

Shaping the learning landscape in neural networks around wide flat minima

Carlo Baldassi,^{1,2} Fabrizio Pittorino,^{1,3} and Riccardo Zecchina^{1,4}

¹*Artificial Intelligence Lab, Institute for Data Science and Analytics, Bocconi University, Milano, Italy*

²*Istituto Nazionale di Fisica Nucleare, Sezione di Torino, Italy*

³*Politecnico di Torino, Italy*

⁴*International Centre for Theoretical Physics, Trieste, Italy*

Learning in Deep Neural Networks (DNN) takes place by minimizing a non-convex high-dimensional loss function, typically by a stochastic gradient descent (SGD) strategy. The learning process is observed to be able to find good minimizers without getting stuck in local critical points, and that such minimizers are often satisfactory at avoiding overfitting. How these two features can be kept under control in nonlinear devices composed of millions of tunable connections is a profound and far reaching open question. In this paper we study basic non-convex one- and two-layer neural network models which learn random patterns, and derive a number of basic geometrical and algorithmic features which suggest some answers. We first show that the error loss function presents few extremely wide flat minima (WFM) which coexist with narrower minima and critical points. We then show that the minimizers of the cross-entropy loss function overlap with the WFM of the error loss. We also show examples of learning devices for which WFM do not exist. From the algorithmic perspective we derive entropy driven greedy and message passing algorithms which focus their search on wide flat regions of minimizers. In the case of SGD and cross-entropy loss, we show that a slow reduction of the norm of the weights along the learning process also leads to WFM. We corroborate the results by a numerical study of the correlations between the volumes of the minimizers, their Hessian and their generalization performance on real data.

CONTENTS

I. Introduction	2
II. HLE/WFM regions exist in non convex neural devices storing random patterns	3
A. The simple example of discrete weights	3
B. Local entropy theory	4
C. Two layer networks with continuous weights	5
D. The existence of wide flat minima	6
E. Not all devices are appropriate: the Parity Machine does not display HLE/WFM regions	8
F. The connection between Local Entropy and Cross-Entropy	8
III. Belief Propagation and focusing Belief Propagation	11
IV. Entropic Least Action Learning	12
V. Numerical studies	12
Conclusions and future directions	15
Acknowledgments	16
References	16
A. High Local Entropy states from the 1-RSB formalism	17
B. Cross-Entropy minima, errors and high local entropy configurations	20
1. Cross-Entropy loss ground states	20
2. The local entropy around CE ground states	23
3. RS stability and zero entropy condition	27
C. Wide Flat Minima for the continuous case	28
1. Large deviation analysis of the Parity Machine	30
D. Belief Propagation on a Tree-like committee machine with continuous weights: equations, local volume around solutions and algorithms for the replicated network	30
1. BP equations for the committee machine	30
2. Exploring the space of solutions around a given configuration	33

3. Focusing-BP	34
E. Numerical experiments details	35
F. Experiments with randomized Fashion-MNIST	35

I. INTRODUCTION

Artificial Neural Networks (ANN), currently also known as Deep Neural Networks (DNN) when they have more than two layers, are powerful nonlinear devices used to perform different types of learning tasks (MacKay 2003). From the algorithmic perspective, learning in ANN is in principle a hard computational problem in which a huge number of parameters, the connection *weights*, need to be optimally tuned. Yet, at least for supervised pattern recognition tasks, learning has become a relatively feasible process in many applications across domains and the performances reached by DNNs have had a huge impact on the field of Machine Learning (ML).

DNN models have evolved very rapidly in the last decade, mainly by an empirical trial and selection process guided by heuristic intuitions. As a result, current DNN are in a sense akin to complex physical or biological systems, which are known to work but for which a detailed understanding of the principles underlying their functioning remains unclear. The tendency to learn efficiently and to generalize with limited overfitting are two properties that often coexist in DNN and yet a unifying theoretical framework is still missing.

Here we provide analytical results on the geometrical structure of the loss landscape of ANN which shed light on the success of Deep Learning (DL) (LeCun et al. 2015) algorithms and allow us to design novel efficient algorithmic schemes.

We focus on non-convex one- and two-layer ANN models that exhibit sufficiently complex behavior and yet are amenable to detailed analytical and numerical studies. Building on methods of statistical physics of disordered systems, we analyze the complete geometrical structure of the minimizers of the loss function of ANN learning random patterns and discuss how the current DNN models are able to exploit such structure, e.g. starting from the choice of the loss function, avoiding algorithmic traps and reaching rare solutions which belong to wide flat regions of the weight space. In our study the notion of flatness is given in terms of the volume of the weights around a minimizer which do not lead to an increase of the loss value. This generalizes the so called Local Entropy of a minimizer (Baldassi et al. 2015), defined for discrete weights as the log of the number of optimal weights assignments within a given Hamming distance from the reference minimizer. We call these regions High Local Entropy (HLE) regions for discrete weights or Wide Flat Minima (WFM) for continuous weights. Our results are derived analytically for the case of random data and corroborated by numerics on real data. In order to eliminate ambiguities which may arise from changes of scale of the weights, we control the norm of the weights in each of the units that compose the network. The outcomes of our study can be summarized as follows.

(i) We show analytically that ANN learning random patterns possess the structural property of having extremely robust regions of optimal weights, namely wide flat minima of the loss, whose existence is important to achieve convergence in the learning process. Though these wide minima are rare compared to the dominant critical points (absolute narrow minima, local minima or saddle points in the loss surface), they can be accessed by a large family of simple learning algorithms. We also show analytically that other learning machines, such as the Parity Machine, do not possess wide flat minima.

(ii) We show analytically that the choice of the cross-entropy loss function has the effect of biasing learning algorithms toward HLE or WFM regions.

(iii) We derive a greedy algorithm – Entropic Least Action Learning (eLAL) – and a message passing algorithm – focusing Belief Propagation (fBP) – which zoom in their search on wide flat regions of minimizers.

(iv) We compute the volumes associated to the minimizers found by different algorithms using Belief Propagation.

(v) We show numerically that the volumes correlate well with the spectra of the Hessian on computationally tractable networks and with the generalization performance on real data. The algorithms which search for WFM display a spectrum which is much more concentrated around zero eigenvalues compared to plain SGD.

Our results on random patterns support the conclusion that the minimizers which are relevant for learning are not the most frequent isolated and narrow ones (which also are computationally hard to sample) but the rare ones which are extremely wide. While this phenomenon was recently disclosed for the case of discrete weights (Baldassi et al. 2015, 2016a), here we demonstrate that it is present also in non convex ANN with continuous weights. Building on these results we derive novel algorithmic schemes and shed light on the performance of SGD with the cross-entropy loss function. Numerical experiments suggest that the scenario generalizes to real data and is consistent with other numerical results on deeper ANN (Keskar et al. 2016).

II. HLE/WFM REGIONS EXIST IN NON CONVEX NEURAL DEVICES STORING RANDOM PATTERNS

In what follows we analyze the geometrical structure of the weights space by considering the simplest non convex neural devices storing random patterns: the single layer network with discrete weights and the two layer networks with both continuous and discrete weights. The choice of random patterns, for which no generalization is possible, is motivated by the possibility of using analytical techniques from statistical physics of disordered systems and by the fact that we want to identify structural features which do not depend on specific correlation patterns of the data.

A. The simple example of discrete weights

In the case of binary weights it is well known that even for the single layer network the learning problem is computationally challenging. Therefore we begin our analysis by studying the so called binary perceptron, which maps vectors of N inputs $\xi \in \{-1, 1\}^N$ to binary outputs as $\sigma(W, \xi) = \text{sign}(W \cdot \xi)$, where $W \in \{-1, 1\}^N$ is the *synaptic* weights vector $W = (w_1, w_2, \dots, w_N)$.

Given a training set composed of αN input patterns ξ^μ with $\mu \in \{1, \dots, \alpha N\}$ and their corresponding desired outputs $\sigma^\mu \in \{-1, 1\}^{\alpha N}$, the learning problem consists in finding a *solution* W such that $\sigma(W, \xi^\mu) = \sigma^\mu$ for all μ . The entries ξ_i^μ and the outputs σ^μ are random unbiased i.i.d. variables. As discussed in [Krauth and Mézard \(1989\)](#) (but see also the rigorous bounds in [Ding and Sun \(2019\)](#)), perfect classification is possible with probability 1 in the limit of large N up to a critical value of α , usually denoted as α_c ; above this value, the probability of finding a solution drops to zero. α_c is called the *capacity* of the device.

The standard analysis of this model is based on the study of the zero temperature limit of the Gibbs measure with a loss (or energy) function \mathcal{L}_{NE} which counts the number of errors (NE) over the training set:

$$\mathcal{L}_{\text{NE}} = \sum_{\mu=1}^{\alpha N} \Theta(-\sigma^\mu \sigma(W, \xi^\mu)) \quad (1)$$

where $\Theta(x)$ is the Heaviside step function, $\Theta(x) = 1$ if $x > 0$ and 0 otherwise. The Gibbs measure is given by

$$P(W) = \frac{1}{Z(\beta)} \exp\left(-\beta \sum_{\mu=1}^{\alpha N} \Theta(-\sigma^\mu \sigma(W, \xi^\mu))\right) \quad (2)$$

where $\beta \geq 0$ is the inverse temperature parameter. For large values of β , $P(W)$ concentrates on the minima of \mathcal{L}_{NE} . The key analytical obstacle for the computation of $P(W)$ is the evaluation of the normalization factor, the partition function Z :

$$Z(\beta) = \sum_{\{w_i = \pm 1\}} \exp\left(-\beta \sum_{\mu=1}^{\alpha N} \Theta(-\sigma^\mu \sigma(W, \xi^\mu))\right) \quad (3)$$

In the the zero temperature limit ($\beta \rightarrow \infty$) and below α_c the partition function simply counts all solutions to the learning problem,

$$Z_\infty = \lim_{\beta \rightarrow \infty} Z(\beta) = \sum_{\{W\}} \mathbb{X}_\xi(W) \quad (4)$$

where $\mathbb{X}_\xi(W) = \prod_{\mu=1}^{\alpha N} \Theta(\sigma^\mu \sigma(W, \xi^\mu))$ is a characteristic function which evaluates to one if all patterns are correctly classified, and to zero otherwise.

Z_∞ is an exponentially fluctuating quantity (in N), and its most probable value is obtained by exponentiating the average of $\log Z_\infty$, denoted by $\langle \log Z_\infty \rangle_\xi$, over the realizations of the patterns

$$Z_{\infty, \text{typical}} \simeq \exp\left(N \langle \log Z_\infty \rangle_\xi\right). \quad (5)$$

The calculation of $\langle \log Z_\infty \rangle_\xi$ has been done in the 80s and 90s by the replica and the cavity methods of statistical physics and, as mentioned above, the results predict that the learning task undergoes a threshold phenomenon at $\alpha_c = 0.833$, where the probability of existence of a solution jumps from one to zero in the large N limit (Krauth and Mézard 1989). This result has been put recently on rigorous grounds by Ding and Sun (2019). Similar calculations predict that for any $\alpha \in (0, \alpha_c)$, the vast majority of the exponentially numerous solutions on the hypercube $W \in \{-1, 1\}^N$ are isolated, separated by a $O(N)$ Hamming mutual distance (Huang and Kabashima 2014). In the same range of α , there also exist an even larger number of local minima at non-zero loss, a result that has been corroborated by analytical and numerical findings on stochastic learning algorithms which satisfy detailed balance (Horner 1992). Recently it became clear that by relaxing the detailed balance condition it was possible to design simple algorithms which can solve the problem efficiently (Baldassi 2009, Baldassi et al. 2007, Braunstein and Zecchina 2006).

B. Local entropy theory

The existence of effective learning algorithms indicates that the traditional statistical physics calculations, which focus on set of solutions that dominate the zero temperature Gibbs measure (i.e. the most numerous ones), are effectively blind to the solutions actually found by such algorithms. Numerical evidence suggests that in fact the solutions found by heuristics are not at all isolated: on the contrary, they appear to belong to regions with a high density of nearby other solutions. This puzzle has been solved very recently by an appropriate large deviations study (Baldassi et al. 2015, 2016a,b,c) in which the tools of statistical physics have been used to study the most probable value of the *local entropy* of the loss function, i.e. a function that is able to detect the existence of regions with an $O(N)$ radius containing a high density of solutions even when the number of these regions is small compared to the number of isolated solutions. For binary weights the local entropy function is the (normalized) logarithm of the number of solutions W' at Hamming distance $D N$ from a reference solution W

$$\mathcal{E}_D(W) = -\frac{1}{N} \ln \mathcal{U}(W, D) \quad (6)$$

with

$$\mathcal{U}(W, D) = \sum_{\{W'\}} \mathbb{X}_\xi(W') \delta(W' \cdot W, N(1 - 2D)) \quad (7)$$

and where δ is the Kronecker delta symbol. In order to derive the typical values that the local entropy can take, one needs to compute the Gibbs measure of the local entropy

$$P_{\text{LE}}(W) = \frac{1}{Z_{\text{LE}}} \exp(-y \mathcal{E}_D(W)) \quad (8)$$

where y has the role of an inverse temperature. For large values of y this probability measure focuses on the W surrounded by an exponential number of solutions within a distance D . The regions of high local entropy (HLE) are then described in the regime of large y and small D . In particular, the calculation of the expected value of the optimal *local entropy*

$$\mathcal{S}(D) \equiv \mathcal{E}_D^{\text{opt}} = \max_{\{W\}} \left\{ -\frac{1}{N} \langle \ln \mathcal{U}(W, D) \rangle_\xi \right\} \quad (9)$$

shows the existence of extremely dense clusters of solutions up to values of α close to α_c (Baldassi et al. 2015, 2016a,b,c).

The probability measure eq. (8) can be written in an equivalent form that generalizes to the non zero errors regime, is analytically simpler to handle, and leads to novel algorithmic schemes (Baldassi et al. 2016a):

$$P_{\text{LE}}(W) \sim P(W; \beta, y, \lambda) = Z(\beta, y, \lambda)^{-1} e^{y \Phi(W, \beta, \lambda)} \quad (10)$$

where $\Phi(W, \beta, \lambda)$ is a “local free entropy” potential in which the distance constraint is forced through a Lagrange multiplier λ

$$\Phi(W, \beta, \lambda) = \ln \sum_{\{W'\}} e^{-\beta \mathcal{L}_{\text{NE}}(W') - \lambda d(W, W')} \quad (11)$$

where $d(\cdot, \cdot)$ is some monotonically increasing function of the distance between configurations, defined according to the type of weights under consideration. In the limit $\beta \rightarrow \infty$ and by choosing λ so that a given distance is selected, this expression reduces to eq. (8).

The crucial property of eq. (10) comes from the observation that by choosing y to be a non-negative integer, the partition function can be rewritten as:

$$\begin{aligned} Z(\beta, y, \lambda) &= \sum_{\{W\}} e^{y\Phi(W, \beta, \lambda)} \\ &= \sum_{\{W\}} \sum_{\{W'^a\}_{a=1}^y} e^{-\beta\mathcal{L}_R(W, W'^a)} \end{aligned} \quad (12)$$

where

$$\mathcal{L}_R(W, W'^a) = \sum_{a=1}^y \mathcal{L}_{\text{NE}}(W'^a) - \frac{\lambda}{\beta} \sum_{a=1}^y d(W, W'^a). \quad (13)$$

These are the partition function and the effective loss of $y + 1$ interacting real replicas of the system, one of which acts as reference system (W) while the remaining y ($\{W'^a\}$) are identical, each being subject to the energy constraint $\mathcal{L}_{\text{NE}}(W'^a)$ and to the interaction term with the reference system. As discussed in Baldassi et al. (2016a), several algorithmic schemes can be derived from this framework by minimizing \mathcal{L}_R . Here we shall also use the above approach to study the existence of WFMs in continuous models and to design message passing and greedy learning algorithms driven by the local entropy of the solutions.

C. Two layer networks with continuous weights

As for the discrete case, we are able to show that in non convex networks with continuous weights the WFMs exist, are rare and yet accessible to simple algorithms. In order to perform an analytic study, we consider the simplest non-trivial two-layer neural network, the committee machine with non-overlapping receptive fields. It consists of N input units, one hidden layer with K units and one output unit. The input units are divided into K disjoint sets of $\tilde{N} = \frac{N}{K}$ units. Each set is connected to a different hidden unit. The input to the ℓ -th hidden unit is given by $x_\ell^\mu = \frac{1}{\sqrt{\tilde{N}}} \sum_{i=1}^{\tilde{N}} w_{\ell i} \xi_{\ell i}^\mu$ where $w_{\ell i} \in \mathbb{R}$ is the connection weight between the input unit i and the hidden unit ℓ , and $\xi_{\ell i}^\mu$ is the i -th input to the ℓ -th hidden unit. As before, μ is a pattern index. We study analytically the pure classifier case in which each unit implements a threshold transfer function and the loss function is the error loss. Other types of (smooth) functions, more amenable to numerical simulation, will be also discussed in a subsequent section. The output of the ℓ -th hidden unit is given by

$$\tau_\ell^\mu = \text{sign}(x_\ell^\mu) = \text{sign}\left(\frac{1}{\sqrt{\tilde{N}}} \sum_{i=1}^{\tilde{N}} w_{\ell i} \xi_{\ell i}^\mu\right) \quad (14)$$

In the second layer all the weights are fixed and equal to one, and the overall output of the network is simply given by a majority vote $\sigma_{\text{out}}^\mu = \text{sign}\left(\frac{1}{\sqrt{K}} \sum_\ell \tau_\ell^\mu\right)$.

As for the binary perceptron, the learning problem consists in mapping each of the random input patterns ($\xi_{\ell i}^\mu$) with $(\ell = 1, \dots, K, i = 1, \dots, \tilde{N}, \mu = 1, \dots, \alpha N)$ onto a randomly chosen output σ^μ . Both $\xi_{\ell i}^\mu$ and σ^μ are independent random variables which take the values ± 1 with equal probability. For a given set of patterns, the volume of the subspace of the network weights which correctly classifies the patterns, the so called *version space*, is given by

$$V = \int \prod_{i\ell} dw_{\ell i} \prod_\ell \delta\left(\sum_i w_{\ell i}^2 - \tilde{N}\right) \prod_\mu \Theta(\sigma^\mu \sigma_{\text{out}}^\mu) \quad (15)$$

where we have imposed a spherical constraint on the weights via a Dirac δ in order to keep the volume finite (though exponential in N). In the case of binary weights the integral would become a sum over all the 2^N configurations and the volume would be the overall number of zero error assignments of the weights.

The committee machine has been studied extensively in the 90s (Barkai et al. 1992, Engel et al. 1992, Schwarze and Hertz 1992). The capacity of the network can be derived by computing the typical weight space volume as a function of the number of correctly classified patterns αN , in the large N limit. As for the binary case, the most probable value of V is obtained by exponentiating the average of $\log V$, $V_{\text{typical}} \simeq \exp\left(N \langle \log V \rangle_\xi\right)$, a difficult task which is achieved by the replica method (Barkai et al. 1990, Mézard et al. 1987).

For the smallest non-trivial value of K , $K = 3$, it has been found that above $\alpha_0 \simeq 1.76$ the space of solutions changes abruptly, becoming clustered into multiple components¹. Below α_0 the geometrical structure is not clustered and can be described by the simplest version of the replica method, known as replica symmetric (RS) solution. Above α_0 the analytical computation of the typical volume requires a more sophisticated analysis that properly describes a clustered geometrical structure. This analysis can be performed by a variational technique which is known in statistical physics as the replica-symmetry-breaking (RSB) scheme, and the clustered geometrical structure of the solution space is known as RSB phase.

The capacity of the network, above which perfect classification becomes impossible, is found to be $\alpha_c \simeq 3.02$. In the limit of large K (but still with $\tilde{N} \gg 1$), the clustering transition occurs at a finite number of patterns per weight, $\alpha_0 \simeq 2.95$ (Barkai et al. 1992) whereas the critical capacity grows with K as $\alpha_c \propto \sqrt{\ln K}$ (Monasson and Zecchina 1995).

