

Modeling the Influence of Data Structure on Learning in Neural Networks: The Hidden Manifold Model

Sebastian Goldt^{1,*}, Marc Mézard^{1,†}, Florent Krzakala^{1,‡} and Lenka Zdeborová^{2,§}

¹Laboratoire de Physique de l'Ecole Normale Supérieure, Université PSL, CNRS, Sorbonne Université, Université Paris-Diderot, Sorbonne Paris Cité, 75005 Paris, France

²Institut de Physique Théorique, CNRS, CEA, Université Paris-Saclay, 91191 Gif-sur-Yvette, France



(Received 3 May 2020; revised 1 September 2020; accepted 8 September 2020; published 3 December 2020)

Understanding the reasons for the success of deep neural networks trained using stochastic gradient-based methods is a key open problem for the nascent theory of deep learning. The types of data where these networks are most successful, such as images or sequences of speech, are characterized by intricate correlations. Yet, most theoretical work on neural networks does not explicitly model training data or assumes that elements of each data sample are drawn independently from some factorized probability distribution. These approaches are, thus, by construction blind to the correlation structure of real-world datasets and their impact on learning in neural networks. Here, we introduce a generative model for *structured* datasets that we call the *hidden manifold model*. The idea is to construct high-dimensional inputs that lie on a lower-dimensional manifold, with labels that depend only on their position within this manifold, akin to a single-layer decoder or generator in a generative adversarial network. We demonstrate that learning of the hidden manifold model is amenable to an analytical treatment by proving a “Gaussian equivalence property” (GEP), and we use the GEP to show how the dynamics of two-layer neural networks trained using one-pass stochastic gradient descent is captured by a set of integro-differential equations that track the performance of the network at all times. This approach permits us to analyze in detail how a neural network learns functions of increasing complexity during training, how its performance depends on its size, and how it is impacted by parameters such as the learning rate or the dimension of the hidden manifold.

DOI: [10.1103/PhysRevX.10.041044](https://doi.org/10.1103/PhysRevX.10.041044)

Subject Areas: Computational Physics,
Statistical Physics

I. INTRODUCTION

The datasets on which modern neural networks are most successful, such as images [1,2] or natural language [3], are characterized by complicated correlations. Yet, most theoretical works on neural networks in statistics or theoretical computer science do not model the structure of the training data at all [4,5], which amounts to assuming that the dataset is chosen in a worst-case (adversarial) manner. A line of theoretical works complementary to the statistics approach emanated from statistical physics [6–9]. These works model inputs as elementwise independent identically

distributed (IID) random variables, with labels that are either random or given by some random but fixed function of the inputs. Despite providing valuable insights, these approaches are by construction blind to even basic statistical properties of real-world datasets such as their covariance structure. This lack of mathematical models for datasets is a major impediment for understanding the effectiveness of deep neural networks.

The structure present in realistic datasets can be illustrated well with classic datasets for image classification, such as the handwritten digits of the MNIST dataset [10] or the images of the CIFAR10 dataset [11]. The inputs that the neural network has to classify are images, so *a priori* the input space is the high-dimensional \mathbb{R}^N , corresponding to the number of pixels, with N large. However, the inputs that can be recognized as actual images rather than random noise span but a lower-dimensional manifold within \mathbb{R}^N ; see Fig. 1. This manifold hence constitutes the actual input space, or the “world,” of our problem. While the manifold is not easily defined, it is tangible: For example, its dimension can be estimated based on the neighborhoods of inputs in the dataset [12–15] and is found to be around

*sebastian.goldt@phys.ens.fr

†marc.mezard@ens.fr

‡florent.krzakala@ens.fr

§lenka.zdeborova@cea.fr

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

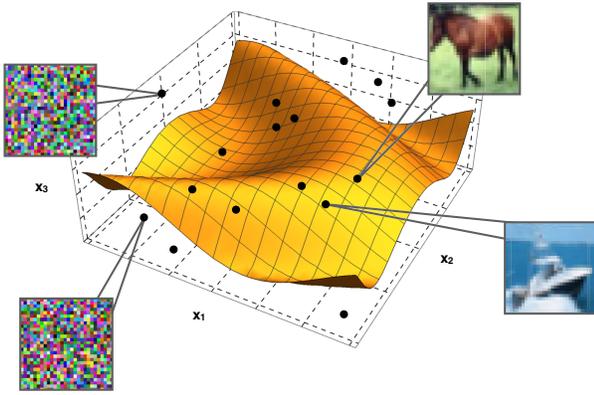


FIG. 1. We illustrate the notion of a hidden manifold in input space using CIFAR10 example images. Each black point indicates a possible input in a high-dimensional input space \mathbb{R}^N . Most points in this space cannot be interpreted as images at all; however, those points that can be interpreted as real images tend to concentrate on a lower-dimensional manifold, here sketched as a two-dimensional curved surface in a three-dimensional space. The intrinsic dimension D of these lower-dimensional manifolds has been measured numerically [12–15].

$D \approx 14$ for MNIST and $D \approx 35$ for CIFAR10, compared to $N = 784$ and $N = 3072$, respectively. We call inputs *structured* if they are concentrated on a lower-dimensional manifold and, thus, have a lower-dimensional latent representation, which consists of the coordinates of the input within that manifold.

A complementary view on the data manifold is provided by today’s most powerful generative models, called generative adversarial networks (GANs) [16]. A GAN \mathcal{G} is a neural network that is trained to take random noise as its input and to transform it into outputs that resemble a given target distribution. For example, GANs can generate realistic synthetic images of human faces [17,18]. From this point of view, the mapping from the hidden manifold to the input space is given by the function that the GAN \mathcal{G} computes.

A. Main results

In this paper, and specifically in Sec. II, we introduce a generative model for structured datasets in the above sense that we call the *hidden manifold model* (HMM) [19]. It is a generative model that produces tuples (x, y^*) of high-dimensional inputs $x \in \mathbb{R}^N$ and their scalar labels y^* . The key idea is to construct the inputs such that they lie on a lower-dimensional manifold; their labels are then a function of only their position within that manifold. The way the inputs are generated is akin to a learned single-layer decoder with random inputs; it can also be viewed as a single-layer generator neural network of a learned GAN. As a result, inputs drawn from the HMM have nontrivial correlations and do not follow a normal distribution, and

their labels y^* cannot be written as a simple function of the inputs x .

Our key theoretical result, presented in Sec. III, is to show that, despite these correlations, the HMM is amenable to an analytical treatment in a thermodynamic limit of large dimensions N and D , large number of samples P , and fixed respective ratios D/N and P/N . We derive the solution by first demonstrating a “Gaussian equivalence property” (GEP) (Proposition III.1). As a first application, we use the GEP to derive a set of integro-differential equations that captures the behavior of two-layer neural networks, with $K = \mathcal{O}(1)$ hidden units, trained using stochastic gradient descent. These equations extend the classical analysis of the dynamics of two-layer neural networks on unstructured data [20–23] to the hidden manifold and provide detailed insight into the dynamics of learning.

We then use these equations to study the dynamics and the performance of two-layer neural networks trained on data generated by the HMM, in Sec. IV. We find the specialization of hidden units, known from the canonical teacher-student model. We analyze the learning for different feature matrices and show that Hadamard matrices perform slightly better than IID Gaussian ones. We show analytically that the generalization performance deteriorates as the manifold dimension D grows. We show that the learning rate has a very minor influence on the asymptotic error and analyze how the final error of the network changes as a function of the width of the hidden layer.

Section V is devoted to a comparison of learning on the HMM and on real datasets such as MNIST [10], fashion-MNIST [24], or CIFAR10 [11]. In particular, we demonstrate that neural networks learn functions of increasing complexity over the course of training on both the HMM and real datasets. We also compare the memorization of some samples during the early stages of training between the HMM to real data. These comparisons provide strong evidence that the HMM captures the properties of learning with one-pass stochastic gradient descent (SGD) and two-layer neural networks on some of the standard benchmark datasets rather faithfully.

B. Further related work

1. The need for models of structured data

Several works recognize the importance of the structure in the datasets used for machine learning and, in particular, the need to go beyond the simple componentwise IID modeling [25–30]. While we focus on the ability of neural networks to generalize from examples, two recent papers study a network’s ability to *store* inputs with lower-dimensional structure and random labels: Chung, Lee, and Sompolinsky [31] study the linear separability of general, finite-dimensional manifolds and their interesting consequences for the training of deep neural networks [32,33], while Rotondo, Cosentino Lagomarsino, and

Gherardi [34] extend Cover’s classic argument [35] to count the number of learnable dichotomies when inputs are grouped in tuples of k inputs with the same label. Recently, Yoshida and Okada [36] analyzed the dynamics of online learning for data having an arbitrary covariance matrix, finding an infinite hierarchy of ordinary differential equations (ODEs). Their study implicitly assumes that inputs have a Gaussian distribution, while our approach handles a more general data structure. The importance of the spectral properties of the data is recognized for learning in deep neural networks by Saxe, McClelland, and Ganguli [37] in the special case of *linear* neurons, where the whole network performs a linear transformation of the data.

2. Relation to random feature learning

The hidden manifold model has an interesting link to random feature learning with unstructured IID input data. The idea of learning with random features goes back to the mechanical perceptron of the 1960s [38] and was extended into the random kitchen sinks of Rahimi and Recht [39,40]. Remarkably, random feature learning in the same scaling limit as used in the theoretical part of this paper is analyzed in several recent and concurrent works, notably in Refs. [41,42] for ridge regression and in Ref. [43] for max-margin linear classifiers. These papers consider full batch learning (i.e., all samples are used at the same time), which makes one difference from our online (one-pass stochastic gradient descent) analysis. Another important difference is that we study learning in a neural network with two layers of learned weights, while the existing works study simpler linear (perceptron-type) architectures where only one layer is learned. Perhaps more importantly, in our analysis, the features do not need to be random but can be chosen deterministically or even be learned from data using a GAN or an autoencoder. The principles underlying the analytic solution of this paper as well as Refs. [41–43] rely on the Gaussian equivalence property, which is stated and used independently in those papers.

3. Gaussian equivalence and random matrix theory

Special cases of the Gaussian equivalence property were, in fact, derived previously using random matrix theory in Refs. [44–47], and this equivalent Gaussian covariates mapping is explicitly stated and used in Refs. [42,43]. This formulation has recently been extended to a broader setting of concentrated vectors encompassing data coming from a GAN in Refs. [48,49], a version closer to our formulation.

C. Reproducibility

We provide the full code to reproduce our experiments as well as an integrator for the equations of motion of two-layer networks online [50].

D. Learning setup

This paper focuses on the dynamics and performance of fully connected two-layer neural networks with K hidden units and first- and second-layer weights $W \in \mathbb{R}^{K \times N}$ and $v \in \mathbb{R}^K$, respectively. Given an input $x \in \mathbb{R}^N$, the output of a network with parameters $\theta = (W, v)$ is given by

$$\phi(x; \theta) = \sum_k^K v_k g(w_k x / \sqrt{N}), \quad (1)$$

where w_k is the k th row of W and $g: \mathbb{R} \rightarrow \mathbb{R}$ is the nonlinear activation function of the network, acting componentwise. We study sigmoidal and rectified linear unit (ReLU) networks with $g(x) = \text{erf}(x/\sqrt{2})$ and $g(x) = \max(0, x)$, respectively.

We train the neural network on datasets with P input-output pairs (x_μ, y_μ^*) , $\mu = 1, \dots, P$, where we use the starred y_μ^* to denote the *true* label of an input x_μ . Networks are trained by minimizing the quadratic training error $E(\theta) = 1/2 \sum_{\mu=1}^P \Delta_\mu^2$ with $\Delta_\mu = [\phi(x_\mu, \theta) - y_\mu^*]$ using stochastic (one-pass, online) gradient descent with constant learning rate η and minibatch size 1:

$$\theta_{\mu+1} = \theta_\mu - \eta \nabla_\theta E(\theta)|_{\theta_\mu, x_\mu, y_\mu^*}. \quad (2)$$

Initial weights for both layers are always taken componentwise IID from the normal distribution with mean 0 and standard deviation 10^{-3} .

The key quantity of interest is the *test error* or *generalization error* of a network, for which we compare its predictions to the labels given in a test set that is composed of P^* input-output pairs (x_μ, y_μ^*) , $\mu = 1, \dots, P^*$, that are *not* used during training:

$$\epsilon_g(\theta) \equiv \frac{1}{2P^*} \sum_\mu^{P^*} [\phi(x_\mu, \theta) - y_\mu^*]^2. \quad (3)$$

The test set in our setting is generated by the same probabilistic model that generated the training data.

1. The canonical teacher-student model

The joint probability distribution of input-output pairs (x_μ, y_μ^*) is inaccessible for realistic datasets such as CIFAR10, preventing analytical control over the test error and other quantities of interest. To make theoretical progress, it is therefore promising to study the generalization ability of neural networks for data arising from a probabilistic generative model.

A classic model for datasets is the *canonical teacher-student setup*, where inputs x_μ are drawn elementwise IID from the standard normal distribution and labels are given by a random, but fixed, neural network with weights θ^* acting

on the inputs: $y_\mu^* = \phi(x_\mu, \theta^*)$. The network that generates the labels is called the teacher, while the network that is trained is called the student. The model was introduced by Gardner and Derrida [6], and its study provides many valuable insights into the generalization ability of neural networks from an average-case perspective, particularly within the framework of statistical mechanics [7–9, 51–56] and also in recent works in theoretical computer science, e.g., Refs. [42, 57–59]. However, it has the notable shortcoming that its analysis crucially relies on the fact that inputs are IID Gaussians and, hence, uncorrelated.

II. THE HIDDEN MANIFOLD MODEL

We now introduce a new generative probabilistic model for structured datasets with correlations. To generate a dataset containing P inputs in N dimensions, we first choose D feature vectors f_r , $r = 1, \dots, D$. These are vectors in N dimensions, and we collect them in a feature matrix $F \in \mathbb{R}^{D \times N}$. Next, we draw P vectors c_μ with random IID components drawn from the normal distribution with mean zero and unit variance and collect them in the matrix $C \in \mathbb{R}^{P \times D}$. The vector c_μ gives the coordinates of the μ th input on the lower-dimensional manifold spanned by the feature vectors in F . We call c_μ the *latent representation* of the input x_μ , which is given by the μ th row of

$$X = f(CF/\sqrt{D}) \in \mathbb{R}^{P \times N}, \quad (4)$$

where f is a nonlinear function acting componentwise. In this model, the “world” of the data on which the true label can depend is a D -dimensional manifold, which is obtained from the linear span of F through a “folding” process induced by the nonlinear function f . We note that the structure of data of the same type arises in a learned variational autoencoder network [60] with a single layer, or, in a learned GAN network [16] with a single-layer generator network, the matrix C then corresponds to the random input, the F to the learned features, and f is the corresponding output activation. The matrix F can be generic with a certain normalization, such that its elements are $\mathcal{O}(1)$. For our analysis to be valid, we later assume the normalization given in Eq. (13) and balance condition given by Eq. (14); other than that, our analysis holds for arbitrary matrices F .

The labels are obtained by applying a two-layer neural network with weights $\tilde{\theta} = (\tilde{W} \in \mathbb{R}^{M \times D}, \tilde{v} \in \mathbb{R}^M)$ within the unfolded hidden manifold according to

$$y_\mu^* = \phi(c_\mu, \tilde{\theta}) = \sum_m^M \tilde{v}^m \tilde{g}(\tilde{w}^m c_\mu / \sqrt{D}). \quad (5)$$

We draw the weights in both layers componentwise IID from the normal distribution with unity variance, unless we note it otherwise. The key point here is the dependency of

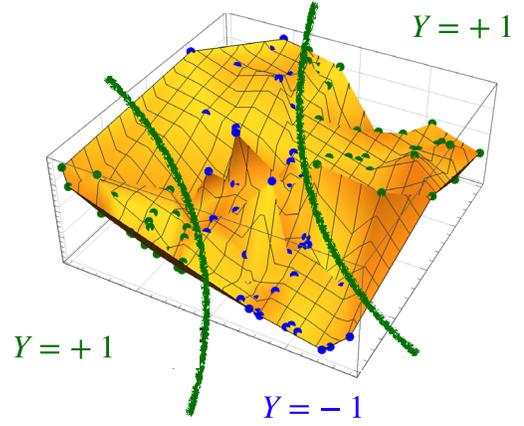


FIG. 2. The hidden manifold model proposed here is a generative model for structured datasets, where inputs x [Eq. (4)] (blue and green balls) concentrate on a lower-dimensional manifold in input space (yellow surface). Their label y^* is a function of their position on the manifold; here, we show the setup of a classification task with two classes $y^* = \pm 1$. In our analysis, the labels are generated according to Eq. (5).

labels y_μ^* on the coordinates of the lower-dimensional manifold c_μ rather than on the high-dimensional data x_μ as illustrated in Fig. 2. We expect the exact functional form of this dependence not to be crucial for the empirical part of this work and that there are other forms that would present the same behavior. Notably, it would be interesting to consider ones where the latent representation is conditioned to the labels as in conditional GANs [61] or the manifold model of Ref. [33].

III. THE SOLUTION OF THE HIDDEN MANIFOLD MODEL

A. The Gaussian equivalence property

The difficulty in analyzing HMM comes from the fact that the various components of one given input pattern, say, $x_{\mu i}$ and $x_{\nu j}$, are correlated. Yet, a key feature of the model is that it is amenable to an analytical treatment. To that end, we study the standard thermodynamic limit of the statistical physics of learning where the size of the input space $N \rightarrow \infty$, together with the number $P \rightarrow \infty$ of patterns that are presented for learning, while keeping the ratio $\alpha \equiv P/N$ fixed. In statistics, this limit corresponds to the challenging high-dimensional limit. The hidden manifold model can then be studied analytically if one assumes that the latent dimension D , i.e., the dimension of the feature space, also scales with N , meaning that it goes to ∞ with a fixed ratio $\delta \equiv D/N$ which is of the order of 1 with respect to N , so that we have

$$N, P, D \rightarrow \infty, \quad \text{with fixed } \alpha \equiv \frac{P}{N} \quad \text{and} \quad \delta \equiv \frac{D}{N}. \quad (6)$$

In this limit, the relevant variables are the “local fields” or preactivations that are acted upon by the neurons in the

hidden layer. They can be shown to follow a Gaussian distribution in the thermodynamic limit (6). We now make this statement precise by formulating a GEP. We demonstrate the power of this equivalence by deriving a set of exact equations for online learning in Sec. III B.

1. Statement of the property

Let $\{C_r\}_{r=1}^D$ be D IID Gaussian random variables distributed as $\mathcal{N}(0, 1)$. In the following, we denote by \mathbb{E} the expectation value with respect to this distribution. Define N variables $u_i, i = 1, \dots, N$, as linear superpositions of the C_r variables:

$$u_i \equiv \frac{1}{\sqrt{D}} \sum_{r=1}^D C_r F_{ir}, \quad (7)$$

and M variables $\nu^m, m = 1, \dots, M$, as other linear superpositions:

$$\nu^m \equiv \frac{1}{\sqrt{D}} \sum_{r=1}^D C_r \tilde{w}_r^m, \quad (8)$$

where \tilde{w}_r^m are the teacher weights [Eq. (5)]. Define K variables λ^k as linear superpositions of $f(u_i)$, where f is an arbitrary function:

$$\lambda^k \equiv \frac{1}{\sqrt{N}} \sum_{i=1}^N w_i^k f(u_i), \quad (9)$$

where \tilde{w}_i^k are the student weights [Eq. (1)]. We occasionally write (λ, ν) to denote the tuple of all local fields λ^k and ν^m . Denoting by $\langle g(u) \rangle$ the expectation of a function $g(u)$ when u is a normal variable with distribution $u \sim \mathcal{N}(0, 1)$, we also introduce for convenience the ‘‘centered’’ variables:

$$\tilde{\lambda}^k \equiv \frac{1}{\sqrt{N}} \sum_{i=1}^N w_i^k [f(u_i) - \langle f(u) \rangle]. \quad (10)$$

Notice that our notation keeps upper indices for indices which take values in a finite range ($k, \ell \in \{1, \dots, K\}$, $m, n \in \{1, \dots, M\}$), and lower indices for those which have a range of the order of N ($i, j \in \{1, \dots, N\}$; $r, s \in \{1, \dots, D\}$).

