = \mathcal{O}(1)$ respectively. For both these limits, we find that solutions of symmetric mixtures become marginally stable. Let us consider the two cases in the following:

- Large s , fixed d : Here we find that the distribution of $|z|$ is

$$P_z^{(s)}(z) = \frac{\mathcal{S}_{d-1}(1) z^{d-1} e^{-\frac{dz^2}{2s}}}{(2\pi \frac{s}{d})^{d/2}} \quad (\text{B18})$$

where finite s corrections are present close to the upper edge $z_{max} = s$. Let us focus on the $T = 0$ case: we have

$$m \simeq \frac{1}{\sqrt{s}} (1 - 1/d) \sqrt{\frac{d}{2}} \frac{\Gamma(\frac{d-1}{2})}{\Gamma(\frac{d}{2})} \equiv \frac{k_d}{\sqrt{s}} \quad (\text{B19})$$

$$\beta I_1 = 1 + \mathcal{O}(1/s) \quad (\text{B20})$$

$$\beta I_2 = 1 - \mathcal{O}(1/\sqrt{s}) \quad (\text{B21})$$

$$\beta I_3 = \frac{\beta I_1}{s-1} = \frac{1}{s} + \mathcal{O}(1/s^2). \quad (\text{B22})$$

The eigenvalues (B15),(B16), (B17) are

$$\lambda_1 = 1 - \frac{c_1}{s}$$

$$\lambda_2 = \frac{c_2}{\sqrt{s}} \quad (\text{B23})$$

$$\lambda_3 = -\frac{c_3}{s}$$

for some constants $c_1, c_2, c_3 > 0$. For any d , in the large s limit, mixtures become marginally stable.

- Large d , fixed s : in this limit, the distribution of $|z|$ is trivial, $P_z^{(s)}(z) \simeq \delta(z - \sqrt{s})$. The distribution of the scalar products is $P_u(u) \simeq \delta(u)$. Thus, we find for $d \rightarrow \infty$ (compare with equation (B11))

$$m = \frac{1}{\sqrt{s}} \sqrt{1 - \frac{1}{\beta}} \quad (\text{B24})$$

$$\beta I_1 = \frac{1}{s(2\beta - 1)} + 1 - \frac{1}{s}$$

$$\beta I_2 = 1 \quad (\text{B25})$$

$$\beta I_3 = \begin{cases} 0 & s = 1 \\ \frac{\beta I_1}{s-1} - \frac{1}{(s-1)(2\beta-1)} & s > 1. \end{cases}$$

Thus we find for the three eigenvalues (B15), (B16), (B17)

$$\lambda_1 = 1 - \frac{1}{2\beta - 1} - \frac{c_1}{d}$$

$$\lambda_2 = \frac{c_2}{d}$$

$$\lambda_3 = \begin{cases} \lambda_1 & s = 1 \\ -\frac{c_3}{d} & s > 1 \end{cases}$$
(B26)

for some constants $c_1, c_2, c_3 > 0$. In the large d limit, symmetric mixtures of s examples become marginally stable.

Appendix C. Scaling limits

We consider the RS solution (21), (24) and the related saddle point equations (22), (25) in the limit $d \rightarrow \infty$. We consider this limit as simultaneous with $N \rightarrow \infty$ and $P \rightarrow \infty$, identifying two scaling regimes:

- (a) $N, P, d \rightarrow \infty$ with $\alpha' = \frac{P}{Nd} = \mathcal{O}(1)$.
- (b) $N, P, d \rightarrow \infty$ with $\alpha'' = \frac{Pd}{N} = \mathcal{O}(1)$.

The first regime corresponds to storing $P' = \mathcal{O}(Nd) \gg \mathcal{O}(N)$ examples, while in the second one instead has $1 \ll P'' = \mathcal{O}(N/d) \ll \mathcal{O}(N)$. Note that in this latter case one can store an infinite number of examples, provided that $N/d \rightarrow \infty$.

When $d \rightarrow \infty$, the correct scaling for the spins norm such that the free energy in equation (21) is well defined is $\sigma = \mathcal{O}(\sqrt{d})$: in particular, $\sigma = \sqrt{d}$ fixes the $\alpha = 0$ critical temperature to unit for any d , so we choose this value for spin norms in our large d limit computations.

a. α' scaling expressing the free energy in (21) in terms of α' and fixing $\sigma = \sqrt{d}$, it reads

$$\frac{f}{d} = \frac{1}{2}m^2 - \frac{1}{\beta d} \int_0^\infty dh P_h(h) |_{\alpha=\alpha'd} \log \left[d^{\frac{d-1}{2}} \mathcal{K}_d(\beta\sigma h) \right]$$

$$+ \frac{\alpha'}{2} \left\{ 1 + \frac{1}{\beta} \log[1 - (1-q)] + \frac{\beta(1-q)q}{[1 - \beta(1-q)]^2} - \frac{q}{1 - \beta(1-q)} \right\}$$
(C1)

The distribution of cavity field in this last equation is

$$P_h(h) |_{\alpha=\alpha'd} = \frac{h^{d-1} e^{-\frac{h^2}{2\alpha'r'}}}{(2\pi\alpha'r')^{d/2}} e^{-\frac{m^2 d}{2\alpha'r'}} \mathcal{K}_d \left(\frac{m\sqrt{d}}{\alpha'r'} h \right)$$
(C2)

where we set $r' = rd$. To perform the limit $d \rightarrow \infty$, we shall evaluate the integral in (C1). The function $\mathcal{K}_d(\cdot)$ defined in (16) has the following asymptotic behavior for large d [70]

$$d^{\frac{d-1}{2}} \mathcal{K}_d(dx) = \int_{\mathcal{S}_{d-1}(\sqrt{d})} d\mathbf{r} e^{\mathbf{r} \cdot x d} \sim \left(\frac{1}{\sqrt{1+4x^2+1}} \right)^{\frac{d-2}{2}} \frac{2^{d-1} e^{\frac{d}{2}} \pi^{\frac{d-1}{2}} e^{\frac{d}{2} \frac{2x^2}{1+\sqrt{1+4x^2}}}}{(1+4x^2)^{1/4}}.$$
(C3)

After rescaling $h \rightarrow \sqrt{d}h$, by using the expression in equation (C3) we can evaluate the integral in (C1) as a saddle point integral for large d , obtaining

$$\begin{aligned} & \lim_{d \rightarrow \infty} \frac{1}{d} \int_0^\infty dh P_h(h) \Big|_{\alpha=\alpha'd} \log \left[d^{\frac{d-1}{2}} \mathcal{K}_d(\beta\sigma h) \right] \\ &= \frac{1}{2} \left[\sqrt{1+4\beta^2(\alpha'r'+m^2)} - \log \left(\sqrt{1+4\beta^2(\alpha'r'+m^2)} + 1 \right) + \log 4\pi \right]. \end{aligned} \quad (\text{C4})$$

Finally, we obtain the free energy density

$$\begin{aligned} & \lim_{d \rightarrow \infty} \frac{f}{d} \Big|_{\alpha=\alpha'd} \\ &= \frac{1}{2} m^2 - \frac{1}{2\beta} \left[\sqrt{1+4\beta^2(\alpha'r'+m^2)} - \log \left(\sqrt{1+4\beta^2(\alpha'r'+m^2)} + 1 \right) + \log 4\pi \right] \\ &+ \frac{\alpha'}{2} \left\{ 1 + \frac{1}{\beta} \ln[1-\beta(1-q)] + \frac{q\beta(1-q)}{[1-\beta(1-q)]^2} - \frac{q}{1-\beta(1-q)} \right\}. \end{aligned} \quad (\text{C5})$$

The saddle point equations in equation (22) converge to the following equations

$$r' = \frac{q}{1-\beta(1-q)} \quad (\text{C6})$$

$$m = \frac{2\beta m}{1 + \sqrt{1+4\beta^2(\alpha'r'+m^2)}} \quad (\text{C7})$$

$$q = \frac{4\beta^2(\alpha'r'+m^2)}{[1 + \sqrt{1+4\beta^2(\alpha'r'+m^2)}]^2} \quad (\text{C8})$$

where we proceeded as before for the integral using the asymptotic expansion of $g_d(\cdot)$:

$$g_d(dx) = \frac{2x}{1 + \sqrt{1+4x^2}} + \mathcal{O}\left(\frac{1}{d}\right). \quad (\text{C9})$$

The only physical solution of equation (C6), (C7), (C8) for $\alpha' > 0$ is the no retrieval solution, with $m=0$ and $q = 1 - \beta_c(\alpha')/\beta$ for $\beta > \beta_c(\alpha') = 1/(1 + \sqrt{\alpha'})$, $q=0$ otherwise. This solution, when casted into equation (C5), returns the equilibrium free energy of the spherical Hopfield model with pairwise interactions [44].

b. α'' scaling by proceeding analogously to the previous case, one gets the limiting free energy

$$\begin{aligned} & \lim_{d \rightarrow \infty} \frac{f}{d} \Big|_{\alpha=\frac{\alpha''}{d}} = \frac{1}{2} (m^2 + \alpha''r'') - \frac{1}{2\beta} \left[\sqrt{1+4\beta^2(\alpha''r''+m^2)} \right. \\ & \left. - \log \left(\sqrt{1+4\beta^2(\alpha''r''+m^2)} + 1 \right) + \log 4\pi \right]. \end{aligned} \quad (\text{C10})$$

where we set $r'' = r/d$. The saddle point equations read

$$r'' = \lim_{d \rightarrow \infty} \frac{1 - 1/\beta}{\left[1 - \beta \left(1 - \langle g_d^2(\beta\sqrt{d}h) \rangle_{h, \alpha=\alpha''/d} \right) \right]^2 d^2} \quad (\text{C11})$$

$$m = \frac{2\beta m}{1 + \sqrt{1+4\beta^2(\alpha''r''+m^2)}} \quad (\text{C12})$$

$$q = 1 - \frac{1}{\beta} = \frac{4\beta^2 (\alpha'' r'' + m^2)}{\left[1 + \sqrt{1 + 4\beta^2 (\alpha'' r'' + m^2)}\right]^2} \quad (\text{C13})$$

Equations (C12), (C13) admit the retrieval solution

$$m(\alpha'', \beta) = \sqrt{1 - \frac{1}{\beta} - \alpha'' r''(\alpha'', \beta)}. \quad (\text{C14})$$

We did not derive the expression of r'' as a function of α'' and β , as it involved next-to-leading-order corrections to the expression in (C3), (C9), resulting computationally quite demanding. However, numerically, by extrapolating the expression in (C11) to $d \rightarrow \infty$ we found that r'' at $T=0$ is real in the range $0 < \alpha'' < \alpha_c''(0) = 4/27 \simeq 0.148181$.

Appendix D. Details on the retrieval dynamics

The retrieval dynamics follows equation (5), denoted by us as *versor rule*. We write it again in this appendix for the convenience of readers

$$\mathbf{S}_i(t+1) = \frac{\sigma \sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j(t)}{|\sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j(t)|} = \frac{\sigma \boldsymbol{\eta}_i(t)}{|\boldsymbol{\eta}_i(t)|}. \quad (\text{D1})$$

For simplicity, let us assume $\sigma = 1$ for the rest of this appendix. We stop the retrieval algorithm when the spins on average have aligned above a certain threshold: the stopping time is defined as

$$t_{\text{stop}} = \min \left\{ t \in \mathbb{N} / \{0\} : \frac{1}{N} \sum_{i=1}^N \mathbf{S}_i(t) \cdot \mathbf{S}_i(t-1) > 1 - \epsilon, \mathbf{S}_i(t) = \frac{\sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j(t)}{|\sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j(t)|} \right\}. \quad (\text{D2})$$

In all our simulations, we always used $\epsilon = 10^{-12}$. In algorithm 1, we report the pseudocode for the retrieval algorithm. Here we consider the retrieval of patterns, but equation (5) is a minimization algorithm for the energy function (1) and thus can be used for arbitrary initial states.

