



A new mechanical approach to handle generalized Hopfield neural networks

Adriano Barra^{a,b,c,*}, Matteo Beccaria^{a,b}, Alberto Fachechi^{a,b}

^a Dipartimento di Matematica e Fisica Ennio De Giorgi, Università del Salento, Lecce, Italy

^b INFN, Istituto Nazionale di Fisica Nucleare, Sezione di Lecce, Italy

^c GNFM-INdAM, Gruppo Nazionale per la Fisica Matematica, Sezione di Lecce, Italy

ARTICLE INFO

Article history:

Received 23 February 2018

Received in revised form 11 July 2018

Accepted 14 July 2018

Available online 21 July 2018

The authors are pleased to dedicate this paper to Carlo Sempì in honor of his seventieth birthday

Keywords:

Hopfield networks

Statistical mechanics

Unlearning

Non-linear PDE

Relativistic mechanics

ABSTRACT

We propose a modification of the cost function of the Hopfield model whose salient features shine in its Taylor expansion and result in more than pairwise interactions with alternate signs, suggesting a unified framework for handling both with deep learning and network pruning. In our analysis, we heavily rely on the Hamilton–Jacobi correspondence relating the statistical model with a mechanical system. In this picture, our model is nothing but the relativistic extension of the original Hopfield model (whose cost function is a quadratic form in the Mattis magnetization and mimics the non-relativistic counterpart, the so-called *classical limit*). We focus on the low-storage regime and solve the model analytically by taking advantage of the mechanical analogy, thus obtaining a complete characterization of the free energy and the associated self-consistency equations in the thermodynamic limit. Further, on the numerical side, we test the performances of our proposal with extensive Monte Carlo simulations, showing that the stability of spurious states (limiting the capabilities of the standard Hebbian construction) is sensibly reduced due to presence of unlearning contributions that *prune* them massively.

© 2018 Elsevier Ltd. All rights reserved.

1. Introduction

In the past few years, thanks both to progresses in hardware (possibly due to the novel generation of GPU computing architectures (Kirk, 2007; Nickolls et al., 2008)), as well as in software (e.g. driven by results from Google and related companies as Deep Mind), we assisted at the rise of a novel and more powerful generation of Artificial Intelligence, whose most popular contributions perhaps lie in the impressive adaptive capability of *Deep Learning* (Le Cun, Bengio, & Hinton, 2015; Schmidhuber, 2015) and in the creative compositional episodes during *sleeping and dreaming* (Hern, 2015; Tubiana & Monasson, 2017).

Regarding the former, deep learning, beyond literally grounding Artificial Intelligence within the theoretical framework of *disordered statistical mechanics* (Amit, 1989; Hopfield, 1982), Hopfield recently offered also a connectionist perspective where the high skills of deep learning machines could possibly be framed (Kroto & Hopfield, 2016): the route he proposes is via many-body extensions of his celebrated pairwise model for feature retrieval and pattern recognition (Amit, 1989; Coolen, Kuhn, & Sollich, 2005;

Hopfield, 1982). The core of the idea is as elegant as simple, as already pointed out by Hopfield and Tank in the past (Hopfield & Tank, 1985): consider N on/off neurons $\sigma_i = \pm 1$ (where i labels neurons from 1 to N), if we want them to retrieve one out of P random patterns ξ^μ (μ labeling the patterns from 1 to P) and we want to describe this property via a cost-function $H(\sigma|\xi)$ that resembles Hamiltonians in Physics (such that the minima of the Hamiltonian would match the patterns themselves (Amit, 1989; Mezard, Parisi, & Virasoro, 1985; Rolls & Treves, 1998)), the simplest and most natural guess would be summing all the squared scalar products between the neurons and the patterns, i.e., $H(\sigma|\xi) \propto -\sum_{\mu}^P (\sigma \cdot \xi^\mu)^2$: for large N , patterns become orthogonal¹ and if the state vector σ is uncorrelated with all of them, each parenthesized term would be very small, but if the state network σ retrieves one of the P patterns (i.e., it becomes highly correlated with the pattern), then its contribution in the summation would be no longer negligible and the pattern would actually act as an attractor for any reasonable stochastic neural dynamics obeying Detailed Balance (Amit, 1989; Coolen et al., 2005).

* Corresponding author at: Dipartimento di Matematica e Fisica Ennio De Giorgi, Università del Salento, Lecce, Italy.

E-mail addresses: adriano.barra@unisalento.it (A. Barra), matteo.beccaria@le.infn.it (M. Beccaria), alberto.fachechi@le.infn.it (A. Fachechi).

¹ With the term *orthogonal* we mean that $\lim_{N \rightarrow \infty} N^{-1} \xi^\mu \xi^\nu = \lim_{N \rightarrow \infty} N^{-1} \sum_i \xi_i^\mu \xi_i^\nu = \delta_{\mu\nu}$. However, at finite N , this is a N -long sum of terms whose probability of being ± 1 is one half: it is a random walk with zero mean and variance N , hence spurious correlations are expected to vanish $\propto 1/\sqrt{N}$.

The robustness of this argument lies in the usage of just local convexity of the cost function and can be generalized straightforwardly beyond the parabolic approximation coded by the pairwise interactions, for instance by including also the even higher order contributions (the so-called *P-spin* terms): remembering that, from the connectionist perspective (Amit, 1989; Fodor & Pylyshyn, 1988; McCloskey & Cohen, 1989), memories lie in the connections (they are stored in the slowly evolving values of the synapses as learning takes place), clearly adding more and more *P-spin* contributions to the Hamiltonian adds more and more synapses where information can be filed.

Regarding the latter, sleeping and dreaming, while the idea that sleep can actually consolidate memories by discarding fake information accidentally stored is more than two centuries old (Hartley, 1791), its inspection has improved significantly since the discovery of REM sleep² in the fifties as pioneered by Aserinsky and Kleitman (Aserinsky & Kleitman, 1953; Stickgold, Hobson, Fosse, & Fosse, 2001). Models of REM sleep have already been framed within a neural network perspective by Crick and Mitchinson (Crick & Mitchinson, 1983) and Hopfield himself (Hopfield, Feinstein, & Palmer, 1983): the whole gave rise to the theory of *unlearning* in neural networks (Dotsenko, Yarunin, & Dorotheyev, 1991; Kleinfeld & Pendergraft, 1987; Nokura, 1996, 1998), a very nice idea to remove spurious states from the landscape of retrievable patterns (and one of the aims of the present work is to show that unlearning in Hebbian models can be used for pruning these states in the network).

Up to nowadays, however, these two branches of neural networks – deep learning and unlearning – never merged and the main aim of the present work is to obtain (in a very natural way) a unified model able to handle both of these features at once.

With this goal in our mind, we now turn to methodologies rather than subjects. Initially with the focus on spin glasses (Barra, Di Biasio, & Guerra, 2010; Guerra, 2001), in the past ten years, several contributions linking the statistical mechanical recipes to analytical mechanical approaches appeared in the Literature (see e.g. (Agliari et al., 2013a, b; Arsie et al., 2015a, b; Barra et al., 2014; Guerra et al., 2014) and references therein). To fix the ideas, let us consider the Curie–Weiss model (Barra, 2008; Genovese et al., 2009) as the simplest example of a pairwise spin Hamiltonian (that can be thought of as a Hopfield model with a unique pattern under the Mattis gauge (Amit, 1989)): it has been proved that its free energy obeys a standard (*i.e. classical*) Hamilton–Jacobi PDE in the space of the coupling and external field (where the former plays as time, and the latter as space in this mechanical analogy): to infer statistical properties on the Curie–Weiss network, we can thus use this mechanical analogy and study a fictitious particle of unitary mass classically evolving in this 1+1 space–time (under the proper PDE derived from the statistical mechanical framework). Its properties, once translated back in the original setting, recover sharply all the results of the standard route (Barra, 2008; Genovese et al., 2009). It is thus possible to perform a whole statistical mechanical analysis of the model (e.g., obtain an explicit expression for the free energy, the self consistent equations for the order parameters, etc.) upon relying solely on techniques typical of *analytical mechanics* (Arnold, Kozlov, & Neishtadt, 2007).

Here, once this framework is adapted to the Hopfield model, we show that such an analogy calls for a very natural generalization of the Hopfield cost-function (or *Hamiltonian* (Amit, 1989)), that is simply its relativistic extension. Indeed, we will see how, within the *mechanical analogy*, moving from the classical kinetic energy to its relativistic counterpart, Hamilton–Jacobi PDE reduces to the

energy–momentum relation (Mandl & Shaw, 2010) and the free energy – that coincides with the *action* in this equivalence – turns out to obey a relativistic Least Action Principle. While the classical expression for the kinetic energy is a second-order monomial (in the Mattis order parameters), *i.e.* the Hopfield cost-function, its relativistic expression is not such but can be Taylor-expanded in all the monomials. We remark that these turn out to be exactly solely the even ones and with alternate signs: the relativistic extension naturally suggested by the mechanical analogy accounts for the higher-order contributions (beyond Hopfield’s pairwise one), hence of potential interest for researcher in Deep Learning, yet being appealing for researches in unlearning, given the alternation of signs in the series. In those regards (*i.e.* focusing on Machine Learning), it is worth pointing out that we will work always with randomly generated pattern’s entries as in standard Amit–Gutfreund–Sompolinsky theory (Amit, 1989): while this choice barely resembles real situations, however, at a more abstract level, by a standard Shannon–Fano compression argument it is immediate to realize that if the network is able to cope with these *P* entirely random patterns, it will be certainly able to handle (at least) the same amount of *structured* patterns (where correlations are expected and thus their compression would eventually save memory for further storage).

As a final remark, we stress that while numerical and heuristic explorations have been largely exploited in the Computer Science Literature in the past decade, here we aim to reach sharp statements, with rigorous and well controllable results. Indeed, along this perspective, important contributions already appeared in the Mathematical and Theoretical Physics Literature (see e.g. Barra et al., 2014; Bovier & Gayrard, 1997, 1998; Bovier, Gayrard, & Picco, 1994, 1995; Choromanska, Henaff, Mathieu, Ben Arous, & Le Cun, 2015; Guerra et al., 2012; Pastur, Shcherbina, & Tirozzi, 1994; Talagrand, 1998, 2000 and references therein). In these regards, in this first paper we force our focus to the low-storage analysis of the model, namely we study properties the model naturally possesses when the number of patterns (or features) to learn and retrieve grows sub-linearly with the number of neurons dedicated to the task: from the mathematical perspective, this regime is much more controllable as the *glassiness* hidden in the model becomes negligible (Amit, 1989; Coolen et al., 2005).

The paper is structured as follows:

We spend the next section (Section 2) for a streamlined introduction to machine learning and neural networks, in particular linking *learning* and *retrieval* in their simplest representation: we use restricted Boltzmann machines (RBM) as prototypes for learning³ and generalized Hopfield models (GHM) as paradigms for retrieval. We revise, by using a standard Bayes argument, how the features learnt during the training stage by a restricted Boltzmann machine play as patterns to be recognized in the future, highlighting how – *de facto* – pattern recognition happens via a standard Hebbian kernel typical of Hopfield retrieval: this is a recently understood bridge glimpsed in Disordered Systems Community (Agliari et al., 2012; Agliari et al., 2013b; Barra, Genovese, Sollich, & Tantari, 2017; Barra et al., 2012; Mezard, 2017; Tubiana & Monasson, 2017) that still deserves to be outlined (especially because, due to this bridge, unlearning in Hopfield network can be connected to pruning in Boltzmann machines). Finally we briefly comment on spurious states and their intrinsic genesis in the Hopfield framework.

Then we move on to build our mathematical framework, *i.e.* the mechanical analogy for neural networks: as a benchmark, in Section 3 we show the analogy at the classical level, namely we consider the standard statistical mechanical package related to the

² R.E.M. is the acronym of Random Eye Movement and *REM sleep* is a unique phase of sleep in mammals and birds, distinguishable by random and rapid movement of the eyes and the propensity of the sleeper to dream vividly.

³ We remark that deep learning structures can be built by properly hierarchically nested chains of RBMs one into another (Salakhutdinov & Hinton, 2009).

original pairwise Hopfield model (whose mechanical counterpart lies in classical mechanics) and we give a complete picture of its properties, re-obtaining all the well-known existing results. In Section 4, we extend the analogy to include higher order (P-spin) contributions to the Hopfield Hamiltonian (whose mechanical counterpart lies in special relativity) and we obtain an exhaustive picture of the resulting proprieties of this generalized model too. Section 5 is due to a numerical analysis of the capabilities of this extended model: through a one-to-one comparison among performances of the classical versus the relativistic Hopfield model, we prove how systematically our extension out-performs its classical limit and, in particular, we show how the spurious states of the standard Hopfield model are almost entirely pruned by its extension.

Finally Section 6 is left for our conclusions and a brief summary of future outlooks.

2. A teaspoon of neural networks from the statistical mechanics perspective: Boltzmann learning and Hopfield retrieval

For long time, *machine learning* (meant as statistical learning of characteristic features – or patterns – from input data (Coolen et al., 2005)) and *machine retrieval* (meant as recognition of previously stored patterns – or features – (Amit, 1989) have been treated separately, the former mainly addressed from a statistical inference perspective, the latter mainly tackled through statistical mechanics techniques. Hereafter instead, we show how Restricted Boltzmann machines learn by a stochastic dynamics (e.g. contrastive divergence (Salakhutdinov & Hinton, 2009)), whose long term relaxation converges to a Gibbs measure of an effective Hamiltonian that turns out to be sharply the Hopfield model for pattern recognition; then, in a streamlined way, we summarize the properties of the latter by means of standard statistical mechanical arguments: this is a preliminary step required to correctly frame the mechanical analogy in the broad world of neural networks.

In a nutshell, a Boltzmann machine is a two-party system (or a *bipartite spin-glass* in the statistical mechanical vocabulary (Agliari et al., 2017; Barra, Genovese, Guerra, & Tantari, 2011; Barra et al., 2017; Guerra et al., 2012)) whose parties (or *layers* to keep the original jargon (Goodfellow, Bengio, & Courville, 2017; Salakhutdinov & Hinton, 2009)) are called *visible* – the one that receives the input from the outside world – and *hidden* – the one dedicated to figure out correlations in the data presented to the visible layer (Ackley, Hinton, & Sejnowski, 1985; Le Roux & Bengio, 2008), (see Fig. 1, left panel).

Keeping Fig. 1 in mind, each layer is composed of spins/neurons (N for the visible and P for the hidden) and these spins can be chosen with high generality, ranging from discrete-valued, e.g., we can select Ising neurons for the visible layer $\sigma_i = \pm 1, i \in (1, \dots, N)$, to real-valued, e.g. we can choose Gaussian neurons for the hidden layer $z_\mu = \mathcal{N}[0, 1], \mu \in (1, \dots, P)$ (Barra et al., 2017; Hinton, 2010).

Analogously, the entries of the weight's matrix ξ_i^μ can be both real or discrete: generally speaking, while continuous weights allow the usage of stronger learning rules (i.e. contrastive divergence with simulated annealing (Kirkpatrick et al., 1983; Salakhutdinov & Hinton, 2009)) w.r.t. their binary counterparts (whose typical learning rule is the old fashioned Hebbian one (Amit, 1989)) and are thus more convenient during the learning stage, binary weights are more convenient in the retrieval stage, when pattern recognition is in order (Barra et al., 2017; Barra, Genovese, Sollich, & Tantari, 2018).