As the C_r are Gaussian, the u_i variables are also Gaussian variables, with mean zero and a matrix of covariance

$$U_{ij} = \mathbb{E}[u_i u_j] = \frac{1}{D} \sum_{r=1}^D F_{ir} F_{jr}. \quad (11)$$

Note that the covariances of the u_i variables scale in the thermodynamic limit as

$$\mathbb{E}[u_i^2] = 1; \quad \mathbb{E}[u_i u_j] = \mathcal{O}(1/\sqrt{D}), \quad i \neq j. \quad (12)$$

We assume that, in the thermodynamic limit, the W, \tilde{W} , and F matrices have elements of $\mathcal{O}(1)$ and that, for $i \neq j$,

$$\frac{1}{\sqrt{D}} \sum_{r=1}^D F_{ir} F_{jr} = \mathcal{O}(1) \quad \text{and} \quad \sum_{r=1}^D (F_{ir})^2 = D. \quad (13)$$

Notice that the only variables which are drawn IID from a Gaussian distribution are the coefficients C_r . Most importantly, the matrices F and W can be arbitrary (and deterministic) as long as they are ‘‘balanced’’ in the sense that $\forall p, q \geq 1, \forall k_1, \dots, k_p, r_1, \dots, r_q$, we have

$$S_{r_1 r_2 \dots r_q}^{k_1 k_2 \dots k_p} = \frac{1}{\sqrt{N}} \sum_i w_i^{k_1} w_i^{k_2} \dots w_i^{k_p} F_{ir_1} F_{ir_2} \dots F_{ir_q} = \mathcal{O}(1), \quad (14)$$

with the q and p distinct. We also have a similar scaling for the combinations involving the teacher weights \tilde{w}_r^m . This assumption is the key behind the Gaussian equivalence property, and we discuss its interpretation immediately after the statement of the GEP.

Property III.1. Gaussian equivalence property (GEP).— In the asymptotic limit when $N \rightarrow \infty, D \rightarrow \infty$, with K, M , and the ratio D/N finite, and under the assumption (14), $\{\lambda^k\}$ and $\{\nu^m\}$ are $K + M$ jointly Gaussian variables, with mean

$$\mathbb{E}[\lambda^k] = a \frac{1}{\sqrt{N}} \sum_{i=1}^N w_i^k; \quad \mathbb{E}[\nu^m] = 0 \quad (15)$$

and covariance

$$Q^{k\ell} \equiv \mathbb{E}[\tilde{\lambda}^k \tilde{\lambda}^\ell] = (c - a^2 - b^2) W^{k\ell} + b^2 \Sigma^{k\ell}, \quad (16)$$

$$R^{km} \equiv \mathbb{E}[\tilde{\lambda}^k \nu^m] = b \frac{1}{D} \sum_{r=1}^D S_r^k \tilde{w}_r^m, \quad (17)$$

$$T^{mn} \equiv \mathbb{E}[\nu^m \nu^n] = \frac{1}{D} \sum_{r=1}^D \tilde{w}_r^m \tilde{w}_r^n. \quad (18)$$

The ‘‘folding function’’ $f(\cdot)$ appears through the three coefficients a, b , and c , which are defined as

$$a \equiv \langle f(u) \rangle, \quad b \equiv \langle u f(u) \rangle, \quad c \equiv \langle f(u)^2 \rangle, \quad (19)$$

respectively, where $\langle \psi(u) \rangle$ denotes the expectation value of the function ψ when $u \sim \mathcal{N}(0, 1)$ is a Gaussian variable.

The covariances are defined in terms of the three matrices

$$S_r^k \equiv \frac{1}{\sqrt{N}} \sum_{i=1}^N w_i^k F_{ir}, \quad (20)$$

$$W^{k\ell} \equiv \frac{1}{N} \sum_{i=1}^N w_i^k w_i^\ell, \quad (21)$$

$$\Sigma^{k\ell} \equiv \frac{1}{D} \sum_{r=1}^D S_r^k S_r^\ell, \quad (22)$$

whose elements are assumed to be of the order of $\mathcal{O}(1)$ in the asymptotic limit. The derivation of the property is given in the Appendix A.

In Sec. III B, we see that the GEP allows us to develop an analytical understanding of learning with the hidden manifold model. We first discuss several aspects of the GEP in detail.

2. Discussion

The Gaussian equivalence property states that the local fields (λ, ν) follow a joint normal distribution if the weights of the student fulfill the balance condition (14). In the simplest case, where the x_i are elementwise IID Gaussian, joint Gaussianity of (λ, ν) follows immediately from the central limit theorem (CLT). The CLT can also be applied directly when input vectors x are drawn from a multidimensional Gaussian with a fixed covariance matrix, which is the setup of Yoshida and Okada [36]. In the hidden manifold model considered here, the inputs $x_i = f(u_i)$ and, thus, the x_i are not normally distributed and have a nontrivial covariance matrix. The GEP can be seen as a central limit theorem for sums of weakly correlated random variables, i.e., $\lambda^k \sim \sum_i w_i^k x_i$. In this case, the GEP establishes that λ^k is Gaussian, provided that the weights of the student w_i^k do not align “too much” with the weights of the generator F_{ir} . More precisely, we require that a sum such as $S_r^k = 1/\sqrt{N} \sum_i w_i^k F_{ir}$ remains of the order of 1 in the thermodynamic limit (6). The balance condition is a generalization of this idea to the higher-order tensors defined in Eq. (14).

The expansion from the hidden manifold in \mathbb{R}^D to the input space \mathbb{R}^N can equivalently be seen as a noisy transformation of the latent variables C . As far as the local fields (λ, ν) are concerned, we can replace the data matrix $X = f(CF/\sqrt{D})$ with the matrix

$$\tilde{X} \simeq a\mathbb{1} + bCF + (c - a^2 - b^2)Z, \quad (23)$$

where Z is a $P \times N$ matrix with entries drawn IID from the normal distribution and $\mathbb{1}$ is a matrix of the same size as X with all entries equal to one. We use the symbol \simeq here to emphasize that the two matrices on the left- and right-hand sides have matching first and second moments and are, hence, equivalent in terms of their low-dimensional

projections but are not the same matrix. We can, thus, think of the inputs X as a noisy transformation of the latent variables, even without any explicit noise in Eq. (4).

We could also add noise to the expansion explicitly, for example, as $X = f(FC/\sqrt{D}) + \zeta$ or $X = f(F(C + \zeta)/\sqrt{D})$, where ζ would be a noise matrix of appropriate dimensions. These noise injections would indeed make the data high dimensional or, if added directly to the latent variables C , result in correlated noise in the input space. In all these cases, the GEP applies and guarantees that the noise ζ would change only the variance of the noise term Z that appears after application of the GEP (23). Our results are, thus, robust to the injection of additional noise.

Finally, the GEP shows that there is a whole family of activation functions $f(x)$ [those that have the same values for a , b , and c from Eq. (19)] that lead to equivalent analytical results for the learning curves studied in this paper.

Related results in random matrix theory.—A related result to the Gaussian equivalence property is, in fact, known in random matrix theory [41,44–48]. These works study quantities that can be written as an integral over the spectral density of the distribution of inputs x , such as the test and training errors for a linear regression problem. However, this spectral density is inaccessible analytically for realistic data. The key idea is then to rewrite these integrals by replacing the intractable spectral density with the spectral density of a Gaussian model with matching first and second moments. Using tools from random matrix theory (RMT), one can show that certain integrals over both spectra coincide. This mapping is explicitly used in Refs. [42,43]. In order to apply tools from RMT, these works have to assume that the weights F of the generator are random. The advantage of the formulation of the GEP above is that it does *not* require the matrix F to be a random one and is valid as well for deterministic or learned weight matrices, as long as the balanced conditions stated in Eqs. (13) and (14) hold. This advantage allows one to generalize these mappings to the case of deterministic features using Hadamard and Fourier matrices, such as the one used in Fastfood [62] or ACDC [63] layers. These orthogonal projections are actually known to be more effective than the purely random ones [64]. It also allows generalization of the analysis in this paper for data coming from a learned GAN, along the lines of Refs. [48,49]. We illustrate this point below by analyzing the dynamics of online learning when the feature matrix F is a deterministic Hadamard matrix (cf. Sec. IV B).

B. The dynamics of stochastic gradient descent for the hidden manifold model

To illustrate the power of the GEP, we now analyze the dynamics of stochastic gradient descent (2) in the case of *online learning*, where, at each step of the algorithm $\mu = 1, 2, \dots$, the student’s weights are updated according

to Eq. (2) using a previously unseen sample (x_μ, y_μ) . This case is also known as one-shot or single-pass SGD. The analysis of online learning has been performed previously for the canonical teacher-student model with IID Gaussian inputs [20–23,65] and has recently been put on a rigorous foundation [55]. Here, we generalize this type of analysis to two-layer neural networks trained on the hidden manifold model.

The goal of our analysis is to track the mean-squared generalization error of the student with respect to the teacher at all times:

$$\epsilon_g(\theta, \tilde{\theta}) \equiv \frac{1}{2} \mathbb{E}[\phi(x, \theta) - \tilde{y}^*]^2, \quad (24)$$

where the expectation \mathbb{E} denotes an average over an input drawn from the hidden manifold model [Eq. (4)] with label $y_\mu^* = \phi(c_\mu, \tilde{\theta}^*)$ given by a teacher network with fixed weights $\tilde{\theta}^*$ acting on the latent representation [Eq. (5)]. Note that the weights of both the student and the teacher, as well as the feature matrix F_{ir} , are held fixed when taking the average, which is an average only over the coefficients $c_{\mu r}$. To keep notation compact, we focus on cases where $a = \mathbb{E}f(u) = 0$ in Eq. (19), which leads to $\tilde{\lambda}^k = \lambda^k$ in Eq. (10). A generalization to the case where $a \neq 0$ is straightforward but lengthy.

We can make progress with the high-dimensional average over x in Eq. (24) by noticing that the input x and its latent representation c enter the expression only via the local fields ν^m and λ^k [Eqs. (8) and (9)]:

$$\epsilon_g(\theta, \tilde{\theta}) = \frac{1}{2} \mathbb{E} \left(\sum_k^K v^k g(\lambda^k) - \sum_m^M \tilde{v}^m \tilde{g}(\nu^m) \right)^2. \quad (25)$$

The average is now taken over the joint distribution of local fields $\{\lambda^{k=1,\dots,K}, \nu^{m=1,\dots,M}\}$. The key step is then to invoke the Gaussian equivalence Property III.1, which guarantees that this distribution is a multivariate normal distribution with covariances $Q^{k\ell}$, R^{km} , and T^{nm} [Eqs. (16)–(18)]. Depending on the choice of $g(x)$ and $\tilde{g}(x)$, this distribution makes it possible to compute the average analytically; in any case, the GEP guarantees that we can express $\epsilon_g(\theta, \tilde{\theta})$ as a function of only the second-layer weights v^k and \tilde{v}^m and the matrices $Q^{k\ell}$, R^{km} , and T^{nm} , which are called *order parameters* in statistical physics [20–22]:

$$\lim_{N,D \rightarrow \infty} \epsilon_g(\theta, \tilde{\theta}) = \epsilon_g(Q^{k\ell}, R^{km}, T^{nm}, v^k, \tilde{v}^m), \quad (26)$$

where, in taking the limit, we keep the ratio $\delta \equiv D/N$ finite [see Eq. (6)].

1. The physical interpretation of the order parameters

The order parameter R^{kn} , defined in Eqs. (17) and (20), measures the similarity between the action of the k th student node on an input x_μ and the n th teacher node acting on the corresponding latent representation c_μ . In the canonical teacher-student setup, where (i) the input covariance is simply $\mathbb{E}x_i x_j = \delta_{ij}$ and (ii) labels are generated by the teacher acting directly on the inputs x , it can be readily verified that the overlap has the simple expression $R^{kn} \equiv \mathbb{E}\lambda^k \nu^n \sim w^k \tilde{w}^n$. It was, hence, called the teacher-student overlap in the previous literature. In the HMM, however, where the teacher and student networks act on different vector spaces, it is not *a priori* clear how to express the teacher-student overlap in suitable order parameters.

The matrix $Q^{k\ell} = [c - b^2]W^{k\ell} + b^2\Sigma^{k\ell}$ quantifies the similarity between two student nodes k and ℓ and has two contributions: the latent student-student overlap $\Sigma^{k\ell}$, which measures the overlap of the weights of two students nodes after they are projected to the hidden manifold, and the ambient student-student overlap $W^{k\ell}$, which measures the overlap between the vectors $w^k, w^\ell \in \mathbb{R}^N$. Finally, we also have that the overlaps of the teacher nodes are collected in the matrix T^{nm} , which is *not* time dependent, as it is a function of the teacher weights only.

2. Statement of the equations of motion

We derive a closed set of equations of motion that describe the dynamics of the order parameters R^{km} , $\Sigma^{k\ell}$, $W^{k\ell}$, and v^k when the student is trained using online SGD (2). We stress at this point that in the online learning, at each step of SGD, a new sample is given to the network. The weights of the network are, thus, uncorrelated to this sample, and, hence, the GEP can be applied at every step. This approach is in contrast with the full-batch learning, where the correlations between weights and inputs have to be taken into account explicitly [43]. Integrating the equations of motion and substituting the values of the order parameters into Eq. (26) gives the generalization error at all times. Here, we give a self-contained statement of the equations and relegate the details of the derivation to Appendix B.

A key object in our analysis is the spectrum of the matrix

$$\Omega_{rs} \equiv \frac{1}{N} \sum_i F_{ir} F_{is}. \quad (27)$$

We denote its eigenvalues and corresponding eigenvectors by ρ and ψ_ρ and write $p_\Omega(\rho)$ for the distribution of eigenvalues. It turns out that it is convenient to rewrite the teacher-student overlap as an integral over a density $r^{km}(\rho, t)$, which is a function of ρ and of the normalized number of steps $t = P/N$, which can be interpreted as a continuous timelike variable. We then have

$$R^{km}(t) = b \int d\rho p_\Omega(\rho) r^{km}(\rho, t) \quad (28)$$

with $b \equiv \langle uf(u) \rangle$ [Eq. (19)]. In the canonical teacher-student model, introducing such a density and the integral that comes with it is not necessary, but in the HMM it is a consequence of the nontrivial correlation matrix $\mathbb{E}x_i x_j$

between input elements. Adopting the convention that the indices $j, k, \ell, \iota = 1, \dots, K$ always denote *student* nodes, while $n, m = 1, \dots, M$ are reserved for teacher hidden nodes.

The equation of motion of the teacher-student density can then be written as

$$\begin{aligned} \frac{\partial r^{km}(\rho, t)}{\partial t} = & -\frac{\eta}{\delta} v^k d(\rho) \left(r^{km}(\rho) \sum_{j \neq k}^K v^j \frac{Q^{jj} I_3(k, k, j) - Q^{kj} I_3(k, j, j)}{Q^{jj} Q^{kk} - (Q^{kj})^2} + \sum_{j \neq k}^K v^j r^{jm}(\rho) \frac{Q^{kk} I_3(k, j, j) - Q^{kj} I_3(k, k, j)}{Q^{jj} Q^{kk} - (Q^{kj})^2} \right. \\ & + v^k r^{km}(\rho) \frac{1}{Q^{kk}} I_3(k, k, k) - r^{km}(\rho) \sum_n^M \tilde{v}^n \frac{T^{nn} I_3(k, k, n) - R^{kn} I_3(k, n, n)}{Q^{kk} T^{nn} - (R^{kn})^2} \\ & \left. - \frac{b\rho}{d(\rho)} \sum_n^M \tilde{v}^n \tilde{T}^{nm} \frac{Q^{kk} I_3(k, n, n) - R^{kn} I_3(k, k, n)}{Q^{kk} T^{nn} - (R^{kn})^2} \right), \end{aligned} \quad (29)$$

where $d(\rho) = (c - b^2)\delta + b^2\rho$. The teacher-teacher overlap $T^{nm} \equiv \mathbb{E}\nu^n \nu^m$ [Eq. (18)], while \tilde{T}^{nm} is the overlap of the teacher weights after rotation into the eigenbasis of Ω_{rs} , weighted by the eigenvalues ρ :

$$\tilde{T}^{nm} \equiv \frac{1}{D} \sum_\tau \rho_\tau \tilde{\omega}_\tau^m \tilde{\omega}_\tau^n, \quad \text{where } \tilde{\omega}_\tau^m = \frac{1}{\sqrt{D}} \sum_r \tilde{w}_r^m \psi_{r\tau}. \quad (30)$$

In writing the equations, we use the following shorthand for the three-dimensional Gaussian averages:

$$I_3(k, j, n) \equiv \mathbb{E}[g(\lambda^k) \lambda^j \tilde{g}(\nu^n)], \quad (31)$$

which was introduced by Saad and Solla [22]. Arguments passed to I_3 should be translated into local fields on the right-hand side by using the convention where the indices j, k, ℓ , and ι always refer to student local fields λ^j , etc., while the indices n and m always refer to teacher local fields ν^n and ν^m , respectively. Similarly,

$$I_3(k, j, j) \equiv \mathbb{E}[g(\lambda^k) \lambda^j g(\lambda^j)], \quad (32)$$

where having the index j as the third argument means that the third factor is $g(\lambda^j)$ rather than $\tilde{g}(\nu^m)$ in Eq. (31). The average in Eq. (31) is taken over a three-dimensional normal distribution with mean zero and covariance matrix

$$\Phi^{(3)}(k, j, n) = \begin{pmatrix} Q^{kk} & Q^{kj} & R^{kn} \\ Q^{kj} & Q^{jj} & R^{jn} \\ R^{kn} & R^{jn} & T^{nn} \end{pmatrix}. \quad (33)$$

For the latent student-student overlap $\Sigma^{k\ell}$, it is again convenient to introduce the density $\sigma^{k\ell}(\rho, t)$ as

$$\Sigma^{k\ell}(t) = \int d\rho p_\Omega(\rho) \sigma^{k\ell}(\rho, t), \quad (34)$$

whose equation of motion is given by

$$\begin{aligned} \frac{\partial \sigma^{k\ell}(\rho, t)}{\partial t} = & -\frac{\eta}{\delta} \left(d(\rho) v^k \sigma^{k\ell}(\rho) \sum_{j \neq k}^K v^j \frac{Q^{jj} I_3(k, k, j) - Q^{kj} I_3(k, j, j)}{Q^{jj} Q^{kk} - (Q^{kj})^2} + v^k \sum_{j \neq k}^K v^j d(\rho) \sigma^{j\ell}(\rho) \frac{Q^{kk} I_3(k, j, j) - Q^{kj} I_3(k, k, j)}{Q^{jj} Q^{kk} - (Q^{kj})^2} \right. \\ & + d(\rho) v^k \sigma^{k\ell}(\rho) v^k \frac{1}{Q^{kk}} I_3(k, k, k) - d(\rho) v^k \sigma^{k\ell}(\rho) \sum_n^M \tilde{v}^n \frac{T^{nn} I_3(k, k, n) - R^{kn} I_3(k, n, n)}{Q^{kk} T^{nn} - (R^{kn})^2} \\ & \left. - b\rho v^k \sum_n^M \tilde{v}^n r^{\ell n}(\rho) \frac{Q^{kk} I_3(k, n, n) - R^{kn} I_3(k, k, n)}{Q^{kk} T^{nn} - (R^{kn})^2} + \text{all of the above with } \ell \rightarrow k, k \rightarrow \ell \right) \\ & + \eta^2 v^k v^\ell \left[(c - b^2)\rho + \frac{b^2}{\delta} \rho^2 \right] \left(\sum_{j, \iota}^K v^j v^\iota I_4(k, \ell, j, \iota) - 2 \sum_j^K \sum_m^M v^j \tilde{v}^m I_4(k, \ell, j, m) + \sum_{n, m}^M \tilde{v}^n \tilde{v}^m I_4(k, \ell, n, m) \right). \end{aligned} \quad (35)$$

This equation involves again the integrals I_3 and a four-dimensional average that we denote

$$I_4(k, \ell, j, n) \equiv \mathbb{E}[g(\lambda^k)g(\lambda^\ell)g(\lambda^j)g(\lambda^n)] \quad (36)$$

using the same notational conventions as for I_3 , so the four-dimensional covariance matrix reads

$$\Phi^{(4)}(k, \ell, j, n) = \begin{pmatrix} Q^{kk} & Q^{k\ell} & Q^{kj} & R^{kn} \\ Q^{k\ell} & Q^{\ell\ell} & Q^{\ell j} & R^{\ell n} \\ Q^{kj} & Q^{\ell j} & Q^{jj} & R^{jn} \\ R^{kn} & R^{\ell n} & R^{jn} & T^{nn} \end{pmatrix}. \quad (37)$$

The equation of motion for the ambient student-student overlap $W^{k\ell}$ can be written directly:

$$\begin{aligned} \frac{dW^{k\ell}(t)}{dt} = & -\eta v^k \left(\sum_j v^j I_3(k, \ell, j) - \sum_n \tilde{v}^n I_3(k, \ell, n) \right) - \eta v^\ell \left(\sum_j v^j I_3(\ell, k, j) - \sum_n \tilde{v}^n I_3(\ell, k, n) \right) \\ & + c\eta^2 v^k v^\ell \left(\sum_{j,a} v^j v^a I_4(k, \ell, j, a) - 2 \sum_j \sum_m v^j \tilde{v}^m I_4(k, \ell, j, m) + \sum_{n,m} \tilde{v}^n \tilde{v}^m I_4(k, \ell, n, m) \right). \end{aligned} \quad (38)$$

Finally, the ODE for the second-layer weights v^k is straightforwardly given by

$$\frac{dv^k}{dt} = \eta \left[\sum_n \tilde{v}^n I_2(k, n) - \sum_j v^j I_2(k, j) \right], \quad (39)$$

where we introduce the final shorthand $I_2(k, j) \equiv \mathbb{E}[g(\lambda^k)g(\lambda^j)]$.

