

Training neural networks with structured noise improves classification and generalization

Marco Benedetti^{1,*} and Enrico Ventura^{1,2,†}

¹*Dipartimento di Fisica, Sapienza Università di Roma, P.le A. Moro 2, 00185 Roma, Italy*

²*Laboratoire de Physique de l'Ecole Normale Supérieure, ENS, Université PSL, F-75005 Paris, France*

The beneficial role of noise-injection in learning is a consolidated concept in the field of artificial neural networks, suggesting that even biological systems might take advantage of similar mechanisms to optimize their performance. The training-with-noise algorithm proposed by Gardner and collaborators is an emblematic example of a noise-injection procedure in recurrent networks, which can be used to model biological neural systems. We show how adding structure to noisy training data can substantially improve the algorithm performance, allowing the network to approach perfect retrieval of the memories and wide basins of attraction, even in the scenario of maximal injected noise. We also prove that the so-called Hebbian Unlearning rule coincides with the training-with-noise algorithm when noise is maximal and data are stable fixed points of the network dynamics.

I. INTRODUCTION

The development of models and tools from theoretical physics has greatly contributed to our understanding of how the brain is able to learn concepts [1]. An initial approach to brain modeling used recurrent networks composed of binary neurons, connected through pairwise interactions. Each neuron could be either *active* or *silent*, based on the signals received from its neighbors. By incorporating realistic rules for updating neuron states (i.e. the neural dynamics), these networks exhibited dynamic multistability, resulting in the formation of many attractors. Such systems, called *attractor neural networks*, are of considerable interest for multiple reasons: they can be modeled through the tools of statistical mechanics [2]; they constitute a simplified version of modern artificial neural networks, and yet they present a resemblance with biological neuronal systems.

The landscape of attractors is shaped by both the dynamics and choice of couplings between neurons. The process of modification of the interactions, necessary to guarantee that concepts presented to the network are stored as attractors, is studied in neuroscience under the name of *synaptic plasticity*. It is often modeled with discrete time *learning algorithms*, that see a set of *training data* (i.e. the concepts themselves or some representations of them) and update the interactions to reinforce memory retrieval. Modern artificial neural networks, based on deep and non-recurrent architectures, exploit the same principle to train connections between neurons from a data-set and perform a task (e.g. classification, generation of new data).

Attractor neural networks, as associative memory devices, share numerous formal analogies with classifiers from statistical learning. An important notion in this

context is *overfitting*, namely when a model becomes too focused on the details of the specific training set, instead of understanding the broader structure highlighted by the data. In attractor neural networks, an overfitting model would translate into a system which is not able to associate corrupted versions of a concept to the concept itself. In this case, one says that *generalization* to unseen samples is lacking in the system. Generalization can be improved by means of regularization techniques, such as injecting noise in the training process. Data-augmentation procedures [3, 4], in particular, act on the training data themselves, by performing transformations (e.g. translating, rotating, blurring) to increase the heterogeneity of the data-set. Even if such procedures are practical and successful in the matter of deep networks, there are only few examples of theory-based criteria for choosing how to engineer noise to inject in the training process [4, 5]. Our work considers a simple learning algorithm for attractor neural networks that employs noise-injection during training. In the context of attractor neural networks, *noise-injection* refers to a random alteration of the neuronal activity encoding the concepts to be learned. The objective of our study is twofold: establishing a theoretical criterion for generating the most effective noise to be utilized in training, thereby optimizing the performance of an associative memory model; showing that injecting a maximal amount of noise, subject to specific constraints, turns the learning process into an unsupervised procedure, which is faster and more biologically plausible. Specifically, we prove that when noise is maximal, our algorithm converges to Hebbian Unlearning, a well known unsupervised learning rule [6].

Section II presents the neural network model considered in this work, as well as some relevant learning algorithms for training this architecture.

Section III introduces and characterizes the *training-with-noise* algorithm [7], emphasizing its relation with Support Vector Machines and with the optimization problem studied by Wong and Sherrington [8].

* The two authors contributed equally

† Corresponding author: enrico.ventura@uniroma1.it

Section IV contains our main results. We derive an analytical characterization of optimal noisy training data, focusing on the extreme case of injecting maximal noise. We demonstrate that generating optimal noise translates into producing training data-points whose features adhere to specific constraints, which we refer to as the *structure* of the noise. When the level of noise is maximal, minima and low saddle points in the energy landscape defined by Hebb's learning rule are shown to be good training configurations according to our analytical characterization.

In Section V we show that, when noise is maximal (and our theoretical criterion applies), the Hebbian Unlearning rule emerges from the training-with-noise procedure. In this setting, the famously good performance of Hebbian Unlearning acts as a confirmation of our analytical characterization of effective training configurations.

II. THE MODEL

We consider a fully connected network of N binary neurons $\{S_i = \pm 1\}$, $i \in [1, \dots, N]$, linked by couplings J_{ij} . The network is endowed with a dynamics

$$S_i(t+1) = \text{sign} \left(\sum_{j=1}^N J_{ij} S_j(t) \right), \quad i = 1, \dots, N \quad (1)$$

which can be run either in parallel (i.e. *synchronously*) or in series (i.e. *asynchronously* in a random order) over the i indices. We will mainly concentrate on asynchronous dynamics, in which case eq. (1) can only converge to fixed points, when they exist [1]. This kind of network can be used as an associative memory device, i.e. a system that is capable of reconstructing a number p of *memories* $\{\xi_i^\mu\} = \pm 1$, $\mu \in [1, \dots, p]$ when the dynamics is initialized to configurations that resemble them. Such neural networks can be studied through the tools of statistical mechanics [2, 9–11], and they can give relevant insights from the biological point of view [1, 12].

In this work, we will concentrate on i.i.d. memories, generated with a probability $P(\xi_i^\mu = \pm 1) = 1/2$. Notice that memories are binary, as the network states are. With an appropriate choice of the couplings, the model can store an extensive number of memories $p = \alpha N$, where α is called *load* of the network. The network performance can be benchmarked in terms of the dynamic stability achieved by the memory vectors, and the ability of the system to retrieve blurry examples of the latter. We define *classification* as the capability to perfectly retrieve each memory when the dynamics is initialized to the memory itself. It is convenient to define, for all site index i and memory index μ , the *stability*

$$\Delta_i^\mu = \frac{\xi_i^\mu}{\sqrt{N} \sigma_i} \sum_{j=1}^N J_{ij} \xi_j^\mu, \quad \sigma_i = \sqrt{\sum_{j=1}^N J_{ij}^2 / N}. \quad (2)$$

When for a given μ all Δ_i^μ , $i \in [1, \dots, N]$ are positive, ξ^μ is *correctly classified*, i.e. it is a fixed point of eq. (1). We define *generalization* as the capability to retrieve the memory, or a configuration that is strongly related to it, by initializing the dynamics on a corrupted version of the memory. This property of the neural network is related to the size of the basins of attraction to which the memories belong, and does not imply classification. A good measure of the performance in this sense is the *retrieval map*

$$m_f(m_0) := \overline{\left\langle \frac{1}{N} \sum_{i=1}^N \xi_i^\mu S_i^\mu(\infty) \right\rangle}, \quad (3)$$

i.e. the overlap between memory ξ^μ and $\bar{S}^\mu(\infty)$, where $\bar{S}^\mu(\infty)$ is a fixed point, reached initializing the dynamics in a state with overlap m_0 with ξ^μ . For the rest of this work, the symbol $\overline{\cdot}$ will denote the average over different realizations of the memories and $\langle \cdot \rangle$ the average over different realizations of the initial state. In the classification regime, one obtains $m_f = 1$ when $m_0 = 1$.

We now provide for some notable examples of learning rules implemented on recurrent neural networks, which will often be mentioned in the rest of this work. All these procedures can be conceived as an iterative modification of the couplings on the basis of an initial choice of J and a set of training data presented to the network.

A. Linear Perceptron & Support Vector Machine (SVM)

Given the fully connected architecture, we want to find a set of couplings that satisfies the constraints

$$\Delta_i^\mu > k, \quad \forall \mu, i, \quad (4)$$

with k being a positive parameter. This problem can be solved by the linear perceptron algorithm developed by [13–15], an iterative procedure leading to dynamically stable memories and tunable generalization performance. Specifically, we are going to refer to *linear perceptron* as the adaptation of the classical perceptron to recurrent neural networks, introduced in [16–18]. Given a value of the load α , condition (4) is satisfiable up to a maximum value of $k_{max}(\alpha)$. The maximum capacity achievable by the network is $\alpha_c^P = 2$, with $k_{max}(2) = 0$. Starting from any initialization of the coupling matrix J such that $J_{ii} = 0 \forall i$, all constraints in (4) will be satisfied after a number of iterations of the following serial update

$$\delta J_{ij}^{(d)} = \lambda \sum_{\mu=1}^p \Theta(k - \Delta_i^\mu) \xi_i^\mu \xi_j^\mu, \quad \delta J_{ii}^{(d)} = 0, \quad (5)$$

where λ is a small positive learning rate, d is the algorithm time step and $\Theta(x)$ is the Heaviside function. Perceptron learning is *supervised*, in the sense that eq. (5)

requires the explicit knowledge of the memories, at each step of training. One can also symmetrize equation (5) as

$$\delta J_{ij}^{(d)} = \frac{\lambda}{2} \sum_{\mu=1}^p \left[\Theta(k - \Delta_i^\mu) + \Theta(k - \Delta_j^\mu) \right] \xi_i^\mu \xi_j^\mu, \quad (6)$$

$$\delta J_{ii}^{(d)} = 0,$$

obtaining a lower critical capacity and a general decreasing of the function $k_{max}(\alpha)$ due to the symmetry constraint [14, 19]. It has been shown numerically that, given α , the larger is k , with $0 < k < k_{max}(\alpha)$, the wider are the basins of attraction. In line with the past literature [16, 20] we call a maximally stable perceptron, such that $k = k_{max}(\alpha)$, a Support Vector Machine (SVM).