D.1. Statistical analysis of stopping time and quality of memory states

We performed numerical experiments to study the relation between the quality of the memory state (high Mattis magnetization) and the stopping time defined in equation (D2). These simulations are apart from those performed in section 3.2 to study the spectrum of memory states. We generated several samples of the $d = 3$ VHM and run the retrieval dynamics as described in algorithm 1, but using as a convergence criterion the stronger condition $\min_i \mathbf{S}_i \cdot \boldsymbol{\xi}_i^{(\mu)} > 1 - \epsilon$. Even though in the different batch of simulations we run to study the spectrum in 3.2 we used the average overlap criterion, with this stronger condition we can characterize spins alignment rates locally and thus compare the average convergence time (which corresponds to the stopping time defined in equation (D2)) and the maximal convergence time.

We focused on capacity values $\alpha = 0.02$ and $\alpha = 0.04$, representing memory states as global ($\alpha = 0.02 < \alpha_m \simeq 0.251$) and local minima ($\alpha = 0.04 > \alpha_m$). Since these simulations are for retrieval at zero temperature, Mattis magnetizations are simply calculated as

$$m_\mu = \frac{1}{N} \underline{\mathbf{S}} \cdot \underline{\boldsymbol{\xi}}^{(\mu)}. \quad (\text{D3})$$

Algorithm 1. Retrieval dynamics algorithm.

- 1: Initialize \underline{S} to a random pattern $\underline{\xi}^{(\mu)}$ sampled uniformly from $\{\underline{\xi}^{(1)}, \dots, \underline{\xi}^{(P)}\}$
- 2: Initialize $q=0$ and $t=0$
- 3: **while** $q < 1 - \epsilon$ **do**
- 4: Save old state $\underline{S}_{old} \leftarrow \underline{S}(t)$
- 5: **for** $i \in \text{RandPermutation}(\{1, \dots, N\})$ **do**
- 6: $S_i(t+1) \leftarrow \frac{\sum_{j:j \neq i} J_{ij} S_j(t)}{|\sum_{j:j \neq i} J_{ij} S_j(t)|}$
- 7: **end for**
- 8: Compute $q \leftarrow \frac{1}{N} \underline{S}(t) \cdot \underline{S}_{old}$
- 9: Increment time $t \leftarrow t + 1$
- 10: **end while**

Statistics of samples magnetizations and convergence times

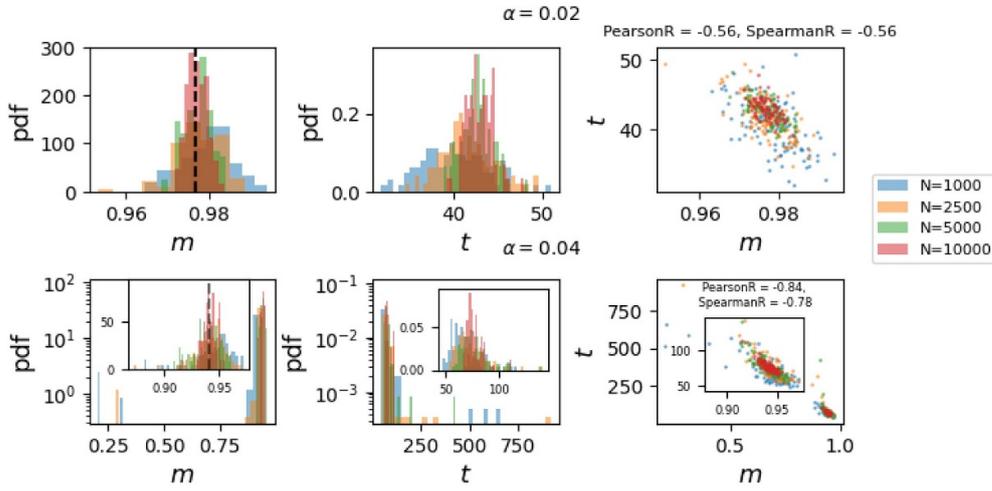


Figure 9. Statistics of Mattis magnetizations and stopping times, for sizes $N = 1000, 2500, 5000, 10000$. The top row is for capacity $\alpha = 0.02$, the bottom row for $\alpha = 0.04$. Left figures are the pdf of Mattis magnetizations, center figures of the stopping times and right figures are scatter plots highlighting the correlation between the two. The correlation coefficients reported in the figure are for the largest size $N = 10000$, while the dashed vertical line on the left figures marks the position of the asymptotic value, predicted by equation (25).

In figure 9, we report our results. For the largest sizes, in both cases the distribution of Mattis magnetizations is concentrated around the asymptotic value, given by the solution of equation (25). For $\alpha = 0.04$, we observe that at the smallest sizes the retrieval dynamics occasionally converges to spurious states, with magnetization $m \simeq 0.25 \ll m_{\text{theor}} \simeq 0.947$. In studying the properties of memory states in section 3.2, we only selected states with high enough magnetization: for the specific case $\alpha = 0.04$, for instance, we found that the threshold $m \simeq 0.8$ was a good compromise.

When retrieval dynamics does not converge to memory states, stopping times increase significantly. Moreover, stopping times are anti-correlated to Mattis magnetizations and thus to the quality of the memory state, as predictable.

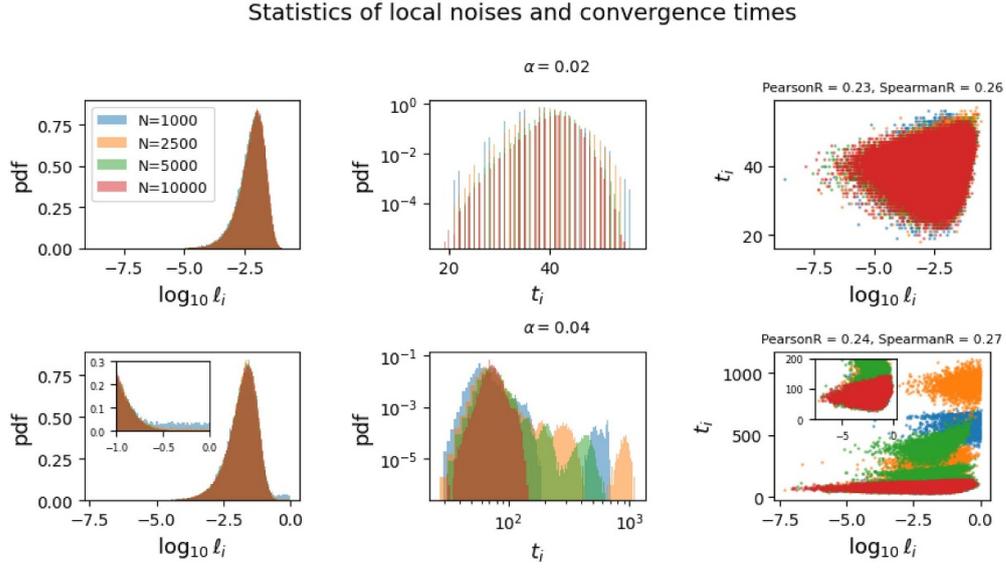


Figure 10. Statistical analysis of local noises and local convergence times. Top row refers to $\alpha = 0.02$, while bottom row to $\alpha = 0.04$. Left plots are the distributions of $\log_{10} \ell_i$, equation (41), central plots the pdfs of local convergence times (D4), right plots are scatter plots of the two quantities. The color legend is common for all the plots. The correlation coefficients in the scatter plots refer to the largest size analyzed $N = 10000$.

D.2. Statistical analysis of local effects: differential alignments and noisy spins

We considered the different rates at which spins converged, defining local stopping times

$$t_i^{\text{stop}} = \min \left\{ t \in \mathbb{N} / \{0\} : \mathbf{S}_i(t) \cdot \mathbf{S}_i(t-1) > 1 - \epsilon, \mathbf{S}_i(t) = \frac{\sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j(t)}{|\sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j(t)|} \right\} \quad (\text{D4})$$

and measured their correlation with the local noises, equation (41), whose definition we report for convenience here

$$\ell_i = \frac{1 - \mathbf{S}_i \cdot \boldsymbol{\xi}_i}{2}. \quad (\text{D5})$$

We report our results in figure 10. First, we comment on the pdfs of the log-local noises. For $\alpha = 0.02$, finite-size effects are absent; for $\alpha = 0.04$, we observe finite size effects on the upper tail of the distributions. These effects are caused by the dynamics occasionally converging to spurious states, as discussed in the previous paragraph (see figure 9). The different features observed for the values $\alpha = 0.02, 0.04$ are reflected in the behavior of the distributions of local convergence times: in the former case, they are single supported, showing that retrieval dynamics consistently converges to memory states; in the latter, for $N < 10000$ we observe multi-supported distributions, an indication that retrieval dynamics not always converges to the closest memory states. Notably, it appears that there are states that are as close to the starting pattern as typical memory states but are more noisy and require longer runs to convergence (see the scatter plot in figure 10 in the lower right corner). Regarding correlation, we observe a positive correlation between local convergence times and local noises, yet weaker than that observed in figure 9 between Mattis magnetizations and convergence times.

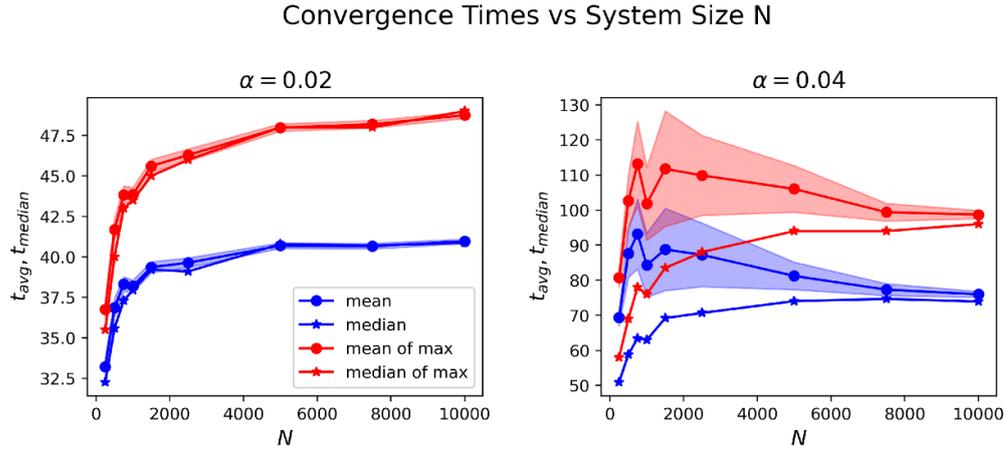


Figure 11. Means and medians of the samples convergence times and means and medians of the maximal convergence times of each sample.

D.3. Computational cost of slow spins

Finally, we characterize the computational cost of slow spins. In the previous analysis, we showed that slow spins correspond to noisy spins. In this paragraph, we want to check the scaling of convergence times with system size. We compare the average and median of the sample convergence times (equation (D2)) with the average and median of the maximal convergence times (the time the slowest spin needs to reach the desired alignment accuracy) of the different samples.