This trade-off gave rise to a broad plethora of variations on theme (see for instance Decelle, Fissore, & Furtlehner, 2017; Hinton, 2010; Huang, 2010; Huang & Toyozumi, 2015; Kappen & Ro-

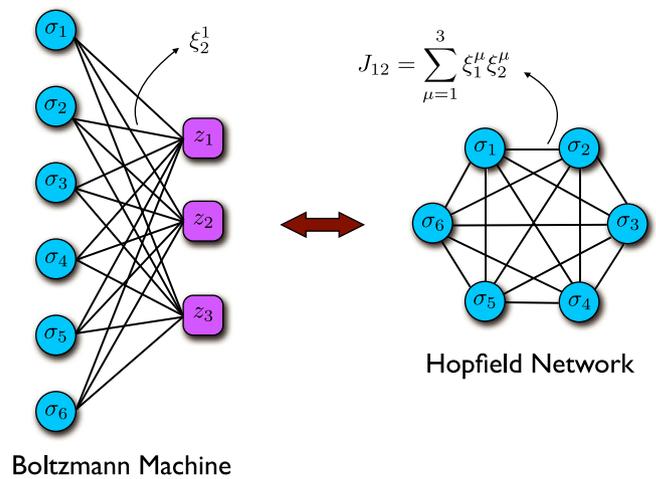


Fig. 1. Restricted Boltzmann machine and associative Hopfield network. Left panel: example of a restricted Boltzmann machine equipped with 6 visible neurons $\sigma_i, i \in (1, \dots, 6)$ in the visible (or input) layer and 3 hidden units $z_\mu, \mu \in (1, \dots, 3)$ in the hidden layer. The weights connecting them form the $N \times P$ matrix ξ_i^μ . Right panel: example of the corresponding Hopfield network, whose six visible neurons $\sigma_i, i \in (1, \dots, 6)$ retrieve as patterns stored in the Hebb matrix $J_{ij} = \sum_{\mu=1}^P \xi_i^\mu \xi_j^\mu$ the three vectors $\xi^\mu, \mu \in (1, \dots, 3)$, each pertaining to a *feature*, i.e. one of the three z_μ hidden variables of the (corresponding) RBM.

driguez, 1998; Salakhutdinov & Hinton, 2009; Tramel, Dremeau, & Krzakala, 2016 and references therein), whose best setting (for our purposes) lies in a Boltzmann machine equipped with a Gaussian hidden layer and a Boolean visible layer (Barra et al., 2017): the structure of this typical restricted Boltzmann machine can thus be coded into the following effective cost-function

$$H_N(\sigma, \mathbf{z}|\xi) = -\frac{1}{\sqrt{N}} \sum_{i,\mu} \xi_i^\mu \sigma_i z_\mu - \sum_i h_i \sigma_i, \quad (2.1)$$

where the N external fields h_i act as *bias terms* from a statistical inference perspective: from the statistical mechanics viewpoint, these *one body* interactions are always mathematically trivial and, while playing a key role in machine learning, they will be discarded soon (with no loss of generality as they can always be reintroduced at any time later).

Crucially, the dynamics of the weights (i.e. the slow variables, also called *synapses* in standard neural network jargon (Amit, 1989; Coolen et al., 2005)) and the dynamics of the spins (i.e. the fast variables, also called *neurons* in standard neural network jargon (Amit, 1989; Coolen et al., 2005)) evolve separately, adiabatically on different timescales (or more properly *epochs* (Le Cun et al., 2015)), such that we can average away the fast neural scale when studying the slow synaptic dynamics (learning) and, likewise, we can keep synapses quenched when interested in the fast neural dynamics (retrieval).

Typically, a set of data vectors is presented to the visible layer of the machine (i.e. the so called *training set*) and, under the assumption that these data have been independently generated by the same probability distribution $Q(\sigma)$, the ultimate goal of the machine is to make an inner representation of $Q(\sigma)$, say $P(\sigma)$, that is as close as possible to the original one.

The way in which $P(\sigma)$ approximates $Q(\sigma)$ is usually achieved by the minimization of the Kullback–Leibler cross entropy $D(Q, P)$, defined as

$$D(Q, P) \equiv \sum_{\sigma} Q(\sigma) \ln \left(\frac{Q(\sigma)}{P(\sigma)} \right), \quad (2.2)$$

as, once a small parameter is introduced ϵ (i.e. the *learning rate*), by imposing $\Delta\xi_i^\mu = -\epsilon(\partial D(Q, P)/(\partial\xi_i^\mu))$ we get

$$\Delta D(Q, P) = \sum_{i,\mu} \frac{\partial D(Q, P)}{\partial \xi_i^\mu} \Delta \xi_i^\mu = -\epsilon \left(\sum_{i,\mu} \left(\frac{\partial D(Q, P)}{\partial \xi_i^\mu} \right)^2 \right) \leq 0, \quad (2.3)$$

and analogously for the response of the Kullback–Leibler cross entropy to a variation in the biases h_i . This is a *secure learning rule* thanks to the definite sign in the last term above.

Weights in the Boltzmann machine are symmetric and this suffices for Detailed Balance to hold (Amit, 1989; Coolen et al., 2005): the latter guarantees that the long term limit of any (non-pathological) stochastic dynamics will always end up in a Gibbs measure, hence the mathematical structure of the probability distribution P is known and this through Eq. (2.3) allows generating explicit algorithms, among which the following *contrastive divergence* criterion is probably the most applied (Coolen et al., 2005):

$$\frac{\partial D(Q, P)}{\partial \xi_i^\mu} = \epsilon (\langle \sigma_i \sigma_j \rangle_{\text{clamped}} - \langle \sigma_i \sigma_j \rangle_{\text{free}}). \quad (2.4)$$

In this equation, the subscript *clamped* means that the averages $\langle \cdot \rangle$ must be evaluated when the visible layer is forced to assume data values, while *free* means that the averages are the standard statistical mechanical ones: this recipe is a very elementary yet powerful learning rule as, roughly speaking, it simply tries to make the theoretical correlation functions as close as possible to the empirical ones.⁴

When the machine is able to reproduce the statistics stored in the training data correctly, the internal weights have been rearranged such that if the machine is now asked to generate vectors according to $P(\sigma)$, their statistical properties will coincide with those of the input data: the machine *has learnt* a representation of the reality it has been fed with (at least at the lowest orders it has been constrained to infer, e.g. averages and variances if we deal with pairwise cost-functions).

Now we question about the retrieval capabilities of such a machine, that is, what kind of features will this machine discover when provided with further data. Since hidden units are independent in RBMs, it is $P(\mathbf{z}|\sigma) = \prod_{\mu=1}^P P(z_\mu|\sigma)$, and, analogously, $P(\sigma|\mathbf{z}) = \prod_{i=1}^N P(\sigma_i|\mathbf{z})$, hence to get their joint and marginal distributions we can use Bayes Theorem in a very elementary way, i.e. $P(\sigma, \mathbf{z}) = P(\mathbf{z}|\sigma)P(\sigma) = P(\sigma|\mathbf{z})P(\mathbf{z})$, thus

$$P(\sigma, \mathbf{z}) = \frac{\exp(-\frac{1}{2} \sum_{\mu} z_{\mu}^2 + \sum_{i,\mu} \xi_i^{\mu} \sigma_i z_{\mu} + \sum_i h_i \sigma_i)}{Z}, \quad (2.5)$$

$$\begin{aligned} P(\sigma) &= \int \prod_{\mu=1}^P dz_{\mu} P(\sigma, \mathbf{z}) \\ &= \frac{\int \prod_{\mu=1}^P dz_{\mu} \exp(-\frac{1}{2} \sum_{\mu} z_{\mu}^2 + \sum_{i,\mu} \xi_i^{\mu} \sigma_i z_{\mu} + \sum_i h_i \sigma_i)}{Z} \\ &= \int \prod_{\mu=1}^P dz_{\mu} P(\mathbf{z}|\sigma) P(\sigma) \\ &= Z^{-1} \exp\left(\frac{1}{2} \sum_{ij} \left(\sum_{\mu} \xi_i^{\mu} \xi_j^{\mu} \right) \sigma_i \sigma_j + \sum_i h_i \sigma_i\right), \quad (2.6) \end{aligned}$$

⁴ This argument can be expanded up to arbitrarily N -points correlation functions by paying the price of adding extra hidden layer and we believe this way of reasoning to lie at the core of Deep Learning when inspected via disordered statistical mechanics (Krotov & Hopfield, 2016; Le Cun et al., 2015).

$$Z = \sum_{\{\sigma\}} \int \prod_{\mu=1}^P dz_{\mu} \exp\left(-\frac{1}{2} \sum_{\mu} z_{\mu}^2 + \sum_{i,\mu} \xi_i^{\mu} \sigma_i z_{\mu} + \sum_i h_i \sigma_i\right) \quad (2.7)$$

namely, carrying out the integrals over the z_{μ} 's to move from (2.5) to (2.6), in the solely visible layer we are left with the Gibbs measure of the following Hopfield Hamiltonian

$$H_N(\sigma|\xi, \mathbf{h}) = -\frac{1}{2N} \sum_{i,j=1}^N \left(\sum_{\mu=1}^P \xi_i^{\mu} \xi_j^{\mu} \right) \sigma_i \sigma_j - \sum_{i=1}^N h_i \sigma_i. \quad (2.8)$$

Note that in Eq. (2.8) the one-body interactions (in statistical mechanics dictionary) show how the biases in the RBM automatically translate into the *thresholds for firing* in the associative neural networks jargon, where the Hopfield model is mostly used, and, from a Bayesian perspective, they account for the prior distribution $P(\sigma)$, while $P(\mathbf{z}|\sigma)$ is handled by the interactions among the two layers: from now on we set $h_i \equiv 0$ for all the N neurons with no loss of generality.

To understand the retrieval capabilities of the above cost function from a statistical mechanical perspective, it is useful to introduce the P Mattis magnetizations, defined as

$$m_{\mu} \equiv \frac{1}{N} \sum_i \xi_i^{\mu} \sigma_i, \quad \mu \in (1, \dots, P). \quad (2.9)$$

They cover the pivotal role of *order parameters* in Hopfield theory as they capture the relaxation of the network. This can happen in essentially three different ways:

(a) network's dynamics ends up in a pure state, namely in an attractor that coincides exactly with one of the P stored patterns (such that its corresponding Mattis magnetization becomes one as the neurons are all aligned with the pattern entries in a pure state).

(b) network's dynamics ends up in a spurious state, namely in an unwanted mixture of pure states (whose amount, spontaneously, grows very quickly as P raises (Amit, 1989)) such that its corresponding Mattis magnetization becomes strictly different from zero, albeit sensibly smaller than the one.

(c) network's dynamics has a failure and ends up in a random (ergodic) sample of the phase space, where the Mattis parameters are all zero (the network relaxed to a fully noisy state).

As we can write the Hopfield Hamiltonian via the Mattis order parameters as

$$H_N(\sigma|\xi) \sim -\frac{1}{2N} \sum_{i,j=1}^N \left(\sum_{\mu=1}^P \xi_i^{\mu} \xi_j^{\mu} \right) \sigma_i \sigma_j = -N \sum_{\mu} m_{\mu}^2, \quad (2.10)$$

by a minimum energy request (i.e. minimizing the cost-function (2.10)) it follows that the candidate minima for the Hopfield model are those states σ that force one of the P Mattis magnetizations to assume value 1 or a combination of Mattis magnetization to be different from zero: remarkably, some equilibrium states, the global minima, do coincide with the P patterns (that can thus be retrieved) and we are then tempted to associate the spontaneous relaxation of the network into the attractor to the cognitive process of information retrieval.

To quantify this statement, the standard route (Amit, 1989; Coolen et al., 2005) is to introduce and study the free energy related to the model, that is, to express and extremize it in terms of the Mattis order parameters, as briefly summarized hereafter for the low-storage case, namely when $\alpha \equiv \lim_{N \rightarrow \infty} P/N = 0$.

The free energy of the Hopfield model, $f(\beta)$ in the low storage regime, reads as

$$\begin{aligned} f(\beta) &= \lim_{N \rightarrow \infty} \frac{1}{N} \ln \sum_{\{\sigma\}} e^{\frac{\beta}{2N} \sum_{ij} (\sum_{\mu=1}^P \xi_i^\mu \xi_j^\mu) \sigma_i \sigma_j} \\ &= \lim_{N \rightarrow \infty} \frac{1}{N} \ln \sum_{\{\sigma\}} e^{\frac{\beta}{2N} \sum_{\mu} (\sum_i \xi_i^\mu \sigma_i)^2} \end{aligned} \quad (2.11)$$

$$= \lim_{N \rightarrow \infty} \frac{1}{N} \ln \sum_{\{\sigma\}} \int \prod_{\mu} \frac{dm^\mu}{\sqrt{2\pi}} e^{-N \left(\frac{\beta}{2} \sum_{\mu} m_\mu^2 + \sum_{\mu} \sum_i m^\mu \xi_i^\mu \sigma_i \right)}, \quad (2.12)$$

where the first line is the bare definition of the free energy, while in the second line we used the Gaussian integral representation: the exponent in (2.12) is extensive due to the N factor multiplying all $O(1)$ terms, thus it can be evaluated with a saddle point argument, whose extremization returns

$$f(\beta) = \ln 2 + \langle \log \cosh(\beta \sum_{\mu} m^\mu \xi_{\mu}) \rangle_{\xi} - \frac{\beta}{2} \langle \sum_{\mu} m_{\mu}^2 \rangle_{\xi}, \quad (2.13)$$

$$\frac{df(\beta)}{d\langle m_{\mu} \rangle} = 0 \Rightarrow \langle m_{\mu} \rangle = \langle \xi_{\mu} \tanh(\beta \sum_{\nu} m^{\nu} \xi_{\nu}) \rangle_{\xi}, \quad \mu \in (1, \dots, P). \quad (2.14)$$

Eq. (2.13) is the explicit expression of the free energy of the model in terms of the Mattis magnetizations: the free energy is the entropy minus the cost function (at a given noise level $\beta = 1/T$, where T plays as the *fast – or thermal – noise* (Ackley et al., 1985; Amit, 1989)), hence extremizing the free energy we obtain the maximum entropy solutions of the cost-function minimization. These are coded in the last line, Eq. (2.14), that is a set of P *self-consistent* equations for the P order parameters m_{μ} , quantifying the strength of the retrieval.

Whatever the route, that is starting from the Hebbian prescription (Hebb, 1949) as in the original Hopfield paper (Hopfield, 1982) or upon marginalization over the hidden layer in Boltzmann machine (hence after learning via contrastive divergence) (Barra et al., 2012), unfortunately, we always end up in a network whose attractors are by far more than the stored patterns (namely more than the solely *pure states* we would see retrieved by the network) (Amit, 1989), the excess stock of (local) minima being constituted by the so-called *spurious states*: the simplest example is a 3-pattern mixture defined as follows

$$\sigma_i = \text{sign}(\xi_i^1 + \xi_i^2 + \xi_i^3). \quad (2.15)$$

The Mattis overlap of this state with any of the three patterns is – for large networks – $m^{\nu} = N^{-1} \sum_i \xi_i^{\nu} \sigma_i = 0.5$ (for $\nu = (1, 2, 3)$) hence, while smaller in amplitude than the Mattis overlap of a pure state (whose amplitude is one), it is still a meta-stable state: if orbiting in the surrounding, the network can be attracted by such spurious states and converge to them rather than to the pure ones. It is worth noticing that these metastable states are more *fragile* than their pure counterparts as, against the fast (thermal) noise, they are stable solely in the range $\beta \in (2, \infty)$ (i.e. $T = (0, 0.5)$), while pure states least up to $\beta \in (1, \infty)$ (i.e. $T = (0, 1)$) (Amit, 1989). Unfortunately, however, as the patterns are added linearly to the memory kernel, there is an exponential (combinatorial) proliferation of these unwanted meta-stable states in the retrieval landscape of the Hopfield network such that Hopfield himself suggested a procedure to prune – or remove – (a part of) them from his coupling matrix (Hopfield et al., 1983). In a nutshell, Hopfield's idea is again transparent, elegant and brilliant: as there are sensibly much more spurious states (i.e. metastable minima)

than pure states (global minima), let us start the system at random and make a quench (e.g. a search for minima with steepest descent rather than conjugate gradient or stochastic algorithms). With high probability, in this way, the system will end up in a spurious state: we can collect this equilibrium configuration – called $\langle \sigma_i \sigma_j \rangle_{\text{spurious}}$ hereafter – and subtract it from the memory kernel, via

$$J_{ij} = \sum_{\mu=1}^P \xi_i^{\mu} \xi_j^{\mu} - \langle \sigma_i \sigma_j \rangle_{\text{spurious}}.$$

We can do this iteratively and check that effectively the network becomes progressively cleared from those nasty attractors: this procedure is called *unlearning* (Dotsenko et al., 1991; Kleinfeld & Pendergraft, 1987; Nokura, 1996, 1998) and it has been linked to REM sleep (Crick & Mitchinson, 1983) (offering a possible intriguing picture for its interpretation) due to the effectiveness of the random starting point setting for the quenching procedure in consolidating memories (phenomenon to be eventually correlated with the rapid eye movements in that part of sleep).