B. Hebb's rule

Hebb's (or Hebbian) learning prescription [12, 21] consists in building up the connections between neurons as an empirical covariance of the memories, i.e.

$$J_{ij}^H = \frac{1}{N} \sum_{\mu=1}^p \xi_i^\mu \xi_j^\mu, \quad J_{ii}^H = 0. \quad (7)$$

This rule allows retrieving memories up to a critical capacity $\alpha_c^H = 0.138$ [9]. Notably, when $\alpha < \alpha_c^H$ memories are not perfectly recalled, but only reproduced with a small number of errors. On the other hand, when $\alpha > \alpha_c^H$ interference between memories makes the learning rule completely ineffective, shifting the system into an oblivion regime. The landscape of attractors given by Hebb's rule is rugged and disseminated with *spurious states*, i.e. stable fixed points of the dynamics that are far from the memories [10].

C. Hebbian Unlearning

Inspired by the brain functioning during REM sleep [22], the Hebbian unlearning algorithm (HU) [6, 19, 23–26] is a training procedure leading to classification and good generalization in a symmetric neural network. Training starts by initializing the connectivity matrix according to Hebb's rule eq. (7) (i.e. $J^{(0)} = J^H$). Then, the following procedure is iterated at each step d :

1. Initialization of the network on a random neural state.
2. Run the asynchronous dynamics (1) until convergence to a stable fixed point $\vec{S}^{(d)}$.
3. Update couplings according to:

$$\delta J_{ij}^{(d)} = -\frac{\lambda}{N} S_i^{(d)} S_j^{(d)}, \quad \delta J_{ii}^{(d)} = 0, \quad (8)$$

where λ is a small positive learning rate. HU is a *unsupervised* learning algorithm, in the sense that each step of training does not require explicit knowledge of the memories, but rather relies on implicit information encoded in the dynamics, i.e. in the current coupling matrix $J^{(d)}$. This algorithm was first introduced to prune the Hebbian energy landscape from proliferating spurious attractors [1, 10]. Classification is achieved by running the algorithm when $\alpha \leq \alpha_c^U$ with $\alpha_c^U \simeq 0.6$. In this regime, HU creates large basins of attraction around the memories, comparable in size to those of a SVM [19].

III. TRAINING WITH NOISE

The concept of learning from noisy examples, introduced in [27], is at the basis of a pioneering study by Gardner and collaborators [7], in an attempt to improve generalization in attractor neural networks. Here, we report the algorithm and characterize, for the first time, its performance over fully connected neural networks. Importantly, we find strong numerical evidence of a connection between the training-with-noise algorithm and the optimization problem studied in [8, 11]. The loss function for this problem, minimized by the training-with-noise algorithm, will be at the core of the analysis contained in our work.

A. Gardner's algorithm

The training-with-noise (TWN) algorithm [7] consists in starting from any initial coupling matrix $J^{(0)}$ with null entries on the diagonal, and updating recursively the couplings according to

$$\delta J_{ij}^{(d)} = \frac{\lambda}{N} \Theta \left(-\xi_i^{\mu_d} \sum_{k=1}^N J_{ik} S_k^{\mu_d} \right) \xi_i^{\mu_d} S_j^{\mu_d}, \quad \delta J_{ii}^{(d)} = 0, \quad (9)$$

where λ is a small learning rate and $\mu_d \in [1, \dots, p]$ is a randomly chosen memory index. In this setting, \tilde{S}^{μ_d} is a noisy data-point, generated as

$$S_i^{\mu_d} = \chi_i^{\mu_d} \xi_i^{\mu_d}, \quad (10)$$

with $\chi_i^{\mu_d}$ i.i.d. variables, sampled with probability

$$P(\chi_i^{\mu_d} = x) = \frac{(1 + m_t)}{2} \delta(x-1) + \frac{(1 - m_t)}{2} \delta(x+1). \quad (11)$$

Notice that \tilde{S}^{μ_d} are binary variables, as the $\xi_i^{\mu_d}$ are. The *training overlap* m_t is a control parameter for the level of *noise* injected during training, corresponding to the expected overlap between \tilde{S}^{μ_d} and ξ^{μ_d} , i.e.

$$m_t = \frac{1}{N} \sum_{i=1}^N \xi_i^{\mu_d} S_i^{\mu_d} + O\left(\frac{1}{\sqrt{N}}\right). \quad (12)$$

The algorithm would converge when every configuration with overlap m_t with a memory generates on each site a local field aligned with the memory itself. TWN is a *supervised* learning algorithm, in the sense that eq. (9) requires the explicit knowledge of the memories, at each step of training.

B. A loss function for TWN

The goal of this subsection is to connect the purely empirical TWN introduced in [7] with the optimization problem studied in [8, 11]. We show that TWN globally minimizes a loss function. This will allow us to get valuable analytical insights on its performance. Let us define the function

$$\mathcal{L}(m, J) = -\frac{1}{\alpha N^2} \sum_{i, \mu}^{N, p} \operatorname{erf} \left(\frac{m \Delta_i^\mu}{\sqrt{2(1-m^2)}} \right). \quad (13)$$

Wong and Sherrington [8, 11] propose an elegant analysis of a network designed to optimize $\mathcal{L}(m, J)$ given m , i.e. whose couplings correspond to the global minimum of $\mathcal{L}(m, J)$. The physical meaning of $\mathcal{L}(m, J)$ comes from the identity

$$\mathcal{L}(m, J) = -\frac{1}{N} \sum_{i=1}^N \left\langle \xi_i^\nu \operatorname{sign} \left(\sum_{j=1}^N J_{ij} S_j^\nu(0) \right) \right\rangle, \quad (14)$$

where $S^\nu(0)$ is a configuration with overlap m with memory ξ^ν . The RHS of eq. (14) is the opposite of the *one-step retrieval map*, defined by selecting one memory, sampling a noisy state having overlap m with the memory, and measuring the overlap obtained after a single step of *synchronous* dynamics (1). Hence, minimizing $\mathcal{L}(m, J)$ means driving the first step of the retrieval dynamics towards the memories, enlarging their attraction basins. The coupling matrix $J(m)$ resulting from the optimization of eq. (13) depends on m . Some of their findings, relevant to this work, are:

1. For any m_0 , the maximum value of the one-step retrieval map $-\mathcal{L}(m_0, J(m))$ is obtained at $m = m_0$.
2. When $m \rightarrow 1^-$, the minimization of the function $\mathcal{L}(m, J(m))$ trains a linear perceptron with maximal stability, i.e. a SVM.
3. When $m \rightarrow 0^+$, the minimization of $\mathcal{L}(m, J(m))$ leads to a Hebbian connectivity matrix $J_{ij} \propto J_{ij}^H$.

We now proceed to establish a connection between the TWN procedure and the theoretical results obtained by Wong and Sherrington. When the network is trained through TWN, the resulting coupling matrix depends on m_t , i.e. $J = J(m_t)$. It is crucial to stress the difference between the variables m and m_t : the former is a parameter of eq. (13), the latter is the level of noise used by the

training algorithm. It comes out that eq. (13) is relevant to the TWN procedure, since iterating eq. (9) leads to a decrease of $\mathcal{L}(m, J(m_t))$, for any value of m and m_t . In fact, considering a small variation of the stabilities induced by the algorithm update

$$\Delta_i^\mu \rightarrow \Delta_i^\mu + \delta \Delta_i^\mu,$$

and performing a Taylor expansion of (13) at first order in $O(N^{-1/2})$, one obtains (see Appendix A 1)

$$\mathcal{L}' = \mathcal{L} + \sum_{i=1}^N \delta \mathcal{L}_i, \quad (15)$$

where

$$\delta \mathcal{L}_i = -\Theta \left(-\xi_i^{\mu_d} \sum_{k=1}^N J_{ik} S_k^{\mu_d} \right) \frac{\lambda}{\alpha \sigma_i N^{5/2}} \cdot \frac{\sqrt{2} m \cdot m_t}{\sqrt{\pi(1-m^2)}} \exp \left(-\frac{m^2 \Delta_i^{\mu_d^2}}{2(1-m^2)} \right). \quad (16)$$

Hence, $\delta \mathcal{L}_i$ is strictly non-positive when $\frac{\lambda}{N}$ is small, so that the Taylor expansion is justified. Moreover, we




FIG. 1: Lines in the main plot report the function $\mathcal{L}(m = 0.5, J(m_t))$ for different training overlaps as functions of the number of algorithm steps d . The *dashed* line represents the theoretical minimum value from [8]. All measures are averaged over 5 realizations of the couplings J . Choice of the parameters: $N = 100$, $\alpha = 0.3$, $\lambda = 1$, the initial couplings are Gaussian with unitary mean, zero variance and $J_{ii}^{(0)} = 0 \forall i$.

numerically find that iterating (9) with a given value of m_t drives $\mathcal{L}(m = m_t, J(m_t))$ to its theoretical absolute minimum computed in [8], as reported in fig. 7. As a




FIG. 2: $m_f(1)$ as a function of m_t and α . Warmer shades of color are associated to higher retrieval performances. The *black dashed* line represents the boundary of the retrieval regime according to the criterion in Appendix C, *white* dots signal the points where basins of attraction to which memories belong are larger than ones obtained from a SVM at $N = 200$.

consequence, the performance of the TWN algorithm can be completely described in the analytical framework of [8] and $\mathcal{L}(m = m_t, J)$ can be considered as the loss function optimized by the TWN algorithm.

