

This is the accepted manuscript made available via CHORUS. The article has been published as:

Properties of the Geometry of Solutions and Capacity of Multilayer Neural Networks with Rectified Linear Unit Activations

Carlo Baldassi, Enrico M. Malatesta, and Riccardo Zecchina

Phys. Rev. Lett. **123**, 170602 — Published 23 October 2019

DOI: [10.1103/PhysRevLett.123.170602](https://doi.org/10.1103/PhysRevLett.123.170602)

Properties of the geometry of solutions and capacity of multi-layer neural networks with Rectified Linear Units activations

Carlo Baldassi, Enrico M. Malatesta,* and Riccardo Zecchina

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

(Dated: September 23, 2019)

Rectified Linear Units (ReLU) have become the main model for the neural units in current deep learning systems. This choice has been originally suggested as a way to compensate for the so called vanishing gradient problem which can undercut stochastic gradient descent (SGD) learning in networks composed of multiple layers. Here we provide analytical results on the effects of ReLUs on the capacity and on the geometrical landscape of the solution space in two-layer neural networks with either binary or real-valued weights. We study the problem of storing an extensive number of random patterns and find that, quite unexpectedly, the capacity of the network remains finite as the number of neurons in the hidden layer increases, at odds with the case of threshold units in which the capacity diverges. Possibly more important, a large deviation approach allows us to find that the geometrical landscape of the solution space has a peculiar structure: While the majority of solutions are close in distance but still isolated, there exist rare regions of solutions which are much more dense than the similar ones in the case of threshold units. These solutions are robust to perturbations of the weights and can tolerate large perturbations of the inputs. The analytical results are corroborated by numerical findings.

Artificial Neural Networks (ANN) have been studied since decades and yet only recently they have started to reveal their potentialities in performing different types of massive learning tasks [1]. Their current denomination is Deep Neural Networks (DNN) in reference to the choice of the architectures, which typically involve multiple interconnected layers of neuronal units. Learning in ANN is in principle a very difficult optimization problem, in which “good” minima of the learning loss function in the high dimensional space of the connection weights need to be found. Luckily enough, DNN models have evolved rapidly, overcoming some of the computational barriers that for many years have limited their efficiency. Important components of this evolution have been the availability of computational power and the stockpiling of extremely rich data sets.

The features on which the various modeling strategies have intersected, besides the architectures, are the choice of the loss functions, the transfer functions for the neural units and the regularization techniques. These improvements have been found to help the convergence of the learning processes, typically based on Stochastic Gradient Descent (SGD) [2], and to lead to solutions which can often avoid overfitting even in over parametrized regimes. All these results pose basic conceptual questions which need to find a clear explanation in term of the optimization landscape.

Here we study the effects that the choice of the Rectified Linear Units (ReLU) for the neurons [3] has on the geometrical structure of the learning landscape. ReLU is one of the most popular non-linear activation functions and it has been extensively used to train DNN, since it is known to dramatically reduce the training time for a typical algorithm [4]. It is also known that another major benefit of using ReLU is that it does not produce the vanishing gradient problem as other common transfer functions (e.g. the tanh) [4, 5]. We study ANN models with one hidden layer storing random patterns, for which we derive analytical results that are corroborated by numerical findings. At variance with what happens in the case

of threshold units, we find that models built on ReLU functions present a critical capacity, i.e. the maximum number of patterns per weight which can be learned, that does not diverge as the number of neurons in the hidden unit is increased. At the same time we find that below the critical capacity they also present wider dense regions of solutions. These regions are defined in terms of the volume of the weights around a minimizer which do not lead to an increase of the loss value (e.g. number of errors) [6]. For discrete weights this notions reduces to the so called Local Entropy [7] of a minimizer. We also check analytically and numerically the improvement in the robustness of these solutions with respect to both weight and input perturbations.

Together with the recent results on the existence of such wide flat minima and on the effect of choosing particular loss functions to drive the learning processes toward them [8], our result contributes to create a unified framework for the learning theory in DNN, which relies on the large deviations geometrical features of the accessible solutions in the overparametrized regime.

