Data-driven effective model shows a liquid-like deep learning

Wenxuan Zou and Haiping Huang

PMI Lab, School of Physics, Sun Yat-sen University, Guangzhou 510275, People's Republic of China

(Received 28 July 2021; accepted 13 September 2021; published 30 September 2021)

The geometric structure of an optimization landscape is argued to be fundamentally important to support the success of deep neural network learning. A direct computation of the landscape beyond two layers is hard. Therefore, to capture the global view of the landscape, an interpretable model of the network-parameter (or weight) space must be established. However, the model is lacking so far. Furthermore, it remains unknown what the landscape looks like for deep networks of binary synapses, which plays a key role in robust and energy efficient neuromorphic computation. Here, we propose a statistical mechanics framework by directly building a least structured model of the high-dimensional weight space, considering realistic structured data, stochastic gradient descent training, and the computational depth of neural networks. We also consider whether the number of network parameters outnumbers the number of supplied training data, namely, over- or under-parametrization. Our least structured model reveals that the weight spaces of the under-parametrization and over-parameterization cases belong to the same class, in the sense that these weight spaces are well connected without any hierarchical clustering structure. In contrast, the shallow-network has a broken weight space, characterized by a discontinuous phase transition, thereby clarifying the benefit of depth in deep learning from the angle of high-dimensional geometry. Our effective model also reveals that inside a deep network, there exists a liquid-like central part of the architecture in the sense that the weights in this part behave as randomly as possible, providing algorithmic implications. Our data-driven model thus provides a statistical mechanics insight about why deep learning is unreasonably effective in terms of the high-dimensional weight space, and how deep networks are different from shallow ones.

DOI: [10.1103/PhysRevResearch.3.033290](https://doi.org/10.1103/PhysRevResearch.3.033290)

I. INTRODUCTION

Artificial deep neural networks have achieved state-of-the-art performance in many industrial and academic domains ranging from pattern recognition and natural language processing [1] to many-body quantum physics and classical statistical physics [2]. However, it remains challenging to reveal the mechanisms underlying the success of deep learning. One key obstacle is building the causal relationship between the high-dimensional nonconvex optimization landscape and the state-of-the-art performance of deep networks [3]. In the past few years, many theoretical and empirical efforts contributed to understanding this fundamental relationship in deep networks of real-valued weights. High-loss local minima of the optimization landscape are rare in the over-parametrization regime where the number of trainable parameters (weights) in a typical deep network is much larger than the number of training data [4]. All minima in the landscape were shown to be global minima, given special requirements for network architectures and neural activation functions [5]. In particular, no substantial barriers between

minima were observed in deep learning [6]. These minima can be even connected via hyper-paths within a unique global valley [7–9]. The geometric structure of local minima was also extensively studied [10–13], in terms of a controversial description of the geometric flatness of the minima [14–18]. The flat minima imply that the weights around them can be perturbed without significantly changing the training cost function. The flatness of such minima was empirically supported by the eigenspectrum of the Hessian matrix evaluated at those minima [11]. The curvature of the landscape was also claimed to correlate with the generalization ability of deep networks for unseen data [18,19]. Therefore, studying deep networks from the perspective of the high-dimensional weight space is fundamentally important, yet the previous studies did not give a global view about the collective behavior of interacting weights.

Physics methods have been applied to investigate the weight space structure of nonconvex high-dimensional problems, e.g., the spherical or binary spin model approximation of deep neural networks [20–23], and the weight space structure of binary perceptron problems [24–26]. However, one drawback of these studies concerns their strong assumptions to build toy models. These models assume unstructured data as well as shallow networks, and moreover the learning behavior is far from the practical training procedure used in a typical deep learning. Therefore, it is computationally hard to directly compute the weight space even beyond two-layer architectures, and moreover it remains elusive from an interpretable-model perspective what the geometric structure

*huanghp7@mail.sysu.edu.cn

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

FIG. 1. Schematic illustration of how to construct an effective physics model analyzing the deep-learning weight space. We consider a four-layer neural network of binary weights to perform a classification task, where the input data pass through the weight regions of the left boundary (LB), the interior (ITR), and the right boundary (RB) in order. After the neural network is trained from different initializations, the weight statistics including magnetizations (first-order moments) $m_i = \langle \sigma_i \rangle$ and pairwise correlations (second-order moments) $C_{ij} = \langle \sigma_i \sigma_j \rangle$ will be computed directly from the collected weight configurations. Based on the weight statistics, we establish an effective data-driven Ising model that is able to reflect the weight space characteristics, and then apply the entropy landscape analysis by introducing a distance-dependent extra term into the Boltzmann distribution of the effective Ising model. The entropy landscape analysis can help to explore the internal structure of the weight parameter space by counting how dense weight parameters σ are at a typical distance d away from the reference one σ^* .

of the deep-learning landscape looks like and which key factor affects the landscape. In this paper, we restrict the theoretical analysis to deep networks of binary synapses, which is particularly important for robust and energy-efficient neuromorphic computation [27,28].

Here, from a conceptually different viewpoint, we propose a data-driven effective model, taking into account interactions among structured data, deep architectures, and training. We consider deep networks of binary weights that are amenable for theoretical analysis and moreover can be efficiently trained with mean-field methods [29]. The weight configurations realizing a high classification accuracy of the structured data are collected as a large ensemble, whose statistics is described by first-order and second-order moments of the weight distribution. To construct an effective model of the weight space, the maximum-entropy principle [30] is applied. This principle was previously used to analyze neural population at the collective network-activity level [31,32]. This effective model of the practical deep learning is then analyzed from an entropy landscape angle, in which the geometric structure of the entire weight space can be characterized in different contexts: over-parameterization, under-parameterization and shallow networks.

Our paper establishes an interpretable model of the weight space focusing on the global view of the landscape, and this model has three important predictions: (i) The deep-learning weight-space is smooth and dominated by a single well-connected component, while the shallow network is characterized by a broken weight space. (ii) For the deep network, the under-parameterization and over-parameterization (the number of training data is more and less than the number of parameters, respectively) belong to the same universal class with a common smooth landscape, which exactly coincides with recent empirical findings of no substantial barriers between minima in the deep-learning loss landscape [6–8], although we focus on the deep networks of binary weights. We remark that the universal class may be related to the absence of the double descent phenomenon for the deep networks of binary weights, at least in the explored range of training parameters. (iii) A most surprising prediction is that a special

interior part of a largest entropy in the weight space emerges after learning, showing a liquid-like property, compared to two more constrained boundaries of the deep network. This part is conjectured to play a key role in fast learning dynamics.

II. DEEP LEARNING SETTING

In order to collect weight configurations in the neural network parameter space, we design a neural network architecture to solve the classification task of the handwritten digits (or the MNIST dataset [33]) first, and then an extensive number of weight solutions that lead to high test accuracy are collected. The test accuracy refers to the network's ability to classify unseen dataset. More precisely, we consider a four-layer fully-connected feedforward network. Each layer has N_l ($l = 0, 1, 2, 3$) neurons, where N_0 equals to the dimension of the input vector, and N_3 equals to the number of total classes. The weight matrix $\mathbf{W}^l$ indicates connections between layer l and layer $l - 1$ (Fig. 1). For simplicity, we do not consider biases for neurons. The neural activity (including the input) is transferred by a nonlinear function, i.e., the leaky ReLU (L-ReLU) function defined by $\max(0, z) + 0.1 \min(0, z)$, where z denotes the preactivation or the weighted sum of input activities. We finally add a softmax layer after the output layer and use the cross entropy as the loss function.

To collect a sufficient number of weight configurations and facilitate theoretical analysis, we require that the neural network architecture should meet two demands. First, the model complexity, namely, the number of the total weights should be small. This avoids a sharp growth of the computational cost consumed to collect a huge number of weight configurations. Second, the weight of the neural network should take binary values $\{+1, -1\}$, as the following analyses are based on the Ising model, in which the weights are analogous to the spins in a physics model.

To fulfill the first demand, we should first create a set of input data with low dimension. Thus, we carry out a principal component analysis (PCA) of the MNIST dataset, retaining

only the top 20 principal components, which explain about 64.5% of the total variance in the original data. After this dimension reduction, a relatively high-quality reconstruction of the digits can be performed. Although the loss of information will degrade the test accuracy, we can still reach a relatively high test accuracy (see Appendix D).

Because of the second demand, the network with binary weights is difficult to train by the standard gradient decent algorithm, due to the fact that the cost function is not differentiable with respect to the discrete synaptic weight. To tackle this challenge, we apply a recently-proposed mean-field training algorithm [29]. More precisely, we consider the situation that each weight w_{ij}^l follows a binomial distribution $P(w_{ij}^l)$ parametrized by an external field θ_{ij}^l as follows:

$$P(w_{ij}^l) = \sigma(\theta_{ij}^l) \delta_{w_{ij}^l, 1} + [1 - \sigma(\theta_{ij}^l)] \delta_{w_{ij}^l, -1}, \quad (1)$$

where $\sigma(\cdot)$ is a sigmoid function. The external fields are clearly continuous. Therefore, the standard backpropagation can be applied to the external fields $\{\theta_{ij}^l\}$, instead of weights (ill-defined in gradients).

In the feedforward transformation, the preactivation can be written as $z_j^l = \sum_i w_{ij}^l a_i^{l-1}$, where a_i^{l-1} is the activation at the layer $l-1$. The weight w_{ij}^l is subject to the aforementioned binomial distribution, and thus its mean and variance, i.e., μ_{ij}^l and $(\sigma_{ij}^l)^2$, are given by

$$\mu_{ij}^l = \langle w_{ij}^l \rangle = 2\sigma(\theta_{ij}^l) - 1, \quad (2a)$$

$$(\sigma_{ij}^l)^2 = \langle (w_{ij}^l)^2 \rangle - \langle w_{ij}^l \rangle^2 = -4\sigma^2(\theta_{ij}^l) + 4\sigma(\theta_{ij}^l). \quad (2b)$$

To avoid a direct sampling process, we use the local reparametrization trick [34]. According to the central limit theorem, z_j^l can then be approximated to follow a Gaussian distribution $z_j^l \sim \mathcal{N}(\sum_i \mu_{ij}^l a_i^{l-1}, \sum_i (\sigma_{ij}^l)^2 (a_i^{l-1})^2)$. Hereafter, we define the mean and variance of z_j^l to be m_j^l and $(v_j^l)^2$, respectively. Taken together, the feedforward transformation can be finally reparametrized as

$$z_j^l = m_j^l + v_j^l \cdot \epsilon_j^l, \quad (3a)$$

$$a_j^l = \frac{1}{\sqrt{N_{l-1}}} \text{L-ReLU}(z_j^l), \quad (3b)$$

where ϵ_j^l is a standard Gaussian random variable, and $\frac{1}{\sqrt{N_{l-1}}}$ is a scaling factor to eliminate the layer-width dependence of the activation a_j^l . Detailed training procedures are given to Appendix A.

