


**Universal mean-field upper bound for the generalization gap of deep neural networks**S. Ariosto,<sup>1,2</sup> R. Pacelli,<sup>3</sup> F. Ginelli,<sup>1,2</sup> M. Gherardi,<sup>4,2</sup> and P. Rotondo <sup>4,2</sup><sup>1</sup>*Dipartimento di Scienza e Alta Tecnologia and Center for Nonlinear and Complex Systems, Università degli Studi dell'Insubria, Via Valleggio 11, 22100 Como, Italy*<sup>2</sup>*I.N.F.N. Sezione di Milano, Via Celoria 16, 20133 Milan, Italy*<sup>3</sup>*Dipartimento di Scienza Applicata e Tecnologia, Politecnico di Torino, 10129 Turin, Italy*<sup>4</sup>*Università degli Studi di Milano, Via Celoria 16, 20133 Milan, Italy*

(Received 17 February 2022; accepted 8 June 2022; published 29 June 2022)

Modern deep neural networks (DNNs) represent a formidable challenge for theorists: according to the commonly accepted probabilistic framework that describes their performance, these architectures should overfit due to the huge number of parameters to train, but in practice they do not. Here we employ results from replica mean field theory to compute the generalization gap of machine learning models with quenched features, in the teacher-student scenario and for regression problems with quadratic loss function. Notably, this framework includes the case of DNNs where the last layer is optimized given a specific realization of the remaining weights. We show how these results—combined with ideas from statistical learning theory—provide a stringent asymptotic upper bound on the generalization gap of fully trained DNN as a function of the size of the dataset  $P$ . In particular, in the limit of large  $P$  and  $N_{\text{out}}$  (where  $N_{\text{out}}$  is the size of the last layer) and  $N_{\text{out}} \ll P$ , the generalization gap approaches zero faster than  $2N_{\text{out}}/P$ , for any choice of both architecture and teacher function. Notably, this result greatly improves existing bounds from statistical learning theory. We test our predictions on a broad range of architectures, from toy fully connected neural networks with few hidden layers to state-of-the-art deep convolutional neural networks.

DOI: [10.1103/PhysRevE.105.064309](https://doi.org/10.1103/PhysRevE.105.064309)**I. INTRODUCTION**

In the last ten years deep neural networks (DNNs) [1] have revolutionized the field of machine learning, outperforming traditional methods in tasks that include image classification, speech recognition, and time series prediction. Despite the enormous success in applications, the size of these architectures represents a puzzle for theorists. When AlexNet won the ImageNet competition in 2012 [2], it had roughly 60 million parameters. In the following years, another DNN, the Visual Geometry Group (VGG) architecture delivered the state-of-the-art performance with more than 138 million parameters [3]. Nowadays, convolutional DNNs such as ResNet or Inception work by training about 10 million weights [4,5]. According to common intuition, models with such a high number of degrees of freedom should overfit the training data and perform poorly on previously unseen data samples. Statistical learning theory (SLT) is regarded as the established probabilistic framework to quantify the generalization performance in machine learning [6,7]. In particular, it provides uniform bounds on performance degradation when generalizing to previously unseen data, based on capability or complexity measures such as the Vapnik-Chervonenkis dimension [8] or the Rademacher complexity [9]. Unfortunately, these bounds do not provide any guarantee that the severely overparametrized state-of-the-art models should have any predictive power on test data.

Overcoming this conceptual puzzle engages computer scientists, mathematicians, and physicists alike [10–12]. In

Ref. [13] the authors provide a mean field view of the stochastic gradient dynamics of one-hidden layer networks by using the theory of gradient flows in Wasserstein spaces [14]. Similar results for one-hidden layer networks can be found in Refs. [15–17]. Unfortunately, it is challenging to extend this approach to deeper networks, although for a recent attempt in a simplified setting one may check Ref. [18]. Other groups are studying the role of overparametrization and related phenomena such as the double descent in the regime of lazy training [19–25]. Also the statistical physics of kernel learning (originally started in Ref. [26]) has undergone a revival in the last few years [27], mainly due to the discovery of the Neural Tangent Kernel (NTK) limit of deep neural networks—a mathematical equivalence between neural networks and a certain kernel that arises in the limit of large layer size [28,29]. Despite all these major conceptual advances in the field, it is fair to say that a unified framework to investigate and understand the generalization performance of DNNs is still missing.

On the mathematical side, it is instructive to rationalize why the theorems proven in the framework of statistical learning theory often yield very loose bounds when applied to practical problems (as brilliantly put forward in Ref. [30] by Bottou or in the recent review [31] by Belkin). The goal of theorems in SLT is to provide *distribution-independent uniform* bounds on the deviation between the generalization and training errors, which are mostly based on the uniform law of large numbers and on techniques for capacity control [31]. The formulation and the derivation of these theorems

reveal a source of possible reasons for their poor quantitative performance: (1) empirically relevant data distributions may lead to smaller typical deviations than the worst possible case [32–36] and (2) uniform bounds hold for all possible functions in the model, but better bounds may hold when one restricts the analysis to functions that perform well on specific (and significant) training sets.

In this paper we build upon these considerations to develop a mean field theory for the generalization gap (GG) of deep neural networks. First, we employ nonrigorous but standard statistical physics tools of disordered systems [37], together with a key Gaussian approximation on the statistics of the features known as Gaussian Equivalence Principle [20], to compute the generalization and training errors of machine learning models with *quenched features*, obtaining simple formulas in the regime of large training dataset size  $P$ . In particular these results hold in the teacher-student scenario for a broad class of input-output distributions when the employed loss function is the mean squared error. Notably, this setup includes the case of DNNs where the  $N_{\text{out}}$  weights in the last layer are optimized given any specific instance of the remaining weights. Analogous results have been derived in the recent literature [25,27,38,39]; here we show how to employ them to derive a universal mean field upper bound for the generalization gap of fully trained DNNs. In the limit  $N_{\text{out}} \ll P$  (a condition satisfied by most state-of-the-art DNNs, even in the overparametrized regime where  $P$  is small compared to the total number of weights), a simple asymptotic upper bound emerges, according to which the gap should approach zero faster than  $2N_{\text{out}}/P$ . This is our central result, which in the large  $P$  limit greatly improves existing estimates. Finally, we check the validity of our mean field bound against several synthetic and empirical data distributions and across a variety of different architectures, ranging from toy DNNs with few fully connected layers to state-of-the-art ones employed for challenging computer vision problems.

Although these results lack the mathematical rigor of formal theorems, they take a concrete step towards understanding why overparametrized DNNs work in practice and may guide us to the formulation of more informed and accurate bounds for the generalization gap of modern DNNs.

## II. GENERALIZATION GAP OF QUENCHED FEATURE MODELS

We start by briefly describing the setting of the supervised learning problem that we will study throughout the paper. Let us consider a training set  $\mathcal{T}$  made of  $P$  independent identically distributed (IID) random observations,  $\mathcal{T} = \{(\mathbf{x}^\mu, y^\mu)\}_{\mu=1}^P$ , where the  $\mathbf{x}^\mu$ 's are  $D$ -dimensional vectors drawn by an input probability distribution  $\rho(\mathbf{x})$  and the  $y^\mu$ 's are scalar outputs provided by a real-valued teacher function  $f_T$ , i.e.,  $y^\mu = f_T(\mathbf{x}^\mu)$ . Under these assumptions, the joint input-output probability distribution  $\rho_{I/O}(\mathbf{x}, y)$  is given by  $\rho_{I/O}(\mathbf{x}, y) = \rho(\mathbf{x})\delta(y - f_T(\mathbf{x}))$ .

Our first goal is to compute the generalization performance of a model (the so-called *student*) of the

following form:

$$f_S(\mathbf{x}) = \sum_{\alpha=1}^N v_\alpha \phi_\alpha(\mathbf{x}) = \mathbf{v} \cdot \phi(\mathbf{x}), \quad (1)$$

where  $\phi$  is an  $N$ -dimensional feature map and  $\mathbf{v}$  is an  $N$ -dimensional vector of real weights to be optimized. The average generalization and training errors are defined as

$$\epsilon_g = \left\langle \int d^D x \rho(\mathbf{x}) [f_T(\mathbf{x}) - f_S^*(\mathbf{x})]^2 \right\rangle_{\mathcal{T}}, \quad (2)$$

$$\epsilon_t = \left\langle \frac{1}{P} \sum_{\mu=1}^P [f_T(\mathbf{x}^\mu) - f_S^*(\mathbf{x}^\mu)]^2 \right\rangle_{\mathcal{T}}, \quad (3)$$

where  $\langle \cdot \rangle_{\mathcal{T}}$  indicates the average over all the possible realizations of a training set of size  $P$ , and the optimized function  $f_S^*$  corresponds to the choice of the vector  $\mathbf{v}^*$  that minimizes the quadratic loss  $P^{-1} \sum_{\mu=1}^P [f_T(\mathbf{x}^\mu) - f_S(\mathbf{x}^\mu)]^2$  for a given instance of the dataset  $\mathcal{T}$ . Although in principle this approach can be developed for arbitrary loss functions [25], here we will consider only quadratic loss and regression problems, which are considerably simpler to deal with analytically [27,40,41].

The generalization power of a machine learning model can be measured by its *generalization gap*,

$$\Delta\epsilon = \epsilon_g - \epsilon_t, \quad (4)$$

which expresses the average performance difference of a trained model between its training dataset and unseen data drawn from the same distribution. Crucially, it is possible to express the generalization and training errors as a function of the features; as we will discuss in the following, this ingredient is fundamental to provide insight on the generalization gap of fully trained DNNs. Here the calculation of the average generalization and training errors is performed using the well-known replica method, a standard statistical physics technique developed to study disordered systems [37]. Optimization is addressed introducing an effective Hamiltonian given by the sum of the training loss with a regularization term,

$$\mathcal{L} = \frac{1}{2} \sum_{\mu=1}^P [f_T(\mathbf{x}^\mu) - f_S(\mathbf{x}^\mu)]^2 + \frac{\lambda}{2} \sum_{\alpha=1}^N (v_\alpha)^2, \quad (5)$$

gauged by the regularization parameter  $\lambda > 0$ . Given the convexity of our problem, we can make use of the replica symmetric ansatz, which is known to deliver the correct result for convex optimization when  $P \gg 1$  [26]. The ground state of the effective Hamiltonian, given by the optimized  $f_S^*$ , is finally evaluated from the replicated partition function in the large  $P$  limit via the saddle-point method. This approach is rather similar to that developed for the random features model (RFM) [42] and for kernel regression [27], and it is discussed in detail in the Appendixes. In order to derive a closed expression for the generalization and training errors, one has to assume the features to be Gaussian distributed, at least for the purpose of the present calculation (further details are found in Appendix A). While this assumption, a sort of generalized central limit theorem, has been proven only for specific types of feature maps [13], empirical observations suggest that it holds in much broader contexts [25].