3. Solving the equations of motion

The equations of motion are valid for any choice of $f(x)$, $g(x)$, and $\tilde{g}(x)$. To solve the equations for a particular setup, one needs to compute the three constants a , b , and c [Eq. (19)] and the averages I_3 and I_4 [Eqs. (31) and (36)]. Choosing $g(x) = \tilde{g}(x) = \text{erf}(x/\sqrt{2})$, they can be computed analytically [21]. Finally, one needs to determine the spectral density of the matrix Ω_{rs} . When drawing the

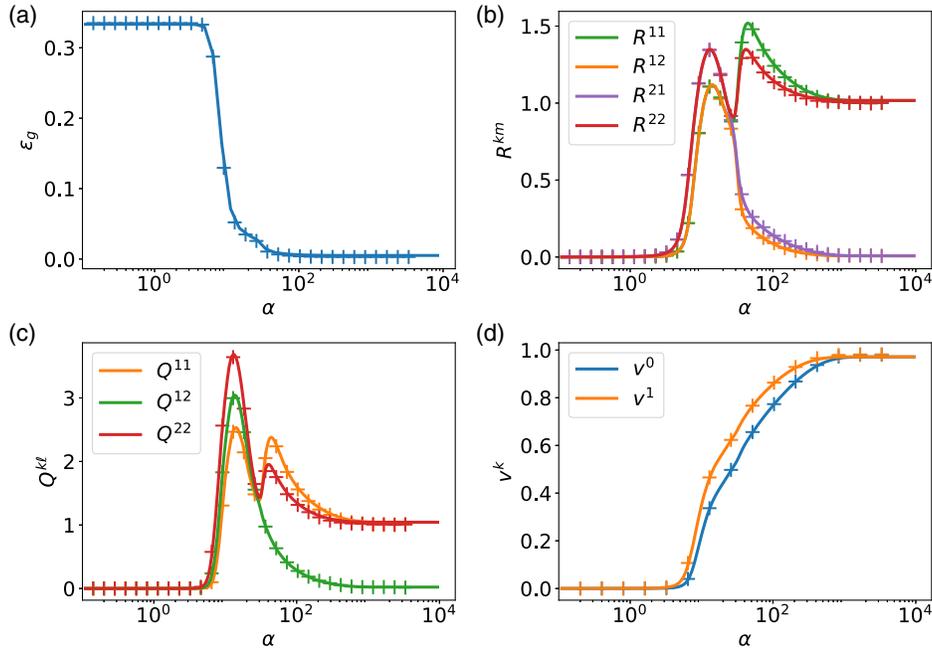


FIG. 3. The analytical description of the hidden manifold generalization dynamics matches experiments even at moderate system size. We plot the time evolution of the generalization error $\epsilon_g(\alpha)$ (a) and the order parameters R^{km} (b), $Q^{k\ell}$ (c) and second-layer weights v^k (d) obtained by integration of the ODEs (solid lines) and from a single run of SGD (crosses). Parameters: $g(x) = \text{erf}(x/\sqrt{2})$, $N = 10\,000$, $D = 100$, $M = 2$, $K = 2$, $\eta = 0.2$, and $\tilde{v}^m = 1$.

entries of the feature matrix F_{ir} IID from some probability distribution with a finite second moment, the limiting distribution of the eigenvalues $p_{\Omega}(\rho)$ in the integral (28) and (34) is the Marchenko-Pastur distribution [66]:

$$p_{\text{MP}}(\rho) = \frac{1}{2\pi\delta} \frac{\sqrt{(\rho_{\max} - \rho)(\rho - \rho_{\min})}}{\rho}, \quad (40)$$

where $\rho_{\min} = (1 - \sqrt{\delta})^2$ and $\rho_{\max} = (1 + \sqrt{\delta})^2$, where we recall that $\delta \equiv D/N$. Note that our theory crucially also applies to nonrandom matrices; we visit such an example in Sec. IV B, where we also discuss the importance of this use case. A complete numerical implementation of the equations of motion is available on GitHub [50].

We illustrate the content of the equations of motion in Fig. 3, where we plot the dynamics of the generalization error, the order parameters R^{km} , $\Sigma^{k\ell}$, and $W^{k\ell}$, and the second-layer weights v^k obtained from a single experiment with $N = 10000$, $D = 100$, and $M = K = 2$, starting from small initial weights (crosses). The elements of the feature matrix are drawn IID from the standard normal distribution, as are the elements of the latent representations c . The solid lines give the dynamics of these order parameters obtained by integrating the equations of motion. The initial conditions for the integration of the ODEs are taken from the simulation. The ODE description matches this single experiment really well even at moderate system sizes. For Fig. 3, our choice of N and D results in $\delta = 0.01$, and we check that the ODEs and simulations agree for various values of δ ; cf. Fig. 5.

4. Discussion

Yoshida and Okada [36] recently analyzed online learning for two-layer neural networks (1) trained on Gaussian inputs, with a two-layer teacher acting directly on the inputs x . Their approach consists of introducing distinct order parameters $R_{(i)}^{km}$, $Q_{(i)}^{kl}$, etc., for each distinct eigenvalue of the input covariance matrix Ω . They analyze their equations for covariance matrices with one and two distinct eigenvalues. Here, we first introduced the GEP (III.1) to show that inputs which are not normally distributed, such as $X = f(CF/\sqrt{D})$, can be reduced to an effective Gaussian model as far as the dynamics of learning are concerned. Furthermore, the description of the learning dynamics we just discussed allows us to analyze inputs with any well-defined spectral density with just a single set of order parameters Q^{kl} , R^{km} , and T^{nm} . This analysis is made possible by introducing the integral over the order parameter densities $r^{km}(\rho)$, etc. As we see below, this integral can actually be solved for small δ , which simplifies the equations of motion considerably and allows for a detailed analysis (cf. Sec. IV C).

We lastly comment on the role of the dimensionality in our setup. Inspection of the test error (25) reveals that a

student has to recover the local fields of the teacher ν^m in order to perform well (if she has the same activation function as the teacher). If the student is trained directly on the latent variables C , she could recover these local fields perfectly and we would be back in the setup of Saad and Solla [22]. In the HMM, the student is given only the high-dimensional inputs X , which can be seen as a noisy projection of the latent variables C (23). The high dimensionality of the student inputs is, thus, a constraint that must be overcome to learn well, because projection to high dimensions is part of the *data-generating* process. This process is to be contrasted with setups like random features [39,40] or certain neural circuits in sensory processing [67,68], where projection of the inputs to higher-dimensional spaces is part of the *analysis* and generally simplifies the subsequent learning problem.

IV. ANALYTICAL RESULTS

The goal of this section is to use the analytic description of online learning to analyze the dynamics and the performance of two-layer neural networks in detail.

A. Specialization of student nodes in the HMM

An intriguing feature of both the canonical teacher-student setup and the hidden manifold model is that they both exhibit a *specialization* phenomenon. Upon closer inspection of the time evolution of the order parameter R^{km} in Fig. 3(b), we see that, during the initial decay of the generalization error up to a time $t = P/N \sim 10$, all elements of the matrix R^{km} are comparable. In other words, the correlations between the preactivation λ^k of any student node and the preactivation ν^m of any teacher node is roughly the same. As training continues, the student nodes “specialize”: The preactivation of one student node becomes strongly correlated with the preactivation of only a single teacher node. In the example shown in Fig. 3, we have strong correlations between the preactivation of the first student and the first teacher node (R^{11}) and similarly between the second student and second teacher node (R^{22}). The specialization of the teacher-student correlations is concurrent to a decorrelation of the student units, as can be seen from the decay of the off-diagonal elements of the latent and ambient student-student overlaps $\Sigma^{k\ell}$ and $W^{k\ell}$, respectively (bottom of Fig. 3). Similar specialization transitions are observed in the canonical teacher-student setup for both online and batch learning [22,69]; see Engel and Van den Broeck [8] for a review.

B. Using nonrandom feature matrices

Our first example of the learning dynamics in Sec. IV A is for a feature matrix F whose entries are taken IID from the normal distribution. The derivation of the ODEs for online learning, however, does not require that the feature matrix F be random; instead, it requires only the balance

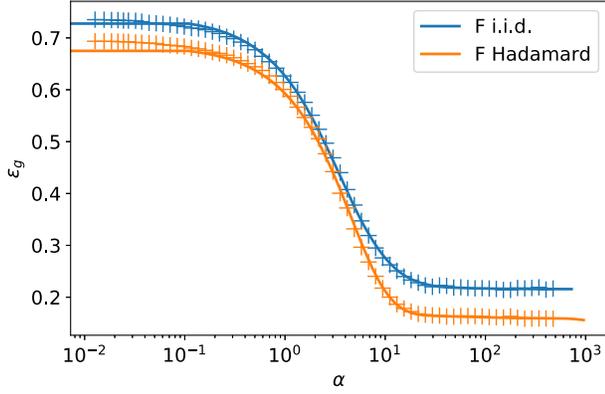


FIG. 4. The ODE analysis is asymptotically correct for non-random feature matrices F . We plot the time evolution of the generalization error ϵ_g obtained by integration of the ODEs (solid lines) and from a single run of SGD (2) (crosses) for two different matrices F : (i) Elements F_{ir} are drawn IID from the standard normal distribution (blue); (ii) F is a Hadamard matrix [70]. $f(x) = \text{sgn}(x)$, $g(x) = \tilde{g}(x) = \text{erf}(x/\sqrt{2})$, $N = 1023$, $D = 1023$, $M = 2$, $K = 2$, $\eta = 0.2$, and $\tilde{v}^m = 1$.

condition stated in Eq. (14) as well as the normalization conditions (13). To illustrate this point, we plot examples of online learning dynamics with $M = K = 2$ in Fig. 4, with the prediction from the ODE as solid lines and the result of a single simulation with crosses. In blue, we show results where the elements of F_{ir} are drawn IID from the standard normal distribution. For the experiment in orange, $F = H_N$, where H_N is a Hadamard matrix [70]. Hadamard matrices are $N \times N$ matrices—hence, $\delta = 1$ —and are popular in error-correcting codes such as the Reed-Muller code [71,72]. They can be defined via the relation

$$H_N H_N^\top = N \mathbb{I}_N, \quad (41)$$

where \mathbb{I}_N is the $N \times N$ identity matrix. As we can see from Fig. 4, the ODEs capture the generalization dynamics of the Hadamard case just as well.

C. The limit of small latent dimension

The key technical challenge in analyzing the analytical description of the dynamics is handling the integro-differential nature of the equations. We can simplify the equations in the limit of small $\delta \equiv D/N$. Numerical integration of the equations reveals that, at convergence, the continuous order parameter densities $r^{km}(\rho)$ and $\sigma^{k\ell}(\rho)$ are approximately constant:

$$r^{km}(\rho) = r^{km}; \quad \sigma^{k\ell}(\rho) = \sigma^{k\ell}. \quad (42)$$

This observation is key, because making the ansatz (42) allows us to transform the integro-differential equations for the dynamics of $r^{km}(\rho, t)$ [Eq. (29)] and $\sigma^{k\ell}(\rho, t)$ [Eq. (35)] into first-order ODEs, provided we can perform the integral

over the eigenvalue distribution $p_\Omega(\rho)$ in Eqs. (28) and (34) analytically. This situation is, for example, the case if we take the elements of the feature matrix F IID from any probability distribution with a bounded second moment, in which case $p_\Omega(\rho)$ is given by the Marchenko-Pastur distribution (40). We focus on this case for the remainder of this section.

Let us note that the regime of small δ is also the relevant regime for image datasets such as MNIST and CIFAR10, whose δ has been estimated previously to be around $\delta_{\text{MNIST}} \sim 14/784$ and $\delta_{\text{CIFAR10}} \sim 35/3072$, respectively [12–15]; cf. our discussion in the introduction.

1. The effect of the latent dimension $D = \delta N$

As a first application of this approach, we analyze the dependence of the asymptotic test error ϵ_g^* on the latent dimension D of the hidden manifold when the teacher and student have the same number of hidden nodes, $K = M$.

From inspection of the form of the order parameters after integrating the full set of ODEs until convergence, we make the following ansatz for the overlap matrices:

$$\Sigma^{k\ell} = \begin{cases} S & k = \ell, \\ s & \text{otherwise,} \end{cases} \quad W^{k\ell} = \begin{cases} W & k = \ell, \\ w & \text{otherwise,} \end{cases} \quad (43)$$

$$T^{nm} = \begin{cases} T & n = m, \\ t & \text{otherwise,} \end{cases} \quad \tilde{T}^{nm} = \begin{cases} \tilde{T} & n = m, \\ \tilde{t} & \text{otherwise,} \end{cases} \quad (44)$$

$$R^{km} = \begin{cases} R & k = m, \\ r & \text{otherwise,} \end{cases} \quad v^k = v; \quad A^m = A. \quad (45)$$

Substituting this ansatz into the ODEs allows us to derive closed-form expressions for ODEs governing the dynamics of seven order parameters R , r , S , s , W , w , and v that are valid for small δ and for any $K = M$. The teacher-related order parameters T , t , \tilde{T} , and \tilde{t} describe the teacher and are constants of the motion. They have to be chosen to reflect the distribution from which the weights of the teacher network are drawn in an experiment. The full equations of motion are rather long, so instead of printing them here in full we provide a *Mathematica* notebook for reference [50].

The key idea of our analytical approach is to look for fixed points of this ODE system and to substitute the values of the order parameters at those fixed points into the expression for the generalization error (26). To understand the structure of the fixed points of the ODEs, we run a numerical fixed point search of the ODEs from 1000 initial values for the order parameters drawn randomly from the uniform distribution. We find two types of solution. First, there exist solutions of the form $R = r$, $S = s$, and $W = w$. This solution is a saddle point of the equations and is, thus, not a stable fixed point of the dynamics. Instead, it corresponds to a well-known “unspecialized” phase, when networks with $K > 1$ hidden nodes have not yet specialized

and, hence, achieve only the performance of a network with $K = 1$ hidden unit (cf. our discussion in Sec. IV A). The learning dynamics approaches this saddle point at an intermediate stage of learning but finally drifts away from it toward a “specialized” solution. This second solution corresponds to the asymptotic fixed point of the learning dynamics where the student has specialized; i.e., we have R large and r small, etc. Substituting the values of the order parameters of this solution into Eq. (26) yields the asymptotic generalization error of a student.

Making this argument rigorous requires a proof of global convergence of the coupled, nonlinear integro-differential equations of motion [(29), (35), (38), (39)] from random initial conditions. This challenging mathematical problem remains open, despite some recent progress for two-layer neural networks with finite N and a large hidden layer [73–77]. Thus, all predictions in this way ultimately need to be compared to simulations to verify their accuracy.

We show the results of this analysis in Fig. 5. The crosses are experimental results for which we train networks with $M = K = 2$ on data from a hidden manifold with latent dimension $D = 25, 50, 100$, and 200 , choosing the input dimension N to obtain the range of δ desired for each curve. We plot the asymptotic error averaged over five runs with dots; error bars indicate two standard deviations. The lowest solid line in Fig. 5 is the theoretical prediction obtained by the procedure just explained when assuming that $T = 1$, $t = 0$, $\tilde{T} = 1$, and $\tilde{t} = 0$.

While the experimental results are approaching the theoretical line as the latent dimension D increases, there

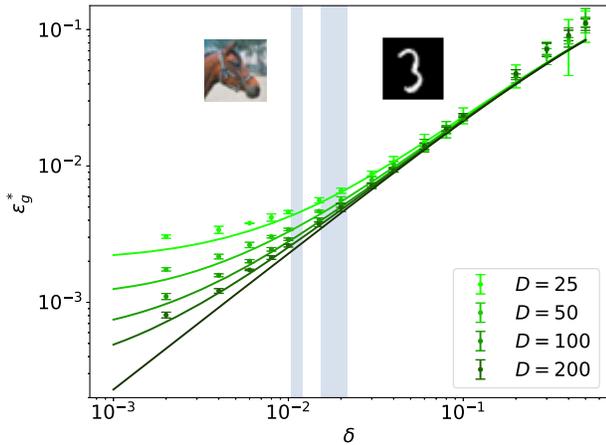


FIG. 5. The impact of the latent dimension $\delta \equiv D/N$. We plot the final test error ϵ_g^* of sigmoidal students trained on the hidden manifold model with three different intrinsic dimensions D as a function of $\delta = D/N$, where N is the input dimension. The average is taken over five runs. The solid lines are the asymptomatic theoretical predictions derived in Sec. IV C 1. The shaded bars indicate experimental estimates [12–15] for δ for the CIFAR10 dataset (left) and the MNIST dataset (right). $f(x) = \text{sgn}(x)$, $g(x) = \tilde{g}(x) = \text{erf}(x/\sqrt{2})$, $M = K = 2$, $\eta = 0.2$, and $\tilde{v}^m = 1$.

are qualitative differences in the shape of the δ dependence for small δ . These differences arise due to the following finite-size effect. While it is numerically easy to enforce $T = 1$, $t = 0$ by orthogonalizing the teacher weight matrix, it is not possible to explicitly control the reweighted teacher-teacher overlap \tilde{T}^{nm} [Eq. (30)]. The deviation of \tilde{T}^{nm} from the identity leads to the deviations we see at small δ . We demonstrate this result in Fig. 5 by also plotting theoretical predictions for $\tilde{T} = 1 - x$ and $\tilde{t} = x$ and choosing $x = 1/D$. These curves match the experiments much better. Plotting the data with a linear y scale (not shown) reveals that the solution obtained making the small- δ ansatz (42) is valid until $\delta \sim 0.2$.

2. Learning rate η

We find that the asymptotic test error ϵ_g^* depends only weakly on the learning rate η , as we show in Fig. 6 for $M = K = 2$ and $M = K = 6$, together with the theoretical prediction for $\tilde{t} = 0$. This theoretical prediction is again obtained by using the ansatz (43) for the order parameters and solving the resulting fixed point equations, as described in the previous section, but this time varying the learning rate η . The weak dependence of ϵ_g on η should be contrasted with the behavior of the canonical teacher-student setup, where the generalization error is proportional to the learning rate in the case of additive Gaussian output noise [55,78].

In the inset in Fig. 6, we plot the generalization dynamics of a neural network trained on the HMM at different learning rates. As expected, the learning rate controls the speed of learning, with increased learning rates leading to

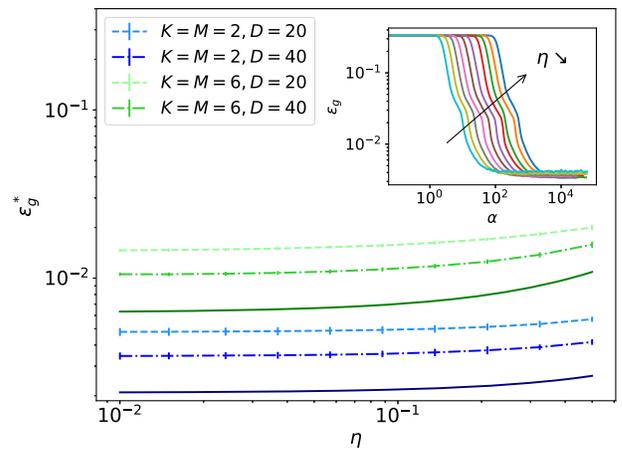


FIG. 6. The impact of the learning rate η . We plot the final test error ϵ_g^* of sigmoidal students trained on the hidden manifold model for a range of learning rates η for sigmoidal networks with $K = M = 2$ (blue lines) and $K = M = 6$ (green lines). We repeat the experiments for two values of D , choosing N such that $\delta = D/N = 0.01$. Inset: Generalization dynamics during training ($K = M = 2$). Parameters: $f(x) = \text{sgn}(x)$, $g(x) = \tilde{g}(x) = \text{erf}(x/\sqrt{2})$, $\delta = 0.01$, $\tilde{v}^m = 1$, and $K = M$.