In practice, we use two hidden layers with 15 neurons for each. The widths of input layer and output layer are determined by the data, which are 20 and 10, respectively. Thus, our network architecture is specified by 20-15-15-10, resulting in the number of weight parameters $N = 675$. We highlight that the total number of possible weights $2^N \approx 10^{203}$ is already very huge for a statistical analysis. As expected, the size of hidden layer makes learning improve over the shallow counterpart (see Appendix F). Unless stated otherwise, we use the entire MNIST dataset (i.e., 60 000 digits examples) for training and the other 10 000 examples for testing. In the following analysis, we denote P as the number of training examples. By fine tuning hyper-parameters of the training

procedure (see Appendix D), the neural network can reach a test accuracy of 85% [Fig. 2(a)], i.e., the trained network can classify correctly 85% of the unseen test dataset. To collect distinct weight configurations, we train the neural network for 50 epochs starting from a random Gaussian initialization of external fields θ during each sampling trial. Then, we use the trained external fields to generate 10 weight configurations and select the one with the optimal test accuracy ($\geq 80\%$). In this way, we obtain two million weight configurations (the corresponding number is specified by M below) with test accuracies ranged from 80% to 83%, for constructing an effective model.

To give a brief picture of the weight parameter space, we apply the t-SNE dimension reduction technique [35] to 20 000 random initialization points and the corresponding learned weight configurations with the same number [Fig. 2(b)]. Here, random initialization points represent the weight parameters sampled from the Gaussian-initialized external fields, from which the external fields are trained to reach the test accuracy threshold. The collected weight configurations are sampled from the final binomial distribution. Note that the dominant part of the weight configuration space (valid weights realizing the classification task) seems to be organized into a large connected component [Fig. 2(b)], which also indicates that compared to the entire parameter space, the weight space realizing the classification task occupies a relatively small region. Next, we design a semi-rigorous statistical mechanics framework to explore a detailed high-dimensional geometric structure of the weight space for the current deep-learning model.

III. DATA-DRIVEN ISING MODEL

We build a data-driven Ising model to describe these M weight configurations, as we are interested in the statistical properties of the deep-learning weight space. More precisely, one weight parameter solution $\mathbf{W} = \{w_1, w_2, \dots, w_N\}$ is transformed to a spin configuration $\sigma = \{\sigma_1, \dots, \sigma_N\}$ in the following analysis. Given the constraints of weight statistics including magnetization (first-order moments) $m_i = \langle \sigma_i \rangle$ and pairwise correlation (second-order moments) $C_{ij} = \langle \sigma_i \sigma_j \rangle$, the least-structured-model probability distribution $P(\sigma)$ to fit the weight statistics follows the maximum entropy principle [30] (see also Fig. 1 for an illustration). According to this principle, $P(\sigma)$ is recast into the form of the Boltzmann distribution:

$$P(\sigma) = \frac{1}{Z} \exp(-\beta E(\sigma)), \quad (4)$$

where Z is the partition function in statistical physics, β is the inverse temperature ($\beta = 1$ throughout the paper as it can be absorbed into the model parameters), and $E(\sigma) = -\sum_{i < j} J_{ij} \sigma_i \sigma_j - \sum_i h_i \sigma_i$ is the energy that is consistent with the Ising-model's Hamiltonian. The model parameter J_{ij} denotes the functional coupling between weight σ_i and weight σ_j interpreting the second-order correlation among weights, where $J_{ij} > 0$ and $J_{ij} < 0$ refer to ferromagnetic and antiferromagnetic interactions, respectively. The model parameter h_i denotes the bias of the weight σ_i , reflected by the external field on the weight σ_i after training.

FIG. 2. Neural network training and properties of weight configurations sampled from independent runs of training ($P = 60\,000$, $N = 675$). (a) Training trajectories of four-layer networks. The cross-entropy (the orange line) and test accuracy (the blue line) can reach 0.61 and 83.5%, respectively at the 200th epoch, where the fluctuation indicated by the shadow regime is computed from ten independent training runs. (b) Two-dimensional representations of 20 000 random initializations (green dots) and 20 000 learned weight configurations (blue dots). The low dimensional projection is carried out by the t-SNE technique [35]. The random initialization denotes the weight configuration (or initial state) sampled by the Gaussian random initial fields θ_0 . The collected weight configurations (or solution states) are obtained at the end of the training starting from the initial fields θ_0 . For the solution states, the deeper their colors are, the higher the corresponding test accuracies are, while the random initial states are all of the same green color of chance-level accuracy.

To find the model parameters $\Omega = \{\mathbf{J}, \mathbf{h}\}$, we use the gradient ascent algorithm corresponding to the maximum likelihood learning principle. By maximizing the log-likelihood $\langle \ln(P_{\Omega}(\sigma)) \rangle_{\text{weight-data}}$ with respect to the model parameters $\Omega = \{\mathbf{J}, \mathbf{h}\}$, we can obtain the following iterative learning rules [36]:

$$J_{ij}^{t+1} = J_{ij}^t + \eta(\langle \sigma_i \sigma_j \rangle_{\text{data}} - \langle \sigma_i \sigma_j \rangle_{\text{model}}), \quad (5a)$$

$$h_i^{t+1} = h_i^t + \eta(\langle \sigma_i \rangle_{\text{data}} - \langle \sigma_i \rangle_{\text{model}}), \quad (5b)$$

where t and η denote the learning step and learning rate, respectively. In the above learning equations, the data expectation terms can be directly computed from the sampled parameter solutions. However, the model expectation terms are commonly difficult to compute due to the $\mathcal{O}(2^N)$ computation complexity. Fortunately, the cavity method in spin glass theory (Appendix C) can be used to approximate the model expectations. At each learning step, to evaluate how well the Ising model fits the weight statistics, we compute the root-mean squared error (or deviation) between the data expectations and model expectations: $\Delta^2 = \frac{1}{N} \sum_i [m_i^{\text{data}} - m_i^{\text{model}}]^2 + \frac{2}{N(N-1)} \sum_{i < j} [C_{ij}^{\text{data}} - C_{ij}^{\text{model}}]^2$, from which a perfect fitting leads to a zero deviation. We stop the learning when $\Delta < 0.01$ or after 300 learning steps are reached.

In Fig. 3(b), we verify that the reconstructed magnetizations and correlations from the data-driven Ising model are in good agreement with the measured ones, which ensures that the data-driven Ising model is an effective model whose model parameters $\Omega = \{\mathbf{J}, \mathbf{h}\}$ capture statistical properties of the deep-learning weight space. Moreover, we also find that the data-driven Ising model can reproduce a significant fraction of three-weight correlations (Appendix B). The histogram of the model parameters is shown in Fig. 3(a). Note that most of the couplings are concentrated around zero value for the full model, despite a relatively less frequent tail. In contrast, the bias is broadly distributed. To get more insights about the model parameters, we plot the histogram for different subparts of the deep network [Fig. 3(d)]. Compared to the boundaries, i.e., LB and RB, the interior's couplings are more concentrated around zero value, and moreover the corresponding biases are

all concentrated around -0.0065 corresponding to the highest peak in the histogram of Fig. 3(a), which clearly implies that weights in the interior of the deep network have weak correlations, and moreover the magnetization gets close to zero, behaving like in a high-temperature phase, as we shall prove semi-rigorously using the concept of entropy landscape. Interestingly, the number of peaks in the LB's (respective RB's) bias distribution approximately corresponds to the number of input (respective output) neurons. This observation demonstrates that the LB's and RB's weights perform redundant encoding and robust decoding, respectively.

Given the model parameters of the constructed Ising model, we can compute the model energy of every weight configuration directly. Figure 3(c) shows the model energy density versus the test accuracy. Although the fluctuation is large, we find that the energies of sampled weight-configurations distribute around the typical energy of the Ising model and they have high accuracies, which helps to build an intuitive relationship between the model energy landscape and the deep-learning landscape.

To cross check the qualitative properties of the constructed Ising model, we also apply the pseudolikelihood maximization method to construct the Ising model, relying on the entire configuration data, rather than only the first two moments of the statistics. This method allows us to control the sparsity of the effective Ising model, i.e., the number of null couplings. Introducing the null couplings in the model is physically intuitive, since weights in different layers may not strongly interact with each other. Therefore, we construct sparse Ising models only for deep networks.