D. The existence of wide flat minima

In order to detect the existence of WFM we use a large deviation measure which is the continuous version of the measure used in the discrete case: each configuration of the weights is re-weighted by a local volume term, analogously to the analysis of section II B. For the continuous case, however, we adopt a slightly different formalism which simplifies the analysis. Instead of constraining the set of y real replicas² to be at distance D from a reference weight vector, we can identify the same WFM regions by constraining them directly to be at a given mutual overlap: for a given value $q_1 \in [-1, 1]$, we impose that $W^a \cdot W^b = Nq_1$ for all pairs of distinct replicas a, b . The overlap q_1 is bijectively related to the mutual distance among replicas (which tends to 0 as $q_1 \rightarrow 1$). That, in turn, determines the distance between each replica and the barycenter of the group $\frac{1}{y} \sum_a W^a$, which takes the role that the reference vector had in the previous treatment. Thus, the regime of small D corresponds to the regime of q_1 close to 1, and apart from this reparametrization the interpretation of the WFM volumes is the same. As explained in Appendix C, the advantage of this novel technique is that it allows to use directly the first-step formalism of the RSB scheme (1-RSB). Similarly to the discrete case, the computation of the maximal WFM volumes leads to the following results: for $K = 3$ and in the large y limit, we find³

$$\begin{aligned} \mathcal{V}(q_1) &= \max_{\{W^a\}_{a=1}^y} \left\{ \frac{1}{Ny} \langle \ln V((W^a)_{a=1}^y, q_1) \rangle_\xi \right\} \\ &= G_S(q_1) + \alpha G_E(q_1) \end{aligned}$$

with

$$\begin{aligned} G_S(q_1) &= \frac{1}{2} [1 + \ln 2\pi + \ln(1 - q_1)] \\ G_E(q_1) &= \int \prod_{\ell=1}^3 Dv_\ell \max_{u_1, u_2, u_3} \left[-\frac{\sum_{\ell=1}^3 u_\ell^2}{2} + \ln \left(\tilde{H}_1 \tilde{H}_2 + \tilde{H}_1 \tilde{H}_3 + \tilde{H}_2 \tilde{H}_3 - 2\tilde{H}_1 \tilde{H}_2 \tilde{H}_3 \right) \right] \end{aligned}$$

where $\tilde{H}_\ell \equiv H\left(\sqrt{\frac{d_0}{1-q_1}} u_\ell + \sqrt{\frac{q_0}{1-q_1}} v_\ell\right)$, $H(x) \equiv \int_x^\infty Dv$, $Dv \equiv dv \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}v^2}$, $d_0 \equiv y(q_1 - q_0)$ and q_0 satisfies a saddle point equation that needs to be solved numerically. $\mathcal{V}(q_1)$ is the logarithm of the volume of the solutions, normalized by Ny , under the spherical constraints on the weights, and with the real replicas forced to be at a mutual overlap q_1 . In analogy with the discrete case, we still refer to $\mathcal{V}(q_1)$ as to the *local entropy*. It is composed by the sum of two terms: the first one, $G_S(q_1)$, corresponds to the log-volume at $\alpha = 0$, where all configurations are solutions and only the geometric constraints are present. This is an upper bound for the local entropy. The second term,

¹ Strictly speaking each cluster is composed of a multitude of exponentially small domains (Monasson and Zecchina 1995).

² Not to be confused with the virtual replicas of the replica method.

³ All the details are reported in Appendix C.

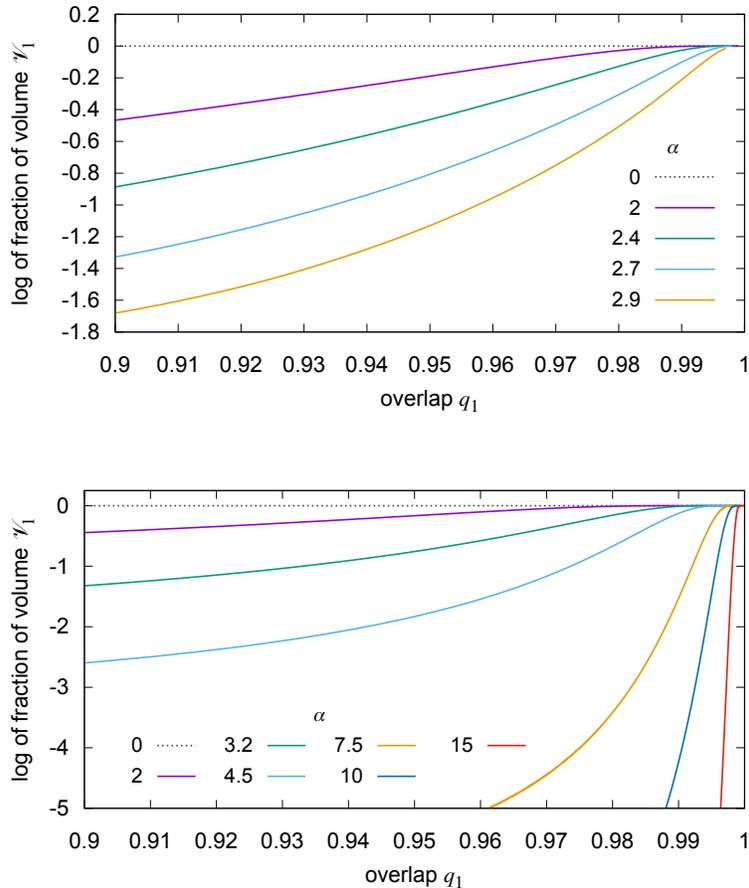


Figure 1. Normalized local entropy \mathcal{V}_1 vs q_1 , i.e. logarithm of the fraction of configurations of y real replicas at mutual overlap q_1 in which all replicas have zero error loss \mathcal{L}_{NE} . The curves are for a tree-like committee machine trained on αN random patterns, for various values of α , obtained from a replica calculation in the limit of large N and large number of real replicas y . When the curves approach 0 as $q_1 \rightarrow 1$ it means that nearly all configurations are minima, thus that the replicas are collectively exploring a wide minimum (any $q_1 < 1$ implies distances of $O(N)$ between replicas). Top: the case of $K = 3$ units. Bottom: the limiting case of a large number of hidden units K .

$\mathcal{V}_1(q_1) \equiv \alpha G_E(q_1)$, is in general negative and it represents the log of the *fraction* of solutions at overlap q_1 (among all configurations at that overlap), and we call it *normalized* local entropy. Following the interpretation given above, we expect (in analogy to the discrete case at small D , cf. fig. 3) that in a WFM this fraction is close to 1 (i.e. \mathcal{V}_1 is close to 0) in an extended region where the distance between the replicas is small, i.e. where q_1 is close to 1; otherwise, WFMs do not exist in the model. In fig. 1 (top panel) we report the values of $\mathcal{V}_1(q_1)$ vs the overlap q_1 , for different values of α . Indeed, one may observe that the behavior is qualitatively similar to that of the binary perceptron: besides the solutions and all the related local minima and saddles predicted by the standard statistical physics analysis (Barkai et al. 1992, Engel et al. 1992, Monasson and Zecchina 1995, Schwarze and Hertz 1992), there exist absolute minima which are flat at relatively large distances. Indeed, reaching such wide minima efficiently is non trivial, and different algorithms can have drastically different behaviors, as we will discuss in detail in sec. V.

The case $K = 3$ is still relatively close to the simple perceptron, though the geometrical structure of its minima is already dominated by non convex features for $\alpha > 1.76$. A case which is closer to more realistic ANNs is $K \gg 1$ (but still $N \gg K$), which, luckily enough, is easier to study analytically. We find:

$$G_S(q_1) = \frac{1}{2} [1 + \ln 2\pi + \ln(1 - q_1)]$$

$$G_E(q_1) = \int Dv \max_u \left[-\frac{u^2}{2} + \ln H \left(\sqrt{\frac{\Delta q_1^e}{1 - q_1^e}} u + \frac{q_0^e}{1 - q_1^e} v \right) \right]$$

where $\Delta q_1^e = q_1^e - q_0^e$ with $q_1^e \equiv 1 - \frac{2}{\pi} \arccos(q_1)$, $q_0^e \equiv 1 - \frac{2}{\pi} \arccos(q_0)$ and q_0 is fixed by a saddle point equation. In fig. 1 (bottom panel) we observe that indeed WFM still exist for all finite values of α .

The results of the above WFM computation may require small corrections due to RSB effects, which however are expected to be very tiny due to the compact nature of the space of solutions at small distances.

A more informative aspect is to study the volumes around the solutions found by different algorithms. This can be done by the Belief Propagation method, similarly to the computation of the weight enumerator function in error correcting codes (Di et al. 2006).

E. Not all devices are appropriate: the Parity Machine does not display HLE/WFM regions

The extent by which a given model exhibits the presence of WFM can vary (see e.g. fig. 1). A direct comparison of the local entropy curves on different models in general does not yet have a well-defined interpretation, although at least for similar architectures it can still be informative Baldassi et al. (2019). On the other hand, the existence of WFM itself is a structural property. For neural networks, its origin relies in the threshold sum form of the nonlinearity characterizing the formal neurons. As a check of this claim, we can analyze a model which is in some sense complementary, namely the so called *parity machine*. We take its network structure to be identical to the committee machine, except for the output unit which performs the product of the K hidden units instead of taking a majority vote. While the outputs of the hidden units are still given by sign activations, eq. (14), the overall output of the network reads $\sigma_{\text{out}}^\mu = \prod_{\ell=1}^K \tau_\ell$. The volume of the weights that correctly classifies a set of patterns is still given by eq. (15).

Parity machines are closely related to error correcting codes based on parity checks. The geometrical structure of the absolute minima of the error loss function is known (Monasson and Zecchina 1995) to be composed by multiple regions, each in one to one correspondence with the internal representations of the patterns. For random patterns such regions are typically tiny and we expect the WFM to be absent. Indeed, the computation of the volume proceeds analogously to the previous case⁴, and it shows that in this case for any distance the volumes of the minima are always bounded away from the maximal possible volume, i.e. the volume one would find for the same distance when no patterns are stored: the log-ratio of the two volumes is constant and equal to $-\alpha \log(2)$. In other words, the minima never become flat, at any distance scale.

F. The connection between Local Entropy and Cross-Entropy

Given that dense regions of optimal solutions exist in non convex ANN, at least in the case of independent random patterns, it remains to be seen which role they play in current models. Starting with the case of binary weights, and then generalizing the result to more complex architectures and to continuous weights, we can show that the most widely used loss function, the so called Cross-Entropy (CE) loss, focuses precisely on such rare regions (see Baldassi et al. (2018) for the case of stochastic weights).

For the sake of simplicity, we consider a binary classification task with one output unit. The CE cost function for each input pattern reads

$$\mathcal{L}_{\text{CE}}(W) = \sum_{\mu=1}^M f_\gamma \left(\frac{\sigma^\mu}{\sqrt{N}} \sum_{i=1}^N w_i \xi_i^\mu \right) \quad (16)$$

where $f_\gamma(x) = -\frac{x}{2} + \frac{1}{2\gamma} \log(2 \cosh(\gamma x))$. The parameter γ allows to control the degree of “robustness” of the training, see the inset in figure 2. In standard machine learning practice γ is simply set to 1, but a global rescaling of the

⁴ It's actually even simpler, see Appendix C 1.

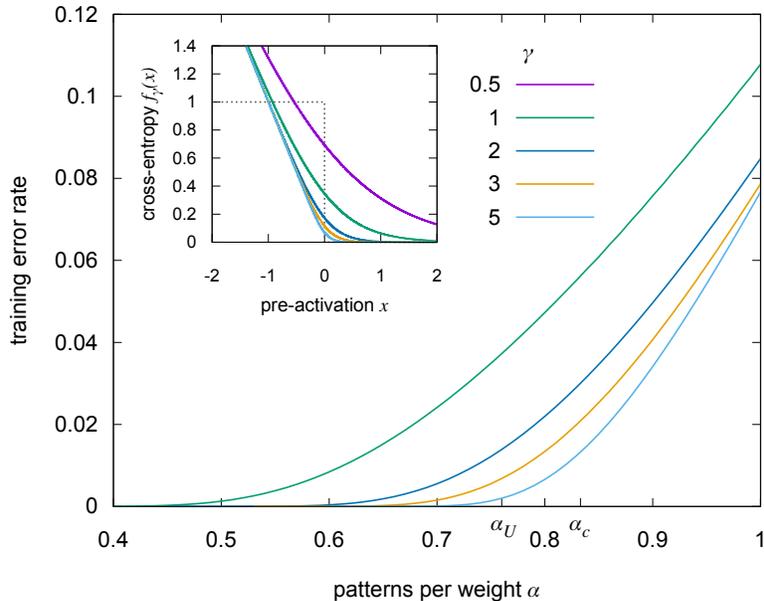


Figure 2. Mean error rate achieved when optimizing the CE loss in the binary single-layer network, as predicted by the replica analysis, at various values of γ (increasing from top to bottom). The figure also shows the points $\alpha_c \approx 0.83$ (up to which solutions exist) and $\alpha_U \approx 0.76$ (up to which non-isolated solutions exist). *Inset:* Binary cross-entropy function $f_\gamma(x)$ for various values of γ (increasing from top to bottom). For low values of γ , the loss is non-zero even for small positive values of the input, and thus the minimization procedure tends to favor more robust solutions. For large values of γ the function tends to $\max(-x, 0)$. The dotted line shows the corresponding NE function, which is just 1 in case of an error and 0 otherwise, cf. eq. (1).

weights W_i can lead to a basically equivalent effect. That setting can thus be interpreted as leaving γ as implicit, letting its effective value, and hence the norm of the weights, to be determined by the initial conditions and the training algorithm. As we shall see, controlling γ explicitly along the learning process plays a crucial role in finding HLE/WFM regions.

For the binary case, however, the norm is fixed and thus we must keep γ as an explicit parameter. Note that since $\lim_{\gamma \rightarrow \infty} f_\gamma(x) = \max(-x, 0)$ the minima of \mathcal{L}_{CE} below α_c at large γ are the solutions to the training problem, i.e. they coincide with those of \mathcal{L}_{NE} .

We proceed by first showing that the minimizers of this loss correspond to near-zero errors for a wide range of values of α , and then by showing that these minimizers are surrounded by an exponential number of zero error solutions.

In order to study the probability distribution of the minima of \mathcal{L}_{CE} in the large N limit, we need to compute its Gibbs distribution (in particular, the average of the log of the partition function, see eq. (5)) as it has been done for the error loss \mathcal{L}_{NE} . The procedure follows standard steps and it is detailed in Appendix B. The method requires to solve two coupled integral equations as functions of the control parameters α , β and γ . In figure 2 we show the behavior of the fraction of errors vs the loading α , for various values of γ . Up to relatively large values of α the optimum of \mathcal{L}_{CE} corresponds to extremely small values of \mathcal{L}_{NE} , virtually equal to zero for any accessible size N .

Having established that by minimizing the CE one ends up in regions of perfect classification where the error loss function is essentially zero, it remains to be understood which type of configurations of weights are found. Does the CE converge to an isolated point-like solution in the weight space (such as the typical zero energy configurations of the error function)⁵ or does it converge to the rare regions of high local entropy?

In order to establish the geometrical nature of the typical minima of the CE loss, we need to compute the average value of $\mathcal{E}_D(W)$ (which tells us how many zero energy configurations of the error loss function can be found within a given radius D from a given W , see eq. (6)) when W is sampled from the minima of \mathcal{L}_{CE} . This can be accomplished by a well known analytical technique (Franz and Parisi 1995) which was developed for the study of the energy landscape in disordered physical systems. The computation is relatively involved, and here we report only the final outcome. For the dedicated reader, all the details of the calculation, which relies on the replica method and includes a double

⁵ A quite unlikely fact given that finding isolated solutions is a well known intractable problem.

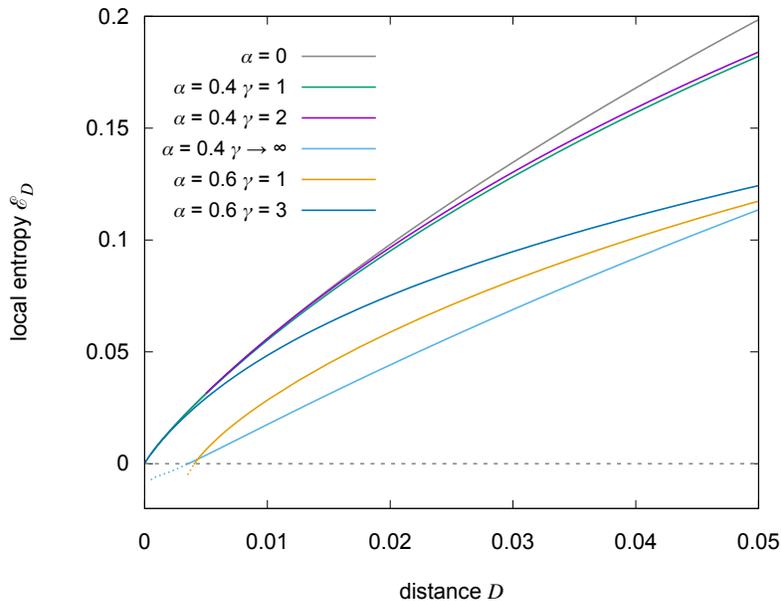


Figure 3. Average local entropy around a typical minimum of the \mathcal{L}_{CE} loss, for various values of α and γ . The gray upper curve, corresponding to $\alpha = 0$, is an upper bound since in that case all configurations are solutions. For $\alpha = 0.4$, the two curves with $\gamma = 1$ and 2 nearly saturate the upper bound at small distances, revealing the presence of dense regions of solutions (HLE regions). There is a slight improvement for $\gamma = 2$, but the curve at $\gamma \rightarrow \infty$ shows that the improvement cannot be monotonic: in that limit, the measure is dominated by isolated solutions. This is reflected by the gap at small d in which the entropy becomes negative, signifying the absence of solutions in that range. For $\alpha = 0.6$ we see that the curves are lower, as expected. We also see that for $\gamma = 1$ there is a gap at small d , and that we need to get to $\gamma = 3$ in order to find HLE regions.

analytic continuation, can be found in Appendix B. As reported in figure 3, we find that the typical minima of the CE loss for small finite γ are indeed surrounded by an exponential number of zero error solutions. In other words, the CE focuses on HLE regions.

It is clear from the figure that γ needs to be sufficiently large for this phenomenon to occur. On the other hand, in the limit $\gamma \rightarrow \infty$ the space of solutions is again dominated by the isolated ones; we thus expect the existence of an optimal value of γ , depending on the parameters α and D . We set $\alpha = 0.4$ and used two values of D , 0.005 and 0.02, and measured the normalized local entropy $\mathcal{E}_D - \mathcal{E}_D^{\text{max}}$ where $\mathcal{E}_D^{\text{max}} = -D \log D - (1 - D) \log (1 - D)$ is the upper bound corresponding to the case $\alpha = 0$ (gray curve in fig. 3). The results are shown in fig. 4. Numerical issues prevented us from reaching the optimal γ at $D = 0.005$ (the main plot, in log scale, shows that the curve is still growing), and the left inset (same data as the main plot, but not in log scale) shows that there is a whole region where the local entropy is extremely close to optimal (what we would call a dense region). At a larger distance, $D = 0.02$, we could find the optimum (denoted with a dot) at a lower local entropy (consistently with fig. 3), but it is also clear that the “good” region of γ is rather large since the curve is very flat. In the limit of $\gamma \rightarrow \infty$ the curves would tend to -0.054 and -0.028 for the $D = 0.02$ and $D = 0.005$ cases, respectively.

The physical significance of these “optimal γ regions” and their relation with the local geometry of the landscape is not obvious. The main reason for this is that a fine description of the local geometry of the HLE regions is still an open problem – circumstantial evidence from theoretical considerations and numerical simulations suggests that they are rather complex structures (Baldassi et al. 2015), possibly with a multifractal nature. At the present time, this is pure speculation. On the practical side, however, a local search algorithm that explores the landscape at low γ and gradually increases it should eventually reach this “good γ ” plateau that leads it toward a HLE region; further increasing γ would then be unlikely to drive it out, unless very strong drift terms are present (Baldassi et al. 2016a).