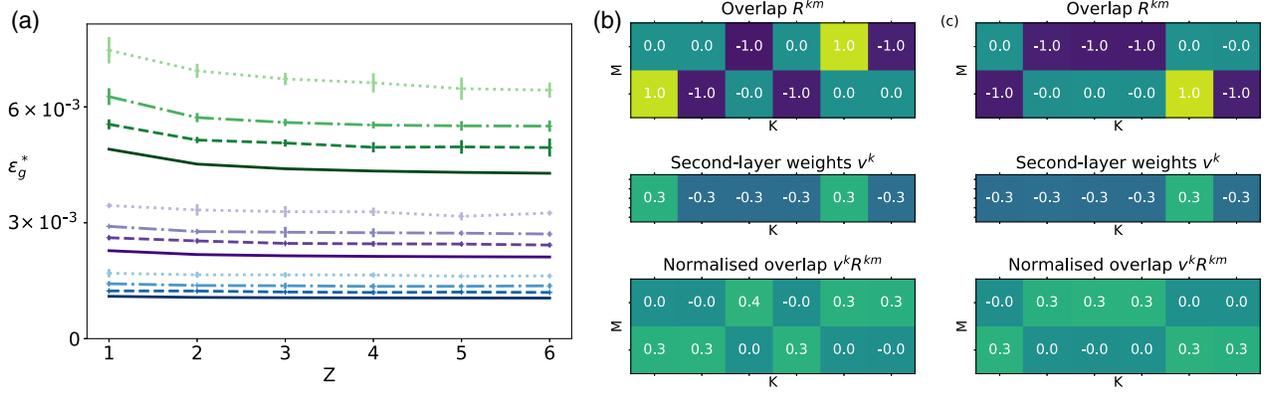


FIG. 7. (a) Asymptotic generalization for online learning of a student with $K = ZM$ hidden nodes learning from a teacher with $M = 1$ (blue lines), $M = 2$ (violet lines), and $M = 4$ (green lines) hidden nodes, respectively. The dotted, dash-dotted, and dashed lines correspond to $D = 50$, 100, and 200, respectively. Error bars indicate two standard deviations over five runs. The solid line is the theoretical prediction obtained for $\bar{t} = 0$. (b),(c) Teacher-student overlap R^{km} [Eq. (17)], second-layer weights v^k , and the normalized overlap $v^k R^{km}$ obtained in two simulations used in the left plot with $M = 2$ and $K = 6$, starting from different initial conditions, all other things being equal. Parameters: In all plots, $f(x) = \text{sgn}(x)$, $g(x) = \text{erf}(x/\sqrt{2})$, $\eta = 0.2$, $\tilde{v}^m = 1$, $\delta = 0.01$, and $\eta = 0.2$.

faster learning until the learning rate becomes so large that learning is not possible anymore; instead, the weights just grow to infinity.

3. The impact of student size

Another key question in our model is how the performance of the student depends on her number of nodes K . Adding hidden units to a student who has less hidden units than her teacher ($K < M$) improves her performance, as would be expected. This result can be understood in terms of the specialization discussed in Sec. IV A: Each additional hidden node of the student specializes to another node of the teacher, leading to improved performance. We see an example of this improvement below in Sec. V A.

But what happens if we give the student more nodes than her teacher has, $K > M$? It is instructive to first study the overlap matrices at the end of training. We show two examples from an experiment with $M = 2$ and $K = 6$ at $\delta = 0.01$ for networks starting from different initial conditions. In particular, we plot the rescaled teacher-student overlap matrix $v^k R^{km}$ in Figs. 7(b) and 7(c). We rescale R^{km} by the second-layer weights to account for two effects: first, the relative influence of a given node to the output of the student, which is determined by the magnitude of the corresponding second-layer weight; and second, we have a symmetry in the output of the student, since for the sigmoidal activation function $v^k g(w^k x/\sqrt{N}) = -v^k g(-w^k x/\sqrt{N})$.

In the two overlap plots for $K > M$ in Fig. 7, the student nodes display *many-to-one* specialization: Several hidden units of the student specialize to the same hidden node of the teacher, essentially providing several estimates of the value of this teacher node. Note that each student node specializes to one and only one of the teacher nodes rather than a combination of two or more teacher nodes. We find

this pattern of activations consistently across all of our runs for various K and M . The fact that student nodes are evenly distributed across teacher nodes is further motivated by the fact that such an arrangement minimizes the generalization error if the second-layer teacher weights \tilde{v}^m have equal magnitude and its first-layer weights \tilde{w}^m have the same norm. We anticipate that this specialization pattern is at least in part due to the sigmoidal form of the activation function $g(x)$. We note that the same many-to-one specialization of hidden units has been previously reported for the same two-layer networks trained on IID inputs [55] and that a similar pattern of specialization is observed for networks with finite input and a wide hidden layer, where this type of specialization is referred to as “distributional dynamics” [74–77].

These observations motivate the following ansatz for the overlaps of a student with $K = ZM$ hidden nodes ($Z \in \mathbb{N}$):

$$R^{km} = \begin{cases} R & k \bmod M = m \bmod M, \\ r & \text{otherwise,} \end{cases} \quad (46)$$

$$\Sigma^{k\ell} = \begin{cases} S & k \bmod M = \ell \bmod M, \\ s & \text{otherwise} \end{cases} \quad (47)$$

and similarly for $W^{k\ell}$, while we use the same parameterization for the teacher order parameters T , t , \tilde{T} , \tilde{t} , A , and v . Searching again for specialized fixed points of the resulting equations for the seven time-dependent order parameters R , r , S , s , W , w , and v and substituting their values into Eq. (26) yields the predictions we indicate by solid lines in Fig. 7, where we plot the asymptotic test error as a function of $Z \equiv K/M$. We can see small performance improvements as the student size increases. We also plot, for the three values of M used, the asymptotic test error measured in experiments with $D = 50$, 100, and 200. As we increase D ,

the experimental results approach the theoretical prediction for $D \rightarrow \infty$.

We finally note that fixed points of the online dynamics with many-to-one specialization have been described previously in the canonical teacher-student setup [55], which found that this behavior leads to a more significant improvement of student performance as K increases for teacher tasks with $y^* = \phi(x)$ compared to the improvement we observe for the HMM. The same type of many-to-one specialization is also found by Mei, Montanari, and Nguyen [74] and Chizat and Bach [76], who consider a complementary regime where the input dimension N stays finite while the size of the hidden layer goes to infinity.

V. COMPARING THE HIDDEN MANIFOLD MODEL TO REAL DATA

We finally turn our attention to the comparison of the hidden manifold model to more realistic datasets, in our cases classic image databases such as CIFAR10 (see Fig. 1 for two examples of images in CIFAR10).

A. Neural networks learn functions of increasing complexity

The specialization transition that we discuss in Sec. IV A has an important consequence for the performance of the neural network, as we show in Fig. 8. As we train increasingly large student networks on a teacher with $M = 10$ hidden units and second-layer weights $\tilde{v}^m = 1/M$, we observe that learning proceeds in two phases. First, there is an initial decay of the generalization error until all students have roughly the same test error as the student with a single hidden unit $K = 1$. In a second phase, students with $K > 1$ break away from this plateau after further training and achieve superior performance, with the larger networks

performing better. These improvements are a result of specialization after approximately 10^3 epochs, which permits the student network to capitalize on their additional hidden nodes.

This way of visualizing specialization not only illustrates its importance for student performance, it is also applicable when training the same two-layer neural networks on more realistic datasets such as MNIST [Fig. 8(b)] or fashion-MNIST [24] and CIFAR (Fig. 10). The plots demonstrate clearly that, in all these cases, the larger networks proceed by first learning functions that are equivalent to the smaller networks.

In all cases, specialization is preceded by a plateau where the generalization error stays constant, because the student is stuck at a saddle point in its optimization landscape, corresponding to the unspecialized solution. This plateau has been discussed extensively in the canonical teacher-student setup [8,23,79,80] and, more recently, in the context of recurrent and deep neural networks [37,81]. By comparing students of different sizes, this plateau can also be demonstrated on image datasets, as we do above. This learning of functions with increasing complexity has also been observed in deep convolutional networks by Kalimeris *et al.* [82], who use quantities from information theory to quantify how well one model explains the performance of another.

These observations are interesting, because they suggest how to explain the ability of neural networks to generalize well from examples when they have *many more* parameters than samples in their training dataset. This explanation is a key open problem in the theory of deep learning, since the intuition from classical statistics suggests that, in these cases, the networks overfit the training data and, thus, generalize poorly [5,83]. It is possible that, by learning functions of increasing complexity, networks are biased

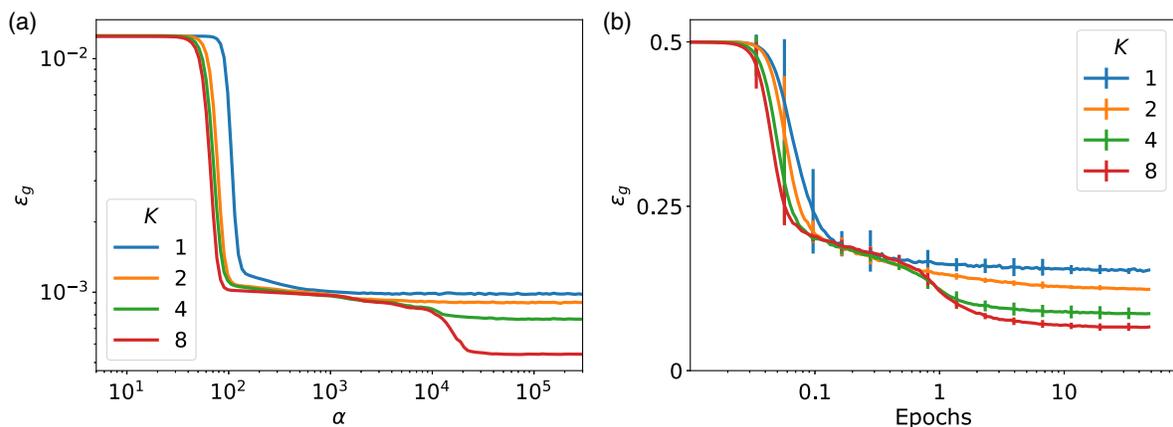


FIG. 8. Two-layer neural networks learn functions of increasing complexity. We plot the generalization error of sigmoidal two-layer networks with an increasing number of hidden nodes K during a single run of online learning with the hidden manifold model with $\delta = 0.05$, $D = 25$, $M = 10$, and $\tilde{v}^m = 1/M$ on the HMM (a) and when trained on odd-versus-even digit classification on MNIST, averaged over ten runs (b). Error bars indicate two standard deviations. For details, see Secs. V A and D. $g(x) = \text{erf}(x/\sqrt{2})$, $\eta = 0.2$, and $N = 784$. (b) Batch size 32.

toward simple classifiers and avoid overfitting if their training is stopped before convergence. This topic is an active research area [84,85].

B. Memorization of random and realistic data

An interesting difference between random and realistic data is demonstrated in a recent paper by Arpit *et al.* [86]. They train 100 two-layer networks ($K = 4096$ hidden units with ReLU activation, ten output units with softmax activation) for a single epoch on the ten-class image classification task on the CIFAR10 dataset, starting from different initial conditions each time. At the end of training, they measure the frequency with which each individual image is classified correctly by the network across runs, which we call the memorability of an image, which should be thought of as a function of the image and the dataset that contains it. We repeat this experiment on CIFAR10 and add three different synthetic datasets (color codes refer to Fig. 9):

- (i) *CIFAR10 (blue)*.— x_μ : CIFAR10 images; $y_\mu^* \in [0, 9]$: CIFAR10 label giving the class of that image.
 - (ii) *Gaussian (orange)*.—Teacher acting on Gaussian inputs. x_μ : IID standard Gaussians; $y_\mu^* = \operatorname{argmax}\phi(x_\mu, \theta^*)$.
 - (iii) *TeacherS (green)*.—Teacher acting on structured inputs. $x_\mu = f(Fc_\mu)$; $y_\mu^* = \operatorname{argmax}\phi(x_\mu, \theta^*)$.
 - (iv) *HMM (red)*.— $x_\mu = f(Fc_\mu)$; $y_\mu^* = \operatorname{argmax}\phi(c_\mu, \theta^*)$.
- The labels for the synthetic datasets are generated by two teacher networks, one with input dimension N for the Gaussian and TeacherS datasets and another with input dimension D for the HMM. The teachers are two-layer

fully connected networks having $M = 2K$ hidden units with ReLU activation function and ten nodes in the last readout layer. Thus, the teacher's output $\phi(\cdot, \theta^*) \in \mathbb{R}^{10}$, and the class for a given input is obtained as the index of the output node with the highest value for that input.

We plot the memorabilities for all images in the training set, sorted by their memorability, in Fig. 9. On the left, we first reproduce the memorability curve for CIFAR10 that was found by Arpit *et al.* [86] (solid blue curve), which demonstrates that many examples are *consistently* classified correctly or incorrectly after a single epoch of training. The memorability curve for a dataset containing the same images with *random* labels (dashed blue curve) demonstrates that randomized CIFAR10 does not contain images that are particularly hard or easy to memorize. The smaller variation in memorability for the randomized dataset is largely due to the fact that it takes it more time to fit randomized datasets [87]. After one epoch, the network thus has a lower training accuracy on the randomized dataset (cf. the inset in Fig. 9), which leads to the smaller area underneath the curve. We verify that no easy or hard samples appear when training the randomized datasets to comparable training accuracy (not shown). In fact, the memorability of datasets with random labels seems to coincide after accounting for differences in the training error, regardless of whether the inputs are CIFAR10 images, Gaussian inputs, or structured inputs $X = f(CF)$ (4) [dashed lines in Fig. 9(a)].

The memorability curves for the Gaussian, TeacherS, and HMM datasets in Fig. 9(b) reveal that hard and easy examples exist for TeacherS and HMM, which both contain structured inputs $X = f(CF)$, but not in the Gaussian

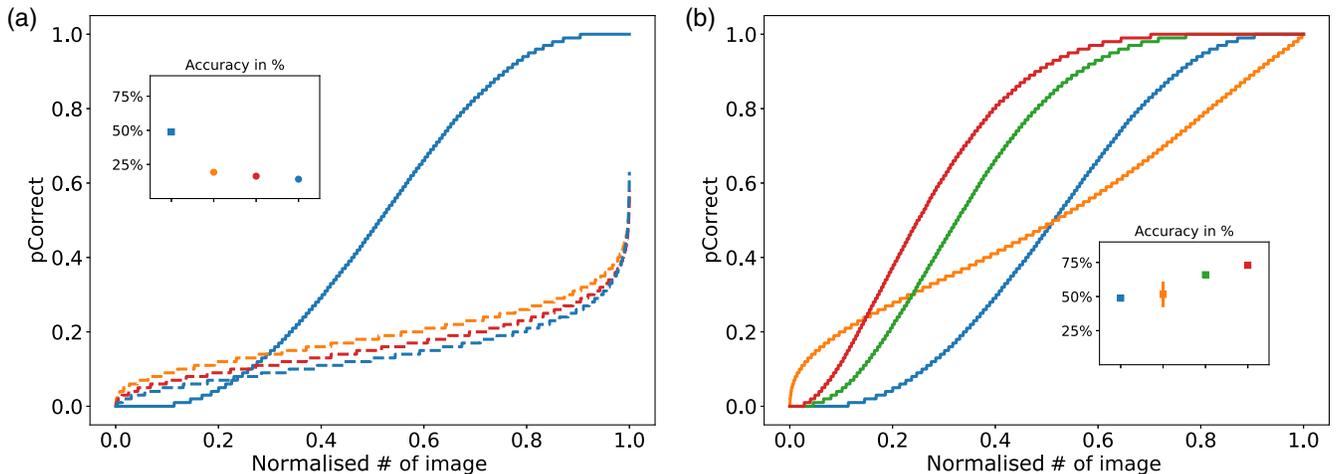


FIG. 9. Neural networks have different memorization patterns for random and structured datasets. We plot the memorability of training images, i.e., the frequency with which an image from the training set is correctly classified by a neural network after training for only a single epoch. In (a) and (b), we reproduce the result of Arpit *et al.* [86] for CIFAR10 (full blue line). This curve demonstrates the existence of hard and easy examples which are never, or always, classified correctly. (a) shows that this property disappears in all models when the labels are reshuffled (dashed lines). The insets indicate the training accuracy after training, using circles for randomized datasets and squares for unmodified data sets. (b) shows that these hard and easy examples also exist in the structured data models, TeacherS (green line) and the HMM (red line), but not in the unstructured Gaussian one (orange line).

dataset. The number of easy examples, but not their existence, correlates well with the training accuracy on these datasets, shown in the inset. In that sense, the hidden manifold model is thus a more realistic model of imagelike data than the canonical teacher-student setup.

Note that, by making the teacher network larger than the students ($M = 2K$), the learning problem is unrealizable for all three synthetic datasets; i.e., there is no set of weights for the student that achieve zero generalization error. The absence of easy examples in the Gaussian dataset, thus, suggests that unrealizability alone is insufficient to obtain a dataset with easy examples. Our results also demonstrate that memorability is not just a function of the input correlations: CIFAR10 images, Gaussian inputs, and structured inputs yield the same memorability curves when their labels are randomized. We leave it to future work to identify some criterion, statistical or otherwise, that predicts either whether a sample (x_μ, y_μ^*) is easy (or hard) to memorize or whether a training set contains easy examples at all.

VI. CONCLUDING PERSPECTIVES

We introduce the hidden manifold model as a generative model for structured datasets that displays some of the phenomena that we observe when training two-layer neural networks on realistic data. The HMM has two key ingredients, namely, high-dimensional inputs which lie on a lower-dimensional manifold and labels for these inputs that depend on the inputs' position within the low-dimensional manifold. We derive an analytical solution of the model for online SGD learning of two-layer neural networks. We, thus, provide a rich test bed for exploring the influence of data structure on learning in neural networks.

Let us close this paper by outlining several important directions in which our work is being (or should be) extended.

Comparison to more deep learning phenomenology.—In the spirit of our experiments in Sec. V, it is of great interest to identify more properties of learning that are consistently reproduced across experiments with realistic datasets and network architectures and to test whether the HMM reproduces these observations as well. Of particular interest will be those cases where learning on realistic data deviates from the HMM and how we can extend the HMM to capture these behaviors.

Beyond online SGD.—Our analytical results on online SGD rely on the assumption that each new sample seen during training is conditionally independent from the weights of the network up to that point. In practice, samples are seen several or even many times during training, giving rise to additional correlations. Taking those correlations into account to analyze those cases is an important future direction. First steps toward a solution to this challenging problem were made using the dynamical replica method [88,89] for two-layer networks and for single-layer neural networks trained using full-batch

gradient descent, where all the samples in the training set are used at every step of the algorithm [42,43,90]. Generalizing these results to two-layer networks is clearly a direction for future work as well.

Learning with a multilayer network.—The present work should be extended to learning with multilayer networks in order to identify how depth helps to deal with structured data. This challenge is serious, and it remains an open problem to find explicitly solvable models of multilayer (nonlinear) networks even in the canonical teacher-student model where inputs are uncorrelated.

Multilayer generative model.—The hidden manifold model is akin to a single-layer generator of a GAN. A natural extension would be to take a generator with an arbitrary number of layers. Multilayer generators are explored in Refs. [41,49], whose results are analogous to the Gaussian equivalence property and suggest that the full solution of the online SGD or of the full-batch gradient descent might also be within reach.

Conditioning the inputs on the labels.—In the HMM, the true label y^* of an input x is conditioned on its latent representation c , i.e., its coordinates in the manifold. It may be more realistic to consider models where, instead, the latent representation is conditioned on the label of the input, i.e., $p(c|y)$. A simple case of such a model that reduces to a Gaussian mixture of two clusters was explored recently [91]. This point of view is also taken implicitly in Ref. [33]. More generally, exploring different approaches to modeling realistic inputs will allow us to better understand how data structure influences learning.

ACKNOWLEDGMENTS

We thank Bruno Loureiro and Federica Gerace for useful discussions. We thank Stanisław Jastrzębski for discussing the experiments of Ref. [86]. We are grateful to the Kavli Institute for Theoretical Physics for its hospitality during an extended stay, during which parts of this work were conceived and carried out. We acknowledge funding from the European Research Council (ERC) under the European Union's Horizon 2020 Research and Innovation Program Grant Agreement No. 714608-SMiLe, from "Chaire de recherche sur les modèles et sciences des données," Fondation CFM pour la Recherche-ENS, and from the French National Research Agency (ANR) grant PAIL. This research was supported in part by the National Science Foundation under Grant No. NSF PHY-1748958.

APPENDIX A: THE GAUSSIAN EQUIVALENCE PROPERTY

1. Nonlinear functions of weakly correlated Gaussian random variables

In order to derive the GEP we first establish some auxiliary lemmas concerning the correlations between nonlinear functions of weakly correlated random variables.

a. Correlations of two functions

Lemma A1.—Given $n + p$ random variables organized in two vectors,

$$x = \begin{pmatrix} x^1 \\ \cdot \\ \cdot \\ x^n \end{pmatrix}, \quad y = \begin{pmatrix} y^1 \\ \cdot \\ \cdot \\ y^p \end{pmatrix}, \quad (\text{A1})$$

with a joint Gaussian distribution, denote by \mathbb{E} the expectation with respect to this distribution. The first moments are supposed to vanish:

$$\mathbb{E}x_i = 0, \quad \mathbb{E}y_j = 0, \quad (\text{A2})$$

and we denote by Q , R , and εS the covariances:

$$\mathbb{E}[x_i x_j] = Q_{ij}, \quad \mathbb{E}[y_i y_j] = R_{ij}, \quad \mathbb{E}[x_i y_j] = \varepsilon S_{ij}. \quad (\text{A3})$$

Let $f(x)$ and $g(y)$ be two functions of x and y , respectively, regular enough so that $\mathbb{E}_x[x_i f(x)]$, $\mathbb{E}_x[x_i x_j f(x)]$, $\mathbb{E}_y[y_i f(y)]$, and $\mathbb{E}_y[y_i y_j f(y)]$ exist, where \mathbb{E}_x denotes the expectation with respect to the distribution $\mathcal{N}(a, Q)$ of x and \mathbb{E}_y denotes the expectation with respect to the distribution $\mathcal{N}(b, R)$ of y .