3. A mechanical formulation of the (classical) Hopfield networks

Aim of this section is to recover formulae (2.13) and (2.14) without relying any longer on the standard statistical mechanical guidance, that is, we will no longer use either the Maximum Entropy or the Minimum Energy principles (i.e. overall the standard free energy extremization). Indeed, by observing that the free energy plays as an action in a suitable mechanical analogy, we can import an arsenal of mathematical tools to investigate its properties (originally developed within the framework of analytical mechanics). In particular, we will show that the free energy – as any proper action – obeys a Hamilton–Jacobi PDE, whose solution returns Eq. (2.13): the variational principle for the free energy minimization will be expressed, in this context, by the Least Action Principle (and will return exactly Eq. (2.14)).

3.1. Preliminary definitions regarding the (classical) Hopfield free energy

In this section we introduce more formally all the quantities and observables we will deal with, starting by the first

Definition 1. The Hopfield model is described by the Hamiltonian

$$H_N(\sigma | \xi) = - \sum_{i < j}^N J_{ij} \sigma_i \sigma_j, \quad (3.1)$$

where $\sigma_i = \pm 1$, $i \in (1, \dots, N)$ are Ising spins (or McCulloch&Pitts neurons), while the coupling matrix J is defined in terms of P patterns ξ^{μ} , $\mu \in (1, \dots, P)$ as follows

$$J_{ij} = \frac{1}{N} \sum_{\mu=1}^P \xi_i^{\mu} \xi_j^{\mu} = \frac{1}{N} \xi_i \cdot \xi_j, \quad (3.2)$$

$$\frac{1}{2} = P(\xi_i^{\mu} = +1) = P(\xi_i^{\mu} = -1). \quad (3.3)$$

Using the label $\beta \in \mathbb{R}^+$ to tune the (inverse) level of noise in the network, the partition function of the model is then introduced as

$$Z_N(\beta) = \sum_{\{\sigma\}} \exp(-\beta H_N(\sigma, \xi)) \sim \sum_{\{\sigma\}} \exp\left(\frac{\beta}{2} \sum_{i,j=1}^N J_{ij} \sigma_i \sigma_j\right), \quad (3.4)$$

where \sim means an equality up to a constant (due to self-interactions $i = j$) that is negligible in the thermodynamic limit $N \rightarrow \infty$, as long as we keep the load of the system sub-linear in the volume (i.e. $P \neq \alpha N$, for some positive constant α), i.e., in the low storage scenario under study.

Definition 2. The (intensive) free-energy $f(\beta)$ is defined⁵ as

$$f(\beta) = \lim_{N \rightarrow \infty} f_N(\beta), \quad f_N(\beta) = \frac{1}{N} \log Z_N(\beta), \quad (3.5)$$

Considering a generic function of the weights and the neurons $\Phi(\sigma, \xi)$, we introduce also the average over the patterns as

$$\langle \Phi(\sigma, \xi) \rangle_\xi = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \Phi(\sigma, \xi_i), \quad (3.6)$$

(where $\xi_i = (\xi_i^1, \dots, \xi_i^P)$), and the standard Boltzmann averages, i.e.

$$\langle \Phi(\sigma, \xi) \rangle = \frac{\sum_{\{\sigma\}} \Phi(\sigma, \xi) \exp(-\beta H_N(\sigma|\xi))}{Z_N(\beta)}. \quad (3.7)$$

Now we start to generalize the former observables in order to make them *flexible enough* for our mechanical construction. To this task we introduce P spatial variables $x_\mu \in \mathbb{R}$, $\mu \in (1, \dots, P)$ and 1 temporal variable $t \in \mathbb{R}^+$ that we use to generalize the partition function (3.4) as follows:

$$Z_N(t, \mathbf{x}) = \sum_{\{\sigma\}} \exp \left(-\frac{t}{2N} \sum_{i,j=1}^N \xi_i \cdot \xi_j \sigma_i \sigma_j + \sum_{i=1}^N \mathbf{x} \cdot \xi_i \sigma_i \right). \quad (3.8)$$

We call $\langle \cdot \rangle_{t, \mathbf{x}}$ the expectation value of any quantity of the network with this new partition function.

Note that, crucially, taking $t = -\beta$ and $\mathbf{x} = \mathbf{0}$, this generalized partition function (Eq. (3.8)) reduces to the standard one of statistical mechanics (Eq. (3.4)).

The next observable will play as the generalized free energy of the Hopfield model from a statistical mechanical perspective, while covers the role of an action in the present mechanical analogy.

Definition 3. By using the partition function (3.8), we introduce the interpolating free energy $f_N(t, \mathbf{x})$ as

$$\begin{aligned} f_N(t, \mathbf{x}) &= \frac{1}{N} \ln Z_N(t, \mathbf{x}) \\ &= \frac{1}{N} \ln \sum_{\{\sigma\}} \exp \left(-\frac{tN}{2} \mathbf{m}_N(\sigma)^2 + N \mathbf{x} \cdot \mathbf{m}_N(\sigma) \right), \end{aligned} \quad (3.9)$$

where

$$\mathbf{m}_N(\sigma) = \frac{1}{N} \sum_{i=1}^N \xi_i \sigma_i, \quad (3.10)$$

is the vector whose components are the Mattis magnetizations m_μ , $\mu \in (1, \dots, P)$.

Finally, it is elementary to check that the space–time derivatives of the interpolating free energy (3.9) read as

$$\begin{aligned} \frac{\partial f_N(t, \mathbf{x})}{\partial t} &= -\frac{1}{2} \langle \mathbf{m}_N^2 \rangle_{t, \mathbf{x}}, \\ \frac{\partial f_N(t, \mathbf{x})}{\partial x^\mu} &= \langle m_N^\mu \rangle_{t, \mathbf{x}}. \end{aligned} \quad (3.11)$$

This is the starting point for the mechanical analogy.

3.2. Construction of the (classical) Hopfield free energy via Hamilton–Jacobi PDE

The mechanical analogy we pursue lies in the Hamilton–Jacobi framework as we show in this section.

⁵ Note that we use here and everywhere the subscript N when we work at fixed volume N , its lacking expressing quantities already evaluated in the thermodynamic limit.

Proposition 1. By construction, $f_N(t, \mathbf{x})$ obeys the following (classical) Hamilton–Jacobi PDE and it plays de facto as an action in this mechanical analogy

$$\frac{\partial f_N}{\partial t} + \frac{1}{2} \left(\frac{\partial f_N}{\partial x^\mu} \right)^2 + V_N(t, \mathbf{x}) = 0, \quad (3.12)$$

$$\begin{aligned} V_N(t, \mathbf{x}) &= \frac{1}{2} \left(\langle \mathbf{m}_N^2 \rangle_{t, \mathbf{x}} - \langle \mathbf{m}_N \rangle_{t, \mathbf{x}}^2 \right) \\ &= \frac{1}{2} \sum_{\mu=1}^P \left(\langle m_{N, \mu}^2 \rangle_{t, \mathbf{x}} - \langle m_{N, \mu} \rangle_{t, \mathbf{x}}^2 \right). \end{aligned} \quad (3.13)$$

This partial differential equation describes – even at finite volume N – the motion of a classical (non relativistic) particle, with unitary mass⁶ in $P + 1$ dimensions.

Remark 1. In the thermodynamic limit, away from critical point, the Mattis magnetizations self-average, i.e.

$$\lim_{N \rightarrow \infty} \left(\langle \mathbf{m}_N^2 \rangle_{t, \mathbf{x}} - \langle \mathbf{m}_N \rangle_{t, \mathbf{x}}^2 \right) = 0, \quad (3.14)$$

such that $\lim_{N \rightarrow \infty} V_N(t, \mathbf{x}) = V(t, \mathbf{x}) = 0$.

Remark 2. In the thermodynamic limit, the motion has space–time symmetries whose Noëther currents, derived respectively from the momentum conservation and from the energy conservation as

$$\lim_{N \rightarrow \infty} \left(\langle \mathbf{m}_N^2 \rangle_{t, \mathbf{x}} - \langle \mathbf{m}_N \rangle_{t, \mathbf{x}}^2 \right) = 0 \quad \text{Momentum Conservation} \quad (3.15)$$

$$\lim_{N \rightarrow \infty} \left(\langle \mathbf{m}_N^4 \rangle_{t, \mathbf{x}} - \langle \mathbf{m}_N^2 \rangle_{t, \mathbf{x}}^2 \right) = 0 \quad \text{Energy Conservation}, \quad (3.16)$$

mirror the classical self-averaging properties in the statistical mechanical jargon.

As a consequence, in the thermodynamic limit, the particle’s motion of the mechanical analogy paints a Galilean trajectory, i.e., a straight line $\mathbf{x} = \mathbf{x}_0 + \langle \mathbf{m} \rangle_{t, \mathbf{x}} \cdot (t - t_0)$.

Further, in this limit, the determination of an explicit expression of the free energy in terms of the Mattis magnetizations reduces to the explicit calculation of the action of this free motion. As Cauchy conditions we choose $t_0 = 0$, \mathbf{x}_0 (note that the choice $t_0 = 0$ kills the coupling between the spins) such that

$$f(t, \mathbf{x}) = f(0, \mathbf{x}_0) + \int_0^t dt' \mathcal{L}(t'), \quad (3.17)$$

where $\mathcal{L} = \frac{1}{2} \langle \mathbf{m} \rangle_{t, \mathbf{x}}^2$ is the Lagrangian. The latter is a constant of motion as $V(t, \mathbf{x}) = 0$, hence the solely calculations required are due to evaluate the (trivial) Cauchy condition

$$f(0, \mathbf{x}_0) = \log 2 + \langle \log \cosh \mathbf{x}_0 \cdot \xi \rangle_\xi. \quad (3.18)$$

Merging the two observations and writing $\mathbf{x}_0 = \mathbf{x}(t) - \langle \mathbf{m} \rangle_{t, \mathbf{x}} t$, we have

Theorem 1. The infinite volume limit of the Hopfield action (3.9) reads as

$$f(t, \mathbf{x}) = \log 2 + \langle \log \cosh(\mathbf{x} - t \langle \mathbf{m} \rangle_{t, \mathbf{x}}) \cdot \xi \rangle_\xi + \frac{t}{2} \langle \mathbf{m} \rangle_{t, \mathbf{x}}^2. \quad (3.19)$$

Remark 3. Note that by setting $t = -\beta$ and $\mathbf{x} = \mathbf{0}$, the (classical) Hopfield free energy is obtained (check Eq. (2.13)).

Corollary 1. Within the mechanical analogy, the expression (3.19) is subjected to the Last Action Principle $\delta f(t, \mathbf{x}) = 0$.

⁶ Note that, according to Eq. (3.12) the classical momentum is $\langle \mathbf{m} \rangle_{t, \mathbf{x}}$.

Indeed we can perform an infinitesimal variation $\langle m^\mu \rangle_{t,\mathbf{x}} \rightarrow \langle m^\mu \rangle_{t,\mathbf{x}} + \delta \langle m^\mu \rangle_{t,\mathbf{x}}$ and verify that

$$\begin{aligned} \delta f(t, \mathbf{x}) &= \frac{\partial f(t, \mathbf{x})}{\partial \langle m^\mu \rangle_{t,\mathbf{x}}} \delta \langle m^\mu \rangle_{t,\mathbf{x}} \\ &= \tanh(x - \langle m^\mu \rangle_{t,\mathbf{x}} t) (-t \delta \langle m^\mu \rangle_{t,\mathbf{x}}) + \langle m^\mu \rangle_{t,\mathbf{x}} \delta \langle m^\mu \rangle_{t,\mathbf{x}} t \\ &= 0, \end{aligned} \quad (3.20)$$

automatically implies the self consistency equations (2.14) to hold.

4. A mechanical formulation of the (relativistic) Hopfield networks

The mechanical analogy can now be used to note pathologies in previous treatment that somehow shine in its transparent calculations and will led us to a very natural extension of the Hopfield cost-function: actually there are two main observations waiting to be done. The first is that, as the free energy plays as an action, we can think at the exponent in the Maxwell–Boltzmann weight (e.g. as expressed in the partition function (3.8)) as the product of the $P + 1$ momentum–energy tensor with the $P + 1$ metric tensor (i.e. $-t \cdot E + \mathbf{x} \cdot \mathbf{m}$), and observe that the underlying metrics is not Euclidean, rather Minkowskian, as in special relativity. The second is that in classical mechanics the velocity is unbounded, while here the magnetization is obviously bounded by one, i.e. $\langle m^\mu \rangle \leq c \equiv 1$: the whole points straight to a natural *relativistic generalization*, that, indeed, is the content of this Section.

4.1. Preliminary definitions regarding the (relativistic) Hopfield free energy

Because in the mechanical analogy the Hopfield cost-function (3.1) reads as the kinetic energy associated to the fictitious particle, a natural extension of Hopfield model is constituted by its relativistic deformation, i.e.

$$-\frac{\mathbf{m}_N^2}{2} \rightarrow -\sqrt{1 + \mathbf{m}_N^2}. \quad (4.1)$$

This calls for the next

Definition 4. The Hamiltonian for the relativistic Hopfield model is defined as

$$H_N(\sigma|\xi) = -N\sqrt{1 + \mathbf{m}_N(\sigma)^2}. \quad (4.2)$$

Two observations are in order here: at first note that Taylor-expanding the Hamiltonian (4.1) in the Mattis magnetizations we collect a series of many-body (P-spin) contributions (hence in the direction suggested by Hopfield recently in regard of Deep Learning (Krotov & Hopfield, 2016)), as

$$\begin{aligned} -\frac{H_N(\sigma|\xi)}{N} &= 1 + \frac{1}{2N^2} \sum_{ij} (\xi_i \cdot \xi_j) \sigma_i \sigma_j \\ &\quad - \frac{1}{8N^4} \sum_{ijkl} (\xi_i \cdot \xi_j)(\xi_k \cdot \xi_l) \sigma_i \sigma_j \sigma_k \sigma_l + \dots \end{aligned} \quad (4.3)$$

further, the r.h.s. of Eq. (4.3) is an alternate-sign series, hence it will have both contributions in learning (those with the minus sign) and in unlearning (those with the plus sign Dotsenko et al., 1991; Nokura, 1996, 1998), with the prevailing contribution played by standard Hebbian learning (as the series converges). We will show in Section 5 that the beneficial role of unlearning lies in destabilizing the retrieval of spurious memories, thus resulting in enhanced network’s performances w.r.t. the classical limit (i.e. the standard Hopfield model).