As an algorithmic check we have verified that while a Simulated Annealing approach gets stuck at very high energies when trying to minimize the error loss function, the very same algorithm with the CE loss is indeed successful up to relatively high values of α , with just slightly worse performance compared to an analogous procedure based on local entropy (Baldassi et al. 2016c). In other words, the CE loss on single-layer networks is a computationally cheap and reasonably good proxy for the LE loss. These analytical results extend straightforwardly to two layer networks with binary weights. The study of continuous weight models can be performed resorting to the Belief Propagation method.

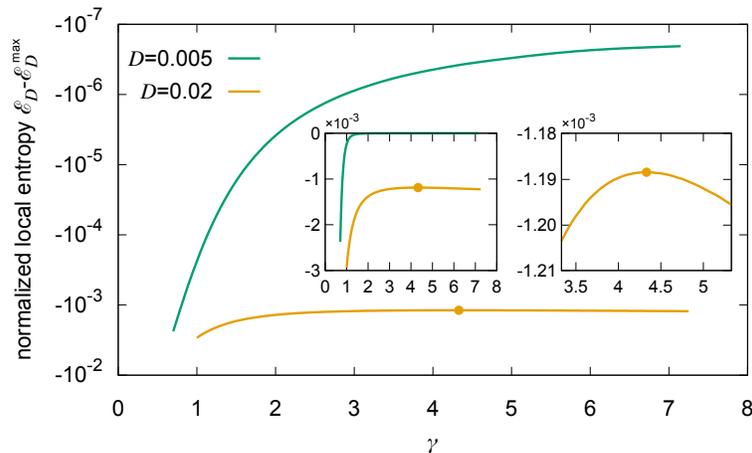


Figure 4. Normalized local entropy as a function of γ , for a binary perceptron trained with the CE loss at $\alpha = 0.4$ and two values of the distance D . The main plot shows the results in logarithmic scale, the left inset is the same plot in linear scale. For $D = 0.005$ we could not find the optimal γ due to numerical issues; for $D = 0.02$ the optimum is shown with a circle, although the curve looks rather flat in an extended region: the right inset shows an enlargement that makes the appearance of a maximum evident.

III. BELIEF PROPAGATION AND FOCUSING BELIEF PROPAGATION

Belief Propagation (BP), also known as sum-product, is an iterative message-passing algorithm for statistical inference. When applied to the problem of training a committee machine with a given set of input-output patterns, it can be used to obtain, at convergence, useful information on the probability distribution, over the weights of the network, induced by the Gibbs measure. In particular, it allows to compute the marginals of the weights as well as their entropy, which in the zero-temperature regime is simply the logarithm of the volume of the solutions, eq. (15), rescaled by the number of variables N . The results are approximate, but (with high probability) they approach the correct value in the limit of large N in the case of random uncorrelated inputs, at least in the replica-symmetric phase of the space of the parameters. Due to the concentration property, in this limit the macroscopic properties of any given problem (such as the entropy) tend to converge to a common limiting case, and therefore a limited amount of experiments with a few samples is sufficient to describe very well the entire statistical ensemble.

We have used BP to study the case of the zero-temperature tree-like committee machine with continuous weights and $K = 3$. We have mostly used $N = 999$, which turns out to be large enough to produce results in quite good agreement with the replica theory analysis. The implementation can be made efficient by encoding each message with only two quantities (see Appendix D). As mentioned above, this algorithm works well in the replica-symmetric phase, which for our case means when $\alpha \leq \alpha_0 \approx 1.76$. Above this value, the (vanilla) algorithm doesn't converge at all.

However, BP can be employed to perform additional analyses as well. In particular, it can be modified rather straightforwardly to explore and describe the region surrounding any given configuration, as it allows to compute the local entropy (i.e. the log-volume of the solutions) for any given distance and any reference configuration (this is a general technique, the details for our case are reported in Appendix D 2). The convergence issues are generally much less severe in this case. Even in the RSB phase, if the reference configuration is a solution in a wide minimum, the structure is locally replica-symmetric, and therefore the algorithm converges and provides accurate results, at least up to a value of the distance where other unconnected regions of the solutions space come into consideration. In our tests, the only other issue arose occasionally at very small distances, where convergence is instead prevented by the loss of accuracy stemming from finite size effects and limited numerical precision.

Additionally, the standard BP algorithm can be modified and transformed into a (very effective) solver. There are several ways to do this, most of which are purely heuristic. However, it was shown in Baldassi et al. (2016a) that adding a particular set of self-interactions to the weight variables could approximately but effectively describe the replicated system of eq. (12): in other words, this technique can be used to analyze the local-entropy landscape instead of the Gibbs one. By using a sufficiently large number of replicas y (we generally used $y = 10$) and following an annealing protocol in the coupling parameter λ (starting from a low value and making it diverge) this algorithm focuses on the maximally dense regions of solutions, thus ending up in wide flat minima. For these reasons, the algorithm was called “focusing-BP” (fBP). The implementation closely follows that of Baldassi et al. (2016a) (complete details are provided in Appendix D 3). Our tests – detailed below – show that this algorithm is the best solver (by a wide margin) among

the several alternatives that we tried in terms of robustness of the minima found (and thus of generalization properties, as also discussed below). Moreover, it also achieves the highest capacity, nearly reaching the critical capacity where all solutions disappear.

IV. ENTROPIC LEAST ACTION LEARNING

Least Action Learning (LAL) (Mitchison and Durbin 1989) is a heuristic greedy algorithm that was designed to extend the well-known Perceptron algorithm to the case of committee-machines with a single binary output and sign activation functions. It takes one parameter, the learning rate η . In its original version, patterns are presented randomly one at a time, and at most one hidden unit is affected at a time. In case of correct output, nothing is done, while in case of error the hidden unit, among those with a wrong output, whose pre-activation was closest to the threshold (and is thus the easiest to fix) is selected, and the standard perceptron learning rule (with rate η) is applied to it. In our tests we simply extended it to work in mini-batches, to make it more directly comparable with stochastic-gradient-based algorithms: for a given mini-batch, we first compute all the pre-activations and the outputs for all patterns at the current value of the weights, then we apply the LAL learning rule for each pattern in turn.

This algorithm proves to be surprisingly effective at finding minima of the NE loss very quickly: in the random patterns case, its algorithmic capacity is higher than gradient-based variants and almost as high as fBP, and it requires comparatively few epochs. It is also computationally very fast, owing to its simplicity. However, as we show in the sec. V, it finds solutions that are much narrower compared to those of other algorithms.

In order to drive LAL toward WFM regions, we add a local-entropy component to it, by applying the technique described in Baldassi et al. (2016a) (see eq. (13)): we run y replicas of the system in parallel and we couple them with an elastic interaction. The resulting algorithm, that we call entropic-LAL (eLAL) can be described as follows. We initialize y replicas randomly with weights W^a and compute their average \tilde{W} . We present mini-batches independently to each replica, using different permutations of the dataset for each of them. At each mini-batch, we apply the LAL learning rule. Then, each replica is pushed toward the group average with some strength proportional to a parameter λ : more precisely, we add a term $\lambda\eta(\tilde{W} - W^a)$ to each of the weight vectors W^a . After this update, we recompute the average \tilde{W} . At each epoch, we increase the interaction strength λ . The algorithm stops when the replicas have collapsed to a single configuration.

This simple scheme proves rather effective at enhancing the wideness of the minima found while still being computationally efficient and converging quickly, as we show in the sec. V. We show tests performed with $y = 20$, but the entropic enhancement is gradual, already providing improvements with very few replicas and progressing at least up to $y = 100$, which is the maximum value that we have tested. Its main limitation is, of course, that it is tailored to committee machines and it is unclear how to extend it to general architectures. On the other hand, the effectiveness of the replication scheme seems very promising from the point of view of improving the quality of the minima found by greedy and efficient heuristics.

V. NUMERICAL STUDIES

We conclude our study by comparing numerically the curvature, the wideness of the minima and the generalization error found by different approaches. We consider two main scenarios: one, directly comparable with the theoretical calculations, where a tree committee machine with $K = 9$ is trained over random binary patterns, and a second one, which allows us to estimate the generalization capabilities, where a fully-connected committee machine with $K = 9$ is trained on a subset of the Fashion-MNIST dataset (Xiao et al. 2017). The choice of using $K = 9$ instead of 3 is intended to enhance the potential robustness effect that the CE loss can have over NE on such architectures (see the inset of fig. 2): for $K = 3$, a correctly classified pattern already requires 2 out for 3 units to give the correct answer, and there is not much room for improvement at the level of the pre-activation of the output unit. On the other hand, since we study networks with a number of inputs of the order of 10^3 , an even larger value of K would either make N/K too small in the tree-like case (exacerbating numerical issues for the BP algorithms and straying too far from the theoretical analysis) or make the computation of the Hessians too onerous for the fully-connected case (each Hessian requiring the computation of $(NK)^2$ terms).

We compare several training algorithms with different settings (see Appendix E): stochastic GD with the CE loss (ceSGD); least-action learning (LAL) and its entropic version (eLAL); focusing BP (fBP). Of these, the non-gradient based ones (LAL, eLAL and fBP) can be directly used with the sign activation functions (14) and the NE loss. On the other hand, ceSGD requires a smooth loss landscape, therefore we used tanh activations, adding a gradually-diverging parameter β in their argument, since $\lim_{\beta \rightarrow \infty} \tanh(\beta x) = \text{sign}(x)$. The γ parameter of the CE loss (16) was also

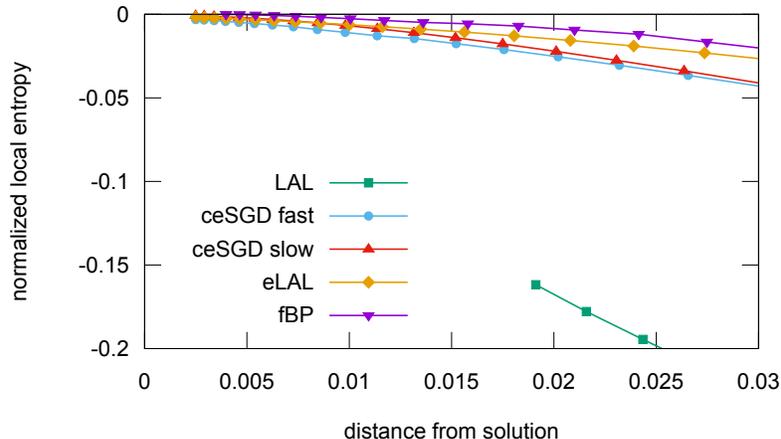


Figure 5. Normalized local entropy as a function of the distance from a reference solution, on a tree-like committee machine with $K = 9$ and $N = 999$, trained on 1000 random patterns. The results were obtained with the BP algorithm, by averaging over 10 samples. Numerical issues (mainly due to the approximations used) prevented BP from converging at small distances for the LAL algorithm, and additionally they slightly affect the results at very small distances. Qualitatively, though, higher curves correspond to larger local entropies and thus wider minima.

increased gradually. As in the theoretical computation, we also constrained the weights of each hidden unit of the network to be normalized. The NE loss with sign activations is invariant under renormalization of each unit’s weights, whereas the CE loss with tanh activations is not. In the latter case, the parameters β and γ can be directly interpreted as the norm of the weights, since they just multiply the pre-activations of the units. In a more standard approach, the norm would be controlled by the initial choice of the weights and be driven by the SGD algorithm automatically. In our tests instead we have controlled these parameters explicitly, which allows us to demonstrate the effect of different schedules. In particular, we show (for both the random and the Fashion-MNIST scenarios) that slowing down the growth of the norm with ceSGD makes a significant difference in the quality of the minima that are reached. We do this by using two separate settings for ceSGD, a “fast” and a “slow” one. In ceSGD-fast both β and γ are large from the onset and grow quickly, whereas in ceSGD-slow they start from small values and grow more slowly (requiring much more epochs for convergence).

In all cases – for uniformity of comparison, simplicity, and consistency with the theoretical analysis – we consider scenarios in which the training error (i.e. the NE loss) gets to zero. This is, by definition, the stopping condition for the LAL algorithm. We also used this as a stopping criterion for ceSGD in the “fast” setting. For the other algorithms, the stopping criterion was based on reaching a sufficiently small loss (ceSGD in the “slow” setting), or the collapse of the replicas (eLAL and fBP).

The analysis of the quality of the results was mainly based on the study of the local loss landscape at the solutions. On one hand, we computed the normalized local entropy using BP as described in a sec. D 2, which provides a description of the NE landscape. On the other hand, we also computed the spectrum of the eigenvalues of a smoothed-out version of the NE loss, namely the mean square error loss computed on networks with tanh activations. This loss depends on the parameters β of the activations: we set β to be as small as possible (maximizing the smoothing and thereby measuring features of the landscape at a large scale) under the constraint that all the solutions under consideration were still corresponding to zero error (to prevent degrading the performance). For the Fashion-MNIST case, we also measured the generalization error of each solution, and the robustness to input noise.

In the random patterns scenario we set $N = 999$ and $\alpha = 1$, and tested 10 samples (the same for all the algorithms). The results are presented in figs. 5 and 6. The two analyses allow to rank the algorithms (for the Hessians we can use the maximum eigenvalue as a reasonable metric) and their results are in agreement. As expected, fBP systematically finds very dense regions of solutions, qualitatively compatible with the theoretical analysis (cf. fig. 5 with fig. 1) and corresponding to the narrowest spectra of the Hessian at all β ; the other algorithms follow roughly in the order eLAL, ceSGD-slow, ceSGD-fast, LAL. The latter is a very efficient solver for this model, but it finds solutions in very narrow regions. On the other hand, the same algorithm performed in parallel on a set of interacting replicas is still efficient but much better at discovering WFMs. These results are for $y = 20$ replicas in eLAL, but our tests show that $y = 10$ would be sufficient to match ceSGD-slow and that $y = 100$ would further improve the results and get closer to fBP. Overall, the results of the random pattern case confirm the existence of WFMs in continuous networks, and suggest that a (properly normalized) Hessian spectrum can be used as a proxy for detecting whether an algorithm has found

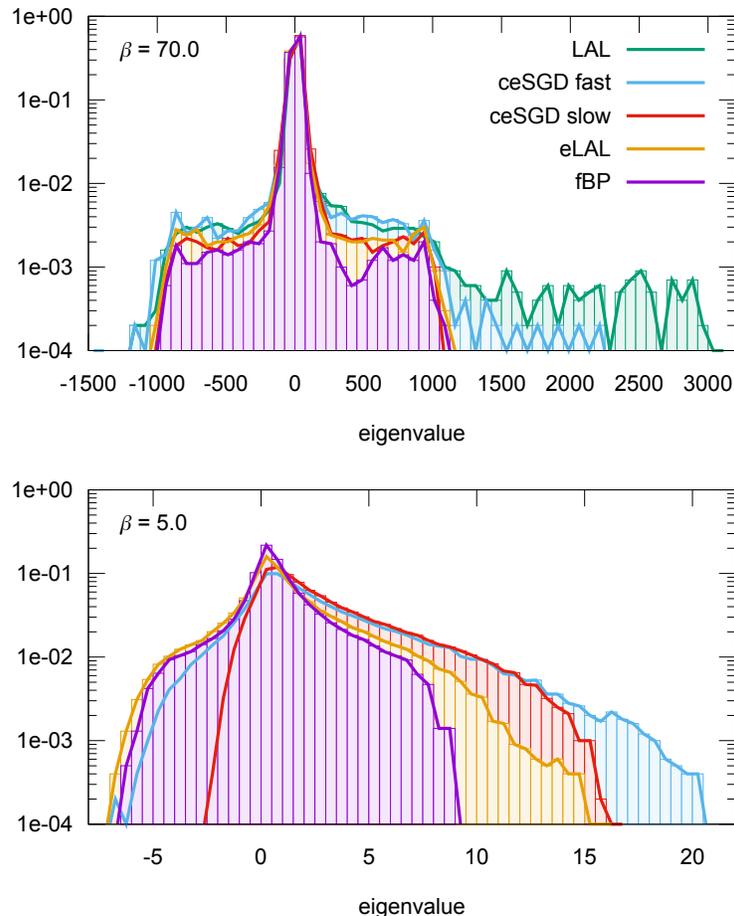


Figure 6. Spectra of the Hessian for the same solutions of fig. 5, for various algorithms. The spectra are directly comparable since they are all computed on the same loss function (MSE; using CE doesn’t change the results qualitatively) and the networks are normalized. The top panel shows the results with the parameter β of the activation functions set to a value such that all solutions of all algorithms are still valid; this value is exclusively determined by the LAL algorithm, and the bottom panel shows the results for a much lower value of β that can be used when removing the LAL solutions, where differences between ceSGD-slow, eLAL and fBP that were not visible at higher β can emerge (the spectrum of LAL would still be the widest by far even at this β).

a WFM region.

We then studied the performance of ceSGD (fast and slow settings), LAL and eLAL on a small fully-connected network which learns to discriminate between two classes of the Fashion-MNIST dataset (we chose the classes Dress and Coat, which are rather challenging to tell apart but also sufficiently different to offer the opportunity to generalize even with a small simple network trained on very few examples). We trained our networks on a small subset of the available examples (500 patterns; binarized to ± 1 by using the median of each image as a threshold on the original grayscale inputs; we filtered both the training and test sets to only use images in which the median was between 0.25 and 0.75 as to avoid too-bright or too-dark images and make the data more uniform and more challenging). This setting is rather close to the one which we could study analytically, except for the patterns statistics and the use of fully-connected rather than tree-like layers, and it is small enough to permit computing the full spectrum of the Hessian. On the other hand, it poses a difficult task in terms of inference (even though finding solutions with zero training error is not hard), which allowed us to compare the results of the analysis of the loss landscape with the generalization capabilities on the test set. Each algorithm was ran 50 times. The results are shown in fig. 7, and they are analogous to those for the random patterns case, but in this setting we can also observe that indeed WFMs tend to generalize better. Also, while we could not run fBP on this data due to the correlations present in the inputs and to numerical problems related to the fully-connected architecture, which hamper convergence, it is still the case that ceSGD can find WFMs if the norms are controlled and increased slowly enough, and that we can significantly improve the (very quick and greedy) LAL algorithm by replicating it, i.e. by effectively adding a local-entropic component.

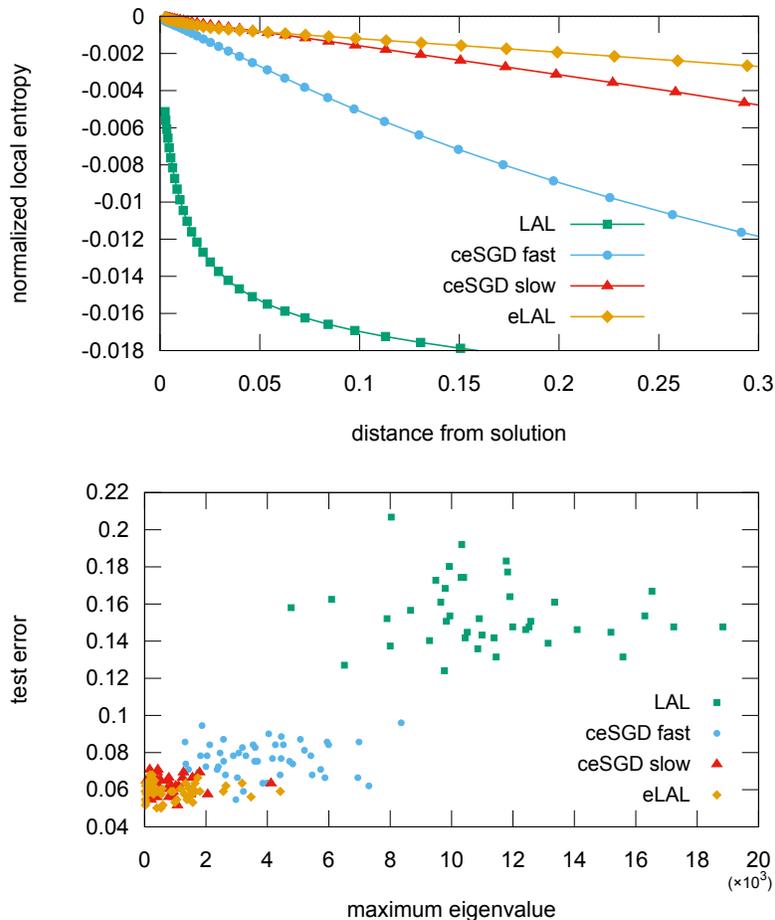


Figure 7. Experiments on a subset of Fashion-MNIST. Top: average normalized local entropies; Bottom: test error vs maximum eigenvalue of the Hessian spectrum at $\beta = 90$. The maximum eigenvalue is correlated to the generalization (WFMs tend to generalize better), and the quality of the minima varies between algorithms.