As an alternative strategy, the pseudolikelihood method requires the entire dataset of the sampled weight configurations $\{\sigma^\mu\}_{\mu=1}^M$, rather than the one- and two-point correlation functions. The objective function to be maximized is written as follows [37]:

$$\mathcal{PL} = \sum_i \mathcal{L}_i, \quad (6a)$$

FIG. 3. Properties of the data-driven Ising model ($P = 60\,000$, $N = 675$). The Ising model is constructed by the maximum entropy method. The total area of the histogram is normalized to one. (a) Histograms of full-model (all) parameters. Couplings J_{ij} and biases h_i (inset) are shown. (b) Reconstructed correlation C_{ij} and magnetization m_i (inset) from the effective Ising model vs the measured ones for the full model. Full lines indicate the equality. (c) Ising energy density ϵ_{Ising} vs test accuracy of the full model. Each accuracy corresponds to a large number of weight configurations, where the blue-solid line indicates expectation values and the light-blue shadow denotes the fluctuation. The blue-dashed line indicates the typical energy of the Ising model. (d) Histograms of model parameters in subparts (see Fig. 1) of the deep hierarchy. Couplings J_{ij} and biases h_i (inset) are shown.

$$\mathcal{L}_i = \frac{1}{M} \sum_{\mu=1}^M \ln P(\sigma_i^\mu | \sigma_{\setminus i}^\mu), \quad (6b)$$

where the conditional probability of the spin σ_i given the rest of the system $P(\sigma_i^\mu | \sigma_{\setminus i}^\mu)$ is given below:

$$P(\sigma_i | \sigma_{\setminus i}) = [1 + e^{-2\beta\sigma_i(h_i + \sum_{j \neq i} J_{ij}\sigma_j)}]^{-1}. \quad (7)$$

Therefore, the pseudolikelihood method simplifies the original N -body problem into N independent one-body problems (but each with $N - 1$ quenched spin variables). The optimization can then be implemented through maximizing each of the N local likelihood functions $\mathcal{L}_i$ separately, and finally the symmetric coupling $J_{ij} = \frac{1}{2}(J_{ij}^i + J_{ij}^j)$, where J_{ij}^x is the coupling value for the maximum of $\mathcal{L}_x$. The maximization can be done by simple gradient ascent procedures [38]. To learn a sparse Ising model, one can rewrite the local likelihood with ℓ_1 regularization as $\tilde{\mathcal{L}}_i = \mathcal{L}_i + \lambda \sum_{j \neq i} |J_{ij}|$, where λ is a freely-tuned parameter controlling the sparsity of the model. The numerical maximization can be efficiently done by the interior point method [39], where λ is the prefactor of λ_{max} [39]. This regularized optimization leads to a sparse effective model where non-zero and null couplings coexist, thereby distinguishing strongly interacting weights from noninteracting ones. To save the computational cost, we fix $M = 100\,000$ in the pseudolikelihood maximization learning.

As shown in Fig. 4, the aforementioned properties are conserved. In fact, the null couplings explain the weakly interacting weights across layers, corresponding to negligible

values of couplings in the dense counterpart. Thus it can be expected that the qualitative behavior does not change for both methods.

IV. ENTROPY LANDSCAPE ANALYSIS

In the previous section, we construct an effective Ising model, whose model parameters preserve the information up to second order correlations of the weight space. In order to explore the internal structure of the weight space of the deep learning, we introduce a distance-dependent term $x \sum_i \sigma_i^* \sigma_i$ in the original Boltzmann distribution [Eq. (4)] as follows [32]:

$$P(\sigma) = \frac{1}{Z} \exp \left(\beta \sum_{i < j} J_{ij} \sigma_i \sigma_j + \beta \sum_i h_i \sigma_i + x \sum_i \sigma_i^* \sigma_i \right), \quad (8)$$

where σ^* is the reference weight configuration taken from the collected weight data, its energy is computed from the Ising Hamiltonian, and x is the coupling field constraining the distance between the configuration σ and the reference one σ^* . Intuitively, $x > 0$ implies that the weight configurations closer to the reference are preferred, while $x < 0$ implies that the weight configuration far away from the reference are preferred (see Fig. 1 for an illustration).

Then the geometric structure of the high-dimensional weight space can be summarized into the constrained free

FIG. 4. Properties of the data-driven Ising model ($P = 60\,000$, $N = 675$). The Ising model is constructed by the pseudolikelihood maximization method ($\lambda = 0.05$). The ℓ_1 regularization training leads to a sparse model where about 73.5% of weights are zero. (a) Histograms of full-model (all) parameters. Couplings J_{ij} and biases h_i (inset) are shown. The total area of the histogram is normalized to one. (b) Reconstructed correlation C_{ij} and magnetization m_i (inset) from the effective Ising model vs the measured ones for the full model. Full lines indicate the equality.

energy function given below:

$$\beta f \equiv -\frac{1}{N} \ln \sum_{\sigma} \exp \left(-\beta E(\sigma) + x \sum_i \sigma_i^* \sigma_i \right), \quad (9)$$

where $E(\sigma)$ is the Hamiltonian of the effective model. We further introduce the energy density $\epsilon = E(\sigma)/N$ and the overlap $q = \sum_i \sigma_i \sigma_i^*/N$. It then follows that

$$\begin{aligned} \beta f &= -\frac{1}{N} \ln \iint d\epsilon dq \exp \left(-N\beta\epsilon + Nxq + \ln \sum_{\sigma} \delta \left(\epsilon - \frac{E(\sigma)}{N} \right) \delta \left(q - \frac{\sum_i \sigma_i \sigma_i^*}{N} \right) \right) \\ &= -\frac{1}{N} \ln \iint d\epsilon dq \exp (-N\beta f(\epsilon, q)), \end{aligned} \quad (10)$$

where $\beta f(\epsilon, q) \equiv \beta\epsilon - xq - s(\epsilon, q)$ is the constrained free energy as a function of energy density ϵ and overlap q . This constrained free energy captures the competition between the model energy and entropy. The entropy is defined by $s(\epsilon, q) = 1/N \ln \sum_{\sigma} \delta(\epsilon - \frac{E(\sigma)}{N}) \delta(q - \frac{\sum_i \sigma_i \sigma_i^*}{N})$, measuring the log-number of the weight configurations with energy $N\epsilon$ and overlap q given a reference.

When the system size N (675 in this paper) is large in Eq. (10), the Laplace method can be used to approximate the integral by taking only the dominant contribution. Hence, $f \approx \min_{\epsilon, q} f(\epsilon, q)$. To estimate the minimal free energy density $f(\epsilon, q)$ and the corresponding saddle-point values (ϵ, q) , we iteratively solve the mean-field equations based on the cavity approximation (Appendix C). Then, the entropy density $s(\epsilon, q)$ is calculated via the double Legendre transformation $s(\epsilon, q) = -\beta f + \beta\epsilon - xq$. For the convenience in the further analyses, we introduce the Hamming distance $(N - \sum_i \sigma_i \sigma_i^*)/2$, which counts the number of weights that are distinct between the target configuration σ and the reference one σ^* . The Hamming distance per weight d can thus be transformed from the overlap by $d = (1 - q)/2$. In the following analysis, we omit the energy density dependence of $s(\epsilon, q)$ and write it as the function of d , i.e., $s(d)$. We also assume $\beta = 1$, as the inverse temperature has been absorbed

into the inferred couplings and biases of the effective model (see Sec. III).

In the case of sparse Ising models learned by the pseudolikelihood maximization method, the cavity approximation may break. However, to estimate the entropy, one can use the heat-capacity integration technique (see Appendix G), despite its demanding computational cost.

To have a complete understanding of the deep-learning landscape, we analyze three representative scenarios. The first one is the under-parametrization case, in which the number of training examples is larger than the network-parameter size N . The second one is the over-parametrization case where the number of training examples is smaller than the network-parameter size N , and the last one is the shallow-network case where all hidden layers are deleted (for more details, see the Appendix D).

We further define the ratio between the number of network parameters and the number of training examples as $\alpha = N/P$. The point of $\alpha = 1$ is a balance point where the number of network parameters matches exactly with the number of training examples. In the deep learning of real-valued weights, with increasing values of α , the test error profile displays the double-descent behavior [40,41], where an interpolation threshold marks the model complexity where the training

FIG. 5. Entropy landscape analyses in three different scenarios of deep learning. (a) Under-parametrization case ($\alpha \approx 0.01$). The curves are the mean results over 20 distinct reference weight configurations σ^* selected from the low energy region (for references coming from high- and medium-energy regions, see Fig. 13 and Fig. 14 in Appendix D). Entropy curves are plotted along the direction of the decreasing coupling field x from 3 to -3 , while the Hamming distance per weight d vs x is plotted by considering two directions of increasing and decreasing x (inset). No hysteresis loops and entropy gaps are observed. (b) Over-parametrization case ($\alpha \approx 1.35$). Other conditions are the same as (a). (c) Shallow network case. The inset shows a hysteresis loop for distance d vs coupling field x , and the entropy curves correspond to the lower-branch of the hysteresis loop. A first-order phase transition arises in this case showing a complex landscape in the high-dimensional weight space.

error starts to vanish, and at the same location, a peak of test error emerges. The interpolation point $\alpha \sim \mathcal{O}(1)$, and the regime around the interpolation threshold is named the critical regime. In contrast, we never observe the double descent behavior around $\alpha \sim \mathcal{O}(1)$ for deep learning of binary weights (see details in Appendix E).

A. Under-parametrization scenario

In the under-parametrization case, the number of supplied training examples is much larger than the network-parameter size (or the model complexity). Interestingly, we do not find any weight-space-structure induced phase-transition. First, the distance d increases or decreases continuously following either of two directions of changing the coupling field [Fig. 5(a)]; both directions collapse into a single curve, without any hysteresis phenomenon. This provides a strong evidence that the weight space is connected in the absence of thermodynamically-dominant barriers. This result of the effective model coincides roughly with the low-dimensional projection evidence shown in Fig. 2(b) as well.

Let us then look at the entropy landscape. According to Eq. (8), the entropy of the system at $x = 0$ gives a rough estimate of the size of the original weight space, which is one advantage of our statistical mechanics analysis (see a similar study of constraint satisfaction problems [42]). By varying the value of x , we found that the entropy landscape is also smooth. For a system with discrete degrees of freedom, one can easily compute the upper bound for the entropy $s_{ub}(d) = -d \ln d - (1-d) \ln(1-d)$ with the unique constraint of the distance between any two weight configurations in the high-dimensional space. One salient feature is that three subparts of the deep architecture display distinct behaviors, despite a common smooth landscape. Surprisingly, the interior region shows an entropy nearly saturating the upper bound. This region in physics behaves like a liquid state, where there exist many realizations of weights to fulfill the classification task of the deep network, and the learning dynamics is thus fast in this region. However, the right or left boundaries are more constrained than the interior region with lower entropy values. The left boundary is the most constrained one among all three

parts of the deep network, while the right boundary behaves similarly to the full network of three parts. An intuitive interpretation is that, two boundaries are directly responsible for encoding or decoding the input-output relationship embedded in the training dataset. Allowing more freedoms in these two boundaries will sacrifice the generalization capability of the network, which is also evidenced in our recent work of learning credit assignments [43]. In sum, the under-parametrization does not lead to a broken weight space, with a liquid-like interior next to two more constrained boundaries.