Under this mean-field-like approximation, also known as Gaussian Equivalence Principle [20], the analytical expressions for the generalization and training errors depend only on second-order moments, given by the following integrals over the input probability distribution:

$$\begin{aligned} J_\alpha &= \int d^D x \rho(\mathbf{x}) f_T(\mathbf{x}) \phi_\alpha(\mathbf{x}), \\ \Phi_{\alpha\beta} &= \int d^D x \rho(\mathbf{x}) \phi_\alpha(\mathbf{x}) \phi_\beta(\mathbf{x}), \\ T &= \int d^D x \rho(\mathbf{x}) f_T^2(\mathbf{x}), \end{aligned} \quad (6)$$

with  $\alpha, \beta = 1, \dots, N$ . The vector  $\mathbf{J}$  and the matrix  $\Phi$  depend on the specific choice of the feature map  $\phi$  and respectively represent a teacher-feature and a feature-feature overlap, whereas the scalar quantity  $T$  is by definition the trivial predictor [43] of the regression problem, and it provides a natural scale to compare different learning problems. Using these definitions, we find the following compact representation for the generalization and training errors (valid in the thermodynamic limit of large  $D, N, P$  as discussed in Appendix A):

$$\begin{aligned} \epsilon_g^R &= \frac{\epsilon_g^R + (\kappa\lambda)^2 \mathbf{J}^T \Phi^{-1} \mathbf{G}^{-2} \mathbf{J}}{1 - P \text{Tr}(\Phi^2 \mathbf{G}^{-2})}, \\ \epsilon_t &= \frac{\epsilon_g^R}{\kappa} + \frac{\epsilon_g}{\kappa} \left( \frac{\kappa - 1}{\kappa} - \frac{N}{P} \right), \end{aligned} \quad (7)$$

where the matrix  $\mathbf{G} = \kappa\lambda\mathbb{1} + P\Phi$  is invertible and the variable  $\kappa$  is self-consistently defined via the following equation:

$$\kappa = 1 + \kappa \text{Tr}(\Phi \mathbf{G}^{-1}). \quad (8)$$

The residual generalization error  $\epsilon_g^R$  corresponds to the best possible performance of the quenched model on the dataset, under the assumption of full knowledge of the input-output probability distribution, and it is given by

$$\epsilon_g^R = T - \mathbf{J}^T \Phi^{-1} \mathbf{J}. \quad (9)$$

It is worth noticing that for strictly infinite size of the dataset  $P$ ,  $\kappa \rightarrow 1$  and it is easy to prove that  $\epsilon_g \rightarrow \epsilon_g^R$  and  $\epsilon_t \rightarrow \epsilon_g$ , which provide a first consistency check of the validity of the mean field theory. Additionally, the self-consistent definition of  $\kappa$  and the way it enters in the expression for the generalization error are the same as in the recent work on kernel regression by Pehlevan's group [27]. This should not come as a surprise, since one could specialize the general quenched features that we employ here to the case of polynomial, Gaussian or NTK kernels. For these particular choices of the quenched features, Eq. (7) just provides a different representation of the generalization error formula given in Ref. [27]. A generalization of Eq. (7) has been recently proved in Ref. [25].

Starting from Eq. (7) it is possible to perform an asymptotic analysis for large size of the training set  $P$ . This is particularly simple if we assume that both  $P$  and  $N$  are large,  $N \ll P$  and the regularization parameter  $\lambda$  is finite. In this case we easily obtain that  $\kappa \sim 1 + N/P$  and  $\mathbf{G} \sim P\Phi$ . Using these asymptotic expressions, and normalizing by the natural scale of the problem—i.e., the trivial predictor  $T$  defined in Eq. (6)—we

can compute the normalized generalization gap:

$$\Delta\tilde{\epsilon} \equiv \frac{\Delta\epsilon}{T} \simeq 2 \frac{\epsilon_g^R}{T} \frac{N}{P}. \quad (10)$$

Note that Eq. (10) does not necessarily imply a linear scaling with  $N$  since  $\epsilon_g^R$  may still retain a dependence on  $N$  [as implied from Eq. (9)]. Moreover, it is worth stressing that the derivation of Eq. (10) assumes that all but a finite number of the  $N$  eigenvalues of the covariance matrix  $\Phi$  are different from zero.

### III. GENERALIZATION GAP OF FULLY TRAINED DNNs

Let us now suppose that the quenched features of the model under consideration are given by a DNN. For instance, in the special case of a fully connected architecture with one hidden layer, we have that the function implemented is  $f_{\text{IHL}}(\mathbf{x}) = \sum_{\alpha=1}^N v_\alpha \sigma(W_\alpha \cdot \mathbf{x})$ , where for simplicity we have set all the biases to zero. The  $W_\alpha$ 's are the  $D$ -dimensional vector weights of the hidden layer and  $\sigma$  is a generic well-behaved activation function. Here the quenched features are given by  $\phi_\alpha^{\text{IHL}}(\mathbf{x}, W) = \sigma(W_\alpha \cdot \mathbf{x})$ . More in general, let us consider a DNN with a fully connected last layer. The specific architecture of the first layers is unimportant. This class includes all the relevant state-of-the-art architectures. Let us fix the dimension of the last layer to  $N = N_{\text{out}}$ , and let us split the weights  $\vartheta$  of the network as  $\vartheta = \{\mathbf{v}, \mathcal{W}\}$  where  $\mathbf{v}$  is the vector of  $N_{\text{out}}$ -dimensional weights of the last fully connected layer, whereas  $\mathcal{W}$  is a short notation for all the remaining weights of the DNN. We introduce the feature map notation  $\phi_\alpha^{\text{DNN}}(\mathbf{x}, \mathcal{W})$  to indicate the corresponding quenched features.

We now reconsider the results provided by Eq. (7) when specialized to the quenched features of a DNN. The crucial observation is that this mean field theory provides the average generalization and training errors for *each* realization of the weights  $\mathcal{W}$ . In other words, given a specific configuration  $\bar{\mathcal{W}}$ , our theory predicts the corresponding average generalization and training errors, supposing that the weights  $\mathbf{v}$  of the last layer are set to the optimal value that minimizes the training loss at fixed  $\bar{\mathcal{W}}$ . From now on, we use the notation  $\epsilon_g(\mathcal{W})$ ,  $\epsilon_g^R(\mathcal{W})$ ,  $\epsilon_t(\mathcal{W})$  to stress that the generalization and training errors depend on  $\mathcal{W}$  via the teacher-feature  $\mathbf{J}$  and feature-feature  $\Phi$  overlaps.

Therefore, the result in Eq. (10) holds for each given realization of the weights  $\mathcal{W}$  of the DNN if we assume perfect training over the last layer. See also the recent conjecture put forward in Ref. [25]. In particular, this equivalence holds for a fully trained configuration  $\theta^* \equiv \{v^*, \mathcal{W}^*\}$  that is a *local* minimum of the loss. Unfortunately, such local minimum may depend on the size  $P$  of the training set, and so it does  $\epsilon_g^R(\mathcal{W})$ . However, since the residual generalization error (9) is positive by definition and bounded by  $T$ , it follows that  $0 \leq \epsilon_g^R(\mathcal{W})/T \leq 1$  for every  $\mathcal{W}$ . As such, this provides us with the following asymptotic mean field upper bound for the (normalized) generalization performance of a DNN:

$$\Delta\tilde{\epsilon}(\mathcal{W}) \leq \frac{2N_{\text{out}}}{P}, \quad (11)$$

which is the central result of this paper. It should be compared with the classic SLT result where the generalization