Then, in the $\varepsilon \rightarrow 0$ limit,

$$\begin{aligned} \mathbb{E}[f(x)g(y)] &= \mathbb{E}_x[f(x)]\mathbb{E}_y[g(y)] \\ &+ \varepsilon \sum_{i=1}^n \sum_{j=1}^p \mathbb{E}_x[x_i f(x)](Q^{-1}SR^{-1})_{ij} \mathbb{E}_y[y_j g(y)] \\ &+ \mathcal{O}(\varepsilon^2). \end{aligned} \quad (\text{A4})$$

Proof.—The result is obtained by a straightforward expansion in ε .

The joint distribution of x and y is

$$P(x, y) = \frac{1}{Z} \exp \left[-\frac{1}{2} \begin{pmatrix} x & y \end{pmatrix} M^{-1} \begin{pmatrix} x \\ y \end{pmatrix} \right], \quad (\text{A5})$$

where

$$M = \begin{pmatrix} Q & \varepsilon S \\ \varepsilon S^T & R \end{pmatrix}. \quad (\text{A6})$$

One can expand the inverse matrix M^{-1} to first order in ε :

$$M^{-1} = \begin{pmatrix} Q^{-1} & 0 \\ 0 & R^{-1} \end{pmatrix} - \varepsilon \begin{pmatrix} 0 & Q^{-1}SR^{-1} \\ R^{-1}S^TQ^{-1} & 0 \end{pmatrix} \quad (\text{A7})$$

and substitute this result into the joint distribution (A5) to find

$$\begin{aligned} P(x, y) &= \frac{1}{Z} \exp \left[-\frac{1}{2} \begin{pmatrix} x & y \end{pmatrix} \begin{pmatrix} Q^{-1} & 0 \\ 0 & R^{-1} \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \right] \\ &\times \left[1 + \varepsilon \sum_{i=1}^n \sum_{j=1}^p x_i (Q^{-1}SR^{-1})_{ij} y_j + \mathcal{O}(\varepsilon^2) \right]. \end{aligned} \quad (\text{A8})$$

Using this expression, the result (A4) follows immediately. ■

An immediate application of the lemma to the case when $n = p = 1$ is the following. Consider two Gaussian random variables u_1 and u_2 with mean zero and covariance

$$\mathbb{E}[u_i^2] = 1; \quad \mathbb{E}[u_1^2] = 1; \quad \mathbb{E}[u_1 u_2] = \varepsilon m_{12}, \quad (\text{A9})$$

and two functions f_1 and f_2 . Define, for $i \in \{1, 2\}$,

$$a_i = \langle f_i(u) \rangle; \quad b_i = \langle u f_i(u) \rangle, \quad (\text{A10})$$

where $\langle \cdot \rangle$ denotes the average over the distribution of the random Gaussian variable u distributed as $\mathcal{N}(0, 1)$.

Then, in the $\varepsilon \rightarrow 0$ limit, the correlation between $f(u_1)$ and $g(u_2)$ is given by

$$\mathbb{E}[f_1(u_1)f_2(u_2)] = a_1 a_2 + \varepsilon m_{12} b_1 b_2 + \mathcal{O}(\varepsilon^2). \quad (\text{A11})$$

This result means that, if we consider centered functions $\tilde{f}_i(u_i) = f_i(u_i) - a_i$, their covariance is

$$\mathbb{E}[\tilde{f}_1(u_1)\tilde{f}_2(u_2)] = +\varepsilon m_{12} b_1 b_2 + \mathcal{O}(\varepsilon^2). \quad (\text{A12})$$

This result generalizes to correlation functions of higher order, as stated in the following lemma.

b. Higher-order correlations

Lemma A2.—Consider m Gaussian random variables u_1, \dots, u_m with mean zero and covariance

$$\forall i: \mathbb{E}[u_i^2] = 1, \quad \forall i \neq j: \mathbb{E}[u_i u_j] = \varepsilon m_{ij}, \quad (\text{A13})$$

and m functions f_1, \dots, f_m . Define as before:

$$a_i = \langle f_i(u) \rangle, \quad b_i = \langle u f_i(u) \rangle, \quad i \in \{1, \dots, m\}, \quad (\text{A14})$$

and define the centered functions as

$$\tilde{f}_i(u) = f_i(u) - a_i; \quad (\text{A15})$$

then,

$$\begin{aligned} & \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon^{m/2}} \mathbb{E} \tilde{f}_1(u_1) \dots \tilde{f}_m(u_m) \\ &= \begin{cases} \sum_{\sigma \in \Pi} m_{\sigma_1 \sigma_2} m_{\sigma_{p-1}} m_{\sigma_p} & \text{if } p \text{ is even,} \\ 0 & \text{if } p \text{ is odd,} \end{cases} \end{aligned} \quad (\text{A16})$$

where Π denotes all the $m!/(2^{m/2}(m/2)!)$ partitions of $\{1, \dots, m\}$ into $m/2$ disjoint pairs. This result means that, for the moments involving only different indices, the random variables $\tilde{f}_1(u_1)/\sqrt{\varepsilon}, \dots, \tilde{f}_m(u_m)/\sqrt{\varepsilon}$ behave, in the $\varepsilon \rightarrow 0$ limit, like Gaussian variables with a covariance matrix $b_i b_j m_{ij}$.

Proof.—The covariance matrix U of the variables u_1, \dots, u_m has elements 1 on the diagonal and elements of the order of ε out of the diagonal: $U = \mathbb{I} + \varepsilon m$. One can expand U^{-1} in powers of ε :

$$U^{-1} = \sum_{p=0}^{\infty} (-\varepsilon)^p m^p. \quad (\text{A17})$$

The integration measure over the variables u_1, \dots, u_m can be expanded as

$$\sqrt{(2\pi)^m \det M} e^{-\frac{1}{2} \sum_i u_i^2} \prod_{p=1}^{\infty} G_p(u_1, \dots, u_m), \quad (\text{A18})$$

where

$$\begin{aligned} G_p(u_1, \dots, u_m) &= 1 + \left(-\frac{\varepsilon}{2}\right)^p \sum_{ij} (m^p)_{ij} u_i u_j \\ &+ \frac{1}{2!} \left(-\frac{\varepsilon}{2}\right)^{2p} \sum_{ijk\ell} (m^p)_{ij} (m^p)_{k\ell} u_i u_j u_k u_\ell \\ &+ \dots \end{aligned} \quad (\text{A19})$$

When we compute the integral of $\tilde{f}_1(u_1) \dots \tilde{f}_m(u_m)$ with the measure (A18), because of the fact that $\langle \tilde{f}_i(u_i) \rangle = 0$, we need to include terms coming from $\prod_p G_p(u_1, \dots, u_p)$ that involve at least one power of each of the variables u_1, \dots, u_m .

When m is even, say, $m = 2r$, for $\varepsilon \rightarrow 0$, the term of this kind with the smallest power of ε is the monomial $u_1 \dots u_{2r}$ that comes from the r th order term in G_1 . This term gives

$$\begin{aligned} \mathbb{E} f_1(u_1) \dots f_{2r}(u_{2r}) &= \frac{1}{r!} \left(\frac{\varepsilon}{2}\right)^r \sum_{i_1 j_1 \dots i_r j_r} m_{i_1 j_1} m_{i_r j_r} \\ &+ \mathcal{O}(\varepsilon^{r+1}), \end{aligned} \quad (\text{A20})$$

where the sum $\sum_{i_1 j_1 \dots i_r j_r}$ runs over all permutations of the indices $1, \dots, 2r$. This result leads to Eq. (A16) for m even.

When m is odd, $m = 2r + 1$, for $\varepsilon \rightarrow 0$, the leading terms coming from $\prod G_p$ that give a nonzero result are

monomials of the type $u_1^1 u_2 \dots u_{2r+1}$. They are of the order of $\mathcal{O}(\varepsilon^{r+1})$. This result proves Eq. (A16) for m odd. ■

Corollary A3.—In the special case $m = 3$, we get

$$\begin{aligned} & \mathbb{E}[f_1(u_1) f_2(u_2) f_3(u_3)] \\ &= a_1 a_2 a_3 + \varepsilon (a_1 m_{23} b_2 b_3 + a_2 m_{13} b_1 b_3 + a_3 m_{12} b_1 b_2). \end{aligned} \quad (\text{A21})$$

2. Derivation of the Gaussian equivalence property

The derivation is based on the computation of moments of the variables λ^k and ν^m , showing that, in the thermodynamic limit, all the moments are those of Gaussian random variables. Here, we make explicit the derivation up to fourth-order moments and leave the daunting higher-order moments for a future formal proof of the GEP.

a. Covariances

We first compute the covariance matrix $G^{k\ell} = \mathbb{E}[\tilde{\lambda}^k \tilde{\lambda}^\ell]$:

$$G^{k\ell} = \frac{1}{N} \sum_{i,j} w_i^k w_j^\ell \mathbb{E}[f(u_i) - a][f(u_j) - a] \quad (\text{A22})$$

$$= (c - a^2) W^{k\ell} + \frac{1}{N} \sum_{i \neq j} w_i^k w_j^\ell \mathbb{E}[f(u_i) - a][f(u_j) - a]. \quad (\text{A23})$$

In the last piece, we need to compute $\mathbb{E}\{[f(u_i) - a][f(u_j) - a]\}$ for two Gaussian random variables u_i and u_j which are weakly correlated in the large N limit. In fact, as $i \neq j$:

$$\mathbb{E} u_i u_j = U_{ij} \quad (\text{A24})$$

is of the order of $1/\sqrt{D}$. In the thermodynamic limit, we can apply Lemma A1, which gives

$$\mathbb{E} f(u_i) f(u_j) = a^2 + b^2 \frac{1}{D} \sum_{r=1}^D F_{ir} F_{jr} \quad (i \neq j). \quad (\text{A25})$$

From Eqs. (A23) and (A25), we get the covariance of λ variables as written in Eq. (16). The covariance $\mathbb{E}[\nu^m \nu^m]$ is analogous.

We now compute the covariance $\mathbb{E}[\tilde{\lambda}^k \nu^m]$, which is equal to

$$\frac{1}{\sqrt{N}} \sum_{i=1}^N w_i^k \frac{1}{\sqrt{D}} \sum_{r=1}^D \tilde{w}_r^m \mathbb{E}[f(u_i) C_r]. \quad (\text{A26})$$

The two variables u_i and c_r are Gaussian random variables with a correlation

$$\mathbb{E}[u_i C_r] = \frac{1}{\sqrt{D}} F_{ir} \quad (\text{A27})$$

$$\mathbb{E}[f(u_i) C_r] = \frac{1}{\sqrt{D}} F_{ir} \langle u f(u) \rangle \langle C_r^2 \rangle = \frac{b}{\sqrt{D}} F_{ir}. \quad (\text{A28})$$

Using this result in Eq. (A26) gives Eq. (17).

which goes to zero as $\mathcal{O}(1/\sqrt{N})$ in the thermodynamic limit. We can, thus, use Lemma A1, and more precisely Eq. (A12), to get

b. Fourth moments of $\tilde{\lambda}^k$ variables

We study the fourth moment defined as

$$G^{k_1 k_2 k_3 k_4} = \langle \tilde{\lambda}^{k_1} \tilde{\lambda}^{k_2} \tilde{\lambda}^{k_3} \tilde{\lambda}^{k_4} \rangle = \frac{1}{N^2} \sum_{i_1, i_2, i_3, i_4} w_{i_1}^{k_1} w_{i_2}^{k_2} w_{i_3}^{k_3} w_{i_4}^{k_4} \langle \tilde{f}(u_{i_1}) \tilde{f}(u_{i_2}) \tilde{f}(u_{i_3}) \tilde{f}(u_{i_4}) \rangle, \quad (\text{A29})$$

where $\tilde{f}(u) = f(u) - a$ is the centered function.

We shall decompose the sum over i_1, i_2, i_3 , and i_4 depending on the number of distinct indices there are.

Distinct indices.—Let us study the first piece of the fourth moment $\langle \lambda^{k_1} \lambda^{k_2} \lambda^{k_3} \lambda^{k_4} \rangle$:

$$G_4^{k_1 k_2 k_3 k_4} = \frac{1}{N^2} \sum_{i_1, i_2, i_3, i_4} w_{i_1}^{k_1} w_{i_2}^{k_2} w_{i_3}^{k_3} w_{i_4}^{k_4} \langle \tilde{f}(u_{i_1}) \tilde{f}(u_{i_2}) \tilde{f}(u_{i_3}) \tilde{f}(u_{i_4}) \rangle, \quad (\text{A30})$$

where the sum runs over four indices i_1, i_2, i_3 , and i_4 which are distinct from each other. We can use the factorization property of the fourth moments of $f(u)$ of Lemma A2. This property gives

$$\begin{aligned} G_4^{k_1 k_2 k_3 k_4} &= \frac{1}{N^2} \sum_{i_1, i_2, i_3, i_4} w_{i_1}^{k_1} w_{i_2}^{k_2} w_{i_3}^{k_3} w_{i_4}^{k_4} [\langle \tilde{f}(u_{i_1}) \tilde{f}(u_{i_2}) \rangle \langle \tilde{f}(u_{i_3}) \tilde{f}(u_{i_4}) \rangle + \text{two perm}] \\ &= \left(\left[\frac{1}{N} \sum_{i_1, i_2} w_{i_1}^{k_1} w_{i_2}^{k_2} \langle f(u_{i_1}) f(u_{i_2}) \rangle \right] \left[\frac{1}{N} \sum_{i_3, i_4} w_{i_3}^{k_3} w_{i_4}^{k_4} \langle f(u_{i_3}) f(u_{i_4}) \rangle \right] - \text{corr} \right) \\ &\quad + \text{two perm}. \end{aligned} \quad (\text{A31})$$

The correction terms come from pieces where the intersection between $\{i_1, i_2\}$ and $\{i_3, i_4\}$ is nonempty. If we first neglect this correction, we find

$$G_4^{k_1 k_2 k_3 k_4} = b^4 [(\sum^{k_1 k_2} - W^{k_1 k_2})(\sum^{k_3 k_4} - W^{k_3 k_4}) + \text{two perm}]. \quad (\text{A32})$$

Now, we show that the corrections are negligible. Consider the term $i_1 = i_3, i_2 \neq i_4$. This term gives a correction

$$-\frac{1}{N^2} \sum_{i_1, i_2, i_4} w_{i_1}^{k_1} w_{i_2}^{k_2} w_{i_1}^{k_3} w_{i_4}^{k_4} [\langle \tilde{f}(u_{i_1}) \tilde{f}(u_{i_2}) \rangle \langle \tilde{f}(u_{i_1}) \tilde{f}(u_{i_4}) \rangle]. \quad (\text{A33})$$

Using Eq. (A12)

$$\langle \tilde{f}(u_{i_1}) \tilde{f}(u_{i_2}) \rangle = b^2 U_{i_1 i_2} = b^2 \frac{1}{D} \sum_{r=1}^D F_{i_1 r} F_{i_2 r}, \quad (\text{A34})$$

we get the expression for the correction

$$-\frac{1}{N^2 D^2} \langle u f(u) \rangle^4 \sum_{i_1, i_2, i_4} w_{i_1}^{k_1} w_{i_2}^{k_2} w_{i_1}^{k_3} w_{i_4}^{k_4} F_{i_1 r} F_{i_2 r} F_{i_1 s} F_{i_4 s} = -\frac{1}{\sqrt{N} D^2} \sum_{r, s} S_{rs}^{k_1 k_3} S_r^{k_2} S_s^{k_4}. \quad (\text{A35})$$

Using our hypothesis on the fact that the quantities S are of the order of one, this correction is clearly at most of the order of $\mathcal{O}(1/\sqrt{N})$ and, therefore, negligible.

The last correction that we need to consider is the term where $i_1 = i_3 = i$ and $i_2 = i_4 = j$. This term gives

$$-\frac{1}{N^2} \sum_{i,j} w_i^{k_1} w_j^{k_2} w_i^{k_3} w_j^{k_4} \langle \tilde{f}(u_i) \tilde{f}(u_j) \rangle^2 = -\frac{1}{NR^2} \langle uf(u) \rangle^4 \sum_{r,s} [S_{rs}^{k_1 k_3} S_{rs}^{k_2 k_4} - S_{rrss}^{k_1 k_2 k_3 k_4}], \quad (\text{A36})$$

which is again negligible in the large N limit.

Three distinct indices.—Let us study the contributions to the fourth moment of λ coming from three distinct indices. We study the case where $i_1 = i_4$:

$$E^{k_1 k_2 k_3 k_4} = \frac{1}{N^2} \sum_{i_1, i_2, i_3} w_{i_1}^{k_1} w_{i_2}^{k_2} w_{i_3}^{k_3} w_{i_1}^{k_4} \langle \tilde{f}(u_{i_1})^2 \tilde{f}(u_{i_2}) \tilde{f}(u_{i_3}) \rangle. \quad (\text{A37})$$

Using the expression for the third moment of functions of u_1, u_2 , and u_3 found in Eq. (A21), we get

$$\begin{aligned} E^{k_1 k_2 k_3 k_4} &= cb^2 \frac{1}{N^2} \sum_{i_1, i_2, i_3} w_{i_1}^{k_1} w_{i_2}^{k_2} w_{i_3}^{k_3} w_{i_1}^{k_4} - \text{corr} \\ &= cb^2 W^{k_1 k_4} [\Sigma^{k_2 k_3} - W^{k_2 k_3}] - \text{corr}. \end{aligned} \quad (\text{A38})$$

The corrections come from cases when $i_1 = i_2$ or $i_1 = i_3$. For instance, the piece with $i_1 = i_2$ gives

$$-cb^2 \frac{1}{ND} \sum_r S_r^{k_1 k_2 k_4} S_r^{k_3}, \quad (\text{A39})$$

which is $\mathcal{O}(1/N)$ at most.

The only pieces that do not vanish in the large N limit are, thus, the pieces similar to the one computed in Eq. (A38). Putting all of them together, we find that the contribution to $\langle \tilde{\lambda}^{k_1} \tilde{\lambda}^{k_2} \tilde{\lambda}^{k_3} \tilde{\lambda}^{k_4} \rangle$ coming from pieces with exactly three distinct indices in i_1, i_2, i_3 , and i_4 is equal to

$$\begin{aligned} G_3^{k_1 k_2 k_3 k_4} &= cb^2 (X^{k_1 k_2; k_3 k_4} + X^{k_1 k_3; k_2 k_4} + X^{k_1 k_4; k_2 k_3} \\ &\quad + X^{k_2 k_3; k_1 k_4} + X^{k_2 k_4; k_1 k_3} + X^{k_3 k_4; k_1 k_2}), \end{aligned}$$

where

$$X^{k_1 k_2; k_3 k_4} = W^{k_1 k_2} [\Sigma^{k_3 k_4} - W^{k_3 k_4}]. \quad (\text{A40})$$

Two distinct indices.—Let us now study the contribution to the fourth moment of λ coming from two distinct indices. We study first one piece of this contribution to the fourth moment, corresponding to $i_1 = i_2 = i$, $i_3 = i_4 = j$:

$$F^{k_1 k_2 k_3 k_4} = \frac{1}{N^2} \sum_{i,j} w_i^{k_1} w_i^{k_2} w_j^{k_3} w_j^{k_4} \langle \tilde{f}(u_i)^2 \tilde{f}(u_j)^2 \rangle. \quad (\text{A41})$$

To leading order in the thermodynamic limit, we can write

$$\langle \tilde{f}(u_i)^2 \tilde{f}(u_j)^2 \rangle = c^2 \quad (\text{A42})$$

and, therefore,

$$F^{k_1 k_2 k_3 k_4} = c^2 W^{k_1 k_2} W^{k_3 k_4} \quad (\text{A43})$$

[the correction coming from $i = j$ being obviously at most $\mathcal{O}(1/N)$].

We study now the second piece of this contribution to the fourth moment, corresponding to $i_1 = i_2 = i_3 = i$, $i_4 = j$. This piece is equal to

$$\frac{1}{N^2} \sum_{i,j} w_i^{k_1} w_i^{k_2} w_i^{k_3} w_j^{k_4} \langle \tilde{f}(u_i)^3 \tilde{f}(u_j) \rangle. \quad (\text{A44})$$

Using

$$\langle \tilde{f}(u_i)^3 \tilde{f}(u_j) \rangle = b \langle u \tilde{f}(u)^3 \rangle \frac{1}{D} \sum_r F_{ir} F_{jr}, \quad (\text{A45})$$

this piece gives

$$b \langle u \tilde{f}(u)^3 \rangle \frac{1}{ND} \sum_r S_r^{k_1 k_2 k_3} S_r^{k_4}, \quad (\text{A46})$$

and it is, therefore, negligible.

Therefore, all the contributions to the fourth moment of λ coming from exactly two distinct indices are of the type (A43). They give a total contribution:

$$G_2^{k_1 k_2 k_3 k_4} = c^2 [W^{k_1 k_2} W^{k_3 k_4} + W^{k_1 k_3} W^{k_2 k_4} + W^{k_1 k_4} W^{k_2 k_3}]. \quad (\text{A47})$$

One distinct index.—The contribution to the fourth moment $\langle \tilde{\lambda}^{k_1} \tilde{\lambda}^{k_2} \tilde{\lambda}^{k_3} \tilde{\lambda}^{k_4} \rangle$ coming from $i_1 = i_2 = i_3 = i_4$ is clearly of $\mathcal{O}(1/N)$ and can be neglected.