The aforementioned properties of the deep-learning weight space are qualitatively conserved, when the effective model is constructed by applying the pseudolikelihood maximization method, where the sparsity of the model is controlled (Fig. 6). When the cavity approximation breaks, the heat-capacity integration method can be used, yielding the smooth landscape as well (data not shown).

The entropy landscape analysis also allows to compute the iso-energy curve (technical details are given in Appendix C). Figure 7 shows the iso-energy entropy curve, where every point on the curve shares the same energy with the reference configuration. This curve is suppressed when $d \in [0.1, 0.9]$, compared to the original entropy landscape. In the regime close to the reference, the entropy is nearly the same as that without the energy constraint. In the regime around $x = 0$, there appears a negligible hysteresis loop, indicating that the roughness of the landscape is weak even if the iso-energy constraint is enforced. Note that there exist two local minima in the energy-distance curve, whose asymmetry is caused by the external bias of weights. This shows that by applying the current physics tool, one can access the structure of the specified landscape around any reference point in the high-dimensional deep-learning space.

B. Over-parametrization scenario

We then ask whether the nice property of the under-parametrization case transfers to the over-parametrization case, which is the popular setting in the current deep-learning era. We thus investigate the over-parametrization case, where the number of training examples is less than the model com-

FIG. 6. Entropy landscape analyses of the sparse Ising model (constructed by the pseudolikelihood maximization method, and $\lambda = 0.05$). Other model settings are the same as Fig. 4. (a) Under-parametrization case ($\alpha \approx 0.01$). The curves are the mean results over 20 distinct reference weight configurations σ^* selected from the low energy region (for references coming from high- and medium-energy regions, the profile holds qualitatively). Entropy curves are plotted along the direction of the decreasing coupling field x from 3 to -3 , while the Hamming distance per weight d vs x is plotted by considering two directions of increasing and decreasing x (inset). No hysteresis loops and entropy gaps are observed. (b) Over-parametrization case ($\alpha \approx 1.35$). Other conditions are the same as (a).

plexity (Appendix D). The standard over-parameterization requires that $\alpha \gg 1$, while our analysis is limited to a light version ($\alpha > 1$) of over-parametrization. We surprisingly find that the qualitative properties of the under-parametrization case also hold in the over-parametrization case [Fig. 5(b)], in an excellent agreement with recent empirical studies of deep learning at a large-scale architecture [6–8]. All these previous studies confirmed that the stochastic gradient descent dynamics does not suffer from high loss barriers [44], and the loss function landscape is characterized by many minima yet separated by low barriers [45]. Our finding adds further evidences by claiming a smooth landscape for deep networks of binary synapses, in both under- and over-parametrization scenarios, which may be connected to the absence of double descent phenomena (Appendix E).

However, we notice a salient difference that the extent the interior entropy gets close to the upper bound is weaker than that of the under-parametrization case, despite the fact that the interior entropy is still highest among three parts of the deep network. Even in the over-parametrization case, a liquid-like

interior inside the network exists, coinciding exactly with a recent study of deep networks in both random inputs/outputs setting and teacher-student generalization setting [46]. This recent paper claimed a glass-liquid-glass/solid-liquid-solid pattern inside the deep-learning architecture. Despite a relatively small scale for the sake of the data-driven analytical study, our paper claims the same phenomenon in the central part of the deep network by constructing an effective data-driven model for a practical deep learning. Although we do not observe a glass/solid like boundaries (unlike in the recent study [46]), these two boundaries are clearly more constrained than the central part. We thus expect that these boundaries may enter a glass/solid phase in the thermodynamic limit, which is a very interesting future direction.

C. Shallow networks

To explore the benefit of depth, an important characteristic of the deep learning, we study a shallow-network architecture without any hidden layers. Other conditions are the same with

FIG. 7. Iso-energy entropy profile for the under-parameterization case. (a) The curves are mean results estimated from ten independent trials where the reference configurations σ^* are selected from the low energy region of ϵ_{Ising} . Both the iso-energy case and the normal case (no fixed-energy constraint) share the same reference configurations. (b) The iso-energy constraint is achieved by fixing the energy to the Ising energy density ϵ_{Ising} of the reference configuration, for which the secant method is applied to search for a compatible inverse temperature β , while the coupling field x still varies from 3 to -3 . In the normal case, the inverse temperature β is set to 1. $d = 0.5$ (or $x = 0$) corresponds to the original unperturbed model.

FIG. 8. Distribution of TAP free energies of the Ising model. The model is constructed by the maximum entropy method. The histogram is obtained by iterating the TAP equation (see Appendix H) from 2 000 random initializations of the magnetizations. (a) Under-parametrization case. (b) Over-parametrization case. (c) Shallow network case.

the other two scenarios. Note that, in terms of test performance, adding hidden layers of increasing width improves over the shallow counterpart (Appendix F). Using the same statistical mechanics framework, we surprisingly find a first-order phase transition in the distance-coupling-field profile [Fig. 5(c)]. The discontinuous transition is characterized by a hysteresis loop and associated entropy gaps, irrespective of references selected from high-, medium- and low-energy regions. For the shallow-network case, the weight space is broken into clusters separated by entropy gaps. In particular, a double-discontinuous transition is even observed for the low-energy references, implying that the weight space around the low-energy references becomes much more complex, as also revealed in a biological data analysis of retina coding [32].

Compared to the deep network, the dynamics in the shallow network can be strongly affected by the discontinuous transition because of metastability in thermodynamic states. When adding more hidden layers, the broken weight space is replaced by a connected smooth subspace, demonstrating the benefit of depth. Our effective model thus distinguishes a shallow network from its deep counterpart, as also supported by the Monte Carlo dynamics in the next section.

D. Distributions of TAP free energy and Monte Carlo dynamics of the effective model

To verify the geometric picture gained from the entropy landscape, we study the TAP free energy distribution of the effective Ising model (see Appendix H for technical details). As shown in Fig. 8 for the under-parametrization setting, the distribution shows a single peak, consistent with the smooth landscape. The over-parameterization setting has similar results. The single peak profile is further corroborated by the Monte Carlo dynamics (like that in Appendix B) starting from a reference solution state. As shown in Fig. 9, the local dynamics is apt to drift towards chance-level configurations, without difficulty of crossing high-energy barriers. In contrast, the Monte Carlo dynamics of the Ising model for the shallow network gets trapped by the basin where the reference state sits on, which coincides exactly with the first-order transition (i.e., emergence of high energy barriers in the phase space). This observation is also supported by the double-peak profile of TAP free energy distributions (Fig. 8). The above prop-

erties are also observed in the sparse Ising model of deep learning.

To sample one high-accuracy weight-configuration, one can instead use Eq. (1), as the probability parameters are directly trained, while the effective model captures only up to pairwise correlations of the weight-space statistics, showing a smooth landscape for deep networks, and thus the local dynamics is easy to escape from any solution state. The effective model could predict a significant fraction of three-weight correlations, and the physics interpretation leads to consistent pictures about the weight space as well as the learning, and even inspire algorithmic designs (shown below). Therefore, the constructed Ising model could be an effective model. It would be very interesting to explore in future works whether higher-than-second order correlations in the data enhance our current understanding.

E. Effects of node-permutation symmetry on the deep-learning landscape

There exists node-permutation symmetry in deep learning, which is able to relate one weight configuration to another one via a permutation of hidden nodes (see Appendix I). To explore effects of node-permutation symmetry, we design a procedure of ranking the pre-activations of hidden layers, and then control the fraction of permutation symmetry in the collected weight configurations. As shown in Fig. 10, increasing the permutation symmetry fraction leads to an enlarged entropy profile, for four different network settings. Two key points must be remarked. First, adding the permutation symmetric configurations for learning the effective model does not affect the qualitative behavior of the entropy profile. Second, the close-by regime of any reference is not affected by the fraction of the permutation symmetric configurations. According to this observation, we conjecture that the underlying picture is that the deep-learning weight space under investigation is smooth as a whole.

F. Algorithmic implication of the effective model

Our theory predicts a liquid core inside the deep network. We ask a last question—does the effective model lead to an algorithmic design, despite the effective model built for

FIG. 9. Monte Carlo dynamics on the landscape of the Ising model. The dynamics starts at a reference weight configuration, and is measured in a time unit of one proposed spin flip (one step in the abscissa). The test accuracy and Ising energy are both tracked during the dynamics. The fluctuation is computed from 100 independent trials. (a) Deep network case. (b) Shallow network case.

small-scale networks? To address this question, we first randomly freeze the connections inside the ITR, i.e., each weight value is chosen with equal probabilities, and then its external field θ is quenched to a positive large value (if $w = +1$) or a negative large value (if $w = -1$). The other parts of the network are normally trained, as we did before. We surprisingly find that this scheme works in practice (Fig. 11), coinciding exactly with the liquid-like core hypothesis drawn from the effective model. An additional salient advantage is that the training would become fast, in that the ITR does not need learning. To conclude, although our effective model is data-driven by the weight interactions of small-scale deep networks, the derived insights could inspire algorithmic designs for large-scale networks.

V. CONCLUSIONS

Key properties of deep networks, including the learning dynamics and the generalization capability, can be affected by the geometry of the weight-space landscape. Here, we design a data-driven effective model, considering all three necessary practical issues of a deep learning—structured data, deep architectures, and training, to explore the geometric structure of the entire weight space of a deep-learning landscape. Our paper provides an interpretable effective model of deep-learning weight space in small-size neural networks, although a theory for finite-size deep networks is rare in previous studies. Remarkably, the theory implies algorithmic designs for large-scale deep networks with fast training.

FIG. 10. Effects of permutation symmetry on the entropy landscape ($P = 60\,000$, $N = 675$). The Ising model is constructed by the maximum entropy method. The procedure of generating permutation symmetric weight configurations is detailed in Appendix I. The fraction of permutation symmetry (r_{ps}) is also specified there, and indicates how large the number of the permutation-dependent weight configurations (used for learning the Ising model) is. The entropy is the mean value over 20 independent references. (a) Full network. (b) Interior part of the network. (c) Left boundary of the network. (d) Right boundary of the network.

FIG. 11. Algorithmic implication of the effective model. Training with randomly-frozen ITR works, consistent with our theory. The training is carried out on the network architecture 784-100-100-10 with the full dataset. The performance of training with trained ITR is compared.