Once this new model is introduced, we extend partition function, free energy and mechanical analogy properly from the previous sections, e.g.

Definition 5. The relativistic partition function and free energy, related to Hamiltonian (4.2), read as

$$Z_N(\beta) = \sum_{\{\sigma\}} \exp\left(\beta N \sqrt{1 + \mathbf{m}_N(\sigma)^2}\right), \quad (4.4)$$

$$f(\beta) = \lim_{N \rightarrow \infty} f_N(\beta), \quad f_N(\beta) = \frac{1}{N} \ln Z_N(\beta). \quad (4.5)$$

As we did in the previous section, we now read the relativistic partition function from a purely mechanical perspective.

Definition 6. The generalized partition function and the action of the relativistic Hopfield model, suitable for the mechanical analogy, read as

$$Z_N(t, \mathbf{x}) = \sum_{\{\sigma\}} \exp\left(N(-t\sqrt{1 + \mathbf{m}_N(\sigma)^2} + \mathbf{x} \cdot \mathbf{m}_N(\sigma))\right), \quad (4.6)$$

$$f_N(t, \mathbf{x}) = \frac{1}{N} \ln Z_N(t, \mathbf{x}). \quad (4.7)$$

Note the minus sign in front of the *spatial part* of the partition function: this is the Minkowskian signature $(+, -, \dots, -)$, useful to represent the exponent in a Lorentzian form as, using the Einstein notation, $-N x_A p^A = -N \eta_{AB} x^A p^B$ with $x^A = (t, \mathbf{x})$, $p^A(\sigma) = (\sqrt{1 + \mathbf{m}_N(\sigma)^2}, \mathbf{m}_N(\sigma))$ and $\eta_{AB} = \text{diag}(+, -, \dots, -)$ is the standard flat metric of Minkowski space-time, $A, B = 0, 1, \dots, p$.

The expectation values $\langle \dots \rangle_\xi$ and $\langle \dots \rangle$ are obviously defined in the same way as (3.6) and (3.7), but with respect to this new partition function (4.6).

Finally, when setting $t = -\beta$ and $\mathbf{x} = \mathbf{0}$ we obtain the *relativistic statistical mechanical framework* we are interested in.

Again, we start our relativistic generalization of the mechanical analogy by computing the space-time derivatives of the free energy. The relevant ones are

$$\begin{aligned} \frac{\partial f_N(t, \mathbf{x})}{\partial t} &= -\langle \sqrt{1 + \mathbf{m}_N^2} \rangle_{t,\mathbf{x}}, \\ \frac{\partial f_N(t, \mathbf{x})}{\partial x^\mu} &= \langle m_N^\mu \rangle_{t,\mathbf{x}}, \end{aligned} \quad (4.8)$$

$$\frac{\partial^2 f_N(t, \mathbf{x})}{\partial t^2} = N(\langle 1 + \mathbf{m}_N^2 \rangle_{t,\mathbf{x}} - \langle \sqrt{1 + \mathbf{m}_N^2} \rangle_{t,\mathbf{x}}^2),$$

$$\nabla^2 f_N(t, \mathbf{x}) = N(\langle \mathbf{m}_N^2 \rangle_{t,\mathbf{x}} - \langle \mathbf{m}_N \rangle_{t,\mathbf{x}}^2).$$

4.2. Construction of the (relativistic) Hopfield free energy via Hamilton–Jacobi PDE

Proposition 2. By construction $f_N(t, \mathbf{x})$ obeys the following (relativistic) Hamilton–Jacobi PDE

$$\partial_t^2 f_N - \nabla^2 f_N = N(1 - (\partial_t f_N)^2 + (\nabla f_N)^2), \quad (4.9)$$

or in the manifestly covariant form

$$(\partial_A f_N)^2 + \frac{1}{N} \square f_N = 1, \quad (4.10)$$

where \square represents the D’Alambertian, i.e., $\square = \partial_A \partial^A$ (still in Einstein notation).

Remark 4. Requiring that the derivatives of the action are regular functions in x^μ , t , in the thermodynamic limit, a.s.⁷

$$(\partial_A f)^2 = 1. \quad (4.11)$$

⁷ Almost surely because when ergodicity breaks down a gradient’s catastrophe prevents regularity even in the infinite volume limit (Barra et al., 2014).

From the mechanical perspective, in the thermodynamic limit, the $P + 1$ -momentum of the particle reads as

$$p^A = -\frac{\partial f}{\partial x_A} = (\langle \sqrt{1 + \mathbf{m}^2} \rangle_{t,\mathbf{x}}, \langle \mathbf{m} \rangle_{t,\mathbf{x}}). \quad (4.12)$$

In terms of this momentum, the equation for the action (4.11) is, using the relativistic mechanics language, the on-shell relation (Mandl & Shaw, 2010) for a free particle with unitary mass.

Note that the fictitious particle of this mechanical analogy moves on the straight lines $\mathbf{x} = \mathbf{x}_0 + \mathbf{v}(t - t_0)$ for arbitrary (t_0, \mathbf{x}_0) , where the velocity \mathbf{v} is related to the spatial part of the momentum as $\langle \mathbf{m} \rangle_{t,\mathbf{x}} = \gamma \mathbf{v}$ and γ is the usual Lorentz factor. In other words:

$$\mathbf{v} = \frac{\langle \mathbf{m} \rangle_{t,\mathbf{x}}}{\sqrt{1 + \langle \mathbf{m} \rangle_{t,\mathbf{x}}^2}} \Leftrightarrow \langle \mathbf{m} \rangle_{t,\mathbf{x}} = \frac{\mathbf{v}}{\sqrt{1 - \mathbf{v}^2}}. \quad (4.13)$$

The Lorentz factor can be therefore written as $\gamma = \sqrt{1 + \langle \mathbf{m} \rangle_{t,\mathbf{x}}^2}$. Summing all these observations together, we end up in the determination of an explicit expression for the relativistic free energy in terms of the Mattis magnetizations as it reduces to the calculation of the action of this free motion. As Cauchy conditions we still keep $t_0 = 0, \mathbf{x}_0$ (where, again, $t_0 = 0$ kills the coupling between the spins), such that

$$\begin{aligned} f(t, \mathbf{x}) &= f(0, \mathbf{x}_0) + \int_0^t dt' \mathcal{L}(t') = \\ &= f(0, \mathbf{x}_0) - \frac{t}{\gamma} = f(0, \mathbf{x}_0) - \frac{t}{\sqrt{1 + \langle \mathbf{m} \rangle_{t,\mathbf{x}}^2}}, \end{aligned} \quad (4.14)$$

since the Lagrangian $\mathcal{L} = -\gamma^{-1}$ is constant on the classical trajectories. Thus, again we have

$$f(0, \mathbf{x}_0) = \log 2 + \langle \log \cosh \mathbf{x}_0 \cdot \boldsymbol{\xi} \rangle_{\xi}. \quad (4.15)$$

Noting that $\mathbf{x}_0 = \mathbf{x} - \mathbf{v}t$ with (4.13), and setting $t = -\beta, \mathbf{x} = \mathbf{0}$ (in order to reobtain the statistical mechanical framework), we can write the next theorem.

Theorem 2. *The free energy density of the relativistic Hopfield network in the thermodynamic limit reads as*

$$f(\beta) = \log 2 + \langle \log \cosh \left(\beta \boldsymbol{\xi} \cdot \frac{\langle \mathbf{m} \rangle}{\sqrt{1 + \langle \mathbf{m} \rangle^2}} \right) \rangle_{\xi} + \frac{\beta}{\sqrt{1 + \langle \mathbf{m} \rangle^2}}. \quad (4.16)$$

In the usual statistical mechanical settings, now by imposing thermodynamic prescriptions (namely requiring the free energy to be extremal with respect to the order parameters), we would get the following self-consistence conditions

$$\frac{df(\beta)}{d\langle \mathbf{m}^\mu \rangle} = 0 \Rightarrow \langle \mathbf{m}^\mu \rangle = \left\langle \xi^\mu \tanh \left(\beta \boldsymbol{\xi} \cdot \frac{\langle \mathbf{m} \rangle}{\sqrt{1 + \langle \mathbf{m} \rangle^2}} \right) \right\rangle_{\xi}. \quad (4.17)$$

Remark 5. Note that Eqs. (4.17) are exactly those prescribed by the Least Action Principle $\delta f(t, \mathbf{x}) = 0$ as

$$\begin{aligned} \delta f(t, \mathbf{x}) &= \frac{\partial f(t, \mathbf{x})}{\partial \langle \mathbf{m}^\mu \rangle_{t,\mathbf{x}}} \delta \langle \mathbf{m}^\mu \rangle_{t,\mathbf{x}} = \\ &= \tanh \left(x^\mu - \frac{\langle \mathbf{m}^\mu \rangle_{t,\mathbf{x}} t}{\sqrt{1 + \langle \mathbf{m} \rangle_{t,\mathbf{x}}^2}} \right) \left(\frac{-t \langle \mathbf{m}^\mu \rangle_{t,\mathbf{x}} \delta \langle \mathbf{m}^\mu \rangle_{t,\mathbf{x}}}{(1 + \langle \mathbf{m} \rangle_{t,\mathbf{x}}^2)^{\frac{3}{2}}} \right) \\ &\quad + \frac{\langle \mathbf{m}^\mu \rangle_{t,\mathbf{x}} t \delta \langle \mathbf{m}^\mu \rangle_{t,\mathbf{x}}}{(1 + \langle \mathbf{m} \rangle_{t,\mathbf{x}}^2)^{\frac{3}{2}}} = 0. \end{aligned} \quad (4.18)$$

Remark 6. Note that, if we take the low momentum limit $|\langle \mathbf{m} \rangle| \ll 1$, we can expand the relativistic model at the lowest order

$$\begin{aligned} \frac{1}{\sqrt{1 + \langle \mathbf{m} \rangle^2}} &= 1 - \frac{\langle \mathbf{m} \rangle^2}{2} + \mathcal{O}(\langle \mathbf{m} \rangle^3), \\ \frac{\mathbf{m}}{\sqrt{1 + \langle \mathbf{m} \rangle^2}} &= \mathbf{m} + \mathcal{O}(\langle \mathbf{m} \rangle^3), \end{aligned} \quad (4.19)$$

so as to recover the classical Hopfield model and results.

5. Numerical Simulations of the classical and relativistic Hopfield networks

Once we have the Mattis self-consistent equations (4.17) we can solve them recursively and obtain the theoretical expectations for the quality of the retrieval of the various pure and spurious states (i.e. the intensities of the various Mattis magnetizations), or we can use them to compare Monte Carlo simulations to network's performances. In order to test the capabilities of the relativistic Hopfield network, we compare its results to those of the classical counterpart: to this task we performed different kind of extensive simulations as follows.

5.1. Stochastic neural dynamics and network's attractors

In order to test the amplitude and stability of pure and spurious attractors, we performed the following series of numerical simulation: once introduced random numbers η_i uniformly sampled from $[-1, +1]$, we set up the following stochastic neural dynamics,

$$\sigma_i(t+1) = \text{sign} [\tanh (\beta h_i(\boldsymbol{\sigma}(t))) + \eta_i(t)] \quad (5.1)$$

$$h_i(\boldsymbol{\sigma}(t)) = \boldsymbol{\xi} \cdot \frac{\langle \mathbf{m}_N \rangle}{\sqrt{1 + \langle \mathbf{m}_N \rangle^2}}, \quad (5.2)$$

or, in its probabilistic formulation for the entire network,

$$\begin{aligned} p_{t+1}(\boldsymbol{\sigma}) &= \prod_{i=1}^N \frac{1}{2} [1 + \sigma_i \tanh (\beta h_i(\boldsymbol{\sigma}(t)))] \\ &= \prod_{i=1}^N \frac{e^{\sigma_i h_i(\boldsymbol{\sigma}(t))}}{2 \cosh (\beta h_i(\boldsymbol{\sigma}(t)))}, \end{aligned} \quad (5.3)$$

$$\begin{aligned} p_{t+1}(\boldsymbol{\sigma}) &= \sum_{\boldsymbol{\sigma}'} W[\boldsymbol{\sigma}' \rightarrow \boldsymbol{\sigma}] p_t(\boldsymbol{\sigma}'), \quad W[\boldsymbol{\sigma}' \rightarrow \boldsymbol{\sigma}] \\ &= \prod_{i=1}^N \frac{e^{\sigma_i h_i(\boldsymbol{\sigma}(t))}}{2 \cosh (\beta h_i(\boldsymbol{\sigma}(t)))}. \end{aligned} \quad (5.4)$$

Note that, also in the dynamical evolution of the system, β tunes the level of noise in the network because it triggers the amplitude of the hyperbolic tangent (containing the signal): at each time step t , for each neuron σ_i , the field $h_i(t)$ is computed and multiplied by β , then the hyperbolic tangent is applied to $\beta h_i(t)$ and this number is compared to $\eta_i(t)$: for $\beta \rightarrow \infty$ – large signal limit – the hyperbolic tangent becomes a ± 1 sign function thus the random numbers play no longer any role and the dynamics becomes deterministic minimization, on the contrary for $\beta \rightarrow 0$ the hyperbolic tangent returns zero whatever the field and the dynamics becomes fully random.

Note further that, in Eq. (5.4), we emphasized the transition rates $W[\boldsymbol{\sigma}' \rightarrow \boldsymbol{\sigma}]$ in order to show their equivalence with the acceptance rates in the Monte Carlo route: indeed, due to the symmetry of the couplings, Detailed Balance ensures that

$$p_\infty(\boldsymbol{\sigma}) W[\boldsymbol{\sigma}' \rightarrow \boldsymbol{\sigma}] = p_\infty(\boldsymbol{\sigma}') W[\boldsymbol{\sigma} \rightarrow \boldsymbol{\sigma}'], \quad (5.5)$$

