

## Full Length Article

# Regularization, early-stopping and dreaming: A Hopfield-like setup to address generalization and overfitting

E. Agliari <sup>a,c,\*</sup>, F. Alemanno <sup>b,c</sup>, M. Aquaro <sup>a,c</sup>, A. Fachechi <sup>a,c</sup>

<sup>a</sup> Dipartimento di Matematica “Guido Castelnuovo”, Sapienza Università di Roma, Italy

<sup>b</sup> Dipartimento di Matematica, Università di Bologna, Italy

<sup>c</sup> GNFM-INdAM, Gruppo Nazionale di Fisica Matematica (Istituto Nazionale di Alta Matematica), Italy



## ARTICLE INFO

## Keywords:

Attractor neural networks  
Overfitting  
Spin-glasses

## ABSTRACT

In this work we approach attractor neural networks from a machine learning perspective: we look for optimal network parameters by applying a gradient descent over a regularized loss function. Within this framework, the optimal neuron-interaction matrices turn out to be a class of matrices which correspond to Hebbian kernels revised by a reiterated unlearning protocol. Remarkably, the extent of such unlearning is proved to be related to the regularization hyperparameter of the loss function and to the training time. Thus, we can design strategies to avoid overfitting that are formulated in terms of regularization and early-stopping tuning. The generalization capabilities of these attractor networks are also investigated: analytical results are obtained for random synthetic datasets, next, the emerging picture is corroborated by numerical experiments that highlight the existence of several regimes (i.e., overfitting, failure and success) as the dataset parameters are varied.

## 1. Introduction

The Hopfield model is probably the best-known example of attractor neural network (Amari, 1972; Hopfield, 1982; Little, 1974): this is constituted by a set of  $N$  binary units, referred to as neurons, that interact pairwise and whose state is iteratively updated by a non-linear activation function, in such a way that the new state of neuron  $i$  depends on the signal acting on  $i$  and stemming from the neighboring neurons. A suitable choice of the neuron interaction matrix, denoted as  $\mathbf{J} \in \mathbb{R}^{N \times N}$ , should ensure the attractivity of a number of patterns, that we want to store and retrieve. More precisely, one initializes the neural configuration by setting the values of the units close to a pattern  $\xi^1 \in \{-1, +1\}^N$ , this configuration represents the input supplied to the machine and may correspond to a corrupted version of  $\xi^1$ ; repeated updates of neurons are then performed until convergence to a fixed point and, if this matches  $\xi^1$ , this state is interpreted as the retrieval of the information codified by  $\xi^1$ . The same is expected to occur for any target pattern, say  $\xi^\mu$ , with  $\mu = 1, \dots, P$ .

The popularity of the Hopfield model is also due to the fact that it is feasible of an analytical treatment and, in particular, it can be recognized as a spin-glass, in such a way that it can benefit from a broad collection of techniques developed to address disordered systems. Indeed, in the last decades, the model has been intensively investigated, and countless variations on the theme have also been accounted for.<sup>1</sup> Remarkably, a significant fraction of these works spotlighted the structure of the neural interaction matrix: in the standard Hopfield model this is based on the so-called Hebb's rule (Hebb, 1949)  $\mathbf{J} = \frac{\xi_i \xi_j^T}{N}$  and suitable revisions of this rule can give rise to better performances of the model in terms of number of storable and retrievable patterns (Amit, 1989). A successful class of these revisions apply unlearning protocols (see e.g., Benedetti, Ventura, Marinari, Ruocco, & Zamponi, 2022; Del Giudice, Franz, & Virasoro, 1989; Dotsenko, Yarunin, & Dorotheyev, 1991; Fachechi, Agliari, & Barra, 2019; Franz, Amit, & Virasoro, 1990; Hopfield, Feinstein, & Palmer, 1983; Kanter & Sompolinsky, 1987; Marinari, 2019; Plakhov, Semenov, & Shuvalova,

\* Corresponding author at: Dipartimento di Matematica “Guido Castelnuovo”, Sapienza Università di Roma, Italy.

E-mail address: [agliari@mat.uniroma1.it](mailto:agliari@mat.uniroma1.it) (E. Agliari).

<sup>1</sup> These include, for instance, the presence of biases or vacancies in pattern entries (see e.g., Agliari, Barra, Galluzzi, Guerra, & Moauro, 2012; Agliari, Leonelli, & Marullo, 2022; Amit, Gutfreund, & Sompolinsky, 1987), generalizations of the Hebbian kernel to include temporal correlations among patterns (see e.g., Agliari, Barra, De Antoni, & Galluzzi, 2013; Agliari, Fachechi, & Marullo, 2020; Cugliandolo, 1993) or synaptic noise (see e.g., Agliari, Alemanno, Barra, Centonze, & Fachechi, 2020; Agliari & De Marzo, 2020; Camilli, Contucci, & Mingione, 2022), non-trivial underlying architectures (see e.g., Agliari, Annibale, Barra, Coolen, & Tantari, 2013; Agliari, Migliozi, & Tantari, 2018; Wemmenhove & Coolen, 2003), or high-order interactions (see e.g., Agliari, Alemanno, Barra, & Fachechi, 2020; Baldi & Venkatesh, 1987; Krotov & Hopfield, 2016). Although these aspects are not directly related to the current work, it is worth stressing the long-standing, broad interest attracted by the model, further, the current approach can be extended to include such features.

1995; Serricchio et al., 2023), whose aim is to impair the attractiveness of configurations that do not correspond to any of the stored patterns (the convergence to those configurations, sometimes referred to as spurious states, is interpreted as a mistake of the machine). More recently, Hebb's rule has also been revised to make it closer to a learning algorithm: the "reality" that we want to retrieve is now not accessible, instead, a corrupted sample, say  $\{\xi^{\mu,1}, \xi^{\mu,2}, \dots\}$  is available and used to build  $\mathbf{J}$  (see e.g., Agliari, Alemanno, Barra, & De Marzo, 2022; Aquaro, Alemanno, Kanter, Barra, & Agliari, 2023; Fontanari, 1990; Negri, Lauditi, Perugini, Lucibello, & Malatesta, 2023). This way, the available sample of data can be interpreted as a training set and the Hopfield model can be employed for generalization tasks. The bridge between a retrieval scenario and a machine learning setting has also been strengthened by leveraging the equivalence between Hopfield model and Boltzmann machines (see e.g., Agliari & Marullo, 2021; Barra, Bernacchia, Santucci, & Contucci, 2012; Cocco, Monasson, & Sessak, 2011; Fachechi, Barra, Agliari, & Alemanno, 2022; Leonelli, Agliari, Albanese, & Barra, 2021; Mézard, 2017). However, a full mapping allowing for the role of regularization parameters, the emergence of overfitting or underfitting phenomena is still under construction (see e.g., Barbier, Camilli, Mondelli, & Saenz, 2023; Camilli, Tieplova, & Barbier, 2023; Ventura, Cocco, Monasson, & Zamponi, 2023; Zamri, Azhar, Mansor, Alway, & Kasihmuddin, 2022; Zamri et al., 2024).

In this paper, we try to contribute in filling this gap, focusing on an unsupervised reconstruction problem: the dataset is made of a set of items belonging to different classes (the label being veiled) and we introduce a loss function for the neuron interaction matrix. The solution of our problem corresponds to a Hebbian kernel subjected to a certain amount  $t_d$  of unlearning iterations and we prove that  $t_d$  is related to the regularization hyperparameter, which, in turn, can be related to the training time in the un-regularized version of the problem. This framework allows us to inspect the emergence of overfitting phenomena and therefore to conceive recipes for an optimal training time. Specifically, the system stores each pattern as a minimum of the Lyapunov function associated to the neural dynamics (this can be interpreted as a cost function or as an energy function); minima corresponding to the same class form a cluster, and – when the number of examples per class is large enough – these minima do coalesce into a single minimum. In this scenario, there emerge both intra-class and inter-class correlations and we find that the role of  $t_d$  (or, equivalently, of the regularization or of the training time) is to disentangle such correlations starting from the lowest ones: as  $t_d$  is increased, the minima corresponding to different classes are shifted, their overlap is reduced and the system gets able to generalize from examples; by further increasing  $t_d$ , minima corresponding to patterns belonging to the same classes get shifted too and the system starts to overfit.

In what follows, we detail these results by first introducing the loss function associated with our problem and showing that the neural interaction matrix that minimizes the loss function corresponds to the revised Hebbian kernel studied in Agliari, Alemanno, Barra, and Fachechi (2019), Fachechi et al. (2019) (Section 2). Subsequently, we find a relation between regularization, training time and unlearning time, and we present numerical experiments on structureless and structured datasets (Section 3). These results constitute the premises for a thorough discussion on the emergence of overfitting phenomena (Section 4) and a corroboration by numerics (Section 5). Lastly, we conclude our paper by offering a concise outlook and final remarks (Section 6). Technical details are collected in the Appendices.

## 2. From the stability condition to the minimization problem

The basic principle underlying an attractor neural network is that each pattern making up the set  $\xi = \{\xi^\mu\}_{\mu=1}^P$  and encoding relevant information is associated to an attracting fixed point for the network dynamics. We assume that patterns are  $N$ -dimensional binary vectors and that the network units  $\sigma = (\sigma_1, \dots, \sigma_N) \in \{-1, +1\}^N$  interact

pairwise as specified by the symmetric<sup>2</sup> coupling matrix  $\mathbf{J}$ , then, we set up the network evolution as

$$\sigma(t+1) = \text{sgn}[\varphi(\sigma(t))], \quad (2.1)$$

where  $\varphi_i(\sigma(t)) \equiv \sum_{j \neq i} J_{ij} \sigma_j(t)$  represents the signal reaching unit  $i$  at time  $t$  and the sign function acts component-wise. This dynamics is applied sequentially and exhibits the following Lyapunov function (see e.g., Coolen et al., 2005)

$$\mathcal{E}(\sigma) = -\frac{1}{2} \sum_{i,j} \sigma_i J_{ij} \sigma_j. \quad (2.2)$$

In this work, we will retain a deterministic<sup>3</sup> evolution and, for the moment, we take  $\mathbf{J}$  as quenched and we exclude the presence of biases. In this framework, the stability condition for a retrieval configuration, e.g.,  $\sigma(t) = \xi^\mu$  without loss of generality, reads as

$$\xi_i^\mu \varphi_i(\xi^\mu) \geq 0, \quad \forall i = 1, \dots, N, \quad (2.3)$$

in such a way that, if at time  $t$  the system is prepared or occurs to be precisely in that state, it will be there trapped for all  $t' \geq t$ . Of course, the fulfillment of (2.3) implies that  $\mathbf{J}$  must be a suitable functional of  $\xi$ . More generally, one is interested in assessing the convergence to a retrieval configuration even when the input quality is relatively low, namely, even when the initial configuration is relatively far (in Hamming sense) from the target pattern. In this case, one asks for a stronger condition, that is

$$\xi_i^\mu \varphi_i(\xi^\mu) \geq \kappa > 0, \quad \forall i = 1, \dots, N, \quad (2.4)$$

which means that, although  $\sigma(t)$  displays some discrepancies with respect to  $\xi^\mu$ , the dynamics (2.1) is still ensured to converge to  $\xi^\mu$ . In this inequality (corresponding to the basic requirement in Gardner's theory Gardner, 1988; Gardner & Derrida, 1989; Personnaz, Guyon, & Dreyfus, 1985),  $\kappa$  controls the width of the attraction basins, that is, if  $\sigma$  belongs to a Hamming ball  $\mathcal{B}(\xi^\mu, R(\kappa))$  centered in  $\xi^\mu$  with radius  $R(\kappa)$ , the network response will be  $f(\sigma) = \xi^\mu$ , with  $f$  being the transfer function  $f(\sigma) = \lim_{n \rightarrow \infty} \mathcal{T}^n(\sigma)$  and  $\mathcal{T}(\sigma) = \text{sgn}(\varphi_i(\sigma))$  the 1-step dynamics. Increasing  $\kappa$ , the stability criterion will be satisfied in a ball with larger and larger radius  $R(\kappa)$  surrounding the patterns, but this goes at cost of a smaller amount of storable information vectors, resulting in a lower storage capacity. In particular, for symmetric networks ( $\mathbf{J} = \mathbf{J}^T$ ) the largest number of patterns that can be retrieved is  $N$  (Gardner, 1988).

In order to satisfy the inequality constraint (2.4), we can impose a (stronger) equality condition requiring that, given  $\gamma \geq \kappa$ ,

$$\xi_i^\mu \sum_{j \neq i} J_{ij} \xi_j^\mu = \gamma, \quad \forall i = 1, \dots, N, \quad (2.5)$$

with  $\gamma$  being the same for all the patterns. The latter point could appear a rather strong assumption, but – at least in the random theory, where patterns are all equivalent, as their entries are i.i.d. – it is reasonable. The requirement (2.5) has important technical consequences. First, if the patterns are Boolean, it can be rewritten in a more transparent form as  $\sum_{j \neq i} J_{ij} \xi_j^\mu = \gamma \xi_i^\mu$ , which is nothing but the Personnaz et al.'s stability

<sup>2</sup> The symmetry constraint on the coupling matrix is traditionally adopted so that, in the statistical mechanical approach, detailed balance principle holds and this directly implies the stochastic relaxation to an equilibrium distribution in Boltzmann-Gibbs form, see e.g., Coolen, Kühn, and Sollich (2005).