The model. We will consider a two-layer neural network with N input units, K neurons in the hidden layer and one output. The mapping from the input to the hidden layer is realized by K non-overlapping perceptrons each having N/K weights. Given $p = \alpha N$ inputs ξ^μ labeled by index $\mu = \{1, \dots, p\}$, the output of the network for each input μ is computed as

$$\sigma_{\text{out}}^\mu = \text{sgn} \left(\frac{1}{\sqrt{K}} \sum_{l=1}^K c_l \tau_l^\mu \right) = \text{sgn} \left(\frac{1}{\sqrt{K}} \sum_{l=1}^K c_l g \left(\lambda_l^\mu \right) \right), \quad (1)$$

where λ_l^μ is the input of the l hidden unit i.e. $\lambda_l^\mu = \sqrt{\frac{K}{N}} \sum_{i=1}^{N/K} W_{li} \xi_{li}^\mu$ and W_{li} is the weight connecting the input unit i to the hidden unit l . c_l is the weight connecting hidden unit l with the output; g is a generic activation function. In the following we will mainly consider two particular choices of activation functions and of the weights c_l . In the first one we take the sign activation $g(\lambda) = \text{sgn}(\lambda)$ and we fix to 1 all

the weights c_l (in general their sign can be absorbed into the weights W_{li}). The $K = 1$ version of this model is the well-known perceptron and it has been extensively studied since the '80s by means of the replica and cavity methods [9–11] used in spin glass theory [12]. The $K > 1$ case is known as the tree-committee machine that has been studied in the '90s [13, 14]. In the second model we will use the ReLU activation function that is defined with $g(\lambda) = \max(0, \lambda)$, and since the output of this transfer function is always non-negative we will fix half of the weights c_l to +1 and the remaining half to -1. Given a training set composed by random i.i.d patterns $\xi^\mu \in \{-1, 1\}^N$ and labels $\sigma^\mu \in \{-1, 1\}$ and defining $\mathbb{X}_{\xi, \sigma}(W) \equiv \prod_\mu \theta\left(\frac{\sigma^\mu}{\sqrt{K}} \sum_{l=1}^K c_l \tau_l^\mu\right)$, the weights that correctly classify the patterns are those for which $\mathbb{X}_{\xi, \sigma}(W) = 1$. Their volume (or number) is therefore [9, 10]

$$Z = \int d\mu(W) \mathbb{X}_{\xi, \sigma}(W), \quad (2)$$

where $d\mu(W)$ is a measure over the weights W . In this study two constraints over the weights will be considered. The spherical constraint where for every $l \in \{1, \dots, K\}$, we have $\sum_i W_{li}^2 = N/K$, i.e. every sub-perceptron has weights that live on the hypersphere of radius $\sqrt{N/K}$. The second constraint we will use is the binary one, where for every $l \in \{1, \dots, K\}$ and $i \in \{1, \dots, N/K\}$ we have $W_{li} \in \{-1, 1\}$. We are interested in the large K limit for which we will be able to compute analytically the capacity of the model for different choices of transfer function, to study the typical distances between absolute minima and to perform the large deviation study giving the local volumes associated to the wider flat minima.

Critical Capacity. We will analyze the problem in the limit of a large number N of input units. The standard scenario in this limit is that there is a sharp threshold α_c such that for $\alpha < \alpha_c$ the probability of finding a solution is 1 while for $\alpha > \alpha_c$ the volume of synapses is empty. α_c is therefore called *critical capacity* since it is the maximum number of patterns per synapses that one can store in a neural network. The critical capacity of the mode, for a generic choice of the transfer function, can be evaluated computing the free entropy $\mathcal{F} \equiv \frac{1}{N} \langle \ln Z \rangle_{\xi, \sigma}$, where $\langle \cdot \rangle_{\xi, \sigma}$ denotes the average over the patterns, using the replica method; one finds