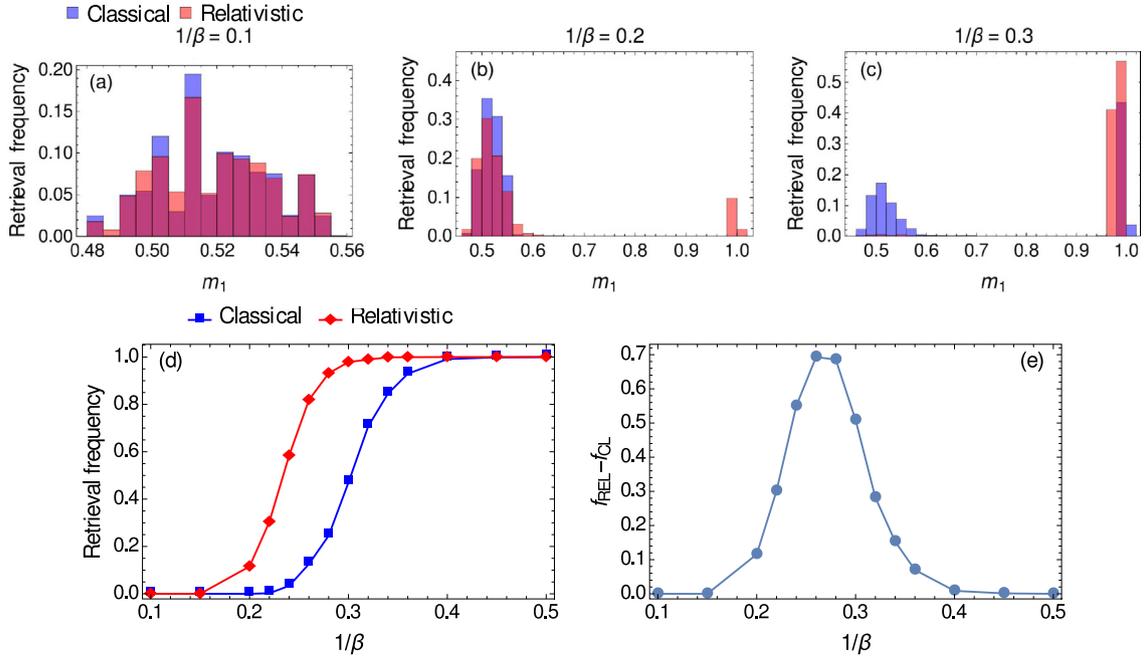


Fig. 2. First row: histograms of the maximum Mattis magnetization starting from a spurious state. Results for the maximum Mattis magnetization m_1 for both classical (blue bars) and relativistic (red bars) Hopfield models for different noise values: $\beta^{-1} = 0.1$ (a), 0.2 (b) and 0.3 (c). Purple bars are the superposition of blue and red ones. The network whose thermalization is under study is built of $N = 1000$ neurons and $P = 3$ stored patterns. Simulations consist in 8000 different runs, organized in 20 different random pattern configurations, 20 randomly selected initial conditions and 20 different stochastic evolutions for each noise level. Second row: Retrieval performances for spurious states initial conditions. Results for the retrieval performances in the noise range $0.1 \leq \beta^{-1} \leq 0.5$. The curves in (d) refer to the number of final states for the classical (blue squares) and the relativistic (red diamonds) models: starting from a spurious state, for lowest thermal energy supply (e.g. $\beta^{-1} < 0.15$), the spurious states are stable. However, for $\beta^{-1} > 0.15$ the relativistic model starts to retrieve properly and at $\beta^{-1} \sim 0.25$ roughly half of the runs end up in a proper pure state. On the contrary, for the classical model, we have to reach $\beta^{-1} > 0.22$ before obtaining a correct thermalization into a pure state and $\beta^{-1} \sim 0.32$ before half of the runs end up correctly. Mattis magnetizations whose magnitude is larger than an arbitrary threshold (set at $m_{threshold} = 0.8$) are coupled to final patterns recognized as “retrieved”: we choose $m_{threshold} = 0.8$ since it best splits spurious state intensities (whose range is close to $m_{spurious} \sim 0.5$) and pure state intensities (whose magnitude is close to one). Plot (e) is obtained subtracting the classical curve from the corresponding relativistic one, giving an empirical estimation about how the latter improves the retrieval performances. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

hence, remembering that p_∞ must assume the maximum entropy Gibbs-expression $p_\infty \propto \exp(-\beta H(\sigma, \xi))$, we have

$$p_t(\sigma_i \rightarrow \sigma'_i) = \frac{1}{1 + e^{\beta[H(\sigma|\xi) - H(\sigma'|\xi)]}}$$

namely the acceptance criterion of the Glauber algorithm used in the successive bulk of simulations (see Section 5.2).

We use the above dynamics with two different starting assumption as follows:

- **Attractors from spurious initial conditions.** We start the system sharply within a spurious state (the 3-patterns mixture for $P = 3$, see Eq. (2.15)) then we let it thermalize at a given noise level β^{-1} under the Markov dynamics (5.3) and we collect the final state of the relaxation process (for both for the classical and the relativistic models), whose existence is ensured by Detailed Balance (5.5). Results are shown in Fig. 2 for $P = 3$, Figs. 3 and 4 for the same analysis respectively for $P = 5$ and $P = 7$. In the first rows we show the histograms counting how many times the system ended up in a pure state at three different noise levels $\beta^{-1} = 0.1, 0.2, 0.3$. In the second rows we plot the frequency (i.e. the percentage) by which a pure state is reached during this thermalization as a function of the noise level β^{-1} . Finally, in Fig. 5 we plot the relative improvement of the relativistic extension with respect to Hopfield model for the values $P = 3, 5, 7$.

Two remarks are in order here: the first is that systematically the relativistic model out-performs its classical counterpart. The second is that such an improvement lies in the range of noise level $0 \leq \beta^{-1} \leq 0.5$, namely exactly in the noise region where spurious states possess an own

stability: for noise levels beyond that threshold spurious states are no longer stable and, likewise, there is generally no more reward in the relativistic extension.

- **Attractors from random initial conditions.** We prepared the system within a fully random initial configuration, then we let it thermalize at a given noise level β^{-1} under the Markov dynamics (5.3) and we collect the final state of the relaxation process (for both for the classical and the relativistic models), whose existence is ensured by Detailed Balance (5.5). Results are shown in Fig. 6. Again, in the first row we show the histograms counting the times the system ended up in a pure state at three different noise levels $\beta^{-1} = 0.1, 0.2, 0.3$. In the second row we plot the frequency (i.e. the percentage) by which a pure state is reached during this thermalization as a function of the noise level β^{-1} .

Again we highlight that, systematically, the relativistic model shows increased performances than its classical counterpart (as long as the spurious states are locally stable).

5.2. Comparison between Monte Carlo simulations and analytical outcomes

At first we performed extensive Monte Carlo runs to check how the evolution of the Mattis magnetizations (versus the noise level) predicted analytically coincides with the numerical one.

Monte Carlo simulations have been implemented iterating a Glauber dynamics within the following scheme:

- Select a neuron at random and compute the difference $\Delta H(\sigma, \xi)$ in the cost function due to its spin flip.

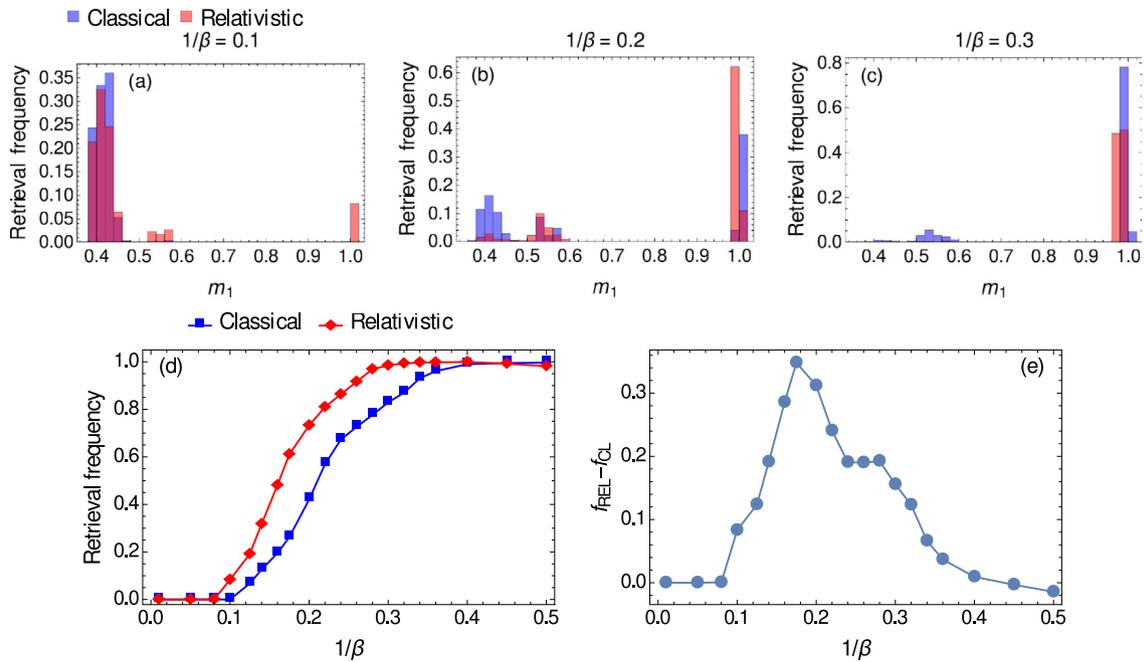


Fig. 3. First row: histograms of the maximum Mattis magnetization starting from a spurious state for $P = 5$. Results for the maximum Mattis magnetization m_1 for both classical (blue bars) and relativistic (red bars) Hopfield models for different noise values: $\beta^{-1} = 0.1$ (a), 0.2 (b) and 0.3 (c). Purple bars are the superposition of blue and red ones. The network whose thermalization is under study is built of $N = 1000$ neurons and $P = 5$ stored patterns. Simulations consist in 8000 different runs, organized in 20 different random pattern configurations, 20 randomly selected initial conditions and 20 different stochastic evolutions for each noise level. Second row: Retrieval performances for spurious states initial conditions. Results for the retrieval performances in the noise range $0.1 \leq \beta^{-1} \leq 0.5$. The curves in (d) refer to the number of final states for the classical (blue squares) and the relativistic (red diamonds) models: starting from a spurious state, for lowest thermal energy supply (e.g. $\beta^{-1} < 0.15$), the spurious states are stable. However, for $\beta^{-1} > 0.15$ the relativistic model starts to retrieve properly and at $\beta^{-1} \sim 0.25$ roughly half of the runs end up in a proper pure state. On the contrary, for the classical model, we have to reach $\beta^{-1} > 0.22$ before obtaining a correct thermalization into a pure state and $\beta^{-1} \sim 0.32$ before half of the runs end up correctly. Mattis magnetizations whose magnitude is larger than an arbitrary threshold (set at $m_{threshold} = 0.8$) are coupled to final patterns recognized as “retrieved”: we choose $m_{threshold} = 0.8$ since it best splits spurious state intensities (whose range is close to $m_{spurious} \sim 0.5$) and pure state intensities (whose magnitude is close to one). Plot (e) is obtained subtracting the classical curve from the corresponding relativistic one, giving an empirical estimation about how the latter improves the retrieval performances. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

- If $\Delta H(\sigma|\xi) < 0$ (hence the flip is convenient), the move is accepted with probability $\propto \exp(\beta \Delta H(\sigma|\xi)) / [1 - \exp(\Delta H(\sigma|\xi))]$, otherwise is rejected (Glauber criterion).
- If $\Delta H(\sigma|\xi) > 0$ (hence the flip is not convenient), the move is rejected.

Again due to Detailed Balance, Hopfield’s Hamiltonian plays also as a Lyapunov function (Coolen et al., 2005) hence (properly tuning the noise level β) we can control its convergence to a minimum (a global one – pure state – or a local one – spurious state) whose attainment is reached when the spins share the same sign of the fields acting on them (and the simulation then stops as any spin flip can only raise the energy value).

In Fig. 7 we show the behavior of the Mattis order parameter as a function of the noise in the network both as predicted by the theory, using the self-consistent Eq. (3.11) for the classical Hopfield model – left panel – and the self-consistent Eq. (4.17) for the relativistic counterpart – right panel – as well as the same behavior as obtained from simulations: these have been obtained via standard Monte Carlo runs for a system built up by $N = 1000$ neurons (whose details are reported in the caption of Fig. 7).

Despite being redundant, we stress that two branches – rather than one – are reported in Fig. 7 due to the unbroken model’s spin-flip symmetry: both the *pattern* ξ^1 and the *anti-pattern* $-\xi^1$ are attractors for the neural dynamics here but clearly this does not happen when the thresholds for firing h_i are re-introduced in the theory (see Eq. (2.8)).

In Fig. 8 we compare the solution at various sizes of the network: as it shines from the plots, both for the classic case (left panel) as well as for the relativistic generalization (right panel), as N grows the analytical scaling is approached.

5.3. Depth of the attractors and energy gaps

Inspired by the pioneering works of the Gardner on the estimates about depth and stability of the basins of attraction of pure and spurious states (Gardner, 1998), we now introduce the percentage of random spin-flips d we impose on the system (such that for e.g. $d = 0$ no random spin-flip at all is performed, while for e.g. $d = 50\%$ we flip randomly one half of the total spins): the underlying idea is to start the system into a known state (e.g. a spurious state) and then we reshuffle it by *kicking* randomly a percentage d of its neurons and checking, as d grows, the return (or the escape) of the network from the initial attractor. Results are shown in Fig. 9 focusing on a 3-mixture spurious state in order to quantify the pruning capabilities of the relativistic model and compare them with those of its classical limit: while, for low noise level ($\beta^{-1} = 0.1$, left panel), we have to reach consistent percentages of spin-flip to allow the network to escape the spurious state ($d \sim 0.40$) for both the models, already at mild noise level ($\beta^{-1} = 0.2$, middle panel) we see that spurious states already become unstable in the relativistic extension – while being stable in the classical limit – and for values of $d \sim 0.35$ a further improvement in the pure state retrieval is obtained (for both the models). Finally in the last plot (right panel) we show that for moderate noise levels ($\beta^{-1} = 0.3$) the relativistic model essentially always escapes from the spurious state regardless of d , tacitly showing that the basins of attraction of those states are corrupted by the unlearning contributions nestled in the relativistic Hopfield cost function (4.2).

A reasonable cause of the increased performances of the relativistic Hopfield model may lie in a decreased energy barrier between the spurious states and the maxima surrounding them. To check this idea, again we can compare these gaps for the classical

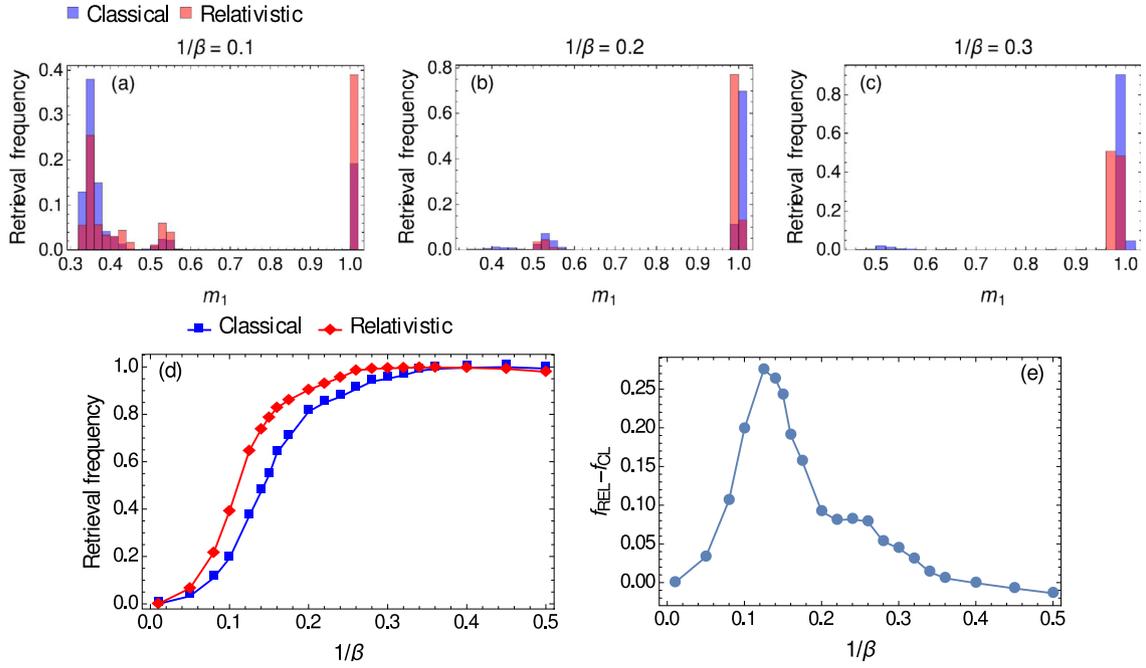


Fig. 4. First row: histograms of the maximum Mattis magnetization starting from a spurious state for $P = 7$. Results for the maximum Mattis magnetization m_1 for both classical (blue bars) and relativistic (red bars) Hopfield models for different noise values: $\beta^{-1} = 0.1$ (a), 0.2 (b) and 0.3 (c). Purple bars are the superposition of blue and red ones. The network whose thermalization is under study is built of $N = 1000$ neurons and $P = 7$ stored patterns. Simulations consist in 8000 different runs, organized in 20 different random pattern configurations, 20 randomly selected initial conditions and 20 different stochastic evolutions for each noise level. Second row: Retrieval performances for spurious states initial conditions. Results for the retrieval performances in the noise range $0.1 \leq \beta^{-1} \leq 0.5$. The curves in (d) refer to the number of final states for the classical (blue squares) and the relativistic (red diamonds) models: starting from a spurious state, for lowest thermal energy supply (e.g. $\beta^{-1} < 0.15$), the spurious states are stable. However, for $\beta^{-1} > 0.15$ the relativistic model starts to retrieve properly and at $\beta^{-1} \sim 0.25$ roughly half of the runs end up in a proper pure state. On the contrary, for the classical model, we have to reach $\beta^{-1} > 0.22$ before obtaining a correct thermalization into a pure state and $\beta^{-1} \sim 0.32$ before half of the runs end up correctly. Mattis magnetizations whose magnitude is larger than an arbitrary threshold (set at $m_{\text{threshold}} = 0.8$) are coupled to final patterns recognized as “retrieved”: we choose $m_{\text{threshold}} = 0.8$ since it best splits spurious state intensities (whose range is close to $m_{\text{spurious}} \sim 0.5$) and pure state intensities (whose magnitude is close to one). Plot (e) is obtained subtracting the classical curve from the corresponding relativistic one, giving an empirical estimation about how the latter improves the retrieval performances. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and the relativistic models: we still start the network’s dynamics in a spurious state and end up in a pure state – the target pattern – but the way we move from the spurious state to the target is via controlled noiseless random walks, *i.e.* a ground-state dynamics: at each step, we select a spin σ_i and if it is already aligned with the target pattern we want the network to reach we just move on, otherwise we flip it and we compute again the energy, then we repeat the process until the pure state is reached. Starting from a spurious state, and approaching a pure one, an energy barrier has to be crossed: we collected these energy gaps and, averaging statistics consisting in 1600 different runs (sampling 40 different pattern choices to generate spurious states, for each of which we performed 40 different stochastic evolutions), and we found that the energy barrier to escape a spurious state is more than halved in the relativistic model if compared to the classical counterpart, *i.e.* $\Delta E_{\text{relativ}} / \Delta E_{\text{classic}} \sim 0.75$, hence confirming that unlearning has an effective role in pruning the network.