<sup>3</sup> Stochastic realizations of the dynamics (2.1) work in similar ways, apart from the fact that the system is not fixed in the precise configuration given by the  $\mu$ -th pattern, but it is free to explore (at some level given by the degree of stochasticity) the associated attraction basin. This is strictly true if two conditions hold: *i.* the initial configuration is close (in the Hamming sense) to the target pattern and *ii.* in the thermodynamic limit, where ergodicity breaking occurs (in the finite-size case, the transition between different attractors is exponentially small in the system size but non-vanishing).

criterion (Kanter & Sompolinsky, 1987; Personnaz et al., 1985). Further, we can remove the constraint on the absence of self-interactions and allow for  $j = i$  in the last sum,<sup>4</sup> thus recasting the previous expression as an eigenvalue problem as  $\mathbf{J} \cdot \xi^\mu = \gamma \xi^\mu$ ; this means that the coupling matrix is designed so that the patterns are eigenvectors with degenerate eigenvalue  $\gamma$ .

In the general case, we can add an external field by replacing the local internal field in (2.4) with the total field, i.e.,  $\varphi_i(\sigma) \rightarrow \sum_j J_{ij} \sigma_j + h_i$ . The previous arguments still hold and, for  $\mu = 1, \dots, P$ , our problem takes the form

$$\begin{cases} \mathbf{J} \cdot \xi^\mu + \mathbf{h} = \gamma \xi^\mu \\ \mathbf{J} = \mathbf{J}^T \end{cases} \quad (2.6)$$

Therefore, in this context, training a network implies finding an arrangement for  $\mathbf{J}$  and  $\mathbf{h}$ , such that (2.6) holds and this can be recast into the minimization of a Mean-Squared Error (MSE) of the form  $(\mathbf{J} \cdot \xi^\mu + \mathbf{h} - \gamma \xi^\mu)^2$ , in such a way that we can set up the minimization procedure for a loss function reading as<sup>5</sup>

$$\begin{aligned} \mathcal{L}_\xi(\mathbf{J}, \mathbf{h}) = & \frac{1}{2P} \sum_{i,\mu} \left( \sum_j J_{ij} \xi_j^\mu + h_i - \gamma \xi_i^\mu \right)^2 \\ & + \frac{1}{2P} \sum_{j,\mu} \left( \sum_i J_{ij} \xi_i^\mu + h_j - \gamma \xi_j^\mu \right)^2 + \epsilon_J \sum_{i,j} J_{ij}^2 + \epsilon_h \sum_i h_i^2, \end{aligned} \quad (2.7)$$

with  $\epsilon_J, \epsilon_h \in [0, +\infty]$ . The second term in the r.h.s. is obtained starting from the first one and reverting the roles of  $i$  and  $j$ ; these two contributions account for the stability criterion and the symmetry constraint<sup>6</sup> in a non-rigid way. The successive two contributions are  $L_2$ -regularization terms for, respectively,  $\mathbf{J}$  and  $\mathbf{h}$ , that protect the norms of these parameters from divergence during training as they are confined by the quadratic potentials.

The explicit form of the solution of the constrained system (2.6) can be achieved via gradient descent method  $\dot{\mathbf{J}} = -\nabla_{\mathbf{J}} \mathcal{L}_\xi$  and  $\dot{\mathbf{h}} = -\nabla_{\mathbf{h}} \mathcal{L}_\xi$ , which yields

$$-\dot{\mathbf{J}} = \mathbf{J} \boldsymbol{\Omega} + \boldsymbol{\Omega} \mathbf{J} + \mathbf{h} \bar{\xi}^T + \bar{\xi} \mathbf{h}^T - 2\gamma \boldsymbol{\Omega} + 2\epsilon_J \mathbf{J}, \quad (2.8)$$

$$-\frac{1}{2} \dot{\mathbf{h}} = \frac{\mathbf{J} + \mathbf{J}^T}{2} \bar{\xi} + (1 + \epsilon_h) \mathbf{h} - \gamma \bar{\xi}, \quad (2.9)$$

where

$$\Omega_{ij} \equiv \frac{1}{P} \sum_{\mu} \xi_i^\mu \xi_j^\mu, \quad \bar{\xi}_i \equiv \frac{1}{P} \sum_{\mu} \xi_i^\mu,$$

are, respectively, the Hebbian kernel<sup>7</sup> and the mean value of the  $i$ th entry over patterns. Notice that the matrices  $\mathbf{J}$  and  $\mathbf{J}^T$  satisfy the same differential equation (2.8) since  $\boldsymbol{\Omega}$  is symmetric, thus, if we consider initial conditions such that  $\mathbf{J}(0) = \mathbf{J}(0)^T$ , the symmetry is preserved for any  $t > 0$  by uniqueness arguments. Hence, we can safely replace  $\mathbf{J}$  to its symmetric part in (2.9). The convergence condition of the discrete form of these dynamical equations is discussed in Appendix A.

Before concluding this Section, it is worth highlighting that, as standard, the neural relaxation (2.1) and the training dynamics (2.8)–(2.9) operate on different time scales, the former being much faster. Biologically, this is motivated by the fact that synaptic plasticity is much slower than neural activation and, in artificial neural networks,

<sup>4</sup> This has quantitative effects on the estimate of the critical storage capacity, that can be neglected here as the focus is rather on the generalization capabilities in an unsupervised scenario, see e.g. Kanter and Sompolinsky (1987).

<sup>5</sup> The case without self-interactions can be recovered by adding a contribution  $\sum_i \theta_i J_{ii}$  in the loss functions, where  $\theta_i$  are Lagrange multipliers ensuring that the diagonal entries are set to zero.

<sup>6</sup> We recall that the symmetry constraint for the coupling matrix  $\mathbf{J}$  is introduced for consistency with the statistical mechanics picture and to ensure that the sequential dynamics (2.1) exhibits fixed points.

<sup>7</sup> With respect to the standard notation here the prefactor  $1/N$  is replaced with  $1/P$ .

by the fact that the machine is first trained and later used for task accomplishment. For consistency with (2.1) one should therefore write  $\dot{\mathbf{J}} = \frac{d\mathbf{J}}{dt} \tau_J^{-1}$  and  $\dot{\mathbf{h}} = \frac{dh}{dt} \tau_h^{-1}$ , with  $\tau_{J,h} \gg 1$ . However, under this adiabatic hypothesis, we can consider the network parameters  $\mathbf{J}, \mathbf{h}$  as planted during the neural dynamics and separate the two dynamical problems in such a way that, when focusing on the synaptic evolution,  $\tau_{J,h}$  can be set as unitary without ambiguity.

### 3. Dreaming as regularization, regularization as early-stopping

The global minimum for  $\mathcal{L}_\xi(\mathbf{J}, \mathbf{h})$  in (2.7) can be obtained by requiring the stability conditions  $\dot{\mathbf{J}} = \mathbf{0}$  and  $\dot{\mathbf{h}} = \mathbf{0}$  in Eqs. (2.8)–(2.9) that give

$$\mathbf{J} \boldsymbol{\Omega} + \boldsymbol{\Omega} \mathbf{J} + \mathbf{h} \bar{\xi}^T + \bar{\xi} \mathbf{h}^T - 2\gamma \boldsymbol{\Omega} + 2\epsilon_J \mathbf{J} = 0, \quad (3.1)$$

$$\mathbf{J} \bar{\xi} + (1 + \epsilon_h) \mathbf{h} - \gamma \bar{\xi} = 0, \quad (3.2)$$

whose solution reads as

$$\mathbf{h} = \frac{1}{1 + \epsilon_h} (\gamma \mathbf{1} - \mathbf{J}) \bar{\xi}, \quad (3.3)$$

$$\mathbf{J} = \frac{1}{P} \bar{\xi} \frac{\gamma}{\mathbf{I} \epsilon_J + \mathbf{C}} \bar{\xi}^T, \quad (3.4)$$

where

$$\hat{\xi}_i^\mu \equiv \xi_i^\mu - \left( 1 - \sqrt{\frac{\epsilon_h}{1 + \epsilon_h}} \right) \bar{\xi}_i, \quad C_{\mu\nu} \equiv \frac{1}{P} \sum_i \hat{\xi}_i^\mu \hat{\xi}_i^\nu.$$

By inspecting Eq. (3.3) one can see that the external field stems from the presence of biases in the input data, i.e.  $\bar{\xi}_i \neq 0$ , in such a way that re-centering the patterns by  $\xi_i^\mu \rightarrow \xi_i^\mu - \bar{\xi}_i$  results in  $\mathbf{h} = \mathbf{0}$ . Thus, as long as data are pre-processed in this way, external fields are not needed. Further, by looking at Eq. (3.4), one can see that our solution recovers the interaction matrix of the ‘‘Dreaming Hopfield model’’ (DHM) (Agliari et al., 2019; Fachechi et al., 2019, 2022)

$$\mathbf{J}^{(D)} \equiv \frac{1}{P} \xi \frac{t_d}{\mathbf{I} + \mathbf{C} t_d} \xi^T, \quad (3.5)$$

upon setting  $\gamma = 1$  and identifying  $t_d$  as the inverse of the hyperparameter  $\epsilon_J$ ,<sup>8</sup> that is,  $t_d = \epsilon_J^{-1}$ . More precisely, in the DHM, the kernel reads as  $\tilde{\mathbf{J}}^{(D)} \equiv \frac{1}{N} \xi \frac{t_d + 1}{\mathbf{I} + \mathbf{C} t_d} \xi^T$  and it was obtained from the standard Hebbian kernel by iteratively applying an unlearning protocol based on an interplay of remotion and consolidation mechanisms inspired by those occurring during sleep in mammals’ brain Crick and Mitchison (1983), Hopfield et al. (1983); because of this analogy the time  $t_d$ , which measures the number of unlearning iterations, is referred to as ‘‘dreaming time’’. The critical storage of the DHM has been shown to increase monotonically with  $t_d$ , reaching, in the  $t_d \rightarrow \infty$  limit, the theoretical upper bound known for symmetric networks and corresponding to a number of retrievable patterns equal to the number of neurons, i.e.,  $P = N$ . Also, the DHM has been proved to outperform the standard Hopfield model as for generalization abilities (Agliari, Alemanno, Aquaro, Barra, Durante, & Kanter, 2024). The difference between  $\mathbf{J}^{(D)}$  and  $\tilde{\mathbf{J}}^{(D)}$  just lays in a pre-factor  $P/N$  and in a shift

<sup>8</sup> The minimization of the regularized-MSE as defined in (2.7) is a special setting of the usual setup of ridge regression theory (Hoerl, 1962; Hoerl & Kennard, 1970), with the target response of the network being the multiplication of the input vectors  $\xi^\mu$  by the constant  $\gamma$ . Neglecting the bias vector, the ridge estimator minimizing the loss-function does coincide with the coupling matrix  $\mathbf{J}^{(D)}$ , where  $t_d$  plays the role of Tichonov regularization parameter (Tikhonov & Arsenin, 1977). Ridge regression, together with their generalization to non-linear regression problems with kernel techniques (Schölkopf & Smola, 2002; Vapnik, 1998; Vovk, 2013), is a central topic in statistical learning theory, focusing in particular on the role of the corresponding hyper parameter (see e.g. Alberti, De Vito, Lassas, Ratti, & Santacesaria, 2021; Meanti, Carratino, De Vito, & Rosasco, 2022; Wu & Xu, 2020) as well as the implicit regularization phenomena emerging in high-dimensional statistics (Bartlett, Long, Lugosi, & Tsigler, 2020; Hastie, Montanari, Rosset, & Tibshirani, 2022).

$t_d \rightarrow t_d + 1$  which, as long as they are finite and non-vanishing, yield only a quantitative correction.<sup>9</sup>

Moreover, for  $t_d \rightarrow \infty$  (or, equivalently, for  $\epsilon_J \rightarrow 0$ ), we recover Kohonen's projector matrix (Kohonen, 1984)  $\mathbf{J}^{(P)} \equiv \frac{1}{P} \xi \mathbf{C}^{-1} \xi^T$ .

Let us now move forward and notice that the solution  $\mathbf{J}^{(D)}$  obtained by a fully-trained ( $t \rightarrow \infty$ )  $L_2$ -regularized ( $\epsilon_J \neq 0$ ) process can be related to the solution of an unregularized ( $\epsilon_J = 0$ ) process which is run up to a finite time  $t^*$ ; as we will see, this relation allows us to map the dreaming time  $t_d$  into the training time and therefore interpret the dreaming mechanism as a training. In order to establish this relation, we resume the dynamical problem with recentered patterns, in such a way that the inferred field is vanishing<sup>10</sup> and Eq. (2.8) simply reads as

$$-\dot{\mathbf{J}} = \mathbf{J}(\mathbf{\Omega} + \epsilon_J \mathbf{1}) + (\mathbf{\Omega} + \epsilon_J \mathbf{1})\mathbf{J} - 2\gamma\mathbf{\Omega}. \quad (3.6)$$

This can be recast in the basis of the eigenvectors of  $\mathbf{\Omega}$  (denoting with  $a, b = 1, \dots, N$  the corresponding indices) as

$$-\dot{J}_{ab} = (\lambda_a + \lambda_b) J_{ab} + 2\epsilon_J J_{ab} - 2\gamma\lambda_a \delta_{ab}, \quad (3.7)$$

where  $J_{ab}$  is the element of  $\mathbf{J}$  in the current basis and  $\sigma(\mathbf{\Omega}) = \{\lambda_a\}_{a=1}^N$  is the  $\mathbf{\Omega}$  spectrum. By solving Eq. (3.7) we find that the non-diagonal terms asymptotically go to zero as  $J_{ab} \sim \exp[-t(\lambda_a + \lambda_b + 2\epsilon_J)]$  for any initial condition, in such a way that, at the equilibrium point, the coupling matrix is diagonal. Here, we choose to prepare the system in a configuration where no information is stored, i.e., a *tabula rasa* setting  $\mathbf{J}(t=0) = 0$ , in this way the off-diagonal entries remain stuck at zero at any time  $t$ . Remarkably, the diagonal structure of  $\mathbf{J}(t)$ , when expressed in the basis of the eigenvectors of  $\mathbf{\Omega}$ , implies that the two matrices share the same eigenvectors. As for the entries on the principal diagonal, given the above-mentioned initial condition, the solution of the associated differential equation is

$$J_{aa}(t) = \frac{\gamma\lambda_a}{\epsilon_J + \lambda_a} \left\{ 1 - \exp[-2t(\lambda_a + \epsilon_J)] \right\}. \quad (3.8)$$

As expected, in the limit of large training-time, we recover  $\mathbf{J}^{(D)}$ , that is

$$J_{aa}(t) \underset{t \rightarrow \infty}{=} \frac{\gamma\lambda_a}{\lambda_a + \epsilon_J} = J_{aa}^{(D)}, \quad (3.9)$$

while, expanding at small  $t$ , we get

$$J_{aa}(t) = \frac{\gamma\lambda_a}{\epsilon_J + \lambda_a} \left\{ 1 - \exp[-2t(\lambda_a + \epsilon_J)] \right\} \underset{t \ll 1}{\approx} 2t\gamma\lambda_a \quad (3.10)$$

which corresponds to  $\mathbf{J} \approx 2t\gamma\mathbf{\Omega}$ : this means that, despite the blank initial condition, at the very start of the training, the kernel  $\mathbf{J}$  is close to an Hebbian structure.