Final result for the four-point correlation function of λ variables.—We can now put together all the contributions to the fourth moment $\langle \tilde{\lambda}^{k_1} \tilde{\lambda}^{k_2} \tilde{\lambda}^{k_3} \tilde{\lambda}^{k_4} \rangle$ coming from pieces with four distinct indices found in Eq. (A32), those with three distinct indices found in Eq. (A40), and those with two distinct indices found in Eq. (A47). Defining

$$Y^{k_1 k_2} = \Sigma^{k_1 k_2} - W^{k_1 k_2} \quad (\text{A48})$$

and recalling the definition (A40) of the X variables, we obtain

$$\begin{aligned} & \langle \tilde{\lambda}^{k_1} \tilde{\lambda}^{k_2} \tilde{\lambda}^{k_3} \tilde{\lambda}^{k_4} \rangle \\ &= b^4 (Y^{k_1 k_2} Y^{k_3 k_4} + Y^{k_1 k_3} Y^{k_2 k_4} + Y^{k_1 k_4} Y^{k_2 k_3}) \\ & \quad + b^2 c (X^{k_1 k_2; k_3 k_4} + X^{k_1 k_3; k_2 k_4} \\ & \quad + X^{k_1 k_4; k_2 k_3} + X^{k_2 k_3; k_1 k_4} + X^{k_2 k_4; k_1 k_3} + X^{k_3 k_4; k_1 k_2}) \\ & \quad + c^2 [W^{k_1 k_2} W^{k_3 k_4} + W^{k_1 k_3} W^{k_2 k_4} + W^{k_1 k_4} W^{k_2 k_3}]. \end{aligned} \quad (\text{A49})$$

We can see that this result is equal to

$$[b^2 Y^{k_1 k_2} + c W^{k_1 k_2}] [b^2 Y^{k_3 k_4} + c W^{k_3 k_4}] + \text{two perm}, \quad (\text{A50})$$

which shows that

$$\langle \tilde{\lambda}^{k_1} \tilde{\lambda}^{k_2} \tilde{\lambda}^{k_3} \tilde{\lambda}^{k_4} \rangle = \langle \tilde{\lambda}^{k_1} \tilde{\lambda}^{k_2} \rangle \langle \tilde{\lambda}^{k_3} \tilde{\lambda}^{k_4} \rangle + \text{two permutations}. \quad (\text{A51})$$

With this result, it is clear how to proceed with the calculation of the fourth moments involving λ and ν variables. We first need to study the moments with three λ and one ν , then moments with two λ and two ν , and, finally, the moments with one λ and three ν variables. In the interest of conciseness, we do not spell out the full details of this calculations here, which proceeds very similarly to the calculations performed hitherto.

The generalization to higher moments of λ variables employs the same combination of repeated use of Lemma A2 and careful decomposition in subsets of distinct indices. As a result, it can be seen that the set of λ variables has a Gaussian distribution in the thermodynamic limit.

APPENDIX B: DERIVATION OF THE EQUATIONS OF MOTION

When we make a step of SGD, we update the weight w_i^k using a new sample, generated using a previously unused sample according to

$$(w_i^k)_{\mu+1} - (w_i^k)_\mu = -\frac{\eta}{\sqrt{N}} v^k \Delta g'(\lambda^k) f(u_i), \quad (\text{B1a})$$

$$v_{\mu+1}^k - v_\mu^k = -\frac{\eta}{N} g(\lambda^k) \Delta, \quad (\text{B1b})$$

where $\Delta = \sum_{j=1}^K v^j g(\lambda^j) - \sum_{m=1}^M \tilde{v}^m \tilde{g}(\nu^m)$. Note the different rescaling of the learning rate for the first- and second-layer weights. From here on out, we drop the index μ on the right-hand side as we work at a fixed iteration time. We keep the convention of Sec. III A, where extensive indices (taking values up to N or D) are below the line, while we use upper indices when they take finite values up to M or K . The challenge of controlling the learning in the

thermodynamic limit is to write closed equations using matrices with only ‘‘upper’’ indices left. Furthermore, we adopt the convention that the indices $j, k, \ell, t = 1, \dots, K$ always denote *student* nodes, while $n, m = 1, \dots, M$ are reserved for teacher hidden nodes.

1. First steps

We start by focusing on the dynamics of the first layer [Eq. (B1)]. When we study the evolution of quantities that are linear in the weights, like S_r^k and the order parameters constructed from it, e.g., $\Sigma^{k\ell}$, we need to study

$$\begin{aligned} & \left[\sum_{j=1}^K v^j g(\lambda^j) - \sum_{m=1}^M \tilde{v}^m \tilde{g}(\nu^m) \right] g'(\lambda^k) f(u_i) \\ &= \sum_{j \neq k}^K v^j a_i^{jk} + v^k \varrho_i^k - \sum_{n=1}^M \tilde{v}^n c_i^{nk}, \end{aligned} \quad (\text{B2})$$

where

$$a_i^{jk} = g(\lambda^j) g'(\lambda^k) f(u_i), \quad (\text{B3})$$

$$\varrho_i^k = g(\lambda^k) g'(\lambda^k) f(u_i), \quad (\text{B4})$$

$$c_i^{nk} = \tilde{g}(\nu^n) g'(\lambda^k) f(u_i), \quad (\text{B5})$$

while we keep the second-layer weights v^k fixed. We can, thus, follow the dynamics of S_r^k [Eq. (20)], which is linear in the weights and enters the definition of the order parameters R^{km} [Eq. (17)] and Σ^{kl} [Eq. (22)]:

$$\begin{aligned} & (S_r^k)_{\mu+1} - (S_r^k)_\mu \\ &= -\frac{\eta}{N} v^k \sum_i F_{ir} \left[\sum_{j \neq k}^K v^j a_i^{jk} + v^k \varrho_i^k - \sum_n^M \tilde{v}^n c_i^{nk} \right]. \end{aligned} \quad (\text{B6})$$

We want to average this update equation over a new incoming sample, i.e., over the c_r variables. Upon contraction with F_{ir} in Eq. (B6), we are thus led to computing the averages

$$A_r^{jk} \equiv \frac{1}{\sqrt{N}} \sum_i \mathbb{E}[F_{ir} a_i^{jk}] = \mathbb{E}[g(\lambda^j) g'(\lambda^k) \beta_r], \quad (\text{B7})$$

$$\mathcal{B}_r^k \equiv \mathbb{E}[g(\lambda^k) g'(\lambda^k) \beta_r], \quad (\text{B8})$$

and

$$C_r^{nk} = \mathbb{E}[\tilde{g}(\nu^n) g'(\lambda^k) \beta_r], \quad (\text{B9})$$

where

$$\beta_r = \frac{1}{\sqrt{N}} \sum_i F_{ir} f(u_i). \quad (\text{B10})$$

The crucial fact that allows for an analytic study of online learning is that, at each step μ of SGD, a previously unseen input x_μ is used to evaluate the gradient. The latent representation c_μ of this input is given by a new set of IID Gaussian random variables $c_{\mu r}$, which are thus independent of the current weights of the student at that time. In the thermodynamic limit, the GEP of the previous section shows that, for one given value of r , the $K + M + 1$ variables $\{\lambda^k\}$, $\{\nu^m\}$, and β_r have a joint Gaussian distribution, making it possible to express the averages over $\{\lambda^k, \nu^m, \beta_r\}$ in terms of only their covariances.

Looking closer, we see that the average of Eqs. (B7)–(B9) over this Gaussian distribution involves two sets of random variables: on the one hand, the $M + K$ local fields $\{\nu^m, \lambda^k\}$, which have correlations of the order of 1, and, on the other hand, the variable β_r (for one given value of r). It turns out that β_r is only weakly correlated with the local fields $\{\nu^m, \lambda^k\}$ [the correlation is $\mathcal{O}(1/\sqrt{N})$]. In Appendix A 1, we discuss how to compute this type of average and prove Lemma A1, which for the averages (B7)–(B9) yields

$$\mathcal{A}_r^{jk} = \frac{1}{Q^{kk}Q^{jj} - (Q^{kj})^2} \{Q^{jj}\mathbb{E}[g'(\lambda^k)\lambda^k g(\lambda^j)]\mathbb{E}[\lambda^k\beta_r] - Q^{kj}\mathbb{E}[g'(\lambda^k)\lambda^j g(\lambda^j)]\mathbb{E}[\lambda^k\beta_r] - Q^{kj}\mathbb{E}[g'(\lambda^k)\lambda^k g(\lambda^j)]\mathbb{E}[\lambda^j\beta_r] + Q^{kk}\mathbb{E}[g'(\lambda^k)\lambda^j g(\lambda^j)]\mathbb{E}[\lambda^j\beta_r]\}, \quad (\text{B11})$$

$$\mathcal{B}_r^k = \frac{1}{Q^{kk}}\mathbb{E}[g'(\lambda^k)\lambda^k g(\lambda^k)]\mathbb{E}[\lambda^k\beta_r], \quad (\text{B12})$$

$$\mathcal{C}_r^{nk} = \frac{1}{Q^{kk}T^{nn} - (R^{kn})^2} \{T^{nn}\mathbb{E}[g'(\lambda^k)\lambda^k \tilde{g}(\nu^n)]\mathbb{E}[\lambda^k\beta_r] - R^{kn}\mathbb{E}[g'(\lambda^k)\nu^n \tilde{g}(\nu^n)]\mathbb{E}[\lambda^k\beta_r] - R^{kn}\mathbb{E}[g'(\lambda^k)\lambda^k \tilde{g}(\nu^n)]\mathbb{E}[\nu^n\beta_r] + Q^{kk}\mathbb{E}[g'(\lambda^k)\nu^n \tilde{g}(\nu^n)]\mathbb{E}[\nu^n\beta_r]\}. \quad (\text{B13})$$

This result yields

$$(S_r^k)_{\mu+1} - (S_r^k)_\mu = -\frac{\eta}{\sqrt{N}}v^k \left[\sum_{j \neq k}^K v^j \mathcal{A}_r^{jk} + v^k \mathcal{B}_r^k - \sum_n^M \tilde{v}^n \mathcal{C}_r^{nk} \right], \quad (\text{B14})$$

with only the single extensive index r left. While this equation would appear to open up a way to write down the equation of motion for the “teacher-student” overlap R^{km} by contracting Eq. (B14) with \tilde{w}_r^m , we show in Appendix C that such a program leads to an infinite hierarchy of equations. To avoid this problem, we rotate the problem to a different basis, as we explain in the next section.

2. Changing the basis to close the equations

We can close the equations for the order parameters by studying their dynamics in the basis given by the eigenvectors of the operator

$$\Omega_{rs} \equiv \frac{1}{N} \sum_i F_{ir} F_{is}, \quad (\text{B15})$$

which is a $D \times D$ symmetric matrix, with diagonal elements $\Omega_{rr} = 1$, and off-diagonal elements of the order of $1/\sqrt{N}$. Consider the orthogonal basis of eigenvectors $\psi_{\tau=1, \dots, D}$ of this matrix, with corresponding eigenvalues ρ_τ , such that

$$\sum_s \Omega_{rs} \psi_{\tau s} = \rho_\tau \psi_{\tau r}. \quad (\text{B16})$$

We suppose that the components of the eigenvectors $\psi_{\tau r}$ are of the order of 1, and we impose the following normalization:

$$\sum_s \psi_{\tau s} \psi_{\tau' s} = D \delta_{\tau \tau'}, \quad \sum_\tau \psi_{\tau r} \psi_{\tau s} = D \delta_{rs}. \quad (\text{B17})$$

In this basis, the teacher-student overlap R^{km} [Eq. (17)] is given by

$$R^{km} = \frac{b}{D} \sum_\tau \Gamma_\tau^k \tilde{\omega}_\tau^m, \quad (\text{B18})$$

where we introduce the projections

$$\Gamma_\tau^k = \frac{1}{\sqrt{D}} \sum_r S_r^k \psi_{\tau r} \quad (\text{B19})$$

and

$$\tilde{\omega}_\tau^m = \frac{1}{\sqrt{D}} \sum_r \tilde{w}_r^m \psi_{\tau r}. \quad (\text{B20})$$

Since $\tilde{\omega}_\tau^m$ is a static variable, the time evolution of Γ_τ^k is given by

$$\begin{aligned} (\Gamma_\tau^k)_{\mu+1} - (\Gamma_\tau^k)_\mu &= -\frac{\eta}{\sqrt{\delta}N} v^k \sum_r \psi_{\tau r} \left[\sum_{j \neq k}^K v^j \mathcal{A}_r^{jk} \right. \\ &\quad \left. + v^k \mathcal{B}_r^k - \sum_n^M \tilde{v}^n \mathcal{C}_r^{nk} \right]. \end{aligned} \quad (\text{B21})$$

As we aim to compute the remaining sum over r , two types of terms appear:

$$\sum_r \psi_{\tau r} \mathbb{E}[\lambda^k \beta_r] = \frac{1}{\sqrt{\delta}} [(c - b^2)\delta + b^2 \rho_\tau] \Gamma_\tau^k = \frac{d_\tau}{\sqrt{\delta}} \Gamma_\tau^k, \quad (\text{B22})$$

where we define $d_\tau = (c - b^2)\delta + b^2 \rho_\tau$, and

$$\sum_r \psi_{\tau r} \mathbb{E}[\nu^n \beta_r] = \frac{b}{\sqrt{\delta}} \rho_\tau \tilde{\omega}_\tau^n. \quad (\text{B23})$$

Putting everything together, the final evolution of Γ_τ^k is

$$\begin{aligned} (\Gamma_\tau^k)_{\mu+1} - (\Gamma_\tau^k)_\mu &= -\frac{\eta}{\delta N} v^k \left[d_\tau \Gamma_\tau^k \sum_{j \neq k} v^j \frac{Q^{jj} \mathbb{E}[g'(\lambda^k) \lambda^k g(\lambda^j)] - Q^{kj} \mathbb{E}[g'(\lambda^k) \lambda^j g(\lambda^j)]}{Q^{kk} Q^{jj} - (Q^{kj})^2} \right. \\ &\quad + \sum_{j \neq k} v^j d_\tau \Gamma_\tau^j \frac{Q^{kk} \mathbb{E}[g'(\lambda^k) \lambda^j g(\lambda^j)] - Q^{kj} \mathbb{E}[g'(\lambda^k) \lambda^k g(\lambda^j)]}{Q^{kk} Q^{jj} - (Q^{kj})^2} + d_\tau v^k \Gamma_\tau^k \frac{1}{Q^{kk}} \mathbb{E}[g'(\lambda^k) \lambda^k g(\lambda^k)] \\ &\quad - d_\tau \Gamma_\tau^k \sum_n \tilde{v}^n \frac{T^{nn} \mathbb{E}[g'(\lambda^k) \lambda^k \tilde{g}(\nu^n)] - R^{kn} \mathbb{E}[g'(\lambda^k) \nu^n \tilde{g}(\nu^n)]}{Q^{kk} T^{nn} - (R^{kn})^2} \\ &\quad \left. - b \rho_\tau \sum_n \tilde{v}^n \tilde{\omega}_\tau^n \frac{Q^{kk} \mathbb{E}[g'(\lambda^k) \nu^n \tilde{g}(\nu^n)] - R^{kn} \mathbb{E}[g'(\lambda^k) \lambda^k \tilde{g}(\nu^n)]}{Q^{kk} T^{nn} - (R^{kn})^2} \right]. \end{aligned} \quad (\text{B24})$$

At this point, we note that the remaining averages appearing in this expression, such as $\mathbb{E}[\lambda^k g'(\lambda^k) \tilde{g}(\nu^m)]$, depend only on subsets of the local fields $\{\lambda^{k=1, \dots, K}, \nu^{m=1, \dots, M}\}$. As discussed above, the GEP guarantees that these random variables follow a multidimensional Gaussian distribution, so these averages depend only on the covariances of the local fields R^{km} , $Q^{k\ell}$, and T^{mn} . To simplify the subsequent equations, we use the shorthand notation for the three-dimensional Gaussian averages

$$I_3(k, j, n) \equiv \mathbb{E}[g'(\lambda^k) \lambda^j \tilde{g}(\nu^n)], \quad (\text{B25})$$

which was introduced by Saad and Solla [22] and that we discuss in the main text. To make this section self-contained, we recall that arguments passed to I_3 should be translated into local fields on the right-hand side by using the convention where the indices j, k, ℓ , and l always refer to student local fields λ^j , etc., while the indices n and m always refer to teacher local fields ν^n and ν^m , respectively. Similarly,

$$I_3(k, j, j) \equiv \mathbb{E}[g'(\lambda^k) \lambda^j g(\lambda^j)], \quad (\text{B26})$$

where having the index j as the third argument means that the third factor is $g(\lambda^j)$ rather than $\tilde{g}(\nu^m)$ in Eq. (B26). The average in Eq. (B26) is taken over a three-dimensional normal distribution with mean zero and covariance matrix

$$\begin{aligned} \Phi^{(3)}(k, j, n) &= \begin{pmatrix} \mathbb{E}[\lambda^k \lambda^k] & \mathbb{E}[\lambda^k \lambda^j] & \mathbb{E}[\lambda^k \nu^n] \\ \mathbb{E}[\lambda^k \lambda^j] & \mathbb{E}[\lambda^j \lambda^j] & \mathbb{E}[\lambda^j \nu^n] \\ \mathbb{E}[\lambda^k \nu^n] & \mathbb{E}[\lambda^j \nu^n] & \mathbb{E}[\nu^n \nu^n] \end{pmatrix} \\ &= \begin{pmatrix} Q^{kk} & Q^{kj} & R^{kn} \\ Q^{kj} & Q^{jj} & R^{jn} \\ R^{kn} & R^{jn} & T^{nn} \end{pmatrix}. \end{aligned} \quad (\text{B27})$$

3. Dynamics of the teacher-student overlap R^{km}

We are now in a position to write the update equation for

$$(R^{km})_{\mu+1} - (R^{km})_\mu = \frac{b}{D} \sum_\tau [(\Gamma_\tau^k)_{\mu+1} - (\Gamma_\tau^k)_\mu] \tilde{\omega}_\tau^m, \quad (\text{B28})$$

where we use that the $\tilde{\omega}_\tau^m$ are static. When performing the last remaining sum over τ , two types of terms appear. First, there is

$$\tilde{T}^{mn} \equiv \frac{1}{D} \sum_\tau \rho_\tau \tilde{\omega}_\tau^m \tilde{\omega}_\tau^n, \quad (\text{B29})$$

which depends only on the choice of the feature matrix F_{ir} and the teacher weights w_{mr}^* and is, thus, a constant of the motion. The second type of term is of the form

$$\frac{1}{D} \sum_{\tau} \rho_{\tau} \Gamma_{\tau}^{\ell} \tilde{\omega}_{\tau}^n. \quad (\text{B30})$$

This sum cannot be reduced to a simple expression in terms of other order parameters. Instead, we are led to introduce the density

$$r^{km}(\rho) = \frac{1}{\varepsilon_{\rho}} \frac{1}{D} \sum_{\tau} \Gamma_{\tau}^k \tilde{\omega}_{\tau}^m \mathbb{1}(\rho_{\tau} \in [\rho, \rho + \varepsilon_{\rho}]), \quad (\text{B31})$$

where $\mathbb{1}(\cdot)$ is the indicator function which evaluates to 1 if the condition given to it as an argument is true and which otherwise evaluates to 0. We take the limit $\varepsilon_{\rho} \rightarrow 0$ after the thermodynamic limit. Then, we can rewrite the order parameter R^{km} as an integral over the density r^{km} , weighted by the distribution of eigenvalues of the operator Ω_{rs} , $p_{\Omega}(\rho)$:

$$R^{km} = b \int d\rho p_{\Omega}(\rho) r^{km}(\rho). \quad (\text{B32})$$

If, for example, we take the elements of the feature matrix F_{ir} to be elementwise IID from the normal distribution with mean zero and unit variance, then the limiting density of eigenvalues of Ω is given by the Marchenko-Pastur law [66]:

$$p_{\text{MP}}(\rho) = \frac{1}{2\pi\delta} \frac{\sqrt{(\rho_{\text{max}} - \rho)(\rho - \rho_{\text{min}})}}{\rho}, \quad (\text{B33})$$

where $\rho_{\text{min}} = (1 - \sqrt{\delta})^2$ and $\rho_{\text{max}} = (1 + \sqrt{\delta})^2$.