We also performed an additional batch of tests on a randomized version of the Fashion-MNIST dataset, in which the inputs were reshuffled across samples on a pixel-by-pixel basis (such that each sample only retained each individual pixel bias while the correlations across pixels were lost). This allowed us to bridge the two scenarios and directly compare the local volumes in the presence or absence of features of the data that can lead to proper generalization. We kept the settings of each algorithm as close as possible to those for the Fashion-MNIST tests. Qualitatively, the results were quite similar to the ones on the original data except for a slight degradation of the performance of eLAL compared to ceSGD. Quantitatively, we observed that the randomized version was more challenging and generally resulted in slightly smaller volumes. Additional measures comparing the robustness to the presence of noise in the input (which measures overfitting and thus can be conceived as being a precursor of generalization) confirm the general picture. The detailed procedures and results are reported in Appendix F.

CONCLUSIONS AND FUTURE DIRECTIONS

In this paper, we have generalized the local entropy theory to continuous weights and we have shown that WFM exists in non convex neural systems. We have also shown that the CE loss spontaneously focuses on WFM. On the algorithmic side we have derived and designed novel algorithmic schemes, either greedy (very fast) or message passing, which are driven by the local entropy measure. Moreover, we have shown numerically that ceSGD can be made to converge in WFM by an appropriate cooling procedure of the parameter which controls the norm of the weights. Our findings are in agreement with recent results showing that ReLU transfer functions also help the learning dynamics to focus on WFM (Baldassi et al. 2019). Future work will be aimed at extending our methods to multiple layers,

trying to reach a unified framework for current DNN models. This is a quite challenging task which has the potential to reveal the role that WFM play for generalization in different data regimes and how that can be connected to the many layer architectures of DNN.

ACKNOWLEDGMENTS

CB and RZ acknowledge ONR Grant N00014-17-1-2569. We thank Leon Bottou for discussions.

-
- Carlo Baldassi. Generalization learning in a perceptron with binary synapses. *Journal of Statistical Physics*, 136(5):902–916, 2009. doi:10.1007/s10955-009-9822-1.
- Carlo Baldassi, Alfredo Braunstein, Nicolas Brunel, and Riccardo Zecchina. Efficient supervised learning in networks with binary synapses. *Proceedings of the National Academy of Sciences of the United States of America*, 104(26):11079–1084, 2007. doi:10.1073/pnas.0700324104.
- Carlo Baldassi, Alessandro Ingrosso, Carlo Lucibello, Luca Saglietti, and Riccardo Zecchina. Subdominant Dense Clusters Allow for Simple Learning and High Computational Performance in Neural Networks with Discrete Synapses. *Physical Review Letters*, 115(12):128101, September 2015. ISSN 0031-9007, 1079-7114. doi:10.1103/PhysRevLett.115.128101.
- Carlo Baldassi, Christian Borgs, Jennifer T. Chayes, Alessandro Ingrosso, Carlo Lucibello, Luca Saglietti, and Riccardo Zecchina. Unreasonable effectiveness of learning neural networks: From accessible states and robust ensembles to basic algorithmic schemes. *Proceedings of the National Academy of Sciences*, 113(48):E7655–E7662, November 2016a. ISSN 0027-8424, 1091-6490. doi:10.1073/pnas.1608103113.
- Carlo Baldassi, Federica Gerace, Carlo Lucibello, Luca Saglietti, and Riccardo Zecchina. Learning may need only a few bits of synaptic precision. *Physical Review E*, 93(5):052313, May 2016b. ISSN 2470-0045, 2470-0053. doi:10.1103/PhysRevE.93.052313.
- Carlo Baldassi, Alessandro Ingrosso, Carlo Lucibello, Luca Saglietti, and Riccardo Zecchina. Local entropy as a measure for sampling solutions in constraint satisfaction problems. *Journal of Statistical Mechanics: Theory and Experiment*, 2016(2):P023301, February 2016c. ISSN 1742-5468. doi:10.1088/1742-5468/2016/02/023301.
- Carlo Baldassi, Federica Gerace, Hilbert J Kappen, Carlo Lucibello, Luca Saglietti, Enzo Tartaglione, and Riccardo Zecchina. Role of synaptic stochasticity in training low-precision neural networks. *Physical review letters*, 120(26):268103, 2018. doi:10.1103/PhysRevLett.120.268103.
- Carlo Baldassi, Enrico M Malatesta, and Riccardo Zecchina. Properties of the geometry of solutions and capacity of multilayer neural networks with rectified linear unit activations. *Physical Review Letters*, 123(17):170602, 2019. doi:10.1103/PhysRevLett.123.170602.
- Eli Barkai, David Hansel, and Ido Kanter. Statistical mechanics of a multilayered neural network. *Physical review letters*, 65(18):2312, 1990. doi:10.1103/PhysRevLett.65.2312.
- Eli Barkai, David Hansel, and Haim Sompolinsky. Broken symmetries in multilayered perceptrons. *Phys. Rev. A*, 45:4146–4161, Mar 1992. doi:10.1103/PhysRevA.45.4146.
- Alfredo Braunstein and Riccardo Zecchina. Learning by message passing in networks of discrete synapses. *Physical Review Letters*, 96:030201, Jan 2006. doi:10.1103/PhysRevLett.96.030201.
- Changyan Di, Thomas J Richardson, and Ruediger L Urbanke. Weight distribution of low-density parity-check codes. *IEEE Transactions on Information Theory*, 52(11):4839–4855, 2006. doi:10.1109/TIT.2006.883541.
- Jian Ding and Nike Sun. Capacity lower bound for the ising perceptron. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 816–827. ACM, 2019. doi:10.1145/3313276.3316383.
- A. Engel, H. M. Köhler, F. Tschepke, H. Vollmayr, and A. Zippelius. Storage capacity and learning algorithms for two-layer neural networks. *Phys. Rev. A*, 45:7590–7609, May 1992. doi:10.1103/PhysRevA.45.7590.
- Andreas Engel and Christian Van den Broeck. *Statistical mechanics of learning*. Cambridge University Press, 2001. doi:10.1017/CB09781139164542.
- Silvio Franz and Giorgio Parisi. Recipes for metastable states in spin glasses. *Journal de Physique I*, 5(11):1401–1415, 1995. doi:10.1051/jp1:1995201.
- Heinz Horner. Dynamics of learning for the binary perceptron problem. *Zeitschrift für Physik B Condensed Matter*, 86(2):291–308, 1992. doi:10.1007/BF01313839.

- Haiping Huang and Yoshiyuki Kabashima. Origin of the computational hardness for learning with binary synapses. *Physical Review E*, 90(5):052813, 2014. doi:10.1103/PhysRevE.90.052813.
- Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016. URL <https://arxiv.org/abs/1609.04836>.
- Werner Krauth and Marc Mézard. Storage capacity of memory networks with binary couplings. *J. Phys. France*, 50:3057–3066, 1989. doi:10.1051/jphys:0198900500200305700.
- Florent Krzakala, Federico Ricci-Tersenghi, Lenka Zdeborova, Riccardo Zecchina, Eric W Tramel, and Leticia F Cugliandolo. *Statistical Physics, Optimization, Inference, and Message-Passing Algorithms: Lecture Notes of the Les Houches School of Physics-Special Issue, October 2013*. Number 2013. Oxford University Press, 2016. doi:10.1093/acprof:oso/9780198743736.001.0001.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015. doi:10.1038/nature14539.
- David JC MacKay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003. URL <http://www.inference.org.uk/mackay/itila/p0.html>.
- Marc Mézard, Giorgio Parisi, and Miguel Virasoro. *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, volume 9. World Scientific Publishing Company, 1987. doi:10.1142/0271.
- Graeme Mitchison and Richard Durbin. Bounds on the learning capacity of some multi-layer networks. *Biological Cybernetics*, 60(5):345–365, 1989. doi:10.1007/BF00204772.
- Rémi Monasson. Structural glass transition and the entropy of the metastable states. *Physical review letters*, 75(15):2847, 1995. doi:10.1103/PhysRevLett.75.2847.
- Rémi Monasson and Riccardo Zecchina. Weight space structure and internal representations: a direct approach to learning and generalization in multilayer neural networks. *Physical review letters*, 75(12):2432, 1995. doi:10.1103/PhysRevLett.75.2432.
- Holm Schwarze and John Hertz. Generalization in a large committee machine. *EPL (Europhysics Letters)*, 20(4):375, 1992. doi:10.1209/0295-5075/20/4/015.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017. URL <https://arxiv.org/abs/1708.07747>.

Appendix A: High Local Entropy states from the 1-RSB formalism

Given a system described by a vector of discrete variables W with an associated energy function $E(W)$, the Boltzmann equilibrium distribution at inverse temperature β reads

$$P(W; \beta) = \frac{1}{Z(\beta)} e^{-\beta E(W)} \quad (\text{A1})$$

where the normalization factor Z is given by the partition function

$$Z(\beta) = \sum_W e^{-\beta E(W)} \quad (\text{A2})$$

In the limit $\beta \rightarrow \infty$, the distribution is just a flat measure over the ground states of the energy; we can denote the ground state energy as $E^* = \min_W E(W)$ and the characteristic function over the ground states as

$$\mathbb{X}(W) = \begin{cases} 1 & \text{if } E(W) = E^* \\ 0 & \text{otherwise} \end{cases} \quad (\text{A3})$$

such that $Z(\infty) = \sum_W \mathbb{X}(W)$ and $\log Z(\infty)$ gives the entropy of the ground states.

In Baldassi et al. (2015), we introduced a large-deviation measure with a modified energy function in which each configuration is reweighted by a “local entropy” term. There, we only considered the $\beta \rightarrow \infty$ limit and defined the local entropy as the number of ground states at a certain normalized distance D from a reference configuration \tilde{W} :

$$\mathcal{S}(\tilde{W}, D) = \log \mathcal{N}(\tilde{W}, D) = \log \sum_W \mathbb{X}(W) \delta(d(W, \tilde{W}) - ND) \quad (\text{A4})$$

where $d(\cdot, \cdot)$ is a suitably defined distance function and $\delta(\cdot)$ is the Kronecker delta. With this definition, we can define a modified partition function as follows:

$$Z(\infty, y, D) = \lim_{\beta \rightarrow \infty} \sum_{\tilde{W}} e^{-\beta E(\tilde{W}) + y S(\tilde{W}, D)} \quad (\text{A5})$$

which (up to irrelevant constant factors) coincides with:

$$Z(\infty, y, D) = \sum_{\tilde{W}} \mathbb{X}(\tilde{W}) \mathcal{N}(\tilde{W}, D)^y \quad (\text{A6})$$

This approach can be shown to be strictly related to the 1-step replica-symmetry-broken (1RSB) formalism of the bare energetic problem. First, let us define a local free entropy at any given inverse temperature β' (i.e. a generalization of eq. (A4)), and use a soft-constraint on the distance through a Lagrange multiplier λ :

$$\phi(\tilde{W}, \lambda, \beta') = \log \sum_W e^{-\beta' E(W) - \frac{\lambda}{2} d(W, \tilde{W})} \quad (\text{A7})$$

Note that the use of a Lagrange multiplier is mostly convenient in order to make the relation with the 1RSB description evident. It is only equivalent to using a hard constraint for the distance in the thermodynamic limit, and depending on the convexity properties of the function ϕ , but we will generally ignore these issues for the time being and come back to them later.

We can then rewrite the large-deviation partition function eq. (A5) in this more general case as:

$$Z(\beta, \beta', y, \lambda) = \sum_{\tilde{W}} e^{-\beta E(\tilde{W}) + y \phi(\tilde{W}, \lambda, \beta')} \quad (\text{A8})$$

Let us now consider the case in which $y \in \mathbb{N}$: this allows us, by simple algebraic manipulations, to rewrite the partition function introducing a sum over all the configurations of y replicas of the system:

$$Z(\beta, \beta', y, \lambda) = \sum_{\tilde{W}, \{W^a\}} e^{-\beta E(\tilde{W}) - \beta' \sum_{a=1}^y E(W^a) - \frac{\lambda}{2} \sum_{a=1}^y d(W^a, \tilde{W})} \quad (\text{A9})$$

This partition function describes a system of $y + 1$ interacting *real* replicas with an interaction which is mediated by the reference configuration \tilde{W} . However we can isolate the sum over the configurations of \tilde{W} to obtain a system of y interacting real replicas. In the special case $\beta = 0$ we obtain:

$$Z(0, \beta', y, \lambda) = \sum_{\tilde{W}, \{W^a\}} e^{-\beta' \sum_{a=1}^y E(W^a) + \log \sum_{\tilde{W}} \exp(-\frac{\lambda}{2} \sum_{a=1}^y d(W^a, \tilde{W}))} \quad (\text{A10})$$

We have stressed the fact that the replicas are real to avoid the confusion with the *virtual* replicas used for the “replica trick”: here, we are not replicating the system virtually in order to compute a free entropy in the limit of zero replicas: instead, we are describing a system of y identical interacting objects. The general case of $y \in \mathbb{R}$ can be obtained by analytic continuation once an expression for all integer y is found.

This description is highly reminiscent of – in fact, almost identical to – the derivation of the ergodicity-breaking scheme used in [Monasson \(1995\)](#): there, an auxiliary symmetry breaking field is introduced (having the same role of \tilde{W} in our notation); then, a free energy expression is introduced in which the role of the energy is taken by a “local free entropy” (the analogous of eq. (A4) for general β), after which the system is replicated y times and the auxiliary field \tilde{W} is traced out, leading to a system of y real replicas with an effective pairwise interaction. Finally, the limit of vanishing interaction ($\lambda \rightarrow 0^+$) is taken in order to derive the equilibrium description. When this system is studied in the replica-symmetric (RS) Ansatz, it results in the 1RSB description of the original system, with y having the role of the Parisi parameter (usually denoted by m). Indeed, in this limit of vanishing interaction and for $\beta = 0$, our equations reduce to the 1RSB case as well.

Therefore, apart from minor differences, the main point of discrepancy between that analysis and our approach is that we don’t restrict ourselves to the equilibrium distribution. Instead, we explore the whole range of values of λ . In this context, we also have no reason to restrict ourselves to the range $y \in [0, 1]$, as it is usually done in order to give a physical interpretation to the 1RSB solution; to the contrary, we are (mostly) interested in the limit of large y , in which only the configurations of maximal local entropy are described.

The relationship between our analysis and the usual 1RSB case can be made even more direct, leading to an alternative – although with very similar results – large deviations analysis: consider, instead of eq. (A10), a partition function in which the interaction among the replicas is pairwise (without the reference configuration \bar{W}) and the constraint on the distance is hard (introduced via a Dirac delta function):

$$Z_{1RSB}(\beta', y, D) = \sum_{\{W^a\}} e^{-\beta' \sum_{a=1}^y E(W^a)} \prod_{a>b} \delta(d(W^a, W^b) - ND) \quad (\text{A11})$$

Suppose then that we study the average free entropy $\langle \log Z_{1RSB}(\beta', y, \gamma) \rangle$ (where $\langle \cdot \rangle$ represents the average over the quenched parameters, if any) in the context of replica theory. Then, we will have n virtual replicas of the whole system, and since each system has y real replicas we end up with ny total replicas. Let's use indices c, d for the virtual replicas and a, b for the real ones, such that a configuration will now have two indices, e.g. W^{ca} . Suppose that we manage to manipulate the expression such that it becomes a function, among other order parameters, of the overlaps $q^{ca,db} = \frac{1}{N} \langle W^{ca}, W^{db} \rangle$, where $\langle \cdot, \cdot \rangle$ represents some inner product, and that the distance function $d(\cdot, \cdot)$ can be expressed in terms of those. Then, as usual, we would introduce auxiliary integrals

$$\int \prod_{(ca,db)} (Ndq^{ca,db}) \prod_{(ca,db)} \delta(Nq^{ca,db} - \langle W^{ca}, W^{db} \rangle) \quad (\text{A12})$$

Using this, we can rewrite the interaction term. Say that $d(W, W') = \langle W, W \rangle + \langle W', W' \rangle - 2 \langle W, W' \rangle$, then:

$$\prod_c \prod_{a>b} \delta(d(W^{ca}, W^{cb}) - ND) = \prod_c \prod_{a>b} \delta(N(q^{ca,ca} + q^{cb,cb} - 2q^{ca,cb} - D)) \quad (\text{A13})$$

By assuming replica symmetry, we seek a saddle point with this structure:

$$\begin{aligned} q^{ca,ca} &= Q \\ q^{ca,cb} &= q_1 \quad (a \neq b) \\ q^{ca,db} &= q_0 \quad (c \neq d) \end{aligned} \quad (\text{A14})$$

with $Q \geq q_1 \geq q_0$. The interaction term eq. (A13) becomes:

$$\prod_c \prod_{a>b} \delta(N(q^{ca,ca} + q^{cb,cb} - 2q^{ca,cb} - D)) = \delta(2N(Q - q_1 - D)) \quad (\text{A15})$$

Therefore, the external parameter D eliminates a degree of freedom in the solution to the saddle point equations for the overlaps. The final step in the replica calculation would have the form

$$\begin{aligned} \langle \log Z_{1RSB}(\beta', y, D) \rangle &= \phi_{1RSB}(\beta', y, Q, q_1, q_0, \dots) \delta(Q - q_1 - D) \\ &= \phi_{1RSB}(\beta', y, Q, Q - D, q_0, \dots) \end{aligned} \quad (\text{A16})$$

where ϕ_{1RSB} is the expression which would have been derived in an equilibrium computation without the interaction term, the dots in the argument represent extra order parameters, and the order parameters are fixed by the saddle point equations

$$\begin{aligned} \partial_Q \phi_{1RSB}(\beta', y, Q, Q - D, q_0, \dots) &= 0 \\ \partial_{q_0} \phi_{1RSB}(\beta', y, Q, Q - D, q_0, \dots) &= 0 \\ &\vdots \end{aligned} \quad (\text{A17})$$

Thus, the difference with respect to the usual 1RSB computation is that the equation for finding the extremum over q_1 is removed, and the one for finding the extremum over Q is modified. Maximizing over D , by solving for $\partial_D \phi = 0$, is then equivalent to the usual 1RSB description (equivalent to the case $\lambda \rightarrow 0$ in the soft-constraint case):

$$Z_{1RSB}(\beta', y) = \max_D Z_{1RSB}(\beta', y, D) \quad (\text{A18})$$

In the common case where Q is fixed (e.g. if the variables W are discrete, or constraints on the norm are introduced) then this representation fixes q_1 ; it is clear then that our large deviations analysis (the alternative one of eq. (A11)) is simply derived by fixing q_1 as an external parameter, and thus omitting the saddle point equation $\partial_{q_1} \phi_{1RSB} = 0$. Note that this wouldn't make physical sense in the standard derivation of the 1RSB equations, since in that context q_1 is only introduced as an overlap between virtual replicas when choosing an Ansatz for the solutions of the saddle point equations; our derivation is only physically meaningful when describing a system of real interacting replicas or, in the case of the original derivation from eq. (A5), a system with a modified energy function.

Appendix B: Cross-Entropy minima, errors and high local entropy configurations

1. Cross-Entropy loss ground states

In order to study analytically the properties of the minima of the CE loss function in the case of i.i.d. random patterns, the key obstacle is to compute the normalization factor of the Gibbs measure, the partition function Z . Once this is done one has access to the the Gibbs measure which concentrates on the minima of the loss in the $\beta \rightarrow \infty$ limit.

Z is an exponentially fluctuating random variable and in order to find its most probable values we need to average its logarithm, a complicated task which we perform by the replica method. Once this is done, the typical value of Z can be recovered by $Z_{\text{typ}} \simeq \exp\left(N \langle \log Z \rangle_\xi\right)$, where $\langle \cdot \rangle_\xi$ stands for the average over the random patterns.