<sup>9</sup> In particular, the two models exhibit statistical-mechanics equivalence as the partition function of the present model  $Z_\beta(\mathbf{J}^{(D)}) := \sum_{\{\sigma\}} \exp(-\beta \sum_{i,j} J_{ij}^{(D)} \sigma_i \sigma_j)$  can be turned into that of the original DHM by rescaling  $\beta \rightarrow \frac{P}{N} \frac{1+t_d}{t_d} \beta$ , where  $\beta$  tunes the degree of stochasticity in the system, that is, it tunes the broadness of the distribution of the neural configurations or, in a physical jargon, it plays as the inverse temperature.

<sup>10</sup> Equivalently, we can choose to work without rescaling the patterns, thus also including the external fields. In this case, we can simplify the analysis of the dynamical problem by requiring that  $\tau_J \dot{\mathbf{J}} = -\nabla_J \mathcal{L}$  and  $\tau_h \dot{\mathbf{h}} = -\nabla_h \mathcal{L}$ , and consider the case  $\tau_h \ll \tau_J$ . Under this assumption, the variation of the external fields is much faster than the typical evolution of the coupling matrix, so – when dealing with the temporal behavior of the latter – the fields do relax instantaneously towards their fixed point at fixed  $\mathbf{J}(t)$ :

$$\mathbf{h}_\infty[\mathbf{J}(t)] = \frac{1}{1 + \epsilon_h} (\gamma \mathbf{1} - \mathbf{J}(t)) \bar{\xi}.$$

As a consequence, the synaptic dynamics is described by the equation

$$-\dot{\mathbf{J}} = \mathbf{J}(\hat{\mathbf{\Omega}} + \epsilon_J \mathbf{1}) + (\hat{\mathbf{\Omega}} + \epsilon_J \mathbf{1})\mathbf{J} - 2\gamma\hat{\mathbf{\Omega}},$$

with  $\hat{\mathbf{\Omega}}_{ij} = \mathbf{\Omega}_{ij} - M_{ij} = P^{-1} \sum_\mu \xi_i^\mu \xi_j^\mu$  and  $M_{ij} = (1 + \epsilon_h)^{-1} \xi_i \bar{\xi}_j$ . When dealing with structured patterns, we will preserve the inferred field in order not to alter the graphical appearance of the data; this is of course completely irrelevant when dealing with a zero-mean dataset.

On the other hand, by setting  $\epsilon_J = 0$  in Eq. (3.8), we find that the diagonal terms evolve as

$$J_{aa}(t) = \gamma[1 - \exp(-2\lambda_a t)]. \quad (3.11)$$

Now, we compare the two explicit forms of the coupling matrix, i.e. the regularized one at  $t \rightarrow \infty$  (3.9) and the  $t$ -dependent one with  $\epsilon_J = 0$  (3.11), and search for the characteristic time  $t^*$  at which the latter is as close as possible to the former. To do this, let us consider the quantity

$$\delta(\lambda, t, \epsilon_J) = \gamma^2 \left[ \frac{\lambda}{\lambda + \epsilon_J} - 1 + \exp(-2t\lambda) \right]^2,$$

measuring the squared difference between the components in the two realizations at fixed eigenvalue  $\lambda$ . Then, we take the average over the  $\mathbf{\Omega}$  spectrum, i.e.

$$\bar{\delta}(t, \epsilon_J) = \int d\lambda \delta(\lambda, t, \epsilon_J) \rho_E(\lambda),$$

where  $\rho_E(\lambda) = \frac{1}{N} \sum_{\lambda_a \in \sigma(\mathbf{\Omega})} \delta(\lambda - \lambda_a)$  is the empirical spectral distribution of  $\mathbf{\Omega}$ . Notice that  $\bar{\delta}$  is nothing but the squared Frobenius distance between the Dreaming kernel and the unregularized coupling matrix. This quantity is minimized for the following

$$t^*(\epsilon_J) = \underset{t}{\operatorname{argmin}} \bar{\delta}(t, \epsilon_J). \quad (3.12)$$

This relation provides an expression for the time  $t^*$  at which the unregularized gradient descent over  $\mathcal{L}_E(\mathbf{J}, \mathbf{h})$  should be interrupted if we want a coupling matrix close to  $\mathbf{J}^{(D)}$  corresponding to the fully-relaxed, regularized gradient-descent. The equivalence between the two scenarios is validated for synthetic, MNIST (Deng, 2012) and Fashion-MNIST (Xiao, Rasul, & Vollgraf, 2017) datasets as reported in Fig. 1.


The functional relation between  $t_d = \epsilon_J^{-1}$  and  $t^*$  highlighted in Eq. (3.12) is depicted in Fig. 2. The logarithmic behavior is justified analytically in Appendix C, where, by expanding (3.12) around the mean eigenvalue, we obtain a first-order approximation of the early-stopping time  $t^*(\epsilon_J)$  which depends on  $t_d$  and on the trace of  $\mathbf{\Omega}$ .

Another way to see the equivalence between regularization and early-stop is the following. Starting from Eq. (3.9), we notice that  $\epsilon_J$  provides a natural scale as


$$J_{aa}^{(D)} \approx \begin{cases} \gamma \left( 1 - \frac{\epsilon_J}{\lambda_a} \right) & \text{if } \lambda_a \gg \epsilon_J \\ \frac{\gamma\lambda_a}{\epsilon_J} & \text{if } \lambda_a \ll \epsilon_J. \end{cases}$$

Therefore, in the regularization approach, the parameter  $\epsilon_J$  prevents the saturation of all the diagonal entries of  $\mathbf{J}^{(D)}$  to the value  $\gamma$  (corresponding to Kohonen's projector  $\mathbf{J}^{(P)}$  Kanter & Sompolinsky, 1987; Kohonen, 1984; Kohonen & Ruohonen, 1973; Personnaz et al., 1985); in fact, as long as  $\epsilon_J > 0$  (or  $t_d$  finite), only the entries corresponding to the top eigenvalues of  $\mathbf{\Omega}$  get close to this limiting value, while the others remain close to the initial condition (i.e.,  $J_{aa} = 0$  for any  $a = 1, \dots, N$ ). On the other hand, the early-stop dynamically accounts for such a filtering, since the time  $t^*$  at which we stop the training is chosen so that only a subset of the diagonal entries saturate to the fixed point  $J_{aa}^{(D)} = \gamma$  of the dynamical system (3.6), while all the others are not changed in a substantial way w.r.t. the initial condition; in fact, as highlighted by (3.11), the characteristic time for saturation is entry-dependent and given by  $(2\lambda_a)^{-1}$ , thus entries corresponding to large eigenvalues of  $\mathbf{\Omega}$  are faster.

What we presented so far does apply to a general set of Boolean vectors, as the coupling matrix  $\mathbf{J}$  naturally arises as the fixed point of a gradient descent algorithm, the only hypothesis that we made on the vectors  $\xi$ , that we want to store as attractors for the neural dynamics, being that they are of the same length and that they share the same ‘‘importance’’  $\gamma$ . In particular, the dataset items could represent (noisy) realizations of some unknown ground-truth patterns to which we have no direct access. In this context, regularization – preventing



**Fig. 1. Retrieval performance of dreaming kernel versus early-stopping.** The three panels show a comparison between the fully-trained solution (3.4) with  $\epsilon_j \neq 0$  and the solution of the early-stopped training procedure with  $\epsilon_j = 0$ ; for the latter the final training time is chosen according to (3.12). In the leftmost panel, the dataset  $\xi$  is made of  $P$  Rademacher vectors that naturally display zero mean, while in the central and in the rightmost panels the dataset  $\xi$  is made of  $P$  items randomly drawn from, respectively, the MNIST and the Fashion-MNIST datasets, and these vectors were pre-processed by Otsu method (Otsu, 1979) to make them binary. The items in these datasets were used to build up the interaction matrices  $J^{(D)}$  and  $J^{(r^*)}$ . For the random dataset  $N = 200$  and  $\gamma = 1$ , whereas for the MNIST and Fashion-MNIST  $N = 784$  and  $\gamma = 1$ , also, different values of the ratio  $P/N$  are considered as reported in the common legend. The performance of the system is measured in terms of the normalized Hamming distance  $d(\xi^\mu, \sigma^{(\infty)})$  between the target pattern  $\xi^\mu$  and the final configuration  $\sigma^{(\infty)}$ , reached by initializing the system in a corrupted version of  $\xi^\mu$  (obtained by flipping the pattern entries with probability 0.1) and iterating (2.1) up to convergence. By averaging over all the  $P$  patterns we get  $\bar{d}(\xi, \sigma^{(\infty)}) = \frac{1}{P} \sum_{\mu} d(\xi^\mu, \sigma^{(\infty)})$ , which is plotted versus the dreaming time. We refer to Appendix B for further details on numerics.



**Fig. 2. Stopping time as a function of dreaming time.** The plot shows the early-stopping time  $t^*$  as a function of the dreaming time  $t_d$ , obtained by a numerical estimate (solid line) from Eq. (3.12) and by a fit (dashed line) based on the functional relation  $r^*(t_d) = a \log(1 + b t_d)$ , suggested by the analytical findings presented in Appendix C. The network parameters for the three cases  $\gamma = 1$ ,  $N = 784$ , and  $P/N = 0.2$ . The couple of coefficients  $(a, b)$  estimated via linear least-squares are  $(a = 0.66, b = 0.54)$ ,  $(a = 0.11, b = 3.25)$ ,  $(a = 0.19, b = 1.67)$  for the random, MNIST and Fashion-MNIST datasets, respectively.

the network parameters from acquiring large norms during learning – also allows for a reduction of the model specialization on the training set. The relation  $t_d = \epsilon_j^{-1}$  therefore suggests that overfitting issues may arise for too large  $t_d$ , as we are going to discuss in more details in Section 4. Further, the dreaming time  $t_d$  can be related, through  $t_d = \epsilon_j^{-1}$  and Eq. (3.12), to the stopping time  $t^*$  in unregularized versions of the gradient descent algorithm. This relation is consistent with the previous remark since early-stop techniques are indeed designed for avoiding overfitting<sup>11</sup> and, again, recalling the monotonic relation between  $t_d$  and  $t^*$ , we expect that overfitting issues may arise for too large  $t_d$ . In the next section we will make use of the framework outlined in this section and, specifically, of the optimal interaction matrix  $J^{(D)}$ , in order to address the generalization capabilities of such models or, conversely, the emergence of overfitting.

<sup>11</sup> This is executed by following the behavior of training and validation losses, abruptly stopping the training procedure when the latter exhibits a general growing behavior, see e.g. Bös (1998).

## 4. Emergence of generalization and overfitting in Hopfield-like models


### 4.1. A synthetic dataset

The results derived in Sections 2 and 3 were obtained without making specific assumptions on the binary vectors  $\{\xi^\mu\}_{\mu=1}^P$ , however, in order to go further in the analytical investigations, some additional hypothesis are in order. In fact, in theoretical studies one usually assumes that pattern entries are extracted according to a prescribed probability distribution that allows working out a controllable theory. For instance, when dealing with the Hopfield model, a common choice is to take pattern entries as i.i.d. Rademacher random variables, and thus treat the patterns as ground truths to be reconstructed starting from a corrupted version of them. However, in practical applications, one has no direct access to the ground-truth patterns, but only to empirical realizations constituting the dataset from which we want to extract information. In a supervised scenario, one knows *a priori* how the different items making up the dataset are partitioned between classes, so that it is possible to define class archetypes (for instance, the average of examples belonging to the same category) which are taken as representative of the ground vectors, see e.g., Agliari et al. (2022), Aquaro et al. (2023), Fontanari (1990). In the unsupervised case, this clearly cannot be done, and the simplest way to proceed is to include all the examples, homogeneously, in the treatment. This is the path that shall be pursued in this section and hereafter we detail this unsupervised setting by considering a synthetic dataset.