Our effective model reveals several interesting phenomena. First, the model can distinguish a deep network from a shallow one in terms of weight-space geometric structure. The deep network weight-space is smooth and dominated by a connected component, while a broken weight space emerges in the shallow network, displaying a discontinuous transition in the Hamming distance profile. Second, for the deep network, the under-parametrization and over-parametrization belong to the same class, in the sense that their weight spaces are identically dominated by a single connected component, which is consistent with recent empirical studies that claimed no substantial barriers between minima in the deep-learning loss landscape [6–8]. Our statistical mechanics analysis thus provides insights towards understanding the success of deep learning despite the NP-hardness (worst case) of the deep learning problem.

A recent paper revealed that as the model complexity of continuous weights increases, the generalization error displays a cusp around a critical value of the complexity; on both sides of this cusp, the generalization error decays [47]. It is thus interesting in future works to see whether there exists a qualitative change in the geometric structure of the weight space around the cusp. We further remark that in our current study, the model complexity is fixed for simplicity. In fact, the model complexity can also be tuned to study the existence of the double descent phenomenon (see Appendix E). We confirm that in the explored range of model complexity, the double descent phenomenon is absent for deep learning of binary synapses. In other words, the generalization capability of the network continuously improves with increasing model complexity.

Third, a special interior part of a largest entropy in the weight space exists, showing a liquid-like property. The liquid phase is defined with a vanishing glass-order-parameter [46], or the absence of ergodicity breaking. In this sense, the learning dynamics in the interior part must be fast, because both directions of each weight are not highly biased, or close to a random initialization. In the particular scenario of under-parametrization, this giant interior entropy even gets very close to the upper bound of entropy set by a completely random-weight space.

The existence of the liquid-like central part coincides exactly with two recent studies, despite the current focus of binary weights. One is the toy random models of deep learning [46], as discussed in Sec. IV B; the other is the ensemble backpropagation applied to practical deep learning [43], which found a peak of model entropy in the central part of the deep network. We remark that these two works studied continuous degrees of freedom and are restricted to only the over-parametrization scenario, while our current effective model focuses on discrete degrees of freedom, reducing a hard-to-describe practical deep network to a theoretically-amenable Ising model, which establishes a bridge between abstract physics model and practical deep learning. Based on our model, the global geometric structure of any specific point in the high-dimensional weight space can be explored, while this is hard for those toy models or numerical experiments.

Our work also provides several significant insights about our current understanding of deep networks with discrete synapses. First, the role of depth can be interpreted in characteristics of the high-dimensional weight space. Second, the under-parameterization and the over-parameterization make no difference in the qualitative property of the weight space. Third, deep learning shapes a liquid-like core in the weight space for the central part of the architecture, which suggests that there exists a fast mode for the learning dynamics along the hierarchy, thereby inspiring further algorithmic development. Although it is still challenging to derive an effective model of large-scale deep learning, we anticipate that the insights gained from this study, including smooth landscapes, liquid-like core, TAP free energies, effects of permutation symmetry, and absence of double descent phenomenon, would inspire derivation of general principles governing the relationship among learning dynamics, geometry of weight space, and the behavior of generalization. In particular, it remains open what the landscape looks like around the test error cusp for modern deep learning of real-valued weights, which our framework may help to clarify, and we will explore this interesting direction in future works.

ACKNOWLEDGMENTS

We would like to thank other members of PMI laboratory for a long-term discussion of this project. This research was supported by the National Key R&D Program of China (2019YFA0706302), the National Natural Science Foundation of China for Grant No. 11805284, and the start-up budget 74130-18831109 of the 100-talent- program of Sun Yat-sen University.

APPENDIX A: MEAN-FIELD TRAINING ALGORITHMS

To train a deep supervised network with binary weights, we apply the mean-field method first introduced in the previous work [29]. In our current setting, each weight w_{ij}^l is sampled from a binomial distribution $P(w_{ij}^l)$ parametrized by an external field θ_{ij}^l as follows:

$$P(w_{ij}^l) = \sigma(\theta_{ij}^l) \delta_{w_{ij}^l, 1} + [1 - \sigma(\theta_{ij}^l)] \delta_{w_{ij}^l, -1}, \quad (\text{A1})$$

FIG. 12. Three-spin correlations of the effective Ising model, in comparison with those estimated from the collected weight ensemble. (a) Predicted three-spin correlations vs the measured ones (data). 200 spins are randomly selected from the effective model to do the comparison, due to a huge number of three-spin correlations. The full line shows equality, and dashed lines correspond to the deviation criterion ϵ , which is the absolute value of the deviation between the predicted three-spin correlation and the true one. The predicted three-spin correlation is estimated by running a Monte Carlo simulation of the effective Ising model. (b) Predicted percentage vs deviation criterion. The predicted percentage in all considered three-spin correlations is the percentage of those three-spin correlations whose absolute deviation from the true one is within the given deviation criterion ϵ .

with mean $\mu_{ij}^l = 2\sigma(\theta_{ij}^l) - 1$ and variance $(\sigma_{ij}^l)^2 = -4\sigma^2(\theta_{ij}^l) + 4\sigma(\theta_{ij}^l)$. According to the central-limit theorem, the feedforward transformation can be reparametrized as

$$z_j^l = m_j^l + v_j^l \cdot \epsilon_j^l, \quad (\text{A2a})$$

$$a_j^l = \frac{1}{\sqrt{N_{l-1}}} \text{L-ReLU}(z_j^l), \quad (\text{A2b})$$

where $m_j^l = \sum_i \mu_{ij}^l a_i^{l-1}$, and $v_j^l = \sqrt{\sum_i (\sigma_{ij}^l)^2 (a_i^{l-1})^2}$.

During the error backpropagation phase, we need to compute the gradient of the loss function $\mathcal{L}$ [e.g., L_{CE} in Fig. 2(a)] with respect to the external field θ , which proceeds as follows:

$$\frac{\partial \mathcal{L}}{\partial \theta_{ij}^l} = \frac{\partial \mathcal{L}}{\partial z_j^l} \frac{\partial z_j^l}{\partial \theta_{ij}^l} = \frac{\partial \mathcal{L}}{\partial z_j^l} \left(\frac{\partial m_j^l}{\partial \theta_{ij}^l} + \epsilon_j^l \frac{\partial v_j^l}{\partial \theta_{ij}^l} \right). \quad (\text{A3})$$

We then define $\Delta_j^l \equiv \frac{\partial \mathcal{L}}{\partial z_j^l}$. On the top layer, $\Delta_j^l = y_j^l - \hat{y}_j^l$, where $y_j^l = \frac{e^{z_j^l}}{\sum_i e^{z_i^l}}$ is the softmax output, and $\hat{y}_j^l$ is the (one-hot) label of the input. On the lower layers, given Δ_k^{l+1} , we can iteratively compute Δ_j^l :

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial z_j^l} &= \sum_k \frac{\partial \mathcal{L}}{\partial z_k^{l+1}} \frac{\partial z_k^{l+1}}{\partial z_j^l} \\ &= \sum_k \Delta_k^{l+1} f'(z_j^l) \left(\mu_{jk}^{l+1} + \epsilon_k^{l+1} \frac{(\sigma_{jk}^{l+1})^2 a_j^l}{v_k^{l+1}} \right). \end{aligned} \quad (\text{A4})$$

Finally, we compute $\frac{\partial m_j^l}{\partial \theta_{ij}^l}$ and $\frac{\partial v_j^l}{\partial \theta_{ij}^l}$, respectively. It then proceeds as follows:

$$\frac{\partial m_j^l}{\partial \theta_{ij}^l} = \frac{\partial m_j^l}{\partial \mu_{ij}^l} \frac{\partial \mu_{ij}^l}{\partial \theta_{ij}^l} = 2a_i^{l-1} \sigma'(\theta_{ij}^l), \quad (\text{A5a})$$

$$\frac{\partial v_j^l}{\partial \theta_{ij}^l} = \frac{\partial v_j^l}{\partial (\sigma_{ij}^l)^2} \frac{\partial (\sigma_{ij}^l)^2}{\partial \theta_{ij}^l} = -2 \frac{(a_i^{l-1})^2 \mu_{ij}^l \sigma'(\theta_{ij}^l)}{v_j^l}. \quad (\text{A5b})$$

Note that ϵ is sampled and quenched for both forward and backward computations in a single mini-batch gradient descent. During training, we add an L_2 regularization to controlling the magnitude of external fields. After the learning is terminated, an effective network with binary weights can be constructed by sampling the binomial distribution parametrized by external fields.

APPENDIX B: THE ISING MODEL PREDICTS THREE-WEIGHT CORRELATIONS

To see whether the data-driven Ising model can reproduce the three-weight correlations, we first run a Monte Carlo experiment on the Ising Hamiltonian, and then carry out a sampling of the steady state space of the Monte Carlo dynamics after a sufficient relaxation. The three-weight correlations are estimated from these samples, and are then compared with those estimated directly from the data (weight ensemble). The result is shown in Fig. 12, which suggests that the Ising model can predict a significant fraction of three-weight correlations (e.g., about 82% provided that the criterion is set to $\epsilon = 0.015$).