The update equation of $r^{km}(\rho)$ can be obtained immediately by substituting the update equation for Γ_{τ}^k [Eq. (B24)] into its definition [Eq. (B31)]. Finally, in the thermodynamic limit, the normalized number of steps $t = P/N$ can be interpreted as a continuous timelike variable, and so we have

$$R^{km}(t) = b \int d\rho p_{\Omega}(\rho) r^{km}(\rho, t), \quad (\text{B34})$$

and we recover the equation of motion for $r^{km}(\rho)$, which we restate here for ease of reading:

$$\begin{aligned} \frac{\partial r^{km}(\rho, t)}{\partial t} = & -\frac{\eta}{\delta} v^k \left(d(\rho) r^{km}(\rho) \sum_{j \neq k} v^j \frac{Q^{jj} I_3(k, k, j) - Q^{kj} I_3(k, j, j)}{Q^{jj} Q^{kk} - (Q^{kj})^2} \right. \\ & + d(\rho) \sum_{j \neq k} v^j r^{jm}(\rho) \frac{Q^{kk} I_3(k, j, j) - Q^{kj} I_3(k, k, j)}{Q^{jj} Q^{kk} - (Q^{kj})^2} \\ & + v^k d(\rho) r^{km}(\rho) \frac{1}{Q^{kk}} I_3(k, k, k) \\ & - d(\rho) r^{km}(\rho) \sum_n \tilde{v}^n \frac{T^{nn} I_3(k, k, n) - R^{kn} I_3(k, n, n)}{Q^{kk} T^{nn} - (R^{kn})^2} \\ & \left. - b\rho \sum_n \tilde{v}^n \tilde{T}^{nm} \frac{Q^{kk} I_3(k, n, n) - R^{kn} I_3(k, k, n)}{Q^{kk} T^{nn} - (R^{kn})^2} \right), \end{aligned} \quad (\text{B35})$$

where $d(\rho) = (c - b^2)\delta + b^2\rho$. Note that, while we drop the explicit time dependence from the right-hand side to keep the equation readable, all the order parameters on the right-hand side are explicitly time dependent, i.e., $Q^{jj} = Q^{jj}(t)$, $r^{km}(\rho) = r^{km}(\rho, t)$, and the averages $I_3(\cdot)$ are also time dependent via their dependence on the order parameters [see Eq. (B26) and the subsequent discussion]. In order to close the equations of motion, we now need to find the equations for the order parameters that are quadratic in the weights.

4. Order parameters that are quadratic in the weights

There are two order parameters that appear when evaluating the covariance of the λ variables:

$$Q^{k\ell} \equiv \mathbb{E}[\lambda^k \lambda^{\ell}] = [c - b^2] W^{k\ell} + b^2 \Sigma^{k\ell}. \quad (\text{B36})$$

We look at both $W^{k\ell}$ and $\Sigma^{k\ell}$ in turn now.

Equation of motion for $W^{k\ell}$.—For the student-student overlap matrix

$$W^{k\ell} = \frac{1}{N} \sum_i^N w_i^k w_i^{\ell}, \quad (\text{B37})$$

we find, after some algebra, that updates read

$$\begin{aligned}
 (W^{k\ell})^{\mu+1} - (W^{k\ell})_{\mu} &= -\frac{\eta}{N^{3/2}} v^k \sum_i w_i^{\ell} \left[\sum_{j \neq k}^K v^j a_i^{jk} + v^k \ell_i^k - \sum_n^M \tilde{v}^n c_i^{nk} \right] \\
 &\quad - \frac{\eta}{N^{3/2}} v^{\ell} \sum_i w_i^k \left[\sum_{j \neq \ell}^K v^j a_i^{j\ell} + v^{\ell} \ell_i^{\ell} - \sum_n^M \tilde{v}^n c_i^{n\ell} \right] \\
 &\quad + \frac{\eta^2}{N^2} v^k v^{\ell} \sum_i f(u_i)^2 g'(\lambda^k) g'(\lambda^{\ell}) \left[\sum_{j,i}^K v^j v^i g(\lambda^j) g(\lambda^i) + \sum_{n,m}^M \tilde{v}^n \tilde{v}^m \tilde{g}(\nu^n) \tilde{g}(\nu^m) \right. \\
 &\quad \left. - 2 \sum_j^K \sum_m^M v^j \tilde{v}^m g(\lambda^j) \tilde{g}(\nu^m) \right]. \tag{B38}
 \end{aligned}$$

For the terms linear in the learning rate η , we can immediately carry out the sum over i , which yields terms of the type

$$\frac{1}{\sqrt{N}} \sum_i w_i^{\ell} \mathbb{E}[g(\lambda^j) g'(\lambda^k) f(u_i)] = \mathbb{E}[g'(\lambda^k) \lambda^{\ell} g(\lambda^j)] = I_3(k, \ell, j), \quad \text{etc.} \tag{B39}$$

The term quadratic in the learning rate η requires the evaluation of terms of the type

$$\frac{1}{N} \sum_i \mathbb{E}[f(u_i)^2 g'(\lambda^k) g'(\lambda^{\ell}) g(\lambda^j) g(\lambda^i)] = c \mathbb{E}[g'(\lambda^k) g'(\lambda^{\ell}) g(\lambda^j) g(\lambda^i)]. \tag{B40}$$

The sum over i thus makes this second-order term contribute to the total variation of $W^{k\ell}$ at leading order, and we are left with an average over four local fields, for which we introduce the shorthand

$$I_4(k, \ell, j, i) \equiv \mathbb{E}[g'(\lambda^k) g'(\lambda^{\ell}) g(\lambda^j) g(\lambda^i)], \tag{B41}$$

where we use the same notation as we did for $I_3(\cdot)$ [Eq. (B26)]. The full equation of motion for $W^{k\ell}$ thus reads

$$\begin{aligned}
 \frac{dW^{k\ell}(t)}{dt} &= -\eta v^k \left(\sum_j^K v^j I_3(k, \ell, j) - \sum_n^M \tilde{v}^n I_3(k, \ell, n) \right) - \eta v^{\ell} \left(\sum_j^K v^j I_3(\ell, k, j) - \sum_n^M \tilde{v}^n I_3(\ell, k, n) \right) \\
 &\quad + c \eta^2 v^k v^{\ell} \left(\sum_{j,a}^K v^j v^a I_4(k, \ell, j, a) - 2 \sum_j^K \sum_m^M v^j \tilde{v}^m I_4(k, \ell, j, m) + \sum_{n,m}^M \tilde{v}^n \tilde{v}^m I_4(k, \ell, n, m) \right). \tag{B42}
 \end{aligned}$$

Equation of motion for $\Sigma^{k\ell}$.—After rotating to the basis ψ_{τ} , we have

$$\Sigma^{k\ell} \equiv \frac{1}{D} \sum_r S_r^k S_r^{\ell} = \frac{1}{D} \sum_{\tau} \Gamma_{\tau}^k \Gamma_{\tau}^{\ell}. \tag{B43}$$

It is then immediate that

$$\begin{aligned}
 (\Sigma^{k\ell})^{\mu+1} - (\Sigma^{k\ell})_{\mu} &= \frac{1}{D} \sum_{\tau} (\Gamma_{\tau}^{\ell})_{\mu} [(\Gamma_{\tau}^k)^{\mu+1} - (\Gamma_{\tau}^k)_{\mu}] + \frac{1}{D} \sum_{\tau} (\Gamma_{\tau}^k)_{\mu} [(\Gamma_{\tau}^{\ell})^{\mu+1} - (\Gamma_{\tau}^{\ell})_{\mu}] \\
 &\quad + \frac{\eta^2}{D^2 N} \sum_{\tau} \sum_{r,s}^R \psi_{\tau r} \psi_{\tau s} \mathbb{E}[\Delta^2 g'(\lambda^k) g'(\lambda^{\ell}) \beta_r \beta_s]. \tag{B44}
 \end{aligned}$$

The terms linear in η can be obtained directly by substituting in the update equation for Γ_{τ}^k [Eq. (B24)] and are similar to the update equation for $r^{km}(\rho)$. As for the term quadratic in η , we have to leading order

$$\begin{aligned} \frac{\eta^2}{DN} \sum_{r,s} \psi_{\tau r} \psi_{\tau s} \mathbb{E}[\Delta^2 g'(\lambda^k) g'(\lambda^\ell) \beta_r \beta_s] &= \frac{\eta^2}{DN} \sum_r (\psi_{\tau r})^2 \mathbb{E}[\Delta^2 g'(\lambda^k) g'(\lambda^\ell)] \mathbb{E}[\beta_r^2] \\ &= \frac{\eta^2}{N} \mathbb{E}[\Delta^2 g'(\lambda^k) g'(\lambda^\ell)] \left[(c - b^2) \rho_\tau + \frac{b^2}{\delta} \rho_\tau^2 \right], \end{aligned} \quad (\text{B45})$$

where we use that covariance of β_r is given by

$$\mathbb{E}[\beta_r^2] = c - b^2 + \frac{b^2}{\delta} \sum_t \Omega_{rt}^2. \quad (\text{B46})$$

To deal with the remaining sum over τ , we again make use of the fact that the equation of motion for $\Sigma^{k\ell}$ depends on the eigenvector index τ only through the eigenvalue ρ_τ . Introducing the density

$$\sigma^{k\ell}(\rho) = \frac{1}{\varepsilon_\rho} \frac{1}{D} \sum_\tau \Gamma_\tau^k \Gamma_\tau^\ell \mathbb{1}(\rho_\tau \in [\rho, \rho + \varepsilon_\rho]), \quad (\text{B47})$$

as we did for $r^{km}(\rho)$ [Eq. (B31)], we have

$$\Sigma^{k\ell}(t) = \int d\rho p_\Omega(\rho) \sigma^{k\ell}(\rho, t) \quad (\text{B48})$$

with

$$\begin{aligned} \frac{\partial \sigma^{k\ell}(\rho, t)}{\partial t} &= -\frac{\eta}{\delta} \left(d(\rho) v^k \sigma^{k\ell}(\rho) \sum_{j \neq k} v^j \frac{Q^{jj} I_3(k, k, j) - Q^{kj} I_3(k, j, j)}{Q^{jj} Q^{kk} - (Q^{kj})^2} \right. \\ &\quad + v^k \sum_{j \neq k} v^j d(\rho) \sigma^{j\ell}(\rho) \frac{Q^{kk} I_3(k, j, j) - Q^{kj} I_3(k, k, j)}{Q^{jj} Q^{kk} - (Q^{kj})^2} \\ &\quad + d(\rho) v^k \sigma^{k\ell}(\rho) v^k \frac{1}{Q^{kk}} I_3(k, k, k) \\ &\quad - d(\rho) v^k \sigma^{k\ell}(\rho) \sum_n \tilde{v}^n \frac{T^{nn} I_3(k, k, n) - R^{kn} I_3(k, n, n)}{Q^{kk} T^{nn} - (R^{kn})^2} \\ &\quad - b \rho v^k \sum_n \tilde{v}^n r^{\ell n}(\rho) \frac{Q^{kk} I_3(k, n, n) - R^{kn} I_3(k, k, n)}{Q^{kk} T^{nn} - (R^{kn})^2} \\ &\quad \left. + \text{all of the above with } \ell \rightarrow k, k \rightarrow \ell \right) \\ &\quad + \eta^2 v^k v^\ell \left[(c - b^2) \rho + \frac{b^2}{\delta} \rho^2 \right] \left(\sum_{j, \iota} v^j v^\iota I_4(k, \ell, j, \iota) \right. \\ &\quad \left. - 2 \sum_j^K \sum_m^M v^j \tilde{v}^m I_4(k, \ell, j, m) + \sum_{n, m}^M \tilde{v}^n \tilde{v}^m I_4(k, \ell, n, m) \right). \end{aligned} \quad (\text{B49})$$

5. Second-layer weights

Finally, we treat each of the second-layer weights of the student v as an order parameter in its own right. Their equations of motion are readily found from their SGD update [Eq. (B1b)] and read

$$\frac{dv^k}{dt} = \eta \left[\sum_n^M \tilde{v}_n I_2(k, n) - \sum_j^K v^j I_2(k, j) \right], \quad (\text{B50})$$

where we introduce the final shorthand

$$I_2(k, j) \equiv \mathbb{E}[g(\lambda^k)g(\lambda^j)], \quad \text{etc.} \quad (\text{B51})$$

where we again use the notation we introduced for $I_3(\cdot)$ [Eq. (B26)].

6. Generalization error

Having introduced the shorthand for the integrals $I_2(k, j)$ [Eq. (B51)], we realize that their form also enters the formula for the generalization error

$$\epsilon_g(\theta, \tilde{\theta}) = \frac{1}{2} \mathbb{E} \left(\sum_k^K v^k g(\lambda^k) - \sum_m^M \tilde{v}^m \tilde{g}(\nu^m) \right)^2 = \frac{1}{2} \sum_{k,\ell} v^k v^\ell I_2(k, \ell) + \frac{1}{2} \sum_{n,m} \tilde{v}^n \tilde{v}^m I_2(n, m) - \sum_{k,n} v^k \tilde{v}^n I_2(k, n). \quad (\text{B52})$$

For example, for a student with $g(\lambda^k) = \text{erf}(\lambda^k/\sqrt{2})$ and a teacher with $\tilde{g}(\nu^m) = \max(0, \nu^m)$, we find that

$$\begin{aligned} \epsilon_g(Q^{k\ell}, R^{kn}, T^{nm}, v^k, \tilde{v}^m) &= \frac{1}{\pi} \sum_{k,\ell} v^k v^\ell \arcsin \frac{Q^{k\ell}}{\sqrt{1+Q^{kk}}\sqrt{1+Q^{\ell\ell}}} - \sum_{k,n} v^k \tilde{v}^n \frac{R^{kn}}{\sqrt{2\pi}\sqrt{1+Q^{kk}}} \\ &+ \frac{1}{4\pi} \sum_{n,m} \tilde{v}^n \tilde{v}^m \left(\sqrt{T^{mm}T^{nn} - (T^{nm})^2} + T^{nm} \left[\pi + 2 \arctan \frac{T^{nm}}{\sqrt{T^{mm}T^{nn} - (T^{nm})^2}} \right] \right). \end{aligned} \quad (\text{B53})$$

7. Summary of the derivation

We have now completed the program that we embarked upon at the beginning of this Appendix: We derive a closed set of equations of motion for the teacher-student overlap R^{km} [Eqs. (B32) and (29)], the student-student overlap $Q^{k\ell} = [c - b^2]W^{k\ell} + b^2\Sigma^{k\ell}$ [Eqs. (38), (34), and (35)], and the student's second-layer weights v^k [Eq. (39)]. These equations give us complete access to the dynamics of a neural network performing one-shot stochastic gradient descent on a dataset generated by the hidden manifold model. We can now integrate these equations and substitute the values of the order parameters at any time into the expression for the generalization error (26), thereby tracking the dynamics of the generalization error at all times. We describe this procedure in more detail next.

8. Explicit form of the integrals I_3 and I_4 for sigmoidal students

The explicit forms of the integrals I_3 and I_4 that appear in the equations of motion for the order parameters and the generalization error for networks with $g(x) = \tilde{g}(x) = \text{erf}(x/\sqrt{2})$ were first given by Refs. [21,22]. Here, we state them to make the paper as self-contained as possible. Denoting the elements of the covariance matrix such as Φ^3 [Eq. (B27)] as ϕ_{ij} , we have

$$I_3(\cdot, \cdot, \cdot) = \frac{2}{\pi} \frac{1}{\sqrt{\Lambda_3}} \frac{\phi_{23}(1 + \phi_{11}) - \phi_{12}\phi_{13}}{1 + \phi_{11}} \quad (\text{B54})$$

with

$$\Lambda_3 = (1 + \phi_{11})(1 + \phi_{33}) - \phi_{13}^2. \quad (\text{B55})$$

For the average I_4 , we have a covariance matrix $\Phi^{(4)}$ that is populated in analogy to $\Phi^{(3)}$ [Eq. (B27)], we have

$$I_4(\cdot, \cdot, \cdot, \cdot) = \frac{4}{\pi^2} \frac{1}{\sqrt{\Lambda_4}} \arcsin \left(\frac{\Lambda_0}{\sqrt{\Lambda_1\Lambda_2}} \right), \quad (\text{B56})$$

where

$$\Lambda_4 = (1 + \phi_{11})(1 + \phi_{22}) - \phi_{12}^2, \quad (\text{B57})$$

$$\begin{aligned} \Lambda_0 &= \Lambda_4\phi_{34} - \phi_{23}\phi_{24}(1 + \phi_{11}) - \phi_{13}\phi_{14}(1 + \phi_{22}) \\ &+ \phi_{12}\phi_{13}\phi_{24} + \phi_{12}\phi_{14}\phi_{23}, \end{aligned} \quad (\text{B58})$$

$$\begin{aligned} \Lambda_1 &= \Lambda_4(1 + \phi_{33}) - \phi_{23}^2(1 + \phi_{11}) - \phi_{13}^2(1 + \phi_{22}) \\ &+ 2\phi_{12}\phi_{13}\phi_{23}, \end{aligned} \quad (\text{B59})$$

$$\begin{aligned} \Lambda_2 &= \Lambda_4(1 + \phi_{44}) - \phi_{24}^2(1 + \phi_{11}) - \phi_{14}^2(1 + \phi_{22}) \\ &+ 2\phi_{12}\phi_{14}\phi_{24}. \end{aligned} \quad (\text{B60})$$

APPENDIX C: THE EQUATIONS OF MOTION DO NOT CLOSE IN THE TRIVIAL BASIS

Here, we give a short demonstration that it is not possible to close the equations for order parameters if we do not rotate their dynamics to the basis given by the eigenvectors of Ω , which is what we do in our derivation in Sec. III B.

1. Order parameters that are linear in the weights

To start with a variable that is linear in the weights, take the time evolution of S_r^k . It is clear that the tensor structure of the result (B14) is of the form

$$(S_r^k)^{\mu+1} - (S_r^k)^\mu = -\frac{\eta}{N} \left[\sum_{\ell} D^{k\ell} \sum_s \Omega_{rs} S_s^\ell + \sum_m E^{km} \sum_s \Omega_{rs} \tilde{w}_s^m \right], \quad (\text{C1})$$

where $D^{k\ell}$ and E^{km} are known quantities, expressed in terms of the matrices Q , T , and R , and we introduce the operator

$$\Omega_{rs} = \frac{1}{N} \sum_i F_{ir} F_{is}, \quad (\text{C2})$$

which has diagonal elements equal to 1 and off-diagonal elements of the order of $1/\sqrt{N}$.

In particular, we can use this evolution to study the evolution of R :

$$(R^{km})^{\mu+1} - (R^{km})^\mu = -\langle uf(u) \rangle \frac{\eta}{N} \left[\sum_{\ell} D^{k\ell} \frac{1}{D} \sum_{rs} \tilde{w}_r^m \Omega_{rs} S_s^\ell + \sum_m E^{km} \frac{1}{D} \sum_{rs} \tilde{w}_s^r \Omega_{rs} \tilde{w}_s^m \right]. \quad (\text{C3})$$

The point of this analysis is to show that the time evolution of S_r^k involves $(\Omega S)_r^\ell$. Therefore, to know the evolution of S , we need the one of ΩS . This process is not innocuous because, in order to have dynamical evolution equations with only “up” indices, we need to contract it. The evolution of R^{km} , which is proportional to the scalar product (in the R -dimensional manifold space) of S^k and \tilde{w}^m , is thus given by the scalar product of ΩS^k and \tilde{w}^m .

It is not difficult to see that the evolution of ΩS will require knowing $\Omega^2 S$, etc. So we have an infinite hierarchy of coupled equations, which would be hard to analyze. Yet, we can find closed equations by changing the basis for S .

APPENDIX D: ADDITIONAL DETAILS ON THE NUMERICAL EXPERIMENTS IN SEC. VA

For the experiments demonstrating the learning of functions of increasing complexity discussed in Sec. VA, we

constructed datasets for binary classification by splitting the image datasets as follows:

- (i) *MNIST*.—Even vs odd numbers.
- (ii) *Fashion-MNIST*.—T-shirt or top, pullover, dress, sandal, and bag vs trouser, coat, shirt, sneaker, and ankle boot.
- (iii) *CIFAR10*.—Airplane, bird, deer, frog, and ship vs automobile, cat, dog, horse, and truck.

We first demonstrate in Fig. 10 that sigmoidal networks show the same learning of functions of increasing complexity discussed in Sec. VA for CIFAR10 when trained on MNIST or fashion-MNIST. Note that for CIFAR10, in particular, we see the effects of overtraining set in after several epochs, when the generalization error starts to increase again (we use plain SGD without any explicit regularization in these experiments).