We refer to [Engel and Van den Broeck \(2001\)](#) for a thorough review of the replica method. Here we just remind the reader that the replica method is an analytic continuation technique which allows in some cases (mean-field models) to compute the expectation of the logarithm of the partition function from the knowledge of its integer moments. The starting point is the following small n expansion

$$Z^n = 1 + n \log Z + O(n^2)$$

This identity may be averaged over the random patterns and gives the average of the log from the averaged n -th power of the partition function

$$\langle \log Z \rangle_\xi = \lim_{n \rightarrow 0} \frac{\langle Z^n \rangle_\xi - 1}{n}$$

The idea of the replica method is to restrict to integer n and to take the analytic continuation $n \rightarrow 0$

$$\langle Z^n \rangle_\xi = \prod_{a=1}^n \langle Z_a \rangle_\xi = \sum_{\{W^1, \dots, W^n\}} \left\langle e^{-\beta \sum_{a=1}^n E(W^a)} \right\rangle_\xi$$

We have n replicas of the initial model. The random patterns in the expression of the energy disappear once the average has been carried out. Eventually one computes the partition function of an effective system of nN variables with a non random energy function resulting from the average. The result may be written formally as

$$\langle Z^n \rangle_\xi = \exp(NF(n))$$

where F is the expression resulting from the sum over all configurations. Once the small n limit is taken, the final expression can be estimated analytically by means of the saddle-point method given that N is assumed to be large.

In the case of our problem we have

$$Z = \sum_{\{w_i = \pm 1\}} \exp\left(-\beta \sum_{\mu=1}^M f\left(\frac{1}{\sqrt{N}} \sum_{i=1}^N w_i \xi_i^\mu\right)\right)$$

Following the replica approach, we need to compute

$$\langle Z^n \rangle = \left\langle \int \prod_{i,a} d\mu(w_i^a) \prod_{\mu,a} \exp\left(-\beta f\left(\frac{1}{\sqrt{N}} \sum_{i=1}^N w_i^a \xi_i^\mu\right)\right) \right\rangle_\xi$$

where the integration measure is just over the binary values of the weights. By enforcing $x^\mu = \frac{1}{\sqrt{N}} \sum_{i=1}^N w_i \xi_i^\mu$ through a delta function, we can linearize the dependence on the randomness of the patterns and perform the average as follows:

$$\begin{aligned}
\langle Z^n \rangle &= \left\langle \int \prod_{i,a} d\mu(w_i^a) \int \prod_{a,\mu} \frac{dx^{a\mu} d\hat{x}^{a\mu}}{2\pi} \prod_{a,\mu} \exp(-\beta f(x^{a\mu})) \prod_{a,\mu} \exp\left(i\hat{x}^{a\mu} x^{a\mu} - i\hat{x}^{a\mu} \sum_{i=1}^N \frac{w_i^a \xi_i^\mu}{\sqrt{N}}\right) \right\rangle_\xi = \\
&= \int \prod_{i,a} d\mu(w_i^a) \int \prod_{a,\mu} \frac{dx^{a\mu} d\hat{x}^{a\mu}}{2\pi} \prod_{a,\mu} \exp(-\beta f(x^{a\mu})) \prod_{a,\mu} \exp(i\hat{x}^{a\mu} x^{a\mu}) \exp\left(-\frac{1}{2N} \sum_{ab} \hat{x}^{a\mu} \hat{x}^{b\mu} \sum_i \frac{w_i^a w_i^b}{N}\right) = \\
&= \int \prod_{i,a} d\mu(w_i^a) \int \prod_{a,\mu} \frac{dx^{a\mu} d\hat{x}^{a\mu}}{2\pi} \int \prod_{a>b} \frac{dq^{ab} d\hat{q}^{ab}}{2\pi} \prod_{a>b} e^{-Nq^{ab} \hat{q}^{ab}} \prod_{a>b} e^{\hat{q}^{ab} \sum_i w_i^a w_i^b} \prod_{a,\mu} \exp(-\beta f(x^{a\mu})) \times \\
&\quad \times \prod_{a,\mu} \exp(i\hat{x}^{a\mu} x^{a\mu}) \exp\left(-\frac{1}{2} \sum_{ab} \hat{x}^{a\mu} \hat{x}^{b\mu} q^{ab}\right) = \\
&= \int \prod_{a>b} \frac{dq^{ab} d\hat{q}^{ab} N}{2\pi} \prod_{a>b} e^{-Nq^{ab} \hat{q}^{ab}} \left(\int \prod_{i,a} d\mu(w_i^a) \prod_{a>b} e^{\hat{q}^{ab} \sum_i w_i^a w_i^b} \right) \times \\
&\quad \times \left(\int \prod_a \frac{dx^a d\hat{x}^a}{2\pi} \prod_a \exp(-\beta f(x^a)) \prod_a \exp(i\hat{x}^a x^a) \exp\left(-\frac{1}{2} \sum_{ab} \hat{x}^a \hat{x}^b q^{ab}\right) \right)^{\alpha N}
\end{aligned}$$

where we have used the delta functions to introduce the order parameters q^{ab} and \hat{q}^{ab} . In order to write the multiple integrals in a form which can be evaluated by saddle point, we restrict to the replica symmetric assumption $q^{ab} = q$ and $\hat{q}^{ab} = \hat{q}$, and perform few simplifications.

First we sum over the weights:

$$\begin{aligned}
\int \prod_{i,a} d\mu(w_i^a) \prod_{a>b} e^{\hat{q} \sum_i w_i^a w_i^b} &= \int \prod_{i,a} d\mu(w_i^a) \prod_i e^{\frac{\hat{q}}{2} (\sum_a w_i^a)^2} e^{-\frac{\hat{q}}{2} N n} = \\
&= e^{-\frac{\hat{q}}{2} N n} \left(\sum_{w=\pm 1} e^{\frac{\hat{q}}{2} (\sum_a w^a)^2} \right)^N = e^{-\frac{\hat{q}}{2} N n} \left(\int Du \sum_{w=\pm 1} e^{\sqrt{\hat{q}} u \sum_a w^a} \right)^N = \\
&\simeq \exp N n \left[-\frac{\hat{q}}{2} + \int Du \log \left(2 \cosh \sqrt{\hat{q}} u \right) \right]
\end{aligned}$$

Second, we simplify the terms which are raised to the power αN :

$$\begin{aligned}
\int \prod_a \frac{dx^a d\hat{x}^a}{2\pi} \prod_a \exp(-\beta f(x^a)) \prod_a \exp(i\hat{x}^a x^a) \exp\left(-\frac{1}{2} \sum_{ab} \hat{x}^a \hat{x}^b q\right) &= \\
&= \int \prod_a \frac{dx^a d\hat{x}^a}{2\pi} \prod_a \exp(-\beta f(x^a)) \prod_a \exp(i\hat{x}^a x^a) \exp\left(-\frac{1}{2} (1-q) \sum_a (\hat{x}^a)^2 - \frac{q}{2} \left(\sum_a \hat{x}^a\right)^2\right) \\
&= \int Du \left(\int \frac{dx d\hat{x}}{2\pi} e^{-\beta f(x)} e^{i\hat{x} x} \exp\left(-\frac{1}{2} (1-q) \hat{x}^2 + iu\sqrt{q}\hat{x}\right) \right)^n \\
&= \int Du \left(\int \frac{dx}{\sqrt{2\pi}} e^{-\beta f(x)} \frac{\exp\left(-\frac{(x+u\sqrt{q})^2}{2(1-q)}\right)}{\sqrt{1-q}} \right)^n
\end{aligned}$$

Finally we can write the saddle point expression for the replicated partition function:

$$\langle Z^n \rangle \simeq \exp \left[N n \left(\frac{q\hat{q}}{2} - \frac{\hat{q}}{2} + \int Du \log \left(2 \cosh \sqrt{\hat{q}} u \right) + \alpha \int Du \log \int \exp \left(-\beta f \left(x \sqrt{1-q} + u\sqrt{q} \right) \right) \right) \right] = \exp(NnG)$$

where it is useful to write the action G as the sum of three terms

$$G = \frac{\hat{q}}{2} (q - 1) + G_S + \alpha G_E$$

The entropic contribution G_S reads

$$G_s = \int Du \log \left(2 \cosh \sqrt{\hat{q}} u \right)$$

and the energetic one G_E

$$G_E = \int Du \log \int \exp \left(-\beta f \left(x \sqrt{1-q} + u \sqrt{q} \right) \right)$$

The replicated partition function can then be computed in the limit $N \rightarrow \infty$ by solving the saddle point equations $\frac{\partial G}{\partial \hat{q}} = 0$ and $\frac{\partial G}{\partial q} = 0$. The derivatives of G_S and G_E read

$$\frac{\partial G_S}{\partial \hat{q}} = \int Du \frac{\int Dx e^{-\beta f(x\sqrt{1-q}+u\sqrt{q})} \left[-\frac{\beta}{2} f' \left(x\sqrt{1-q} + u\sqrt{q} \right) \left(\frac{u}{\sqrt{q}} - \frac{x}{\sqrt{1-q}} \right) \right]}{\int Dx e^{-\beta f(x\sqrt{1-q}+u\sqrt{q})}}$$

and

$$\frac{\partial G_E}{\partial q} = \int Du \frac{u}{2\sqrt{\hat{q}}} \tanh \left(u \sqrt{\hat{q}} \right) = \frac{1}{2} \left(1 - \int Du \tanh^2 \left(u \sqrt{\hat{q}} \right) \right)$$

Setting to zero these derivatives we get the saddle point equations for q and \hat{q}

$$\begin{aligned} q &= \int Du \tanh^2 \left(u \sqrt{\hat{q}} \right) \\ \hat{q} &= -\alpha \beta \int Du \frac{\int Dx e^{-\beta f(x\sqrt{1-q}+u\sqrt{q})} \left[-\frac{\beta}{2} f' \left(x\sqrt{1-q} + u\sqrt{q} \right) \left(\frac{u}{\sqrt{q}} - \frac{x}{\sqrt{1-q}} \right) \right]}{\int Dx e^{-\beta f(x\sqrt{1-q}+u\sqrt{q})}} = \\ &= -\frac{\alpha}{\sqrt{1-q}} \int Du \frac{\int Dx e^{-\beta f(x\sqrt{1-q}+u\sqrt{q})} \left[\frac{1}{\sqrt{1-q}} + x \left(\frac{u}{\sqrt{q}} - \frac{x}{\sqrt{1-q}} \right) \right]}{\int Dx e^{-\beta f(x\sqrt{1-q}+u\sqrt{q})}} \end{aligned}$$

In the limit of large β we need to rescale the order parameters to obtain finite quantities. By setting $q = 1 - \frac{\delta q}{\beta}$, we find for the last equation

$$\hat{q} = \frac{\alpha \beta^2}{\delta q} \int Du \left[\operatorname{argmax}_x \left(-\frac{x^2}{2} - \log \left(1 + \exp \left(-2\gamma \left(x \sqrt{\delta q} + u \right) \right) \right) \right) \right]^2$$

Once the saddle point equations are solved numerically, we can compute the minimum energy (minimum loss) and the entropy at low temperature. We have:

$$E = -\frac{\partial G}{\partial \beta} = -\alpha \frac{\partial G_E}{\partial \beta} = \alpha \int Du \frac{\int Dx e^{-\beta f(x\sqrt{1-q}+u\sqrt{q})} f \left(x\sqrt{1-q} + u\sqrt{q} \right)}{\int Dx e^{-\beta f(x\sqrt{1-q}+u\sqrt{q})}}$$

In the limit of large β , with $q = 1 - \frac{\delta q}{\beta}$, we find

$$E = \alpha \int Du f \left(x^* \left(u \right) \right)$$

where

$$x^*(u) \equiv \operatorname{argmax}_x \left(-\frac{x^2}{2} - \log \left(1 + \exp \left(-2\gamma \left(x\sqrt{\delta q} + u \right) \right) \right) \right)$$

We can compute the entropy using the relation $S = G + \beta E$.

In figure 2 we show the behavior of the energy vs the loading α . As one may observe, up to relatively large values of α the energy is extremely small, virtually equal zero for any accessible size N .

Having established that by minimizing the the cross-entropy one ends up in regions of perfect classification where the error loss function is zero, it remains to be understood which type of configurations of weights are found. Does the CE converge to a typical zero energy configuration of the error function (i.e. an isolated point-like solution in the weight space) or does it converge to the rare regions of high local entropy?

The answer to this question is that the CE does in fact focus on the HLE subspaces.

2. The local entropy around CE ground states

In order to show this analytically we need to be able to count how many zero error configurations exist in a region close to a typical minima of the CE loss. Following [Franz and Parisi \(1995\)](#), this computation can be done by averaging with the CE Gibbs measure the entropy of the Number of Errors loss function.

Let's call f and g the cross-entropy and error loss functions per pattern, respectively. We need to evaluate the most probable value of

$$\log Z_{FP} = \log \frac{\int \prod_i d\mu(w_i) \prod_\mu \exp \left(-\beta f \left(\sum_i \frac{w_i \xi_i^\mu}{\sqrt{N}} \right) \right) \log \left[\int \prod_i d\mu(v_i) \prod_\mu \exp \left(-\beta' g \left(\sum_i \frac{v_i \xi_i^\mu}{\sqrt{N}} \right) \right) \delta \left(pN - \sum_i w_i v_i \right) \right]}{\int \prod_i d\mu(w_i) \prod_\mu \exp \left(-\beta f \left(\sum_i \frac{w_i \xi_i^\mu}{\sqrt{N}} \right) \right)}$$

which can be computed by replica approach as we have done for Z , yielding the local entropy $\mathcal{E}_D = \langle \log Z_{FP} \rangle$. In the above expression pN is the constrained overlap between the different minima, which is trivially related to the Hamming distance DN by $D = \frac{1}{2}(1 - p)$. We will need to perform twice the replica trick, one to extract the most probable value of $\log Z_{FP}$ (index n) and one to linearize the log inside the integral (index r),

$$\begin{aligned} \langle Z_{FP}^{n,r} \rangle = & \left\langle \int \prod_{ai} d\mu(w_i^a) \prod_{\mu a} \exp \left(-\beta f \left(\sum_i \frac{w_i^a \xi_i^\mu}{\sqrt{N}} \right) \right) \times \right. \\ & \left. \times \frac{1}{r} \left(\int \prod_{ci} d\mu(v_i) \prod_{\mu c} \exp \left(-\beta' g \left(\sum_i \frac{v_i^c \xi_i^\mu}{\sqrt{N}} \right) \right) \delta \left(pN - \sum_i w_i^{a=1} v_i^c \right) - 1 \right) \right\rangle_\xi \end{aligned}$$

This quantity is computed for n, r integer ($a, b = 1 \dots, n$ and $c, d, = 1, \dots, r$) and eventually the analytic continuation $n, r \rightarrow 0$ is taken. We will do this under the replica symmetric (RS) assumption which for small distances D is expected to be exact. For the sake of completeness, we report hereafter all the main steps of the calculation.

We need to compute

$$\begin{aligned} \langle Z_{FP}^{n,r} \rangle = & \int \prod_{ai} d\mu(w_i^a) \int \prod_{a\mu} \frac{dx^{a\mu} d\hat{x}^{a\mu}}{2\pi} \prod_{\mu a} \exp(-\beta f(x^{a\mu})) \prod_{a,\mu} e^{i\hat{x}^{a\mu} x^{a\mu}} \times \\ & \times \int \prod_{ci} d\mu(v_i^c) \int \prod_{c\mu} \frac{dy^{c\mu} d\hat{y}^{c\mu}}{2\pi} \prod \exp(-\beta' g(y^{c\mu})) \times \\ & \times \prod_\mu \left\langle e^{i\hat{x}^{a\mu} \sum_{ia} \frac{w_i^a \xi_i^\mu}{\sqrt{N}} + i\hat{y}^{c\mu} \sum_{ic} \frac{v_i^c \xi_i^\mu}{\sqrt{N}}} \right\rangle_\xi \prod_c \delta \left(pN - \sum_i w_i^{a=1} v_i^c \right) \end{aligned}$$

The average over the patterns is factorized and can be easily performed. Upon expanding the results for large N , the term in the brackets reads

$$\begin{aligned} & \exp \left(-\frac{1}{2N} \sum_i \left(\sum_a w_i^a \hat{x}^{a\mu} + \sum_c v_i^c \hat{y}^{c\mu} \right)^2 \right) = \\ & = \exp \left(-\frac{1}{2} \sum_{ab} \hat{x}^{a\mu} \hat{x}^{b\mu} \frac{1}{N} \sum_i w_i^a w_i^b - \frac{1}{2} \sum_{cd} \hat{x}^{c\mu} \hat{x}^{d\mu} \frac{1}{N} \sum_i v_i^c v_i^d - \sum_{ac} \hat{x}^{a\mu} \hat{y}^{c\mu} \frac{1}{N} \sum_i w_i^a v_i^c \right) \end{aligned}$$

Introducing the order parameters corresponding to the different overlaps we find for the total expression

$$\begin{aligned} \langle Z_{FP}^{n,r} \rangle &= \int \prod_{a>b} \frac{dq^{ab} d\hat{q}^{ab} N}{2\pi} \int \prod_{c>d} \frac{ds^{cd} d\hat{s}^{cd} N}{2\pi} \prod_{a>b} e^{-N\hat{q}^{ab} q^{ab}} \prod_{c>d} e^{-N\hat{s}^{cd} s^{cd}} \int \prod_c \frac{d\hat{p}^c N}{2\pi} \prod_c e^{-N\hat{p}^c p} \times \\ & \times \int \prod_{a>1,c} \frac{dt^{ac} d\hat{t}^{ac}}{2\pi} N \prod_{a>1,c} e^{-N\hat{t}^{ac} t^{ac}} \times \\ & \times \int \prod_{ai} d\mu(w_i^a) \int \prod_{ci} d\mu(v_i^c) \prod_{a>b} e^{\hat{q}^{ab} \sum_i w_i^a w_i^b} \prod_{c>d} e^{\hat{s}^{cd} \sum_i v_i^c v_i^d} \prod_c e^{\hat{p}^c \sum_i w_i^{a=1} v_i^c} \prod_{a>1,c} e^{\hat{t}^{ac} \sum_i w_i^a v_i^c} \times \\ & \times \left(\int \prod_a \frac{dx^a d\hat{x}^a}{2\pi} \int \prod_c \frac{dy^c d\hat{y}^c}{2\pi} \prod_a e^{-\beta f(x^a)} \prod_c e^{-\beta' g(y^c)} \prod_a e^{ix^a \hat{x}^a} \times \right. \\ & \left. \times \prod_c e^{iy^c \hat{y}^c} e^{-\frac{1}{2} \sum_{ab} \hat{x}^a \hat{x}^b q^{ab} - \frac{1}{2} \sum_{cd} \hat{y}^c \hat{y}^d s^{cd} - \sum_c \hat{x}^1 \hat{y}^c p - \sum_{a>1,c} \hat{x}^a \hat{y}^c t^{ac}} \right)^{\alpha N} \end{aligned}$$

In order to proceed, we search the solutions of the saddle point equations in the RS subspace, $q^{ab} = q$, $\hat{q}^{ab} = \hat{q}$, $s^{ab} = s$, etc. The various factors can be simplified as follows:

$$\begin{aligned} \prod_i e^{\sum_{a>b} \hat{q}^{ab} w_i^a w_i^b} &= \prod_i e^{\frac{\hat{q}}{2} [(\sum_a w_i^a)^2 - \sum_a (w_i^a)^2]} = e^{-\frac{Nn\hat{q}}{2}} \prod_i e^{\frac{\hat{q}}{2} (\sum_a w_i^a)^2}, \\ \prod_i e^{\sum_{c>d} \hat{s}^{cd} v_i^c v_i^d} &= e^{-\frac{Nr\hat{s}}{2}} \prod_i e^{\frac{\hat{s}}{2} (\sum_c v_i^c)^2}, \\ \prod_i e^{\sum_c \hat{p}^c w_i^{a=1} v_i^c} &= \prod_i e^{\hat{p} \sum_c w_i^{a=1} v_i^c}, \\ \prod_i e^{\sum_{a>1,c} \hat{t}^{ac} w_i^a v_i^c} &= \prod_i e^{\hat{t} (\sum_{a>1} w_i^a) (\sum_c v_i^c)} = \prod_i e^{\hat{t} (\sum_a w_i^a) (\sum_c v_i^c) - \hat{t} w_i^{a=1} (\sum_c v_i^c)}, \end{aligned}$$

and

$$\begin{aligned} -\sum_{a>b} \hat{q}^{ab} q^{ab} &= -\frac{n(n-1)}{2} \hat{q} q \simeq \frac{n}{2} \hat{q} q, \\ -\sum_{c>d} \hat{s}^{cd} s^{cd} &= -\frac{r(r-1)}{2} \hat{s} s \simeq \frac{r}{2} \hat{s} s, \\ -\sum_c \hat{p}^c p &= -r \hat{p} p, \\ -\sum_{a>1,c} \hat{t}^{ac} t^{ac} &= -(n-1) r \hat{t} t \simeq r \hat{t} t. \end{aligned}$$