C. TWN performance

The same analytical techniques applied in [8] can be used to determine whether the TWN algorithm is capable to reach perfect classification. We find that the distribution of stabilities in the trained network (see equation (B6)) has always a tail in the negative values, implying that classification is never reached. The only exception to this statement are the trivial cases of $m_t = 1^-$ (SVM limit) and $m_t = 1$ (linear perceptron with $k = 0$), where all memories are stable for $\alpha \leq 2$.

The generalization properties of a network trained through TWN are now discussed. The color map in fig. 2 reports the numerical estimate of the retrieval map m_f at $m_0 = 1$, extrapolated in the $N \rightarrow \infty$ limit, as a measure of the distance between a given memory and the closest attractor. In Appendix C we propose an empirical criterion to separate a *retrieval* phase from a *non-retrieval* one. In the first region, the memory appears to close to an attractor when $N \rightarrow \infty$, while in the second region the closest attractor is typically orthogonal in the thermodynamic limit. Such criterion is based on assuming that $m_f(m_0)$, measured with respect to the attractor of the basin to which a given memory belongs and not the memory itself, always develops a

plateau starting in $m_0 = 1$ and ending in some $m_c < 1$ when $N \rightarrow \infty$. The behavior of the basin radius is then observed numerically as a function of m_t : when the plateau disappears (i.e. $m_c = 1$) then one can suppose that basins get shattered in the configurations space due to the interference with the other attractors, and this occurs at some value of m_t . The transition between the two regions is reported as a dashed line in fig. 2.

In the retrieval region, we compute the typical size of the basins of attraction, as described in Appendix C. White dots in fig. 2 signal the combinations of (m_t, α) where the basins obtained by TWN algorithm resulted larger than the ones shaped by a SVM at the same value of α . One can conclude that for most of the retrieval region the generalization performance is worse than the SVM, which maintains larger basins of attraction; on the other hand, at higher values of α the trained-with-noise network sacrifices its classification property to achieve a basin that appears wider than the SVM one. We want to stress the importance of a comparison between the TWN and the corresponding SVM, since numerical investigations have shown the latter to achieve extremely large basins of attraction, presumably due to the maximization of the stabilities [19, 28].

In conclusion, the TWN algorithm never outperforms the corresponding SVM in terms of attraction basins size without reducing its classification capabilities. This observation moves our research for a modification to the TWN algorithm, which can ensure a performance that is similar to the SVM one even at lower values of m_t and higher loads α .

IV. THE OPTIMAL STRUCTURE OF NOISE

As previously stated, SVMs are considered to be highly efficient associative memory models, due to their very good classification and generalization capabilities. Wong and Sherrington's analytical argument proves that minimization of $\mathcal{L}(m = 1^-, J)$ trains a SVM. The TWN procedure proposed by Gardner and collaborators, relying on a Bernoulli process to generate noise, accomplishes this task only when $m_t = 1^-$ (see Section III). Injecting a larger amount of noise during training would deteriorate the performance: specifically, $m_t = 0^+$ trains a Hebbian neural network. Nevertheless, training a network with examples that are nearly uncorrelated with the memories can significantly speed up the sampling process, since such states can be generated in an unsupervised fashion, without knowledge of the memories (as seen for the HU algorithm in Section II C). Such unsupervised processes are also considered more biologically plausible.

In this Section, we show that it is possible to use maximally noisy configurations (i.e. $m_t = 0^+$) to train a network that resembles a SVM, by means of the TWN

algorithm. This translates into approaching the global minimum of $\mathcal{L}(m = 1^-, J(0^+))$. For this purpose, one must change the way noisy data are generated: they need to meet specific constraints which lead to internal dependencies among the features (i.e. a *structure*). We derive a theoretical condition characterizing the optimal structure of noise, and show that specific configurations in the Hebbian energy landscape, including local minima, match well the theoretical requirements.

A. The optimal noise condition

It will be helpful for our purposes to implement a symmetric version of rule (9), i.e.

$$\begin{aligned} \delta J_{ij}^{(d)} = \frac{\lambda}{N} & \left(\Theta \left(-\xi_i^{\mu_d} \sum_{k=1}^N J_{ik} S_k^{\mu_d} \right) \xi_i^{\mu_d} S_j^{\mu_d} + \right. \\ & \left. + \Theta \left(-\xi_j^{\mu_d} \sum_{k=1}^N J_{jk} S_k^{\mu_d} \right) \xi_j^{\mu_d} S_i^{\mu_d} \right). \end{aligned} \quad (17)$$

Henceforth in the paper, TWN will be understood as the symmetric update rule eq. (17). Since SVMs show a high degree of symmetry in the J matrix [16, 18], we extend the results presented in Section III to the current case. This assumption is well-supported by further numerical analysis. Equation (17) can be reformulated explicitly by rewriting the Heaviside function as $\Theta(x) = (1 + \text{sign}(x))/2$, leading to

$$\begin{aligned} \delta J_{ij}^{(d)} = \frac{\lambda}{2N} & (\xi_i^{\mu_d} S_j^{\mu_d} + S_i^{\mu_d} \xi_j^{\mu_d}) \\ & - \frac{\lambda}{2N} (S_i^{1,\mu_d} S_j^{\mu_d} + S_i^{\mu_d} S_j^{1,\mu_d}), \end{aligned} \quad (18)$$

where $S_i^{1,\mu_d} := \text{sign}(\sum_{k=1}^N J_{ik} S_k^{\mu_d})$ corresponds to one step of the dynamics. In the maximal noise case $m_t = 0^+$, the variation of $\mathcal{L}(m, J)$ in one step of the algorithm is then (see Appendix A 2)

$$\delta \mathcal{L} \propto \frac{m}{\sqrt{2\pi(1-m^2)}} \sum_{i,\mu} \omega_i^\mu \exp \left(-\frac{m^2 \Delta_i^{\mu^2}}{2(1-m^2)} \right), \quad (19)$$

where the weights ω_i^μ are given by

$$\omega_i^\mu = \frac{1}{2\sigma_i} \left(m_\mu \chi_i^{1,\mu} + m_{1,\mu} \chi_i^\mu \right), \quad (20)$$

with

$$\chi_i^\mu = \xi_i^\mu S_i^{\mu_d}, \quad \chi_i^{1,\mu} = \xi_i^\mu S_i^{1,\mu_d}, \quad (21)$$

and

$$m_\mu = \frac{1}{N} \sum_{j=1}^N S_j^{\mu_d} \xi_j^\mu, \quad m_{1,\mu} = \frac{1}{N} \sum_{j=1}^N S_j^{1,\mu_d} \xi_j^\mu. \quad (22)$$

Notice that definitions eq. (21) and eq. (22) mirror eq. (10) and eq. (12), but describe the relation between training data points and a *generic* memory $\tilde{\xi}^\mu$, not just the specific $\tilde{\xi}^{\mu_d}$ picked for the training step. Since our goal is to approach the performance of a SVM, we are interested in the limit $m \rightarrow 1^-$. When $m \rightarrow 1^-$, the Gaussian factors in eq. (19) are tightly peaked around zero. Since we want $\delta \mathcal{L}$ to be negative, we need, for most of the pairs i, μ ,

$$\omega_i^\mu < 0 \text{ if } \Delta_i^\mu \sim 0. \quad (23)$$

The more negative ω_i^μ is when $\Delta_i^\mu \sim 0$, the more powerful is its contribution to approach the SVM performances. Equation (23) is the main result of our work: it gives the condition that maximally noisy states \tilde{S}^{μ_d} must satisfy to train an attractor neural network approaching the performance of a SVM. Selecting training data that satisfy equation (23) amounts to imposing specific internal dependencies among the noise units $\tilde{\chi}$, which are no more i.i.d. random variables, as it was in [7]. We refer to such dependencies as *structure* of the noise. One should also bear in mind that training is a dynamic process: to reduce $\mathcal{L}(m = 1^-, J(0^+))$ condition (23) should hold throughout training.