We also repeat these experiments for ReLU networks with activation function $g(x) = \max(0, x)$. While the

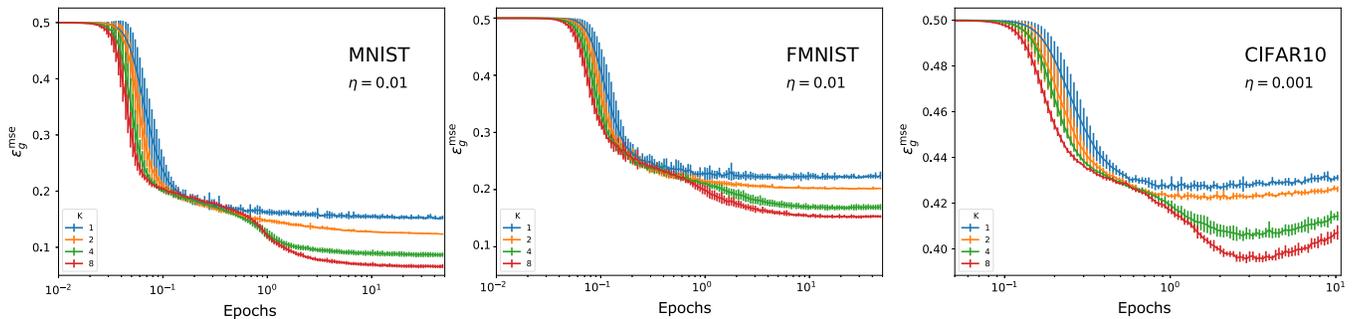


FIG. 10. Two-layer sigmoidal neural networks learn functions of increasing complexity on different datasets. We plot the mean-squared error as a function of training time for sigmoidal networks with an increasing hidden layer when trained on three different datasets. The curves are obtained by averaging ten different runs, starting from different initial weights. Error bars indicate one standard deviation. For all plots, $g(x) = \text{erf}(x/\sqrt{2})$, Gaussian initial weights with standard deviation 10^{-3} , and batch size 32.

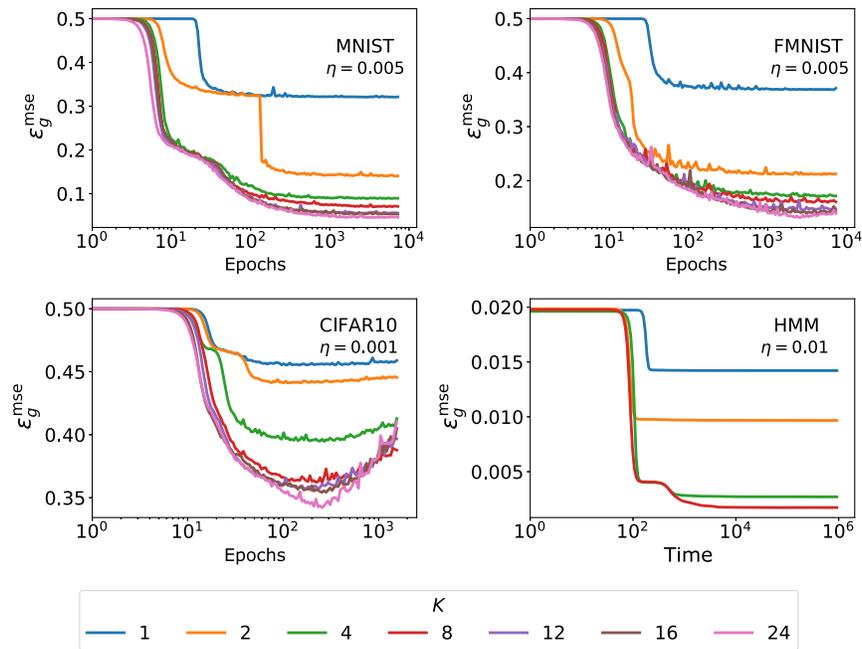


FIG. 11. Two-layer ReLU neural networks also learn functions of increasing complexity and show more run-to-run variance. We plot the mean-squared error as a function of training time for sigmoidal networks with an increasing number of nodes K when trained on three different datasets. For all plots, $g(x) = \max(0, x)$, Gaussian initial weights with standard deviation 10^{-3} , and batch size 32. For online learning, we choose a teacher with $\tilde{g}(x) = \max(x, 0)$, $M = 10$, and $\tilde{v}^m = 1/M$.

dynamics of ReLU students also show a progression from simple to more complex classifiers, the run-to-run fluctuations are much larger than for the sigmoidal students. This result is true both quantitatively, but also qualitatively: For example, networks sometimes get stuck in really suboptimal minimizers for a long time. Hence, plotting the mean trajectories is not as informative, as the standard variations would be very high, so in Fig. 11 we instead show representative curves for individual runs of ReLU students for all three datasets and for online learning from a teacher with $\tilde{g}(x) = \max(x, 0)$, $M = 10$, and $\tilde{v}^m = 1/M$.

[1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, *Imagenet Classification with Deep Convolutional Neural Networks*, in *Advances in Neural Information Processing Systems* (2012), pp. 1097–1105.
 [2] Y. LeCun, Y. Bengio, and G. E. Hinton, *Deep Learning*, *Nature (London)* **521**, 436 (2015).
 [3] I. Sutskever, O. Vinyals, and Q. V. Le, *Sequence to Sequence Learning with Neural Networks*, in *Advances in Neural Information Processing Systems 27*, edited by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Curran Associates, Inc., 2014), pp. 3104–3112.
 [4] V. Vapnik, *The Nature of Statistical Learning Theory* (Springer Science, New York, 2013).

[5] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of Machine Learning* (MIT, Cambridge, MA, 2012).
 [6] E. Gardner and B. Derrida, *Three Unfinished Works on the Optimal Storage Capacity of Networks*, *J. Phys. A* **22**, 1983 (1989).
 [7] H. S. Seung, H. Sompolinsky, and N. Tishby, *Statistical Mechanics of Learning from Examples*, *Phys. Rev. A* **45**, 6056 (1992).
 [8] A. Engel and C. Van den Broeck, *Statistical Mechanics of Learning* (Cambridge University Press, Cambridge, England, 2001).
 [9] L. Zdeborová and F. Krzakala, *Statistical Physics of Inference: Thresholds and Algorithms*, *Adv. Phys.* **65**, 453 (2016).
 [10] Y. LeCun and C. Cortes, *The MNIST database of handwritten digits*, 1998.
 [11] A. Krizhevsky, G. Hinton *et al.*, *Learning Multiple Layers of Features from Tiny Images*, <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>.
 [12] P. Grassberger and I. Procaccia, *Measuring the Strangeness of Strange Attractors*, *Physica D (Amsterdam)* **9D**, 189 (1983).
 [13] J. A. Costa and A. O. Hero, *Learning Intrinsic Dimension and Intrinsic Entropy of High-Dimensional Datasets*, in *Proceedings of the 12th European Signal Processing Conference, 2004* (2004), pp. 369–372.
 [14] E. Levina and P. J. Bickel, *Maximum Likelihood Estimation of Intrinsic Dimension*, in *Advances in Neural Information Processing Systems 17* (2004).
 [15] S. Spigler, M. Geiger, and M. Wyart, *Asymptotic Learning Curves of Kernel Methods: Empirical Data vs. Teacher-Student Paradigm*, [arXiv:1905.10843](https://arxiv.org/abs/1905.10843).

- [16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, *Generative Adversarial Nets*, in *Advances in Neural Information Processing Systems* (2014), pp. 2672–2680.
- [17] A. Radford, L. Metz, and S. Chintala, *Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks*, [arXiv:1511.06434](https://arxiv.org/abs/1511.06434).
- [18] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, *Analyzing and Improving the Image Quality of Stylegan*, [arXiv:1912.04958](https://arxiv.org/abs/1912.04958).
- [19] Not to be confused with the hidden Markov models that are also commonly abbreviated as HMM but which are not used in the present paper.
- [20] W. Kinzel and P. Ruján, *Improving a Network Generalization Ability by Selecting Examples*, *Europhys. Lett.* **13**, 473 (1990).
- [21] M. Biehl and H. Schwarze, *Learning by On-Line Gradient Descent*, *J. Phys. A* **28**, 643 (1995).
- [22] D. Saad and S. A. Solla, *Exact Solution for On-Line Learning in Multilayer Neural Networks*, *Phys. Rev. Lett.* **74**, 4337 (1995).
- [23] D. Saad and S. A. Solla, *On-Line Learning in Soft Commitment Machines*, *Phys. Rev. E* **52**, 4225 (1995).
- [24] 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).
- [25] B. Derrida, E. Gardner, and A. Zippelius, *An Exactly Solvable Asymmetric Neural Network Model*, *Europhys. Lett.* **4**, 167 (1987).
- [26] J. Bruna and S. Mallat, *Invariant Scattering Convolution Networks*, *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 1872 (2013).
- [27] A. B. Patel, M. T. Nguyen, and R. Baraniuk, *A Probabilistic Framework for Deep Learning*, in *Advances in Neural Information Processing Systems 29*, edited by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., 2016), pp. 2558–2566.
- [28] M. Mézard, *Mean-Field Message-Passing Equations in the Hopfield Model and Its Generalizations*, *Phys. Rev. E* **95**, 022117 (2017).
- [29] M. Gabrié, A. Manoel, C. Luneau, J. Barbier, N. Macris, F. Krzakala, and L. Zdeborová, *Entropy and Mutual Information in Models of Deep Neural Networks*, in *Advances in Neural Information Processing Systems 31* (2018), pp. 1826–1836.
- [30] E. Mossel, *Deep Learning and Hierarchical Generative Models*, [arXiv:1612.09057](https://arxiv.org/abs/1612.09057).
- [31] S. Chung, D. D. Lee, and H. Sompolinsky, *Classification and Geometry of General Perceptual Manifolds*, *Phys. Rev. X* **8**, 031003 (2018).
- [32] S. Y. Chung, U. Cohen, H. Sompolinsky, and D. D. Lee, *Learning Data Manifolds with a Cutting Plane Method*, *Neural Comput.* **30**, 2593 (2018).
- [33] U. Cohen, S. Y. Chung, D. D. Lee, and H. Sompolinsky, *Separability and Geometry of Object Manifolds in Deep Neural Networks*, *Nat. Commun.* **11**, 746 (2020).
- [34] P. Rotondo, M. C. Lagomarsino, and M. Gherardi, *Counting the Learnable Functions of Structured Data*, *Phys. Rev. Research* **2**, 023169 (2020).
- [35] T. M. Cover, *Geometrical and Statistical Properties of Systems of Linear Inequalities with Applications in Pattern Recognition*, *IEEE Trans. Electron. Comput.* **EC-14**, 326 (1965).
- [36] Y. Yoshida and M. Okada, *Data-Dependence of Plateau Phenomenon in Learning with Neural Network—Statistical Mechanical Analysis*, in *Advances in Neural Information Processing Systems 32* (2019), pp. 1720–1728.
- [37] A. M. Saxe, J. L. McClelland, and S. Ganguli, *Exact Solutions to the Nonlinear Dynamics of Learning in Deep Linear Neural Networks*, in *ICLR* (2014).
- [38] F. Rosenblatt, *Principles of Neurodynamics* (Spartan, 1962).
- [39] A. Rahimi and B. Recht, *Random Features for Large-Scale Kernel Machines*, in *Adv. Neural Inf. Process. Syst.* (2008), pp. 1177–1184.
- [40] A. Rahimi and B. Recht, *Weighted Sums of Random Kitchen Sinks: Replacing Minimization with Randomization in Learning*, in *Advances in Neural Information Processing Systems* (2009), pp. 1313–1320.
- [41] C. Louart, Z. Liao, and R. Couillet, *A Random Matrix Approach to Neural Networks*, *Ann. Appl. Probab.* **28**, 1190 (2018).
- [42] 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).
- [43] A. Montanari, F. Ruan, Y. Sohn, and J. Yan, *The Generalization Error of Max-Margin Linear Classifiers: High-Dimensional Asymptotics in the Overparametrized Regime*, [arXiv:1911.01544](https://arxiv.org/abs/1911.01544).
- [44] W. Hachem, P. Loubaton, and J. Najim, *Deterministic Equivalents for Certain Functionals of Large Random Matrices*, *Ann. Appl. Probab.* **17**, 875 (2007).
- [45] X. Cheng and A. Singer, *The Spectrum of Random Inner-Product Kernel Matrices*, *Random Matrices Theory Appl.* **02**, 1350010 (2013).
- [46] Z. Fan and A. Montanari, *The Spectral Norm of Random Inner-Product Kernel Matrices*, *Probab. Theory Relat. Fields* **173**, 27 (2019).
- [47] J. Pennington and P. Worah, *Nonlinear Random Matrix Theory for Deep Learning*, in *Advances in Neural Information Processing Systems* (2017), pp. 2637–2646.
- [48] M. E. A. Seddik, M. Tamaazousti, and R. Couillet, *Kernel Random Matrices of Large Concentrated Data: The Example of GAN-Generated Images*, in *Proceedings of the 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (IEEE, New York, 2019), pp. 7480–7484.
- [49] M. E. A. Seddik, C. Louart, M. Tamaazousti, and R. Couillet, *Random Matrix Theory Proves That Deep Learning Representations of GAN-Data Behave as Gaussian Mixtures*, [arXiv:2001.08370](https://arxiv.org/abs/2001.08370).
- [50] An ODE integrator and source code for all experiments can be found at <https://github.com/sgoldt/hidden-manifold-model>.
- [51] T. L. H. Watkin, A. Rau, and M. Biehl, *The Statistical Mechanics of Learning a Rule*, *Rev. Mod. Phys.* **65**, 499 (1993).
- [52] M. S. Advani and A. M. Saxe, *High-Dimensional Dynamics of Generalization Error in Neural Networks*, [arXiv:1710.03667](https://arxiv.org/abs/1710.03667).

- [53] B. Aubin, A. Maillard, J. Barbier, F. Krzakala, N. Macris, and L. Zdeborová, The Committee Machine: Computational to Statistical Gaps in Learning a Two-Layers Neural Network, in *Advances in Neural Information Processing Systems 31* (2018), pp. 3227–3238.
- [54] J. Barbier, F. Krzakala, N. Macris, L. Miolane, and L. Zdeborová, *Optimal Errors and Phase Transitions in High-Dimensional Generalized Linear Models*, *Proc. Natl. Acad. Sci. U.S.A.* **116**, 5451 (2019).
- [55] S. Goldt, M. S. Advani, A. M. Saxe, F. Krzakala, and L. Zdeborová, *Dynamics of Stochastic Gradient Descent for Two-Layer Neural Networks in the Teacher-Student Setup*, in *Advances in Neural Information Processing Systems 32* (2019).
- [56] Y. Yoshida, R. Karakida, M. Okada, and S.-I. Amari, *Statistical Mechanical Analysis of Learning Dynamics of Two-Layer Perceptron with Multiple Output Units*, *J. Phys. A* **52**, 184002 (2019).
- [57] R. Ge, J. D. Lee, and T. Ma, *Learning One-Hidden-Layer Neural Networks with Landscape Design*, [arXiv:1711.00501](https://arxiv.org/abs/1711.00501).
- [58] Y. Li and Y. Yuan, *Convergence Analysis of Two-Layer Neural Networks with Relu Activation*, in *Advances in Neural Information Processing Systems* (2017), pp. 597–607.
- [59] S. Arora, N. Cohen, W. Hu, and Y. Luo, *Implicit Regularization in Deep Matrix Factorization*, in *Advances in Neural Information Processing Systems 33* (2019).
- [60] D. P. Kingma and M. Welling, *Auto-encoding Variational Bayes*, [arXiv:1312.6114](https://arxiv.org/abs/1312.6114).
- [61] D. Michelsanti and Z. Tan, *Conditional Generative Adversarial Networks for Speech Enhancement and Noise-Robust Speaker Verification*, in *Proceedings of Interspeech 2017* (2017), pp. 2008–2012.
- [62] Q. Le, T. Sarló, and A. Smola, *Fastfood-Approximating Kernel Expansions in Loglinear Time*, in *Proceedings of the International Conference on Machine Learning* (2013), Vol. 85.
- [63] M. Moczulski, M. Denil, J. Appleyard, and N. d. Freitas, *ACDC: A Structured Efficient Linear Layer*, in *International Conference on Learning Representations (ICLR)* (2016).
- [64] F. X. Yu, A. T. Suresh, K. M. Choromanski, D. N. Holtmann-Rice, and S. Kumar, *Orthogonal Random Features*, in *Adv. Neural Inf. Process. Syst.* (2016), pp. 1975–1983.
- [65] D. Saad, *On-Line Learning in Neural Networks* (Cambridge University Press, Cambridge, England, 2009), Vol. 17.
- [66] V. A. Marchenko and L. A. Pastur, *Distribution of Eigenvalues for Some Sets of Random Matrices*, *Mat. Sb.* **114**, 507 (1967).
- [67] B. Babadi and H. Sompolinsky, *Sparseness and Expansion in Sensory Representations*, *Neuron* **83**, 1213 (2014).
- [68] J. Kadmon and H. Sompolinsky, *Optimal Architectures in a Solvable Model of Deep Networks*, in *Adv. Neural Inf. Process. Syst.* (2016), pp. 4781–4789.
- [69] P. Riegler and M. Biehl, *On-Line Backpropagation in Two-Layered Neural Networks*, *J. Phys. A* **28**, L507 (1995).
- [70] J. Hadamard, *Resolution d'une Question Relative aux Determinants*, *Bull. Sci. Math.* **2**, 240 (1893).
- [71] D. E. Muller, *Application of Boolean Algebra to Switching Circuit Design and to Error Detection*, *Trans. IRE Profess. Group Electron. Comput.* **3**, 6 (1954).
- [72] I. S. Reed, *A Class of Multiple-Error-Correcting Codes and the Decoding Scheme*, Massachusetts Institute of Technology, Lexington Lincoln Lab, technical report, 1953.
- [73] M. Soltanolkotabi, A. Javanmard, and J. D. Lee, *Theoretical Insights into the Optimization Landscape of Over-parameterized Shallow Neural Networks*, *IEEE Trans. Inf. Theory* **65**, 742 (2018).
- [74] S. Mei, A. Montanari, and P. Nguyen, *A Mean Field View of the Landscape of Two-Layer Neural Networks*, *Proc. Natl. Acad. Sci. U.S.A.* **115**, E7665 (2018).
- [75] G. M. Rotskoff and E. Vanden-Eijnden, *Parameters as Interacting Particles: Long Time Convergence and Asymptotic Error Scaling of Neural Networks*, in *Advances in Neural Information Processing Systems 31* (2018), pp. 7146–7155.
- [76] L. Chizat and F. Bach, *On the Global Convergence of Gradient Descent for Over-parameterized Models Using Optimal Transport*, in *Advances in Neural Information Processing Systems 31* (2018), pp. 3040–3050.
- [77] J. Sirignano and K. Spiliopoulos, *Mean Field Analysis of Neural Networks: A Central Limit Theorem*, *Stoch. Proc. Appl.* **130**, 1820 (2020).
- [78] D. Saad and S. A. Solla, *Learning with Noise and Regularizers Multilayer Neural Networks*, in *Advances in Neural Information Processing Systems 9* (1997), pp. 260–266.
- [79] M. Biehl, P. Riegler, and C. Wöhler, *Transient Dynamics of On-Line Learning in Two-Layered Neural Networks*, *J. Phys. A* **29**, 4769 (1996).
- [80] M. Rattray, D. Saad, and S.-I. Amari, *Natural Gradient Descent for On-Line Learning*, *Phys. Rev. Lett.* **81**, 5461 (1998).
- [81] Y. N. Dauphin, R. Pascanu, G. Gulcehre, K. Cho, S. Ganguli, and Y. Bengio, *Identifying and Attacking the Saddle Point Problem in High-Dimensional Non-convex Optimization*, in *Adv. Neural Inf. Process. Syst.* (2014), pp. 2933–2941.
- [82] D. Kalimeris, G. Kaplun, P. Nakkiran, B. Edelman, T. Yang, B. Barak, and H. Zhang, *SGD on Neural Networks Learns Functions of Increasing Complexity*, in *Advances in Neural Information Processing Systems 32* (2019), pp. 3496–3506.
- [83] V. Vapnik, *Statistical Learning Theory* (Springer, New York, 1998), pp. 156–160.
- [84] F. Farnia, J. Zhang, and D. Tse, *A Spectral Approach to Generalization and Optimization in Neural Networks*, in *ICLR* (2018).
- [85] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. A. Hamprecht, Y. Bengio, and A. Courville, *On the Spectral Bias of Neural Networks*, in *Proceedings of the 36th International Conference on Machine Learning* (2019).
- [86] D. Arpit and S. Jastrzębski, M. S. Kanwal, T. Maharaj, A. Fischer, A. Courville, and Y. Bengio, *A Closer Look at Memorization in Deep Networks*, in *Proceedings of the 34th International Conference on Machine Learning* (2017).
- [87] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, *Understanding Deep Learning Requires Rethinking Generalization*, in *ICLR* (2017).

- [88] A. C. C. Coolen, D. Saad, and Y.-S. Xiong, *On-Line Learning from Restricted Training Sets in Multilayer Neural Networks*, *Europhys. Lett.* **51**, 691 (2000).
- [89] A. C. C. Coolen and D. Saad, *Dynamics of Learning with Restricted Training Sets*, *Phys. Rev. E* **62**, 5444 (2000).
- [90] F. Gerace, B. Loureiro, F. Krzakala, M. Mézard, and L. Zdeborová, *Generalisation Error in Learning with Random Features and the Hidden Manifold Model*, in *Proceedings of the 37th International Conference on Machine Learning* (2020).
- [91] F. Mignacco, F. Krzakala, Y.M. Lu, and L. Zdeborová, *The role of Regularization in Classification of High-Dimensional Noisy Gaussian Mixture*, in *Proceedings of the 37th International Conference on Machine Learning* (2020).