Let  $\{\xi^\mu\}_{\mu=1}^K \in \{-1, +1\}^{N \times K}$  be the ground patterns to which we have access only through empirical realizations referred to as  $\xi^{\mu,A}$  with  $A = 1, \dots, M$  for each  $\mu$ . We assume that the examples in this training dataset are obtained from the ground patterns with a multiplicative noise, that is,  $\xi^{\mu,A} = \chi^{\mu,A} \xi^\mu$  (with entry-wise multiplication), with

$$\text{Prob}[\chi_i^{\mu,A} = \pm 1] = \frac{1 \pm r}{2},$$

where  $r \in [0, 1]$  is the parameter quantifying the *quality* of the example (i.e., it measures the correlation between the example and the corresponding ground-pattern). Our *training* dataset is therefore constituted by the set  $\mathcal{S} = \{\xi^{\mu,A}\}_{\mu=1, \dots, M}^{A=1, \dots, K}$  and we distinguish two loads:  $\alpha := KM/N$ , i.e., the ratio between the number of examples in the training set and the network size, and  $\eta := K/N \leq 1$ , i.e., the ratio between the number of classes in the dataset (which is in principle unknown) and the network size. With this kind of available information, we want to train the system in order to make it able to generalize, namely to reconstruct the hidden ground-truths  $\xi^\mu$ , starting from an input data



**Fig. 3. Schematic representation of training points, spurious combinations and ground-pattern.** The figure sketches the organization of attracting configurations within each class in the dataset. The class is represented by a ground-pattern  $\zeta$  (the red dot in the center), while the training points are located at distance  $(1-r)/2$  from it (i.e., they have correlation  $r$ ). Spurious combinations of training points are themselves attracting points, and their correlation  $c_L(r)$  increases with the number  $L$  of training points involved in the combination. For large enough  $M$ , the resulting landscape consists in many local minima very close to each other, so that they coalesce and form flat valleys around the ground pattern.

$\sigma^{(0)}$  that is a corrupted representation of  $\zeta^\mu$ . Since we have no direct access to the ground-truths, a direct error minimization procedure is not feasible in this case. However, we can include each single item in our loss function and take advantage of emergent phenomena in Hopfield-like models: as we will see, for a sufficiently large dataset, a plethora of spurious states appear and, depending on the control parameters of the system, these can favor the appearance of a generalization phase. In this scenario, regularization mechanism plays a key role, preventing the solution of the system to trivialize or overspecialize on the training set. In this unsupervised setting, the interaction matrix is therefore obtained by plugging in (3.5), that is still solution of the gradient descent Eq. (2.6), the empirical realization of the Hebbian kernel with entries  $\Omega_{ij} = \frac{1}{P} \sum_{\mu,A} \xi_i^{\mu,A} \xi_j^{\mu,A}$ , and the empirical correlation matrix, whose size is  $(KM) \times (KM)$  and whose entries are  $C_{(\mu A),(\nu B)} = \frac{1}{KM} \sum_i \xi_i^{\mu,A} \xi_i^{\nu,B}$ .

To evaluate the performance of the network, we generate a test set  $\tilde{\mathcal{S}} = \{\tilde{\xi}^{\mu,A}\}_{\mu=1,\dots,K}^{A=1,\dots,M}$ , sampled in the same way as the training set, and initialize the network with the configurations of the test set, say  $\sigma^{(0)} = \tilde{\xi}^{\mu,A}$ , the latter being, by construction, a noisy version of  $\zeta^\mu$  with quality  $r$ . Next, we check whether the network response is  $f(\sigma^{(0)}) = \zeta^\mu$ , an outcome that we interpret as a correct generalization; conversely, the retrieval of one of the training items, say  $f(\sigma^{(0)}) = \xi^{\mu,A}$ , is interpreted as overfitting. In the following subsection, we discuss the role of spurious states in the emergence of generalization and overfitting.

#### 4.2. Spurious states of training data enable generalization

In the classical Hopfield setup, spurious states (i.e., configurations that are combinations of stored data) are known to impair the retrieval capabilities of the model and should be suitably treated in order to reduce their attractiveness, see e.g., Agliari et al. (2019), Christos (1996), Crick and Mitchison (1983), Dotsenko et al. (1991), Fachechi et al. (2019, 2022), Hopfield et al. (1983), Plakhov et al. (1995). In fact, the dreaming mechanisms mentioned in Section 3 are precisely aimed at this purpose and their implementation improves the retrieval capabilities of the network. On the other hand, when dealing with a dataset made of unlabeled examples, the situation is quite different, and spurious attractors can be helpful for the emergence of generalization capabilities of the network, as we are going to discuss.

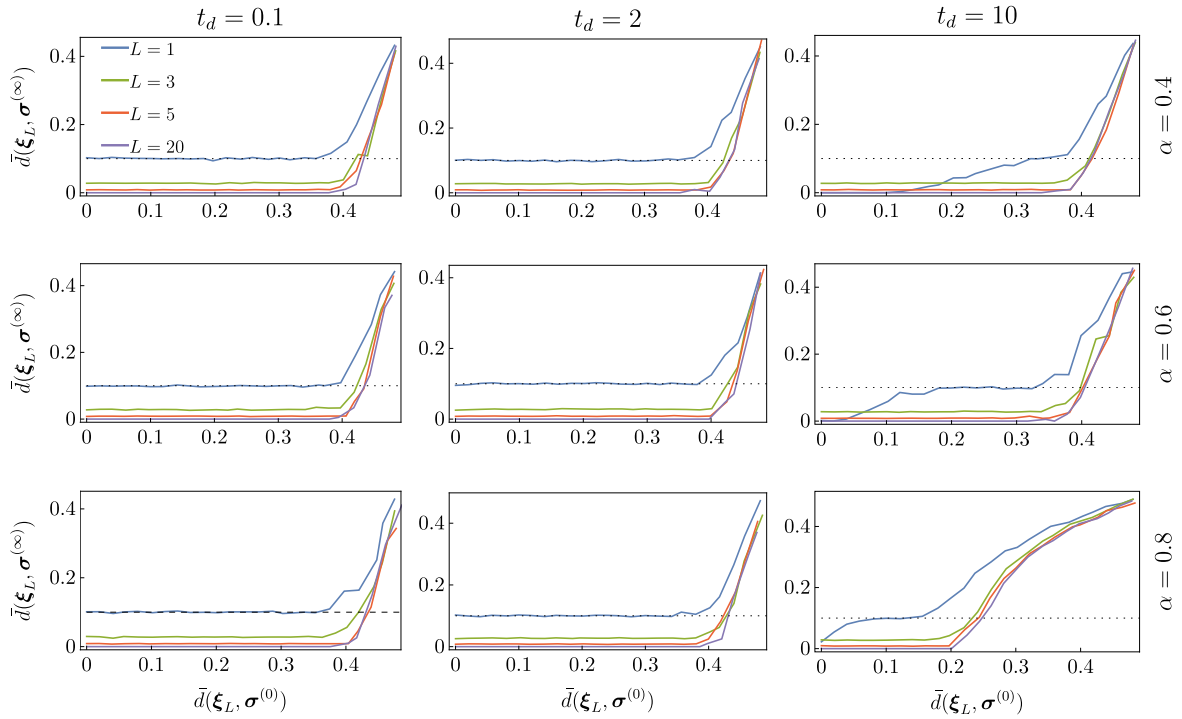
Let us suppose that the training set is made of a large number  $M$  of data for each class, and let us consider a spurious configuration given by a symmetric combination of  $L$  examples pertaining to the same class, that is,

$$\xi_L^\mu = \text{sgn}\left(\sum_{l=1}^L \xi^{\mu,A_l}\right), \quad (4.1)$$

being  $A_1, \dots, A_L \in \{1, \dots, M\}$  the indices of the examples that we are mixing. Denoting with  $c_L(r) := \xi_L^\mu \cdot \zeta^\mu / N$  the correlation between  $\xi_L^\mu$  and the related ground truth, we notice that, as long as  $L$  is relatively large,  $\xi_L^\mu$  displays a correlation with the ground pattern  $\zeta^\mu$  that is larger than the correlation  $r$  displayed by any training item, that is  $r < c_L(r) \xrightarrow{L \gg 1} 1$ , see Appendix D for more details.

Now, spurious configurations of the form (4.1) in Hopfield-like models can be stable attractors, so that running the dynamics (2.1) we could end in one of these minima and reach a relatively fair retrieval. Indeed, each of these configurations originates as a combination of attractors associated to stored vectors. In our scenario, increasing  $M$ , would result in an increasing number of intra-class spurious configurations which, as  $L$  is increased, do exhibit larger and larger correlation with the corresponding ground-truths as sketched in Fig. 3. For sufficiently large  $M$ , it is then reasonable to expect that minima of training examples and spurious configurations coalesce together, so that the resulting landscape consists in a wide minima centered in the ground-truth  $\zeta^\mu$ , favoring the reconstruction of the hidden patterns.

Given such a landscape, which is the role of  $t_d$  (or, equivalently, of  $\epsilon_j$ )? As recalled at the beginning of this subsection, in a Hopfield model where we store ground patterns, dreaming mechanisms reduce (and ultimately remove, if the load is not too high) the stability of spurious mixtures between independent patterns and this is obtained by shrinking and lifting the attraction basins associated to the patterns. Moving to an unsupervised setting, we realize that there are two kinds of mixtures, according to whether they involve examples belonging to different classes or examples belonging to the same class; the former, just like in the Hopfield model, impair retrieval and should be removed, while the latter, as mentioned above, can be beneficial for generalization. Therefore, in this case, the dreaming mechanism should operate in removing only the first type of correlation. In fact, by increasing  $t_d$  we are disentangling the minima corresponding to the stored patterns and this process affects progressively minima with larger and larger overlap. Inter-class correlation is typically smaller than intra-class correlation – their extents being related to, respectively,  $K/N$  and  $r$  – and relatively small values of  $t_d$  can be sufficient to detach the attraction basins related to different ground-patterns. Yet, if we let the dreaming mechanism operate for too long, the mélange of intra-class minima can be separated as well and they get fragmented in many energetic minima, each associated to a single example. As a consequence, we cannot retrieve spontaneously-formed archetypes, but only the single examples: the system is specialized over the training set, thus ending in an overfitting regime. This picture is corroborated by the numerical results reported in Fig. 4, where we focus on the behavior of the system when prepared in the neighborhood of a training example or of an intra-class spurious configuration. First, we notice that, for low enough  $t_d$ , even when the system is initialized in a configuration consisting in a mild perturbation of one of the stored examples ( $L = 1$ ), the neural dynamics will drive it towards a final configuration with the relative distance w.r.t. the reference example being  $(1-r)/2$  (precisely the distance between the stored examples and the corresponding ground-truth). For spurious states, the situation is similar, with the relative distance between the final state and the reference mixture getting lower and lower as  $L$  is increased (as the correlation  $c_L(r)$  with the ground-truth increases monotonically with  $L$ ). A moderately larger  $t_d$  yields a larger attractivity of the ground patterns. This picture is consistent with our claim about the coalescence of the population of attraction basins in a wide minima centered at  $\zeta^\mu$ . By further increasing  $t_d$ , the training points get attractive, with a basin



**Fig. 4. Relaxation to fixed points from perturbed training examples and spurious states.** The plots show the retrieval capabilities of the model initialized in a configuration consisting in a perturbed version of the training examples (i.e.,  $L = 1$ ) or in an intra-class spurious configuration  $\xi_L$  (as given by Eq. (4.1)). The network size is fixed to  $N = 500$ , the number of classes is  $K = 10$ , and the quality of the dataset is  $r = 0.8$ , while different values of dreaming time (from left to right  $t_d = 0.1, 2, 10$ ) and of load (from top to bottom  $\alpha = 0.4, 0.6, 0.8$ , that is,  $M = 20, 30, 40$ ) are considered. The analysis is performed by taking a reference configurations  $\xi_L$  (with for  $L = 1, 3, 5, 20$ , as explained by the legend) and applying a perturbation that consists in randomly flipping a fraction  $q$  of the entries; preparing the system in this configuration  $\sigma^{(0)}$ , we update the network up to convergence towards the fixed point  $\sigma^{(\infty)}$ . Then, we compare the average distances  $\bar{d}(\xi_L, \sigma^{(0)})$  and  $\bar{d}(\xi_L, \sigma^{(\infty)})$  between the reference configurations and, respectively, the initial and the final configurations. The dashed black lines correspond to the distance between the training examples used to build  $\mathbf{J}^{(D)}$  and the associated ground-truths. The results are averaged over 50 different realizations of the dataset.

width depending on the number of examples per class. This signals that, according to the setting, the dreaming mechanism can either enhance generalization or favor overfitting.

## 5. Numerical experiments

In this Section, we provide numerical evidences to our theoretical findings. We first inspect the regions in the space of parameters  $(\alpha, t_d, r)$  where the system equipped with the interaction matrix  $\mathbf{J}^{(D)}$  can successfully generalize, that is, when tested with examples not included in the training set (but sharing with them the same underlying statistics), it is able to fairly reconstruct the ground pattern. Next, we corroborate this picture by applying a clustering algorithm to the network outputs and showing that, in the region of the parameter space where the system is expected to generalize (resp. overfit), the number of classes is nicely estimated (overestimated). Details on numerics are collected in Appendix B.