In figure 11, we show our results. We first remark that the difference between the two metrics (average and median) is apparent for $\alpha = 0.04$, in agreement with the fact that in this case the dynamics occasionally converges to spurious states that require longer times to convergence. In both cases, we find that the gap between typical and maximal times seems to saturate, or at most to grow very slowly: with the current accuracy and range of sizes, we are not yet able to single out one of the two options. Anyway, for any practical purposes, the gap essentially stabilizes at large enough sizes. This means that the additional computational cost of slow spins comes only from the cost related to the additional iterations. Since our retrieval dynamics in equation (5) can be parallelized, slow spins carry an additional $\mathcal{O}(nN^2/b)$ computational cost compared to the average, where n is the number of additional iterations and b the number of batches of parallelization. Therefore, the optimal choice to avoid unnecessary slowdowns is to stop the retrieval dynamics according to the average overlap, as we did in the simulations we run for the spectral analysis in 3.2. For sufficiently small thresholds ϵ , the loss of accuracy that one has on consecutive alignments of slow spins is negligible.

Appendix E: Further results for the hessian spectrum (I)

This appendix contains several derivations for the Hessian spectrum that were not included in the main body of the text.

E.1. The distribution of local fields

The local fields

$$\eta_i = \left| \sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j^* \right|. \quad (\text{E1})$$

are distributed according to

$$P_\eta(\eta) = \vartheta(\eta - p) P_h(\eta - p), \quad p = \frac{\alpha\sigma}{d} \frac{\chi}{\frac{d}{\sigma^2} - \chi} \quad (\text{E2})$$

where p is the Onsager reaction term, with χ solving the $T = 0$ saddle points equations in (25), and $P_h(h)$ is the pdf of cavity fields defined by equation (23).

We measured the empirical pdfs of local fields over numerical samples

$$P_\eta^{(\text{emp})}(\eta) = \frac{1}{N} \sum_{i=1}^N \delta \left(\eta_i - \left| \sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j^* \right| \right) \quad (\text{E3})$$

and compared it with the prediction yielded by (E2), for both Mattis states and spurious spin glass states. In figure we show our results, for the $d = 3$ VHM at capacity $\alpha = 0.04$. The left panel shows the pdf for Mattis states, the right one that for no-retrieval states. In the former case, we observe a perfect agreement between the theoretical prediction (E2) and the empirical one (E3); in the latter case, it can be seen that our prediction is not exact. This is due to the RS ansatz used for our solution in (24), (25),(23), that in the case of spurious spin glass states is not correct because of the RSB transition.

E.2. Derivation of the Hessian

The Hessian in (28) can be derived straightforwardly from the Hamiltonian (1). Consider a configuration of spins $\underline{\mathbf{S}}_*$ that is a stationary point of \mathcal{H} : any configuration sufficiently close to it can be parametrized as follows

$$\mathbf{S}_i = \sigma \boldsymbol{\pi}_i + \sqrt{1 - \pi_i^2} \mathbf{S}_i^* \quad (\text{E4})$$

where $|\mathbf{S}_i^*| = \sigma$ and $\boldsymbol{\pi}_i \cdot \mathbf{S}_i^* = 0$. For any spin, we introduced an orthogonal ‘pion’ vector: it is clear that (E4) describes for any site i a random rotation of the original spin \mathbf{S}_i^* within its half of hypersphere: indeed, $0 \leq \frac{1}{\sigma^2} \mathbf{S}_i \cdot \mathbf{S}_i^* = \sqrt{1 - \pi_i^2} \leq 1$. By plugging (E4) into (1) we get

$$\begin{aligned} \mathcal{H}(\underline{\mathbf{S}}) &= -\frac{1}{2} \sum_{i \neq j} \sqrt{(1 - \pi_i^2)(1 - \pi_j^2)} \mathbf{S}_i^* \cdot \mathbb{J}_{ij} \mathbf{S}_j^* - \frac{\sigma^2}{2} \sum_{i \neq j} \boldsymbol{\pi}_i^\top \mathbb{J}_{ij} \boldsymbol{\pi}_j - \sigma \sum_{i \neq j} \sqrt{1 - \pi_j^2} \boldsymbol{\pi}_i^\top \mathbb{J}_{ij} \mathbf{S}_j^* \\ &= \mathcal{H}(\underline{\mathbf{S}}_*) - \sigma \sum_{i=1}^N \boldsymbol{\pi}_i \cdot \sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j^* + \frac{\sigma^2}{2} \sum_{ij} \boldsymbol{\pi}_i^\top \left[-\mathbb{J}_{ij} + \delta_{ij} \mathbb{I}_d \frac{\mathbf{S}_i^*}{\sigma^2} \cdot \sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j^* \right] \boldsymbol{\pi}_j + \mathcal{O}(\boldsymbol{\pi}^3) \\ &\simeq \mathcal{H}(\underline{\mathbf{S}}_*) + \frac{\sigma^2}{2} \sum_{ij} \boldsymbol{\pi}_i^\top \left(-\mathbb{J}_{ij} + \delta_{ij} \frac{\eta_i}{\sigma} \mathbb{I}_d \right) \boldsymbol{\pi}_j \end{aligned}$$

where \mathbb{I}_d is the d -dimensional identity matrix and between the third and the fourth side of this last equation we applied the condition of stationarity (3): local field vectors $\boldsymbol{\eta}_i = \sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j^*$

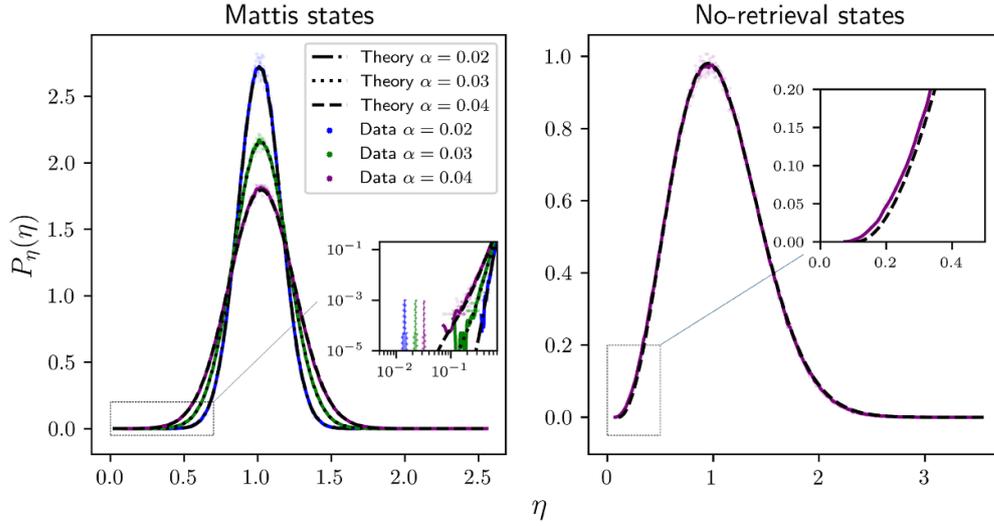


Figure 12. The distribution of local fields of the $d = 3$ VHM. We compare our theoretical prediction (31) with empirical pdf from numerical simulations of $N = 4000$ spins. Left: the distribution for Mattis states. We show theoretical and empirical pdfs for capacities $\alpha = 0.02, 0.03, 0.04$. There is a good agreement between theory and measurements. As shown in the inset in log-log scale, Onsager reactions (whose value is indicated by colored arrows) are quite small: one has $p \simeq 0.014, 0.023, 0.032$ respectively for $\alpha = 0.02, 0.03, 0.04$. Right: the pdf for no-retrieval states only for capacity $\alpha = 0.04$ (the curves of different capacities would be too close). As expected, there is no good agreement between the RS theory and experiment: the inset highlights that RSB corrections seems to move low local fields to lower values. Here the RS Onsager reaction is $p \simeq 0.11$.

are parallel to spins on a stationary point. The Hessian obtained translates into (28) when extended to the whole space of configurations:

$$\mathbb{M}_{ij}(\underline{\mathbf{S}}_*) = \mathbb{P}_\perp(\mathbf{S}_i^*) \left(-\mathbb{J}_{ij} + \frac{\eta_i}{\sigma} \delta_{ij} \mathbb{I}_d \right) \mathbb{P}_\perp(\mathbf{S}_j^*)$$

where $\mathbb{P}_\perp(\mathbf{x}) = \mathbb{I}_d - \frac{\mathbf{x} \mathbf{x}^\top}{x}$ is an orthogonal projector.

E.3. The spectrum of the interaction matrix

We consider the off-diagonal part of the Hessian, having elements

$$W_{ij}^{\gamma\delta} = -(\mathbb{P}_i^\perp \mathbb{J}_{ij} \mathbb{P}_j^\perp)_{\gamma\delta} (1 - \delta_{ij}). \tag{E5}$$

We compute its resolvent function

$$g(z) = \frac{1}{Nd} \text{Tr}(\underline{\mathbb{W}} - z \underline{\mathbb{I}}_{Nd})^{-1} \tag{E6}$$

through simplification of the exact Schur formula

$$(\underline{\mathbb{G}}_{11}^{-1})_{\beta\gamma} = -z\delta_{\beta\gamma} - \sum_{j,k>1} \sum_{\delta\epsilon}^{1,d} W_{1j}^{\beta\delta} \tilde{\mathbb{G}}_{jk}^{\delta\epsilon} W_{k1}^{\epsilon\gamma}. \tag{E7}$$

We define $\widetilde{\mathbb{G}}$ as the cavity resolvent (the resolvent in the system with a cavity on site 1). We found that the spectrum of the interaction matrix does not depend on the type of minimum (pattern-correlated or spin glass). In the following, the calculation: introducing $v_{i,\gamma}^\mu = (\mathbb{P}_i^\perp \xi_i^\mu)_\gamma$, we expand equation (E7)

$$(\mathbb{G}_{11}^{-1})_{\beta\gamma} = -z\delta_{\beta\gamma} - \sum_{j,k>1} \sum_{\delta\epsilon} \frac{1}{N^2} \sum_{\mu,\nu}^{1,M} v_{1,\beta}^\mu v_{j,\delta}^\mu \widetilde{\mathbb{G}}_{jk}^{\delta\epsilon} v_{k,\epsilon}^\nu v_{1,\gamma}^\nu. \quad (\text{E8})$$

The second term in the r.h.s. can be easily manipulated (note that in the following $A_{\mu\nu}$ is uncorrelated to $v_{1,\beta}^\mu, v_{1,\gamma}^\nu$ by construction)

$$\begin{aligned} & \frac{1}{N} \sum_{\mu,\nu}^{1,M} v_{1,\beta}^\mu \underbrace{\left(\frac{1}{N} \sum_{j,k>1} \sum_{\delta\epsilon} v_{j,\delta}^\mu \widetilde{\mathbb{G}}_{jk}^{\delta\epsilon} v_{k,\epsilon}^\nu \right)}_{A_{\mu\nu}} v_{1,\gamma}^\nu \\ &= \frac{1}{N} \sum_{\mu,\nu} v_{1,\beta}^\mu A_{\mu\nu} v_{1,\gamma}^\nu = \frac{1}{N} \sum_{\mu} v_{1,\beta}^\mu v_{1,\gamma}^\mu A_{\mu\mu} + O(1/\sqrt{N}) \\ &\simeq \frac{(\mathbb{P}_1^\perp)_{\beta\gamma}}{Nd} \sum_{j,k>1} \sum_{\delta\epsilon} \left(\frac{1}{N} \sum_{\mu} v_{j,\delta}^\mu v_{k,\epsilon}^\mu \right) \widetilde{\mathbb{G}}_{jk}^{\delta\epsilon} \\ &= \frac{(\mathbb{P}_1^\perp)_{\beta\gamma}}{Nd} \sum_{j,k>1} \sum_{\delta\epsilon} (-W_{jk}^{\delta\epsilon} + K_j^{\delta\epsilon} \delta_{jk}) \widetilde{\mathbb{G}}_{jk}^{\delta\epsilon} \\ &= \frac{(\mathbb{P}_1^\perp)_{\beta\gamma}}{Nd} \left(-\text{Tr} \widetilde{\mathbb{W}} \widetilde{\mathbb{G}} + \sum_{j>1} \text{Tr} \mathbb{K}_j \widetilde{\mathbb{G}}_j \right) \\ &\simeq -(\mathbb{P}_1^\perp)_{\beta\gamma} - (\mathbb{P}_1^\perp)_{\beta\gamma} z g(z) + \frac{(\mathbb{P}_1^\perp)_{\beta\gamma}}{Nd} \sum_{j>1} \text{Tr} \mathbb{K}_j \widetilde{\mathbb{G}}_{jj}. \end{aligned} \quad (\text{E9})$$