A series of further simplifications are needed in order to write the $\langle Z_{FP}^{n,r} \rangle$ in the appropriate saddle point form. The terms containing the integrals over w and v become factorized

$$\begin{aligned}
& \left(\int \prod_a d\mu(w^a) \int \prod_c d\mu(v^c) e^{-\frac{n}{2}\hat{q}-\frac{r}{2}\hat{s}} \int Dz e^{z\sqrt{\hat{q}-\hat{t}}\sum_a w^a} \int Du e^{u\sqrt{\hat{s}-\hat{t}}\sum_c v^c} \int Dx e^{x\sqrt{\hat{t}}(\sum_a w^a+\sum_c v^c)} \times \right. \\
& \left. \times e^{(\hat{p}-\hat{t})w^{a=1}\sum_c v^c} \right)^N = \\
& = \left(\int Dz Du Dx \left(\int d\mu(w) e^{-\frac{\hat{q}}{2}+w(z\sqrt{\hat{q}-\hat{t}}+x\sqrt{\hat{t}})} \right)^{n-1} \times \right. \\
& \left. \times \int d\mu(w^{a=1}) e^{-\frac{\hat{q}}{2}+w^{a=1}(z\sqrt{\hat{q}-\hat{t}}+x\sqrt{\hat{t}})} \left(\int d\mu(v) e^{(u\sqrt{\hat{s}-\hat{t}}+x\sqrt{\hat{t}})v+(\hat{p}-\hat{t})w^{a=1}v} \right)^r \right)^N
\end{aligned}$$

where we have kept the notation $w^{a=1}$ just for the sake of clarity. Being an integration variable we now drop it. By summing over w and v and with some straightforward change of variables we get

$$\begin{aligned}
& \int Dz Du Dx \frac{\sum_{w=\pm 1} e^{w(z\sqrt{\hat{q}-\hat{t}}+x\sqrt{\hat{t}})} \left(-\frac{\hat{s}}{2} + \log 2 \cosh \left((u\sqrt{\hat{s}-\hat{t}}+x\sqrt{\hat{t}})v + (\hat{p}-\hat{t})w \right) \right)}{\sum_{w=\pm 1} e^{w(z\sqrt{\hat{q}-\hat{t}}+x\sqrt{\hat{t}})}} \\
& = -\frac{\hat{s}}{2} + \int Dz \frac{\sum_{w=\pm 1} e^{wz\sqrt{\hat{q}}} \left(\int Du Dx \log \cosh \left(u\sqrt{\hat{s}-\hat{t}} + \sqrt{\hat{t}} \left(\sqrt{\frac{\hat{t}}{\hat{q}}}z + \sqrt{\frac{\hat{q}-\hat{t}}{\hat{q}}}x \right) + (\hat{p}-\hat{t})w \right) \right)}{\sum_{w=\pm 1} e^{wz\sqrt{\hat{q}}}} = \\
& = -\frac{\hat{s}}{2} + \int Dz \frac{\sum_{w=\pm 1} e^{wz\sqrt{\hat{q}}} \int D\phi \log \cosh \left(\phi \sqrt{\frac{(\hat{s}-\hat{t})\hat{q}+(\hat{q}-\hat{t})\hat{t}}{\hat{q}}} + \frac{\hat{t}}{\sqrt{\hat{q}}}z + (\hat{p}-\hat{t})w \right)}{2 \cosh z\sqrt{\hat{q}}} =
\end{aligned}$$

For the integral containing the dependence on f and g we find similar simplifications.

$$\begin{aligned}
& \int Dz Dh \left(\int \frac{dx d\hat{x}}{2\pi} e^{-\beta f(x) - \frac{1-2q}{2}\hat{x}^2 + i\hat{x}(z\sqrt{q-t}+h\sqrt{t+x})} \right)^{n-1} \times \\
& \times \left(\int \prod_c \frac{dy^c d\hat{y}^c}{2\pi} \int \frac{dx d\hat{x}}{2\pi} e^{-\beta f(x) - \frac{1-2q}{2}\hat{x}^2 + i\hat{x}(z\sqrt{q-t}+h\sqrt{t+x} + i(p-t)\sum_c \hat{y}^c)} \times \right. \\
& \left. \times \prod_c e^{-\beta' g(y^c) - \frac{1-s}{2}\sum_c (\hat{y}^c)^2 - \frac{s-t}{2}(\sum_c \hat{y}^c)^2 + i\sum_c \hat{y}^c (h\sqrt{t+y^c})} \right) = \\
& = \int Dz Dh \frac{1}{Z_0} \int \prod_c \frac{dy^c d\hat{y}^c}{2\pi} \int \frac{dx}{\sqrt{2\pi}} \frac{e^{-\beta f(x)}}{\sqrt{1-q}} e^{-\frac{1}{2(1-q)}[(z\sqrt{q-t}+h\sqrt{t+x})^2 - (p-t)^2(\sum_c \hat{y}^c)^2 + 2i(z\sqrt{q-t}+h\sqrt{t+x})(p-t)\sum_c \hat{y}^c]} \times \\
& \times \prod_c e^{-\beta' g(y^c) - \frac{1-s}{2}\sum_c (\hat{y}^c)^2 - \frac{s-t}{2}(\sum_c \hat{y}^c)^2 + i\sum_c \hat{y}^c (h\sqrt{t+y^c})} = \dots
\end{aligned}$$

where we have used the notation

$$Z_0 \equiv \int \frac{dx}{\sqrt{2\pi}} \frac{e^{-\beta f(x)}}{\sqrt{1-q}} e^{-\frac{1}{2(1-q)}(z\sqrt{q-t}+h\sqrt{t+x})^2}$$

Continuing the computation, we can linearize the terms in \hat{y}^c with an auxiliary integral, factor the terms with the

index c , take the $r \rightarrow 0$ limit and perform explicitly two integrals:

$$\begin{aligned}
\cdots &= \int Dz Dh \frac{1}{Z_0} \int Du \int \frac{dx}{\sqrt{2\pi}} \frac{e^{-\beta f(x)}}{\sqrt{1-q}} e^{-\frac{1}{2(1-q)}(z\sqrt{q-t}+h\sqrt{t+x})^2} \times \\
&\times \left(\int \frac{dy d\hat{y}}{2\pi} e^{-\beta' g(y) - \frac{1-s}{2} \hat{y}^2 + i\hat{y} \left(\sqrt{s-t - \frac{(p-t)^2}{1-q}} u + h\sqrt{t+y} - \frac{(z\sqrt{q-t}+h\sqrt{t+x})(p-t)}{1-q} \right)} \right)^r = \\
&= \int Dz Dh \frac{1}{Z_0} \int Du \int \frac{dx}{\sqrt{2\pi}} \frac{e^{-\beta f(x)}}{\sqrt{1-q}} e^{-\frac{1}{2(1-q)}(z\sqrt{q}+x)^2} \times \\
&\times \log \int \frac{dy}{\sqrt{2\pi}} \frac{e^{-\beta' g(y)}}{\sqrt{1-s}} e^{-\frac{1}{\sqrt{2(1-s)}} \left[u\sqrt{s-t - \frac{(p-t)^2}{1-q}} + \sqrt{t} \left(z\sqrt{\frac{t}{q}} + h\sqrt{\frac{q-t}{q}} \right) + y - \frac{z\sqrt{q}(p-t)}{1-q} \right]^2} = \\
&= \int Dz \frac{1}{Z'_0} \int \frac{dx}{\sqrt{2\pi}} \frac{e^{-\beta f(x)}}{\sqrt{1-q}} e^{-\frac{1}{2(1-q)}(z\sqrt{q}+x)^2} \times \\
&\times \int Du \log \int \frac{dy}{\sqrt{2\pi}} \frac{e^{-\beta' g(y)}}{\sqrt{1-s}} e^{-\frac{1}{\sqrt{2(1-s)}} \left[u\sqrt{s-t - \frac{(p-t)^2}{1-q} - \frac{t(q-t)}{q} + \frac{t}{\sqrt{q}} z - \frac{z\sqrt{q}(p-t)}{1-q} + y \right]^2}
\end{aligned}$$

where

$$Z'_0 \equiv \int \frac{dx}{\sqrt{2\pi}} \frac{e^{-\beta f(x)}}{\sqrt{1-q}} e^{-\frac{1}{2(1-q)}(z\sqrt{q}+x)^2}$$

The local entropy $\mathcal{E}_D = \langle \log Z_{FP} \rangle$ is nothing but the total exponent for the saddle point equations, which after some additional changes of variables can eventually be written as

$$\mathcal{E}_D = \frac{1}{2} \hat{s}s - \hat{p}p + t\hat{t} - \frac{\hat{s}}{2} + G_S + \alpha G_E$$

where

$$G_S = \int Dz \frac{\sum_{w=\pm 1} e^{wz\sqrt{\hat{q}}} \int D\phi \log \left[2 \cosh \left(\left(\sqrt{\hat{s} - \frac{\hat{t}^2}{q}} \right) \phi + \frac{\hat{t}}{\sqrt{\hat{q}}} z + (\hat{p} - \hat{t}) w \right) \right]}{2 \cosh(z\sqrt{\hat{q}})}$$

and

$$G_E = \int Dz \frac{\int Dx e^{-\beta f(x\sqrt{1-q}+z\sqrt{q})} \int Du \log \left[\int Dy e^{-\beta' g \left(y\sqrt{1-s} + \frac{t}{\sqrt{q}} z + \frac{p-t}{\sqrt{1-q}} x + u\sqrt{\omega} \right)} \right]}{\int Dx e^{-\beta f(x\sqrt{1-q}+z\sqrt{q})}}$$

where we have defined $\omega = s - t - \frac{(p-t)^2}{1-q} + \frac{t(q-t)}{q}$.

If we now take the limit $\beta' \rightarrow \infty$ and plug in the expression for the error loss function $g(x) = \Theta(-x)$, we can eliminate one integral

$$\begin{aligned}
\lim_{\beta' \rightarrow \infty} \int Dy e^{-\beta' \Theta \left(- \left(y\sqrt{1-s} + \frac{t}{\sqrt{q}} z + \frac{p-t}{\sqrt{1-q}} x + u\sqrt{\omega} \right) \right)} &= \\
&= \int Dy \Theta \left(y\sqrt{1-s} + \frac{t}{\sqrt{q}} z + \frac{p-t}{\sqrt{1-q}} x + u\sqrt{\omega} \right) = H \left(- \frac{\frac{t}{\sqrt{q}} z + \frac{p-t}{\sqrt{1-q}} x + u\sqrt{\omega}}{\sqrt{1-s}} \right)
\end{aligned}$$

and the expression of G_E simplifies to

$$G_E = \int Dz \frac{\int Dx e^{-\beta f(x\sqrt{1-q}+z\sqrt{q})} \int Du \log H \left(- \frac{\frac{t}{\sqrt{q}} z + \frac{p-t}{\sqrt{1-q}} x + u\sqrt{\omega}}{\sqrt{1-s}} \right)}{\int Dx e^{-\beta f(x\sqrt{1-q}+z\sqrt{q})}}$$

In order to compute the entropy for a given distance $D = (1 - p)/2$, we need to solve the saddle point equations with respect to $s, \hat{s}, t, \hat{t}, \hat{p}$ with the values of q and \hat{q} obtained by solving the equations for the CE loss function. The saddle point equations can be written as

$$\begin{aligned} 0 &= \frac{\partial \mathcal{E}_D}{\partial \hat{s}} = \frac{1}{2}(s - 1) + \frac{\partial G_S}{\partial \hat{s}} \\ 0 &= \frac{\partial \mathcal{E}_D}{\partial \hat{p}} = -p + \frac{\partial G_S}{\partial \hat{p}} \\ 0 &= \frac{\partial \mathcal{E}_D}{\partial \hat{t}} = t + \frac{\partial G_S}{\partial \hat{t}} \\ 0 &= \frac{\partial \mathcal{E}_D}{\partial s} = \frac{\hat{s}}{2} + \alpha \frac{\partial G_E}{\partial s} \\ 0 &= \frac{\partial \mathcal{E}_D}{\partial t} = t + \alpha \frac{\partial G_E}{\partial t} \end{aligned}$$

with

$$\begin{aligned} \frac{\partial G_S}{\partial \hat{s}} &= \int Dz \frac{\sum_{w=\pm 1} e^{wz\sqrt{\hat{q}}} \int D\phi \tanh\left(\phi\sqrt{\hat{s} - \frac{\hat{t}^2}{\hat{q}}} + z\frac{\hat{t}}{\sqrt{\hat{q}}} + (\hat{p} - \hat{t})w\right) \frac{\phi}{2\sqrt{\hat{s} - \frac{\hat{t}^2}{\hat{q}}}}}{2 \cosh z\sqrt{\hat{q}}}, \\ \frac{\partial G_S}{\partial \hat{p}} &= \int Dz \frac{\sum_{w=\pm 1} e^{wz\sqrt{\hat{q}}} \int D\phi \tanh\left(\phi\sqrt{\hat{s} - \frac{\hat{t}^2}{\hat{q}}} + z\frac{\hat{t}}{\sqrt{\hat{q}}} + (\hat{p} - \hat{t})w\right) w}{2 \cosh z\sqrt{\hat{q}}}, \\ \frac{\partial G_S}{\partial \hat{t}} &= \int Dz \frac{\sum_{w=\pm 1} e^{wz\sqrt{\hat{q}}} \int D\phi \tanh\left(\phi\sqrt{\hat{s} - \frac{\hat{t}^2}{\hat{q}}} + z\frac{\hat{t}}{\sqrt{\hat{q}}} + (\hat{p} - \hat{t})w\right)}{2 \cosh z\sqrt{\hat{q}}} \left[-\frac{\hat{t}\phi}{\hat{q}\sqrt{\hat{s} - \frac{\hat{t}^2}{\hat{q}}}} + z\frac{\hat{t}}{\sqrt{\hat{q}}} - w \right], \\ \frac{\partial G_E}{\partial s} &= \int Dz \frac{\int Dx e^{-\beta f(x\sqrt{1-q}+z\sqrt{q})} \int Du \mathcal{G}\left(\frac{\frac{t}{\sqrt{q}}z + \frac{p-t}{\sqrt{1-q}}x + u\sqrt{\omega}}{\sqrt{1-s}}\right) \left[\frac{u}{2\sqrt{1-s}\sqrt{\omega}} + \frac{\frac{t}{\sqrt{q}} + x\frac{p-t}{\sqrt{1-q}} + u\sqrt{\omega}}{2(1-s)^{3/2}} \right]}{\int Dx e^{-\beta f(x\sqrt{1-q}+z\sqrt{q})}}, \\ \frac{\partial G_E}{\partial t} &= \int Dz \frac{\int Dx e^{-\beta f(x\sqrt{1-q}+z\sqrt{q})} \int Du \mathcal{G}\left(\frac{\frac{t}{\sqrt{q}}z + \frac{p-t}{\sqrt{1-q}}x + u\sqrt{\omega}}{\sqrt{1-s}}\right) \left[-\frac{x}{\sqrt{1-q}} + \frac{z}{\sqrt{q}} + u\frac{1}{\sqrt{\omega}} \frac{pq-t}{q(1-q)} \right] \frac{1}{\sqrt{1-s}}}{\int Dx e^{-\beta f(x\sqrt{1-q}+z\sqrt{q})}} \end{aligned}$$

where we have defined $\mathcal{G}(-x) \equiv \frac{\partial}{\partial x} \log H[-x]$.

The results are reported in fig. 3. We may observe that the minima of the CE are indeed surrounded by an exponential number of zero error solutions. In other words, the CE focuses on HLE regions.

3. RS stability and zero entropy condition

In order to corroborate the validity of the RS solution, we need to check two necessary conditions: the entropy of the CE model is positive and the replica symmetric solution is stable. In other words we can focus on values of the parameters α, γ, β such that both conditions are met.

Following [Engel and Van den Broeck \(2001\)](#), the stability is verified if the following condition is met

$$\alpha\gamma_E\gamma_S < 1$$

where γ_E and γ_S are the two eigenvalues of the Hessian matrix computed at the RS saddle point. We find

$$\gamma_S = \int Dz \left[1 - \tanh(z\sqrt{\hat{q}}) \right]^2$$

and

$$\gamma_E = \frac{1}{(1-q)^2} \int Dz \left[1 - \left(\langle x^2 \rangle_z - \langle x \rangle_z^2 \right) \right]^2$$

where the averages can be expressed through the quantity

$$\overline{\tilde{x}^k(z)} = \frac{\int \frac{dx d\tilde{x}}{2\pi} \tilde{x}^k e^{-\frac{1}{2}(1-q)\tilde{x}^2 + i\tilde{x}(z\sqrt{q}+x)} e^{-\beta f(x)}}{\int \frac{dx d\tilde{x}}{2\pi} e^{-\frac{1}{2}(1-q)\tilde{x}^2 + i\tilde{x}(z\sqrt{q}+x)} e^{-\beta f(x)}}$$

One finds:

$$\overline{\tilde{x}(z)} = \frac{i}{1-q} \frac{\int \frac{dx}{\sqrt{2\pi}} e^{-\frac{(z\sqrt{q}+x)^2}{2(1-q)}} (z\sqrt{q}+x) e^{-\beta f(x)}}{\int \frac{dx}{\sqrt{2\pi}} e^{-\frac{(z\sqrt{q}+x)^2}{2(1-q)}} e^{-\beta f(x)}} = \frac{i}{\sqrt{1-q}} \frac{\int Dx x e^{-\beta f(z\sqrt{q}+x\sqrt{1-q})}}{\int Dx e^{-\beta f(z\sqrt{q}+x\sqrt{1-q})}} = \frac{i}{\sqrt{1-q}} \langle x \rangle_z$$

and similarly

$$\overline{\tilde{x}^2(z)} = \frac{1}{(1-q)^2} \frac{\int Dx (1-q - (1-q)x^2) e^{-\beta f(z\sqrt{q}+x\sqrt{1-q})}}{\int Dx e^{-\beta f(z\sqrt{q}+x\sqrt{1-q})}} = \frac{1}{1-q} - \frac{\langle x^2 \rangle_z}{(1-q)}$$

For each α , we can thus identify the values of γ and β for which both the entropy is positive and the solution is stable. In particular, β can be chosen to be quite large, corresponding to energies which are extremely small (RSB is expected to have relatively minor effects at zero temperature).

Appendix C: Wide Flat Minima for the continuous case

Following [Barkai et al. \(1992\)](#) and the technique described in sec. [A](#) we may analyze the existence of WFM by studying the 1-RSB saddle point equations, with q_1 and y (usually called m in the 1-RSB context) used as control parameters.