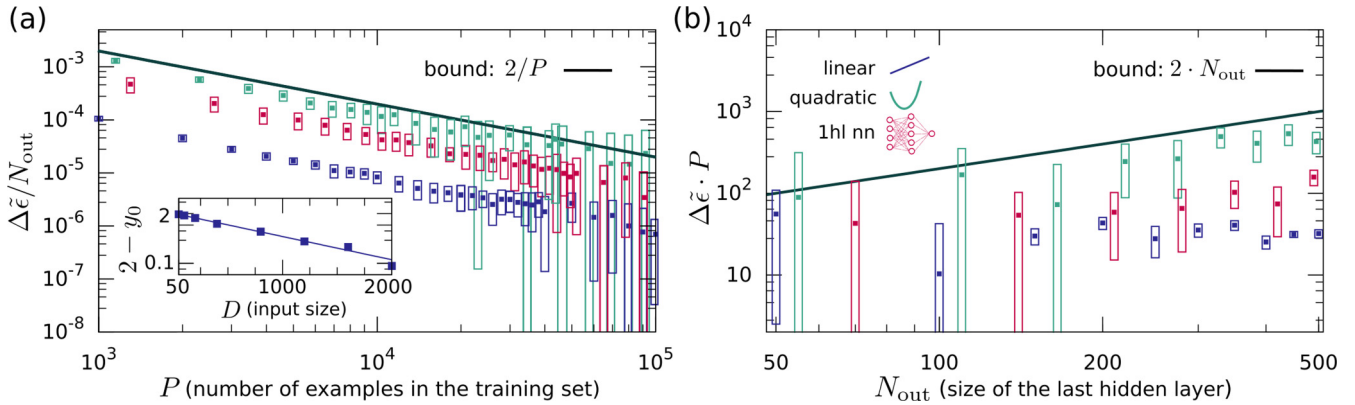


FIG. 1. Generalization gap in the lazy-training regime. The behavior of the normalized GG for a one-hidden layer student architecture is displayed for the three different classes of teacher outlined in the main text: linear (blue symbols), quadratic (green), and one-hidden layer (red). The solid black line marks our mean field upper bound. (a) Normalized GG (rescaled by  $N_{\text{out}}$ ) as a function of the training dataset  $P$ . Data points are the result of an average over 50 different realizations of the teacher and of the input (of dimension  $D = 50$ ) with  $N_{\text{out}} = 400$ . We observe that the functional form of the rescaled GG  $\Delta\bar{\epsilon}$  is compatible with  $y_0/P$  over two decades. In the inset we consider different input dimension for the linear teacher case, showing that by increasing  $D$ , the residual generalization gap converges exponentially fast to the bound, i.e., the prefactor  $y_0$  (which is obtained by a standard fitting procedure) converges exponentially from below to 2 as  $D \rightarrow \infty$  with a rate  $\gamma \approx 0.0014$  (see also Appendix B). (b) The normalized GG (multiplied by  $P$ ) is shown as a function of the size of the hidden layer  $N_{\text{out}}$ . Simulations are performed with  $P$  equal to  $2 \times 10^4$  (linear teacher),  $4 \times 10^4$  (quadratic), and  $8 \times 10^4$  (1HL), and by averaging over 20 different teacher and input realizations ( $D = 50$ ). Error bars in both panels correspond to one standard error. Typical (rescaled by  $T$ ) training errors in the lazy-training regime are of order  $10^{-1}$  and are systematically smaller for the linear teacher.

gap is bounded by  $\sqrt{d_{\text{VC}}/P}$ , where we have introduced the Vapnik-Chervonenkis dimension  $d_{\text{VC}}$ , a compact measure of the architecture capacity. In the case of deep neural networks it is roughly proportional to the total number of parameters  $N_{\text{tot}}$  of the architecture.<sup>1</sup>

## IV. NUMERICAL EXPERIMENTS

### A. Toy DNNs with synthetic datasets

We start by testing our bound on a fully connected architecture with one-hidden layer and ReLU activations. We have chosen three different teacher classes of increasing complexity: (1) a linear function  $f_{\text{T}}(\mathbf{x}) = \mathbf{t} \cdot \mathbf{x}$ , (2) a quadratic polynomial  $f_{\text{T}}(\mathbf{x}) = \mathbf{t} \cdot \mathbf{x} + (\mathbf{t} \cdot \mathbf{x})^2$ , (3) a fully connected one-hidden layer (1HL) architecture with ReLU activations,  $f_{\text{T}}(\mathbf{x}) = \sum_{\alpha=1}^M q_{\alpha} \text{ReLU}(\mathbf{S}_{\alpha} \cdot \mathbf{x})$ . See also Appendix B for details on these architectures and on the input choice.

We first consider DNNs where only the last layer is trained and the remaining weights are kept fixed to their initialization values  $\bar{\mathcal{W}}$ , i.e., we consider the lazy training regime [45]. Since the  $\bar{\mathcal{W}}$ 's do not change during training, the residual generalization error  $\epsilon_{\text{g}}^{\text{R}}(\bar{\mathcal{W}})$  is independent of  $P$ , and one expects not only the bound (11) to hold but also the generalization gap to scale precisely as  $1/P$  for  $P$  large enough. This is verified in Fig. 1(a) for the three different teacher classes introduced above. On the other hand, as already noted,  $\epsilon_{\text{g}}^{\text{R}}(\bar{\mathcal{W}})$  may still retain a dependence on the last layer size  $N_{\text{out}}$ . In particular, one may expect that increasing  $N_{\text{out}}$  will decrease the residual

generalization error, as this increases the number of functions available to approximate the target  $f_{\text{T}}(\mathbf{x})$ . Therefore, for large  $P$  and  $N_{\text{out}}$  we expect  $\Delta\bar{\epsilon}$  to increase *at most* linearly as a function of  $N_{\text{out}}$ . The numerical behavior of the GG as a function of  $N_{\text{out}}$  is shown in Fig. 1(b). Once again, our mean field bound holds, but different scaling with  $N_{\text{out}}$  can be appreciated. In particular, the GG is almost constant for the linear teacher, whereas it has an approximately linear behavior for the quadratic one, reflecting different dependencies of the residual generalization error from the last layer size. Note also that in both panels of Fig. 1 the GG is systematically lower for the linear teacher case, confirming the intuitive expectation that the linear problem should be the easiest to learn.

It is worth remarking that the input dimension  $D$  only enters the theory through the residual generalization error (9). Interestingly, as one increases the input dimension  $D$ , the normalized generalization gap seems to saturate the unit bound with an exponential convergence in  $D$ , as shown in the inset of Fig. 1(a) for the linear teacher case (see more details in Appendix B). Currently we have no theory for this.

We next consider fully trained DNNs. Here the weights  $\mathcal{W}$  are trained and the residual generalization error  $\epsilon_{\text{g}}^{\text{R}}(\mathcal{W})$  may depend on the training set size  $P$ . Suppose for instance that there exists a configuration of the weights  $\mathcal{W}^{\dagger}$  such that the residual generalization error  $\epsilon_{\text{g}}^{\text{R}}(\mathcal{W}^{\dagger}) = 0$ , i.e., the DNN can learn the teacher function perfectly. Intuitively, as the size of the dataset  $P$  grows, we expect that the DNN will be capable of finding configurations  $\mathcal{W}$  that are closer to the optimal one  $\mathcal{W}^{\dagger}$ . This means that the residual generalization error will decrease in some way towards zero as a function of  $P$  for  $P$  large enough, and that the GG will decrease faster than  $1/P$ .

In the complementary case where the DNN is not capable of learning the target function, there will be a nonzero residual

<sup>1</sup>Classic Vapnik-Chervonenkis dimension is defined for classification problems, but generalization to regression (the so-called pseudo-dimension) is straightforward [44].

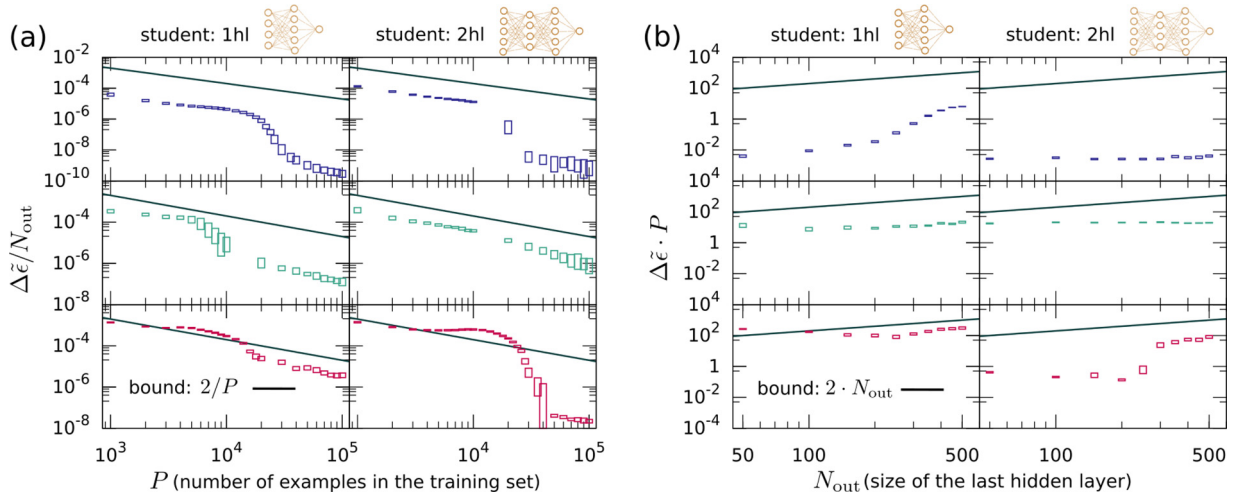


FIG. 2. Generalization gap of fully trained toy DNNs with one [left columns in (a) and (b)] and two (right columns) fully connected hidden layers for the three different teacher classes outlined in the main text. From top to bottom: linear (blue symbols), quadratic (green), and 1HL teacher architectures (red). The solid black line marks the mean field upper bound. (a) Normalized GG (rescaled by  $N_{\text{out}}$  as a function of the size of the training set  $P$ . Data points are the result of an average over 50 realizations of the teacher and of the input ( $D = 50$ ) with  $N_{\text{out}} = 100$ . (b) Normalized GG (multiplied by  $P$  vs the size of the last hidden layer  $N_{\text{out}}$ . Data are averaged over 20 different teacher and input ( $D = 50$ ) realizations with  $P = 4 \times 10^4$ . Error bars in both panels correspond to one standard error. Typical training errors are of the order of  $10^{-6}$  for each teacher-student pair, except in the case 1HL/1HL, where the training error is of order  $10^{-3}$ .

generalization error  $\epsilon_g^R(\mathcal{W}) = \epsilon^R$  even for  $P \rightarrow \infty$ : the  $1/P$  scaling of the GG will thus be restored asymptotically.

Numerical simulations for fully trained DNNs are shown in Fig. 2. As expected from the considerations raised above, and differently from the lazy training regime, here we do not observe a simple  $1/P$  scaling of the gap [Fig. 2(a)]. On the contrary, the GG curves display two learning stages as a function of  $P$ , with the GG falling systematically below the mean field bound above  $P \sim 10^4$ . Once we enter in the second learning stage (for  $P$  of the order of  $10^4$ ) the gap, however, seems to approach zero as fast as  $1/P$  for both one and two hidden layer architectures and across the different synthetic datasets, suggesting a fine constant residual generalization error in this second learning stage. In Fig. 2(b) we analyze the generalization performance as the width of the last layer  $N_{\text{out}}$  grows. According to our predictions, the bound holds, and a linear or sublinear degradation of the generalization performance is systematically observed across the different student and teacher architectures.