The tensor $K_j^{\gamma\delta}$ is

$$K_j^{\gamma\delta} = \frac{1}{N} \sum_{\mu=1}^M v_{j,\gamma}^\mu v_{j,\delta}^\mu \simeq \dots \simeq \frac{\alpha}{d} (\mathbb{P}_j^\perp)_{\gamma\delta}. \quad (\text{E10})$$

The last term in the last r.h.s. of (E9) is

$$\frac{1}{Nd} \sum_{j>1} \text{Tr} \mathbb{K}_j \widetilde{\mathbb{G}}_{jj} = \frac{\alpha}{Nd^2} \sum_{j>1} \text{Tr} \mathbb{P}_j^\perp \widetilde{\mathbb{G}}_{jj} \simeq \alpha \frac{(d-1)}{d^2} g_\perp(z) \quad (\text{E11})$$

where we used the idem-potency of the Projector and the properties of trace to write $(1/Nd) \sum \text{Tr} \mathbb{P}_j^\perp \widetilde{\mathbb{G}}_{jj} = (1/Nd) \sum \text{Tr} (\mathbb{P}_j^\perp \widetilde{\mathbb{G}}_{jj} \mathbb{P}_j^\perp) = (1/Nd) \sum \text{Tr} \widetilde{\mathbb{G}}_{jj}^\perp \simeq (1-1/d) g_\perp(z)$. Collecting everything and using $g(z) = (1-1/d) g_\perp(z) + (1/d) g_\parallel(z)$, we end up with the expression

$$\begin{aligned}
 (\mathbb{G}_{11}^{-1})_{\beta\gamma} &= -z\delta_{\beta\gamma} - \left[-1 - \frac{z}{d}g_{\parallel}(z) - z\left(1 - \frac{1}{d}\right)g_{\perp}(z) + \frac{\alpha}{d}\left(1 - \frac{1}{d}\right)g_{\perp}(z) \right] \\
 &\times (\mathbb{P}_j^{\perp})_{\beta\gamma} + O(N^{-1/2}).
 \end{aligned}
 \tag{E12}$$

Or equivalently, introducing $C(z) = 1 + \frac{z}{d}g_{\parallel}(z) + (z - \frac{\alpha}{d})(1 - \frac{1}{d})g_{\perp}(z)$

$$\mathbb{G}_{11} = - (z\mathbb{I}_d - C(z)\mathbb{P}_1^{\perp})^{-1}.
 \tag{E13}$$

The resolvent function is thus

$$g(z) = -\frac{1}{Nd} \sum_{i=1}^N \sum_{\beta=1}^d (z\mathbb{I}_d - C(z)\mathbb{P}_i^{\perp})_{\beta\beta}^{-1}
 \tag{E14}$$

We can easily compute the r.h.s. of this last expression, knowing that the orthogonal projector has 1 null eigenvalue and $d - 1$ unit eigenvalues:

$$g(z) = -\frac{d-1}{d} \frac{1}{z - C(z)} - \frac{1}{dz} = (1 - 1/d)g_{\perp}(z) + (1/d)g_{\parallel}(z).
 \tag{E15}$$

The last equation shows that there are N null eigenvalues out of $N \times d$ (there is a first order pole in $z = 0$ with weight $1/d$). Let us consider from now on only $g_{\perp}(z)$, the resolvent function in the sector orthogonal to \underline{S} . From equation (E15), we obtain

$$g_{\perp}(z) = -\frac{1}{z - (1 - \frac{1}{d}) - (z - \frac{\alpha}{d})(1 - 1/d)g_{\perp}(z)}
 \tag{E16}$$

which yields a second order algebraic equation in $g(z)$ whose correct¹² solution is the one with the ‘+’

$$\begin{aligned}
 g_{\perp}(z) &= \frac{z - (1 - \frac{1}{d})}{2(z - \alpha/d)(1 - \frac{1}{d})} \\
 &+ \frac{\sqrt{[z - (1 - \frac{1}{d})]^2 + 4(z - \alpha/d)(1 - 1/d)}}{2|z - \alpha/d|(1 - \frac{1}{d})}.
 \end{aligned}
 \tag{E17}$$

The spectrum is given from the first order poles of the resolvent function and from its cuts on the real line. It is not hard to show that the resolvent has a pole on the point $\lambda_p = +\alpha/d$ when $\alpha < d - 1$, with weight $1 - \alpha/(d - 1)$. For the continuous part of the spectrum, we shall compute the imaginary part of $g_{\perp}(z)$. Let $z = \lambda + i0_+$: one has that the argument of the square root in (E17) becomes negative when $\lambda_- < \lambda < \lambda_+$, with

$$\lambda_{\pm} = - \left[1 - 1/d \mp 2\sqrt{(1 - 1/d)(\alpha/d)} \right].
 \tag{E18}$$

The final expression for the spectral density is a shifted Marchenko-Pastur law

¹² It is the solution having the expected analytical properties for the resolvent function defined. In particular, it should behave as $g(z) \sim -1/z$ for large $|z|$.

$$\begin{aligned} \rho_{\perp}(\lambda) &= (1 - \alpha/(d-1)) \vartheta(d-1-\alpha) \delta(\lambda - \lambda_p) \\ &+ \frac{\sqrt{(\lambda_+ - \lambda)(\lambda - \lambda_-)}}{2\pi|\lambda - \lambda_p|(1-1/d)} \vartheta(\lambda - \lambda_-) \vartheta(\lambda_+ - \lambda). \end{aligned} \quad (\text{E19})$$

Interestingly, we recover the spectrum of a shifted (with no diagonal) Wishart matrix in the large d limit if and only if one fixes $\alpha' = \alpha/d$.

E.4. Derivation of the spectral equation

We derive the spectral equation (29) using random matrix techniques. In the large N limit, we treat the Hessian (28) as an instance of a random matrix ensemble, in this case the deformed Wishart ensemble, where a sort of Wishart matrix (the off-diagonal part or interaction matrix in (28)) is summed to a disordered diagonal matrix (the local fields in (28)).

We assume that the two random matrices that compose Hessian (28) are free: two random matrices are free when their respective eigenbasis are uncorrelated [45]. Free random matrices have the so-called R-transform of the resolvent function $\mathcal{R}(\mathcal{G}) = z(\mathcal{G}) + \frac{1}{\mathcal{G}}$ additive, where $z(\mathcal{G})$ is the inverse of the resolvent function. Given a matrix \mathbb{A} and a perturbation $\mathbb{B} = \text{diag } \mathbf{v}$, one can show that the resolvent of the sum $\mathbb{A} + \mathbb{B}$, thanks to $\mathcal{R}_{\mathbb{A}+\mathbb{B}} = \mathcal{R}_{\mathbb{A}} + \mathcal{R}_{\mathbb{B}}$, satisfies a self-consistent equation [45]

$$\mathcal{G}(z) = \int \frac{d\mathbf{v} P_{\mathbf{v}}(\mathbf{v})}{v - z + R_{\mathbb{A}}(\mathcal{G}(z))} \quad (\text{E20})$$

Therefore, if one knows the distribution of the diagonal elements (for our Hessian, P_{η}), it suffices to know the expression of the R-transform of matrix \mathbb{A} (the off-diagonal part of (28) in our case).

The self-consistent equation for the Green function $\mathcal{G}(z)$ of the off-diagonal part of the Hessian is equation (E16): from it, we see that the inverse of $\mathcal{G}(z)$ is

$$z(\mathcal{G}) = -\frac{1}{[1 - (1-1/d)\mathcal{G}]\mathcal{G}} + \frac{(1-1/d)(1-\alpha\mathcal{G}/d)}{1 - (1-1/d)\mathcal{G}} \quad (\text{E21})$$

and consequently for the R-transform

$$R(\mathcal{G}) = z(\mathcal{G}) + \frac{1}{\mathcal{G}} = -\frac{(\alpha/d)(1-1/d)\mathcal{G}}{1 - (1-1/d)\mathcal{G}}. \quad (\text{E22})$$

It then follows, plugging this last equation into (E20), that the resolvent function is given by

$$\mathcal{G}(z) = \int \frac{d\eta P_{\eta}(\eta)}{\frac{\eta}{\sigma} - \lambda - \frac{(\alpha/d)(1-1/d)\mathcal{G}(z)}{1 - (1-1/d)\mathcal{G}(z)}}. \quad (\text{E23})$$

The resolvent in this last equation is defined as $\mathcal{G}(z) = \frac{\sigma^2}{Nd} \text{Tr}(\underline{\mathbb{M}} - z \mathbb{I}_{Nd})^{-1}$: by changing definition into $\mathcal{G}(z) = \frac{d-1}{Nd} \text{Tr}(\underline{\mathbb{M}} - z \mathbb{I}_{Nd})^{-1}$, one gets

$$\mathcal{G}(z) = \frac{d-1}{\sigma^2} \int \frac{d\eta P_{\eta}(\eta)}{\frac{\eta}{\sigma} - \lambda - \frac{(\alpha/d)\mathcal{G}(z)}{\frac{d}{\sigma^2} - \mathcal{G}(z)}}$$

which is equation (29).