The computation of the average of $\langle \log V \rangle_\xi$ over the patterns by the replica methods leads to the following saddle point expression in the large N limit

$$\frac{1}{N} \langle \log V \rangle_\xi = \text{extr}_{q_\ell^{ab}, \hat{q}_\ell^{ab}, E_\ell^a} G \left(\left\{ q_\ell^{ab}, \hat{q}_\ell^{ab}, E_\ell^a \right\} \right)$$

where

$$G \left(\left\{ q_\ell^{ab}, \hat{q}_\ell^{ab}, E_\ell^a \right\} \right) = G_S \left(\left\{ q_\ell^{ab}, \hat{q}_\ell^{ab}, E_\ell^a \right\} \right) + \alpha G_E \left(\left\{ q_\ell^{ab} \right\} \right).$$

Given that the distribution of the input patterns is the same for each hidden unit, averages are expected to be independent of ℓ and the dependency on ℓ of the order parameters can be dropped $\left\{ q_\ell^{ab}, \hat{q}_\ell^{ab}, E_\ell^a \right\} \rightarrow \left\{ q^{ab}, \hat{q}^{ab}, E^a \right\}$. In 1-RSB scheme, once the conjugate order parameters $\left\{ \hat{q}^{ab}, E^a \right\}$ are integrated out, the expressions for G_S and G_E read

$$G_S(q_0, q_1, y) = \frac{1}{2} \left[\frac{1 + (y-1)\Delta q_1}{1 - q_1 + y\Delta q_1} + \ln 2\pi + \left(1 - \frac{1}{y} \right) \ln(1 - q_1) + \frac{1}{y} \ln(1 - q_1 + y\Delta q_1) \right]$$

$$G_E(q_0, q_1, y) = \frac{1}{y} \int \prod_{\ell=1}^K Dv_\ell \ln \left(\int \prod_{\ell=1}^K Du_\ell (\Sigma_{(K)})^y \right)$$

where $\Delta q_1 = q_1 - q_0$ and $\Sigma_{(K)}$ is a complicated function of the order parameters which for $K = 3$ reads

$$\Sigma_{(3)} = H_1 H_2 + H_1 H_3 + H_2 H_3 - 2H_1 H_2 H_3$$

with

$$H_\ell = H \left[\sqrt{\frac{\Delta q_1}{1 - q_1}} u_\ell + \sqrt{\frac{q_0}{1 - q_1}} v_\ell \right].$$

For our WFM computation, q_1 is the constrained overlap between the weight vectors of the m real replicas. q_0 is the only parameter for which we have to solve the saddle point equation. In order to look for the WFM of maximum volume we are interested in the large y limit. In this case the expressions simplify substantially

$$G_S(q_1) = \frac{1}{2} [1 + \ln 2\pi + \ln(1 - q_1)]$$

$$G_E(q_0, q_1) = \int \prod_{\ell=1}^K Dv_\ell \max_{u_1, u_2, u_3} \left[-\frac{\sum_{\ell=1}^3 u_\ell^2}{2} + \log \left(\tilde{H}_1 \tilde{H}_2 + \tilde{H}_1 \tilde{H}_3 + \tilde{H}_2 \tilde{H}_3 - 2\tilde{H}_1 \tilde{H}_2 \tilde{H}_3 \right) \right]$$

where $\tilde{H}_\ell \equiv H \left[\sqrt{\frac{d_0}{1 - q_1}} u_\ell + \sqrt{\frac{q_0}{1 - q_1}} v_\ell \right]$ and $d_0 \equiv y\Delta q_1$. The pre-factor $\frac{1}{y}$ of G_E has been eliminated by a change of variables $u'_\ell = \frac{u_\ell}{\sqrt{y}}$ (with u' then renamed u).

Notice that $G_S(q_1)$ corresponds to the volume at $\alpha = 0$, i.e. to the volume of the weight space just under the spherical constraints, with the real replicas forced to be at an overlap q_1 . If WFM exist for positive α , we expect to observe the normalized local entropy $\mathcal{V}_1(q_1) = \alpha G_E(q_1)$ to approach 0 for q_1 sufficiently close to one.

In fig. 1 (top panel) we report the values of the WFM volumes vs the overlap q_1 , for different values of α . Indeed one may observe that the behavior is qualitatively similar to that of the binary perceptron, WFM exist deep into the RSB region, i.e. besides the RSB states and all the related local minima and saddles, there exist absolute minima which are flat at large distances. We mention that evaluating the max inside the integral is a quite challenging task as the function to be maximized may present multiple maxima. We have tackled this problem by an appropriate sampling technique.

The case $K = 3$ is still relatively close to the perceptron, though the geometrical structure of its minima is already dominated by non convex RSB features for $\alpha > 1.76$. A case which is closer to more realistic NN is $K \gg 1$, which, luckily enough, is easier to study analytically (Barkai et al. 1992).

The 1-RSB expression for G , simplifies to

$$G_S(q_0, q_1, y) = \frac{1}{2} \left[\frac{1 + (y - 1)\Delta q_1}{1 - q_1 + y\Delta q_1} + \ln 2\pi + \left(1 - \frac{1}{y}\right) \ln(1 - q_1) + \frac{1}{y} \ln(1 - q_1 + y\Delta q_1) \right]$$

$$G_E(q_0, q_1, y) = \frac{1}{y} \int Dv \ln \int Du \left(H \left[\sqrt{\frac{\Delta q_{1\text{eff}}}{1 - q_{1\text{eff}}}} u + \sqrt{\frac{q_{0\text{eff}}}{1 - q_{1\text{eff}}}} v \right] \right)^y$$

where $q_{0\text{eff}} = 1 - \frac{2}{\pi} \arccos q_0$, $q_{1\text{eff}} = 1 - \frac{2}{\pi} \arccos q_1$ and $\Delta q_{1\text{eff}} = q_{1\text{eff}} - q_{0\text{eff}}$. While the critical capacity diverges with K as $\sqrt{\ln K}$ (Monasson and Zecchina 1995), the value of α at which RSB sets in and the landscape of the minima becomes non trivial remains finite, $\alpha_{RSB} \simeq 2.95$.

As we have done for the $K = 3$, we study the large y limit. We find

$$G_S(q_1) = \frac{1}{2} [1 + \ln 2\pi + \ln(1 - q_1)]$$

$$G_E(q_0, q_1) = \int Dv \max_u \left\{ -\frac{u^2}{2} + \log \left[H \left(\sqrt{\frac{\Delta q_1^e}{1 - q_1^e}} u + \frac{q_0^e}{1 - q_1^e} v \right) \right] \right\}$$

In fig. 1 (bottom panel) we observe that WFM are indeed still present.

In order to check the validity of the WFM computation one should check for the stability of the solutions of the saddle point equations by a stability analysis or a 2-RSB computation. For numerical reasons this is a quite difficult task. We thus decided to follow a different path which has also algorithmic interest, namely to study the problem by Belief Propagation.

1. Large deviation analysis of the Parity Machine

Here we report the results of the large deviation analysis on the so-called parity machine. Its network structure is identical to the committee machine, except for the output unit which performs the product of the K hidden units instead of taking a majority vote. The outputs of the hidden units are still given by sign activations. Thus, the overall output of the network reads:

$$\sigma_{\text{out}}^\mu = \prod_{\ell=1}^K \tau_\ell.$$

For a given set of patterns, the volume of the weights which correctly classifies the patterns is then given by

$$V = \int \prod_{i\ell} dw_{\ell i} \prod_{\ell} \delta \left(\sum_i w_{\ell i}^2 - \tilde{N} \right) \prod_{\mu} \Theta \left(\sigma^\mu \prod_{\ell=1}^K \tau_\ell^\mu \right).$$

The computation proceeds as for the committee machine case, until we find the following expressions for the 1-RSB volume:

$$G_S(q_1) = \frac{1}{2} \left[\frac{1 + (y-1)\Delta q_1}{1 - q_1 + y\Delta q_1} + \ln 2\pi + \left(\frac{y-1}{y} \right) \ln(1 - q_1) + \frac{1}{y} \ln(1 - q_1 + y\Delta q_1) \right]$$

$$G_E(q_1) = \frac{1}{y} \int \prod_{\ell=1}^K Dv_\ell \ln \left[\int \prod_{\ell=1}^K Du_\ell \left(\sum_{\{\tau_\ell\}} \prod_{\ell} H_\ell(\tau_\ell \omega_\ell) \Theta \left(\prod_{\ell=1}^K \tau_\ell \right) \right)^y \right]$$

where $\Delta q_1 = q_1 - q_0$ and where $\omega_\ell = \sqrt{\frac{\Delta q_1}{1 - q_1}} u_\ell + \sqrt{\frac{q_0}{1 - q_1}} v_\ell$ and q_0 is fixed by a saddle point equation. The sum over the internal states be computed for general K , leading to the following final expression for G_E :

$$G_E(q_1) = \frac{1}{y} \int \prod_{\ell=1}^K Dv_\ell \ln \left[\int \prod_{\ell=1}^K Du_\ell \frac{1}{2^y} (1 + \zeta_K(\{\omega_\ell\}))^y \right]$$

where $\zeta_K(\{\omega_\ell\}) = (-1)^K \prod_{\ell=1}^K (1 - 2H_\ell(\omega_\ell))$. In the large K limit, $\zeta_K(\{\omega_\ell\})$ converges rapidly to zero and the expression for G_E simplifies. We can thus compute the volume by optimizing over q_0 for arbitrary y and q_1 , and compare it to the volume that one would find for the same distance when no patterns are stored: the log-ratio of the two volumes is constant and equal to $-\alpha \log(2)$. This shows that the minima never become flat, at any distance scale.

Appendix D: Belief Propagation on a Tree-like committee machine with continuous weights: equations, local volume around solutions and algorithms for the replicated network

1. BP equations for the committee machine

We can use Belief Propagation (see e.g. [Baldassi et al. \(2016a\)](#), [Braunstein and Zecchina \(2006\)](#), [Krzakala et al. \(2016\)](#)) to study the space of the solutions of a tree-like committee machine with continuous weights and random inputs (the outputs can either be random or generated from a rule). The messages in this case are probability density distributions over \mathbb{R} , and we will need to ensure normalization by using an additional constraints over the norm of the weights vectors.

The basic factor graph is thus composed of N continuous variable nodes x_i , divided into K groups of N/K variables each; for each pattern, we will have K factor nodes, each one involving one group of N/K variables and an auxiliary binary output variable τ_{ka} (with two indices, one for the hidden unit and one for the pattern), and another factor node connected to the τ variables and controlling the final output. We will enforce the normalization of the weights by adding an extra field for each variable x_i , as explained below.

As a general notation scheme, we will use the letter h to denote messages from variable nodes to factor nodes, the letter u for messages from factor nodes to variable nodes, and the letter m for non-cavity marginals over variables.

Let's then call $h_{ki \rightarrow ka}(x_{ki})$ the cavity message from variable ki (representing a weight with hidden unit index $k \in \{1, \dots, K\}$, weight index $i \in \{1, \dots, N/K\}$, whose value is x_{ki}) to the factor node ka (representing the part of an input pattern a which involves the hidden unit k). The BP equation reads:

$$h_{ki \rightarrow ka}(x_{ki}) \propto u_n(x_{ki}) \prod_{b \neq a} u_{kb \rightarrow ki}(x_{ki}) \quad (\text{D1})$$

where $u_{kb \rightarrow ki}(x_{ki})$ represents a message from another pattern node kb to the variable ki , while $u_n(x_{ki})$ is an external field enforcing the normalization constraint, which is the same for all variables:

$$u_n(x) = e^{-\frac{\chi}{2}x^2} \quad (\text{D2})$$

This method of enforcing normalization is equivalent to using a Dirac delta in the limit of large N , but it's easier to implement in the BP algorithm. In order to set the parameter χ we will need to evaluate the norm of the weights at convergence (detailed below) and adjust χ until the norm matches the desired value, $\sum_{ki} x_{ki}^2 = N$.

The other messages read:

$$u_{ka \rightarrow ki}(x_{ki}) = \sum_{\tau_{ka} = \pm 1} h_{ka}(\tau_{ka}) \int \prod_{j \neq i} h_{kj \rightarrow ka}(x_{kj}) dx_{kj} \Theta \left(\tau_{ka} \left(\sum_{j \neq i} x_{kj} \xi_{kj}^{ka} + x_{ki} \xi_{ki}^{ka} \right) \right) \quad (\text{D3})$$

where the message h_{ka} goes from the auxiliary output variable τ_{ka} to node ka . Notice that the messages h are assumed to be normalized, but the messages u aren't, because the integral of expression (D3) might diverge.

In the limit of large N , eq. (D3) can be approximated using the central limit theorem, since we're assuming that the input pattern entries are random i.i.d. variables. Therefore, we don't need the full distributions $h_{kj \rightarrow ka}$ to perform the integral, only their first and second moments. For simplicity of implementation, though, instead of the moments it is actually more convenient to use the inverse of the variance and the mean rescaled by the variance, which we will denote with ζ and $\tilde{\mu}$, respectively. We will use the following notation:

$$\zeta_{ki \rightarrow ka} = \left(\langle x_{ki}^2 \rangle_{h_{ki \rightarrow ka}} - \langle x_{ki} \rangle_{h_{ki \rightarrow ka}}^2 \right)^{-1} \quad (\text{D4})$$

$$\tilde{\mu}_{ki \rightarrow ka} = \zeta_{ki \rightarrow ka} \langle x_{ki} \rangle_{h_{ki \rightarrow ka}} \quad (\text{D5})$$

With these, we can compute the two auxiliary quantities

$$c_{ka \rightarrow ki} = \sum_{j \neq i} \frac{\tilde{\mu}_{kj \rightarrow ka}}{\zeta_{kj \rightarrow ka}} \xi_{kj}^{ka} \quad (\text{D6})$$

$$v_{ka \rightarrow ki} = \sum_{j \neq i} \zeta_{kj \rightarrow ka}^{-1} (\xi_{kj}^{ka})^2 \quad (\text{D7})$$

which can be simplified by noting that $(\xi_{kj}^{ka})^2 = 1$. We can now rewrite eq. (D3) using the central limit theorem:

$$\begin{aligned} u_{ka \rightarrow ki}(x_{ki}) &= \sum_{\tau_{ka} = \pm 1} h_{ka}(\tau_{ka}) \int Dz \Theta \left(\tau_{ka} (c_{ka \rightarrow ki} + z \sqrt{v_{ka \rightarrow ki}} + x_{ki} \xi_{ki}^{ka}) \right) \\ &= \frac{1}{2} \left(1 + h_{ka} \operatorname{erf} \left(\frac{c_{ka \rightarrow ki} + x_{ki} \xi_{ki}^{ka}}{\sqrt{2v_{ka \rightarrow ki}}} \right) \right) \end{aligned} \quad (\text{D8})$$

where in the last line we have introduced the shorthand notation

$$h_{ka} \equiv h_{ka}(+1) - h_{ka}(-1)$$

because we can represent any message over a binary variable with a single quantity (in this case, the "magnetization" often employed in spin glass literature) and therefore we abuse the notation and identify that message with a single parameter. The context and the presence or absence of the argument will suffice to disambiguate the notation.

In the limit of large N , we observe that $c_{ka \rightarrow ki} = O(\sqrt{N})$ and $v_{ka \rightarrow ki} = O(N)$, and therefore $\frac{c_{ka \rightarrow ki}}{\sqrt{v_{ka \rightarrow ki}}} = O(1)$, but $\frac{x_{ki} \xi_{ki}^{ka}}{\sqrt{v_{ka \rightarrow ki}}} = O\left(\frac{1}{\sqrt{N}}\right)$ since we are assuming $x_{ki} = O(1)$ due to the normalization constraint. We can therefore expand $u_{ka \rightarrow ki}$ to the second order, obtaining (up to an irrelevant factor):

$$u_{ka \rightarrow ki}(x_{ki}) \propto 1 + U_{ka \rightarrow ki} x_{ki} - \frac{1}{2} (V_{ka \rightarrow ki} - U_{ka \rightarrow ki}^2) x_{ki}^2 \quad (\text{D9})$$

where

$$U_{ka \rightarrow ki} = \frac{\xi_{ki}^{ka}}{\sqrt{v_{ka \rightarrow ki}}} \frac{2h_{ka} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} \frac{c_{ka \rightarrow ki}^2}{v_{ka \rightarrow ki}}\right)}{1 + h_{ka} \operatorname{erf}\left(\frac{c_{ka \rightarrow ki}}{\sqrt{2v_{ka \rightarrow ki}}}\right)} \quad (\text{D10})$$

$$V_{ka \rightarrow ki} = \frac{\xi_{ki}^{ka}}{\sqrt{v_{ka \rightarrow ki}}} \frac{c_{ka \rightarrow ki}}{\sqrt{v_{ka \rightarrow ki}}} U_{ka \rightarrow ki} + U_{ka \rightarrow ki}^2 \quad (\text{D11})$$

Now we can substitute this into eq. (D1):

$$h_{ki \rightarrow ka}(x_{ki}) = \frac{1}{z_{ki \rightarrow ka}} e^{-\frac{\chi}{2} x_{ki}^2} \prod_{b \neq a} \left(1 + U_{kb \rightarrow ki} x_{ki} - \frac{1}{2} (V_{ka \rightarrow ki} - U_{ka \rightarrow ki}^2) x_{ki}^2 \right)$$

where $z_{ki \rightarrow ka}$ is a normalization constant. Expanding this to the second order we finally obtain:

$$h_{ki \rightarrow ka}(x_{ki}) = \frac{1}{z_{ki \rightarrow ka}} e^{-\frac{1}{2} (\chi + \sum_{b \neq a} V_{kb \rightarrow ki}) x_{ki}^2 + \sum_{b \neq a} U_{kb \rightarrow ki} x_{ki}} \quad (\text{D12})$$

and thus we obtain the following simple expression for the parameters of the distribution eqs. (D4)-(D5):

$$\zeta_{ki \rightarrow ka} = \chi + \sum_{b \neq a} V_{kb \rightarrow ki} \quad (\text{D13})$$

$$\tilde{\mu}_{ki \rightarrow ka} = \sum_{b \neq a} U_{kb \rightarrow ki} \quad (\text{D14})$$

Also, the normalization constant is:

$$z_{ki \rightarrow ka} = \frac{\sqrt{2\pi}}{\sqrt{\lambda_{ki \rightarrow ka}}} e^{\frac{1}{2} \frac{\tilde{\mu}_{ki \rightarrow ka}^2}{\zeta_{ki \rightarrow ka}}} \quad (\text{D15})$$

One can immediately write the corresponding formulas for the non-cavity parameters of the marginal $m_{ki}(x_{ki})$:

$$\zeta_{ki} = \chi + \sum_b V_{kb \rightarrow ki} \quad (\text{D16})$$

$$\tilde{\mu}_{ki} = \sum_b U_{kb \rightarrow ki} \quad (\text{D17})$$

$$z_{ki} = \frac{\sqrt{2\pi}}{\sqrt{\lambda_{ki}}} e^{\frac{1}{2} \frac{\tilde{\mu}_{ki}^2}{\zeta_{ki}}} \quad (\text{D18})$$

This form of the equations shows that we can, in fact, parametrize all the distributions as Gaussian, and also that the updates in the equations can be performed efficiently for each node by keeping in memory the non-cavity versions of the parameters ζ and $\tilde{\mu}$ and the parameters U and V , such that updating $\zeta, \tilde{\mu}$ is only a matter of subtracting the previous value of U, V and adding the new one.

We can also note that the computation of U, V in eqs. (D10)-(D11) can be made computationally more efficient by writing $c_{ka \rightarrow ki}$ and $v_{ka \rightarrow ki}$ in terms of their non-cavity counterparts

$$c_{ka} = \sum_j \frac{\tilde{\mu}_{kj \rightarrow ka}}{\zeta_{kj \rightarrow ka}} \xi_{kj}^{ka} \quad (\text{D19})$$

$$v_{ka} = \sum_j \zeta_{kj \rightarrow ka} (\xi_{kj}^{ka})^2 \quad (\text{D20})$$

and expanding. However, in practice this further approximation does not work well at large values of α when using the Focusing-BP protocol (detailed below, sec. D3) and at moderate values of N : by avoiding it, we can find solutions at slightly larger values of α (e.g. we can reach $\alpha = 2.7$ instead of 2.6 at $N = 999, K = 3$). Relatedly, it is also slightly problematic at large values of the polarization field λ when exploring the space of configurations (see below, sec. D2). In general terms, the problems arise when the fields become very polarized and the normalization constants become very small. Therefore, in our implementation this approximation is optional.

In order to set the normalization parameter χ , we simply compute the following quantity at the end of each iteration

$$\frac{1}{N} \sum_{ki} \langle x_{ki}^2 \rangle = \frac{1}{N} \sum_{ki} (\zeta_{ki}^{-1} + \zeta_{ki}^{-2} \tilde{\mu}_{ki}^2) \quad (\text{D21})$$

and adjust χ so that this becomes 1. Of course, we also add the criterion that this adjustment needs to be sufficiently small in evaluating whether the BP equations have converged.