### B. State-of-the-art architectures

As an additional and more challenging test, we present the results for the generalization gap obtained by training three different state-of-the-art convolutional architectures (ResNet18, DenseNet121 and VGG-11) on the MNIST dataset of handwritten digits [46]. Notice that this problem is in principle a classification problem, but our theory has been formulated for regression. For this reason we implemented the learning problem as a regression task: for each digit vector  $\mathbf{x}$ , the associated output is simply the integer number, between 0 and 9, corresponding to its class. Coherently, the performance of the network is not measured using the standard accuracy (i.e., the fraction of correctly classified digits), but as the

mean square deviation between the network's output and the class index.

A summary of the simulations is found in Fig. 3 (details on the learning protocols are provided in Appendix B). Remarkably, our bound is also satisfied in the case of state-of-the-art architectures trained on a dataset of practical relevance for computer vision. Moreover, notice that in the regime we are exploring the generalization gap approaches zero faster than  $1/P$ .

### V. DISCUSSION AND FUTURE PERSPECTIVES

Our mean field analysis, while lacking the full rigor of theorems, establishes a much more stringent bound for the generalization gap w.r.t. the ones obtained in the strict context of statistical learning theory. For instance, SLT predicts an upper bound roughly proportional to  $\sqrt{N_{\text{tot}}/P}$  with  $N_{\text{tot}}$  being the *total* number of network parameters [47] (in the setting of regression problems this can be found using the Rademacher complexity [9]), while our upper bound depends on the number of parameters of the network only through the width of the last layer; this may lead to speculate that the width of the last layer plays a special role in DNN architectures. The reason why we have been able to obtain this improved nonrigorous upper bound on the generalization gap is related to taking optimization into account, at least in the last layer. Whereas SLT attempts to find results for the generalization gap that hold for every function in the model class, here we are restricting to special (but significant) elements of the class. In particular, for the specific case of a DNN with parameters  $\vartheta = \{\mathbf{v}, \mathcal{W}\}$ , classic SLT bounds hold for every realization of  $\vartheta$ . On the contrary, our approach assumes that at fixed  $\mathcal{W}$  the weights  $\mathbf{v}$  of the last layer are optimized w.r.t. the training set.

It is fair to stress the major limitations of our mean field bound. (1) Crucially the bounds in SLT hold for any size of the




FIG. 3. Generalization gap for three state-of-the-art architectures trained on the MNIST dataset of handwritten digits. The dependence of the normalized generalization gap on the size of the training set  $P$  is qualitatively similar for ResNet18 (blue symbols), DenseNet121 (green), and VGG-11 (red). The solid black line marks the mean field upper bound. Notice that by rescaling the GG by  $N_{\text{out}}$  we can better compare architectures with different last layer size. Averages have been performed over three different initial conditions for the architecture weights, and error bars measure one standard error. Typical training errors are of order  $10^{-4}$ .

dataset and of the architecture (although it loses predictivity when the latter is larger than the former). On the contrary, we have no control on finite-size corrections, since our results hold only in the limit of large  $P$  and  $N_{\text{out}}$ , but our bound retains predictive power in the overparametrized regime. (2) As observed in Fig. 2, there are cases in which the bound starts to hold only after a threshold  $P_t \approx 10^4$ – $10^5$ , which the theory does nothing to predict. However, our numerical experiments (e.g., those in Fig. 3) show that the threshold is rather small in many empirically relevant cases. (3) A Gaussian approximation was performed at the level of the replicated partition function (for details see Appendix A). Despite similar assumptions have been successfully employed in the past to study the statistical physics of kernels and of random feature models (where excellent agreement with numerics has also been found) [26,27], we cannot guarantee that this approximation is always quantitatively correct. (4) These results hold for regression only. It would be interesting to understand whether the same nonrigorous tools could also be used to study classification problems [25]. (5) In the spirit of SLT, we have investigated the generalization gap, but it would be interesting to understand if our results have some interpretation in terms of the classic bias-variance decomposition of the generalization error, as done, for instance, in Ref. [19].

Some of these drawbacks may be addressed by a more rigorous approach, for instance, by the use of random matrix

theory to avoid replicas (as done, for instance, in the case of kernel learning in Refs. [38,48–51]). This may shed light on the dependence of the GG on the input dimension  $D$ , which we have shown numerically in the lazy training regime to saturate exponentially to the mean field bound.

In conclusion, we would like to point out an apparently intriguing consequence of our findings: since our mean field bound suggests a (linear or sublinear) degradation of the generalization performance with the last layer size  $N_{\text{out}}$ , we might be led to surmise that, to improve generalization performance, it may be convenient to design architectures with a small last layer. A more systematic investigation on state-of-the-art architectures is needed to understand whether this insight may lead to design more performant deep neural networks in the future.

### ACKNOWLEDGMENTS

The authors would like to thank V. Erba, B. Loureiro, and M. Pastore for feedback on the manuscript and for pointing out recent related work. R.P. would like to thank her colleagues at Bocconi University for discussions. P.R. acknowledges funding from the Fellini program under the H2020-MSCA-COFUND action, Grant Agreement No. 754496, INFN (IT).

### APPENDIX A: SKETCH OF THE REPLICA SYMMETRIC CALCULATION

We now discuss the salient aspects of the replica calculation. Further technical details can be found in Appendix C. Our goal is to evaluate the generalization and training errors defined in Eqs. (2) and (3) for arbitrary teacher function  $f_T(\mathbf{x})$  and using as a student a function of the form (1).

#### 1. Replicated partition function

In order to evaluate these observables, one introduces a Gibbs distribution  $p_G(\mathbf{v}) = \frac{1}{Z} e^{-\beta \mathcal{L}(\mathbf{v})}$ , where  $\mathcal{L}$  is the effective Hamiltonian defined in Eq. (5) and  $\beta$  can be thought as the inverse of an effective temperature. The partition function  $Z$  is given by

$$Z = \int d^N \mathbf{v} e^{-\frac{\beta}{2} \sum_{\mu} [f_T(\mathbf{x}^{\mu}) - f_S(\mathbf{v}, \mathbf{x}^{\mu})]^2 - \frac{\beta}{2} \lambda \|\mathbf{v}\|^2}. \quad (\text{A1})$$

In the  $\beta \rightarrow \infty$  limit, the Gibbs measure is dominated by the minimum of the Hamiltonian, which corresponds to the minimum of the training loss. As for many other problems in disordered systems, meaningful results can be obtained only by *quenched* averages, that is, by averaging the logarithm of the partition function over all the possible realizations of the training set  $\mathcal{T}$ . Physically, this amounts to optimizing first the student weights for any given instance of the dataset  $\mathcal{T}$  and then averaging over all dataset realizations [52]. In order to perform this quenched average we exploit the standard replica method [37],

$$\langle \log Z \rangle_{\mathcal{T}} = \lim_{m \rightarrow 0} \frac{\langle Z^m \rangle_{\mathcal{T}} - 1}{m}, \quad (\text{A2})$$

where one first computes the average of  $Z^m$  for an integer number of replica  $m$  and only later one performs the

analytical continuation to  $m \rightarrow 0$ . Since the inputs are drawn as iid variables, the integral over the training set factorizes to yield

$$\langle Z^m \rangle_{\mathcal{T}} = \int \prod_{a=1}^m d^N \mathbf{v}^a e^{-\frac{\beta}{2} \lambda \sum_a \|\mathbf{v}^a\|^2} \left[ \int d^D x \rho(\mathbf{x}) e^{-\frac{\beta}{2} \sum_a (q^a)^2} \right]^P, \quad (\text{A3})$$

where  $a$  is a replica index and we introduced a set of auxiliary random variables  $q^a \equiv f_{\mathcal{T}}(\mathbf{x}) - \mathbf{v}^a \cdot \phi(\mathbf{x})$  with mean

$\mu_q^a(\{\mathbf{v}^a\}) = \langle q^a \rangle_{\rho}$  and covariance matrix:

$$Q_{ab}(\{\mathbf{v}^a\}) = \langle q^a q^b \rangle_{\rho} = T + (\mathbf{v}^a)^T \Phi \mathbf{v}^b - \mathbf{J}^T \cdot (\mathbf{v}^a + \mathbf{v}^b) \quad (\text{A4})$$

(with  $\mathbf{J}$ ,  $\Phi$ , and  $T$  given by Eq. (6)). To proceed further we note that each of the random variables  $q^a$  is the sum of  $N$  random variables. For a large last layer size  $N$  and input dimension  $D$  we approximate their probability distribution with a multivariate Gaussian with mean  $\mu_q^a$  and covariance matrix  $Q_{ab}$  (an order parameter which measures the overlap between replica  $a$  and  $b$ ). This allows us to perform the integration in the square brackets in Eq. (A3) to get

$$\begin{aligned} \int d^D x \rho(\mathbf{x}) e^{-\frac{\beta}{2} \sum_a (q^a)^2} &= \int \prod_a d q^a e^{-\frac{\beta}{2} \sum_a (q^a)^2} \int d^D x \rho(\mathbf{x}) \delta\left(q_a - \sum_{\alpha} v_{\alpha}^a \phi_{\alpha}(\mathbf{x}) + f_{\mathcal{T}}(\mathbf{x})\right) \\ &\simeq \int \prod_a d q^a \frac{e^{-\frac{\beta}{2} \sum_a (q^a)^2 - \frac{1}{2} \sum_{ab} (q^a)^T Q_{ab}^{-1} q^b}}{\sqrt{(2\pi)^m \det \mathbf{Q}}} = \left( \det(\mathbb{I} + \beta \mathbf{Q}) \right)^{-\frac{1}{2}}, \end{aligned} \quad (\text{A5})$$

where we have assumed  $\mu_q^a = 0$  without loss of generality [27].

It is worth noticing that this Gaussian approximation is crucial in order to make progress, but rather uncontrolled, as we lack a formal result demonstrating its validity. Nonetheless similar nonrigorous approximations are quite standard in the literature on kernel learning, which include the seminal work on support vector machines by Dietrich, Oppen, and Sompolinsky [26], more recent findings on kernel regression [27] or works on the so-called random feature model, where this approximation goes under the name of *Gaussian equivalence principle* [20]. These ideas imply that the  $N$  feature maps  $\phi_{\alpha}(\mathbf{x})$  are somehow mutually weakly correlated, an assumption deemed reasonable for a large class of relevant architectures in the thermodynamic limit of large  $N$  and  $D$  with finite  $N/D$  [20].