5.4. Signal-to-noise analysis

In this subsection, we want to better analyze the role of the various contributions that build up the generalized Hopfield cost function (4.2), as they appear in its expansion (4.3). From now on, we will write it in a symbolic and visually compact way as

$$H(\sigma) \sim -c_2 \cdot \sigma\sigma + c_4 \cdot \sigma\sigma\sigma\sigma - c_6 \cdot \sigma\sigma\sigma\sigma\sigma\sigma + \dots \quad (5.6)$$

Here, the main contribution (weighted by c_2) is the standard pairwise classical limit (*i.e.* the Hopfield model). The next one (weighted by c_4) is a four-spin contribution preventing retrieval

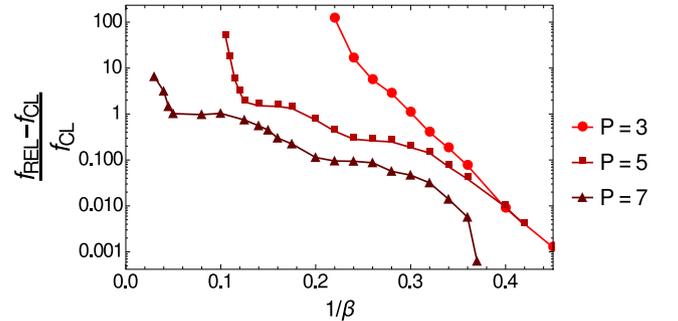


Fig. 5. Relative improvement as a function of the temperature. The plot shows the empirical curves for the empirical relative performances $(f_{\text{REL}} - f_{\text{CL}}) / f_{\text{CL}}$ of relativistic Hopfield model versus its classical counterpart for $P = 3, 5, 7$, where f is the retrieval frequency of a given model (recall that we say that a pattern is retrieved if the associated Mattis magnetization is higher than a fixed value $m_{\text{threshold}} = 0.8$). The starting initial condition is aligned to the maximal spurious state (obtained by a majority rule applied to the sum of all patterns, *i.e.* $\sigma_{0,i} = \text{sign}(\sum_{\mu} \xi_i^{\mu})$ for each i). Increasing the value of P , such spurious attractors become unstable also for the classical model, therefore causing the fast fall-off of the relative improvement for increasing P . For the particular value $P = 3$, the curve falls also outside the range and, for low temperature values ($\beta^{-1} \sim 0.1$), the relative improvement is very high *i.e.* $(f_{\text{REL}} - f_{\text{CL}}) / f_{\text{CL}} \sim 120$, since in this case the retrieval frequency for the relativistic model is finite, while for the classical one is almost zero.

(*i.e.* it has the opposite sign w.r.t. the c_2 term), while the successive one again favors retrieval and so on.

In particular, we consider the first four terms in the Taylor expansion (4.3), which we write here as the following four truncated

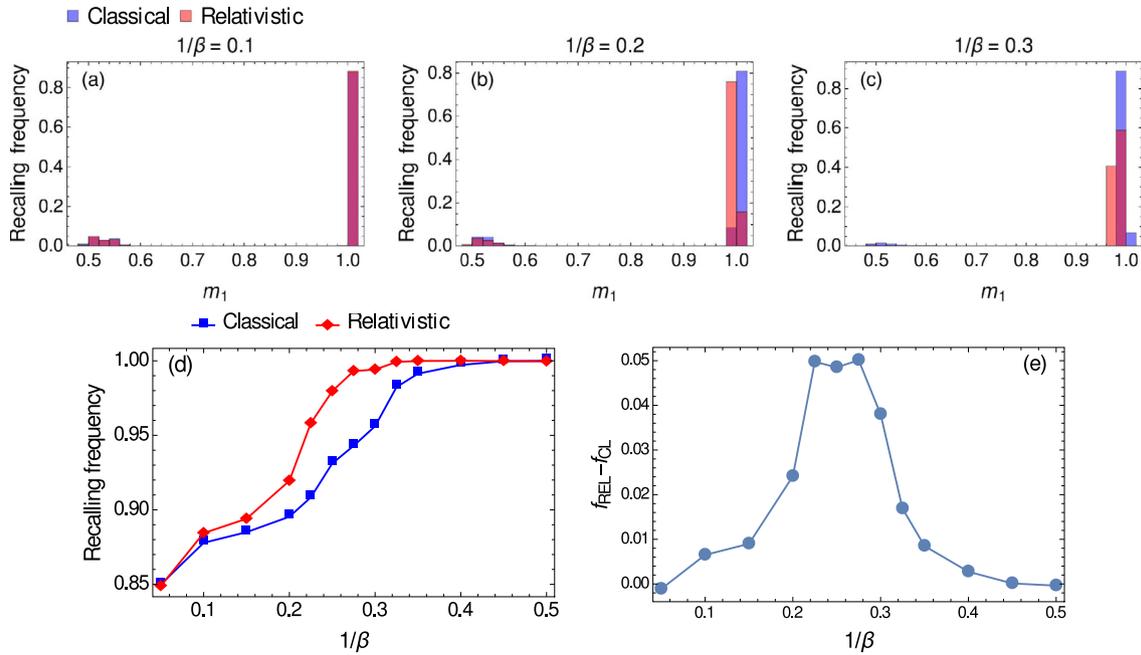


Fig. 6. First row: histograms of the Mattis magnetization from random initial conditions for $P = 3$. Results for the maximum Mattis magnetization m_1 for both classical (blue bars) and relativistic (red bars) Hopfield models for different noise values: $\beta^{-1} = 0.1$ (a), 0.2 (b) and 0.3 (c). Purple bars are the superposition of blue and red ones. The network whose thermalization is under study is built of $N = 1000$ neurons and $P = 3$ stored patterns. Monte Carlo simulations consist in 8000 different runs, organized in 20 different random pattern configurations, 20 randomly selected initial conditions and 20 different stochastic evolutions for each noise level. Second row: retrieval performances for random initial conditions. Results for the retrieval performances in the noise range $0.05 \leq \beta^{-1} \leq 0.5$. The curves in (d) refer to the number of final states for the classical (blue squares) and the relativistic (red diamonds) models: Mattis magnetizations whose magnitude is larger than an arbitrary threshold (set at $m_{threshold} = 0.8$) are coupled to final patterns recognized as “retrieved”: we choose $m_{threshold} = 0.8$ since it best splits spurious state intensities (whose range is close to $m_{spurious} \sim 0.5$) and pure state intensities (whose magnitude is close to one). Plot (e) is obtained subtracting the classical curve from the corresponding relativistic one, giving an empirical estimation about how the latter improves the retrieval performances. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

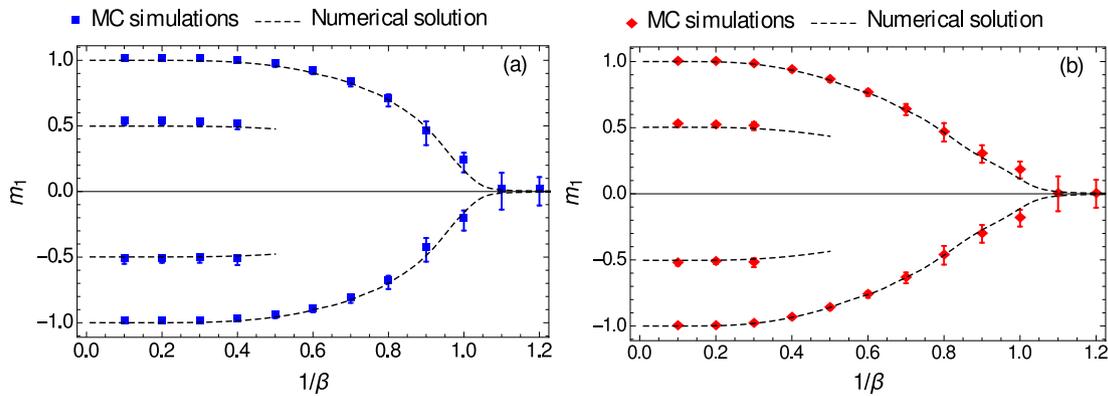


Fig. 7. Monte Carlo simulations vs self-consistency prediction for the Mattis order parameter. Results on the evolution of $\langle m^{\mu} \rangle$ vs β are presented for Monte Carlo simulations of the classical (a, blue squares) and relativistic (b, red diamonds) models for a network with $N = 1000$, where $P = 3$ (orthogonal) patterns are stored. The data points are averages over 20 different pattern realizations, for each of which we sampled 20 random initial conditions (i.e. 8000 runs at any given noise level). Monte Carlo results are compared to the main branches (pure states) of the theoretical solution (dashed black curves) of the self-consistency Eqs. (3.11) and (4.17). Note the presence of the spurious states (evidenced by the two segments of magnetization’s values $m_1 \sim \pm 0.5$) for noise level $\beta^{-1} \leq 0.45$ and note further that the variances of the data-points spread at the bifurcation point (i.e. at $\beta^{-1} \sim 1$), as typical in a second order phase transition.

cost-functions:

$$H_2(\sigma) = -c_2 \cdot \sigma \sigma, \quad (5.7)$$

$$H_4(\sigma) = -c_2 \cdot \sigma \sigma + c_4 \cdot \sigma \sigma \sigma \sigma, \quad (5.8)$$

$$H_6(\sigma) = -c_2 \cdot \sigma \sigma + c_4 \cdot \sigma \sigma - c_6 \cdot \sigma \sigma \sigma \sigma \sigma \sigma, \quad (5.9)$$

$$H_8(\sigma) = -c_2 \cdot \sigma \sigma + c_4 \cdot \sigma \sigma - c_6 \cdot \sigma \sigma \sigma \sigma \sigma \sigma + c_8 \cdot \sigma \sigma \sigma \sigma \sigma \sigma \sigma \sigma. \quad (5.10)$$

From the theoretical side, we compare the first two cost-functions with a standard signal-to-noise analysis. One of the main advantages of this technique – for the purpose we are using it now, namely to deepen the role of these various contributions of the expansion (5.6) – is that it works in the noiseless $\beta \rightarrow \infty$ limit, therefore stripping any genuine effect of the model without the risk of confusing a thermal fluctuation with a novel feature. Further, we compare the truncated cost-functions numerically by looking at the internal field distributions in the final states of the stochastic evolution for various thermal noise levels, see Figs.10 and 11.

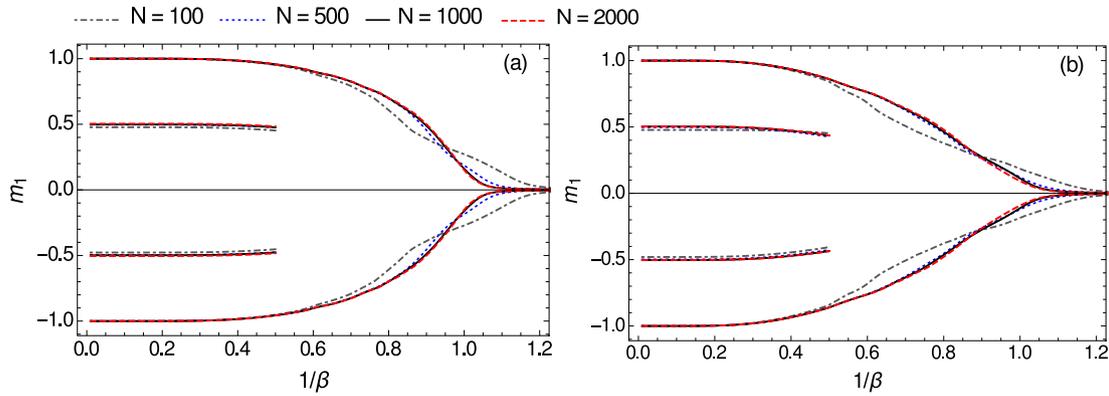


Fig. 8. Finite size dependence. Results for different networks of size $N = 100, 500, 1000, 2000$ are shown for the classical Hopfield model (left panel) and for the relativistic extension (right panel): as expected, as N grows, the curves approach the analytic solution (obtained in the thermodynamic limit) while for the smallest values of N (e.g. $N = 100$) there is no trace of the phase transition as expected.

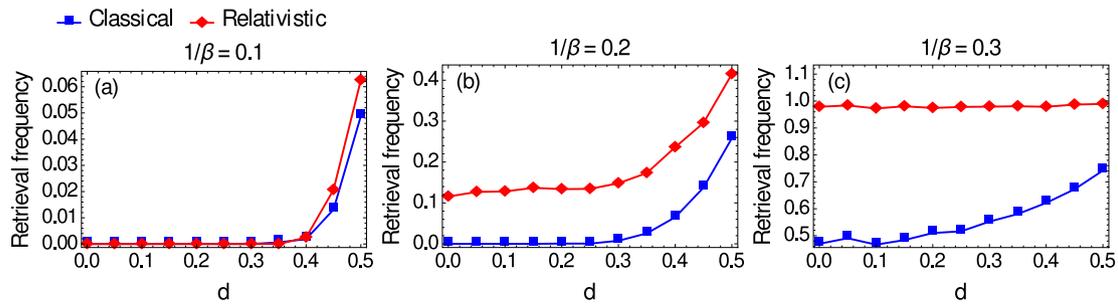


Fig. 9. Stability of retrieval performances. Results for retrieval frequency as a function of the initial spin-flip fraction d for $\beta^{-1} = 0.1$ (a), 0.2 (b) and 0.3 (c) in the range $0 \leq d \leq 0.5$ (higher spin-flips fraction would take the initial condition immediately too far away from the spurious attractor, so they are not relevant in this discussion). In the left panel we see that for too low noise levels we have to significantly flip the neural states before both the models become capable to escape from the spurious state, however in the middle panel we see that already for mild noise level the relativistic model has the ability to drift away from the spurious toward the pure attractor, finally, increasing further the noise level in the network, in the right panel we can appreciate how the relativistic model out-performs its classical counterpart in the entire analyzed range.