APPENDIX C: CAVITY METHOD FOR LANDSCAPE ANALYSIS

To analyze the statistical properties of the least structured model, we apply the cavity method in the spin glass theory [48]. Note that the weight configuration follows the Boltzmann distribution $P(\sigma) = \exp(-\beta E(\sigma))/Z$, where the energy $E(\sigma) = -\sum_{i<j} J_{ij} \sigma_i \sigma_j - \sum_i h_i \sigma_i$, and Z is the partition function. An exact computation of the free energy $(-T \ln Z)$ is an NP-hard problem. To get an approximate free energy, we first compute the free energy change ΔF_i when adding a weight (e.g., σ_i) and its associated interactions. We then calculate the free energy change ΔF_a when adding an interaction [e.g., $J_a \sigma_i \sigma_j$, and $a \equiv (ij)$]. In all these derivations, we make an assumption that when an interaction is

removed, its neighboring weights become independent such that the joint distribution of them becomes factorized, which is exact if the factor graph of the model is locally tree like. The factor graph is a bipartite graph where two kinds of nodes (weight nodes and interaction nodes) are present. For detailed derivations, interested readers are recommended to go through the standard textbook [49], or the Appendix of a recent paper [50]. In the cavity approximation, the free energy can be constructed as below:

$$F \equiv -\frac{1}{\beta} \ln Z = \sum_i \Delta F_i - \sum_a (|\partial a| - 1) \Delta F_a, \quad (C1a)$$

$$\beta \Delta F_i = -\ln \sum_{x=\pm 1} \mathcal{H}_i(x), \quad (C1b)$$

$$\mathcal{H}_i(x) \equiv e^{x\beta h_i} \prod_{b \in \partial i} \cosh \beta J_b (1 + x \hat{m}_{b \rightarrow i}), \quad (C1c)$$

$$\beta \Delta F_a = -\ln \cosh \beta J_a - \ln \left(1 + \tanh \beta J_a \prod_{i \in \partial a} m_{i \rightarrow a} \right), \quad (C1d)$$

where $|\partial a|$ denotes the number of weights connected to the interaction a , $i \in \partial a$ denotes the neighboring weights i of the interaction a , and $a \in \partial i$ denotes the neighboring interactions a of the weight i . $m_{j \rightarrow b}$ is the cavity magnetization of weight j in the absence of the interaction a , and $\hat{m}_{b \rightarrow i} = \tanh \beta J_b \prod_{j \in \partial b \setminus i} m_{j \rightarrow b}$ is the conjugate magnetization.

According to the variational principle, the free energy should be stationary with respect to $\{m_{i \rightarrow a}\}$. We can then obtain a set of self-consistent equations called message passing equations as follows:

$$m_{i \rightarrow a} = \tanh \left(\beta h_i + \sum_{b \in \partial i \setminus a} \tanh^{-1} \hat{m}_{b \rightarrow i} \right), \quad (C2a)$$

$$\hat{m}_{b \rightarrow i} = \tanh \beta J_b \prod_{j \in \partial b \setminus i} m_{j \rightarrow b}, \quad (C2b)$$

where $\setminus i$ means that the weight i should be excluded from the set. $m_{i \rightarrow a}$ is interpreted as the message passing from weight i to interaction a , and $\hat{m}_{b \rightarrow i}$ is interpreted as the message passing from interaction b to weight i . By iterating these equations from some random initialization, the iteration will converge to a fixed point $\{m_{i \rightarrow a}^*\}$, which can be used to compute thermodynamic moments for learning (introduced below) and entropy (used for landscape analysis). We remark that the fixed point corresponds to a local minimum of the free energy function. The algorithm may not converge when the model becomes glassy.

We can directly compute the magnetization $m_i = \langle \sigma_i \rangle$, and correlation $C_a = \langle \sigma_i \sigma_j \rangle$ from the fixed point $\{m_{i \rightarrow a}^*\}$ as follows:

$$m_i = \tanh \left(\beta h_i + \sum_{b \in \partial i} \tanh^{-1} \left[\tanh \beta J_b \prod_{j \in \partial b \setminus i} m_{j \rightarrow b}^* \right] \right), \quad (C3a)$$

$$C_a = \frac{\tanh \beta J_a + \prod_{i \in \partial a} m_{i \rightarrow a}^*}{1 + \tanh \beta J_a \prod_{i \in \partial a} m_{i \rightarrow a}^*}. \quad (C3b)$$

These moments of the effective model can be used for updating the parameters of the model, approximating the model expectation terms in the Boltzmann learning rule [Eq. (5)], which are faster in our current setting than other methods, e.g., Monte Carlo samplings.

Given the free energy, we can estimate the energy from the standard thermodynamics relation, $E = \frac{\partial(\beta F)}{\partial \beta}$, which is given by

$$E = - \sum_i \Delta E_i + \sum_a (|\partial a| - 1) \Delta E_a, \quad (C4a)$$

$$\Delta E_i = \frac{\beta h_i \sum_{x=\pm 1} x \mathcal{H}_i(x) + \sum_{x=\pm 1} \mathcal{G}_i(x)}{\sum_{x=\pm 1} \mathcal{H}_i(x)}, \quad (C4b)$$

$$\Delta E_a = \beta J_a \frac{\tanh \beta J_a + \prod_{i \in \partial a} m_{i \rightarrow a}}{1 + \tanh \beta J_a \prod_{i \in \partial a} m_{i \rightarrow a}}, \quad (C4c)$$

where

$$\begin{aligned} \mathcal{G}_i(x) = & \sum_{b \in \partial i} e^{x\beta h_i} \left[\beta J_b \sinh \beta J_b (1 + x \hat{m}_{b \rightarrow i}) \right. \\ & \left. + x \beta J_b \cosh \beta J_b (1 - \tanh^2 \beta J_b) \prod_{j \in \partial b \setminus i} m_{j \rightarrow b} \right] \\ & \times \prod_{a \in \partial i \setminus b} \cosh \beta J_a (1 + x \hat{m}_{a \rightarrow i}). \end{aligned} \quad (C5)$$

The entropy of the model is defined as $S = -\sum_{\sigma} P(\sigma) \ln P(\sigma)$, which can be also estimated from the standard thermodynamics relation $S = -\beta F + \beta E$.

In the entropy landscape analysis, we introduce a perturbed probability distribution to explore the internal structure of the network parameter space, which is given by

$$P(\sigma) = \frac{1}{Z} \exp \left(\beta \sum_{i < j} J_{ij} \sigma_i \sigma_j + \beta \sum_i h_i \sigma_i + x \sum_i \sigma_i^* \sigma_i \right), \quad (C6)$$

where β is the inverse temperature, and x is the coupling field tuning the distance between the configuration σ and the reference σ^* . With the Laplace method, the free energy density f can be approximated as $f \approx \min_{\epsilon, q} f(\epsilon, q)$, which is given by $\beta f = \beta \epsilon - xq - s(\epsilon, d)$, where ϵ is the energy density (E/N), q is the overlap $\sum_i \sigma_i \sigma_i^* / N$, $s(\epsilon, d)$ is the entropy density S/N , and d is the Hamming distance per weight related to the overlap q by $d = (1 - q)/2$. At the same time, (x, q) should obey the following equations: $\partial s(\epsilon, d) / \partial d = 2x$ and $\partial s(\epsilon, d) / \partial \epsilon = \beta$, because of the Laplace method. According to the double Legendre transform, the entropy density $s(\epsilon, d)$ is given by $s(\epsilon, q) = -\beta f + \beta \epsilon - xq$. Note that, by fixing d or ϵ , one can construct iso-distance or iso-energy curves just through finding compatible x or β for the following equations:

$$\frac{\partial(\beta f)}{\partial \beta} = \epsilon, \quad (C7a)$$

$$\frac{\partial(\beta f)}{\partial x} = -q, \quad (C7b)$$

due to the Legendre transform.

FIG. 13. Entropy landscape analyses in the high-energy reference case. The reference weight configurations are randomly selected from the high-energy region, which displays no significant difference from the low-energy reference case (Fig. 5). (a) Under-parametrization case. (b) Over-parametrization case.

Under the perturbed probability measure, the cavity method is still applicable, with an only change for the bias h_i as $\beta h_i \rightarrow \beta h_i + x\sigma_i^*$. By iterating the same message passing equations to a fixed point, we can finally obtain the free energy density, energy density, and entropy density, etc. Note that q is computed via $\sum_i m_i \sigma_i^*/N$, where m_i is the fixed-point magnetization.

APPENDIX D: SIMULATION DETAILS

The MNIST handwritten digits dataset is a classification benchmark dataset, which contains 60 000 training images and 10 000 test images. The image is a 28×28 gray-scale handwriting digit with a label taken from one of 0 to 9. For the sake of simplicity, the dataset we use for training is re-generated by applying the PCA to the original dataset. More precisely, we apply the PCA on training images to compute the eigenvectors first, and then project training and test images with these eigenvectors to finally obtain the low dimensional (20 pixels) images.

To both have a high test accuracy and avoid a large network-parameter size at the same time, we choose the network architecture as 20-15-15-10, where each number denotes the corresponding layer-width. We add a softmax layer onto the output layer and use the cross entropy as the loss function. One training epoch involves the whole training dataset, which is divided into 300 mini-batches, each of which contains 200 images. The loss is optimized by Adam [51] and the initial

learning rate is set to 0.3. During the training, we also introduce an L_2 regularization to the loss function to penalize large external fields. The regularization strength is set to 10^{-5} , which is optimal for our experiments.

The stimulation details presented above apply to the case of under-parametrization. We also design two other kinds of training scenarios. The first case is the over-parametrization scenario, where the number of training images is smaller than the number of network parameters. In this case, we only train 500 images, which are divided into 25 mini-batches with the mini-batch size of 20. Other hyper-parameters remain unchanged. The second case is the shallow-network scenario, where we delete all hidden layers, and thus the network architecture changes to 20-10. Other hyper-parameters remain unchanged as well. Following the same weight-collection procedure as the under-parametrization case except for setting the accuracy threshold to 70%, we collect one million and half a million weight configurations for the over-parameterization and shallow-network scenarios, respectively, in order to construct effective models.

Finally, we show additional landscape analysis when the reference weight configuration is selected from the high and medium energy regions. A qualitative same behavior is identified (Figs. 13 and 14). In addition, the histograms of model parameters inferred from the weight data are also shown in Fig. 15 for the shallow-network setting. Our framework can also be applied to other more challenging types of dataset,

FIG. 14. Entropy landscape analyses in the medium-energy reference case. The reference weight configurations are randomly selected from the medium-energy region, which displays no significant difference from the low-energy reference case (Fig. 5). (a) Under-parametrization case. (b) Over-parametrization case.