$$\mathcal{F} = \mathcal{G}_S + \alpha \mathcal{G}_E. \quad (3)$$

\mathcal{G}_S is the entropic term, which represents the logarithm of the volume at $\alpha = 0$, where there are no constraints induced by the training set; this quantity is independent on K and it is affected only by the binary or spherical nature of the weights. \mathcal{G}_E is the energetic term and it represents the logarithm of the fraction of solutions. Moreover it depends on the order parameters $q_l^{ab} \equiv \frac{K}{N} \sum_i W_{li}^a W_{li}^b$ which represent the overlap between sub-perceptrons l of two different replicas a and b of the machine. Using a replica-symmetric (RS) ansatz, in which we assume $q_l^{ab} = q$ for all a, b, l , and in the large K limit, \mathcal{G}_E

is

$$\mathcal{G}_E = \int Dz_0 \ln H\left(-\sqrt{\frac{\Delta - \Delta_{-1}}{\Delta_2 - \Delta}} z_0\right) \quad (4)$$

where $Dz \equiv \frac{dz}{\sqrt{2\pi}} e^{-z^2/2}$ and $H(x) \equiv \int_x^\infty Dz$. This expression is equivalent to that of the perceptron (i.e. $K = 1$), the only difference being that the order parameters are replaced by effective ones that depend on the general activation function used in the machine [13]. In (4) we have called these effective order parameters as Δ_{-1} , Δ and Δ_2 , see the supplementary information (SI) for details, and in the perceptron they are 0, q and 1 respectively.

In the binary case the critical capacity is always smaller than 1 and it is identified with the point where the RS free entropy \mathcal{F} vanishes. This condition requires

$$\alpha_c = \frac{\frac{\hat{q}}{2}(1-q) - \int Du \ln\left(2 \cosh\left(\sqrt{\hat{q}}u\right)\right)}{\int Dz_0 \ln H\left(-\sqrt{\frac{\Delta - \Delta_{-1}}{\Delta_2 - \Delta}} z_0\right)} \quad (5)$$

where \hat{q} is the conjugated parameter of q . q and \hat{q} are found by solving their associated saddle point equations (details in the SI). Solving these equations one finds for the ReLU case $\alpha_c = 0.9039(9)$ which is a smaller value than in the sign activation function case, where one gets $\alpha_c = 0.9469(5)$ as shown in [13].

In the spherical case the situation is different since the capacity is not bounded from above. Previous works [13, 14] have shown, in the case of the sign activations, that the RS estimate of the critical capacity diverges with the number of neurons in the hidden layer as $\alpha_c \simeq \left(\frac{72K}{\pi}\right)^{1/2}$, violating the Mitchison-Durbin bound [15]. The reason of this discrepancy is due to the fact that the Gardner volume disconnects before α_c and therefore replica-symmetry breaking (RSB) takes place. Indeed, the instability of the RS solution occurs at a finite value $\alpha_{AT} \simeq 2.988$ at large K . A subsequent work [16] based on multifractal techniques derived the correct scaling of the capacity with K as $\alpha_c \simeq \frac{16}{\pi} \sqrt{\ln K}$, which saturates the Mitchison-Durbin bound.

In the case of the ReLU functions the RS estimate of the critical capacity is obtained simply performing the $q \rightarrow 1$ limit, as for the perceptron. Quite surprisingly, if the activation function is such that $\Delta_2 - \Delta \simeq \delta\Delta(1-q)$ for $q \rightarrow 1$, with $\delta\Delta$ a finite proportionality term, the RS estimate of the critical capacity is finite (for the same reason of the finiteness of the capacity of the perceptron where one exactly has $\Delta_2 - \Delta = 1-q$). Contrary to the sign activation (where the effective parameters are such that $\Delta_2 - \Delta \simeq \delta\Delta\sqrt{1-q}$), the ReLU activation function happens to belong to this class (with $\delta\Delta = \frac{1}{2}$). The RS estimate of the critical capacity is therefore given by

$$\alpha_c^{\text{RS}} = \frac{2\delta\Delta}{\Delta_2 - \Delta_{-1}}. \quad (6)$$