Regarding the signal-to-noise analysis, it is relatively simple (despite being quite lengthy) to prove the next theorem.

Theorem 3. *The cost function $H_4(\sigma)$, compared with the standard Hopfield cost function $H_2(\sigma)$ paints a phase space whose spurious states are destabilized w.r.t. to the pure ones, i.e.*

$$\frac{\text{noise}(H_4(\sigma))}{\text{noise}(H_2(\sigma))} \Big|_{\text{spurious state}} > \frac{\text{noise}(H_4(\sigma))}{\text{noise}(H_2(\sigma))} \Big|_{\text{pure state}} \quad (5.11)$$

Remark 7. We stress that, the c_4 term – beyond destabilizing spurious states (see Eq. (5.11)) – plays a role even as a regularizer for the amplitudes of the fields experienced by the neurons, as shown in Figs. 10 and 11.

Further, still in the same figures, we also show H_6 and H_8 at work in order to allow, by a direct comparison of the outputs produced by the various cost functions, how the balance between attractive and repulsive higher orders is crucial to stabilize the retrieval of the pure states and destabilize that of the spurious ones. Indeed, looking the distributions relative to H_2 (i.e. the standard Hopfield reference) and those returned by the truncated models, the suppression of spurious states by the relativistic extension is manifest.

The proof of Theorem 3 is standard and works by brute force, thus we will report solely the key points for the sake of simplicity. We require that the fields acting over the neurons are aligned with the neurons themselves, namely we require that $h_i \sigma_i \geq 0, \forall i \in (1, \dots, N)$ and we check this stability criterion both for the pure

state as well as for the 3-mixture-state (note that this is the most robust spurious state, hence these results act as rigorous bounds for all the higher order – more fragile – spurious states).

It is well known that for H_2 , the signal-to-noise ratio returns (for large N)

- pure state: signal + noise = $1 + \sqrt{(P-1)/N}$,
- spurious state: signal + noise = $\frac{1}{2} + \sqrt{(P-3)/N}$.

The same analysis for the H_4 model instead reads

- pure state: signal + noise = $\frac{1}{2} + \frac{1}{2}\sqrt{(P-1)/N}$ + subleading contributions,
- spurious state: signal + noise = $\frac{5}{16} + \frac{5}{8}\sqrt{(P-3)/N}$ + sub-leading contributions.

Thus, by properly combining these partial results together and noticing that those are the prevailing ones as the series (5.6) converges to the relativistic Hamiltonian (4.6), in the thermodynamic limit we obtain the statement of the theorem.

5.5. A 1-parametric generalization of relativistic Hopfield model

We would like to conclude this work with a possible extension of our investigation in these regards: the relativistic extension of the Hopfield model has a nice justification in terms of the mechanical analogy through Hamilton–Jacobi formalism, i.e. it acted as a guide in the formulation of the present theory somehow as the two-body Hamiltonian in Statistical Mechanics drove the

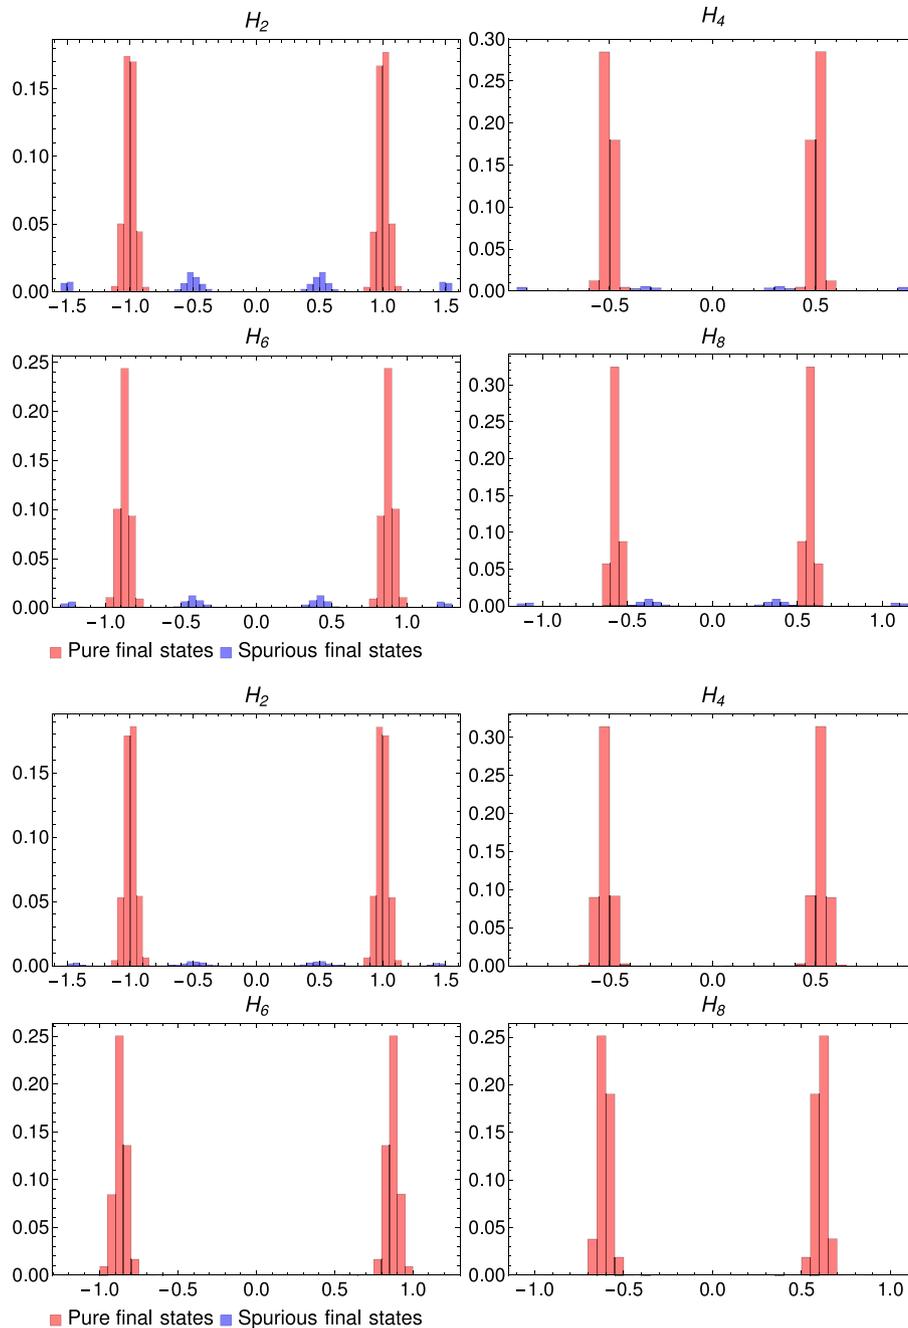


Fig. 10. Internal field distributions (random initial condition). Distribution for internal fields in the final configuration of the network's truncated expansion (see Eq. (5.6)), respectively for $H_2(\sigma)$, $H_4(\sigma)$, $H_6(\sigma)$, $H_8(\sigma)$ with random initial configurations and two thermal noise choices: $T \equiv \beta^{-1} = 0.2$ (upper box) and 0.3 (lower box). The network parameters are $N = 1000$ and $P = 3$, and the statistics used consists in 8000 different runs, organized in 20 different random patterns configurations, 20 randomly selected initial conditions and 20 different stochastic evolutions for each noise level.

first formulation of standard Hopfield networks. However, on the concrete side of neural networks, our model is only one of the possible choices regarding a cost-function incorporating the principles of deep learning and network pruning (*i.e.* whose expansion has contributions in alternate signs), so it is reasonable to extend it with more general cost-functions (whose sole request is to recover the Hopfield associative memory framework as the leading contribution). The reward in this further generalization has emerged in the previous section, in particular as coded by Eq. (5.11) and in the plots of Fig. 12, as it can be convenient to have a tunable regularizing parameter to optimize the network (that is to remove most possible spurious states without affecting the retrieval of pure ones). A straightforward proposal to the case is

thus the following 1-parametric Hamiltonian

$$H_N(\sigma|\xi) = -N \frac{\sqrt{1 + \lambda \mathbf{m}_N(\sigma)^2}}{\lambda}, \quad (5.12)$$

whose Taylor expansion

$$\frac{H_N}{N} = -\frac{1}{\lambda} - \frac{\mathbf{m}_N^2}{2} + \frac{\lambda \mathbf{m}_N^4}{8} + O(\mathbf{m}_N^6), \quad (5.13)$$

gives the Hopfield model at the leading order.⁸ In order to have pruning corrections with alternate signs, we have to choose λ to

⁸ Of course, the zero-point energy λ^{-1} does not affect the thermodynamic properties of the system.

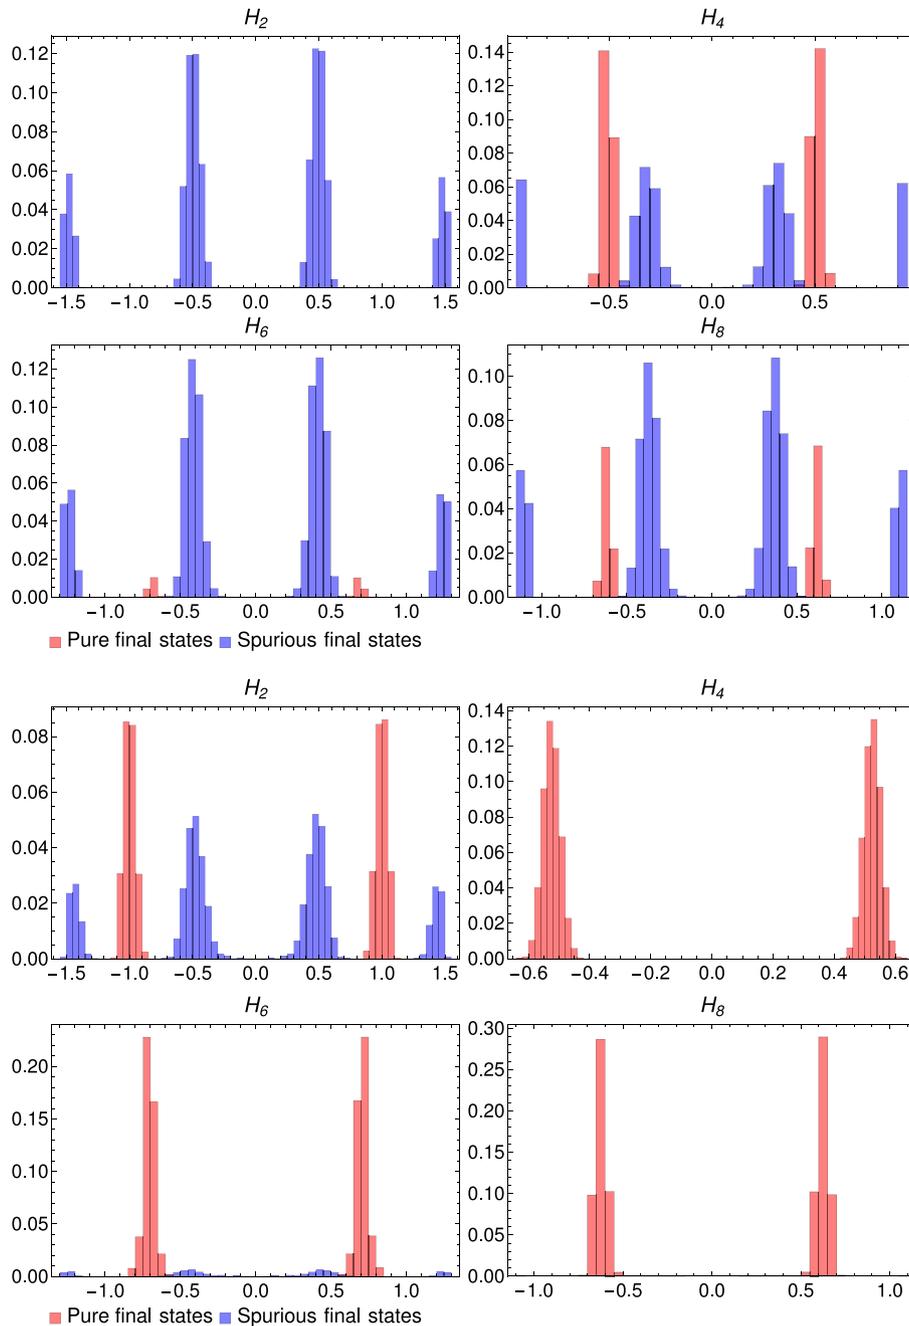


Fig. 11. Internal field distributions (spurious initial condition). Distribution for internal fields in the final configuration of the network’s truncated expansion (see Eq. (5.6)), respectively for $H_2(\sigma)$, $H_4(\sigma)$, $H_6(\sigma)$, $H_8(\sigma)$ with random initial configurations and two thermal noise choices: $T \equiv \beta^{-1} = 0.2$ (upper box) and 0.3 (lower box). The network parameters are $N = 1000$ and $P = 3$, and the statistics used consists in 1600 different runs, organized in 40 different random patterns configurations and 40 different stochastic evolutions for each noise level. The initial condition is given by $\sigma_{0,i} = \text{sign}(\sum_{\mu} \xi_i^{\mu})$.

be a positive real number. To give a sketch of the performances of this new model, we compared the retrieval frequency curves for different λ values for a network with $N = 1000$ and $P = 3$ and choosing the initial condition to be aligned to the spurious state, see Fig. 12. As a result, λ higher than 1 (corresponding to the relativistic Hopfield model) leads to models performing better in retrieving when the network is prepared in a spurious state configuration. This is a good point which should encourage future study in this context. In particular, a more detailed investigation of this 1-parametric model (as well as more general suitable choices) is required, both numerically and analytically. We leave this point

and the study of the high storage regime (i.e. replica trick analysis and the realization of the phase diagram) open for future works.

6. Conclusions and further developments

Regarding the unlearning phenomenon, quoting Hassibi and Stork (Hassibi & Stork, 1993), a central problem in machine learning and pattern recognition is to minimize the system complexity consistent with the training data. (...) If we begin with a trained network having too many weights, the question becomes: which weights should be eliminated? Answering this question gave rise to pruning in machine learning. However, while pruning algorithms

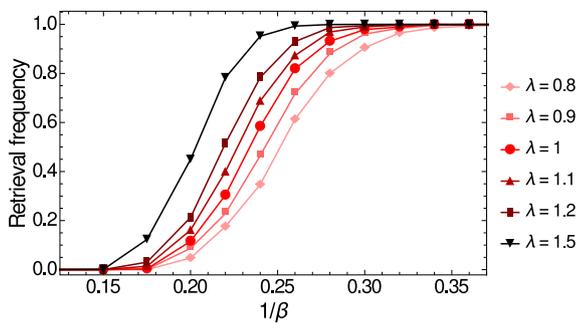


Fig. 12. Retrieval performances for the 1-parametric model. Results for 1-parametric model for a network of size $N = 1000$ and $P = 3$ patterns stored. The initial condition is aligned to the spurious state: $\sigma_{0,i} = \text{sign}(\xi_i^1 + \xi_i^2 + \xi_i^3)$. For values of the control parameter λ lower than 1, performances gradually worsen with respect to the relativistic Hopfield model. On the other side, increasing the value of λ beyond 1, the retrieval frequency accordingly grows and the network dynamics more likely ends in a pure state. As always for spurious state initial conditions, we averaged the results of the Glauber dynamics over 40 different pattern realizations and, for each of them, over 40 different stochastic evolution.

have been extensively exploited in neural networks away from detailed balance, e.g. radial-basis-function and/or feed-forward neural networks (Hassibi, Stork, & Wolff, 1993; Huang, Saratchandran, & Sundararajan, 2005; Le Cun, Denker, & Solla, 1990), not so much has been said for stochastic neural dynamics obeying detailed balance, namely for associative neural networks as Boltzmann machines and/or Hopfield-like models. For these models, however, very powerful tricks appeared in the past to remove *spurious states* (locally stable states, pure pattern's mixtures that do not match exactly any of the stored patterns): an unlearning procedure inspired by REM-sleep has been proposed (Aserinsky & Kleitman, 1953; Crick & Mitchinson, 1983; Hopfield et al., 1983) but it has always been seen as an *a fortiori* algorithmic procedure rather than an intrinsic property of a model.