FIG. 15. Histograms of model parameters in the Ising model of the shallow-network case. (a) Histogram of inferred couplings. The couplings are distributed relatively broader compared with non-shallow networks. (b) Histogram of the inferred biases. The biases are also distributed broadly.

e.g., Fashion-MNIST consisting of fashion items images belonging to ten classes [52], as displayed in Fig. 16 for the over-parameterization case. Unlike the case of the MNIST dataset, the entropy landscape of the full network gets close to that of the central part of the network, while other qualitative properties remain. All codes to obtain the results of this paper are available on GitHub [53]

APPENDIX E: DOUBLE DESCENT PHENOMENON

In the traditional statistical learning theory, a classical U-shaped test error curve is predicted when the training is performed with increasing model complexity. However, deep neural networks are able to achieve the state-of-the-art performance using over-parametrized architectures. More precisely,

FIG. 16. Entropy landscape analysis on the fashion-MNIST dataset. The network structure remains 20-15-15-10 and the training data size is 500 corresponding to the over-parameterization case. One million weight configurations after learning are collected to calculate the weight statistics, which should reach the test accuracy threshold equal to 70%. The root-mean squared errors Δ are less than 0.085 to obtain an effective Ising model. To get the entropy curves, 20 weight configurations are selected from the low energy region as the reference configurations σ^* . Other conditions are the same as in the MNIST case.

when the number of network parameters are sufficiently large to perfectly learn the training data (i.e., at the so-called interpolation threshold), further increasing the model complexity will cause a second drop of the test error. Therefore, the learning behavior in the over-parametrization regime is called the double descent phenomenon [40,47,54,55]. Note that for ridge regression problems (implemented via a two-layer neural network for an asymptotic study, where only the second layer weights are trained while the first layer weights are random for the sake of theoretical analysis [40,54–56]), the test error may diverge as the interpolation boundary is approached [54,55], because of the divergence of network weights [57]. However, these studies are restricted to deep neural networks with real-valued weights and standard stochastic gradient descent training. In contrast, our study focuses on bounded synapses (e.g., ± 1), showing intrinsic differences. In this section, we provide details to check if the double descent phenomenon remains in our current setting, i.e., training deep networks of binary weights.

We consider the architecture of K - h - h - x , where K means the dimension of the input image (the image can be pre-processed by PCA), h means the size of hidden layers (i.e., $N_1 = N_2 = h$ in the main text), and $x = 2$ or 10 means binary classification (parity of the digits, see also Ref. [47]) or ten-class classification respectively. For the parity classification task, we use the hinge loss, while we use the cross-entropy loss for the ten-class classification task. The value of h is used to tune the model complexity. The value of K takes from 10 to 50 (preservation of 82% original information). In all these settings, we found double descent phenomenon for networks of real-valued weights [see Figs. 17(a) and 17(b)]. In contrast, as shown in Figs. 17(c) and 17(d), for networks of binary weights, the test error keeps decreasing with the model complexity, while the training error never reaches zero (or decreases extremely slowly). Compared to the real-valued-weight network, it is very difficult to train the network to perfectly fit the target function using the low-precision binary weights, which is evident in our simulation results [Figs. 17(c) and 17(d)]. We thus conclude that no interpolation threshold exists in our current setting, and correspondingly the peak of test error is not observed. This conclusion is robust for binary classification and other values of input dimensionality.

FIG. 17. Detection of double descent phenomenon in deep learning of continuous weights or binary weights. The network structure is specified by 20-h-h-10, and the training data size is 2 000. The ratio between the network complexity and the training data size is defined by α . We choose the cross entropy as the cost function. [(a),(b)] Deep learning of continuous weights. The point where the training error vanishes defines the interpolation point, indicating a peak in test error as well. [(c),(d)] Deep learning of binary weights. The double descent phenomenon is absent. We actually train the network of binary weights for a total number of 10^5 epochs, therefore the abscissa should be multiplied by a factor of 100.

APPENDIX F: PERFORMANCE COMPARISON BETWEEN DEEP AND SHALLOW NETWORKS

In this section, we show that the depth in our current setting can improve the performance. We consider the architecture 20-h-h-10, where h can be tuned. For the shallow network, two hidden layers are removed. We train the network using the entire dataset of MNIST, and observe that when the size of the hidden layer crosses a threshold (e.g., 10), the test performance keeps being improved over the shallow architecture. The results are summarized in Fig. 18.

APPENDIX G: HEAT-CAPACITY INTEGRATION FOR ENTROPY ESTIMATION

When the Bethe approximation breaks down, one can use heat-capacity integration to estimate the model entropy. According to the thermodynamics relationship, we write the entropy as follows,

$$S(T = 1, N) = \int_0^1 \frac{C(T, N)}{T} dT, \quad (\text{G1})$$

where the heat capacity $C(T, N)$ can be computed through the Monte Carlo sampling procedure [58]. More precisely, the

variance of the energy, $\langle E^2 \rangle - \langle E \rangle^2$, equals to $T^2 C(T, N)$ (we have assumed unit Boltzmann constant). The energy variance can be obtained by drawing a large number of samples (e.g., 2×10^6) at every intermediate temperature ($T \in [0, 1]$). In fact, taking into account the coupling field term, the entropy of landscape is computed as $S - Nxq$.

FIG. 18. Performance comparison between deep and shallow networks. Ten independent trainings are considered for the fluctuation.

FIG. 19. Schematic illustration of permutation symmetry transformation. The transformation is carried out by a permutation of hidden neurons according to descending levels of pre-activation (denoted by $\mathbf{a}^l$), i.e., $\mathbf{a}^l_{\text{new}} = \mathbf{S}_h \mathbf{a}^l_{\text{old}}$, which further leads to a new weight configuration σ_{gr} , related to the original one σ via $\sigma_{\text{gr}} = \mathcal{T}\sigma$, where the matrix entry takes either zero or one. The thick edges represent the involved weights for the transformation, i.e., the incoming and outgoing synapses for the permuted neurons.

APPENDIX H: TAP FREE ENERGY DISTRIBUTION

The well-known Thouless-Anderson-Palmer (TAP) equation for a disordered Ising model reads as follows:

$$m_i^{t+1} = \tanh \left(\beta h_i + \beta \sum_{j(\neq i)} J_{ij} m_j^t - \beta^2 \sum_{j(\neq i)} J_{ij}^2 (1 - (m_j^t)^2) m_i^{t-1} \right), \quad (\text{H1})$$

where we have put the correct time indexes for iteration [59]. The fixed points of TAP are the stationary points of the TAP free energy [60], as given by

$$-\beta F_{\text{TAP}} = - \sum_i \left(\frac{1+m_i}{2} \ln \frac{1+m_i}{2} + \frac{1-m_i}{2} \ln \frac{1-m_i}{2} \right) + \frac{1}{2} \beta \sum_{i \neq j} J_{ij} m_i m_j + \frac{\beta^2}{4} \sum_{i \neq j} J_{ij}^2 (1 - m_i^2) (1 - m_j^2). \quad (\text{H2})$$

At low temperatures, The TAP equation has many solutions with $m_i \neq 0$, which can be interpreted as stable or metastable thermodynamic states. Note that in our current setting, the temperature effect has been absorbed into the parameters (Ω) of the effective Ising model. To estimate the free energy distribution, we start the iteration of Eq. (H1) from 2000 random initializations. The results show that for deep networks, a single peak profile is observed, in contrast to the double peak profile for the shallow network (see Fig. 8 in the main text). This observation is consistent with our entropy landscape analysis.

APPENDIX I: NODE-PERMUTATION SYMMETRY

It is clear that there exists node-permutation symmetry for hidden layers, e.g., exchange of an arbitrary pair of nodes

(including their associated incoming and outgoing weights) does not cause the change of the network output. This raises an interesting question—how much this permutation affects the weight configuration data and further the landscape of the effective model. To address this question, we first select all digits of label 1 in the test set as the input, and then each sample of weight configuration produces the hidden preactivations, whose mean values over the input ensemble are subsequently ranked in the descending order for each hidden layer (called $\mathbb{H}_D^\ell$). The ranked preactivation result thus corresponds to an element of the node-permutation group, which transforms the original weight configuration σ to a permuted one σ_{gr} by a permutation symmetry transformation $\mathcal{T}$: $\sigma_{\text{gr}} = \mathcal{T}\sigma$ (see also Fig. 19). Note that each element in the transformation matrix $\mathcal{T}$ takes either zero or one. We thus call σ_{gr} as a generating weight-configuration, from which permutation-dependent configurations can be generated by carrying out one of the permutation transformations. The total number of these transformations is given by $h! \times h!$, where h is the width of hidden layer ($N_1 = N_2 = h$). We remark that the result does not depend on the selected input ensemble, e.g., other digits.

According to this rule, if two weight configurations lead to the same σ_{gr} (identical $\{\mathbb{H}_D^\ell\}$ as well), we say they are permutation-dependent; otherwise, they are permutation-independent. In fact, our original training data of weight configurations are completely permutation-independent, probably due to the limited-size weight-data yet the extremely large size of the deep-learning weight space. To introduce the permutation symmetry among weight configurations, we use M_s of σ_{gr} (by a random selection) to generate as many as $2M_s$ permutation-dependent configurations (one σ_{gr} is paired with the corresponding generated configuration). These generated configurations are combined with the other $M - M_s$ ones (M is the number of the collected weight configurations). Therefore, the permutation symmetry ratio is calculated as $r_{\text{ps}} = \frac{2M_s}{M - M_s + 2M_s}$. The effects of r_{ps} on the landscape are discussed in the main text.