E.5. Numerical solution of spectral equation

In this appendix we explain how to solve numerically the spectral equation (29). Written in its real and imaginary part, equation (29) translates into

$$\Re \mathcal{G}(\lambda) = \frac{(d-1)}{\sigma} \int \frac{d\eta P_\eta(\eta) \left[\eta - \lambda\sigma - \frac{\alpha\sigma}{d} \frac{(\Re \mathcal{G}(\lambda) \frac{d}{\sigma^2} - |\mathcal{G}(\lambda)|^2)}{|\frac{d}{\sigma^2} - \mathcal{G}(\lambda)|^2} \right]}{\left| \eta - \lambda\sigma - \frac{\alpha\sigma}{d} \frac{\mathcal{G}(\lambda)}{\frac{d}{\sigma^2} - \mathcal{G}(\lambda)} \right|_{\mathbb{C}}^2}. \quad (\text{E24})$$

$$1 = \frac{\alpha(d-1)}{\sigma^2 |\frac{d}{\sigma^2} - \mathcal{G}(\lambda)|^2} \int \frac{d\eta P_\eta(\eta)}{\left| \eta - \lambda\sigma - \frac{\alpha\sigma}{d} \frac{\mathcal{G}(\lambda)}{\frac{d}{\sigma^2} - \mathcal{G}(\lambda)} \right|_{\mathbb{C}}^2}. \quad (\text{E25})$$

On principle, one has two solve equations (E24) and (E25) simultaneously in order to determine $\mathcal{G}(\lambda)$. We used a simplification that reduces the number of equations to be solved to one: we introduce the auxiliary resolvent $\tilde{\mathcal{G}}(\lambda) = \frac{\mathcal{G}(\lambda)}{\frac{d}{\sigma^2} - \mathcal{G}(\lambda)}$, which satisfy the self consistent equation

$$\tilde{\mathcal{G}}(\lambda) = \frac{\frac{(d-1)}{\sigma} \int \frac{dh P_h(h)}{h+x(\lambda)}}{\frac{d}{\sigma^2} - \frac{(d-1)}{\sigma} \int \frac{dh P_h(h)}{h+x(\lambda)}} \quad (\text{E26})$$

where we used (31) to express integrals in terms of cavity fields pdf $P_h(h)$ and set $x(\lambda) = \frac{\alpha\sigma}{d} \tilde{\mathcal{G}}(0) - \lambda\sigma - \frac{\alpha\sigma}{d} \tilde{\mathcal{G}}(\lambda)$. The equations for the real and imaginary parts of (E26) are

$$\Re \tilde{\mathcal{G}}(\lambda) = \tilde{\mathcal{G}}(0) - \frac{d}{\alpha\sigma} (\lambda\sigma + \Re x) = \frac{\Re \mathcal{I}(x) \frac{d}{\sigma^2} - |\mathcal{I}(x)|_{\mathbb{C}}^2}{|\frac{d}{\sigma^2} - \mathcal{I}(x)|_{\mathbb{C}}^2} \quad (\text{E27})$$

$$1 = \frac{\alpha(d-1) \mathcal{J}(x)}{\sigma^2 |\frac{d}{\sigma^2} - \mathcal{I}(x)|_{\mathbb{C}}^2} \quad (\text{E28})$$

$$\mathcal{I}(x) = \int \frac{dh P_h(h)}{h+x} \quad (\text{E29})$$

$$\mathcal{J}(x) = \int \frac{dh P_h(h)}{|h+x|_{\mathbb{C}}^2}. \quad (\text{E30})$$

For any value of $\Re x$ (respectively $\Im x$), one can solve (E28) to determine the correspondent $\Im x$ (respectively $\Re x$) and then compute λ through (E27): in this way one is able to obtain the auxiliary resolvent $\tilde{\mathcal{G}}(\lambda)$ for any given $\lambda > 0$. Finally, the original resolvent can be obtained through

$$\Re \mathcal{G}(\lambda) = \frac{\left(\Re \tilde{\mathcal{G}}(\lambda) + |\tilde{\mathcal{G}}(\lambda)|_{\mathbb{C}} \right) d}{\left(1 + \Re \tilde{\mathcal{G}}(\lambda) \right)^2 + \left(\Im \tilde{\mathcal{G}}(\lambda) \right)^2} \quad (\text{E31})$$

$$\Im \mathcal{G}(\lambda) = \frac{d}{\sigma^2} \frac{\Im \tilde{\mathcal{G}}(\lambda)}{\left(1 + \Re \tilde{\mathcal{G}}(\lambda) \right)^2 + \left(\Im \tilde{\mathcal{G}}(\lambda) \right)^2}. \quad (\text{E32})$$

We solved numerically equation (E28) using Mathematica™, using as input values of $\Re x$ ranging in $-5 < \Re x < 0$. The theoretical curves shown in figure 3 were obtained using the method described in this appendix to solve (29).

Appendix F: Further results for the hessian spectrum (II)

F.1. Crossover from non-trivial to trivial bulk

The spectrum of the Hessian is composed of a trivial and nontrivial bulk (see equation (33)). The nontrivial bulk spectral density ρ_{nt} is generated by the convolution of the Marchenko-Pastur term in (E19) with $P_\eta(\eta)$ in (31), while the trivial contribution is given by the free convolution of the Dirac delta term in (E19) with (31).

Let us show how to obtain the trivial bulk from the spectral equation (29). In order to consider the effect of the pole $\lambda_p = \alpha/d$ in (E19) to the solution of (29), let us consider the behavior of the equation for $\Re g \rightarrow \infty$: by expanding, one finds

$$g(\lambda) \simeq (d-1) \int \frac{d\eta P_\eta(\eta)}{\eta - (\lambda - \alpha/d) - i\epsilon} \quad (\text{F1})$$

where we individuated the small parameter $\epsilon = \alpha \frac{\Im g}{(\Re g)^2} \ll 1$. In this limit, then one finds $(\frac{\epsilon}{x^2 + \epsilon^2} \simeq \pi \delta(x))$ and $\rho = \frac{1}{(d-1)\pi} \Im g$

$$\rho(\lambda) \approx P_\eta(\lambda - \alpha/d). \quad (\text{F2})$$

This regime is expected to hold with very good approximation for sufficiently large λ , being more accurate the larger the λ .

There are two ways to define the crossover value: the first is based on the weight born by the singular term in (E19): by comparison, it is easy to see that the equation for $\hat{\lambda}$ reads

$$1 - \frac{\alpha}{d-1} = \int_{\lambda_*}^{\infty} d\lambda P_\eta(\lambda - \alpha/d). \quad (\text{F3})$$

Equation (F3) assigns a weight equal to that of the pole to the total probability of being in the region where the approximation (F1) holds. Alternatively, one can estimate the crossover value from the breaking of the regime where (F3) holds, i.e. $\epsilon = 1$ or

$$1 = \alpha \frac{\Im g(\hat{\lambda}_*)}{(\Re g(\hat{\lambda}_*))^2}. \quad (\text{F4})$$

F.2. Asymptotic expansions of the spectrum close to the lower edge

F.2.1. Expansion of $\rho(\lambda)$ for $\lambda \rightarrow 0_+$ and $d > 2$. Let us derive the behavior of the spectral density and the eigenvector moments in the vicinity of the lower edge $\lambda = 0$. We consider the spectrum in the case of both Mattis states and spin-glass states. In the following paragraphs, we will use the notation $\langle \cdot \rangle \equiv \int dh P_h(h) \langle \cdot \rangle$ for averages over cavity fields. Before doing the expansion, we show that the lower spectral edge must be $\lambda = 0$. We consider only the case of Mattis states, as for spin-glass states of vector spin glass systems it is well known that the spectrum is gapless.

We can easily show for Mattis states that $\lambda_{\min} = 0$ in the thermodynamic limit, using a *ab absurdum* short proof: suppose that $\rho(\lambda) = \pi \Im \mathcal{G}(\lambda) = 0$ for small enough $\lambda > 0$, which is equivalent to asking that the spectrum be gapped. By expanding equation (29) for $\lambda \rightarrow 0$, we get

$$\Re \mathcal{G}(\lambda) \simeq \frac{(d-1)}{\sigma} \int_0^{\infty} \frac{dh P_h(h)}{h - \frac{\lambda\sigma}{\Lambda_0}} \quad (\text{F5})$$

where Λ_0 is the $T = 0$ Replicon defined in equation (27). Since $\Lambda_0 > 0$ for a Mattis state, the integral in (F5) exists if and only if $\lambda = 0$, contradicting the initial condition $\Im\mathcal{G}(\lambda) = 0$ for $\lambda > 0$. It follows that one must necessarily have $\rho(\lambda) = \pi\Im\mathcal{G}(\lambda) > 0$ for any $\lambda > 0$.

We proceed with the expansion.

- Mattis states: here the Replicon, as defined by equation (27), is strictly positive. Let us begin: by plugging equation (E25) into (E24), we get

$$\Re\mathcal{G}(\lambda) = \frac{(d-1)}{\sigma} \left\langle \frac{h}{|h+x|_{\mathbb{C}}^2} \right\rangle + \frac{\sigma \left| \frac{d}{\sigma^2} - \mathcal{G}(\lambda) \right|_{\mathbb{C}}^2 \Re x(\lambda)}{\alpha} \quad (\text{F6})$$

where

$$x(\lambda) = \frac{\alpha\sigma}{d} \frac{\chi}{\frac{d}{\sigma^2} - \chi} - \lambda\sigma - \frac{\alpha\sigma}{d} \frac{\mathcal{G}(\lambda)}{\frac{d}{\sigma^2} - \mathcal{G}(\lambda)}. \quad (\text{F7})$$

Now, knowing that $\Re\mathcal{G}(\lambda) = \chi + \chi_{\text{SG}}\lambda + \mathcal{O}(\lambda^2)$, and $\chi = \frac{(d-1)}{\sigma} \langle 1/h \rangle$, after a few straightforward simplifications equation (F6) rewrites

$$- \left(\chi_{\text{SG}}\lambda + \frac{\sigma \left(\frac{d}{\sigma^2} - \chi \right)^2}{\alpha} \Re x(\lambda) \right) = \frac{(d-1)}{\sigma} |x(\lambda)|_{\mathbb{C}}^2 \left\langle \frac{1}{h|h+x|_{\mathbb{C}}^2} \right\rangle. \quad (\text{F8})$$

Since the r.h.s. of this last equation is positive, the l.h.s. is positive if and only if $\Re x < 0$: this is indeed the case. Indeed, $x(0) = 0$ and one can show that

$$\chi_{\text{SG}} = \frac{d\mathcal{G}}{d\lambda}(0) = -\frac{(d-1)}{\sigma} \left\langle \frac{1}{h^2} \right\rangle \frac{d\Re x}{d\lambda}(0) \quad (\text{F9})$$

so that

$$\frac{d\Re x}{d\lambda}(0) = -\frac{\sigma\chi_{\text{SG}}}{(d-1)\langle \frac{1}{h^2} \rangle} = -\frac{\sigma}{\Lambda}. \quad (\text{F10})$$

We have then that the l.h.s. of equation (F8) is

$$\begin{aligned} -\chi_{\text{SG}}\lambda - \frac{\sigma \left(\frac{d}{\sigma^2} - \chi \right)^2}{\alpha} \Re x &\simeq -\frac{\sigma^2 \left(\frac{d}{\sigma^2} - \chi \right)^2}{\alpha\Lambda} (1 - \Lambda)\lambda + \frac{\sigma^2 \left(\frac{d}{\sigma^2} - \chi \right)^2}{\alpha\Lambda} \lambda \\ &= \frac{\sigma^2 \left(\frac{d}{\sigma^2} - \chi \right)^2}{\alpha} \lambda. \end{aligned} \quad (\text{F11})$$

With this, equation (F8) rewrites

$$1 = \frac{\alpha(d-1)|x|_{\mathbb{C}}^2}{\sigma^3 \left(\frac{d}{\sigma^2} - \chi \right)^2 \lambda} \left\langle \frac{1}{h|h+x|_{\mathbb{C}}^2} \right\rangle. \quad (\text{F12})$$

To evaluate the integral in the r.h.s. of this last expression, we assume that $|\Im x| \ll |\Re x|$: by doing this, we can evaluate the integral easily using $\frac{\epsilon}{x^2 + \epsilon^2} \simeq \pi \delta(x)$

$$\int \frac{dh P_h(h)}{h|h+x|^2} = \frac{\pi}{|\Im x|} \int \frac{|\Im x| P_h(h) dh}{\pi h \left[(h + \Re x)^2 + (\Im x)^2 \right]} \simeq \frac{\pi}{|\Im x|} \frac{P_h(|\Re x|)}{|\Re x|}. \quad (\text{F13})$$

By combining the last two equations and expanding until leading order, we finally get ($|\Im x| \simeq \frac{\alpha}{\sigma^2(\frac{d}{\sigma^2}-\chi)^2} \Im \mathcal{G}$ and $\rho = \frac{\Im \mathcal{G}}{\pi(d-1)}$)

$$\rho(\lambda) \simeq \frac{\sigma}{\Lambda} P_h \left(\frac{\lambda\sigma}{\Lambda} \right) = a_d \lambda^{d-1} \tag{F14}$$

$$a_d = \frac{\mathcal{S}_{d-1}(\sigma) \sigma e^{-\frac{m^2(d-1)\langle 1/h^2 \rangle \sigma^2}{2(1-\Lambda)}} (1-\Lambda)^{d/2}}{[2\pi(d-1)\langle 1/h^2 \rangle]^{d/2} \Lambda^d}$$

where we used the identity $\alpha r_0 = \frac{1-\Lambda}{(d-1)\langle 1/h^2 \rangle}$ coming from equation (27). The prefactor a_d is exponentially suppressed as $\Lambda \rightarrow 1$, which corresponds to the sub-linear capacity limit $\alpha \rightarrow 0$: this corresponds to the trivialization of the spectrum in that limit, $\rho(\lambda) \rightarrow \delta(\lambda - 1)$.