### 5.1. Generalization diagrams

In the first part of the numerical experiments we consider both structureless and structured datasets, to confirm and check the robustness of our theoretical results. The structureless datasets are built synthetically as follows: we initially generate a set of  $K$  Rademacher ground patterns  $\mathcal{G} = \{\zeta^\mu\}_{\mu=1, \dots, K}$ , whence we obtain a set of training examples  $S = \{\xi^{\mu, A}\}_{\mu=1, \dots, K}^{A=1, \dots, M}$  (characterized by a quality  $r$  as specified in Section 4.1), which are used to build  $\mathbf{J}^{(D)}$ , according to Eq. (3.5). Next, we generate a test set  $\tilde{S} = \{\tilde{\xi}^{\mu, A}\}_{\mu=1, \dots, K}^{A=1, \dots, M}$ , applying the same procedure used for the training set, that is, each item  $\tilde{\xi}^{\mu, A}$  exhibits a quality  $r$  w.r.t. the related ground pattern  $\zeta^\mu$ . For structured datasets, we consider the MNIST (Deng, 2012) and the Fashion-MNIST (Xiao et al.,

2017) benchmarks and define the ground patterns as the class averages, then, the training and the test sets are made of  $M$  items, drawn from the whole datasets (overall made of, respectively, 60000 and 10000 instances), in such a way that the two sets have null intersection.

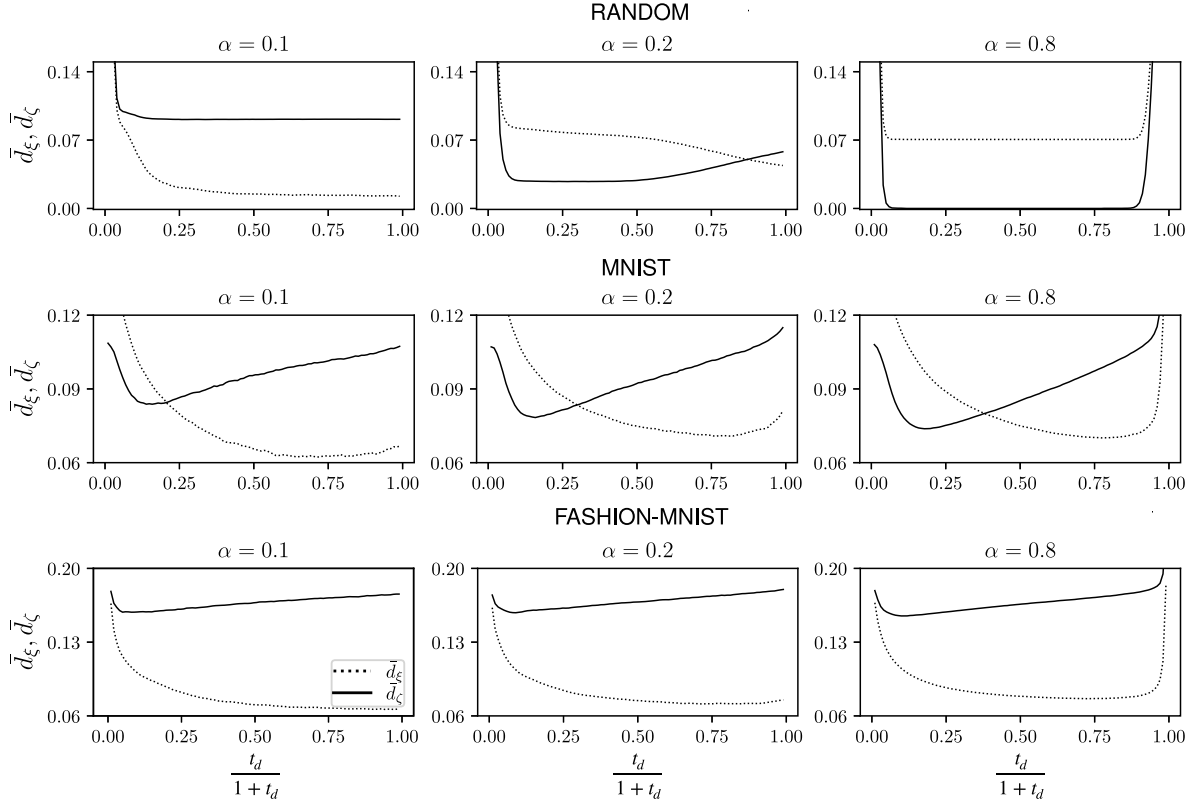
Whatever the dataset, we initialize the system in a configuration  $\sigma^{(0)}$  belonging to the test set, we run the dynamics (2.1) and collect the final configuration  $\sigma^{(\infty)} = f(\sigma^{(0)})$ . In other words,  $\sigma^{(0)}$  and  $\sigma^{(\infty)}$  represent, respectively, the input and the output of the system. Next, we evaluate the following quantities:

$$d_\xi = \min_{\xi \in S} d(\xi, \sigma^{(\infty)}), \quad (5.1)$$

$$d_\zeta = \min_{\zeta \in \mathcal{G}} d(\zeta, \sigma^{(\infty)}), \quad (5.2)$$

where  $d : \{-1, +1\}^N \times \{-1, +1\}^N \rightarrow [0, 1]$  is the normalized Hamming distance, measuring the fraction of misaligned entries among the two configurations that are compared. We stress that, despite not highlighted in Eqs. (5.1)–(5.2), these quantities depend on the initial configuration as  $\sigma^{(\infty)}$  does depend on  $\sigma^{(0)}$ . These distances are then averaged over the sample  $\tilde{S}$ , to get, respectively,  $\bar{d}_\xi$  and  $\bar{d}_\zeta$ .

In Fig. 5, we compare the behavior of  $\bar{d}_\xi$  and  $\bar{d}_\zeta$  versus  $t_d$ , for different choices of  $M$ , while  $K$  is fixed. We find that, in any case,  $\bar{d}_\xi$  and  $\bar{d}_\zeta$  are monotonically decreasing with  $t_d$  as long as  $t_d$  is relatively small, next, their behavior depends on  $M$ . In particular, for the random dataset, when  $M$  is small, we always have  $\bar{d}_\xi < \bar{d}_\zeta$  that evidences poor generalization capabilities; when  $M$  is larger we can leverage  $t_d$  to enhance the generalization capabilities and get  $\bar{d}_\zeta < \bar{d}_\xi$ , yet when  $t_d$  is too large the curves cross; finally, when  $M$  is large, at intermediate dreaming time,  $\bar{d}_\xi$  and  $\bar{d}_\zeta$  exhibit a plateau and at large values of  $t_d$  they grow, the height of the plateau (respectively,  $\approx (1-r)/2$  and  $\approx 0$ ) suggests that there the final configuration is close to the ground pattern, while the final growth suggests a possible harmful



**Fig. 5. Retrieval on synthetic and structured datasets.** The retrieval performance is measured in terms of the normalized Hamming distance  $d$  between the final configuration  $\sigma^{(\infty)}$  and the nearest training example  $\xi$  (dotted curve, see Eq. (5.1)) and the nearest ground-truth  $\zeta$  (solid curve, see Eq. (5.2)); the results presented have been averaged over the  $K \times M$  different initial configurations which constitute the test set (see Appendix B for further details). The network parameters for the random dataset are  $N = 200$ ,  $K = 10$  and  $r = 0.8$ , whereas for the structured datasets they are  $N = 784$ ,  $K = 10$ . For all the datasets, we reported results for different choices of  $\alpha = 0.1, 0.2$  and  $0.8$ , retaining  $\eta = K/N$  fixed, therefore, recalling that  $\alpha = KM/N$ , we varied  $\alpha$  by increasing  $M$ .

effect of a large dreaming time. As for structured datasets,  $\bar{d}_\zeta$  exhibits a minimum at intermediate values of  $t_d$ , corresponding to an optimal generalization performance and, again, for large dreaming times,  $\bar{d}_\zeta$  grows. In particular, for the MNIST dataset, for relatively small (resp. large) values of  $t_d$ , we get  $\bar{d}_\zeta < \bar{d}_\xi$  (resp.  $\bar{d}_\zeta > \bar{d}_\xi$ ), suggesting good (impoverished) generalization capabilities. Before proceeding, we also emphasize that, for all the datasets considered, when  $t_d \gg 1$  and when  $M$  is relatively large, both  $\bar{d}_\zeta$  and  $\bar{d}_\xi$  grow. This is due to the fact that the number of examples is relatively large to give rise to spurious attractors, but not large enough to make these attractors close to the ground truths; this point is further examined in the following.


We now focus on the synthetic dataset and summarize the information processing capabilities of the system into “phase diagrams”. To this aim, we distinguish between different outcomes as follows:

- **Success:** this corresponds to  $\bar{d}_\zeta < \bar{d}_\xi$  and  $\bar{d}_\zeta < \frac{1-r}{2}$ . The first requirement ensures that the system relaxes in a configuration which is more correlated with the ground pattern than with the training points; the second condition, instead, guarantees that the dynamics ends up within the Hamming ball centered in the ground-truth with radius  $(1-r)/2$  and therefore that the system has moved closer to the ground-truth.
- **Overfitting:** this corresponds to  $\bar{d}_\zeta \geq \bar{d}_\xi$  and  $\bar{d}_\zeta < \frac{1-r}{2}$ . The first condition states that the final configuration is closer to one of the training points than to the ground; the second condition, guarantees that, this time, the dynamics ends up within the Hamming ball centered in the nearest training points with radius  $(1-r)/2$  and therefore that the system has moved closer to a specific training item.
- **Failure:** otherwise. In this case the system is neither sufficiently close to a ground-pattern nor to a training item.

For a given choice of  $K$  and  $r$ , moving within the  $(\alpha, t_d)$  plane, we thus depict the generalization diagrams for synthetic datasets. The results are reported in Fig. 6 and discussed hereafter.

Starting from panel *a* ( $K = 10$  and  $r = 0.7$ ), we find that, at very low  $t_d$ , the system is always in a failure regime. This is not surprising since, there, the interaction matrix is very close to the Hebbian prescription and, for this choice of the dataset parameters, interferences among examples are significant enough for the system to be likely to end up in inter-class spurious states, thus we expect that equilibrium configurations are non-retrieval states. Increasing  $t_d$ , such interferences are (possibly) removed. If the load parameter  $\alpha$  is low (meaning that the number of examples is low), the system enters in an overfitting regime because the minima corresponding to examples are sparse enough to be easily disentangled. On the other hand, by increasing  $\alpha$  (namely by increasing the number of examples per class), the attractors coalesce and configurations corresponding to the ground patterns become more and more attractive, in such a way that the system starts to well-generalize. When  $t_d \gg 1$ , as  $\alpha$  is increased, the transition from overfitting to success is no longer direct as, for intermediate values of  $\alpha$ , we can end up into spurious states that are still too sparse to ensure a sound generalization; this region can be shrunk by enhancing the dataset quality. In fact, by increasing the dataset quality as in panel *b* ( $K = 10$  and  $r = 0.8$ ), the qualitative picture is the same, with just an expansion of the success region due to the fact that intra-class examples are now closer to each other. Next, we move to panels *c* and *d*, where a larger  $\eta$  (namely, a larger number of classes  $K = 30$ , with fixed  $N$ ) benefits the failure region, since in this case the clusters of minima associated to each class are closer and clusters of minima corresponding to different classes now present non-trivial overlaps, in such a way that the relaxation of the system could end far away from the class from which the initial condition was generated.





**Fig. 6. Generalization diagrams.** The four diagrams show the generalization outcomes of the neural network where the interaction matrix  $J^{(D)}$  is built on a sample of random, synthetic examples  $\{\xi^{\mu,A}\}_{\mu=1,\dots,M}^{A=1,\dots,K}$  with  $N = 200$  and  $M$  tunable ( $M = \alpha N/K$ ). In the  $(\alpha, t_d)$  plane, for various values of  $K$  and  $r$ , we outline three regions: success (S), overfitting (O), failure (F). In any case the initial conditions  $\sigma^{(0)}$  are taken as perturbed versions of the ground-truths sampled with the same quality  $r$  as the training examples.

## 5.2. An analogy with a clustering algorithm

In this section, we use another approach to check the emergence of overfitting and generalization regime as the system parameters are tuned. The idea is to use an unsupervised clustering algorithm to partition the final configurations  $\sigma^{(\infty)}$  obtained by applying the dynamics (2.1) to the test configurations; here unsupervised refers to the fact that the clustering algorithm is unaware of the number of effective clusters.