One correctly recovers $\alpha_c = 2$ in the case of the perceptron [9] whereas for the committee machine with ReLU activation one has $\alpha_c = 2 \left(1 - \frac{1}{\pi}\right)^{-1} \approx 2.934$. As for the sign activation, one expects also for the ReLU activation that RS is unstable before α_c . Indeed we have computed the stability of the RS solution in the large K limit and we found $\alpha_{AT} \approx 0.615$ which is far smaller than the corresponding value of the sign activation. This suggests that strong RSB effects are at play.

We have therefore used a 1RSB ansatz to better estimate the critical capacity in the ReLU case. This can be obtained by taking the limits $q_1 \rightarrow 1$ for intra-block overlap and $m \rightarrow 0$ for the Parisi parameter. We found $\alpha_c^{1RSB} \approx 2.92$, which is not too far from the RS result.

For the sake of brevity, we mention that the results on the non divergent capacity with K generalize to other monotone smooth functions such as the sigmoid.

Typical distances. In order to get a quantitative understanding of the geometrical structure of the weight space, we have also derived the so called Franz-Parisi entropy for the committee machine with a generic transfer function. This framework was originally introduced in [17] to study the metastable states of mean-field spin glasses and only recently it was used to study the landscape of the solutions of the perceptron [18]. The basic idea is to sample a solution \tilde{W} from the equilibrium Boltzmann measure and to study the entropy landscape around it. In the binary setting, it turns out that the equilibrium solutions of the learning problem are isolated; this means that, for any positive value of α , one must flip an extensive number of weights to go from an equilibrium solution to another one. The Franz-Parisi entropy is defined as

$$\mathcal{F}_{FP}(S) = \frac{1}{N} \left\langle \frac{\int d\mu(\tilde{W}) \mathbb{X}_{\xi,\sigma}(\tilde{W}) \ln \mathcal{N}_{\xi,\sigma}(\tilde{W}, S)}{\int d\mu(\tilde{W}) \mathbb{X}_{\xi,\sigma}(\tilde{W})} \right\rangle_{\xi,\sigma} \quad (7)$$

where $\mathcal{N}_{\xi,\sigma}(\tilde{W}, S) = \int d\mu(W) \mathbb{X}_{\xi,\sigma}(W) \prod_{l=1}^K \delta(W_l \cdot \tilde{W}_l - \frac{N}{K} S)$ counts the number of solutions at a distance $d = \frac{1-S}{2}$ from a reference \tilde{W} . The distance constraint is imposed by fixing the overlap between every sub-perceptron of W and \tilde{W} to $\frac{S}{K}$. The quantity defined in eq. (7) can again be computed by the replica method. However, the expression for K finite is quite difficult to analyze numerically since the energetic term contains $4K$ integrals. The large K limit is instead easier and, again, the only difference with the perceptron expression is that the order parameters are replaced with effective ones in the energetic term (see SI for details).

In Fig. 1 we plot the Franz-Parisi entropy \mathcal{F}_{FP} as a function of the distance $d = \frac{1-S}{2}$ from a typical reference solution \tilde{W} for the committee machine with both sign and ReLU activations. The numerical analysis shows that, as in the case of the binary perceptron, also in the 2-layer case solutions are isolated since there is a minimal distance d^* below which the entropy becomes negative. This minimal distance increases with the constraint density α . However at a given value of α , typical solutions of the committee machine with ReLU activations are less isolated than the ones of the sign counterpart. The same




Figure 1. (Left panel) Typical solutions are isolated in 2 layer neural networks with binary weights. We plot the Franz-Parisi entropy \mathcal{F}_{FP} as a function of the distance from a typical reference solution \tilde{W} for the committee machine with ReLU activations (red line) and sign activations (blue line) in the limit of large number of neurons in the hidden layer K . We have used $\alpha = 0.6$. For both curves there is a value of the distance for which the entropy becomes negative. This signals that typical solutions are isolated. (Right panel) Franz-Parisi entropy as a function of distance, normalized with respect to the unconstrained $\alpha = 0$ case, for spherical weights and for $\alpha = 1$.