Let us check *a posteriori* the validity of the assumption $|\Im x| \ll |\Re x|$: combining equation (F13) and $\Re x \simeq -\lambda/\Lambda$, we get

$$\frac{|\Im x|}{|\Re x|} \simeq b_d \lambda^{d-2} \ll 1 \implies \lambda_{cr} = \frac{1}{b_d^{d-2}} \tag{F15}$$

$$b_d = \frac{\pi\alpha(d-1)a_d\Lambda}{\sigma^4\left(\frac{d}{\sigma^2}-\chi\right)^2}.$$

The prefactor b_d is proportional to a_d , therefore the crossover eigenvalue λ_{cr} grows exponentially fast as $\alpha \rightarrow 0$, consistently with the trivialization of the spectrum in this limit $\rho(\lambda) \rightarrow \delta(\lambda - 1)$ (with a delta peak spectrum, the imaginary part of x is identically null for any λ).

We finally complete this section with the estimation of the pseudogap eigenvalue. To do that, we rearrange the terms in equation (F7), obtaining for small $|x|_{\mathbb{C}}$

$$-\lambda - \Lambda_0 x = \left(\frac{d-1}{\sigma}\right) \alpha r_0 \left[\left\langle \frac{1}{h^3} \right\rangle + \left(\frac{d-1}{\sigma}\right) \sigma \sqrt{r_0} \left\langle \frac{1}{h^2} \right\rangle^2 \right] x^2. \tag{F16}$$

By neglecting the imaginary part of x , we can solve equation (F16) as a second order algebraic equation for its real part: to estimate the pseudogap, we compute the value of λ that cancels the discriminant of said equation, obtaining

$$\lambda_{pg} = \frac{\Lambda_0^2}{4W} \tag{F17}$$

where W is the coefficient of the x^2 term in equation (F16). In the $d=3$ case, the term $\langle 1/h^3 \rangle$ is divergent and the coefficient is replaced by a singular term $W \propto |\log \Lambda_0|$.

- Spin glass states: although we do not know the fRSB expression of the pdf of local fields on a local minimum of the energy, we can altogether characterize the scaling of $\rho(\lambda)$ at small eigenvalues. In order to obtain a physical solution for the spectrum, we need to impose $\Lambda_0 = 0$ for any $\alpha > 0$: in fact, the unphysical solution with a spectral gap in the right panel of figure 3 is a consequence of having $\Lambda_0 < 0$ from equation (27) when using the RS $P_h(h)$. In the following lines, the average over the cavity fields is assumed to be over the fRSB pdf.

Let us consider equation (F6): by assuming $|\Re x| \ll |\Im x|$ and $|\Re x| = \mathcal{O}(\Im x)^2$, the integral can be expanded and we get at leading order

$$\Re \mathcal{G}(\lambda) = \chi - \frac{2\Re x}{\alpha r_0 \sigma} - \frac{d-1}{\sigma} \left\langle \frac{1}{h^3} \right\rangle (\Im x)^2 + \frac{\sigma}{\alpha} \left| \frac{d}{\sigma^2} - \mathcal{G}(\lambda) \right|^2 \Re x. \quad (\text{F18})$$

In deriving this last expression we imposed the nullity of the Replicon in (27) and thus set $1 = \alpha r(d-1)\langle 1/h^2 \rangle$. Here, $\langle \cdot \rangle$ denotes an average over the unknown $P_h^{\text{RSB}}(h)$. Note that the relation between the distribution of local fields and that of cavity fields in (31) still holds, though with a value of the Onsager reaction dictated by the unknown fRSB solution and a different expression of the fRSB distribution of cavity fields.

If we further assume that $\Re \mathcal{G}(\lambda) = \chi + o(\Im x)^2$ and thus $\Re x = -\lambda + o(\lambda)$, we get at leading order

$$\frac{d-1}{\sigma} \left\langle \frac{1}{h^3} \right\rangle (\Im x)^2 = \frac{\lambda}{\alpha r_0 \sigma} \implies \rho(\lambda) = \frac{1}{\pi} \sqrt{\frac{\lambda}{(\alpha r_0)^3 (d-1) \langle 1/h^3 \rangle}} \equiv \hat{a}_d \sqrt{\lambda} \quad (\text{F19})$$

where we used $|\Im x| = \alpha r_0 |\Im \mathcal{G}|$ and $\rho = \frac{1}{\pi} |\Im \mathcal{G}|$. The prefactor \hat{a}_d is well defined if $P_h^{\text{RSB}}(h) = o(h^2)$ for small h . On contrary, it should be replaced by a λ -dependent singular term. Assuming that $P_h^{\text{RSB}}(h)$ retains the low cavity field of the RS solution, $P_h^{\text{RSB}}(h) \approx h^{d-1}$, we get the following scaling for the spectral density at low eigenvalues

$$\rho(\lambda) \approx \begin{cases} \sqrt{\lambda} & d > 3 \\ \sqrt{\frac{\lambda}{|\log \lambda|}} & d = 3. \end{cases} \quad (\text{F20})$$

F.2.2. Expansion of $\rho(\lambda)$ for $\lambda \rightarrow 0_+$ for $d=2$. In the $d=2$ case the integral defining the Replicon in equation (27) is logarithmically divergent to $-\infty$, if close to $h=0$ the scaling $P_h(h) \propto h$ is assumed. Given that the retrieval phase has fRSB as well, we expect both for Mattis and spin glass states to have $\Lambda = 0$ and thus $P_h^{\text{RSB}}(h) \ll h$ close to $h=0$.

Here we cannot find the scaling of the spectrum without knowing the fRSB distribution of fields. We conjecture that in the case of Mattis states the behavior of the exponent is continuous, so that one has by comparing with equation (F14)

$$\rho(\lambda) \approx \lambda \quad d = 2. \quad (\text{F21})$$

For spin glass minima, instead we conjecture a behavior akin to that of the $d=3$ case, with a square root behavior with logarithmic corrections.

F.2.3. Expansion of $\kappa(\lambda)$ for $\lambda \rightarrow 0_+$. In this section, we expand equation (38) for the rescaled IPR $\kappa(\lambda) = NI(\lambda)$ close to the lower edge.

Preliminary to that, we derive equation (38) starting from equation (40), which we report here for the convenience of the readers

$$N \langle |\psi_i^{(k)}|^2 \rangle = \frac{\alpha/\sigma^2}{\left| \frac{d}{\sigma^2} - \mathcal{G}(\lambda_k) \right|_{\mathbb{C}}^2} \frac{d-1}{\left| \eta_i - \lambda_k \sigma - \frac{\alpha \sigma}{d} \frac{\mathcal{G}(\lambda_k)}{\frac{d}{\sigma^2} - \mathcal{G}(\lambda_k)} \right|_{\mathbb{C}}^2} \quad (\text{F22})$$

where $|\cdot|_{\mathbb{C}}$ is the complex norm. Equation (40) should be understood as an equation for the average value of eigenvector weights at fixed local fields η_k . Indeed, eigenvector components in deformed ensembles $\psi_{i,a}^j$ for a fixed configuration of η_k are Gaussian variables with variance

given by (40) with a factor $1/d$ (see [62] for the case of deformed Wigner ensembles). The rescaled IPR $\kappa(\lambda) = NI(\lambda)$ can be computed combining equations (37) and (40) and exploiting the Gaussian nature of the $\psi_{i,a}^j$

$$\kappa(\lambda) = N \sum_{i=1}^N \langle |\psi_i(\lambda)|^4 \rangle = 3N \sum_{i=1}^N \langle |\psi_i(\lambda)|^2 \rangle^2 \tag{F23}$$

$$\stackrel{N \rightarrow \infty}{\approx} 3 \left(1 - \frac{1}{d}\right) \left[\frac{(d-1)\alpha}{\sigma^2 \left| \frac{d}{\sigma^2} - \mathcal{G}(\lambda) \right|_{\mathbb{C}}^2} \right]^2 \int \frac{d\eta P_\eta(\eta)}{\left| \eta - \lambda\sigma - \frac{\alpha\sigma}{d} \frac{\mathcal{G}(\lambda)}{\frac{d}{\sigma^2} - \mathcal{G}(\lambda)} \right|_{\mathbb{C}}^4}.$$

We now proceed to expand (F23) in the limit $\lambda \rightarrow 0_+$. We consider the case of Mattis states of $d > 2$ VHM. By proceeding similarly to the previous section, we find

$$\kappa(\lambda) \approx \frac{3(d-1)^3}{d} (\alpha r_0)^2 \frac{\pi P_h(-\Re x)}{2(\Im x)^3} \propto \lambda^{-2(d-1)} \tag{F24}$$

where we defined x in (F7). The scaling in (F24) leads to $I(\lambda) \propto N\lambda^{2(d-1)} \propto N^{1-\frac{2}{d}}$ for the IPR, where we used the scaling $\lambda \sim N^{-1/d}$, valid for lower edge eigenvalues. Since by definition the IPR is bounded by unit (compare with the definition in (37)), formula (F23) cannot be true for the softest modes. To highlight the origin of this paradox, consider equation (40). Applying it to the normalization of eigenvectors $\sum_k |\psi_k^j|^2 = 1$, it returns

$$1 = \frac{\alpha(d-1)}{\sigma^2 \left| \frac{d}{\sigma^2} - \mathcal{G}(\lambda) \right|_{\mathbb{C}}^2} \int \frac{d\eta P_\eta(\eta)}{\left| \eta - \lambda\sigma - \frac{\alpha\sigma}{d} \frac{\mathcal{G}(\lambda)}{\frac{d}{\sigma^2} - \mathcal{G}(\lambda)} \right|_{\mathbb{C}}^2}. \tag{F25}$$

For any $\lambda > 0$, this identity is valid: it is just the imaginary part of the spectral equation (29). However, at $\lambda = 0$ the r.h.s. of (F25) is equal to $1 - \Lambda_0$, where $\Lambda_0 < 1$ is the $T = 0$ Replicon introduced by equation (27). Since for $d > 2$ one has $\Lambda_0 > 0$ for Mattis states, it follows that (F25) is wrong in the limit $\lambda \rightarrow 0_+$. The solution of such an impasse is that soft modes of Mattis states undergo condensation: by assuming the existence of a condensate weight $|\psi_{i_*}(0)|^2 = \mathcal{O}(1)$ for some node i_* , we get by using equation (40) with the condensate weight

$$1 = |\psi_{i_*}(0)|^2 + \frac{\alpha(d-1)}{\sigma^4 \left[\frac{d}{\sigma^2} - \mathcal{G}(0) \right]^2} \int \frac{d\eta P_\eta(\eta)}{\left[\eta - \frac{\alpha}{d} \frac{\mathcal{G}(0)}{\frac{d}{\sigma^2} - \mathcal{G}(0)} \right]^2}. \tag{F26}$$

The second addendum in the rhs of this last equation is just $1 - \Lambda_0$ (see equations (31) and (27)); therefore, we find that the condensate is equal to the Replicon $|\psi_{i_*}(0)|^2 = \Lambda_0$.