- [1] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- [2] G. Carleo, J. I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborova, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [3] T. J. Sejnowski, The unreasonable effectiveness of deep learning in artificial intelligence, *Proce. Natl. Acad. Sci. U.S.A.* **117**, 30033 (2020).
- [4] M. Geiger, S. Spigler, S. d'Ascoli, L. Sagun, M. Baity-Jesi, G. Biroli, and M. Wyart, The jamming transition as a paradigm to understand the loss landscape of deep neural networks, *Phys. Rev. E* **100**, 012115 (2019).
- [5] Q. Nguyen and M. Hein, The loss surface of deep and wide neural networks, in *Proceedings of the 34th International Conference on Machine Learning*, edited by D. Precup and Y. Whye Teh, Proceedings of Machine Learning Research Vol. 70 (PMLR, 2017), pp. 2603–2612.
- [6] F. Dräxler, K. Veschgini, M. Salmhofer, and F. A. Hamprecht, Essentially no barriers in neural network energy landscape, in *Proceedings of the 35th International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 80, edited by J. Dy and A. Krause (PMLR, Sweden, 2018), pp. 1308.
- [7] T. Garipov, P. Izmailov, D. Podoprikin, D. Vetrov, and A. Wilson, Loss surfaces, mode connectivity, and fast ensembling of DNNs, in *NIPS'18: Proceedings of the 32nd Annual Conference on Neural Information Processing Systems* (Montreal, Canada, 2018), p. 8803.
- [8] Q. Nguyen, On connected sublevel sets in deep learning, in *Proceedings of the 36th International Conference on Machine Learning* (Long Beach, California, PMLR 97, 2019), p. 4790.
- [9] S. Fort and S. Jastrzebski, Large scale structure of neural network loss landscapes, in *33rd Conference on Neural Information Processing Systems* (NeurIPS, Curran Associates, Inc., 2019), p. 6706.
- [10] A. J. Ballard, R. Das, S. Martiniani, D. Mehta, L. Sagun, J. D. Stevenson, and D. J. Wales, Perspective: Energy landscapes for machine learning, *Phys. Chem. Chem. Phys.* **19**, 12585 (2017).
- [11] L. Sagun, U. Evci, V. U. Güney, Y. Dauphin, and L. Bottou, Empirical analysis of the Hessian of over-parametrized neural networks, [arXiv:1706.04454](https://arxiv.org/abs/1706.04454).
- [12] H. Li, Z. Xu, G. Taylor, and T. Goldstein, Visualizing the loss landscape of neural nets, in *32nd Conference on Neural Information Processing Systems* (NeurIPS, 2018), (Montreal, Canada).
- [13] S. Fort and S. Ganguli, Emergent properties of the local geometry of neural loss landscapes, [arXiv:1910.05929](https://arxiv.org/abs/1910.05929).
- [14] S. Hochreiter and J. Schmidhuber, Flat minima, *Neural Comput.* **9**, 1 (1997).
- [15] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang, On large-batch training for deep learning: Generalization gap and sharp minima, in *5th International Conference on Learning Representations, ICLR 2017* (2017).
- [16] L. Dinh, R. Pascanu, S. Bengio, and Y. Bengio, Sharp minima can generalize for deep nets, in *Proceedings of the 34th International Conference on Machine Learning* (Sydney, Australia, PMLR 70, 2017).
- [17] C. Baldassi, F. Pittorino, and R. Zecchina, Shaping the learning landscape in neural networks around wide flat minima, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 161 (2019).
- [18] Y. Feng and Y. Tu, The inverse variance–flatness relation in stochastic gradient descent is critical for finding flat minima, *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2015617118 (2021).
- [19] P. Chaudhari, A. Choromanska, S. Soatto, Y. LeCun, C. Baldassi, C. Borgs, J. T. Chayes, L. Sagun, and R. Zecchina, Entropy-SGD: Biasing gradient descent into wide valleys, *J. Stat. Mech.* (2019) 124018.
- [20] Y. V. Fyodorov and I. Williams, Replica symmetry breaking condition exposed by random matrix calculation of landscape complexity, *J. Stat. Phys.* **129**, 1081 (2007).
- [21] A. Choromanska, M. Henaff, M. Mathieu, G. B. Arous, and Y. LeCun, The loss surfaces of multilayer networks, *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics* (PMLR, San Diego, 2015).
- [22] H. Huang, and A. Goudarzi, Random active path model of deep neural networks with diluted binary synapses, *Phys. Rev. E* **98**, 042311 (2018).
- [23] S. Becker, Y. Zhang, and Alpha A. Lee, Geometry of Energy Landscapes and the Optimizability of Deep Neural Networks, *Phys. Rev. Lett.* **124**, 108301 (2020).
- [24] H. Huang, K. Y. Michael Wong, and Y. Kabashima, Entropy landscape of solutions in the binary perceptron problem, *J. Phys. A* **46**, 375002 (2013).
- [25] H. Huang, and Y. Kabashima, Origin of the computational hardness for learning with binary synapses, *Phys. Rev. E* **90**, 052813 (2014).
- [26] C. Baldassi, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina, Subdominant Dense Clusters Allow for Simple Learning and High Computational Performance in Neural Networks with Discrete Synapses, *Phys. Rev. Lett.* **115**, 128101 (2015).
- [27] S. K. Esser, P. A. Merolla, J. V. Arthur, A. S. Cassidy, R. Appuswamy, A. Andreopoulos, D. J. Berg, J. L. McKinstry, T. Melano, R. Davis, C. di Nolfo, P. Datta, A. Amir, B. Taba, M. D. Flickner, and D. S. Modha, Convolutional networks for fast, energy-efficient neuromorphic computing, *Proc. Natl. Acad. Sci. U.S.A.* **113**, 11441 (2016).
- [28] K. Roy, A. Jaiswal, and P. Panda, Towards spike-based machine intelligence with neuromorphic computing, *Nature (London)* **575**, 607 (2019).
- [29] O. Shayer, D. Levi, and E. Fetaya, Learning discrete weights using the local reparameterization trick, in *ICLR 2018 : International Conference on Learning Representations 2018* (2018).
- [30] E. T. Jaynes, Information theory and statistical mechanics, *Phys. Rev.* **106**, 620 (1957).
- [31] E. Schneidman, M. J. Berry, R. Segev, and W. Bialek, Weak pairwise correlations imply strongly correlated network states in a neural population, *Nature (London)* **440**, 1007 (2006).
- [32] H. Huang, and T. Toyozumi, Clustering of neural code words revealed by a first-order phase transition, *Phys. Rev. E* **93**, 062416 (2016).
- [33] Y. LeCun, The MNIST database of handwritten digits, retrieved from <http://yann.lecun.com/exdb/mnist>.
- [34] D. P. Kingma, T. Salimans, and M. Welling, Variational dropout and the local reparameterization trick, in *Advances in Neural Information Processing Systems*, Vol. 28, p 2575 (Curran Associates, Inc., 2015).

- [35] L. van der Maaten and G. Hinton, Visualizing data using t-sne, *J. Mach. Learn. Res.* **9**, 2579 (2008).
- [36] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, A learning algorithm for Boltzmann machines, *Cogn. Sci.* **9**, 147 (1985).
- [37] P. Ravikumar, M. J. Wainwright, and J. D. Lafferty, High-dimensional Ising model selection using ℓ_1 -regularized logistic regression, *Ann. Statist.* **38**, 1287 (2010).
- [38] E. Aurell and M. Ekeberg, Inverse Ising Inference Using All the Data, *Phys. Rev. Lett.* **108**, 090201 (2012).
- [39] K. Koh, S.-J. Kim, and S. Boyd, An interior-point method for large-scale ℓ_1 -regularized logistic regression, *J. Mach. Learn. Res.* **8**, 1519 (2007).
- [40] M. Belkin, D. Hsu, S. Ma, and S. Mandal, Reconciling modern machine-learning practice and the classical bias-variance trade-off, *Proc. Natl. Acad. Sci. U.S.A.* **116**, 15849 (2019).
- [41] B. Neal, S. Mittal, A. Baratin, V. Tantia, M. Scicluna, S. Lacoste-Julien, and I. Mitliagkas, A modern take on the bias-variance tradeoff in neural networks, [arXiv:1810.08591](https://arxiv.org/abs/1810.08591).
- [42] H. Huang and H. Zhou, Counting solutions from finite samplings, *Phys. Rev. E* **85**, 026118 (2012).
- [43] C. Li and H. Huang, Learning Credit Assignment, *Phys. Rev. Lett.* **125**, 178301 (2020).
- [44] M. Baity-Jesi, L. Sagun, M. Geiger, S. Spigler, G. B. Arous, C. Cammarota, Y. LeCun, M. Wyart, and G. Biroli, Comparing dynamics: Deep neural networks versus glassy systems, *J. Stat. Mech.* (2019) 124013.
- [45] P. C. Verpoort, A. A. Lee, and David J. Wales, Archetypal landscapes for deep neural networks, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 21857 (2020).
- [46] H. Yoshino, From complex to simple: Hierarchical free-energy landscape renormalized in deep neural networks, *SciPost Phys. Core* **2**, 005 (2020).
- [47] S. Spigler, M. Geiger, S. d'Ascoli, L. Sagun, G. Biroli, and M. Wyart, A jamming transition from under- to over-parametrization affects generalization in deep learning, *J. Phys. A* **52**, 474001 (2019).
- [48] M. Mezard and G. Parisi, The Bethe lattice spin glass revisited, *Eur. Phys. J. B* **20**, 217 (2001).
- [49] M. Mézard and A. Montanari, *Information, Physics, and Computation* (Oxford University Press, Oxford, 2009).
- [50] H. Huang, Theory of population coupling and applications to describe high order correlations in large populations of interacting neurons, *J. Stat. Mech.* (2017) 033501.
- [51] D. P. Kingma and J. L. Ba, Adam: A method for stochastic optimization, in *ICLR 2015 : International Conference on Learning Representations 2015* (2015).
- [52] H. Xiao, K. Rasul, and R. Vollgraf, Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms, [arXiv:1708.07747](https://arxiv.org/abs/1708.07747).
- [53] <https://github.com/ZouWenXuan/Liquid-like-deep-learning>.
- [54] T. Hastie, A. Montanari, S. Rosset, and R. J. Tibshirani, Surprises in high-dimensional ridgeless least squares interpolation, [arXiv:1903.08560](https://arxiv.org/abs/1903.08560).
- [55] S. Mei and A. Montanari, The generalization error of random features regression: Precise asymptotics and double descent curve, [arXiv:1908.05355](https://arxiv.org/abs/1908.05355).
- [56] A. Engel and C. Van den Broeck, *Statistical Mechanics of Learning* (Cambridge University Press, Cambridge, 2001).
- [57] M. S. Advani, A. M. Saxe, and H. Sompolinsky, High-dimensional dynamics of generalization error in neural networks, *Neural Networks* **132**, 428 (2020).
- [58] G. Tkacik, E. Schneidman, M. J. Berry, and W. Bialek, Spin glass models for a network of real neurons, [arXiv:0912.5409](https://arxiv.org/abs/0912.5409).
- [59] E. Bolthausen, An iterative construction of solutions of the tap equations for the Sherrington–Kirkpatrick model, *Commun. Math. Phys.* **325**, 333 (2014).
- [60] A. Georges and J. Yedidia, How to expand around mean-field theory using high-temperature expansions, *J. Phys. A* **24**, 2173 (1991).