The integration of Eq. (A3) over the weights  $\mathbf{v}^a$ , while rather convoluted, now follows a standard replica scheme as detailed in the more technical Appendix C. Due to Eq. (A5) and making use of the identities

$$\begin{aligned} 1 &= \int d Q_{ab} \delta(Q_{ab} - \langle q^a q^b \rangle_{\rho}) \\ &= \int d Q_{ab} \frac{d \hat{Q}_{ab}}{2\pi} e^{-i \hat{Q}_{ab} (Q_{ab} - \langle q^a q^b \rangle_{\rho})} \end{aligned} \quad (\text{A6})$$

one may express Eq. (A3) as an integral over the replica order parameter  $Q_{ab}$  and its conjugated variables  $\hat{Q}_{ab}$ . Working in the replica symmetric ansatz, which holds for our convex problem,  $Q_{ab} = Q_0 \delta_{ab} + Q(1 - \delta_{ab})$  (the same symmetry holding for the conjugated variable), one may compute the leading, linear order contribution in  $m$  which determines the limit (A2), to get

$$\langle Z^m \rangle_{\mathcal{T}} \approx \int d Q_0 d Q d \hat{Q}_0 d \hat{Q} e^{-\frac{mP}{2} S_{\beta}(Q_0, Q, \hat{Q}_0, \hat{Q})} \quad (\text{A7})$$

with the rather complicated expression for the action  $S_{\beta}(Q_0, Q, \hat{Q}_0, \hat{Q})$  given by Eq. (C12) of Appendix C.

We can finally solve this last integral by saddle-point method in the large  $P$  limit. One has to solve a set of four saddle-point equations to find the action minimum that determines the two order parameters of the problem and their conjugated variables (see Appendix C). In the  $\beta \rightarrow \infty$  limit the order parameter has the rather simple saddle-point solution

$$Q^* = Q_0^* = \frac{T - P \mathbf{J}^T (2\kappa \lambda \mathbb{I} + P \Phi) \mathbf{G}^{-2} \mathbf{J}}{1 - P \text{Tr}[\Phi^2 \mathbf{G}^{-2}]}, \quad (\text{A8})$$

where the variable  $\kappa$  and the invertible matrix  $\mathbf{G}$  have been introduced in the main text.

## 2. Generalization and training errors

The generalization  $\epsilon_g$  and training  $\epsilon_t$  errors can be easily related to the saddle-point replica order parameters. To see this, first consider the generalization error. From Eqs. (2) and (1),

$$\begin{aligned} \epsilon_g &= \int d^D x \rho(\mathbf{x}) [f_{\mathcal{T}}(\mathbf{x}) - \mathbf{v}^* \cdot \phi(\mathbf{x})]^2 \\ &= T - 2 \mathbf{J}^T \mathbf{v}^* + \mathbf{v}^{*T} \Phi \mathbf{v}^* = Q^*, \end{aligned} \quad (\text{A9})$$

where in the last equality we used Eq. (A4) and the replica symmetric ansatz. In the limit  $P \rightarrow \infty$ , it is easy to show that  $\kappa \rightarrow 1$  and  $\mathbf{G} \sim P \Phi$  so that the generalization error converges to the *residual generalization error* introduced in the main text, Eq. (9),

$$\epsilon_g \rightarrow \epsilon_g^{\text{R}} = T - \mathbf{J}^T \Phi^{-1} \mathbf{J}. \quad (\text{A10})$$

Notice that this result is not surprising, and it provides a first consistency check of our replica mean field theory: one could also obtain it by directly minimizing Eq. (A9) with respect to the parameters  $\mathbf{v}$ . In fact,  $\partial_{\mathbf{v}} \epsilon_g = -2 \mathbf{J}^T + 2 \Phi \mathbf{v}$  implies  $\mathbf{v}^* = \mathbf{J}^T \Phi^{-1}$ , so that  $\epsilon_g(\mathbf{v}^*) \equiv \epsilon_g^{\text{R}} = T - \mathbf{J}^T \Phi^{-1} \mathbf{J}$ . Furthermore, by isolating the residual generalization error in Eq. (A8), we finally find the compact formula for the generalization error quoted in Eq. (7) of the main text.

Using the theory developed so far, we can also have access to the average value of the training error. We notice that the average training error is by definition the average loss function defined in Eq. (5), evaluated in  $\lambda = 0$  (up to a factor  $P$ ). This means that one can extract its value by evaluating the action on the saddle-point solution and performing the limit  $\beta \rightarrow \infty$ :

$$\epsilon_t = \lim_{\beta \rightarrow \infty} \frac{1}{\beta} S_\beta(Q^*, Q_0^*, \hat{Q}^*, \hat{Q}_0^*)|_{\lambda=0}. \quad (\text{A11})$$

After some straightforward but lengthy algebraic manipulations, one recovers the training error given in Eq. (7).

## APPENDIX B: NUMERICAL EXPERIMENT DETAILS

In this section we give a detailed report of all our numerical procedures. The code to replicate our experiments can be found at [53].

### 1. Teacher-student architectures

*Student architectures.* We considered six types of student architectures: two toy networks with one and two hidden layers (size of the second hidden layer  $N_{\text{hid}} = 200$ ) and three state-of-the-art convolutional ones (ResNet18, DenseNet121, and VGG11). The toy architectures have fully connected layers and ReLU activation functions at every layer but the last; the convolutional networks are the standard PyTorch [54] models modified to yield a scalar output suitable for regression through a last fully connected linear layer with parameters  $\mathbf{v}$ , instead of the *LogSoftMax* that is employed for classification. All these architectures have several convolutional layers before a last fully connected one that counts respectively  $N_{\text{out}} = 512, 1024, 4096$  hidden units. The total number of trainable parameters in the first two networks (weights and biases) is approximately 10 million, while vgg11 counts 10 times as many.

*Teacher architectures and inputs for toy DNNs.* Each linear or quadratic teachers (see Results) is defined by a random uniform vector  $\mathbf{t} \in \mathbb{R}^D$  of unitary norm. 1HL teachers are defined by parameters  $q_\alpha \in \mathbb{R}$  and  $\mathbf{S}_\alpha \in \mathbb{R}^D$  ( $\alpha = 1, \dots, M = 200$ ). They are drawn from a normal distribution with zero mean and variance (respectively)  $1/M$  and  $1/D$ . Inputs  $\mathbf{x} \in \mathbb{R}^D$  are also drawn from normal distribution with zero mean and unit variance.

### 2. Learning and generalization

*Learning algorithm.* All the architectures are trained with the Adam optimizer [55] and a small weight decay ( $w_d = 10^{-5}$ ), while the learning rate  $\alpha$  is set to  $10^{-3}$  for the toy architectures and  $10^{-4}$  for the convolutional ones. The weight decay  $w_d$  is related to the regularization parameter  $\lambda$  via  $w_d = \lambda\alpha$ , and we have verified that our results do not change quantitatively by varying  $\lambda$  in a reasonable range ( $\lambda \leq 10^{-1}$ ). The regression loss employed is the standard mean-squared error,

$$\text{MSE} = \frac{1}{P} \sum_{\mu=1}^P [f_T(x^\mu) - f_S(x^\mu)]^2. \quad (\text{B1})$$




FIG. 4. Train (blue) and test (magenta) loss of different teacher-student tasks as a function of the training epochs. The test loss reaches a plateau even if the training loss is noisy, due to the different order of magnitude reached by the two. It is worth remarking that the losses are normalized with the trivial predictor and therefore at epoch 1 are  $O(1)$ .

For the MNIST dataset, we consider the labels as integers:  $f_S(x^\mu) = [0, 1, 2, 3, 4, 5, 6, 7, 8, 9]$ , and the MSE loss is computed as in other cases.

*Estimation of the generalization gap.* We first bring the training procedure to convergence, i.e., we ensure that the train and test loss have reached a plateau. For synthetic datasets, this requires a large number of training epochs ( $3 \times 10^4$ ), while MNIST is learnt faster (100 epochs). In one epoch the network is fed all the dataset, arranged in minibatches only when full batch learning is prohibited by memory limitations. Once the train loss is steady with respect to the test loss, we retrieve the generalization gap as the average over the last 100 epochs (50 for MNIST). The size of the train set is  $P_{\text{test}} = 10^4$  is all cases.

Several plots of train and test loss vs the number of epochs are shown in Fig. 4: for different teacher-student pairs the plateau in the test loss is always reached: even if some noise is still visible in the train loss, its oscillations are too small to affect the test loss and therefore the generalization gap.

### 3. Trivial predictor

To compare results obtained from different teacher-student pairs, we need to normalise the loss by its natural scale, i.e., the trivial predictor  $T$  defined in Eq. (6). By doing this, we make sure that the train and test loss are always of  $O(1)$  for a random architecture (i.e., at epoch 0).  $T$  is a property of the dataset, and its computation changes accordingly.

*MNIST.* The MNIST dataset is intrinsically suitable for classification problems, since it has 10 classes with discrete labels, that are simply the first 10 integers. For this reason,






FIG. 5. Numerical estimate of the trivial predictor  $T$  as a function of the number of points  $P_{\text{norm}}$  used in the approximation. This values describe a single realization of a random teacher function (linear, quadratic, 1HL).  $T$  converges to a fixed value around  $P_{\text{norm}} = 10^6$ , that is, the value chosen for all our numerical experiments.

the integral that describes the trivial predictor here becomes a summation over the classes. Another simplification can be made considering that MNIST is balanced over the training labels: the  $P = 6 \times 10^4$  training examples are equally distributed over the 10 classes. In this case the computation of the trivial predictor is simple and can be performed analytically:

$$T = \frac{1}{P} \sum_{n=0}^9 \frac{P}{\text{no. classes}} n^2 = \frac{1}{\text{no. classes}} \sum_{n=0}^9 n^2 = 28.5. \quad (\text{B2})$$

*Synthetic datasets.* For the other teachers that we employed  $T$  cannot be computed analytically. We therefore perform a numerical estimation in the following approximation:

$$T \sim \frac{1}{P_{\text{norm}}} \sum_{\mu=1}^{P_{\text{norm}}} f_{\text{T}}^2(x^\mu). \quad (\text{B3})$$

Operatively, we draw an independent data sample of  $P_{\text{norm}}$  elements  $\{x^\mu\}_{\mu=1, \dots, P_{\text{norm}}}$ , and average the respective squared true labels  $f_{\text{T}}^2(x^\mu)$ .  $P_{\text{norm}} = 10^6$  was chosen as a safe compromise between computational time and consistency of the estimate. The convergence of  $T$  to a fixed value when  $P_{\text{norm}}$  grows is shown in Fig. 5.