The remaining BP equations involve the nodes connecting the auxiliary variables τ_{ka} and enforcing the desired outputs from the committee:

$$u_{ka}(\tau_{ka}) = \frac{1}{2} \left(1 + \tau_{ka} \operatorname{erf} \left(\frac{c_{ka}}{\sqrt{2v_{ka}}} \right) \right) \quad (\text{D22})$$

$$h_{ka}(\tau_{ka}) = \sum_{\{\tau_{la}\}_{l \neq k}} \prod_{l \neq k} u_{la}(\tau_{la}) \Theta \left(\sigma_a \left(\sum_{l \neq k} \tau_{la} + \tau_{ka} \right) \right) \quad (\text{D23})$$

where σ_a is the desired output for pattern a . For small K , we compute h_{ka} exactly (i.e. without using the central limit theorem), which can be done in $O(K^2)$ time.

Equations (D6), (D7), (D10), (D11), (D13), (D14), (D19), (D20), (D22) and (D23), together with the normalization requirement obtained through eq. (D21), form the full set of BP equations.

The free entropy (also sometimes known as action) of the system can be computed from the usual BP formulas. We get:

$$\phi = \frac{1}{N} \left(\sum_{ki} f_{ki} - \sum_{ka} f_{ka} - \sum_a f_a \right) \quad (\text{D24})$$

where:

$$f_{ki} = \log z_{ki} \quad (\text{D25})$$

$$f_{ka} = \sum_i (\log z_{ki \rightarrow ka} - \log z_{ki}) \quad (\text{D26})$$

$$f_a = -\log \left(1 + \frac{1}{2} \sigma_a \operatorname{erf} \left(\frac{\sum_k h_{ka}}{\sqrt{2 \sum_k (1 - h_{ka}^2)}} \right) \right) \quad (\text{D27})$$

From this, we can compute the entropy by simply accounting for the energetic contribution introduced by the normalization constraint. We also shift it by subtracting the log-volume of the normalized sphere, so that its value is upper bounded by 0

$$S = \phi + \frac{\chi}{2} - \log(\sqrt{2\pi e}) \quad (\text{D28})$$

2. Exploring the space of solutions around a given configuration

Given a particular configuration \tilde{W} , which we assume normalized as $\sum_{ki} \tilde{w}_{ki}^2 = N$, we are interested in exploring the space of solutions at a given distance D from it (for a suitable definition of the distance). Analogously to the norm, in the limit of large N we can control the distance by just adding an extra field to each node (putting it as an extra factor in eq. (D1)), of the form:

$$u_D(x_{ki}) = e^{-\frac{\lambda}{2}(x_{ki} - \tilde{w}_{ki})^2}$$

By varying the auxiliary parameter λ between 0 and ∞ we can restrict ourselves to smaller and smaller regions around \tilde{w} . Adding this extra field in practice just amounts at adding two terms $\lambda \tilde{w}_{ki}$ and λ to the expressions of $\tilde{\mu}_{ki \rightarrow ka}$ and $\zeta_{ki \rightarrow ka}$, respectively (eqs. (D13)-(D14)).

For convenience, we actually abuse the terminology and use a squared distance in our definition:

$$d(x, \tilde{w}) = \frac{1}{2N} \|x - \tilde{w}\|_2^2 \quad (\text{D29})$$

which can be computed from the BP messages at convergence, as follows. First define the auxiliary quantities (representing the cavity variance and mean of each variable without the u_D field):

$$r_{ki \rightarrow D} = \langle x_{ki}^2 \rangle_{h_{ki \rightarrow D}} - \langle x_{ki} \rangle_{h_{ki \rightarrow D}}^2 = (\zeta_{ki} - \lambda)^{-1} \quad (\text{D30})$$

$$\mu_{ki \rightarrow D} = \langle x_{ki} \rangle_{h_{ki \rightarrow D}} = \frac{\tilde{\mu}_{ki} - \lambda \tilde{w}_{ki}}{r_{ki \rightarrow D}} \quad (\text{D31})$$

Then the expression of the average distance reads:

$$\begin{aligned} \langle d(x, \tilde{W}) \rangle &= \frac{1}{N} \sum_{ki} \frac{\int dx_{ki} u_D(x_{ki}) h_{ki \rightarrow D}(x_{ki})^{\frac{1}{2}} (x_{ki} - \tilde{w}_{ki})^2}{\int dx_{ki} u_D(x_{ki}) h_{ki \rightarrow D}(x_{ki})} \\ &= \frac{1}{N} \sum_{ki} \frac{1}{2\lambda} \left(\frac{r_{ki \rightarrow D}}{r_{ki \rightarrow D} + \lambda^{-1}} + \frac{1}{\gamma} \left(\frac{\mu_{ki \rightarrow D} - \tilde{w}_{ki}}{r_{ki \rightarrow D} + \lambda^{-1}} \right)^2 \right) \end{aligned} \quad (\text{D32})$$

The free entropy of the system can be written as before, eq. (D24), but now in the computation of the entropy we also need to account for the energy of the extra field:

$$S = \phi + \frac{\chi + \lambda}{2} - \frac{\lambda}{N} \sum_{ki} \tilde{w}_{ki} \frac{\tilde{\mu}_{ki}}{\lambda_{ki}} - \log(\sqrt{2\pi e}) \quad (\text{D33})$$

By varying λ and using eqs. (D32) and (D33), we can obtain a plot of the local entropy as a function of the distance around any given configuration \tilde{W} , as long as the BP equations converge. As a general rule of thumb, the equations don't converge when λ is too low in the 1-RSB phase, and when λ is too large and \tilde{W} is not a solution. The equations do converge however even in the 1-RSB phase for large enough λ if \tilde{W} is a solution, which can be understood as the external field breaking the symmetry. When \tilde{W} is not a solution, going to the limit $\lambda \rightarrow \infty$ eventually restricts the BP equations to a region of the configuration space without solutions, leading to non-normalizable messages (this could of course be amended e.g. by just working at non-zero temperature). Besides these situations, other numerical problems may arise under certain circumstances when N is not large enough, due to the approximations in the messages.

As a consistency check, it can be verified numerically that formula (D33) yields the expected result for the $\alpha = 0$ case, $S = \log(\sqrt{D(2-D)})$.

3. Focusing-BP

In order to implement the Focusing-BP protocol we need to add a new type of node to each variable. These nodes do not directly represent energy terms in the usual sense, and therefore they are not factor nodes; rather, their role is that of effectively representing an interaction of y identical replicas of the original system with an extra auxiliary configuration x^* . The derivation (in a discrete setting) can be found in [Baldassi et al. \(2016a\)](#).

We denote as $u_{\star \rightarrow ki}(x_{ki})$ the new extra field to be multiplied in eq. (D1). This field, like all others over the x_{ki} variables, is parametrized by two quantities $U_{\star \rightarrow ki}$ and $V_{\star \rightarrow ki}$ which get added to $\tilde{\mu}_{ki}$ and ζ_{ki} (eqs. (D13) and (D14)). The update equation is rather involved, but it can be simplified by breaking it down in steps and adopting a new notation for messages composition, leading to this expression:

$$u_{\star \rightarrow ki} = (((h_{ki \rightarrow \star} * \lambda) \uparrow (y-1)) \otimes u_{n^*}) * \lambda \quad (\text{D34})$$

where the notation is as follows:

- $u = h * \lambda$ represents (intuitively speaking) the effect of passing a message h through a Gaussian interaction with strength λ . In formulas:

$$u(\tilde{x}) = \int dx e^{-\frac{\lambda}{2}(\tilde{x}-x)^2} h(x)$$

- $h = u_1 \otimes u_2$ represents the ‘‘composition’’ of two messages, in formulas:

$$h(x) \propto u_1(x) u_2(x)$$

which, using the internal representation in terms of rescaled mean $\tilde{\mu}$ and inverse variance ζ for the h , and the corresponding quantities for the $u_{1/2}$ fields, simply translates to:

$$\begin{aligned}\tilde{\mu} &= U_1 + U_2 \\ \zeta &= V_1 + V_2\end{aligned}$$

- $\hat{u} = u \uparrow y$ with integer y represents the composition of u with itself (i.e. $\hat{u} = u \otimes u \otimes \dots \otimes u$) y times. This can be trivially extended to non-integer y and yields:

$$\begin{aligned}\hat{U} &= yU \\ \hat{V} &= yV\end{aligned}$$

- u_{n^*} is a normalization field, required for normalization of the x^* variables analogously to the u_n field for the x variables; analogously to that field, it is parametrized with a single parameter χ^* and can be represented as $U_{n^*} = 0$, $V_{n^*} = \chi^*$ (cf. equations (D1), (D2) and (D13)).
- $h_{ki \rightarrow \star}$ is the cavity field from a variable x_{ki} to the corresponding variable x_{ki}^* , which can be defined from the relation $m_{ki} = h_{ki \rightarrow \star} \otimes u_{\star \rightarrow ki}$.

Therefore formula (D34) can be intuitively understood as follows: the cavity messages coming from the replicated variable nodes ($h_{ki \rightarrow \star}$) are passed through their interaction (with strength λ) with the corresponding variable x_{ki}^* ; the resulting messages are composed together (there are $y - 1$ identical messages) and with the normalization field u_{n^*} ; the resulting cavity message $h_{\star \rightarrow ki}$ is again passed through the interaction λ to yield the field $u_{\star \rightarrow ki}$.

The parameter χ^* must be set to a value that normalizes the x^* variables; the norm of the x^* variables can be obtained from their marginals with the same formula used for the x , eq. (D21); the marginals can be obtained with this formula:

$$m_{\star} = ((h_{ki \rightarrow \star} * \lambda) \uparrow y) \otimes u_{n^*}$$

Appendix E: Numerical experiments details

Here we provide the details and the settings used for the numerical experiments reported in sec. V.

In all the experiments and for all algorithms except fBP we have used a mini-batch size of 100. The mini-batches were generated by randomly shuffling the datasets and splitting them at each epoch. For eLAL, the permutations were performed independently for each replica. Also, for all algorithms except fBP the weights were initialized from a uniform distribution and then normalized for each unit. The learning rate η was kept fixed throughout the training. The parameters γ and β for the ceSGD algorithm were initialized at some values γ_0 , β_0 and multiplied by $1 + \gamma_1$, $1 + \beta_1$ after each epoch. Analogously, the parameter λ for the eLAL algorithm was initialized to λ_0 and multiplied by $1 + \lambda_1$ after each epoch. The parameter λ for the fBP algorithm ranged in all cases between 0.5 and 30 with an exponential schedule divided into 30 steps; at each step, the algorithm was run until convergence or at most 200 iterations. We used $y = 20$ for eLAL and $y = 10$ for fBP. The stopping criterion for ceSGD-fast was that a solution (0 errors with $\beta = \infty$) was found; for ceSGD-slow, that the CE loss reached 10^{-7} ; for eLAL, that the sum of the squared distances between each replica and the average replica \tilde{W} reached 10^{-7} . We also report here the average and standard deviation of the number of epochs \bar{T} for each algorithm.

Parameters for the case of random patterns. ceSGD-fast: $\eta = 10^{-2}$, $\gamma_0 = 3$, $\beta_0 = 1$, $\gamma_1 = 1$, $\beta_1 = 10^{-3}$ ($\bar{T} = 770 \pm 150$). ceSGD-slow: $\eta = 3 \cdot 10^{-3}$, $\gamma_0 = 0.1$, $\beta_0 = 0.5$, $\gamma_1 = 4 \cdot 10^{-4}$ ($\bar{T} = (1.298 \pm 0.007) \cdot 10^4$). LAL: $\eta = 5 \cdot 10^{-3}$ ($\bar{T} = 76 \pm 15$). eLAL: $\eta = 10^{-2}$, $\lambda_0 = 0.5$, $\lambda_1 = 10^{-4}$ ($\bar{T} = 861 \pm 315$).

Parameters for the Fashion-MNIST experiments. ceSGD-fast: $\eta = 2 \cdot 10^{-4}$, $\gamma_0 = 5$, $\beta_0 = 2$, $\gamma_1 = 1$, $\beta_1 = 10^{-4}$ ($\bar{T} = 460 \pm 334$). ceSGD-slow: $\eta = 3 \cdot 10^{-5}$, $\gamma_0 = 0.5$, $\beta_0 = 0.5$, $\gamma_1 = 10^{-3}$, $\beta_1 = 10^{-3}$ ($\bar{T} = (3.57 \pm 0.05) \cdot 10^3$). LAL: $\eta = 10^{-4}$ ($\bar{T} = 61 \pm 21$). eLAL: $\eta = 2 \cdot 10^{-3}$, $\lambda_0 = 30$, $\lambda_1 = 5 \cdot 10^{-3}$ ($\bar{T} = 190 \pm 40$).

Appendix F: Experiments with randomized Fashion-MNIST

Our protocol to produce a randomized versions of the Fashion-MNIST dataset was as follows. Starting with the binarized, two-class dataset described in the main text, new input patterns were derived from the original ones as

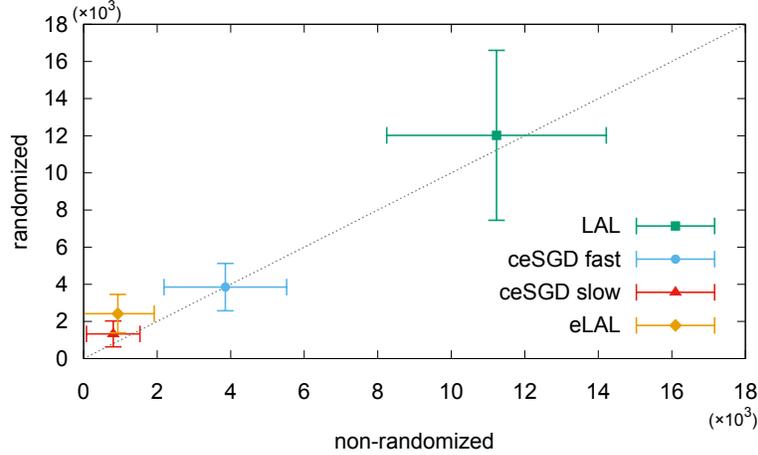


Figure A1. Comparison of the distributions of the maximum eigenvalues (means and standard deviations) between the original Fashion-MNIST dataset and its randomized version, for various algorithms. The “non-randomized” values use the same data as the clouds of points of fig. 7, bottom panel).

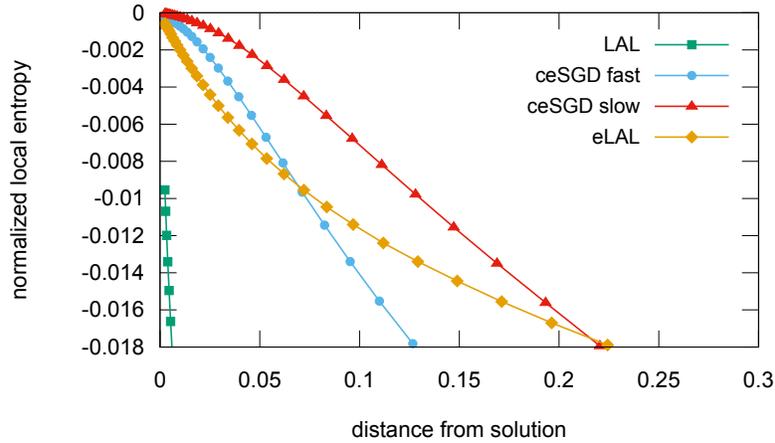


Figure A2. Normalized local entropies as a function of the distance for solutions found by various algorithms, using the randomized-Fashion-MNIST dataset. This is the analogous of fig. 7, top panel, on the randomized dataset.

$\xi_i^{\mu'} = \xi_i^{\pi_i(\mu)}$ where $\pi_i(\cdot)$ is a random permutation, different for each i (i.e. we shuffled each pixel across samples). In this way, each pixel i has the same bias as in the original dataset, $\mathbb{E}_\mu[\xi_i'] = \mathbb{E}_\mu[\xi_i]$, but the (connected) correlations are destroyed, $\mathbb{E}_\mu[\xi_i'\xi_j'] \approx \mathbb{E}_\mu[\xi_i']\mathbb{E}_\mu[\xi_j']$, and furthermore the patterns $\xi^{\mu'}$ no longer carry information about the target label σ^μ , so that no generalization is possible. As a result, the randomized patterns carry more information per pixel that needs to be stored by the device, which in turn can be expected to make the learning problem harder.

We produced 5 such shuffled datasets and performed 10 tests on each with the same algorithms as for the Fashion-MNIST tests, using the same parameters when possible. The exception was eLAL, which we had to tweak slightly to avoid divergencies: we used $\eta = 5 \cdot 10^{-4}$, $\lambda_0 = 20$, $\lambda_1 = 1 \cdot 10^{-2}$.

We directly compared the results with those obtained on the original dataset. Fig. A1 shows that the maximum eigenvalues (cf. fig. 7, bottom panel) hardly change between the two tests, with only a slight degradation for the eLAL algorithm.

Fig. A2 shows the volumes around the solutions, computed with the BP algorithm; this is the analogous of fig. 7, top panel. We also kept the same scale for easiness of comparison. Again, we observe a slight degradation of the eLAL algorithm compared to ceSGD, albeit only at short distances. It is still the case that LAL is by far the worst algorithm, and that ceSGD slow is better than ceSGD fast. The volumes are overall smaller than for the original dataset.

We also performed an additional experiment on both datasets, measuring the robustness towards the presence of

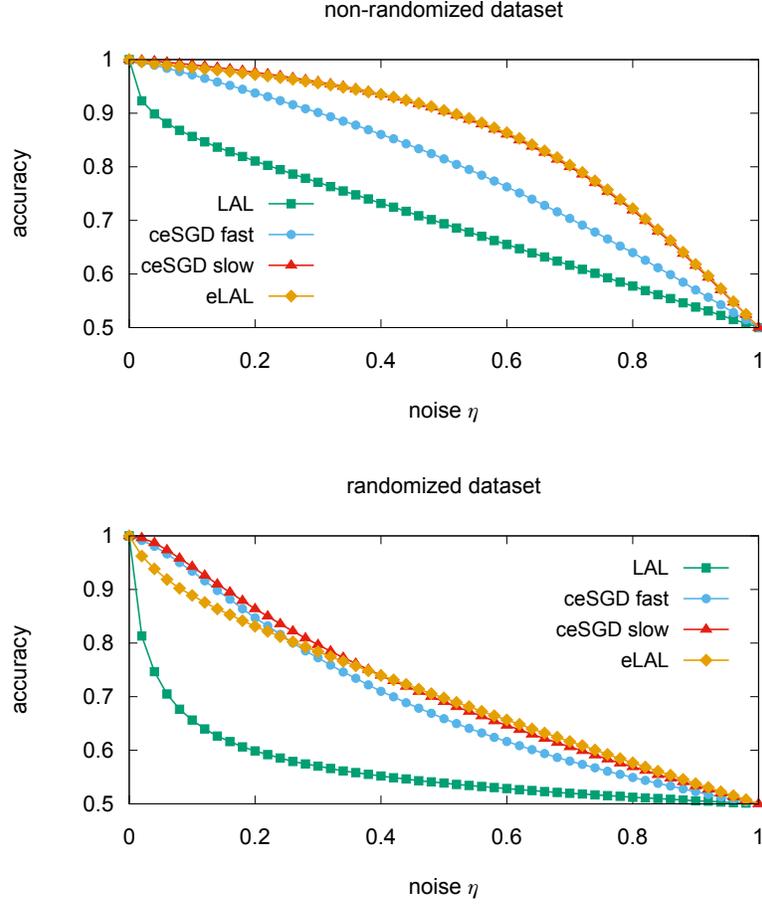


Figure A3. Training accuracy in the presence of (biased) noise, for the original Fashion-MNIST dataset (top panel) and for its randomized version (bottom panel). In the top panel, eLAL and ceSGD slow are hardly distinguishable.

noise in the input. For each input image, the noise was added by replacing a randomly selected fraction η of pixels with random binary values. Each new value of a pixel i was extracted with the same bias observed in the original dataset for that pixel, $\mathbb{E}_\mu[\xi_i]$: in this way, the new corrupted images were still rather close to the original input distribution (although the correlation with the desired output label, and the internal correlations between pixels, were degraded) and the networks would be able to operate in the same regime in which they were trained. This model of noise is intended to provide a rough proxy for the generalization capabilities of the networks without the need for a validation set, since it measures the amount of overfitting within a manifold which should approximate the distribution of the data to be classified. The results are shown in fig. A3, and they are fully consistent with the picture emerging from the study of the local volumes.