We start the experiment by generating a random synthetic dataset  $\mathcal{G}$  made of  $K$  Rademacher ground patterns, hence we build a training set  $S$  and a test set  $\tilde{S}$ , both characterized by a quality  $r$  and a size  $M$ . We use the former to construct  $J^{(D)}$  and the latter to initialize the neural configuration. We collect the final configurations  $\sigma^{(\infty)}$  obtained by iterating the neural dynamics and we expect that, if the network correctly generalizes, the clustering algorithm applied to the final configurations will return an estimated number of clusters  $\hat{K}$  which is (approximately)  $K$  and each cluster contains a number of items which is (approximately)  $M$ . Conversely, if the network overfits, we expect that  $\hat{K} > K$ . In order to quantify the likelihood of these outcomes we introduce the accuracy  $\frac{\hat{M}}{M \times K} \in [0, 1]$ , where  $\hat{M}$  is the total number of examples correctly clustered by the algorithm. The unsupervised clustering algorithm considered here is based on the Disjoint Set Union data structure (Yadav, Shokeen, & Yadav, 2021), which works as follows. Initially, it associates to each item  $\sigma^{(\infty)}$  a different class, in such a way that, at this stage, the number of estimated classes is  $M \times K$ . Next, we consider all the  $\binom{M}{2}$  couples of configurations  $\sigma^{(\infty)}$  and check whether their normalized Hamming distance is smaller than a threshold  $d^*(r)$  and, if so, the two items are merged in the same class. Once all the couples have been examined the algorithm stops. The threshold value


is chosen equal to the minimum of the normalized Hamming distance between all pairs of examples belonging to the test set:

$$d^*(r) = \min_{\substack{(\sigma_1, \sigma_2) \in \tilde{S} \times \tilde{S} \\ \sigma_1 \neq \sigma_2}} d(\sigma_1, \sigma_2).$$

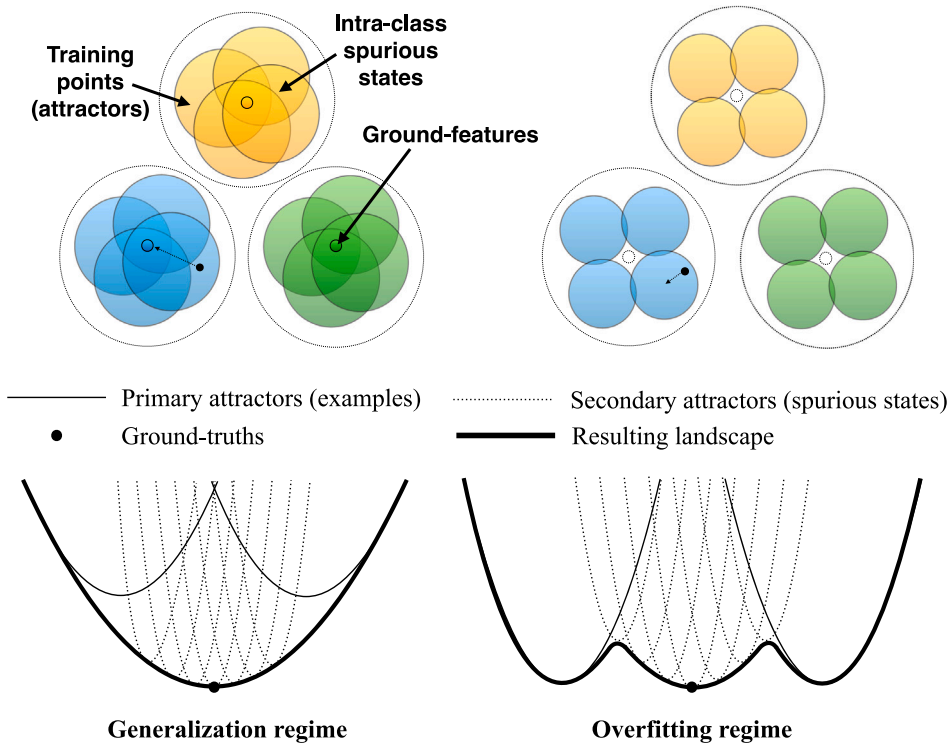
The idea underlying this choice is that, after applying the dynamics to the test examples, if the network performs well, the examples belonging to the same class get closer to the common ground pattern, their distance is reduced and expected to be smaller than  $d^*(r)$ ; on the other hand, for two examples belonging to different classes the distance is expected not to vary significantly and remain larger than  $d^*(r)$ .

The results of this experiment, repeated for different loads and different dreaming times, are reported in Fig. 7. The panels in the first row show the accuracy (left axis) and the difference between  $\hat{K}$  and the true number of classes  $K$  (right axis) as a function of  $\frac{t_d}{1+t_d}$ , while the panels in the second row show the average distances  $\bar{d}_\zeta$  and  $\bar{d}_\xi$  between  $\sigma^{(\infty)}$  and, respectively, the ground truth and the nearest example, as a function of  $\frac{t_d}{1+t_d}$ . The colors in the background correspond to the different regimes of the network and we used the same colormap previously adopted in Fig. 6, to highlight the consistency. In fact, in the failure region the accuracy is low and the number of clusters is underestimated; in the success region the accuracy is unitary and  $\hat{K} = K$ ; in the overfitting region the accuracy is suboptimal and the number of clusters is overestimated. Further, these outcomes are nicely mirrored by the behavior of  $d_\zeta$  and  $d_\xi$ : the transition from failure to success corresponds to an abrupt decrease of  $d_\zeta$  that leaves  $d_\xi$  behind; the transition between success to overfit corresponds to  $d_\zeta$  outpacing  $d_\xi$ .

We stress the clustering procedure described here acts on the dataset items by grouping those with a smaller Hamming distance, in such



**Fig. 7. Clustering of the test set.** Application of the clustering algorithm in case of a random synthetic dataset with parameters  $K = 10, r = 0.8, N = 200$  for two different load values:  $\alpha = 0.15$  and  $\alpha = 0.3$  keeping fixed the number of ground-truths  $K$ . In the simulation with  $\alpha = 0.15$  the overfitting region occurs when  $\bar{d}_\xi$  is overcome by  $\bar{d}_\zeta$ , in this region the clustering algorithm is no longer able to correctly cluster the examples and the number of clusters estimated starts to grow. In the simulation with  $\alpha = 0.3$ , the number of examples per ground truth in the training set is double that of the previous simulation,  $\bar{d}_\zeta$  is always lower than that of the nearest training example  $\bar{d}_\xi$  and the overfitting region is no longer present.



**Fig. 8. Emergence of fixed points for a system trained without supervision.** This schematic picture shows the evolution of the landscape generated by  $J^{(D)}$  as  $t_d$  is varied. When  $t_d$  is relatively low (left), the mélange of minima has been partially disentangled: the weak inter-class interference is removed, while the minima corresponding to examples of the same class are still clustered; when  $t_d$  is relatively large (right) the interference among training examples has been completely shifted. This picture is compatible with our numerical results reported in Fig. 4.

a way that the items belonging to the same cluster are arranged around the related centroid. Therefore, this procedure turns out to be effective when a clustering based on the notion of distance between items is meaningful. In structured datasets (such as MNIST and Fashion-MNIST), this does not necessarily identifies the “natural” clusters (i.e., those corresponding to the ten digits, or to the different fashion categories).

## 6. Conclusions

The main results obtained in this work are listed hereafter:

- We introduced a regularized loss function, whose minimum provides the interaction matrix  $J$  of an associative neural network; according to the dataset provided, the neural network equipped

with the solution  $\mathbf{J}$  is able to retrieve a set of stored ground patterns or to generalize starting from a corrupted version of unknown ground patterns.

- We proved that the solution  $\mathbf{J}$  of the loss function corresponds to a Hebbian-like kernel  $\mathbf{J}^{(D)}$ , known as dreaming Hebbian kernel and parametrized by the “dreaming time”  $t_d$ , as long as  $t_d$  is identified with the inverse of the regularization parameter  $\epsilon_J$ .
- In the absence of regularization ( $\epsilon_J = 0$ ) the solution  $\mathbf{J}^{(D)}$  can be recovered by applying an early-stop strategy to the gradient descent over the loss function. This suggests that  $t_d$  (or, equivalently,  $\epsilon_J = 0$ ) plays a role in preventing overspecialization on the training set.
- Focusing on the case of a training set made of corrupted versions of some unknown ground patterns, we found robust numerical evidence that relatively large values of  $t_d$  and relatively sparse training sets can yield overfitting.
- The emergence of overfitting is related to the structure of the Lyapunov function associated to the neural dynamics and this picture allowed us to speculate on optimal settings for the loss hyperparameters and/or for the training time. This picture is sketched in Fig. 8.

To conclude, our results highlight the relevant mechanism allowing for the emergence of generalization capabilities of Hopfield-like networks: this is identified as the coalescence of attractors associated to training points giving rise to wide minima around the underlying ground truths (which is, indeed, a crucial ingredient for general models exhibiting robust generalization properties, see for example Baldassi, Lauditi, Malatesta, Perugini, & Zecchina, 2021; Baldassi, Pittorino, & Zecchina, 2020 and references therein). In this scenario we give a comprehensive characterization of generalization and overfitting for synthetic random datasets. Developments of the present work would require the extensions to structured data, as well as a statistical mechanics characterization of relevant collective phenomena.

#### CRedit authorship contribution statement

**E. Agliari:** Writing – review & editing, Writing – original draft, Validation, Supervision, Investigation, Formal analysis, Conceptualization. **F. Alemanno:** Formal analysis, Conceptualization. **M. Aquaro:** Writing – original draft, Validation, Investigation, Formal analysis, Data curation. **A. Fachechi:** Writing – review & editing, Writing – original draft, Validation, Investigation, Formal analysis, Data curation, Conceptualization.

#### Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Elena Agliari, Alberto Fachechi reports financial support was provided by Ministry of University and Research. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Data availability

The data used was taken from public repositories.

#### Acknowledgments

E.A. and A.F. acknowledge financial support from PNRR MUR project PE0000013-FAIR and from Sapienza University of Rome, Italy (RM120172B8066CB0, AR2221815D7192C1, AR1221815EA97525).

F.A. and A.F. have been fully supported by PNRR MUR project PE0000013-FAIR.

E.A., M.A., F.A., A.F. acknowledge the stimulating research environment provided by the Alan Turing Institute’s Theory and Methods Challenge Fortnights event “Physics-informed Machine Learning”.

## Appendix A. Convergence of the gradient descent procedure

The discretization of the dynamical equation (3.6) reads as

$$\mathbf{J}(n+1) = \mathbf{J}(n) - \epsilon [\mathbf{J}(n)(\mathbf{\Omega} + \epsilon_J \mathbf{1}) + (\mathbf{\Omega} + \epsilon_J \mathbf{1})\mathbf{J}(n) - 2\gamma \mathbf{\Omega}], \quad (\text{A.1})$$

with initial condition  $\mathbf{J}(n=0) = 0$ . For simplicity, we fix  $\gamma = 1$  without loss of generality, as it only accounts for a global rescaling of the coupling matrix. The dynamical equation Eq. (A.1) can be rewritten as a fixed-point equation:

$$\mathbf{J}(n+1) = G_\epsilon(\mathbf{J}(n)),$$

with

$$G_\epsilon(\mathbf{J}) = \mathbf{J} - \epsilon(\mathbf{J}(\mathbf{\Omega} + \epsilon_J \mathbf{1}) + (\mathbf{\Omega} + \epsilon_J \mathbf{1})\mathbf{J} - 2\gamma \mathbf{\Omega}).$$

Given  $\mathbf{J}$  and  $\mathbf{K}$  two  $N \times N$  matrices, we have

$$G_\epsilon(\mathbf{J}) - G_\epsilon(\mathbf{K}) = \frac{1}{2}(\mathbf{J} - \mathbf{K})(1 - 2\epsilon(\mathbf{\Omega} + \epsilon_J \mathbf{1})) + \frac{1}{2}(1 - 2\epsilon(\mathbf{\Omega} + \epsilon_J \mathbf{1}))(\mathbf{J} - \mathbf{K}).$$

Taking the (operator) norm of both sides, we have

$$\|G_\epsilon(\mathbf{J}) - G_\epsilon(\mathbf{K})\| \leq \|\mathbf{J} - \mathbf{K}\| \cdot \|\mathbf{1} - 2\epsilon(\mathbf{\Omega} + \epsilon_J \mathbf{1})\| = \|\mathbf{J} - \mathbf{K}\|(1 - 2\epsilon(\lambda_1 + \epsilon_J)),$$

where  $\lambda_1$  is the largest eigenvalue of  $\mathbf{\Omega}$ . This means that the function  $G_\epsilon(\mathbf{J})$  is a contraction map, provided that

$$\epsilon \leq \frac{1}{2(\epsilon_J + \lambda_1)}. \quad (\text{A.2})$$

and, by exploiting Banach Fixed Point Theorem, the algorithm converges to the solution of the system. Further, by exploiting Gershgorin’s theorem we can get the following bound

$$\lambda_1 \leq \sum_{ij} |\Omega_{ij}|. \quad (\text{A.3})$$

Thus, for any  $\epsilon$  such that

$$\epsilon \leq \frac{1}{2(\epsilon_J + \sum_{ij} |\Omega_{ij}|)}, \quad (\text{A.4})$$

the convergence requirement (A.2) is trivially satisfied.

## Appendix B. Methods

This appendix is devoted to a detailed description of numerical experiments. We recall that our experiments encompass a training phase, in which we train the coupling matrix of the network with a training set, and a reconstruction phase, in which we perform a sequential dynamics on the neural configuration starting from an item belonging to a test set. Training and test sets are described in Sections 4.1 and 5: synthetic (Rademacher ground patterns) and benchmark (MNIST and Fashion-MNIST) datasets are considered. To facilitate the reproducibility of the research, here we present algorithms and pseudocodes of the numerical experiments (Appendix B.1) along with the performance metrics to evaluate the quality of the final neuron states (Appendix B.2); further evidence on the emergence of overfitting is also provided (Appendix B.3). Moreover, we point out that all the simulations were carried out using the high-performance, dynamic programming language Julia and were run on a personal computer with an Intel Core i7 processor.

### B.1. Training design and parameters assignment

In the training phase we find the expression of the coupling matrix  $\mathbf{J}$  and of the external field  $\mathbf{h}$  which minimize the loss function  $\mathcal{L}_\xi(\mathbf{J}, \mathbf{h})$  given in Eq. (2.7) via gradient descent method. The parameters appearing in the loss function are summarized and described in Table 1. Specifically,  $N$  is the size of the input and is equal to the length of the training patterns,  $P$  is the size of the training set (or total number of training patterns),  $\epsilon_J$  is the regularization constant for the coupling matrix,  $\epsilon_h$  is the regularization constant for the field  $\mathbf{h}$  and, finally,  $t$  is

Table 1

Training parameters.