#### 4. Higher dimensional synthetic inputs

In the main text we show the generalization gap scaling for fixed input size  $D = 50$ , with the exception of the inset in Fig. 1. Here we report the full analysis of higher dimensional input size and clarify the procedure used to obtain the inset. In Fig. 6 we plot the generalization gap of a lazy training architecture learning the three function considered in the main text (linear, quadratic, and 1HL), as a function of the number of training examples and for different sizes of the input  $D$ . Increasing  $D$  does not change the asymptotic behavior of the curves, but has the effect of pushing them closer to the bound given the same trainset size  $P$ . To assert how fast the curves approach the bound, we have performed a linear fit of the generalization gap as a function of  $P$ :

$$\log \Delta \tilde{\epsilon} = m \log P + y_0. \quad (\text{B4})$$

Note that for the bound  $m \equiv -1$  and  $y_0 \equiv \log 2N_{\text{out}}$ . For different values of  $D$ , it is verified that  $m \sim -1$ , while  $y_0 - 2N_{\text{out}}$  approaches 0 exponentially fast (see also inset in Fig. 1).

#### APPENDIX C: REPLICIA SYMMETRIC ANSATZ AND SADDLE-POINT EQUATIONS

Due to the Gaussian approximation (A5) and making use of the identities (A6), the replicated partition function (A3) may be recast in the following form:

$$\begin{aligned} \langle Z^m \rangle_{\mathcal{T}} &= \int dQ_{ab} d\hat{Q}_{ab} K(\beta, \mathbf{Q})^{-\frac{P}{2}} e^{iQ_{ab}\hat{Q}_{ab}} \\ &\times \int d^N \mathbf{v}^a e^{-\frac{\beta \lambda}{2} \sum_a \|\mathbf{v}^a\|^2 - i \sum_{a \leq b} \hat{Q}_{ab}} \\ &\times [T + (\mathbf{v}^a)^T \Phi \mathbf{v}^b - \mathbf{J}^T (\mathbf{v}^a + \mathbf{v}^b)], \end{aligned} \quad (\text{C1})$$

where

$$K(\beta, \mathbf{Q}) \equiv \det(\mathbb{I} + \beta \mathbf{Q}). \quad (\text{C2})$$

The idea is now to explicitly perform the integration over the replicated weights  $\mathbf{v}^a$  and to evaluate the integrals over  $Q_{ab}$  and  $\hat{Q}_{ab}$  with the saddle-point method. As occurs in standard spin glass models, we are left with the complication of




FIG. 6. Generalization gap of a lazy training architecture as a function of the number of examples in the training set for different input sizes  $D$ . From left to right the teacher functions are, respectively, linear, quadratic, and 1HL. The size of the hidden layer is kept fixed to  $N_{\text{out}} = 400$ .

performing the tricky limit  $m \rightarrow 0$ , which may lead to the breaking of the so-called replica symmetry of the matrix  $Q_{ab}$ . However, in this specific case the underlying optimization problem that we are dealing with is convex and this guarantees that the simplest replica symmetric ansatz for the matrix  $Q_{ab}$

will provide the exact solution to the saddle-point equations in the limit  $m \rightarrow 0$  [26].

The assumption of replica symmetry amounts to require that the matrices  $Q_{ab}/\hat{Q}_{ab}$  take the following form:

$$Q_{ab} = \begin{cases} Q_0 & a = b \\ Q & a \neq b \end{cases} \quad \hat{Q}_{ab} = \begin{cases} \hat{Q}_0 & a = b \\ \hat{Q} & a \neq b. \end{cases} \quad (\text{C3})$$

With a slight abuse of notation, we can specify our analysis to the replica symmetric ansatz already at the level of Eq. (C1), thus obtaining

$$\begin{aligned} \langle Z^m \rangle_{\mathcal{T}} &= \int dQ_0 dQ d\hat{Q}_0 d\hat{Q} K^{-\frac{P}{2}} e^{im[\hat{Q}_0(Q_0-T) + \frac{m-1}{2}\hat{Q}(Q-T)]} \\ &\times \int d^N \mathbf{v}^a e^{-\frac{\beta\lambda}{2} \sum_{\alpha,a} (v_a^\alpha)^2 - i(\hat{Q}_0 - \hat{Q}) \sum_a (\mathbf{v}^a)^T \Phi \mathbf{v}^a} e^{-i\frac{\hat{Q}}{2} \sum_{ab} (\mathbf{v}^a)^T \Phi \mathbf{v}^b + 2i[\hat{Q}_0 - (1-m)\hat{Q}] \sum_a \mathbf{J}^T \mathbf{v}^a}. \end{aligned} \quad (\text{C4})$$

By performing the Gaussian integrals over the replicated weights and using standard mathematical manipulations, we finally get the following result:

$$\begin{aligned} \langle Z^m \rangle_{\mathcal{T}} &= \int dQ_0 dQ d\hat{Q}_0 d\hat{Q} K^{-\frac{P}{2}} e^{-\frac{m}{2}[\hat{Q}_0(Q_0-T) + (m-1)\hat{Q}(Q-T)]} \\ &\times (2\pi)^{\frac{m}{2}} e^{-\frac{m-1}{2} \text{Tr}[\log[\beta\lambda \mathbb{1} - (\hat{Q}_0 - \hat{Q})\Phi]] - \frac{1}{2} \text{Tr}[\log\{\beta\lambda \mathbb{1} - [\hat{Q}_0 - (1-m)\hat{Q}]\Phi\}]} e^{\frac{m}{2}[\hat{Q}_0 - (1-m)\hat{Q}]^2 \mathbf{J}^T [\beta\lambda \mathbb{1} - (\hat{Q}_0 - (1-m)\hat{Q})\Phi]^{-1} \mathbf{J}}. \end{aligned} \quad (\text{C5})$$

At this point we have managed to integrate both on the dataset  $\mathcal{T}$  and on the replicated weights. To find the solution of the last integrals over the order parameters, we exploit the saddle-point method, which will deliver the exact solution in the limit  $P \rightarrow \infty$ .

### 1. Replica symmetric ansatz

First we notice that assuming replica symmetry, the contribution  $K$  to the partition function simplifies in the following way:

$$\begin{aligned} K(\beta, Q, Q_0) &= \det[\mathbb{1} + \beta Q] \\ &= [1 + \beta(Q_0 - Q)]^{m-1} \{1 + \beta[Q_0 - (1-m)Q]\}. \end{aligned} \quad (\text{C6})$$

In order to extract the  $m \rightarrow 0$  limit, we have to keep only those terms that are linear in  $m$ . This is easily done term by

term:

$$\begin{aligned} &e^{-\frac{P}{2} \log\{(1 + \beta(Q_0 - Q))^{m-1} \{1 + \beta[Q_0 - (1-m)Q]\}\}} \\ &\approx e^{-\frac{P}{2} m \left\{ \frac{\beta Q}{1 + \beta(Q_0 - Q)} + \log[1 + \beta(Q_0 - Q)] \right\}}, \end{aligned} \quad (\text{C7})$$

$$e^{-\frac{m}{2}[\hat{Q}_0(Q_0-T) + (m-1)\hat{Q}(Q-T)]} \approx e^{-\frac{m}{2}[\hat{Q}_0(Q_0-T) - \hat{Q}(Q-T)]}, \quad (\text{C8})$$

$$\begin{aligned} &e^{-\frac{m-1}{2} \text{Tr}[\log[\beta\lambda \mathbb{1} - (\hat{Q}_0 - \hat{Q})\Phi]] - \frac{1}{2} \text{Tr}[\log\{\beta\lambda \mathbb{1} - [\hat{Q}_0 - (1-m)\hat{Q}]\Phi\}]} \\ &\approx e^{-\frac{m}{2} \text{Tr}[\log[\beta\lambda \mathbb{1} - (\hat{Q}_0 - \hat{Q})\Phi]] + \frac{mQ}{2} \text{Tr}\{\Phi[\beta\lambda \mathbb{1} - (\hat{Q}_0 - \hat{Q})\Phi]^{-1}\}}, \end{aligned} \quad (\text{C9})$$

and

$$\begin{aligned} &e^{\frac{m}{2}[\hat{Q}_0 - (1-m)\hat{Q}]^2 \mathbf{J}^T [\beta\lambda \mathbb{1} - (\hat{Q}_0 - (1-m)\hat{Q})\Phi]^{-1} \mathbf{J}} \\ &\approx e^{\frac{m}{2}(\hat{Q}_0 - \hat{Q})^2 \mathbf{J}^T [\beta\lambda \mathbb{1} - (\hat{Q}_0 - \hat{Q})\Phi]^{-1} \mathbf{J}}. \end{aligned} \quad (\text{C10})$$