The presence of condensation in soft modes of continuous disordered systems is mathematically akin to a Bose–Einstein condensation, and is sometimes referred in literature as Random Matrix Bose–Einstein Condensation (RMBEC) [48]. The RMBEC condition of our model can be obtained from equation (40) by imposing $|\psi_{i_*}^{(k)}|^2 = \mathcal{O}(1)$:

$$\frac{\alpha}{\sigma^2 \left| \frac{d}{\sigma^2} - \mathcal{G}(\lambda_k) \right|_{\mathbb{C}}^2} \frac{d-1}{\left| \eta_{i_*} - \lambda_k\sigma - \frac{\alpha\sigma}{d} \frac{\mathcal{G}(\lambda_k)}{\frac{d}{\sigma^2} - \mathcal{G}(\lambda_k)} \right|_{\mathbb{C}}^2} = \mathcal{O}(N) \tag{F27}$$

that is valid in the regime $\lambda_k \sim N^{-1/d}$. It can be shown that equation (F27) implies an order statistics between low-rank eigenvalues and low-rank cavity fields of the kind $\lambda \sim \Lambda_0 h$ [27, 46].

We can estimate the scale of eigenvalues in N at which formula (F23) does not need to be rectified with RMBEC. Let us consider lower edge modes with subextensive rank $k = \tilde{k}N^a$ for $0 < a < 1$: then $\lambda_k = \tilde{\lambda}N^{-\frac{1-a}{d}}$ and from equation (F24) it follows that RMBEC disappears at ranks with an exponent

$$\lambda^{-2(d-1)} \leq N \Rightarrow a_* = \frac{d-2}{2(d-1)}. \quad (\text{F28})$$

For $a > a_*$, one has $I(\lambda) = \kappa(\lambda)/N \leq 1$ and therefore for modes with rank $k \gg N^{a_*}$ equation (F23) is correct.

F.2.4. Limit $d \rightarrow \infty$ for the spectrum. Let us consider the large d limit of the spectral equation (29) in the two regimes $P \propto Nd$ and $P \propto N/d$ described in appendix C. Let us begin with the first, where the capacity is $\alpha' = \frac{\alpha}{d}$. In this case, one has $P_h(h) \rightarrow \delta(h - \sqrt{m^2 + \alpha' r'})$, and the spectral equation becomes (again, we use $\sigma = \sqrt{d}$ in the derivation)

$$\mathcal{G}(\lambda) = \frac{1}{\sqrt{m^2 + \alpha' r'_0 + \frac{\alpha' \chi}{1-\chi} - \lambda - \frac{\alpha' \mathcal{G}(\lambda)}{1-\mathcal{G}(\lambda)}}}. \quad (\text{F29})$$

Considering the no-retrieval physical solution $m = 0$, $\chi = \frac{1}{1+\sqrt{\alpha'}}$, knowing that $r_0 = \frac{1}{(1-\chi)^2}$, the solution of (F29) reads

$$\mathcal{G}(\lambda) = \frac{2(1 + \sqrt{\alpha'}) - \lambda}{2(1 + 2\sqrt{\alpha'} + \alpha' - \lambda)} + \frac{\sqrt{\lambda(\lambda - 4\sqrt{\alpha'})}}{2(1 + 2\sqrt{\alpha'} + \alpha' - \lambda)}. \quad (\text{F30})$$

The spectrum is given by the continuous bulk (the imaginary part of (F30)) and the pole in $\lambda = 1 + 2\sqrt{\alpha'} + \alpha'$, which has weight $1 - \alpha'$ and exists for $0 \leq \alpha' \leq 1$. The final expression of the spectral density is thus

$$\rho(\lambda) = \frac{\sqrt{\lambda(4\sqrt{\alpha'} - \lambda)}}{2(1 + 2\sqrt{\alpha'} + \alpha' - \lambda)} + (1 - \alpha') \vartheta(1 - \alpha') \delta(\lambda - 1 - 2\sqrt{\alpha'} - \alpha') \quad (\text{F31})$$

This is the spectrum that one obtains for a $p = 2$ spherical Hopfield model: this system is marginally stable having $\Lambda_0 = 0$, as it should be since the lower tail in equation (F31) is $\propto \sqrt{\lambda}$. As a consequence, the corresponding eigenvectors are fully delocalized with $\kappa(\lambda) = \mathcal{O}(1)$ for all $0 \leq \lambda \leq 4\sqrt{\alpha'} + \alpha'$ or $\lambda = 1 + 2\sqrt{\alpha'} + \alpha'$.

In the second case the capacity reads $\alpha'' = \alpha d$. The spectrum becomes trivial for $d \rightarrow \infty$: again, $P_h(h) \rightarrow \delta(h - \sqrt{m^2 + \alpha'' r''})$, but the spectral equation is simply

$$\mathcal{G}(\lambda) = \frac{1}{\sqrt{m'^2 + \alpha'' r''_0} - \lambda} \quad (\text{F32})$$

In the SubLin regime, one has at equilibrium $\sqrt{m^2 + \alpha'' r''_0} = 1$ and therefore $\rho(\lambda) = \delta(\lambda - 1)$, as in the $\alpha = 0$ case for finite d .

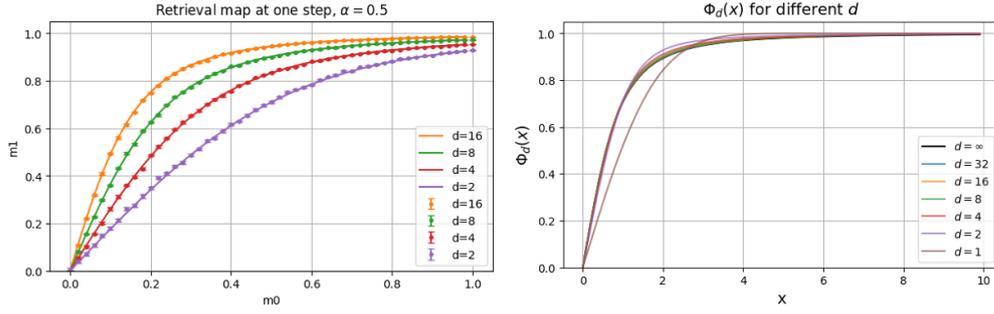


Figure 13. Left Overlap m_1 after one step, starting from an initial overlap m_0 : comparison between numerical results [points] and analytical prediction from STN [solid lines]. Right Function $\Phi_d(x)$ direct dependence on d is weak already from modest values of d : the main dependence is through the argument $x = m_0/\sqrt{\alpha'}$.

Appendix G. Details on signal-to-noise (STN) method to predict first-step denoising

(a) *Patterns.* Let us consider a fixed point \underline{S} of equation (5) with magnetization m with the pattern $\underline{\xi}^1$: the local field vector on this point reads

$$\frac{1}{\sigma} \sum_{j:j \neq i} \mathbb{J}_{ij} S_j = \frac{1}{N\sigma} (\underline{\xi}^1 \cdot \underline{S}) \xi_i^1 + \frac{1}{N\sigma} \sum_{\mu=2}^P (\underline{\xi}^\mu \cdot \underline{S}) \xi_i^\mu + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right) \simeq m \xi_i^1 + \mathbf{n}_i. \quad (\text{G1})$$

The vector \mathbf{n}_i is an isotropic Gaussian vector with zero mean and variance $\frac{\alpha r_0}{\sigma^2}$ (compare with def. (9)), where r_0 is the order parameter defined in equation (25). Thanks to the stationarity condition (3), the spins are $S_i = \frac{(m \xi_i^1 + r_i) \sigma}{|m \xi_i^1 + r_i|}$, hence the Mattis magnetization $m = \frac{1}{N\sigma} \underline{S} \cdot \underline{\xi}^1$ for large N satisfies

$$\begin{aligned} m &= \frac{1}{N\sigma} \sum_{i=1}^N \xi_i^1 \cdot S_i = \int \frac{d\mathbf{n}}{(2\pi \alpha r / \sigma^2)^{d/2}} e^{-\frac{n^2 \sigma^2}{2\alpha r_0}} \int \frac{d\xi \delta(|\xi| - 1)}{\mathcal{S}_{d-1}(1)} \frac{m \xi + \mathbf{n}}{|m \xi + \mathbf{n}|} \cdot \xi \\ &= \int_0^\infty dn P_h(n) g_d\left(\frac{m\sigma n}{\alpha r_0}\right) = \Phi_d\left(\frac{m}{\sqrt{\alpha r_0}}\right) \end{aligned} \quad (\text{G2})$$

and we derived again the third of equation (25). We introduced a function $\Phi_d(\cdot)$ to represent in a compact fashion the integral in the last r.h.s. of (G2). For the standard Hopfield model, when $d = 1$, one can show that $\Phi_1(\cdot) = \text{Erf}\left(\frac{\cdot}{\sqrt{2}}\right)$.

We remark that the noise vectors \mathbf{n}_i have the same statistics as the cavity fields of the RS solution, but they are not exactly cavity vectors since in cavity theory the cavity fields \mathbf{h}_i vectors must be parallel to the spins S_i .

(b) *Mixture states.* We proceed in a similar way as in the previous paragraph, but now consider as initial condition $\frac{1}{\sigma} \mathbf{S}_i(0) = m_0 \frac{\mathbf{z}_i}{z_i} + \sqrt{1 - m_0^2} \boldsymbol{\pi}_i$, where $\boldsymbol{\pi}_i$ are unit vectors satisfying $\boldsymbol{\pi}_i \cdot \mathbf{z}_i = 0$ and $\mathbf{z}_i = \sum_{\mu=1}^s c_\mu \boldsymbol{\xi}_i^\mu$ with $s = \mathcal{O}(1)$. We found

$$\begin{aligned} \frac{1}{\sigma} \sum_{j:j \neq i} \mathbb{J}_{ij} \mathbf{S}_j(0) &= \tilde{m}_0 \mathbf{z}_i + \mathbf{n}_i \\ m_1 &= \frac{1}{N} \sum_{i=1}^N \frac{\tilde{m}_0 \mathbf{z}_i + \mathbf{n}_i}{|\tilde{m}_0 \mathbf{z}_i + \mathbf{n}_i|} \cdot \frac{\mathbf{z}_i}{z_i} \\ &= \mathcal{S}_{d-1}(1) \int_0^s dz z^{d-1} \rho_s(z) \Phi_d \left(\frac{\tilde{m}_0 z}{\sqrt{\alpha/d}} \right). \end{aligned} \quad (\text{G3})$$

where the noise vector is again an isotropic Gaussian vector with zero mean and variance α/d^2 , $\rho_s(z)$ is the pdf of the vector $\mathbf{z} = \sum_{\mu=1}^s c_\mu \boldsymbol{\xi}_\mu$ and $\tilde{m}_0 = \langle 1/z \rangle m_0$. We discuss how to compute the distributions $\rho_s(z)$ in appendix B.