B. Numerical Analysis

As seen in section III, traditional TWN relies on fully random states, i.e. very high states in the energy landscape defined by the function

$$E[\tilde{S}|J] = -\frac{1}{2} \sum_{i,j} S_i J_{ij} S_j. \quad (24)$$

Since E plays the role of the Lyapunov function of the dynamics when couplings J_{ij} are symmetric [1], the dynamic rule in eq. (1) always converges to local minima of E , in the same framework. Inspired by other learning algorithms that make use of particular states in the energy landscape, e.g. the HU procedure described in section IIC, we want to test numerically whether configurations sampled at different altitudes of the landscape are characterized by an internal structure that satisfies condition (23). To do so, we sample training configurations according to the Gibbs-Boltzmann measure by means of a Monte Carlo routine at temperature T . Temperature acts as a control parameter: when $T = 0$ training configurations are stable fixed points of eq. (1). Higher values of T progressively reduce the structure of noise in training configurations, and in the limit $T \rightarrow \infty$, training configurations are the same as in the TWN algorithm. The Monte Carlo of our choice is of the Kawasaki kind [29], to ensure that all training

configurations are at overlap¹ $m_t = 0^+$. We are going to use this technique to probe the states across two types of landscapes: the one resulting from a Hebbian initialization and the one resulting from a Sherrington-Kirkpatrick (SK) model [30], where couplings are i.i.d. Gaussian with zero mean.

Figure 3 reports numerical results for a network initialized according to Hebb's prescription eq. (7), for four different temperatures. Each panel shows the distribution of $(\Delta_i^\mu, \omega_i^\mu)$. Data points are collected over fifteen realizations of the network, then plotted and smoothed to create a density map. We are interested in the *typical* behavior of ω_i^μ when $\Delta_i^\mu = 0$, which can be estimated by a linear fit of the data. We consider the intercept $\omega_{emp}(0)$ of the best fit line as an *indicator* of the typical value of ω_i^μ around $\Delta_i^\mu = 0$. We find that at high temperature $\omega_{emp}(0)$ gets closer to zero, suggesting low quality in terms of training performance. On the other hand, at lower temperatures the sampled training configurations favor both classification and generalization, since $\omega_{emp}(0)$ is more negative.

For a comparison, one can study the distribution of ω_i^μ in the case of a random initialization of the coupling matrix J . We chose the SK model [30] as a case of study. Panels (a) and (b) in fig. 4 report the smoothed distribution of $(\Delta_i^\mu, \omega_i^\mu)$ showing a different scenario with respect to the Hebbian one. The distribution looks anisotropic, as in the Hebbian case, yet the stabilities are centered Gaussians, so $\omega_{emp}(0)$ is positive. In particular, things appear to improve when T increases, in contrast with the previous case of study. Panel (c) shows the Pearson coefficient between ω_i^μ and Δ_i^μ , at various T : both Hebb's and SK show some mutual dependence, but ω_i^μ and Δ_i^μ in the Hebbian landscape appear to be more correlated. Finally, panel (d) shows $\omega_{emp}(0)$ in both cases. This measure is consistent with the indication coming from the Pearson coefficient: while for the Hebb's initialization $\omega_{emp}(0)$ remains significantly negative and reaches the lowest values at low temperatures², the random case shows the opposite trend, with the estimated $\omega_{emp}(0)$ staying generally close to zero.

Going back to the Hebbian energy landscape, one can also study how the distribution of $(\Delta_i^\mu, \omega_i^\mu)$ evolves during the training process. Fig. 5a shows the value of $\omega_{emp}(0)$ at different time steps of the training process, for different values of α , when training configurations are sampled at $T = 0$. We find that $\omega_{emp}(0) < 0$ throughout training for $\alpha \leq 0.6$. The progressive increase of

$\omega_{emp}(0)$ means that fixed points of the initial Hebbian landscape are better training configurations compared to fixed points at intermediate stages of training. In the last part of training, points reacquire more negative values, but this is not a reliable indication of good performance: as shown in fig. 5c, in this part of the process the standard deviation of the couplings σ_i is comparable to $O(\lambda \cdot N^{-1/2})$, and the expansion of the \mathcal{L} in eq. (19) is not valid. The last part of the training, where $\sigma_i \simeq 0 \forall i$, has been neglected from the plot. This picture is confirmed by the study of the Pearson correlation coefficient between ω_i^μ and the associated stabilities Δ_i^μ (see fig. 5b). High values of the Pearson coefficient show a strong dependence of the structure of noise on the relative stabilities. For all α , the Pearson coefficient is highest at $d = 0$, and progressively decreases during training, suggesting that the quality of the training configurations is deteriorating. The final increase in the coefficient is, again, due to the vanishing of the standard deviations σ_i of the couplings, and does not indicate good performance.

V. HEBBIAN UNLEARNING IS TRAINING WITH STRUCTURED NOISE

The total update to the couplings at time D by the symmetric TWN algorithm in eq. (9) can be decomposed as a sum of two contributions

$$\Delta J_{ij}(D) = \Delta J_{ij}^N(D) + \Delta J_{ij}^U(D). \quad (25)$$

The first, which will be referred to as *noise* contribution, is expressed in terms of noise units as

$$\Delta J_{ij}^N(D) = \frac{\lambda}{2N} \sum_{d=1}^D \xi_i^{\mu_d} \xi_j^{\mu_d} \chi_j^{\mu_d} + \frac{\lambda}{2N} \sum_{d=1}^D \xi_j^{\mu_d} \xi_i^{\mu_d} \chi_i^{\mu_d}, \quad (26)$$

while the second, which will be referred to as *unlearning* contribution, is given by

$$\Delta J_{ij}^U(D) = -\frac{\lambda}{2N} \sum_{d=1}^D \left(S_i^{1,\mu_d} S_j^{\mu_d} + S_i^{\mu_d} S_j^{1,\mu_d} \right). \quad (27)$$

Notice that, in the maximal noise case $m_t = 0^+$, $\Delta J_{ij}^N(D)$ averages to zero over the process. When the number of steps D is proportional to N/λ , its variance is of order $O(\lambda/N)$, leading to

$$\Delta J_{ij}^N(D) = O\left(\sqrt{\frac{\lambda}{N}}\right). \quad (28)$$

Hence, when $m_t = 0^+$ and $\lambda/N \rightarrow 0$, the noise contribution is negligible, and the TWN update rule reduces to

$$\delta J_{ij}^{(d)} = -\frac{\lambda}{2N} \left(S_i^{1,\mu_d} S_j^{\mu_d} + S_i^{\mu_d} S_j^{1,\mu_d} \right). \quad (29)$$

¹ For the sake of the experiment, we say that a state has overlap $m_t = 0^+$ with one memory ξ if $m_t \in (0, N^{-1/2})$.

² Specifically, the lowest value of $\omega_{emp}(0)$ in the Hebbian case is reached at low, but non-zero temperature. This suggests that saddles in the energy landscape also can be good training configurations: the study of this particular case is reported in Appendix D.




FIG. 3: Distribution of ω_i^μ as a function of Δ_i^μ for training configurations sampled with a Monte Carlo at temperature $T = 0$ i.e. stable fixed points only (a), $T = 0.5$ (b), $T = 1$ (c), $T = 8$ (d), on a Hebbian network. Warmer colors represent denser region of data points. The *full black* line is the non-weighted best fit line for the points, the *dotted white* line represents $\omega = 0$, the *red dot* is the value of the best fit line associated with $\Delta = 0$. Sub-panels to each panel report a zoom of the line around $\Delta = 0$. Measures have been collected over 15 samples of the network, and observations show that finite size effects are negligible. Choice of the parameters: $N = 500$, $\alpha = 0.5$.




FIG. 4: (a), (b): Distribution of ω_i^μ as a function of Δ_i^μ for training configurations sampled with a Monte Carlo at temperature $T = 0$ i.e. stable fixed points only (a), and $T = 8$ (b) on a SK model. Warmer colors represent denser region of data points. The *full black* line is the non-weighted best fit line for the points, the *dotted white* line represents $\omega = 0$, the *red spot* is the value of the best fit line associated with $\Delta = 0$. Sub-panels to each panel report a zoom of the line around $\Delta = 0$. (c), (d): Comparison between the Hebbian initialization and the Random one through evaluation of: the Pearson coefficient between ω_i^μ and Δ_i^μ (c) and the estimated value of $\omega_{emp}(0)$ from the dispersion plots (d). Measures have been collected over 15 samples of the network. Choice of the parameters: $N = 500$, $\alpha = 0.5$.

In this limit, if fixed points of the dynamics are used for training, i.e. $\tilde{S}^{\mu_d} = \tilde{S}^{1,\mu_d}$, TWN and the HU converge to the same update rule.