Parameter	Description
$N$	Size of the input
$P$	Size of the training set
$\epsilon_J$	$L_2$ -regularization term for $\mathbf{J}$
$\epsilon_h$	$L_2$ -regularization term for $\mathbf{h}$
$t$	Training time

**Algorithm 1** Training of the coupling matrix  $\mathbf{J}$ 

**Input:** Training set  $\xi = \{\xi^\mu\}_{\mu=1}^P \in \{-1, +1\}^{N \times P}$ , stopping time  $t$ , regularizer  $\epsilon_J$

Settings for  $t$  and  $\epsilon_J$ :

$L_2$ -regularization:  $t = \infty$  and  $\epsilon_J > 0$

Early-stop regularization:  $t = t^*$  as given by (3.12) and  $\epsilon_J = 0$

**Output:**  $\mathbf{J}(t)$ ,  $\mathbf{h}(t)$

- 1:  $\mathbf{J}(0) = \mathbf{0}_{N,N}$  initialization of the coupling matrix
- 2:  $\xi^\mu = \xi^\mu - \frac{1}{P} \sum_{\mu=1}^P \xi^\mu$  pattern centering
- 3:  $\hat{\Omega} = \frac{1}{P} (\hat{\xi} \cdot \hat{\xi}^T)$  computation of the Hebbian kernel
- 4:  $\Delta_t = \frac{1}{2} [\epsilon_J + \sum_{i,j=1}^N |\hat{\Omega}_{ij}|]^{-1}$  discretization time as given by (A.4)
- 5:  $iters = \lceil \frac{t}{\Delta_t} \rceil$  number of iterations to be run
- 6:  $n=0$
- 7: **repeat**
- 8:  $\mathbf{J} = \mathbf{J} - \Delta_t (2\epsilon_J \mathbf{J} + \mathbf{J}' \hat{\Omega} + \hat{\Omega}' \mathbf{J} - 2\gamma \hat{\Omega})$
- 9:  $n = n + 1$
- 10: **until**  $n = iters$  or  $\mathbf{J}$  has reached a fixed point.
- 11: **return**  $\mathbf{J}, \mathbf{h} = (\gamma \mathbf{1} - \mathbf{J}) \frac{1}{P} \sum_{\mu=1}^P \xi^\mu$

the training time. In the simulations with structured datasets such as MNIST and Fashion-MNIST, since we want to reconstruct the original dataset and not its centered version, we preserve the field  $\mathbf{h}$  in the dynamics and, for the sake of convenience and simplicity we set  $\epsilon_h = 0$  in all the simulations. The pseudo-code presented in Algorithm 1 highlights the steps followed to train the model.

If we run Algorithm 1 up to convergence, the resulting interaction matrix recovers the dreaming kernel, denoted as  $\mathbf{J}^{(D)}$  and reported in Eq. (3.5), as long as we set  $\epsilon_J = t_d^{-1}$  and  $\gamma = 1$ . Thus, if we are interested in the regularized, full-trained model, we can directly pose  $\mathbf{J} = \mathbf{J}^{(D)}$  without the need of running the training procedure.

**B.2. Formulation of the performance metrics**

Once the training is over and we have the desired expression of the coupling matrix  $\mathbf{J}$ , the retrieval capabilities of the machine are investigated. The initial configuration  $\sigma^{(0)}$  is taken as a corrupted version of one of the training patterns (with a fraction of flipped entries w.r.t. the reference configuration) or as an item of a test set (whose elements are statistically analogous to the training patterns, but were not involved in the training procedure); in general, we denote with  $\mathcal{Q}$  the sample of initial configurations. Then, the system relaxes according to the evolution rule introduced in Eq. (2.1) and reported hereafter in a discrete-time notation

$$\sigma_i^{(n+1)} = \text{sgn} \left( \sum_j J_{ij} \sigma_j^{(n)} + h_i \sigma_i^{(n)} \right). \quad (\text{B.1})$$

The pseudo-code for the dynamics can be found in Algorithm 2. After the relaxation towards an equilibrium configuration  $\sigma^{(\infty)}$ , we check its proximity to a specific configuration  $\xi^*$  by exploiting the normalized Hamming distance as a performance metrics. The latter is the number of misaligned entries between two configurations  $\sigma^1, \sigma^2 \in \{-1, 1\}^N$

**Algorithm 2** Sequential dynamic

**Input:** Couplings  $\mathbf{J} \in \mathbb{R}^{N \times N}$ , fields  $\mathbf{h} \in \mathbb{R}^N$ , input  $\sigma^{(0)} \in \{-1, 1\}^N$

**Output:** Final neural configuration  $\sigma^{(\infty)}$

- 1: Remove the diagonal terms from  $\mathbf{J}$
- 2: Set  $n = 0$
- 3: **repeat**
- 4: sample a random integer  $i$  uniformly in the set  $\{1, 2, \dots, N\}$
- 5: update the  $i$ -th spin  $\sigma_i$  according to  $\sigma_i^{(n+1)} = \text{sgn} \left( \sum_{j=1}^N J_{ij} \sigma_j^{(n)} + h_i \sigma_i^{(n)} \right)$
- 6:  $n = n + 1$
- 7: **until**  $\sigma$  has reached a fixed point.

divided by  $N$ , that is

$$d(\sigma^1, \sigma^2) := \frac{1}{2N} \sum_{i=1}^N |\sigma_i^1 - \sigma_i^2|.$$

Then, we evaluate  $d(\sigma^{(\infty)}, \xi^*)$ , which depends on the initial configuration, due to the fact that  $\sigma^{(\infty)}$  depends on  $\sigma^{(0)}$ . Next, we average  $d(\sigma^{(\infty)}, \xi^*)$  over the  $|\mathcal{Q}|$  different realizations of the initial condition; when  $\sigma^{(0)}$  is meant as an item of the test set, this operation corresponds to a batch average. This way, we get the average distance, defined as

$$\bar{d}_{\xi^*} = \frac{1}{|\mathcal{Q}|} \sum_{\sigma^{(0)} \in \mathcal{Q}} d(\sigma^{(\infty)}(\sigma^{(0)}), \xi^*), \quad (\text{B.2})$$

that is expected to depend on the reference point  $\xi^*$  and on the system parameters.

Finally, we notice that  $d(\sigma^{(\infty)}, \xi^*)$  is nothing but the (normalized) absolute error made by the machine outputting  $\sigma^{(\infty)}$ , when asked to reconstruct  $\xi^*$ .

**B.3. Numerical signatures of overfitting**

In this section, we will give more numerical indications of overfitting emergence for the model we are dealing with. Specifically, we check that overfitting can take place in the non-regularized training procedure. To do this, we follow the evolution of the following learning and validation loss functions over training time, in particular the loss functions have the following structure

$$\begin{aligned} \mathcal{L}_{\xi}(\mathbf{J}(t), \mathbf{h}(t)) &= \frac{1}{2P} \sum_{i,\mu} \left( \sum_j J_{ij}(t) \xi_j^\mu + h_i(t) - \xi_i^\mu \right)^2 \\ &+ \frac{1}{2P} \sum_{j,\mu} \left( \sum_i J_{ij}(t) \xi_i^\mu + h_j(t) - \xi_j^\mu \right)^2, \end{aligned} \quad (\text{B.3})$$

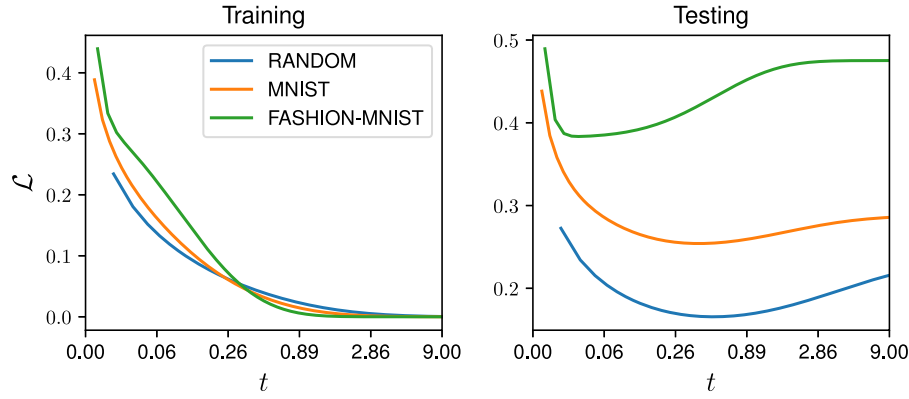
where

$$h_i(t) = \frac{1}{P} \sum_{\mu=1}^P (\xi_i^\mu - \sum_{j=1}^N J_{ij}(t) \xi_j^\mu) \quad (\text{B.4})$$


In the following numerical simulations,  $J_{ij}(t)$  evolves according to Algorithm 1 with parameters  $\epsilon_J$ . Eq. (B.3) is the loss given in Eq. (2.7) of the main text with  $\epsilon_J = \epsilon_h = 0, \gamma = 1$ . The learning and validation loss functions are obtained by substituting into the previous equations  $\{\xi\}_{\mu=1}^P$  with the features of the training and test dataset respectively. For the synthetic random dataset and MNIST and Fashion-MNIST cases we found that, even if the training loss goes to zero, the validation one do exhibit a global minimum at finite training time, and then it starts to increase, thus signaling a worsening in generalization performances. These results are reported in Fig. 9.

**Appendix C. A faster way to compute the early-stopping time**

The estimate of the early-stopping time reported in Eq. (3.12) is based on the computation of the empirical spectral distribution  $\rho_E$  of  $\mathcal{Q}$ , which is a  $N \times N$  matrix. For high-dimension datasets, this can



**Fig. 9.** Training and test losses as a function of the training time. The figure shows the comparison between the training loss function and the validation ones for synthetic dataset (blue curve), MNIST (orange curve) and Fashion-MNIST (green curve) for a training procedure without regularization term.



**Fig. 10.** Comparison between early-stopping and dreaming kernel with approximated time. The three plots show the comparison between the retrieval performances of the Dreaming kernel and the early-stopped training procedure. The content is perfectly specular to Fig. 1, with the only exception that the early-stopping time here is computed with the first-order approximation (C.1).

constitute a bottleneck in the training procedure, so operative criteria needs to be provided. To do this, we can Taylor expand the quantity  $\delta(\lambda, t, \epsilon_J)$  around the average eigenvalue of the empirical distribution  $\bar{\lambda}$ :

$$\delta(\lambda, t, \epsilon_J) = \sum_{k=0}^{\infty} \frac{1}{k!} \delta(\bar{\lambda}, t, \epsilon_J)^{(k)} (\lambda - \bar{\lambda})^k,$$

then we take the average over the  $\Omega$  spectrum, i.e.

$$\begin{aligned} \bar{\delta}(t, \epsilon_J) &= \sum_{k=0}^{\infty} \frac{1}{k!} \delta(\bar{\lambda}, t, \epsilon_J)^{(k)} \int d\lambda (\lambda - \bar{\lambda})^k \rho_E(\lambda) \\ &= \sum_{k=0}^{\infty} \frac{1}{k!} \delta(\bar{\lambda}, t, \epsilon_J)^{(k)} \frac{1}{N} \mathbb{E} \text{Tr}(\Omega - \bar{\lambda} \mathbf{1})^k. \end{aligned}$$

Stopping at the first order in  $\lambda - \bar{\lambda}$  and then solving the minimization problem, we are left with the prescription

$$t^*(\epsilon_J) \approx \frac{1}{2 \frac{1}{N} \mathbb{E} \text{Tr} \Omega} \log \left( 1 + \frac{1}{\epsilon_J N} \mathbb{E} \text{Tr} \Omega \right). \quad (\text{C.1})$$

The results are reported in Fig. 10, which shows again a substantial agreement between the early-stopping procedure and the Dreaming kernel scenario. Notice that this criterion precisely accounts for the correct logarithmic dependence of the early-stopping time w.r.t. the dreaming time  $t_d$ . Clearly, this prescription can provide, in the general case, a rough estimate of the early-stopping time. In that case, one can decide also to work out sub-leading orders in  $\lambda - \bar{\lambda}$ : regardless of the order at which the computation is performed, the numerical estimate of  $t^*$  is based on the computation of low-order moments of the Hebbian

kernel  $\Omega$ , which are far easier to compute than the whole empirical spectral distribution.

#### Appendix D. Details on spurious states

In this Appendix, we report some details about spurious configurations of training examples in the synthetic random dataset. We recall that  $\xi^{\mu, A} = \chi^{\mu, A} \zeta^{\mu}$  (with entry-wise multiplication), with the  $\chi_i^{\mu, A}$  variables extracted as

$$\text{Prob}[\chi_i^{\mu, A} = \pm 1] = \frac{1 \pm r}{2}.$$

Relevant spurious configurations in this setup are of Hopfield-type, so that we can consider combination of the form

$$\xi_L^{\mu} = \text{sgn} \left( \sum_{l=1}^L \xi^{\mu, A_l} \right). \quad (\text{D.1})$$

The correlation of this new configuration with the ground-pattern  $\zeta^{\mu}$  is

$$c_L = \frac{1}{N} \sum_i \xi_{L,i}^{\mu} \zeta_i^{\mu} = \frac{1}{N} \sum_i \text{sgn} \left( \sum_{l=1}^L \xi_i^{\mu, A_l} \zeta_i^{\mu} \right) = \frac{1}{N} \sum_i \text{sgn} \left( \sum_{l=1}^L \chi_i^{\mu, A_l} \right).$$

For large  $L$ , the random variable in the sign function is, by central limit theorem (CLT), Gaussian distributed, with mean  $r$  and variance  $\sqrt{(1-r^2)/L}$ , thus

$$c_L \approx \frac{1}{N} \sum_i \text{sgn} \left( 1 + \sqrt{\frac{1-r^2}{Lr^2}} z_i \right),$$

with  $z \sim \mathcal{N}(0, 1)$ , where we also used  $r > 0$ . The argument of the sign function is positive with probability

$$p := P(z_i \geq -\sqrt{\frac{Lr^2}{1-r^2}}) = 1 - \frac{1}{2} \operatorname{erfc}\left[\sqrt{\frac{Lr^2}{2(1-r^2)}}\right].$$

Thus, the quantity  $c_L$  is nothing but a random walk of unitary steps with probability  $p$  to jump on the right. In the large  $N$  limit, we thus have

$$c_L(r) = \frac{1}{N} \sum_i \operatorname{sgn}\left(1 + \sqrt{\frac{1-r^2}{Lr^2}} z_i\right) \approx 2p - 1 = \operatorname{erf}\left[\sqrt{\frac{Lr^2}{2(1-r^2)}}\right],$$

which yields

$$c_L(r) > r \Rightarrow L > 2 \frac{[\operatorname{erf}^{-1}(r)]^2 (1-r^2)}{r^2}.$$

Now, the r.h.s. of the last inequality ranges in  $[0, \pi/2]$ , while the approximated expression for  $c_L(r)$  is obtained under the CLT, requiring  $L \gg 1$ , which is therefore sufficient for obtaining  $c_L(r) > r$ . For fixed  $L$  there are  $\binom{M}{L}$  possible configurations of the form (4.1). Just to give an example of the magnitudes here involved, for  $M = 50$  and  $r = 0.6$ , we have 1225 spurious combinations of  $L = 2$  examples displaying a correlation with ground-feature  $c_2 \approx 0.8884$ , we have  $\sim 2 \cdot 10^6$  combinations with  $L = 5$  displaying a correlation  $c_5 \approx 0.9993$ , and so on.