As such, the leading contribution to the average replicated partition function in the  $m \rightarrow 0$  limit reads

$$\begin{aligned} \langle Z^m \rangle_{\mathcal{T}} &\sim \int dQ_0 dQ d\hat{Q}_0 d\hat{Q} e^{-\frac{P}{2} m \left\{ \frac{\beta Q}{1 + \beta(Q_0 - Q)} + \log[1 + \beta(Q_0 - Q)] \right\}} e^{-\frac{m}{2}[\hat{Q}_0(Q_0-T) - \hat{Q}(Q-T)]} \\ &\times e^{-\frac{m}{2} \text{Tr}[\log[\beta\lambda \mathbb{1} - (\hat{Q}_0 - \hat{Q})\Phi]] + \frac{mQ}{2} \text{Tr}\{\Phi[\beta\lambda \mathbb{1} - (\hat{Q}_0 - \hat{Q})\Phi]^{-1}\}} e^{\frac{m}{2}(\hat{Q}_0 - \hat{Q})^2 \mathbf{J}^T [\beta\lambda \mathbb{1} - (\hat{Q}_0 - \hat{Q})\Phi]^{-1} \mathbf{J}}, \end{aligned} \quad (\text{C11})$$

and by rescaling the parameters  $\hat{Q} \rightarrow P\hat{Q}$  and  $\hat{Q}_0 \rightarrow P\hat{Q}_0$ , we recast the partition function in the form  $\langle Z^m \rangle_{\mathcal{T}} = \int e^{-\frac{mP}{2} S_\beta}$ , where the action  $S_\beta$  is defined as

$$S_\beta = \frac{\beta Q}{1 + \beta(Q_0 - Q)} + \log[1 + \beta(Q_0 - Q)] + \hat{Q}_0(Q_0 - T)$$

$$\begin{aligned} &-\hat{Q}(Q - T) + \frac{1}{P} \text{Tr}[\log[\beta\lambda \mathbb{1} - P(\hat{Q}_0 - \hat{Q})\Phi]] \\ &-\hat{Q} \text{Tr}\{\Phi[\beta\lambda \mathbb{1} - P(\hat{Q}_0 - \hat{Q})\Phi]^{-1}\} \\ &-P(\hat{Q}_0 - \hat{Q})^2 \mathbf{J}^T [\beta\lambda \mathbb{1} - P(\hat{Q}_0 - \hat{Q})\Phi]^{-1} \mathbf{J}. \end{aligned} \quad (\text{C12})$$

## 2. Saddle-point equations

We can now move to the derivation of the saddle-point equations. First, we notice that direct differentiation with respect to  $Q$  and  $Q_0$  allows us to find the explicit expressions for  $\hat{Q}$  and  $\hat{Q}_0$ :

$$0 = \frac{\partial S_\beta}{\partial Q} = -\hat{Q} + \frac{\beta^2 Q}{[1 + \beta(Q_0 - Q)]^2}, \quad (\text{C13})$$

which implies

$$\hat{Q} = \frac{\beta^2 Q}{[1 + \beta(Q_0 - Q)]^2} \quad (\text{C14})$$

and

$$0 = \frac{\partial S_\beta}{\partial Q_0} = \hat{Q}_0 + \frac{-\beta Q(\beta)}{[1 + \beta(Q_0 - Q)]^2} + \frac{\beta}{1 + \beta(Q_0 - Q)}, \quad (\text{C15})$$

which gives

$$\begin{aligned} \hat{Q}_0 &= \frac{\beta^2 Q}{[1 + \beta(Q_0 - Q)]^2} - \frac{\beta}{1 + \beta(Q_0 - Q)} \\ &= \hat{Q} - \frac{\beta}{1 + \beta(Q_0 - Q)}. \end{aligned} \quad (\text{C16})$$

Let us look at the derivative of the action w.r.t.  $\hat{Q}$ :

$$0 = \frac{\partial S_\beta}{\partial \hat{Q}} = -(Q - T) + \frac{1}{P} \partial_{\hat{Q}} \text{Tr}[\log(\tilde{\mathbf{G}})] - \text{Tr}[\Phi \tilde{\mathbf{G}}^{-1}] - \hat{Q} \partial_{\hat{Q}} \text{Tr}[\Phi \tilde{\mathbf{G}}^{-1}] + P \mathbf{J}^T \frac{P \Phi (\hat{Q}_0 - \hat{Q})^2 + 2 \tilde{\mathbf{G}} (\hat{Q}_0 - \hat{Q})}{\tilde{\mathbf{G}}^2} \mathbf{J}, \quad (\text{C17})$$

which gives

$$\begin{aligned} Q &= T + \frac{1}{P} \partial_{\hat{Q}} \text{Tr}[\log(\tilde{\mathbf{G}})] - \text{Tr}[\Phi \tilde{\mathbf{G}}^{-1}] - \hat{Q} \partial_{\hat{Q}} \text{Tr}[\Phi \tilde{\mathbf{G}}^{-1}] + P \mathbf{J}^T \frac{P \Phi (\hat{Q}_0 - \hat{Q})^2 + 2 \tilde{\mathbf{G}} (\hat{Q}_0 - \hat{Q})}{\tilde{\mathbf{G}}^2} \mathbf{J} \\ &= T + P \hat{Q} \text{Tr}[(\Phi \tilde{\mathbf{G}}^{-1})^2] + P \mathbf{J}^T \frac{P \Phi (\hat{Q}_0 - \hat{Q})^2 + 2 \tilde{\mathbf{G}} (\hat{Q}_0 - \hat{Q})}{\tilde{\mathbf{G}}^2} \mathbf{J}, \end{aligned} \quad (\text{C18})$$

where for convenience we have defined the  $N \times N$  matrix  $\tilde{\mathbf{G}} = \beta \lambda \mathbb{1} - P(\hat{Q}_0 - \hat{Q})\Phi$ . Finally we obtain the saddle-point equation for  $\hat{Q}_0$ :

$$0 = \frac{\partial S_\beta}{\partial \hat{Q}_0} = (Q_0 - T) + \frac{1}{P} \partial_{\hat{Q}_0} \text{Tr}[\log(\tilde{\mathbf{G}})] - \hat{Q} \partial_{\hat{Q}_0} \text{Tr}[\Phi \tilde{\mathbf{G}}^{-1}] - P \mathbf{J}^T \frac{P \Phi (\hat{Q}_0 - \hat{Q})^2 + 2 \tilde{\mathbf{G}} (\hat{Q}_0 - \hat{Q})}{\tilde{\mathbf{G}}^2} \mathbf{J}, \quad (\text{C19})$$

giving us

$$\begin{aligned} Q_0 &= T + \text{Tr}[\Phi \tilde{\mathbf{G}}^{-1}] + P \hat{Q} \text{Tr}[\Phi \tilde{\mathbf{G}}^{-1} \Phi \tilde{\mathbf{G}}^{-1}] \\ &\quad + P \mathbf{J}^T \frac{P \Phi (\hat{Q}_0 - \hat{Q})^2 + 2 \tilde{\mathbf{G}} (\hat{Q}_0 - \hat{Q})}{\tilde{\mathbf{G}}^2} \mathbf{J} \quad (\text{C20}) \\ &= Q + \text{Tr}[\Phi \tilde{\mathbf{G}}^{-1}]. \end{aligned}$$

Let us now consider the special combination  $\kappa = 1 + \beta(Q_0 - Q)$ . By using Eq. (C20), we obtain

$$\begin{aligned} \kappa &= 1 + \beta(Q_0 - Q) = 1 + \beta(Q + \text{Tr}[\Phi \tilde{\mathbf{G}}^{-1}] - Q) \\ &= 1 + \beta \text{Tr}[\Phi \tilde{\mathbf{G}}^{-1}], \end{aligned} \quad (\text{C21})$$

whereas by considering the difference between Eq. (C14) and Eq. (C16), we easily show that the difference  $\hat{Q}_0 - \hat{Q}$  depends on  $\kappa$  only as

$$\begin{aligned} \hat{Q}_0 - \hat{Q} &= \hat{Q} - \frac{\beta}{1 + \beta(Q_0 - Q)} - \hat{Q} = -\frac{\beta}{1 + \beta(Q_0 - Q)} \\ &= -\frac{\beta}{\kappa}. \end{aligned} \quad (\text{C22})$$

By inserting this result into the definition of  $\tilde{\mathbf{G}}$  we can define the rescaled matrix  $\mathbf{G} = \kappa \tilde{\mathbf{G}}/\beta$ :

$$\mathbf{G} = \frac{\kappa}{\beta} [\beta \lambda \mathbb{1} - P(\hat{Q}_0 - \hat{Q})\Phi] = \kappa \lambda \mathbb{1} + P\Phi. \quad (\text{C23})$$

These observations allow us to show that the new variable  $\kappa$  satisfies the following self-consistency equation:

$$\begin{aligned} \kappa &= 1 + \beta \text{Tr} \left[ \Phi \left( \frac{\beta}{\kappa} \mathbf{G} \right)^{-1} \right] = 1 + \kappa \text{Tr}[\Phi \mathbf{G}^{-1}] \\ &= 1 + \kappa \text{Tr} \left[ \frac{\Phi}{\kappa \lambda \mathbb{1} + P\Phi} \right] \end{aligned} \quad (\text{C24})$$

and allow us to recast the solution of the saddle-point equations in the following very convenient form (notice that from now on the solutions of the saddle-point equations will be indicated with an asterisk):

$$\begin{aligned} \hat{Q}_0^* &= \hat{Q}^* - \frac{\beta}{\kappa}, \\ \hat{Q}^* &= \frac{\beta^2 Q^*}{\kappa^2}, \\ Q_0^* &= Q^* + \frac{\kappa - 1}{\beta}, \\ Q^* &= \frac{T - P \mathbf{J}^T \frac{2\kappa\lambda + P\Phi}{\mathbf{G}^2} \mathbf{J}}{1 - P \text{Tr}[\Phi^2 \mathbf{G}^{-2}]}. \end{aligned} \quad (\text{C25})$$

It is worth noticing that since  $\mathbf{G}$  and  $\kappa$  are independent on the inverse temperature  $\beta$ , the solution for the order parameter  $Q^*$  is also independent on the temperature. Moreover, we easily get that  $Q^* = Q_0^*$  in the limit  $\beta \rightarrow \infty$ .