Hence, the two algorithms must perform equally when we employ stable fixed points of the dynamics as training data-points. In order to test numerically this statement we proceed in the following manner: the network is initialized according to Hebb's learning rule; we train J by means of: the TWN update rule (18) using stable

fixed points of the dynamics with overlap ³ $m = 0^+$; the HU update rule (8) using stable fixed points of the dynamics randomly sampled in the landscape (as described in section II C). The resulting networks are compared after a fixed amount of iterations D , verifying that they overlap significantly. More precisely, fig. 6 compares the

³ Fixed points are sampled by initializing the network with an overlap in $(0, N^{-1/2})$ with one memory, and then running the dynamics to a fixed point. Typically, such an initial overlap implies the same order of magnitude for the final overlap.

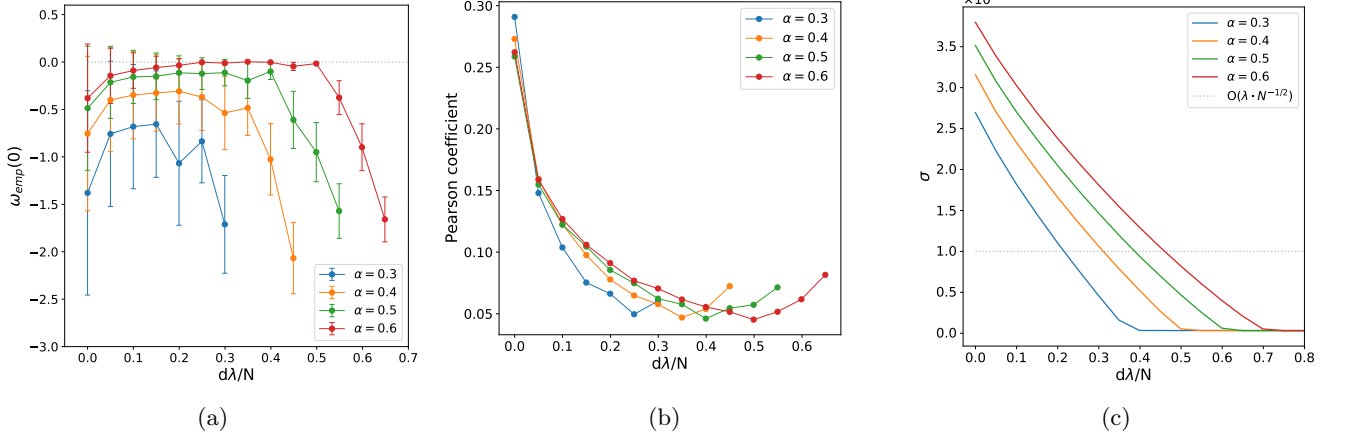


FIG. 5: The TWN algorithm is implemented by sampling stable fixed points of the network dynamics with $m_t = 0^+$. (a) The empirical measure of ω_i^μ around $\Delta_i^\mu = 0$ for the case of stable fixed points as a function of the rescaled number of iterations of the learning algorithm. Error bars are given by the standard deviations of the measures. (b) Pearson coefficient measured between ω_i^μ and Δ_i^μ . (c) The standard deviation of the couplings during learning, defined as $\sigma = \frac{1}{N} \sum \sigma_i$. Points are averaged over 50 samples and the choice of the parameters is: $N = 100$, $\lambda = 10^{-2}$.




FIG. 6: Elements of the J matrix obtained from HU *vs.* the one resulting from TWN for one realization of the memories after the same amount of training steps. Choice of the parameters: $N = 500$, $\alpha = 0.5$, $\lambda = 5 \cdot 10^{-3}$.

elements of the two matrices for one of our experiments, reporting a strong correlation between the two.

Numerically, it is known that HU performs as well as SVMs when $\alpha < \alpha_c^{HU}$ (see [19]). Our analysis of the optimal structure of noise can explain this excellent performance. In fact, as showed by fig. 4d, local minima of an energy landscape initialized according to Hebb's learning rule yield very negative ω_i^μ when $\Delta_i^\mu \sim 0$, driving the network towards the SVM solution.

Interestingly, the weights ω_i^μ have a precise geometric interpretation in the case of HU, where $\omega_i^\mu = m_\mu \chi_i^\mu$. Define the following N -dimensional vectors: \vec{J}_i as the collection of the elements contained in the i^{th} row of the connectivity matrix, $\vec{\eta}_i^\mu = \xi_i^\mu \vec{\xi}^\mu$ and $\vec{\eta}_i^{\mu d} = S_i^{\mu d} \vec{S}^{\mu d}$, where $\vec{S}^{\mu d}$ is the training configuration. Then, at iteration d , one has

$$\omega_i^\mu = (\vec{\eta}_i^{\mu d} \cdot \vec{\eta}_i^\mu) / N. \quad (30)$$

On the other hand, stabilities can be written as

$$\Delta_i^\mu = \frac{\vec{J}_i \cdot \vec{\eta}_i^\mu}{\sqrt{N} \sigma_i}, \quad (31)$$

and the HU rule becomes

$$\delta \vec{J}_i^{(d)} = -\frac{\lambda}{N} \vec{\eta}_i^{\mu d}. \quad (32)$$

It follows that $\delta \Delta_i^\mu \propto -\omega_i^\mu$: a negative value of ω_i^μ will increase the degree of alignment between \vec{J}_i and $\vec{\eta}_i^\mu$, improving the stability. Such an anisotropic distribution of the sampled $\vec{\eta}_i^{\mu d}$ around $\vec{\eta}_i^\mu$ is exactly what condition (23) prescribes for good training configurations. Notice that this anisotropy is hidden in the space of configurations, where minima $\vec{S}^{\mu d}$ of the energy landscape surround memories $\vec{\xi}^\mu$ isotropically. The fact that ω_i^μ are not i.i.d. variables (as one would have for χ_i^μ generated as a consequence of eq. (10)), but rather their statistics depends (i, μ) , is the real physical interpretation of the data-structure. Interestingly, considering only the transformed memories $\vec{\eta}_i^\mu$ that have $\Delta_i^\mu \sim 0$ is sufficient to drag all stabilities above zero.

Finally, it is essential to note that, since the training overlap under consideration is close to 0, stable fixed

points can be sampled in a full unsupervised fashion: when the dynamics is initialized at random (which implies having an overlap $m_t = 0^+$ with some memory) it will typically conserve a small overlap (with the same memory) even at convergence [31]. As a consequence, the argument presented above holds, and a supervised algorithm as TWN can be reduced to a more biologically plausible and faster unsupervised learning rule, as the HU one.

VI. CONCLUSIONS

Gardner’s training-with-noise (TWN) algorithm [7] showed that noise can be injected during training to increase generalization, i.e. the capability of neural networks to retrieve memories by receiving corrupted stimuli [32–35]. In Section III we showed that TWN is well described by Wong and Sherrington’s calculations in [8, 11]. This implies convergence to a Hebbian matrix or a Support Vector Machine (SVM) when learning random configurations with, respectively, maximal (i.e. $m_t = 0^+$) or minimal (i.e. $m_t = 1^-$) amount of noise. The training overlap m_t can be then tuned to interpolate between these two models.

In Section IV we consider the maximal noise scenario represented by $m_t = 0^+$, and derived the condition to be satisfied by the noise in order to approach the performance of a SVM. Specifically, the entries of training configurations (i.e. the data features) are constrained by internal dependencies that we call *structure*. We have explored the landscape of attractors numerically, and showed that in the lower levels of a Hebbian energy landscape there exist a band of maximally noisy saddle states which are excellent training data. Specifically such states, among which we also find local minima of the energy, are able to train a network which shares the same properties of a SVM.

In Section V we proved that, when stable fixed points of the dynamics are used as training data, the TWN algorithm and the HU algorithm are equivalent. This sheds light over why basins of attraction generated by HU are extremely large, as shown in [19]: training data are such as to approach the global minimum of the loss function of a SVM, which is known in literature to have high associativity properties.

As neuroscientists become more convinced of the importance of sleep in memory consolidation [36, 37], connections arise between unsupervised training algorithms on machines and synaptic plasticity processes that occur outside the wakeful hours [22, 38–40]. In light of these remarks and the work presented in this article, one might conceive natural learning as a *two-phase* process. During a first *online* phase, external stimuli are processed by the network through the standard TWN algorithm. One can imagine the stimuli to be maximally noisy versions of some unknown *archetypes* embedded into the environment. As a consequence, training will shape a pure Hebbian landscape of attraction out of the retrieval regime. In a second *offline* phase, the early formed network samples structured noisy neural configurations, still weakly correlated with the archetypes, from the landscape of attractors. Such states could be, for instance, stable fixed points of the neural dynamics. These neural configurations are then processed by the same kind of TWN algorithm, and memory is consolidated by centering the unconscious archetypes in the middle of large basins of attraction.

In this context, the results contained in this work make progress on two fronts. On one side, they shed light over the structure of noise that is optimal for learning in neural networks, helping to develop a finer theory behind the very empirical techniques of noise injection implemented in training deep networks [3, 32, 41, 42]. On the other side, they draw a connection between unsupervised, and thus more biologically relevant, learning processes and the supervised ones, from which most of the modern theory of neural networks derive.