## References

Agliari, E., Alemanno, F., Aquaro, M., Barra, A., Durante, F., & Kanter, I. (2024). Hebbian dreaming for small datasets. *Neural Networks*, Article 106174.

Agliari, E., Alemanno, F., Barra, A., Centonze, M., & Fachechi, A. (2020). Neural networks with a redundant representation: Detecting the undetectable. *Physical Review Letters*, 124, Article 028301.

Agliari, E., Alemanno, F., Barra, A., & De Marzo, G. (2022). The emergence of a concept in shallow neural networks. *Neural Networks*, 148, 232–253.

Agliari, E., Alemanno, F., Barra, A., & Fachechi, A. (2019). Dreaming neural networks: Rigorous results. *Journal of Statistical Mechanics*, Article 083503.

Agliari, E., Alemanno, F., Barra, A., & Fachechi, A. (2020). Generalized Guerra's interpolation schemes for dense associative neural networks. *Neural Networks*, 128, 254–267.

Agliari, E., Annibale, A., Barra, A., Coolen, A. C. C., & Tantari, D. (2013). Immune networks: Multitasking capabilities near saturation. *Journal of Physics A*, 46, Article 415003.

Agliari, E., Barra, A., De Antoni, A., & Galluzzi, A. (2013). Parallel retrieval of correlated patterns: From Hopfield networks to Boltzmann machines. *Neural Networks*, 38, 52–63.

Agliari, E., Barra, A., Galluzzi, A., Guerra, F., & Moauro, F. (2012). Multitasking associative networks. *Physical Review Letters*, 109, Article 268101.

Agliari, E., & De Marzo, G. (2020). Tolerance versus synaptic noise in dense associative memories. *European Physical Journal Plus*, 135, 883.

Agliari, E., Fachechi, A., & Marullo, C. (2020). The relativistic Hopfield model with correlated patterns. *Journal of Mathematical Physics*, 61(123301).

Agliari, E., Leonelli, F. E., & Marullo, C. (2022). Storing, learning and retrieving biased patterns. *Applied Mathematics and Computation*, 415, Article 126716.

Agliari, E., & Marullo, C. (2021). Boltzmann machines as generalized Hopfield networks: A review on recent results and outlooks. *Entropy*, 23(1), 34.

Agliari, E., Migliozi, D., & Tantari, D. (2018). Non-convex multi-species Hopfield models. *Journal of Statistical Physics*, 172, 1247–1269.

Alberti, G. S., De Vito, E., Llassas, M., Ratti, L., & Santacesaria, M. (2021). Learning the optimal Tikhonov regularizer for inverse problems. *Advances in Neural Information Processing Systems*, 34, 25205–25216.

Amari, S.-I. (1972). Learning patterns and pattern sequences by self-organizing nets of threshold elements. *IEEE Transactions*, 21, 1197–1206.

Amit, D. J. (1989). *Modeling brain function: the world of attractor neural networks*. Cambridge University Press.

Amit, D. J., Gutfreund, H., & Sompolinsky, H. (1987). Information storage in neural networks with low levels of activity. *Physical Review A*, 35(5), 2293.

Aquaro, M., Alemanno, F., Kanter, I., Barra, A., & Agliari, E. (2023). Supervised Hebbian learning. *Europhysics Letters - Perspective*, 141, 11001.

Baldassi, C., Lauditi, C., Malatesta, E., Perugini, G., & Zecchina, R. (2021). Unveiling the structure of wide flat minima in neural networks. *Physical Review Letters*, 127, Article 278301.

Baldassi, C., Pittorino, F., & Zecchina, R. (2020). Shaping the learning landscape in neural networks around wide flat minima. *Proceedings of the National Academy of Sciences*, 117(1), 161–170.

Baldi, P., & Venkatesh, S. S. (1987). Number of stable points for spin-glasses and neural networks of higher orders. *Physical Review Letters*, 58.

Barbier, J., Camilli, F., Mondelli, M., & Saenz, M. (2023). Fundamental limits in structured principal component analysis and how to reach them. *Proceedings of the National Academy of Sciences*, 120(30), Article e2302028120.

Barra, A., Bernacchia, A., Santucci, E., & Contucci, P. (2012). On the equivalence of Hopfield networks and Boltzmann machines. *Neural Networks*, 34, 1–9.

Bartlett, P. L., Long, P. M., Lugosi, G., & Tsigler, A. (2020). Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48), 30063–30070.

Benedetti, M., Ventura, E., Marinari, E., Ruocco, G., & Zamponi, F. (2022). Supervised perceptron learning vs unsupervised Hebbian unlearning: Approaching optimal memory retrieval in Hopfield-like networks. *Journal of Chemical Physics*, 156, Article 104107.

Bös, S. (1998). Statistical mechanics approach to early stopping and weight decay. *Physical Review E*, 58(1), 833.

Camilli, F., Contucci, P., & Mingione, E. (2022). An inference problem in a mismatched setting: A spin-glass model with Mattis interaction. *SciPost Physics*, 12, 125.

Camilli, F., Tieplova, D., & Barbier, J. (2023). Fundamental limits of overparametrized shallow neural networks for supervised learning. arXiv:2307.05635.

Christos, G. A. (1996). Investigation of the crick-mitchison reverse-learning dream sleep hypothesis in a dynamical setting. *Neural Networks*, 9(3), 427–434.

Cocco, S., Monasson, R., & Sessak, V. (2011). High-dimensional inference with the generalized Hopfield model: Principal component analysis and corrections. *Physical Review E*, 83, Article 051123.

Coolen, A. C. C., Kühn, R., & Sollich, P. (2005). *Theory of neural information processing systems*. Oxford University Press.

Crick, F., & Mitchison, G. (1983). The function of dream sleep. *Nature*, 304(5922), 111–114.

Cugliandolo, L. (1993). Correlated attractors from uncorrelated stimuli. *Neural Computation*, 6, 220.

Del Giudice, P., Franz, S., & Virasoro, M. A. (1989). Perceptron beyond the limit of capacity. *Journal of Physics France*, 50, 121–134.

Deng, L. (2012). The mnist database of handwritten digit images for machine learning research [best of the web]. *IEEE Signal Processing Magazine*, 29(6), 141–142.

Dotsenko, V., Yarunin, N. D., & Dorotheev, E. A. (1991). Statistical mechanics of Hopfield-like neural networks with modified interactions. *Journal of Physics A*, 24, 2419.

Fachechi, A., Agliari, E., & Barra, A. (2019). Dreaming neural networks: Forgetting spurious memories and reinforcing pure ones. *Neural Networks*, 112, 24–40.

Fachechi, A., Barra, A., Agliari, E., & Alemanno, F. (2022). Outperforming RBM feature-extraction capabilities by dreaming mechanism. *IEEE Transactions on Neural Networks and Learning Systems*, 1–10.

Fontanari, J. F. (1990). Generalization in a Hopfield network. *Journal of Physics France*, 51, 2421–2430.

Franz, S., Amit, D. J., & Virasoro, M. A. (1990). Prosopagnosia in high capacity neural networks storing uncorrelated classes. *Journal of Physics France*, 51, 387–408.

Gardner, E. (1988). The space of interactions in neural network models. *Journal of Physics A*, 21(1), 257.

Gardner, E., & Derrida, B. (1989). Three unfinished works on the optimal storage capacity of networks. *Journal of Physics A*, 22(12), 1983.

Hastie, T., Montanari, A., Rosset, S., & Tibshirani, R. J. (2022). Surprises in high-dimensional ridgeless least squares interpolation. *Annals of Statistics*, 50(2), 949.

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. New York, NY: John Wiley & Sons.

Hoerl, A. E. (1962). Application of ridge analysis to regression problems. *Chemical Engineering Progress*, 58, 54–59.

Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1), 55–67.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America*, 79, 2554–2558.

Hopfield, J. J., Feinstein, D. I., & Palmer, R. G. (1983). Unlearning has a stabilizing effect in collective memories. *Nature Letters*, 304, Article 280158.

Kanter, I., & Sompolinsky, H. (1987). Associative recall of memory without errors. *Physical Review A*, 35(1), 380.

Kohonen, T. O. (1984). *Self-organization and associative memory*. Berlin: Springer.

Kohonen, T. O., & Ruohonen, M. (1973). Representation of associated data by matrix operators. *IEEE Transactions on Computers*.

Krotov, D., & Hopfield, J. J. (2016). Dense associative memory for pattern recognition. *Advances in Neural Information Processing Systems*, 1180–1188.

Leonelli, F. E., Agliari, E., Albanese, L., & Barra, A. (2021). On the effective initialization for restricted Boltzmann machines via duality with Hopfield model. *Neural Networks*, 143, 314–326.

Little, W. A. (1974). The existence of persistent states in the brain. *Mathematical Biosciences*, 19(1–2), 101–120.

Marinari, E. (2019). Forgetting memories and their attractiveness. *Neural Computation*, 31(3), 503–516.

Meanti, G., Carratino, L., De Vito, E., & Rosasco, L. (2022). Efficient hyperparameter tuning for large scale kernel ridge regression. In *International conference on artificial intelligence and statistics* (pp. 6554–6572). PMLR.

- Mézard, M. (2017). Mean-field message-passing equations in the Hopfield model and its generalizations. *Physical Review E*, 95.
- Negri, M., Lauditi, C., Perugini, G., Lucibello, C., & Malatesta, E. (2023). Storage and learning phase transitions in the random-features Hopfield model. *Physical Review Letters*, 131(25), Article 257301.
- Otsu, N. (1979). A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1), 62–66.
- Personnaz, L., Guyon, I., & Dreyfus, G. (1985). Information storage and retrieval in spin-glass like neural networks. *Journal of Physics Letters*, 46, 359–365.
- Plakhov, A. Y., Semenov, S. A., & Shuvalova, I. B. (1995). Convergent unlearning algorithm for the Hopfield neural network. *IEE Computation Society Press*, 2(95), 30.
- Schölkopf, B., & Smola, A. J. (2002). *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT Press.
- Serricchio, L., Chilin, C., Bocchi, D., Marino, R., Negri, M., Cammarota, C., et al. (2023). Daydreaming Hopfield networks and their surprising effectiveness on correlated data. In *Associative memory & Hopfield networks in 2023*.
- Tikhonov, A. N., & Arsenin, V. I. A. (1977). *Solutions of ill-posed problems*. Winston: Halsted Press book.
- Vapnik, V. N. (1998). *Statistical learning theory*. A Wiley-Interscience publication. Wiley.
- Ventura, E., Cocco, S., Monasson, R., & Zamponi, F. (2023). Unlearning regularization for Boltzmann machines. (pp. 1065–1095). <https://arxiv.org/pdf/2311.09418.pdf>, 16.
- Vovk, V. (2013). Kernel ridge regression. In *Empirical inference: festschrift in honor of Vladimir N. Vapnik* (pp. 105–116). Springer.
- Wemmenhove, B., & Coolen, A. C. C. (2003). Finite connectivity attractor neural networks. *Journal of Physics A*, 36(9617).
- Wu, D., & Xu, J. (2020). On the optimal weighted  $l_2$  regularization in overparameterized linear regression. *Advances in Neural Information Processing Systems*, 33, 10112–10123.
- Xiao, H., Rasul, K., & Vollgraf, R. (2017). Fashion-mnist: A novel image dataset for benchmarking machine learning algorithms. arXiv preprint arXiv:1708.07747.
- Yadav, A., Shokeen, H., & Yadav, J. (2021). Disjoint set union for trees. In *2021 12th international conference on computing communication and networking technologies ICCCNT*, (pp. 1–6). IEEE.
- Zamri, N. E., Azhar, S. A., Mansor, M. A., Alway, A., & Kasihmuddin, M. S. M. (2022). Weighted random k satisfiability for  $k = 1, 2$  (r2sat) in discrete Hopfield neural network. *Applied Soft Computing*, 126, Article 109312.
- Zamri, N. E., Mansor, M. A., Kasihmuddin, M. S. M., Sidik, S. S., Alway, A., Romli, N. A., et al. (2024). A modified reverse-based analysis logic mining model with weighted random 2 satisfiability logic in discrete Hopfield neural network and multi-objective training of modified niched genetic algorithm. *Expert Systems with Applications*, 240, Article 122307.