References

- [1] Peretto P and Niez J 1986 Long term memory storage capacity of multiconnected neural networks *Biol. Cybern.* **54** 53
- [2] Gardner E 1987 Multiconnected neural network models *J. Phys. A: Math. Gen.* **20** 3453
- [3] Krotov D and Hopfield J J 2016 Dense associative memory for pattern recognition *Advances in Neural Information Processing Systems* vol 29
- [4] Demircigil M, Heusel J, Löwe M, Uppgang S and Vermet F 2017 On a model of associative memory with huge storage capacity *J. Stat. Phys.* **168** 288–99
- [5] Lucibello C and Mézard M 2024 Exponential capacity of dense associative memories *Phys. Rev. Lett.* **132** 077301
- [6] Ramsauer H et al 2020 Hopfield networks is all you need (arXiv:2008.02217)
- [7] Ambrogioni L 2024 In search of dispersed memories: generative diffusion models are associative memory networks *Entropy* **26** 381
- [8] Krotov D 2023 A new frontier for Hopfield networks *Nat. Rev. Phys.* **5** 366
- [9] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez A N, Kaiser L u and Polosukhin I 2017 Attention is all you need *Advances in Neural Information Processing Systems* vol 30, ed I Guyon, U V Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan and R Garnett (Curran Associates, Inc.)
- [10] Hoover B, Liang Y, Pham B, Panda R, Strobel H, Chau D H, Zaki M and Krotov D 2023 Energy transformer *Advances in Neural Information Processing Systems* vol 36, ed A Oh, T Naumann, A Globerson, K Saenko, M Hardt and S Levine (Curran Associates, Inc.) pp 27532–59
- [11] Devlin J, Chang M-W, Lee K and Toutanova K 2019 Bert: pre-training of deep bidirectional transformers for language understanding (arXiv:1810.04805 [cs])
- [12] Jumper J et al. 2021 Highly accurate protein structure prediction with alphafold *Nature* **596** 583
- [13] D’Amico F and Negri M 2024 Self-attention as an attractor network: transient memories without backpropagation 2024 *IEEE Workshop on Complexity in Engineering (COMPENG)* pp 1–6
- [14] Smart M, Bietti A and Sengupta A M 2025 In-context denoising with one-layer transformers: connections between attention and associative memory retrieval (arXiv:2502.05164 [cs])
- [15] Kalaj S, Lauditi C, Perugini G, Lucibello C, Malatesta E M and Negri M 2024 Random features Hopfield networks generalize retrieval to previously unseen examples (arXiv:2407.05658 [cond-mat.dis-nn])
- [16] OpenAI 2022 Introducing ChatGPT (available at: <https://openai.com/index/chatgpt/>)
- [17] Bal M 2023 Spin-model transformers (available at: <https://mcbal.github.io/post/spin-model-transformers>)

- [18] Hopfield J J 1982 Neural networks and physical systems with emergent collective computational abilities *Proc. Natl Acad. Sci.* **79** 2554
- [19] Zhou C and Liu L 1999 Vector-product Hopfield model *Opt. Commun.* **168** 445–55
- [20] Kymn C, Bybee C, Yun Z, Kleyko D and Olshausen B 2025 Oscillator associative memories facilitate high-capacity, compositional inference *ICLR Workshop on New Frontiers in Associative Memories*
- [21] Noest A 1987 Phasor neural networks *Neural Information Processing Systems* ed D Anderson (American Institute of Physics)
- [22] Frady E P, Kent S J, Olshausen B A and Sommer F T 2020 Resonator networks: I. An efficient solution for factoring high-dimensional, distributed representations of data structures *Neural Comput.* **32** 2311
- [23] Kent S J., Frady E P., Sommer F T. and Olshausen B A. 2020 Resonator networks: II. Factorization performance and capacity compared to optimization-based methods *Neural Comput.* **32** 2332
- [24] Miyato T, Löwe S, Geiger A and Welling M 2025 Artificial Kuramoto oscillatory neurons (arXiv:2410.13821 [cs.LG])
- [25] Amit D J, Gutfreund H and Sompolinsky H 1987 Statistical mechanics of neural networks near saturation *Ann. Phys., NY* **173** 30
- [26] Clark D G 2025 Transient dynamics of associative memory models (arXiv:2506.05303 [cond-mat.dis-nn])
- [27] Franz S, Nicoletti F, Parisi G and Ricci-Tersenghi F 2022 Delocalization transition in low energy excitation modes of vector spin glasses *SciPost Phys.* **12** 016
- [28] Franz S, Nicoletti F and Ricci-Tersenghi F 2022 Linear low energy excitations in fully-connected models of glasses *J. Stat. Mech.* **053302**
- [29] Amit D J, Gutfreund H and Sompolinsky H 1985 Spin-glass models of neural networks *Phys. Rev. A* **32** 1007
- [30] Amit D J, Gutfreund H and Sompolinsky H 1985 Storing infinite numbers of patterns in a spin-glass model of neural networks *Phys. Rev. Lett.* **55** 1530
- [31] Cicuta G M, Krausser J, Milkus R and Zaccone A 2018 Unifying model for random matrix theory in arbitrary space dimensions *Phys. Rev. E* **97** 032113
- [32] Benetti F P C, Parisi G, Pietracaprina F and Sicuro G 2018 Mean-field model for the density of states of jammed soft spheres *Phys. Rev. E* **97** 062157
- [33] Patil N, Aguirre-López F and Bouchaud J-P 2024 The spectral boundary of block structured random matrices *J. Phys. Complex.* **5** 035001
- [34] Amit D J 1989 *Modeling Brain Function: The World of Attractor Neural Networks* (Cambridge University Press)
- [35] Negri M, Lauditi C, Perugini G, Lucibello C and Malatesta E 2023 Storage and learning phase transitions in the random-features Hopfield model *Phys. Rev. Lett.* **131** 257301
- [36] Mézard M, Parisi G and Virasoro M A 1987 *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications* vol 9 (World Scientific Publishing Company)
- [37] Talagrand M 2006 The Parisi formula *Ann. Math.* **163** 221
- [38] Mézard M and Montanari A 2009 *Information, Physics and Computation* (Oxford University Press)
- [39] Mézard M 1989 The space of interactions in neural networks: Gardner’s computation with the cavity method *J. Phys. A: Math. Gen.* **22** 2181
- [40] Mézard M 2017 Mean-field message-passing equations in the Hopfield model and its generalizations *Phys. Rev. E* **95** 022117
- [41] Cugliandolo L F and Dean D S 1995 Full dynamical solution for a spherical spin-glass model *J. Phys. A: Math. Gen.* **28** 4213
- [42] Niedda J, Tonolo T and Gradenigo G 2024 Probing marginal stability in the spherical $p = 2$ model (arXiv:2403.15819 [cond-mat.dis-nn])
- [43] Amit D J, Gutfreund H and Sompolinsky H 1987 Information storage in neural networks with low levels of activity *Phys. Rev. A* **35** 2293
- [44] Bollé D, Nieuwenhuizen T M, Castillo I P and Verbeiren T 2003 A spherical Hopfield model *J. Phys. A: Math. Gen.* **36** 10269
- [45] Bouchaud J-P and Potters M 2020 *A First Course in Random Matrix Theory: for Physicists, Engineers and Data Scientists* (Cambridge University Press)
- [46] Lee J O and Schnell K 2016 Extremal eigenvalues and eigenvectors of deformed Wigner matrices *Probab. Theory Relat. Fields* **164** 165

- [47] Rainone C, Urbani P, Zamponi F, Lerner E and Bouchbinder E 2021 Mean-field model of interacting quasilocated excitations in glasses *SciPost Phys. Core* **4** 008
- [48] Ikeda H 2023 Bose-Einstein-like condensation of deformed random matrix: a replica approach *J. Stat. Mech.* **023302**
- [49] Anderson P W 1958 Absence of diffusion in certain random lattices *Phys. Rev.* **109** 1492
- [50] Venturelli D, Cugliandolo L F, Schehr G and Tarzia M 2023 Replica approach to the generalized Rosenzweig-Porter model *SciPost Phys.* **14** 110
- [51] Turner E, Dabholkar K and Barak O 2021 Charting and navigating the space of solutions for recurrent neural networks (arXiv:2111.09356 [quant-ph])
- [52] Brown T B *et al* 2020 Language models are few-shot learners (arXiv:2005.14165)
- [53] Touvron H *et al* 2023 Llama: Open and efficient foundation language models (arXiv:2302.13971 [cs.CL])
- [54] D'Amico F, Rossi S, del Bono L M and Negri M 2025 Pseudo-likelihood produces associative memories able to generalize, even for asymmetric couplings *New Frontiers in Associative Memories*
- [55] Serricchio L, Bocchi D, Chilin C, Marino R, Negri M, Cammarota C and Ricci-Tersenghi F 2025 Daydreaming Hopfield networks and their surprising effectiveness on correlated data *Neural Netw.* **186** 107216
- [56] Benedetti M, Ventura E, Marinari E, Ruocco G and Zamponi F 2022 Supervised perceptron learning vs unsupervised hebbian unlearning: approaching optimal memory retrieval in hopfield-like networks *J. Chem. Phys.* **156** 104107
- [57] Agliari E, Alessandrelli A, Barra A, Centonze M S and Ricci-Tersenghi F 2025 Generalized hetero-associative neural networks *J. Stat. Mech.* **013302**
- [58] Agliari E, Alessandrelli A, Barra A, Centonze M S and Ricci-Tersenghi F 2025 Networks of neural networks: more is different (arXiv:2501.16789)
- [59] Lerner E and Bouchbinder E 2021 Low-energy quasilocated excitations in structural glasses *J. Chem. Phys.* **155** 200901
- [60] Moriel A, Lerner E and Bouchbinder E 2024 Experimental evidence for the ω^4 tail of the non-phononic spectra of glasses *J. Appl. Phys.* **136** 225105
- [61] Schirmacher W, Paoluzzi M, Mocanu F C, Khomenko D, Szamel G, Zamponi F and Ruocco G 2024 The nature of non-phononic excitations in disordered systems *Nat. Commun.* **15** 3107
- [62] Truong K and Ossipov A 2016 Eigenvectors under a generic perturbation: non-perturbative results from the random matrix approach *Europhys. Lett.* **116** 37002
- [63] Bapst V and Semerjian G 2011 Lifshitz tails on the bethe lattice: a combinatorial approach *J. Stat. Phys.* **145** 51
- [64] Slanina F 2012 Localization of eigenvectors in random graphs *Eur. Phys. J. B* **85** 1
- [65] Kühn R, Bös S and van Hemmen J L 1991 Statistical mechanics for networks of graded-response neurons *Phys. Rev. A* **43** 2084
- [66] Franz S, Lupo C, Nicoletti F, Parisi G and Ricci-Tersenghi F 2025 Soft modes in vector spin glass models on sparse random graphs *Phys. Rev. B* **111** 014203
- [67] Kanter I 1988 Potts-glass models of neural networks *Phys. Rev. A* **37** 2739
- [68] Naim M, Boboeva V, Kang C J and Treves A 2018 Reducing a cortical network to a Potts model yields storage capacity estimates *J. Stat. Mech.* **043304**
- [69] Boboeva V, Brasselet R and Treves A 2018 The capacity for correlated semantic memories in the cortex *Entropy* **20** 824
- [70] Abramowitz M and 1965 *Handbook of Mathematical Functions With Formulas, Graphs and Mathematical Tables* ed I A Stegun (Dover Publications)