ACKNOWLEDGMENTS

The authors are particularly grateful to their mentors Enzo Marinari, Giancarlo Ruocco and Francesco Zamponi for precious suggestions and support. They also thank Fabian Aguirre Lopez, Aldo Battista, Simona Cocco, Giampaolo Folena, Rémi Monasson and Mauro Pastore for useful discussions, as well as an anonymous referee for useful comments that have improved the manuscript.

-
- [1] D.J. Amit. *Modeling Brain Function*. Cambridge University Press, 1989.
 - [2] M. Mezard, G. Parisi, and M. Virasoro. *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications*, volume 9 of *World Scientific Lecture Notes in Physics*. World Scientific, 1986.
 - [3] C. Shorten and T. M. Khoshgoftaar. A survey on Image Data Augmentation for Deep Learning. *Journal of Big Data*, 6(1):1, 2019.
 - [4] L. Zhao, T. Liu, X. Peng, and D. Metaxas. Maximum-Entropy Adversarial Data Augmentation for Improved Generalization and Robustness. In *Advances in Neural Information Processing Systems*, volume 33, pages 14435–14447. Curran Associates, Inc., 2020.
 - [5] A. Achille and S. Soatto. Information dropout: Learning optimal representations through noisy computation. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2897–2905, 2018.

- [6] J.J. Hopfield, D.I. Feinstein, and R.G. Palmer. ‘Unlearning’ has a stabilizing effect in collective memories. *Nature*, 304(5922):158, 1983.
- [7] E. J. Gardner, D. J. Wallace, and N. Stroud. Training with noise and the storage of correlated patterns in a neural network model. *Journal of Physics A: Mathematical and General*, 22(12):2019, 1989.
- [8] K.Y.M. Wong and D. Sherrington. Neural networks optimally trained with noisy data. *Physical Review E*, 47(6):4465, 1993.
- [9] D. J. Amit, H. Gutfreund, and H. Sompolinsky. Storing Infinite Numbers of Patterns in a Spin-Glass Model of Neural Networks. *Physical Review Letters*, 55(14):1530, 1985.
- [10] E. Gardner. Structure of metastable states in the Hopfield model. *Journal of Physics A: Mathematical and General*, 19(16):L1047, 1986.
- [11] K.Y.M. Wong and D. Sherrington. Optimally adapted attractor neural networks in the presence of noise. *Journal of Physics A: Mathematical and General*, 23(20):4659, 1990.
- [12] D.O. Hebb. *The Organization of Behavior : A Neuropsychological Theory*. John Wiley and Sons, 1949.
- [13] E. Gardner. The space of interactions in neural network models. *Journal of Physics A: Mathematical and General*, 21(1):257, 1988.
- [14] E. Gardner, H. Gutfreund, and I. Yekutieli. The phase space of interactions in neural networks with definite symmetry. *Journal of Physics A: Mathematical and General*, 22(12):1995, 1989.
- [15] M. Minsky and P. Seymour. *Perceptrons: an introduction to computational geometry*. MIT Press, 1969.
- [16] A. Battista and R. Monasson. Capacity-resolution trade-off in the optimal learning of multiple low-dimensional manifolds by attractor neural networks. *Physical Review Letters*, 124(4):048302, 2020.
- [17] N. Brunel, V. Hakim, P. Isope, J.P. Nadal, and B. Barbour. Optimal Information Storage and the Distribution of Synaptic Weights: Perceptron versus Purkinje Cell. *Neuron*, 43(5):745–757, 2004.
- [18] N. Brunel. Is cortical connectivity optimized for storing information? *Nature Neuroscience*, 19(5):749–755, 2016.
- [19] M. Benedetti, E. Ventura, E. Marinari, G. Ruocco, and F. Zamponi. Supervised perceptron learning vs unsupervised Hebbian unlearning: Approaching optimal memory retrieval in Hopfield-like networks. *The Journal of Chemical Physics*, 156(10):104107, 2022.
- [20] B. Schölkopf and A.J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. The MIT Press, 2018.
- [21] J.J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554, 1982.
- [22] F. Crick and G. Mitchison. The function of dream sleep. *Nature*, 304(5922):111, 1983.
- [23] J.L. Van Hemmen, L.B. Ioffe, R. Kühn, and M. Vaas. Increasing the efficiency of a neural network through unlearning. *Physica A: Statistical Mechanics and its Applications*, 163(1):386, 1990.
- [24] J.L. van Hemmen and N. Klemmer. Unlearning and Its Relevance to REM Sleep: Decorrelating Correlated Data. In J.G. Taylor, C.L.T. Mannion, J.G. Taylor, E.R. Cianiello, R.M.J. Cotterill, and J.W. Clark, editors, *Neural Network Dynamics*, page 30. Springer London, London, 1992.
- [25] E. Ventura, S. Cocco, R. Monasson, and F. Zamponi. Unlearning regularization in Boltzmann Machines. *arXiv:2311.09418*, 2024.
- [26] M. Benedetti, L. Carillo, E. Marinari, and M. Mézard. Eigenvector dreaming. *Journal of Statistical Mechanics: Theory and Experiment*, 2024(1):013302, 2024.
- [27] Y. Le Cun. Learning Process in an Asymmetric Threshold Network. In *Disordered Systems and Biological Organization*, page 233. Springer, Berlin, Heidelberg, 1986.
- [28] B.M. Forrest. Content-addressability and learning in neural networks. *Journal of Physics A: Mathematical and General*, 21(1):245, 1988.
- [29] M. Newman and G. Barkema. *Monte Carlo Methods in Statistical Physics*. Oxford University Press, 1999.
- [30] D. Sherrington and S. Kirkpatrick. Solvable Model of a Spin-Glass. *Physical Review Letters*, 35(26):1792, 1975.
- [31] T.B. Kepler and L.F. Abbott. Domains of attraction in neural networks. *Journal de Physique*, 49(10):1657, 1988.
- [32] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929, 2014.
- [33] T. Tadros, G. Krishnan, R. Ramyaa, and Bazhenov M. Biologically inspired sleep algorithm for increased generalization and adversarial robustness in deep neural networks. *International Conference on Learning Representations*, 2019.
- [34] D. Saad and S. Solla. Learning with Noise and Regularizers in Multilayer Neural Networks. In *Advances in Neural Information Processing Systems*, volume 9. MIT Press, 1996.
- [35] B. Schottky and U. Krey. Phase transitions in the generalization behaviour of multilayer perceptrons: II. The influence of noise. *Journal of Physics A: Mathematical and General*, 30(24):8541, 1997.
- [36] J.D. Creery, D.J. Brang, J.D. Arndt, A. Bassard, V.L. Towle, J.X. Tao, S. Wu, S. Rose, P.C. Warnke, N.P. Issa, and K.A. Paller. Electrophysiological markers of memory consolidation in the human brain when memories are reactivated during sleep. *Proceedings of the National Academy of Sciences*, 119(44):e2123430119, 2022.
- [37] N. Maingret, G. Girardeau, R. Todorova, M. Goutier, and M. Zugaro. Hippocampo-cortical coupling mediates memory consolidation during sleep. *Nature Neuroscience*, 19(7):959, 2016.
- [38] G. Girardeau and V. Lopes-dos Santos. Brain neural patterns and the memory function of sleep. *Science*, 374(6567):560–564, 2021.
- [39] G. Hinton. The Forward-Forward Algorithm: Some Preliminary Investigations. *arXiv:2212.13345*, 2022.
- [40] J. Hinton and T.J. Sejnowski. *Unsupervised Learning: Foundations of Neural Computation*. The MIT Press, 1999.
- [41] U.M. Tomasini, L. Petrini, F. Cagnetta, and M. Wyart. How deep convolutional neural networks lose spatial information with training. *Machine Learning: Science and Technology*, 4(4):045026, 2023.
- [42] L. Bonnasse-Gahot and J.P. Nadal. Categorical Perception: A Groundwork for Deep Learning. *Neural Computation*, 34(2):437, 2022.
- [43] S. Diamond and S. Boyd. CVXPY: A Python-Embedded Modeling Language for Convex Optimization. *Journal of Machine Learning Research*, 17:1, 2016.

- [44] V.S. Dotsenko, N.D. Yarunin, and E.A. Dorotheyev. Statistical mechanics of Hopfield-like neural networks with modified interactions. *Journal of Physics A: Mathematical and General*, 24(10):2419–2429, 1991.
- [45] T. Aspelmeier, R.A. Blythe, A.J. Bray, and M.A. Moore.

- Free energy landscapes, dynamics and the edge of chaos in mean-field models of spin glasses. *Physical Review B*, 74(18):184411, 2006.
- [46] S. Kirkpatrick, C.D. Gelatt, and M.P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671, 1983.