Figure 2. Numerical evidence of the greater robustness of the minima of ReLU transfer function (red line) compared with the sign one (blue line) using the large deviation analysis in the binary (left panel) setting and spherical setting (right panel). The exchange in the curves in both settings for α sufficiently large is due to the fact that the algorithmic threshold of the ReLU is reached before than the corresponding one of the sign case.

framework applies to the case of spherical weights where we find that the minimum distance between typical solutions is smaller for the ReLU case.

Large deviation analysis. The results of the previous section show that the Franz-Parisi framework does not capture the features of high local entropy regions. These regions indeed exist since algorithms can be observed to find solutions belonging to large connected clusters of solutions. In order to study the properties of wide flat minima or high local entropy regions one needs to introduce a large deviation measure [7], which favors configurations surrounded by an exponential number of solutions at small distance. This amounts to study a system with y real replicas constrained to be at a distance d from a reference configuration \tilde{W} . The high local entropy region is found around $d \approx 0$ in the limit of large y . As shown in [8], an alternative approach can be obtained by directly constraining

the set of y replicas to be at a given mutual distance d , that is:

$$Z_{LD}(d, y) = \int \prod_{a=1}^y d\mu(W^a) \prod_{a=1}^y \mathbb{X}_{\xi, \sigma}(W^a) \times \prod_{a < b} \delta\left(W_l^a \cdot W_l^b - \frac{N(1-2d)}{K}\right) \quad (8)$$

This last approach has the advantage of simplifying the calculations, since it is related to the standard 1RSB approach on the typical Gardner volume [19] given in equation (2): the only difference is that the Parisi parameter m and the intra-block overlap q_1 are fixed as external parameters, and play the same role of y and $1-2d$ respectively. Therefore m is not limited anymore to the standard range $[0, 1]$; indeed, the large m regime is the significant one for capturing high local entropy regions. In the large m and K limit the large deviation free entropy $\mathcal{F}_{LD} \equiv \frac{1}{N} \langle \ln Z_{LD} \rangle_{\xi, \sigma}$ reads

$$\mathcal{F}_{LD}(q_1) = \mathcal{G}_S(q_1) + \alpha \mathcal{G}_E(q_1) \quad (9)$$

where, again, the entropic term has a different expression depending on the constraint over the weights W . Its expression, together with the corresponding energetic term, is reported in the SI.

We report in Fig. 2 the numerical results for both binary and spherical weights of the large deviation entropy (normalized with respect to the unconstrained $\alpha = 0$ case) as a function of q_1 . For both sign and ReLU activations, the region $q_1 \simeq 1$ is flat around zero. This means that there exist \tilde{W} references around which the landscape of solutions is basically indistinguishable from the $\alpha = 0$ case where all configurations are solutions. We also find that the ReLU curve, in the vicinity of $q_1 \simeq 1$, is always more entropic than the corresponding one of the sign. This picture is valid for sufficiently low values of α ; for α greater of a certain value α^* the two curves switch. This is due to the fact that the two models have completely different critical capacities (divergent in the sign case, finite in the ReLU case) so that one expects that clusters of solutions disappear at a smaller constrained capacity when ReLU activations are used.