Focusing on Deep Learning instead, since the seminal review by LeCun, Bengio and Hinton (Le Cun et al., 2015), where deep learning impressive skills have been listed alongside our limitation in understanding how deep learning machines achieve these scores, a plethora of contributions quickly appeared on the theme, among which is the Hopfield's idea to use higher-order interactions (many-body terms) in the cost-function to be optimized. However, as for the unlearning, this intriguing proposal did not come out as a natural property of a model.

In this paper we tried to merge these two major breakthroughs in Artificial Intelligence (*i.e.* pruning/unlearning and deep learning) by re-obtaining them as particular features of a unique and very natural model, the relativistic extension of the Hopfield paradigm.

To accomplish this task, at first, we had to develop a full mechanical analogy of the statistical mechanical treatment of the standard Hopfield model (Hopfield, 1982), checking that the mechanical analogy (largely in use in spin-glass Literature) does work correctly even when dealing with neural networks: in this mirror, the free energy of the model plays as a mechanical action and it obeys a classic (*i.e.* not relativistic) Hamilton–Jacobi equation in the space of parameters (where the two-body coupling – coding for two points correlations – tunes the time, while the one-body coupling – coding for one point correlations – covers the role of space). Extremizing the action by the least square action principle we re-obtain the correct expression for the self-consistency of the Mattis order parameter, as obtained through the standard route by Amit–Gutfreund–Sompolinsky (Amit, 1989) in the low storage case.

Once checked that the mechanical analogy correctly recovers all the details of the celebrated Hopfield picture, we noticed two

pathologies in the classical mechanical treatment, namely an underlying Minkowskian metric tensor and a bounded velocity for the fictitious particle of the mechanical motion. Both these observations suggested to extend the mechanical analogy to a relativistic treatment: this naturally introduced a novel cost-function, the relativistic generalization of the original Hopfield proposal.

This new model has been studied in all details, confined to the low storage (namely for regimes where there is abundance of available neurons w.r.t. patterns to be stored) and an analytical expression for its free energy, as well as a prescription for the evolution of the Mattis order parameters, has been explicitly obtained and shown to be in total agreement with numerical simulations. We also considered a 1-parametric generalization of the relativistic model, giving a numerical sketch of how performances increase with a suitable choice of the control parameter. This result opens the possibility to the study of more general cost-function choices.

Remarkably such Hamiltonians, once Taylor expanded, are shown to include all the (even) higher order monomials (beyond the two-body interactions of the original Hopfield framework) and these monomials succeed one another with alternate signs, such that the leading term is the standard pairwise Hopfield model (ensuring retrieval), the next one is a four-neuron coupling with reverse sign (hence appealing for unlearning/pruning the network), while the third term is a six-neuron coupling with the correct sign (hence appealing for deep learning/increasing memory storage).

These two features are recovered as *emergent properties* in the relativistic extension of the Hopfield model and clearly deserve detailed inspection: in this paper we deepened solely the former (namely the unlearning capability that the network uses for removing spurious states from the retrieval landscape), postponing an analysis of the storage capacity of the network (the high storage case) to a forthcoming paper. The reason behind this choice is that the required mathematical treatment is completely different. When dealing with the high storage (*i.e.* to check storage capabilities and deep learning skills through disordered statistical mechanics) further concepts borrowed from the statistical mechanics of spin glasses are required (mainly replicas and overlaps among replicas). We plan to report soon also results on the high-storage regime.

It is also worth pointing out the *possible role* of this new *physical tool* - *i.e.*, the “mechanical analogy” - we introduced in the context of Artificial Intelligence, even beyond the particular case addressed in this paper: it is indeed our genuine credo that concepts borrowed from Theoretical Physics deeply helped so far in AI. For instance, the whole statistical mechanics theory of Amit–Gutfreund–Sompolinsky has been a guide for pattern recognition and content-addressable memory, while the contrastive divergence by Ackley–Hinton–Sejnowski for the restricted Boltzmann machine played a major role in Statistical Learning Theory (not to point out that Boltzmann in *restricted Boltzmann machine* is Ludwig Boltzmann, the father of Statistical Mechanics, one of the main branches of Theoretical Physics). Thus, from this perspective, extending the package of techniques that we can borrow from Theoretical Physics by including Classical and Relativistic Mechanics could be seen as something potentially very useful in Neural Networks in complete general terms.

Finally, a natural conclusive remark is that while we restricted our analysis to binary weights/patterns (mainly for the sake of continuity with the mainstream in Literature on Hopfield networks), our approach holds for real-valued variable too (as can be easily understood noticing that the pattern's average is performed at the end of the whole calculation, en route for the explicit expression of the model's free energy).

Finally, we conclude this paper by observing also that the mechanical analogy can perform as a useful tool possibly beyond the associative memory theory, for other problems in Artificial Intelligence, Hopfield networks remaining its first benchmark.

Acknowledgments

AB acknowledges partial financial support by GNFM-INdAM (via AGLIARI2018), by MIUR (via research basic fundings) and by the Rete Match: Progetto Pythagoras (CUP:J48C17000250006).

AF would like to thank Dr. Elena Agliari for the useful discussions, especially concerning the numerical analysis.

References

- Ackley, D. H., Hinton, G. E., & Sejnowski, T. J. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, 9(1), 147–169.
- Agliari, E., et al. (2012). Multitasking associative networks. *Physical Review Letters*, 109, 268101.
- Agliari, E., et al. (2013a). Multitasking attractor networks with neuronal threshold noises. *Neural Networks*, 49, 19.
- Agliari, E., et al. (2013b). Parallel retrieval of correlated patterns: From Hopfield networks to Boltzmann machines. *Neural Networks*, 38, 52.
- Agliari, E., et al. (2017). Neural Networks retrieving binary patterns in a sea of real ones. *Journal of Statistical Physics*, 168, 1085.
- Amit, D. J. (1989). *Modeling brain functions*. Cambridge Univ. Press.
- Arnold, V. I., Kozlov, V. V., & Neishtadt, A. I. (2007). *Mathematical aspects of classical and celestial mechanics*. Heidelberg: Springer.
- Arsie, A., et al. (2015a). Integrable viscous conservation laws. *Nonlinearity*, 28(6), 1859.
- Arsie, A., et al. (2015b). On integrable conservation laws. *Proceedings of the Royal Society of London, Series A (Mathematical and Physical Sciences)*, 471(2173), 20140124.
- Aserinsky, E., & Kleitman, N. (1953). Regularly occurring periods of eye motility, and concomitant phenomena, during sleep. *Science*, 118(3062), 273–274.
- Barra, A. (2008). The mean field Ising model through interpolating techniques. *Journal of Statistical Physics*, 132(5), 787–809.
- Barra, A., Di Biasio, A., & Guerra, F. (2010). Replica symmetry breaking in mean field spin glasses through Hamilton–Jacobi technique. *Journal of Statistical Mechanics: Theory and Experiment*, P09006.
- Barra, A., Genovese, G., Guerra, G., & Tantari, D. (2011). Equilibrium statistical mechanics of bipartite spin systems. *Journal of Physics A*, 44, 245002.
- Barra, A., Genovese, G., Sollich, P., & Tantari, D. (2017). Phase transitions of restricted Boltzmann machines with generic priors. *Physical Review E*, 96, 042156.
- Barra, A., Genovese, G., Sollich, P., & Tantari, D. (2018). Phase diagram of restricted Boltzmann machines & generalized Hopfield models. *Physical Review E*, 97, 022310.
- Barra, A., et al. (2012). On the equivalence among Hopfield neural networks and restricted Boltzmann machines. *Neural Networks*, 34, 1–9.
- Barra, A., et al. (2014). On quantum and relativistic mechanical analogues in mean field spin models. *Proceedings of the Royal Society of London, Series A*, 470, 20140589.
- Bovier, A., & Gayrard, V. (1997). The retrieval phase of the Hopfield model: a rigorous analysis of the overlap distribution. *Probability Theory and Related Fields*, 107(1), 61–98.
- Bovier, A., & Gayrard, V. (1998). Hopfield models as generalized random mean field models. In *Mathematical aspects of spin glasses and neural networks*. Boston: Birkhauser.
- Bovier, A., Gayrard, V., & Picco, P. (1994). Gibbs states of the Hopfield model in the regime of perfect memory. *Probability Theory and Related Fields*, 100(3), 329–363.
- Bovier, A., Gayrard, V., & Picco, P. (1995). Gibbs states of the Hopfield model with extensively many patterns. *Journal of Statistical Physics*, 79(1), 395–414.
- Choromanska, A. M., Henaff, M., Mathieu, M., Ben Arous, G., & Le Cun, Y. (2015). The loss surfaces of multilayer networks. *Artificial Intelligence and Statistics*.
- Coolen, A. C. C., Kuhn, R., & Sollich, P. (2005). *Theory of neural information processing systems*. Oxford Press.
- Crick, F., & Mitchinson, G. (1983). The function of dream sleep. *Nature*, 304, 111.
- Decelle, A., Fissore, G., Furtlehner, C. (2017). Spectral learning of restricted Boltzmann machines. arXiv preprint arXiv:1708.02917.
- Dotsenko, V., Yarinun, N. D., & Dorotheyev, E. A. (1991). Statistical mechanics of Hopfield-like neural networks with modified interactions. *Journal of Physics A*, 24, 2419.
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1), 3–71.
- Gardner, E. (1998). The space of interactions in neural network models. *Journal of Physics A*, 21, 257.
- Genovese, G., et al. (2009). A mechanical approach to mean field spin models. *Journal of Mathematical Physics*, 50(5), 053303.
- Goodfellow, I., Bengio, Y., & Courville, A. (2017). *Deep learning*. MIT Press.
- Guerra, F. (2001). Sum rules for the free energy in the mean field spin glass model. *Fields Institute Communications*, 30, 161.
- Guerra, F. (2012). How glassy are neural networks?. *Journal of Statistical Mechanics: Theory and Experiment*, P07009.
- Guerra, F., et al. (2014). Mean Field bipartite spin models treated with mechanical techniques. *European Physical Journal B*, 87, 74.
- Hartley, D. (1791). *Observation of man, his frame, his duty and his expectations*. London: Johnson Press.
- Hassibi, B., & Stork, D. G. (1993). Second order derivatives for network pruning: optimal brain surgeon. In *Adv. neural. inform. proc. sys.* (pp. 164–171).
- Hassibi, B., Stork, D.G., Wolff, G. (1993). Optimal brain surgeon and general network pruning. In *ICNN proc.*, Vol. 93 (pp. 293–299). IEEE Press.
- Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. Psychology Press.
- Hern, A. (2015). Yes, androids do dream of electric sheep. In *The guardian, technology and artificial intelligence*.
- Hinton, G. (2010). A practical guide to training restricted Boltzmann machines. Available at <http://learning.cs.toronto.edu>.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8), 2554–2558.
- Hopfield, J. J., Feinstein, D. I., & Palmer, R. G. (1983). Unlearning has a stabilizing effect in collective memories. *Nature Letters*, 304, 280158.
- Hopfield, J.J., & Tank, D. W. (1985). Neural computation of decisions in optimization problems. *Biological Cybernetics*, 52(3), 141–152.
- Huang, H. (2010). Reconstructing the Hopfield network as an inverse Ising problem. *Physical Review E*, 81(3), 036104.
- Huang, G. B., Saratchandran, P., & Sundararajan, N. (2005). A generalized growing and pruning RBF neural network for function approximation. *IEEE Transactions on Neural Networks and Learning Systems*, 16(1), 1045.
- Huang, H., & Toyozumi, T. (2015). Advanced mean-field theory of the restricted Boltzmann machine. *Physical Review E*, 91(5), 050101.
- Kappen, H. J., & Rodriguez, F. (1998). Efficient learning in Boltzmann machines using linear response theory. *Neural Computation*, 10(5), 1137–1156.
- Kirk, D. (2007). NVIDIA CUDA software and GPU parallel computing architecture. *Institute of Supply and Materials Management*, 7, 103.
- Kirkpatrick, S., et al. (1983). Optimization by simulated annealing. *Science*, 220, 671–680.
- Kleinfeld, D., & Pendergraft, D. B. (1987). Unlearning increases the storage capacity of content addressable memories. *Biophysical Journal*, 51, 47–53.
- Krotov, D., & Hopfield, J. J. (2016). Dense associative memory for pattern recognition. In *Adv. neur. inform. proc. sys.* (pp. 1172–1180).
- Le Cun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444.
- Le Cun, Y., Denker, J., & Solla, S. (1990). Optimal brain damage. In *Adv. neur. inform. proc. sys. 2* (pp. 598–605).
- Le Roux, N., & Bengio, Y. (2008). Representational power of restricted Boltzmann machines and deep belief networks. *Neural Computation*, 20(6), 1631–1649.
- Mandl, F., & Shaw, G. (2010). *Quantum field theory*. Manchester: John Wiley & Sons.
- McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of Learning and Motivation*, 24, 109–165.
- Mezard, M. (2017). Mean-field message-passing equations in the Hopfield model and its generalizations. *Physical Review E*, 95, 022117.
- Mezard, M., Parisi, G., & Virasoro, M. A. (1985). *Spin glass theory and beyond*. World Scientific.
- Nickolls, J., et al. (2008). Scalable parallel programming with CUDA. *Queue*, 6(2), 40–53.
- Nokura, K. (1996). Paramagnetic unlearning in neural network models. *Physical Review E*, 54(5), 5571.
- Nokura, K. (1998). Spin glass states of the anti-Hopfield model. *Journal of Physics A*, 31, 7447.
- Pastur, L., Shcherbina, M., & Tirozzi, B. (1994). The replica-symmetric solution without replica trick for the Hopfield model. *Journal of Statistical Physics*, 74(5), 1161–1183.
- Rolls, E. T., & Treves, A. (1998). *Neural networks and brain function*. Oxford University Press.
- Salakhutdinov, R., & Hinton, G. (2009). Deep Boltzmann machines. *Artificial Intelligence and Statistics*.
- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61, 85–117.
- Stickgold, R., Hobson, J. A., Fosse, R., & Fosse, M. (2001). Sleep, learning and dreams: Off-line memory reprocessing. *Science*, 294, 1052.

- Talagrand, M. (1998). Rigorous results for the Hopfield model with many patterns. *Probability Theory and Related Fields*, 110(2), 177–275.
- Talagrand, M. (2000). Exponential inequalities and convergence of moments in the replica-symmetric regime of the Hopfield model. *The Annals of Probability*, 1393–1469.
- Tramel, E. W., Dremeau, A., & Krzakala, F. (2016). Approximate message passing with restricted Boltzmann machine priors. *Journal of Statistical Mechanics: Theory and Experiment*, 073401.
- Tubiana, J., & Monasson, R. (2017). Emergence of compositional representations in restricted Boltzmann machines. *Physical Review Letters*, 118(13), 138301.