Appendix A: Descending over the \mathcal{L} function

1. Training with noise

At each step of the algorithm, a memory label μ_d is sampled at random and the update (9) is performed over the couplings. The new value of the \mathcal{L} function (13) is

$$\mathcal{L}' = -\frac{1}{\alpha N^2} \sum_{i,\mu}^{N,p} \operatorname{erf} \left(\frac{m\Delta_i^\mu}{\sqrt{2(1-m^2)}} + \frac{\lambda m}{N\sigma_i \sqrt{2N(1-m^2)}} \epsilon_i^{\mu_d} \xi_i^\mu \xi_i^{\mu_d} \sum_{j \neq i} S_j^{\mu_d} \xi_j^\mu \right), \quad (\text{A1})$$

where

$$\epsilon_i^{\mu_d} = \Theta \left(\xi_i^{\mu_d} \sum_{k=1}^N J_{ik} S_k^{\mu_d} \right). \quad (\text{A2})$$

Since $\delta\sigma_i \propto \frac{\lambda}{N} \frac{J_i}{\sigma_i} + O\left(\frac{\lambda}{N}\right)^3$ and the mean J_i of the couplings along line i equals zero by initialization and it is naturally maintained null during the algorithm, we have considered $\sigma'_i \simeq \sigma_i$. Then \mathcal{L}' can be rewritten as

$$\begin{aligned} \mathcal{L}' = & -\frac{1}{\alpha N^2} \sum_{i,\mu \neq \mu_d}^{N,p} \operatorname{erf} \left(\frac{m\Delta_i^\mu}{\sqrt{2(1-m^2)}} + O\left(\frac{1}{N}\right) \right) - \\ & -\frac{1}{\alpha N^2} \sum_i^N \operatorname{erf} \left(\frac{m\Delta_i^{\mu_d}}{\sqrt{2(1-m^2)}} + \frac{\lambda m \cdot m_t}{\sigma_i \sqrt{2N(1-m^2)}} \epsilon_i^{\mu_d} + O\left(\frac{1}{N}\right) \right), \end{aligned} \quad (\text{A3})$$

where we have used that $\frac{1}{N} \sum_{j \neq i} \xi_j^\mu S_j^{\mu_d} = O(N^{-1/2})$ when $\mu \neq \mu_d$ and $m_t = O(1)$. We thus expand the errorfunction at the first order in $O(N^{-1/2})$ obtaining the variations to \mathcal{L} in equation (16).

As a technical comment, note that standard deviation along one row of the couplings matrix σ_i (see eq. (2)) is a variable quantity over time, and numerics suggest it is slowly decreasing. As a result, the expansion performed to determine the variation of \mathcal{L} (see eq. (A3) in Appendix A 1) might not be justified after a certain number of steps, leading to a non-monotonic trend of the loss function. The non-monotonic trend of $\mathcal{L}(m_t, J(m_t))$ due to this effect is shown in the inset of fig. 7. However, this inconvenience can be overcome by rescaling the learning rate λ into $\lambda_i = \lambda \cdot \sigma_i$ at each iteration, obtaining the curves in fig. 7.

2. Training with structured noise

The new value of \mathcal{L} is derived by using equation (17) to evaluate the variation of stabilities

$$\begin{aligned} \mathcal{L}' = & -\frac{1}{\alpha N^2} \sum_{i,\mu}^{N,p} \operatorname{erf} \left(\frac{m\Delta_i^\mu}{\sqrt{2(1-m^2)}} + \frac{\lambda m \xi_i^\mu \xi_i^{\mu_d}}{2N\sigma_i \sqrt{2N(1-m^2)}} \sum_{j=1}^N S_j^{\mu_d} \xi_j^\mu + \frac{\lambda m \xi_i^\mu S_i^{\mu_d}}{2N\sigma_i \sqrt{2N(1-m^2)}} \sum_{j=1}^N \xi_j^{\mu_d} \xi_j^\mu - \right. \\ & \left. - \frac{\lambda m \xi_i^\mu S_i^{1,\mu_d}}{2N\sigma_i \sqrt{2N(1-m^2)}} \sum_{j=1}^N S_j^{\mu_d} \xi_j^\mu - \frac{\lambda m \xi_i^\mu S_i^{\mu_d}}{2N\sigma_i \sqrt{2N(1-m^2)}} \sum_{j=1}^N S_j^{1,\mu_d} \xi_j^\mu \right), \end{aligned} \quad (\text{A4})$$

where $\sigma'_i \simeq \sigma_i$ as in the previous paragraph. We now redefine $\chi_i^\mu = \xi_i^\mu S_i^{\mu_d}$, $\chi_i^{1,\mu} = \xi_i^\mu S_i^{1,\mu_d}$, $m_\mu = \frac{1}{N} \sum_{j=1}^N S_j^{\mu_d} \xi_j^\mu$ and $m_{1,\mu} = \frac{1}{N} \sum_{j=1}^N S_j^{1,\mu_d} \xi_j^\mu$ and expand the errorfunction at the first order in $O(N^{-1/2})$ obtaining

$$\delta\mathcal{L} = \frac{m\lambda}{\sqrt{2\pi\alpha^2 N^5 (1-m^2)}} \sum_{i,\mu}^{N,p} \frac{1}{2\sigma_i} \left[(m_\mu \chi_i^{1,\mu} + m_{1,\mu} \chi_i^\mu) - (m_\mu \xi_i^\mu \xi_i^{\mu_d} + M_\mu^{\mu_d} \chi_i^\mu) \right] \exp \left(-\frac{m^2 \Delta_i^{\mu^2}}{2(1-m^2)} \right), \quad (\text{A5})$$




FIG. 7: The *blueish* lines in the main plot report the function $\mathcal{L}(m = 0.5, J(m_t))$ for different training overlaps as functions of the number of algorithm steps d . The *dotted* line represents the theoretical minimum value from [8]. The learning strength λ has been rescaled by the standard deviation of the couplings as described in the text. The subplot reports the case $m_t = 0.5$ when the learning strength is not rescaled: \mathcal{L} is in *blue*, while a measure of the standard deviation of the couplings, defined as $\sigma = \frac{1}{N} \sum \sigma_i$, is reported in *red*: the value $\lambda \cdot N^{-1/2}$ of the standard deviation is depicted in *light gray* to properly signal the moment when equation (16) loses its validity. All measures are averaged over 5 realizations of the couplings J . Choice of the parameters: $N = 100$, $\alpha = 0.3$, $\lambda = 1$, the initial couplings are Gaussian with unitary mean, zero variance and $J_{ii}^{(0)} = 0 \forall i$.

where $M_\mu^{\mu_d} = \frac{1}{N} \sum_{i=1}^N \xi_i^\mu \xi_i^{\mu_d}$. Equation (A5) can be decomposed in

$$\delta \mathcal{L} = \delta \mathcal{L}_N + \delta \mathcal{L}_U, \quad (\text{A6})$$

where $\delta \mathcal{L}_U$ contains the weight

$$\omega_i^\mu = \frac{1}{2\sigma_i} \left(m_\mu \chi_i^{1,\mu} + m_{1,\mu} \chi_i^\mu \right), \quad (\text{A7})$$

while $\delta \mathcal{L}_N$ contains

$$\Omega_i^\mu = \frac{1}{2\sigma_i} \left(m_\mu \xi_i^\mu \xi_i^{\mu_d} + M_\mu^{\mu_d} \chi_i^\mu \right). \quad (\text{A8})$$

We study the two contributions numerically, on a Hebbian network, i.e. with no learning going on, for the case of $m_t = 0^+$. The Pearson coefficient is measured between the vector of the stabilities Δ_i^μ and the weights ω_i^μ as well as with Ω_i^μ separately. This quantity should underline an eventual reciprocal dependence between $\omega_i^\mu, \Omega_i^\mu$ and Δ_i^μ . The test is repeated over states sampled by a Monte Carlo at different temperatures T . Results are reported in fig. 8a where it is evident that Ω_i^μ does not have any correlation with Δ_i^μ , while the dependence of ω_i^μ on the stabilities is evident. Moreover, we measured the indicator $\omega_{emp}(0)$, signaling the typical values of the weights ω_i^μ and Ω_i^μ when $\Delta_i^\mu \sim 0$ (see Section IV for further details), as reported in fig. 8b. The plot clearly shows that Ω_i^μ is small and generally fluctuating around zero. These aspects hold during the training procedure also, as it can be observed by performing the same measure at different step of the TWN procedure over states with $m_t = 0^+$. We will thus refer to $\delta \mathcal{L}_U$ as the relevant contribution to the variation of the function \mathcal{L} .




FIG. 8: (a) Pearson coefficient between ω_i^μ , Ω_i^μ and the stabilities Δ_i^μ , (b) $\omega_{emp}(0)$ for the case of a Hebbian network at different temperatures T . Configurations at a given temperature T have been sampled by a Monte Carlo of the Kawasaki kind, in order to choose only maximally noisy states ($m_t = 0^+$). Points are collected from 15 samples of the network. Choice of the parameters: $N = 500$, $\alpha = 0.5$.