Stability distribution and robustness. To corroborate our previous results, we have also computed (details in the SI), for various models and types of solutions W , the distribution of the *stabilities* $\Xi = \left\langle \frac{\sigma}{\sqrt{K}} \sum_{l=1}^K c_l \tau_l^\mu \right\rangle_{\xi, \sigma}$, which measure the distance from the threshold at the output unit in the direction of the correct label σ , cf. eq. (1). Previous calculations [20] have shown that in the simple case of the spherical perceptron at the critical capacity the stability distribution around a typical solution W develops a delta peak in the origin $\Xi \simeq 0$; we confirmed that even in the two-layer case the stability distribution of a typical solution, being isolated, also has its mode at $\Xi \simeq 0$ even at lower α , see the dashed lines in Fig. 3 (left). A solution surrounded by an exponential number of other solutions, instead, should be more robust and be centered away from 0. Our calculations show that this is indeed the case both for the sign and




Figure 3. (Left panel) Dashed lines: theoretical stability curves for the typical solutions, for binary weights at $\alpha = 0.4$. Solid lines: comparison between the numerical and theoretical stability distributions in the large deviation scenario, same α . We have used $q_1 = 0.85$, $y = 20$. (Right panel) Robustness of the reference configuration found by replicated simulated annealing when one pattern is perturbed by flipping a certain fraction η of entries.

for the ReLU activations, and we have confirmed the results by numerical simulations. In Fig. 3 (left) we show the analytical and numerical results for the case of binary weights at $\alpha = 0.4$ with $y = 20$ replicas at $q_1 = 0.85$. For the numerical results, we have used simulated annealing on a system with $K = 32$ ($K = 33$) for the ReLU (sign) activations (respectively), and $N = K^2 \simeq 10^3$. We have simulated a system of y interacting replicas that is able to sample from the local-entropic measure [6] with the RRR Monte Carlo method [21], ensuring that the annealing process was sufficiently slow such that at the end of the simulation all replicas were solutions, and controlling the interaction such that the average overlap between replicas was equal to q_1 within a tolerance of 0.01. The results were averaged over 20 samples. As seen in Fig. 3 (left), the agreement with the analytical computations is remarkable, despite the small values of N/K and K and the approximations introduced by sampling with simulated annealing.

The stabilities for the sign and ReLU activations are qualitatively similar, but quantitatively we observe that in all cases the curves for the ReLU case have a peak closer to 0 and a smaller variance. These are not, however, directly comparable, and it is difficult to tell from the stability curves alone which choice is more robust. We have thus directly measured, on the results of the simulations, the effect of introducing noise in the input patterns. For each trained group of y replicas, we used the configuration of the reference \tilde{W} (which lays in the middle of the cluster of solutions) and we measured the probability that a pattern of the training set would be misclassified if perturbed by flipping a fraction η of randomly chosen entries. We explored a wide range of values of the noise η and sampled 50 perturbations per pattern. The results are shown in Fig. 3 (right), and they confirm that the networks with ReLU activations are more robust than those with sign activations for this α , in agreement with the results of Fig. 2. We also verified that the reference configuration is indeed more robust than the individual replicas. The results for other choices of the parameters are qualitatively identical. Our preliminary tests show that the same phenomenology is maintained when the network architecture is changed to a fully-connected scheme, in which

each hidden unit is connected to all of the input units.

The architecture of the model that we have analyzed here is certainly very simplified compared to state of the art deep neural networks used in applications. Investigating deeper models would certainly be of great interest, but extremely challenging with current analytical techniques, and is thus an open problem. Extending our analysis to a one-hidden-layer fully-connected model, on the other hand, would in principle be feasible (the additional complication comes from the permutation symmetry of the hidden layer). However, based on the existing literature (e.g. [13, 22]), and our preliminary numerical experiments mentioned above, we do not expect that such extension would result in qualitatively different outcomes compared to our tree-like model.