- [1] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, ImageNet classification with deep convolutional neural networks, *Adv. Neural Inf. Proc. Syst.* **25**, 1097 (2012).
- [3] K. Simonyan and A. Zisserman, Very deep convolutional networks for large-scale image recognition, [arXiv:1409.1556](https://arxiv.org/abs/1409.1556).
- [4] K. He, X. Zhang, S. Ren, and J. Sun, Deep residual learning for image recognition, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June (IEEE, 2016), pp. 770–778.
- [5] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, Going deeper with convolutions, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June (IEEE, 2015), pp. 1–9.
- [6] V. Vapnik, An overview of statistical learning theory, *IEEE Trans. Neural Netw.* **10**, 988 (1999).
- [7] O. Bousquet, S. Boucheron, and G. Lugosi, *Introduction to Statistical Learning Theory* (Springer, Berlin, 2004), pp. 169–207.
- [8] V. Vapnik, *The Nature of Statistical Learning Theory* (Springer Science & Business Media, New York, 1999).
- [9] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms* (Cambridge University Press, Cambridge, 2014).
- [10] L. Zdeborová, Understanding deep learning is also a job for physicists, *Nat. Phys.* **16**, 602 (2020).
- [11] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, Understanding deep learning (still) requires rethinking generalization, *Commun. ACM* **64**, 107 (2021).
- [12] Y. Bahri, J. Kadmon, J. Pennington, S. S. Schoenholz, J. Sohl-Dickstein, and S. Ganguli, Statistical mechanics of deep learning, *Annu. Rev. Condens. Matter Phys.* **11**, 501 (2020).
- [13] S. Mei, A. Montanari, and P.-M. Nguyen, A mean field view of the landscape of two-layer neural networks, *Proc. Natl. Acad. Sci. USA* **115**, E7665 (2018).
- [14] F. Santambrogio, Euclidean, metric, and Wasserstein gradient flows: An overview, *Bull. Math. Sci.* **7**, 87 (2017).
- [15] G. M. Rotskoff and E. Vanden-Eijnden, Trainability and accuracy of neural networks: An interacting particle system approach, [arXiv:1805.00915](https://arxiv.org/abs/1805.00915).
- [16] L. Chizat and F. Bach, On the global convergence of gradient descent for over-parameterized models using optimal transport, *Adv. Neural Inf. Proc. Syst.* **31**, (2018).
- [17] J. Sirignano and K. Spiliopoulos, Mean field analysis of neural networks: A law of large numbers, *SIAM J. Appl. Math.* **80**, 725 (2020).
- [18] D. Araújo, R. I. Oliveira, and D. Yukimura, A mean-field limit for certain deep neural networks, [arXiv:1906.00193](https://arxiv.org/abs/1906.00193).
- [19] S. D’Ascoli, M. Refinetti, G. Biroli, and F. Krzakala, Double trouble in double descent: Bias and variance(s) in the lazy regime, in *Proceedings of the 37th International Conference on Machine Learning*, edited by H. Daumé III and A. Singh, *Proceedings of Machine Learning Research* Vol. 119 (PMLR, 2020), pp. 2280–2290.
- [20] S. Goldt, M. Mézard, F. Krzakala, and L. Zdeborová, Modeling the Influence of Data Structure on Learning in Neural Networks: The Hidden Manifold Model, *Phys. Rev. X* **10**, 041044 (2020).
- [21] S. Mei and A. Montanari, The generalization error of random features regression: Precise asymptotics and the double descent curve, *Comm. Pure Appl. Math.* **75**, 667 (2021).
- [22] F. Gerace, B. Loureiro, F. Krzakala, M. Mezard, and L. Zdeborova, Generalisation error in learning with random features and the hidden manifold model, in *Proceedings of the 37th International Conference on Machine Learning*, edited by H. Daumé III and A. Singh, *Proceedings of Machine Learning Research* Vol. 119 (PMLR, 2020), pp. 3452–3462.
- [23] C. Baldassi, C. Lauditi, E. M. Malatesta, R. Pacelli, G. Perugini, and R. Zecchina, Learning through atypical phase transitions in overparameterized neural networks, [arXiv:2110.00683](https://arxiv.org/abs/2110.00683).
- [24] C. Baldassi, C. Lauditi, E. M. Malatesta, G. Perugini, and R. Zecchina, Unveiling the Structure of Wide Flat Minima in Neural Networks, *Phys. Rev. Lett.* **127**, 278301 (2021).
- [25] B. Loureiro, C. Gerbelot, H. Cui, S. Goldt, F. Krzakala, M. Mezard, and L. Zdeborová, Learning curves of generic features maps for realistic datasets with a teacher-student model, in *Advances in Neural Information Processing Systems*, edited by M. Ranzato, A. Beygelzime, Y. Dauphin, P. S. Liang, and J. Wortman Vaughan (Curran Associates, Inc., 2021), Vol. 34, pp. 18137–18151.
- [26] R. Dietrich, M. Opper, and H. Sompolinsky, Statistical Mechanics of Support Vector Networks, *Phys. Rev. Lett.* **82**, 2975 (1999).
- [27] A. Canatar, B. Bordelon, and C. Pehlevan, Spectral bias and task-model alignment explain generalization in kernel regression and infinitely wide neural networks, *Nat. Commun.* **12**, 2914 (2021).
- [28] A. Jacot, F. Gabriel, and C. Hongler, Neural tangent kernel: Convergence and generalization in neural networks, in *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, Red Hook, NY (Curran Associates, 2018), pp. 8580–8589.
- [29] J. Lee, L. Xiao, S. Schoenholz, Y. Bahri, R. Novak, J. Sohl-Dickstein, and J. Pennington, Wide neural networks of any depth evolve as linear models under gradient descent, in *Advances in Neural Information Processing Systems* Vol. 32, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Curran Associates, 2019).
- [30] L. Bottou, Making Vapnik–Chervonenkis bounds accurate, in *Measures of Complexity*, edited by V. Vovk, H. Papadopoulos, and A. Gammernan (Springer, Cham, 2015), pp. 143–155.
- [31] M. Belkin, Fit without fear: Remarkable mathematical phenomena of deep learning through the prism of interpolation, *Acta Numerica* **30**, 203 (2021).
- [32] P. Rotondo, M. C. Lagomarsino, and M. Gherardi, Counting the learnable functions of geometrically structured data, *Phys. Rev. Research* **2**, 023169 (2020).
- [33] P. Rotondo, M. Pastore, and M. Gherardi, Beyond the Storage Capacity: Data-Driven Satisfiability Transition, *Phys. Rev. Lett.* **125**, 120601 (2020).
- [34] M. Pastore, P. Rotondo, V. Erba, and M. Gherardi, Statistical learning theory of structured data, *Phys. Rev. E* **102**, 032119 (2020).
- [35] M. Gherardi, Solvable model for the linear separability of structured data, *Entropy* **23**, 305 (2021).
- [36] M. Pastore, Critical properties of the SAT/UNSAT transitions in the classification problem of structured data, *J. Stat. Mech.: Theory Exp.* (2021) 113301.

- [37] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications* Vol. 9 (World Scientific, 1986).
- [38] E. Dobriban and S. Wager, High-dimensional asymptotics of prediction: Ridge regression and classification, *Ann. Stat.* **46**, 247 (2018).
- [39] T. Hastie, A. Montanari, S. Rosset, and R. J. Tibshirani, Surprises in high-dimensional ridgeless least squares interpolation, *Ann. Stat.* **50**, 949 (2022).
- [40] G. Mel and S. Ganguli, A theory of high dimensional regression with arbitrary correlations between input features and target functions: Sample complexity, multiple descent curves and a hierarchy of phase transitions, in *Proceedings of the 38th International Conference on Machine Learning*, edited by M. Meila and T. Zhang, *Proceedings of Machine Learning Research* Vol. 139 (PMLR, 2021), pp. 7578–7587.
- [41] A. C. C. Coolen, M. Sheikh, A. Mozeika, F. Aguirre-Lopez, and F. Antenucci, Replica analysis of overfitting in generalized linear regression models, *J. Phys. A: Math. Theor.* **53**, 365001 (2020).
- [42] A. Rahimi and B. Recht, Random features for large-scale kernel machines, *Adv. Neural Inf. Proc. Syst.* **20**, 1177 (2007).
- [43] B. Ghorbani, S. Mei, T. Misiakiewicz, and A. Montanari, Linearized two-layers neural networks in high dimension, *Ann. Stat.* **49**, 1029 (2021).
- [44] D. Pollard, *Convergence of Stochastic Processes* (Springer Science & Business Media, New York, 2012).
- [45] L. Chizat, E. Oyallon, and F. Bach, On lazy training in differentiable programming, *Adv. Neural Inf. Proc. Syst.* **32**, (2019).
- [46] Y. LeCun, The MNIST database of handwritten digits, <http://yann.lecun.com/exdb/mnist/> (1998).
- [47] P. L. Bartlett, N. Harvey, C. Liaw, and A. Mehrabian, Nearly-tight VC-dimension and pseudodimension bounds for piecewise linear neural networks, *J. Machine Learn. Res.* **20**, 2285 (2019).
- [48] B. Bordonon, A. Canatar, and C. Pehlevan, Spectrum dependent learning curves in kernel regression and wide neural networks, in *Proceedings of the 37th International Conference on Machine Learning*, edited by H. Daumé III and A. Singh, in *Proceedings of Machine Learning Research* Vol. 119 (PMLR, 2020), pp. 1024–1034.
- [49] A. Jacot, B. Simsek, F. Spadaro, C. Hongler, and F. Gabriel, Kernel alignment risk estimator: Risk prediction from training data, in *Advances in Neural Information Processing Systems* Vol. 33, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, (Curran Associates, 2020), pp. 15568–15578.
- [50] S. Goldt, B. Loureiro, G. Reeves, F. Krzakala, M. Mézard, and L. Zdeborová, The Gaussian equivalence of generative models for learning with shallow neural networks, in *Proceedings of the 2nd Mathematical and Scientific Machine Learning Conference*, edited by J. Bruna, J. Hesthaven, and L. Zdeborova, *Proceedings of Machine Learning Research* (PMLR, 2022), Vol. 145, pp. 426–471.
- [51] H. Hu and Y. M. Lu, Universality laws for high-dimensional learning with random features, [arXiv:2009.07669](https://arxiv.org/abs/2009.07669).
- [52] E. Gardner, The space of interactions in neural network models, *J. Phys. A: Math. Gen.* **21**, 257 (1988).
- [53] [https://github.com/rosalba-p/Generalisation\\_DNN](https://github.com/rosalba-p/Generalisation_DNN).
- [54] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, Pytorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems* Vol. 32 (Curran Associates, 2019), pp. 8024–8035.
- [55] D. Kingma and L. Ba, A method for stochastic optimization, in *International Conference on Learning Representations (ICLR)* (ICLR, San Diego, 2015).