Appendix B: Training with noise - Classification analysis

We now resume the computations performed in [8, 11] to analyze the classification performance of TWN. We restrict ourselves to a spherical space of interactions such that

$$\sum_{j \neq i} J_{ij}^2 = N \quad \forall i, \quad (B1)$$

We want to compute the probability distribution of the stabilities that now become

$$\Delta_i^\mu = \xi_i^\mu \sum_{j=1}^N \frac{J_{ij}}{\sqrt{N}} \xi_j^\mu. \quad (B2)$$

The partition function of the model is then given by

$$Z = \int \prod_{i,j} dJ_{ij} \delta \left(\sum_{j \neq i} J_{ij}^2 - N \right) \exp \left(-\beta \mathcal{L}(m_t, J) \right), \quad (B3)$$

where β is the inverse annealing temperature of the problem and the loss function \mathcal{L} is defined in eq. (13). The distribution of the stabilities is

$$\rho_{m_t}(\Delta) = \overline{\frac{1}{Z} \int \prod_j dJ_{\bullet j} \delta \left(\sum_j J_{\bullet j}^2 - N \right) \exp \left(\beta \sum_\mu \text{erf} \left(\frac{m_t \xi_\bullet^\mu \sum_j J_{\bullet j} \xi_j^\mu}{\sqrt{2N(1-m_t^2)}} \right) \right) \delta \left(\xi_\bullet^1 \sum_j \frac{J_{\bullet j}}{\sqrt{N}} \xi_j^1 - \Delta \right)}, \quad (B4)$$

where $\overline{\cdot}$ denotes the average over the realizations of the memories and we have neglected the factorization over i , since we treat the optimization process as independent along the lines of the J_{ij} matrix. Index i has been nevertheless substituted by \bullet for the sake of completeness. Replicas can be used to evaluate the normalization, i.e.

$$1/Z = \lim_{n \rightarrow 0} Z^{n-1}. \quad (B5)$$

The replica calculation in the replica symmetric ansatz [2] when $\beta \rightarrow \infty$ leads to the following expression for the density of the stabilities

$$\rho_{m_t}(\Delta) = \frac{1}{\sqrt{2\pi}} \left(1 + \sqrt{\frac{2}{\pi}} \frac{m_t^3 \chi}{(1 - m_t^2)^{3/2}} \Delta \exp \left(-\frac{m_t^2 \Delta}{2(1 - m_t^2)} \right) \right) \exp -\frac{w(\Delta)^2}{2}, \quad (\text{B6})$$

where χ and w are derived by solving the following two equations

$$w = x - \frac{\sqrt{2} m_t \chi}{\sqrt{\pi(1 - m_t^2)}} \exp \left(-\frac{m_t^2 x^2}{2(1 - m_t^2)} \right), \quad (\text{B7})$$

$$\int_{-\infty}^{+\infty} Dw (x^*(w, \chi) - w)^2 = \alpha^{-1}, \quad (\text{B8})$$

with x^* being the solution of (B7) and Dw being the standard Gaussian measure with zero mean and unitary variance.

Appendix C: Training with noise - Generalization analysis

We report here a general experimental procedure to measure the average size of the basins of attraction of a fully connected neural network of finite size N and a given choice of the control parameters.

The network is firstly trained according to an algorithm of our choice. Once the couplings have been found, the asynchronous version of dynamics (1) is initialized in one of the memories. The dynamics is run until convergence onto the attractor associated to the basin of belonging of the memory. Now the retrieval map $m_f(m_0)$ is measured with respect to that particular attractor, and the procedure is repeated over different memories and realizations of the network. The average radius of the basin of attraction is then measured as the value of $1 - m_0$ where $m_f(m_0)$ equals a reference value. In our case such value is $m_f = 0.98$. We have applied this procedure on networks trained either as SVMs and with the TWN algorithm. In the former case a convex algorithm contained in the *cvxpy* Python domain [43] is implemented to train the network. To be more specific, N independent machines are trained to correctly classify $p = \alpha N$ binary memories of the kind of $\tilde{\xi}^\mu \in \{-1, +1\}^N$ having as labels ξ_i^μ with $i \in [1, \dots, N]$.

Regarding the dynamics, the stability of fixed points is in general implied by some properties of the couplings, mainly their degree of symmetry. For the case of the TWN algorithm we start from a random symmetric matrix, as done in [31]: the update of the couplings will only perturb the initial symmetry yet allowing the measures to be still consistent with the theory. On the other hand, numerics show that SVMs are sufficiently symmetric to let the asynchronous dynamics converge. The comparison between the retrieval maps obtained for the two algorithms with $\alpha = 0.45$ and $N = 200$ is reported in fig. 9. The curve relative to the SVM is fixed, while the one associated with the TWN is changing with respect to the training overlap m_t .

Even if the SVM always reaches classification of the memories for $\alpha < 2$ [13], a network trained with noise might show two different behaviors: one associated to *retrieval* where each memory is close to an attractor, and one related to *non-retrieval* where the memory is far from its attractor and the basins of attraction might contain orthogonal configurations with respect to the central attractor. In particular network models where couplings are assembled according to particular rules (e.g. [21, 44]) the transition between these two regimes can be computed analytically. In the case of TWN this cannot be done. It is then important to find an empirical criterion to divide the two behaviors as a function of (m_t, α) . Let us assume that when $N \gg 1$ the retrieval map $m_f(m_0)$ develops a plateau starting from $m_0 = 1$ and ending in some limit value $m_0 = m_c < 1$ such that $m_f = 1$ along all this interval (as in fig. 9). Hence, one can associate the formation of such a plateau with the existence of a cohesive basin of attraction, where close configurations in hamming distance to the attractor converge to the attractor. One then wants to measure the value of m_t at which such property of the basin disappears. As a possible estimate, it is convenient to consider the m_t^* such that

$$\left. \frac{dm_f}{dm_0}(m_t^*) \right|_{m_0=1} = 1. \quad (\text{C1})$$

The numerical extrapolation of the overlap in fig. 2 at different values of N shows a good agreement between the approximate separation between the two regions showed in fig. 2 and the line estimated by condition (C1).




FIG. 9: Retrieval map $m_f(m_0)$ in the case of networks trained through SVM and TWN algorithms. Curves are shown as a function of m_t and compared with the bisector, indicated with a *dashed black* line. Points are averaged over 10 samples. Choice of the parameters: $N = 200$, $\alpha = 0.45$.

Appendix D: Training with Saddles

Notice, from both panels (c), (d) of fig. 4, the existence of an optimum which does not coincide with the stable fixed points of the dynamics (i.e. $T = 0$). As noted in previous studies on spin glasses [45], one can associate configurations probed by a Monte Carlo at finite temperatures with configurations which are typically saddles in the energy landscape with a given *saddle index*. The *saddle index* is defined as the ratio between the number of unstable sites under the dynamics (see eq. (1)) and the total number of directions N . Stable fixed points have $f = 0$, while random configurations are expected to have $f = 1/2$. In order to check whether a particular f is capturing relevant features of the virtuous training configurations, we sampled training data according to the requirement that their saddle fraction assumes a specific value f and $m_t = 0^+$. Saddles are then employed for training the network according to eq. (9). Sampling is performed by randomly initializing the network on a configuration having training overlap $m_t = 0^+$ with a reference memory, and performing a zero temperature dynamics on the landscape defined by the new energy

$$\mathcal{E}(\vec{S}|f, J) = \frac{1}{2} \left(\frac{1}{N} \sum_{i=1}^N \Theta(-S_i \sum_{j=1}^N J_{ij} S_j) - f \right)^2, \quad (\text{D1})$$

where $\Theta(x)$ is the Heaviside function. Yet again, the value of m_t was maintained constant during the descent. The left panel in fig. 10 shows how the minimum stability evolves during the training process while a TWN algorithm is initialized in the Hebbian matrix and learns saddles of different indices. For a network of $N = 100$ and $\alpha = 0.35$, we found that classification (i.e. positive minimum stability) is reached until a certain value of f , suggesting that saddles belonging to this band are indeed good training data. The band of saddles that are suitable for learning shrinks when α increases until such states do not significantly satisfy eq. (23) anymore. Such limit capacity is located around the critical one for HU. It should be stressed that the precise performance as a function of f is quite sensitive to the sampling procedure. Simulated annealing routines [46] have also been employed to minimize (D1), obtaining qualitatively similar results yet not coinciding with the ones reported in fig. 10. A qualitative study of the basins of attraction of the network has been performed and reported in the right panel in fig. 10. Specifically, the retrieval map $m_f(m_0)$ has been measured relatively to the saddle indices f at the first time they reached classification, in analogy to what has been measured in [19]. The curves coincide quite well, suggesting that finite sized networks trained with

different f assume similar volumes of the basins of attraction when they are measured at the very first instant they reach classification. The plot also shows that the generalization performance is comparable with the one of a SVM trained with the same choice of the control parameters.




FIG. 10: Left: Minimum stability $\Delta_{\min} = \min_{i,\mu}(\Delta_i^\mu)$ as a function of the algorithm steps on a network trained with the TWN routine that learns saddles of various indices f . The initial matrix is assembled according to the Hebb's rule. Full dots report the amount of iterations needed to accomplish classification. Right: the retrieval map $m_f(m_0)$ as measured on the positions of the colored dots from the right panel, with the same color code being applied. A comparison with a SVM trained with the same choice of the parameters is also presented through the *dashed blue* line. All measures are averaged over 5 samples with the shaded region indicating the experimental errors. The choice of the parameters is: $N = 100$, $\alpha = 0.35$, $\lambda = 10^{-3}$.