* enrico.malatesta@unibocconi.it

- [1] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015. doi:[10.1038/nature14539](https://doi.org/10.1038/nature14539).
- [2] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer, 2010. doi:[10.1007/978-3-7908-2604-3_16](https://doi.org/10.1007/978-3-7908-2604-3_16).
- [3] Richard HR Hahnloser, Rahul Sarpeshkar, Misha A Mahowald, Rodney J Douglas, and H Sebastian Seung. Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature*, 405(6789):947, 2000. doi:[10.1038/35016072](https://doi.org/10.1038/35016072).
- [4] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In Geoffrey Gordon, David Dunson, and Miroslav Dudík, editors, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 315–323, Fort Lauderdale, FL, USA, 11–13 Apr 2011. PMLR. URL: <http://proceedings.mlr.press/v15/glorot11a.html>.
- [5] Sepp Hochreiter. Untersuchungen zu dynamischen neuronalen netzen. *Diploma, Technische Universität München*, 91(1), 1991.
- [6] 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 2016. doi:[10.1073/pnas.1608103113](https://doi.org/10.1073/pnas.1608103113).
- [7] 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. *Phys. Rev. Lett.*, 115:128101, Sep 2015. doi:[10.1103/PhysRevLett.115.128101](https://doi.org/10.1103/PhysRevLett.115.128101).
- [8] Carlo Baldassi, Fabrizio Pittorino, and Riccardo Zecchina. Shaping the learning landscape in neural networks around wide flat minima. *arXiv preprint arXiv:1905.07833*, 2019. URL: <https://arxiv.org/abs/1905.07833>.
- [9] Elizabeth Gardner. The space of interactions in neural network models. *Journal of Physics A: Mathematical and General*, 21(1):257–270, jan 1988. doi:[10.1088/0305-4470/21/1/030](https://doi.org/10.1088/0305-4470/21/1/030).
- [10] Elizabeth Gardner and Bernard Derrida. Optimal storage properties of neural network models. *Journal of Physics A: Mathematical and General*, 21(1):271–284, jan 1988. doi:[10.1088/0305-4470/21/1/031](https://doi.org/10.1088/0305-4470/21/1/031).
- [11] Marc Mézard. The space of interactions in neural networks: Gardner’s computation with the cavity method. *Journal of Physics A: Mathematical and General*, 22(12):2181, 1989. doi:[doi:doi.org/10.1088/0305-4470/22/12/018](https://doi.org/10.1088/0305-4470/22/12/018).
- [12] Hidetoshi Nishimori. *Statistical Physics of Spin Glasses and Information Processing: An Introduction*. International series of monographs on physics. Oxford University Press, 2001. doi:[10.1093/acprof:oso/9780198509417.001.0001](https://doi.org/10.1093/acprof:oso/9780198509417.001.0001).
- [13] 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](https://doi.org/10.1103/PhysRevA.45.4146).
- [14] A. Engel, H. M. Köhler, F. Tschepeke, H. Vollmayr, and A. Zipfelius. Storage capacity and learning algorithms for two-layer neural networks. *Phys. Rev. A*, 45:7590–7609, May 1992. doi:[10.1103/PhysRevA.45.7590](https://doi.org/10.1103/PhysRevA.45.7590).
- [15] 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](https://doi.org/10.1007/BF00204772).
- [16] 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](https://doi.org/10.1103/PhysRevLett.75.2432).
- [17] 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](https://doi.org/10.1051/jp1:1995201).
- [18] Haiping Huang and Yoshiyuki Kabashima. Origin of the computational hardness for learning with binary synapses. *Physical Review E*, 90(5):052813, 2014. doi:<https://doi.org/10.1103/PhysRevE.90.052813>.
- [19] 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](https://doi.org/10.1051/jphys:0198900500200305700).
- [20] Thomas B Kepler and Laurence F Abbott. Domains of attraction in neural networks. *Journal de Physique*, 49(10):1657–1662, 1988. doi:[10.1051/jphys:0198800490100165700](https://doi.org/10.1051/jphys:0198800490100165700).
- [21] Carlo Baldassi. A method to reduce the rejection rate in monte carlo markov chains. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(3):033301, 2017. doi:[10.1088/1742-5468/aa5335](https://doi.org/10.1088/1742-5468/aa5335).
- [22] R Urbanczik. Storage capacity of the fully-connected committee machine. *Journal of Physics A: Mathematical and General*, 30(11):L387–L392, jun 1997. URL: <https://doi.org/10.1088/0305-4470/30/11/007>.